Analyse spectrale à haute résolution et décompositions non-négatives appliquées au traitement des signaux de musique

Roland Badeau

2011D001

janvier 2011

Département Traitement du Signal et des Images
Groupe AAO : Audio, Acoustique et Ondes

Mémoire

présenté à l'Université Pierre et Marie Curie

pour obtenir

l'Habilitation à Diriger des Recherches

Analyse spectrale à haute résolution et décompositions non-négatives appliquées au traitement des signaux de musique

Roland BADEAU

Institut Télécom, Télécom ParisTech, CNRS LTCI
46 rue Barrault - 75634 Paris Cedex 13 - France

Habilitation soutenue le mardi 30 novembre 2010 devant le jury composé de :

Rapporteurs

Frédéric Bimbot, Directeur de Recherche CNRS - IRISA, Rennes
Philippe Depalle, Professeur - Université McGill, Montréal
Bruno Torrésani, Professeur - Université de Provence, Marseille

Examinateurs

Olivier Adam, Professeur - Université Pierre et Marie Curie, Paris 6
Laurent Daudet, Professeur - Université Paris Diderot, Paris 7
Sylvie Marcos, Directrice de Recherche CNRS - LSS, Gif-sur-Yvette

Résumé en français

Mes travaux de recherche portent sur le traitement des sons et de la musique, dont ils visent à extraire une représentation révélant leurs structures temporelles et spectrales. Diverses applications musicales ont été abordées, allant de la transcription automatique de musique à la séparation de sources, en passant par le codage audio. Une partie importante de ces recherches a été menée dans le cadre du projet ANR DESAM (Décompositions en Éléments Sonores et Applications Musicales), que j'ai coordonné de novembre 2006 à février 2010.

Analyse spectrale à haute résolution. La première partie de mes travaux a été initiée dans le cadre de ma thèse, démarrée en septembre 2001 et soutenue en avril 2005, dont la finalité a été l'application des méthodes d'analyse spectrale à haute résolution au traitement des signaux audio et musicaux. Nous avons d'abord généralisé le modèle de signal sous-jacent et la méthode d'estimation et proposé une nouvelle technique de sélection de l'ordre du modèle. Par ailleurs, le nombre élevé de composantes fréquentielles et la forte dynamique spectrale des signaux audio ont nécessité de mettre en œuvre des prétraitements adéquats et leur non-stationnarité a exigé de développer des méthodes adaptatives permettant de suivre les variations temporelles de chaque composante, en utilisant des algorithmes de poursuite de sous-espace. Ces outils ont été intégrés à la boîte à outils Matlab du projet DESAM et appliqués à la représentation temps-fréquence à haute résolution de signaux audio, à la séparation des composantes tonales et bruitées et à la synthèse avec effets sonores. A la suite de ma thèse, diverses applications musicales de ces méthodes ont été réalisées dans le cadre de collaborations avec des chercheurs et doctorants de laboratoires français ou étrangers : estimation de hauteur de notes de piano (thèse de Valentin Emiya à Télécom ParisTech), estimation de tempo musical (thèse de Miguel Alonso à Télécom ParisTech), analyse des modes couplés d'une harpe de concert (thèse de Jean-Loïc Le Carrou au LAUM, Le Mans) et codage audio (projet DESAM, collaboration avec le LMA, Marseille). D'autres applications ont également été réalisées, comme l'estimation de canal dans le domaine des communications numériques (collaboration avec l'université de Mondragón en Espagne) ou la factorisation de tenseurs structurés (collaboration avec le LSS, Gif-sur-Yvette).

Décompositions non-négatives. La deuxième partie de mes travaux constitue une thématique nouvelle depuis ma thèse, dont les différents aspects ont notamment été abordés avec les cinq doctorants que j'ai co-encadrés (Valentin Emiya, Nancy Bertin, Romain Hennequin, Benoît Fuentes et Antoine Liutkus). Elle s'intéresse à la décomposition de signaux audio à partir d'une représentation spectrale ou d'une représentation temps-fréquence, dans le but d'en extraire des éléments sonores possédant un sens musical, comme des notes de musique ou des accords. Ces décompositions prennent en compte différents critères : non-négativité de la représentation et des éléments sonores, critères psycho-acoustiques, harmonicité du spectre, régularité de l'enveloppe spectrale, régularité de l'enveloppe temporelle, variations temporelles de l'enveloppe spectrale, variations temporelles de la fréquence fondamentale. Elles s'appuient sur des extensions de la factorisation en matrices non-négatives, exprimées dans le formalisme de l'estimation bayésienne, qui permet de contraindre la décomposition en introduisant des modèles paramétriques appropriés ou des distributions *a priori* des paramètres. Les décompositions sont calculées à l'aide d'algorithmes multiplicatifs, dont nous avons étudié et amélioré les propriétés de convergence, ou d'algorithmes de type espérance-maximisation. Ces outils ont été appliqués à l'estimation de fréquences fondamentales multiples et à la transcription automatique de la musique. Ils ont également été intégrés à la boîte à outils Matlab du projet DESAM et leur évaluation a conduit au développement de la base de données MAPS (*MIDI aligned piano sounds*). Ils ont enfin été appliqués à la séparation informée des différentes sources sonores mixées dans un enregistrement stéréophonique.

English abstract

My research works focus on audio and music signal processing. They aim to extract from the signal a representation which reveals its temporal and spectral structures. Various musical applications have been addressed, ranging from automatic transcription of music to source separation and audio coding. An important part of this research has been carried out within the framework of the ANR DESAM (Decompositions into sound elements and musical applications) project, that I have managed from November 2006 to February 2010.

High resolution spectral analysis. The first part of my research works was initiated within the framework of my Ph.D. thesis, started in September 2001 and defended in April 2005, whose purpose was the application of high resolution spectral analysis methods to audio and music signal processing. We first generalized the underlying signal model and the estimation method, and proposed a new model order selection technique. In other respects, the high number of frequency components and the high spectral dynamics of audio signals have required to develop adequate preprocessing techniques, and their non-stationarity has required to develop adaptive methods in order to track the temporal variations of each component, by using subspace tracking algorithms. These tools have been integrated into the Matlab toolbox of the DESAM project, and applied to the high resolution time-frequency representation of audio signals, to the separation of tonal and noisy components, and to audio synthesis with sound effects. After my Ph.D. thesis, various musical applications of these methods have been carried out in collaboration with researchers and Ph.D. students from French and foreign laboratories: pitch estimation of piano sounds (Ph.D. thesis of Valentin Emiya at Telecom ParisTech), music beat tracking (Ph.D. thesis of Miguel Alonso at Telecom ParisTech), analysis of harp's sympathetic modes (Ph.D. thesis of Jean-Loïc Le Carrou at LAUM, Le Mans), and audio coding (DESAM project, collaboration with the LMA, Marseilles). Other applications have also been carried out, such as blind channel estimation in digital communications (collaboration with the University of Mondragón in Spain) or structured tensors factorization (collaboration with the LSS, Gif-sur-Yvette).

Non-negative decompositions. The second part of my research works has formed a new theme since my Ph.D. thesis, whose various aspects have been addressed with the five Ph.D. students that I jointly supervised (Valentin Emiya, Nancy Bertin, Romain Hennequin, Benoît Fuentes, and Antoine Liutkus). It focuses on the decomposition of audio signals from a spectral or a time-frequency representation, in order to extract sound elements having a musical meaning, such as music notes or chords. The decompositions take different criteria into account: non-negativity of the representation and the sound elements, psycho-acoustic criteria, spectral harmonicity, spectral envelope smoothness, temporal envelope smoothness, temporal variations of the spectral envelope, or temporal variations of the fundamental frequency. They rely on extensions of non-negative matrix factorization, expressed in the framework of Bayesian estimation, which permits to constrain the decomposition by introducing appropriate parametric models, or *a priori* distributions of the parameters. The decompositions are computed by means of multiplicative algorithms, whose convergence properties have been studied and enhanced, or variants of the expectation-maximization algorithm. These tools have been applied to multiple fundamental frequency estimation and to automatic transcription of music. They have been integrated into the Matlab toolbox of the DESAM project and their evaluation has leaded to the development of the MAPS (MIDI aligned piano sounds) database. They have also been applied to the informed separation of various sound sources mixed in a stereophonic recording.

Remerciements

Je remercie avant tout les rapporteurs (Frédéric Bimbot, Philippe Depalle, Bruno Torrésani) et les examinateurs (Olivier Adam, Laurent Daudet, Sylvie Marcos) qui m'ont fait l'honneur d'accepter de participer à mon jury d'Habilitation à Diriger des Recherches.

Je dois beaucoup à mes collègues Gaël Richard, Professeur et responsable du groupe Audio, Acoustique et Ondes à Télécom ParisTech, et Bertrand David, Maître de Conférences à Télécom ParisTech, qui après avoir dirigé mon doctorat, ont encadré et continuent à encadrer avec moi les thèses dont je suis co-directeur.

Bien sûr une partie importante des travaux de recherche présentés dans ce mémoire n'aurait pu être menée à bien sans la contribution active de mes doctorants : dans l'ordre chronologique, Valentin Emiya, actuellement en post-doctorat à l'IRISA, Nancy Bertin, Chargée de Recherche CNRS à l'IRISA, ainsi que Romain Hennequin, Benoît Fuentes et Antoine Liutkus, toujours doctorants à Télécom ParisTech.

Je dois à mes collaborateurs la réussite du projet ANR DESAM (Décompositions en Éléments Sonores et Applications Musicales), dont j'ai eu l'honneur d'être le coordinateur pendant un peu plus de trois ans : dans l'ordre alphabétique, Laurent Daudet, Professeur à l'Université Paris Diderot, Olivier Derrien, Maître de Conférences à l'Université de Toulon et du Var, Mathieu Lagrange, alors post-doctorant à Télécom ParisTech et aujourd'hui Chargé de Recherche CNRS à l'IRCAM, et Sylvain Marchand, Maître de Conférences à l'Université Bordeaux 1.

Je remercie tous mes collègues enseignants-chercheurs du département de Traitement du Signal et des Images de Télécom ParisTech, notamment Yves Grenier, Professeur et responsable du département, Karim Abed-Meraim, Maître de Conférences, et Cédric Févotte, Chargé de Recherche CNRS, qui ont co-signé plusieurs publications avec moi.

Je tiens aussi à remercier mes autres co-auteurs, avec qui j'ai pu établir des collaborations fortes hors de tout contexte institutionnel, notamment Emmanuel Vincent, Chargé de Recherche CNRS à l'IRISA, Rémy Boyer, Maître de Conférences à l'Université Paris XI, Jean-Loïc Le Carrou, Maître de Conférences à l'Université Pierre et Marie Curie, et François Gautier, Professeur à l'Université du Maine.

Au delà des personnes, je voudrais également remercier les institutions qui ont soutenu mes recherches. En premier lieu Télécom ParisTech, qui m'a d'abord formé puis offert un environnement de travail extrêmement dynamique, et le LTCI (Laboratoire de Traitement et Communication de l'Information), unité mixte de recherche CNRS regroupant l'ensemble des activités de recherche de Télécom ParisTech. Ensuite l'ANR et le CNRS, qui ont respectivement financé et géré le projet DESAM, et auprès de qui Cécile Guion et Christine Barba ont été mes interlocutrices.

Enfin, parvenir au terme de cette aventure aurait été plus difficile sans le soutien bienveillant et chaleureux de ma famille et de mes proches. Je remercie tout particulièrement ma chère Sonia, dont le soutien et les encouragements quotidiens m'auront porté jusqu'au bout de la rédaction de ce mémoire.

Table des matières

Table des figures

Liste des tableaux

Acronymes

ANR	Agence Nationale de la Recherche
AR	autorégressif
ARMA	autorégressif à moyenne ajustée
ATIAM	Acoustique, Traitement du signal, Informatique Appliqués à la Musique
Bi-SVD	*Bi-iteration SVD*
BPM	battements par minute
DAP	*Discrete All-pole Modeling*
DESAM	Décompositions en Éléments Sonores et Applications Musicales
dB	decibels
EM	Espérance-Maximisation
ENS	École Normale Supérieure
ENST	École Nationale Supérieure des Télécommunications
ERB	*Equivalent Rectangular Bandwidth*
ESM	*Exponential Sinusoidal Model*
ESPRIT	*Estimation of Signal Parameters via Rotational Invariance Techniques*
ESTER	*ESTimation ERror*
EUC	euclidienne
FAPI	*Fast Approximated Power Iteration*
FDPM	*Fast Data Projection Method*
GIPSA-Lab	Grenoble-Images-Parole-Signal-Automatique
GPL	*GNU General Public License*
HOSVD	*Higher-Order SVD*
HR	à Haute Résolution
HR-ogramme	Spectrogramme à Haute Résolution
Hz	Hertz
IRCAM	Institut de Recherche et Coordination Acoustique/Musique
IRISA	Institut de Recherche en Informatique et Systèmes Aléatoires
IS	Itakura-Saito
ISS	*Informed Source Separation*
JCJC	Jeunes Chercheuses, Jeunes Chercheurs
KL	Kullback-Leibler
LaBRI	Laboratoire Bordelais de Recherche en Informatique
LAM	Lutheries-Acoustique-Musique
LAUM	Laboratoire d'Acoustique de l'Université du Maine

LIRIS	Laboratoire d'InfoRmatique en Image et Systèmes d'information
LMA	Laboratoire de Mécanique et Acoustique
LORAF	*LOw Rank Adaptive Filters*
LSS	Laboratoire des Signaux et Systèmes
LTCI	Laboratoire Traitement et Communication de l'Information
MA	à moyenne ajustée
MAC	*multiply / accumulate*
MAP	maximum *a posteriori*
MAPS	*MIDI Aligned Piano Sounds*
MDCT	*Modified Discrete Cosine Transform*
MFCC	*Mel-Frequency Cepstral Coefficients*
MIDI	*Musical Instrument Digital Interface*
MIR	*Music Information Retrieval*
MIREX	*Music Information Retrieval Evaluation eXchange*
MP3	*MPEG-1/2 audio layer 3*
MP4	*MPEG-2/4 AAC*
MPEG	*Moving Picture Experts Group*
MUSIC	*MUltiple Signal Classification*
NKSVD	*Non-negative K-SVD*
NMF	*Non-negative Matrix Factorization*
NP	*Natural Power method*
ODG	*Objective Difference Grade*
OPAST	*Orthonormal PAST*
PACE	*Polynomial Amplitude Complex Exponentials*
PARAFAC	*PARAllel FACtor analysis*
PAST	*Projection Approximation Subspace Tracking*
PEMO-Q	*PErception MOdel Quality*
PQMF	*Pseudo-Quadrature Mirror Filters*
RBF	*Radial Basis Function*
RSB	rapport signal sur bruit
RSBQ	rapport signal sur bruit de quantification
RWC	*Real World Computing*
SAGE	*Space Alternating Generalized EM*
SP	*Subspace Projection*
STAP	*Space-Time Adaptive Processing*
SVD	*Singular Value Decomposition*
SW-PAST	*Sliding Window PAST*
SW-OPAST	*Sliding Window OPAST*
SWASVD	*Sliding Window Adaptive SVD*
TFCT	Transformée de Fourier à Court Terme
TSI	Traitement du Signal et des Images
UPMC	Université Pierre et Marie Curie
X	École Polytechnique
YAST	*Yet Another Subspace Tracker*

Notations

Ensembles

$\mathbb{Z}$, $\mathbb{N}$	ensembles des entiers, respectivement relatifs et naturels
$\mathbb{R}$, $\mathbb{C}$	ensembles des nombres réels et des nombres complexes
$(.)^*$	ensemble privé de zéro
$(.)_+$	ensemble privé des nombres strictement négatifs

Variables

z	scalaire (caractère maigre minuscule)
$\boldsymbol{v}$	vecteur (caractère gras minuscule)
v_i	coefficient du vecteur $\boldsymbol{v}$ d'indice i
$\boldsymbol{v}(z)$	vecteur contenant les puissances successives du scalaire z : $\boldsymbol{v}(z) = [1, z, z^2, \ldots]^T$
$\boldsymbol{M}$	matrice (caractère gras majuscule)
m_{ij}	coefficient de la matrice $\boldsymbol{M}$ d'indices (i, j)

Matrices

$\boldsymbol{I}_n$	matrice identité de dimensions $n \times n$
diag(.)	matrice diagonale construite à partir du vecteur en argument
$(.)^T$	transposé
$(.)^*$	conjugué
$(.)^H$	conjugué hermitien
$(.)^{\frac{1}{2}}$	racine carrée matricielle
$(.)^\dagger$	pseudo-inverse d'une matrice
Im(.)	espace image d'une matrice
trace(.)	trace d'une matrice
$\|.\|_2$	norme euclidienne d'un vecteur ou norme spectrale d'une matrice
$\otimes$	produit matriciel terme à terme
$\frac{(.)}{(.)}$	division matricielle terme à terme
$(.)_\uparrow$	matrice extraite en supprimant la première ligne
$(.)_\downarrow$	matrice extraite en supprimant la dernière ligne

Probabilités

$\mathbb{E}(.)$	espérance mathématique
$p(x\|y)$	probabilité de x sachant y
$\widehat{(.)}$	estimateur d'un paramètre
$\sim \mathcal{N}(\boldsymbol{m}, \boldsymbol{R})$	vecteur aléatoire gaussien de moyenne $\boldsymbol{m}$ et de matrice de covariance $\boldsymbol{R}$

Divers

$\overset{c}{=}$	égalité à une constante additive près
$\propto$	égalité à un facteur multiplicatif près
$O(.)$	grandeur dominée par la fonction en argument
$d(.\|.)$	mesure de divergence scalaire
$\boldsymbol{\nabla}(.)$	vecteur gradient
$\Gamma(.)$	fonction Gamma ($\Gamma(z) = \int_0^{+\infty} t^{z-1} e^{-t} dt$)

Chapitre I

Avant-propos

Les récents progrès des technologies de l'information ont conduit à un accroissement considérable de la quantité de données à traiter (sons, images, vidéos, textes, *etc.*). L'archivage et la transmission de ces données ont nécessité la mise au point de techniques avancées de compression et de communications numériques. Leur multiplication a également motivé le développement de traitements automatisés : requête dans une base de données, indexation, classification, extraction de caractéristiques, résumé, restauration ou débruitage.

D'autres applications sont spécifiques aux types de données traitées. Ainsi le traitement d'antenne (formation de voies, localisation de sources, séparation de sources) et l'analyse spectrale (analyse modale, spectroscopie) s'intéressent aux phénomènes vibratoires, comme les ondes électromagnétiques (ondes lumineuses, ondes radio) ou les ondes mécaniques (ondes sonores, ondes sismiques, cordes vibrantes). Ils sont utilisés en astronomie, physique, chimie, mécanique, optique ou encore en acoustique. Le traitement d'images, statiques ou animées, a des applications très variées : défloutage, segmentation (détection de contours, modélisation de textures), reconnaissance de formes (visages, caractères écrits ou imprimés, empreintes digitales, *etc.*). Le génie biomédical a pour finalité l'aide au diagnostic et au traitement des patients et s'applique à des signaux comme les électrocardiogrammes (ECG), les électroencéphalogrammes (EEG), les séquences ADN (gènes) ou encore les images produites par des systèmes d'imagerie par résonance magnétique (IRM). Le traitement des séries financières a connu un véritable essor ces dernières années. Il a pour objectif d'identifier les tendances des marchés boursiers, avec des applications comme la prédiction du risque ou l'optimisation de portefeuille. Enfin, le traitement des signaux audiofréquences porte sur divers types de sons : parole, musique ou sons environnementaux, avec des applications comme l'analyse / synthèse, la transcription, la reconnaissance de l'instrument ou du locuteur et l'annulation d'écho. Dans ce contexte pluridisciplinaire, mes travaux de recherche ont été plus particulièrement appliqués au traitement des sons et de la musique.

Décompositions de signaux audiofréquences

Les signaux audiofréquences obtenus en enregistrant une performance musicale sont généralement composés de plusieurs éléments sonores, comme les notes de musique, les accords, les sons percussifs et transitoires ou les sons bruités, qui sont émis par différentes sources sonores (comme la voix du chanteur ou les divers instruments de musique). Parvenir à identifier automatiquement les éléments et les sources sonores présents dans un signal représente un enjeu important pour réaliser diverses applications, comme la transcription automatique de musique (qui nécessite en particulier de reconnaître les instruments et d'estimer la hauteur des notes, les instants d'attaque et la mesure rythmique), la séparation de sources sonores, le codage audio ou encore la synthèse avec modifications (modification de hauteur tonale, de durée, suppression ou ajout d'un effet de vibrato ou de trémolo, *etc.*). Une telle opération s'avère cependant très délicate à partir de la seule représentation temporelle du signal, particulièrement lorsque plusieurs sources sonores se superposent. Il est toutefois possible de séparer des sons superposés lorsque leurs spectres ne se recouvrent pas (par exemple dans le cas de deux sons harmoniques dont les fréquences fondamentales ne sont pas en relation d'harmonicité). On utilise alors des représentations temps-fréquence comme le spectrogramme, calculé à partir de la Transformée de Fourier à Court Terme (TFCT) [Allen et Rabiner, 1977], ou ses variantes de la classe de Cohen [Cohen, 1995]. Sur ces représentations, il devient possible de distinguer les différents éléments sonores, reconnaissables grâce à leurs propriétés temps-fréquences particulières (fréquences des partiels, harmonicité, enveloppes spectrale et temporelle, *etc.*). Il reste néanmoins difficile de séparer des éléments sonores dont les supports dans le plan temps-fréquence ne sont

pas disjoints. Par ailleurs, même s'il existe des méthodes de réassignement permettant d'améliorer la précision temporelle et fréquentielle des représentations temps-fréquences de la classe de Cohen [Auger et Flandrin, 1995], leur résolution spectrale reste limitée par la taille de la fenêtre utilisée.

Représentations parcimonieuses

De nombreuses méthodes ont été proposées pour détecter et extraire automatiquement les éléments sonores composant un signal de musique. Dans la littérature, une approche très répandue s'appuie sur la propriété de parcimonie [Mallat, 2008] : la description d'un signal de musique, constitué d'un très grand nombre d'échantillons, en un nombre limité d'éléments sonores distincts constitue en effet une représentation parcimonieuse de ce signal. Les décompositions parcimonieuses visent ainsi à fournir une approximation du signal, exprimée comme une combinaison linéaire d'un nombre réduit d'éléments extraits d'un dictionnaire. Cette décomposition peut être obtenue à l'aide d'un algorithme itératif comme le *Matching Pursuit* [Mallat et Zhang, 1993], très utilisé pour sa rapidité et sa résistance au bruit additif. Cet algorithme glouton ne garantissant pas l'optimalité de la solution, des méthodes d'optimisation globale plus coûteuses ont également été proposées [Chen *et al.*, 1998, Tibshirani, 1996]. Enfin, la connaissance des propriétés particulières des éléments sonores recherchés peut être intégrée dans la structure du dictionnaire utilisé [Gribonval et Bacry, 2003], qui peut être également appris à partir d'une base de données ou adapté au signal à décomposer [Daudet et Torrésani, 2006].

En traitement des signaux audiofréquences, les applications usuelles des décompositions parcimonieuses incluent la compression, le débruitage, la séparation de sources, la transcription de musique [Plumbley *et al.*, 2010] et la reconnaissance d'instruments [Leveau *et al.*, 2007]. Les décompositions parcimonieuses ont également été utilisées dans des applications aussi variées que l'acquisition comprimée de signaux (*compressed sensing*) [Donoho, 2006], la classification automatique [Huang et Aviyente, 2007], l'extraction de caractéristiques et la reconnaissance de formes [Wright *et al.*, 2009], le traitement d'images (restauration, défloutage, débruitage, compression) [Elad, 2010], la génétique [Wu *et al.*, 2009] ou encore les neurosciences [Fadili et Bullmore, 2005].

Approximation par réduction de rang

Par rapport aux décompositions parcimonieuses, les méthodes d'analyse du signal que j'ai étudiées s'appuient également sur la notion de parcimonie, qui est toutefois abordée sous un angle différent. Il s'agit d'utiliser des techniques d'analyse matricielle consistant à calculer une approximation d'une matrice de données, s'exprimant comme un produit de matrices de rang inférieur. La réduction de rang est en effet une stratégie très utilisée pour extraire les principales caractéristiques d'un ensemble de données. Un des outils les plus répandus est l'analyse en composantes principales [Jolliffe, 2002], qui fait appel à la décomposition en valeurs singulières - *Singular Value Decomposition* (SVD) [Golub et Van Loan, 1996]. Selon le type des données figurant dans la matrice à décomposer, différentes caractéristiques pourront être extraites. Dans mes travaux de recherche, j'ai ainsi abordé deux approches d'analyse de signaux dont le dénominateur commun est la notion de réduction de rang : l'**analyse spectrale à Haute Résolution (HR)** et les **décompositions non-négatives**.

Méthodes à haute résolution

L'analyse spectrale à haute résolution [Marcos, 1998] s'appuie sur une représentation matricielle des données obtenue à partir d'une suite d'observations vectorielles, constituées d'échantillons successifs du signal. On montre en effet que la matrice obtenue à partir des échantillons d'un signal constitué de composantes fréquentielles pures, dont l'amplitude est modulée exponentiellement, est théoriquement de rang égal au nombre de ces composantes. Ainsi les méthodes à haute résolution de type "sous-espace", comme *MUltiple Signal Classification* (MUSIC) [Schmidt, 1986] et *Estimation of Signal Parameters via Rotational Invariance Techniques* (ESPRIT) [Roy et Kailath, 1989], exploitent la factorisation de cette matrice pour estimer les paramètres de ces composantes. Ces méthodes permettent d'atteindre une résolution fréquentielle supérieure à celles basées sur une détection de pics dans une représentation de Fourier. L'application des méthodes HR aux signaux audiofréquences s'est avérée difficile en raison de leur complexité algorithmique et de la variabilité temporelle des paramètres (liée aux non-stationnarités) ; nous avons donc développé des algorithmes rapides de poursuite des paramètres, basés sur des algorithmes de poursuite de sous-espace [Badeau *et al.*, 2005a, Badeau *et al.*, 2008b].

Les domaines classiques d'application des méthodes à haute résolution sont l'analyse spectrale et le traitement d'antenne [Marcos, 2009]. Dans le cadre de l'analyse spectrale de signaux sonores, diverses applications des méthodes HR ont été proposées, comme le codage audio [Jensen *et al.*, 2004], la séparation de pistes de batterie [Gillet et Richard, 2008], l'estimation de hauteur [Christensen *et al.*, 2006], la modélisation de sons percussifs ou de contact [Reissell et Pai, 2007a, Lagrange et Scherrer, 2008] ou encore la compression [Haddad et Noga, 2007]

et le débruitage [Hermus *et al.*, 2007] de signaux de parole. Par ailleurs, l'analyse spectrale à haute résolution a également été appliquée à des problèmes variés comme l'analyse modale de structures [Ege *et al.*, 2009b], l'analyse de signaux biomédicaux [Laudadio, 2005] et même le contrôle de l'état de systèmes d'embrayage [Ompusunggu *et al.*, 2009]. En traitement d'antenne, les applications les plus courantes des méthodes HR sont la formation de voies [Foutz et Spanias, 2007], la localisation de sources [Guo *et al.*, 2006, Valizadeh *et al.*, 2007] et la séparation de sources [Zhang *et al.*, 2010]. Ces méthodes ont par exemple été utilisées en acoustique sous-marine (sonar, tomographie des océans), en traitement des ondes radio (radar, radiogoniométrie, radio-astronomie), et pour séparer des ondes sismiques [Marcos, 2009]. Enfin, dans le domaine des communications numériques, les méthodes HR sont couramment utilisées dans des applications d'identification et d'égalisation de canaux [Tong et Perreau, 1998, Cardoso *et al.*, 1995].

Décompositions non-négatives

L'analyse spectrale à haute résolution constitue une analyse "bas niveau" du signal : pour identifier les éléments sonores mentionnés plus haut, il est nécessaire de regrouper les composantes fréquentielles appartenant à un même élément, ce qui nécessite par exemple de faire appel à un algorithme de classification. La seconde approche que j'ai étudiée présente l'avantage de fournir directement une décomposition en éléments sonores de plus haut niveau. Il s'agit de la factorisation en matrices non-négatives - *Non-negative Matrix Factorization* (NMF) [Lee et Seung, 1999] et de ses diverses variantes qui fournissent également une décomposition du signal en facteurs non-négatifs. La représentation matricielle utilisée est alors une représentation temps-fréquence, typiquement un spectrogramme d'amplitude ou de puissance. Par rapport à la SVD, on ajoute ici la contrainte de positivité des coefficients de la décomposition. Le signal est alors décomposé comme une combinaison linéaire de formes spectrales (par exemple des spectres de notes isolées) multipliées par des activations temporelles (permettant d'identifier les moments où ces notes sont jouées). Par rapport à l'approche précédente, on perd la haute résolution spectrale (limitée par la représentation temps-fréquence utilisée), mais le modèle sous-jacent est plus général, puisqu'il permet aussi de représenter des éléments non-sinusoïdaux comme des sons percussifs ou du bruit. En outre, de récents travaux ont conduit à formuler le problème de NMF comme un problème d'estimation statistique [Virtanen *et al.*, 2008, Schmidt et Laurberg, 2008, Févotte *et al.*, 2009, Shashanka *et al.*, 2008]. Cette approche a ouvert de nombreuses perspectives, puisqu'elle permet d'incorporer aisément des informations sur les propriétés particulières des éléments sonores à extraire, à l'aide de modèles paramétriques ou de distributions *a priori* dans le cadre de l'inférence bayésienne.

Dans le domaine du traitement des sons, les applications de la NMF incluent la séparation de sources [Virtanen, 2007, FitzGerald *et al.*, 2008], l'estimation de hauteurs [Raczyński *et al.*, 2007, Smaragdis, 2009], la transcription de musique [Smaragdis et Brown, 2003, Févotte *et al.*, 2009] ou encore la classification d'instruments de musique [Cichocki *et al.*, 2009]. En plus de l'analyse de signaux audiofréquences, la NMF a également été appliquée dans des domaines aussi variés que l'analyse d'images [Lee et Seung, 1999, Cichocki *et al.*, 2009], la fouille de texte [Lee et Seung, 1999, Berry *et al.*, 2007], la spectroscopie [Berry *et al.*, 2007, Cichocki *et al.*, 2009], l'analyse de signaux biomédicaux [Chen *et al.*, 2006, Cichocki *et al.*, 2009], la finance [Drakakis *et al.*, 2007], et même la classification automatique d'eaux-de-vie [Young *et al.*, 2006].

Les méthodes d'analyse spectrale à haute résolution et de décompositions non-négatives présentées ci-dessus ont été appliquées à des signaux de musique dans le cadre du projet Décompositions en Éléments Sonores et Applications Musicales (DESAM), que j'ai coordonné du mois de novembre 2006 au mois de février 2010. Ce projet a également abordé l'utilisation de décompositions parcimonieuses pour le codage audio scalable et les transformations sonores [Ravelli *et al.*, 2008], la poursuite de partiels pour la modélisation de mélanges de sinusoïdes non-stationnaires et son application aux transformations sonores [Marchand et Depalle, 2008] ou encore la séparation de la mélodie principale et de son accompagnement dans un enregistrement musical [Durrieu *et al.*, 2008]. Un sous-ensemble des outils développés dans le cadre du projet a été intégré dans une boîte à outils baptisée *DESAM Toolbox* [Lagrange *et al.*, 2010b]. Il s'agit d'une bibliothèque de fonctions Matlab et de scripts de démonstration distribués sous les termes de la *GNU General Public License* (GPL).

Ce mémoire est structuré de la façon suivante : la partie I est consacrée aux méthodes d'analyse spectrale à haute résolution, aux algorithmes adaptatifs rapides de poursuite des paramètres et à diverses applications de ces méthodes. La partie II est consacrée aux décompositions non-négatives de spectres et de représentations temps-fréquences et à leurs applications en estimation de hauteur, en transcription automatique et en séparation de sources. Les annexes (partie III) contiennent mon curriculum vitae, ma bibliographie personnelle et les copies des dix publications les plus représentatives de mes travaux de recherche.

Première partie

Analyse à haute résolution

Chapitre II

Introduction de la Première partie

Les résultats présentés dans cette première partie sont basés sur ma thèse de doctorat [Badeau, 2005], démarrée en septembre 2001 et soutenue en avril 2005, dont le sujet était "Méthodes à haute résolution pour l'estimation et le suivi de sinusoïdes modulées. Application aux signaux de musique". Ces travaux ont été approfondis par la suite, notamment concernant les aspects de poursuite et les applications des méthodes HR, dont les publications les plus récentes datent de 2009.

En analyse spectrale à haute résolution, on utilise le modèle de signal baptisé *Exponential Sinusoidal Model* (ESM), qui représente un signal réel $s(t)$ à temps discret $t \in \mathbb{Z}$ comme la somme de $r \in \mathbb{N}^*$ composantes sinusoïdales modulées exponentiellement [Badeau, 2005] : $s(t) = \sum_{k=1}^{r} a_k \, e^{-\delta_k t} \cos(2\pi\nu_k t + \phi_k)$, où $a_k \in \mathbb{R}_+^*$ est l'amplitude réelle, $\delta_k \in \mathbb{R}$ est le facteur d'atténuation (ou taux d'amortissement, éventuellement négatif dans le cas d'une amplification), $\nu_k \in]-\frac{1}{2}, \frac{1}{2}]$ est la fréquence réduite et $\phi_k \in]-\pi, \pi]$ est la phase de la composante k. Pour des raisons de commodité d'écriture, nous utiliserons plutôt le modèle complexe suivant, qui inclut le modèle réel et en fournit une représentation analytique :

$$s(t) = \sum_{k=1}^{r} \alpha_k \, {z_k}^t, \tag{II.1}$$

où l'on a introduit les amplitudes complexes $\alpha_k = a_k \, e^{i\phi_k} \in \mathbb{C}^*$ et les pôles complexes $z_k = e^{-\delta_k + i2\pi\nu_k} \in \mathbb{C}^*$, que l'on suppose tous distincts. On remarque que les pôles de même angle polaire et de modules différents sont associés à la même fréquence. Le modèle ESM permet donc pour chaque fréquence de représenter des modulations d'amplitudes qui sont des combinaisons linéaires d'exponentielles. Par rapport à une représentation de Fourier discrète, on gagne ainsi la possibilité de représenter des composantes de fréquences quelconques et d'amplitudes variables. En pratique, le signal observé $x(t)$ ne satisfait jamais exactement le modèle ESM, c'est pourquoi on le représente comme la somme du signal déterministe $s(t)$ défini dans l'équation (II.1) et d'un bruit blanc gaussien complexe $b(t)$, centré et de variance σ^2 : $x(t) = s(t) + b(t)$.

II.1 Analyse spectrale (chapitre III)

Les méthodes HR [Marcos, 1998] permettant d'estimer les pôles z_k peuvent être essentiellement classées en deux catégories : les méthodes de type prédiction linéaire, comme la méthode de Prony [Riche de Prony, 1795], et les méthodes de type sous-espace, comme MUSIC [Schmidt, 1986] et ESPRIT [Roy et Kailath, 1989]. Nous choisissons ici de présenter la méthode ESPRIT, dont on introduit une généralisation dans la section III.1 page 7, et sur laquelle une part importante de nos travaux se sont appuyés.

Soit un entier $n > r$. Pour tout $t \geq 0$, on définit le vecteur $\boldsymbol{s}(t) = [s(t), s(t+1), \ldots, s(t+n-1)]^T$ de dimension n (où la notation ^{T} désigne la transposition). Alors la matrice $\boldsymbol{S}(t) = [\boldsymbol{s}(0), \boldsymbol{s}(1), \ldots, \boldsymbol{s}(t)]$ est une matrice de Hankel de dimensions $n \times (t+1)$, qui contient les échantillons $s(0), \ldots, s(t+n-1)$ du signal défini dans l'équation (II.1). Pour tout $z \in \mathbb{C}$, on définit également le vecteur $\boldsymbol{v}(z) = [1, z, \ldots, z^{n-1}]^T$ de dimension n. Alors la matrice $\boldsymbol{V}^n = [\boldsymbol{v}(z_1) \ldots \boldsymbol{v}(z_r)]$ est une matrice de **Vandermonde** de dimensions $n \times r$. On démontre que $\boldsymbol{S}(t)$ se factorise sous la forme

$$\boldsymbol{S}(t) = \boldsymbol{V}^n \, \boldsymbol{A} \, {\boldsymbol{V}^{t+1}}^T, \tag{II.2}$$

où $\boldsymbol{A}$ est une matrice diagonale de coefficients $\alpha_1, \ldots, \alpha_r$, ce que nous notons $\boldsymbol{A} = \text{diag}(\alpha_1 \ldots \alpha_r)$, et $\boldsymbol{V}^{t+1}$ est une matrice de Vandermonde de dimensions $(t+1) \times r$ [Badeau, 2005]. Si $t \geq r-1$ et si les r pôles $\{z_1, \ldots, z_r\}$ sont distincts, alors la matrice $\boldsymbol{S}(t)$ est de rang $r < n$, et son espace image est engendré par la matrice $\boldsymbol{V}^n$. Comme cet espace image caractérise les composantes fréquentielles de $s(t)$, il est couramment nommé **espace signal**. En présence du bruit $b(t)$, on définit pour tout $t \geq 0$ les vecteurs $\boldsymbol{x}(t) = [x(t), x(t+1), \ldots, x(t+n-1)]^T$, contenant les échantillons du signal $x(t) = s(t) + b(t)$, et les matrices $\boldsymbol{X}(t) = [\boldsymbol{x}(0), \boldsymbol{x}(1), \ldots, \boldsymbol{x}(t)]$. On introduit également la matrice de "corrélation"

$$\boldsymbol{C}_{xx}(t) = \sum_{\tau=0}^{t} \gamma_\tau \, \boldsymbol{x}(t-\tau) \, \boldsymbol{x}(t-\tau)^H = \boldsymbol{X}(t) \, \boldsymbol{\Gamma}(t) \, \boldsymbol{X}(t)^H, \tag{II.3}$$

où $(\gamma_\tau)_{\tau \in \mathbb{N}}$ est une suite de coefficients positifs, sommable, dont les r premières valeurs sont non-nulles, $\boldsymbol{\Gamma}(t) = \text{diag}(\gamma_t, \gamma_{t-1}, \ldots, \gamma_0)$ et la notation ^{H} désigne le conjugué hermitien. Comme $b(t)$ est un bruit blanc centré de variance σ^2, en substituant l'équation (II.2) dans l'équation (II.3), on obtient l'espérance mathématique

$$\mathbb{E}(\boldsymbol{C}_{xx}(t)) = \boldsymbol{V}^n \boldsymbol{P}(t) \boldsymbol{V}^{n\,H} + \sigma^2 \gamma(t) \boldsymbol{I}_n,$$

où $\boldsymbol{P}(t) = \boldsymbol{A} \boldsymbol{V}^{t+1\,T} \boldsymbol{\Gamma}(t) \boldsymbol{V}^{t+1\,*} \boldsymbol{A}^*$ est une matrice à symétrie hermitienne de dimensions $r \times r$ (la notation * désigne la conjugaison complexe) définie positive pour tout $t \geq r-1$, $\gamma(t) = \sum_{\tau=0}^{t} \gamma_\tau$, et $\boldsymbol{I}_n$ est la matrice identité de dimensions $n \times n$. Soit alors $\{\boldsymbol{w}_1(t), \ldots, \boldsymbol{w}_n(t)\}$ une base orthonormée de vecteurs propres de la matrice $\mathbb{E}(\boldsymbol{C}_{xx}(t))$, associés aux valeurs propres $\lambda_1(t), \ldots, \lambda_n(t)$ rangées par ordre décroissant. Si $t \geq r-1$, on a $\lambda_1(t) \geq \ldots \geq \lambda_r(t) > \lambda_{r+1}(t) = \ldots = \lambda_n(t) = \sigma^2 \gamma(t)$. Par conséquent, le **sous-espace principal** de dimension r de la matrice $\mathbb{E}(\boldsymbol{C}_{xx}(t))$, engendré par la matrice $\boldsymbol{W}_r(t) = [\boldsymbol{w}_1(t), \ldots, \boldsymbol{w}_r(t)]$, coïncide avec l'espace signal engendré par la matrice de Vandermonde $\boldsymbol{V}^n$. Réciproquement, les $n-r$ dernières valeurs propres de $\mathbb{E}(\boldsymbol{C}_{xx}(t))$ étant toutes égales à $\sigma^2 \gamma(t)$, le **sous-espace mineur** de dimension $n-r$ de cette matrice, engendré par $\boldsymbol{W}_r^\perp(t) = [\boldsymbol{w}_{r+1}(t), \ldots, \boldsymbol{w}_n(t)]$, est appelé **espace bruit**. En pratique, on peut estimer les vecteurs $\boldsymbol{w}_1(t), \ldots, \boldsymbol{w}_n(t)$ en diagonalisant la matrice $\boldsymbol{C}_{xx}(t)$ (au lieu de son espérance mathématique) ou de façon équivalente en calculant la SVD de la matrice $\boldsymbol{X}'(t) = \boldsymbol{X}(t) \, \boldsymbol{\Gamma}(t)^{\frac{1}{2}}$.

Une fois que l'espace signal ou l'espace bruit a été estimé, il est possible d'en déduire une estimation des pôles z_k. Ainsi la méthode MUSIC [Schmidt, 1986] est basée sur l'espace bruit et consiste à chercher les solutions de l'équation $\|\boldsymbol{W}_r^\perp(t)^H \boldsymbol{v}(z)\|^2 = 0$, dont on sait qu'elles sont au nombre de r et égales aux pôles z_k. La méthode ESPRIT [Roy et Kailath, 1989] est basée sur l'espace signal et exploite la propriété d'**invariance rotationnelle** de la matrice de Vandermonde $\boldsymbol{V}^n$, qui s'énonce de la façon suivante : $\boldsymbol{V}_\uparrow^n = \boldsymbol{V}_\downarrow^n \boldsymbol{D}_r$, où $\boldsymbol{D}_r = \text{diag}(z_1, \ldots, z_r)$, $\boldsymbol{V}_\downarrow^n$ est la matrice extraite de $\boldsymbol{V}^n$ en supprimant sa dernière ligne et $\boldsymbol{V}_\uparrow^n$ la matrice extraite de $\boldsymbol{V}^n$ en supprimant sa première ligne. Puisque $\boldsymbol{W}_r(t)$ et $\boldsymbol{V}^n$ engendrent le même sous-espace, il existe une matrice inversible $\boldsymbol{G}_r(t)$ de dimensions $r \times r$ telle que $\boldsymbol{V}^n = \boldsymbol{W}_r(t) \, \boldsymbol{G}_r(t)$. Par substitution, on montre alors que $\boldsymbol{W}_r(t)$ vérifie une propriété d'invariance rotationnelle semblable à celle de $\boldsymbol{V}^n$:

$$\boldsymbol{W}_r(t)_\uparrow = \boldsymbol{W}_r(t)_\downarrow \, \boldsymbol{\Phi}_r(t) \tag{II.4}$$

où la matrice $\boldsymbol{\Phi}_r(t)$, appelée **matrice spectrale**, est définie par sa décomposition en valeurs propres : $\boldsymbol{\Phi}_r(t) = \boldsymbol{G}_r(t) \, \boldsymbol{D}_r \, \boldsymbol{G}_r(t)^{-1}$. La méthode ESPRIT s'appuie sur cette propriété remarquable pour estimer les pôles du signal. Elle se compose des trois étapes suivantes :

1. Calcul de la base $\boldsymbol{W}_r(t)$ du sous-espace principal de la matrice $\boldsymbol{C}_{xx}(t)$,
2. Calcul de la matrice spectrale $\boldsymbol{\Phi}_r(t) = (\boldsymbol{W}_r(t)_\downarrow)^\dagger \boldsymbol{W}_r(t)_\uparrow$ (où le symbole † désigne la pseudo-inverse),
3. Calcul des valeurs propres $\hat{z}_k(t)$ de la matrice $\boldsymbol{\Phi}_r(t)$, qui définissent les estimateurs des pôles z_k à l'instant t.

Bien sûr les performances des méthodes MUSIC et ESPRIT dépendent de l'estimation des bases $\boldsymbol{W}_r^\perp(t)$ et $\boldsymbol{W}_r(t)$ et donc du choix des coefficients γ_τ. Dans le cas d'une **fenêtre rectangulaire** de longueur $l \in \mathbb{N}$ ($\gamma_\tau = 1_{[0,l-1]}(\tau)$), l'horizon d'observation du signal est de longueur $N = n + l - 1$. Les performances des méthodes MUSIC et ESPRIT ont alors été analysées dans la littérature en rapportant les variances des estimateurs à leurs bornes de Cramér-Rao, sous des hypothèses asymptotiques : $N \to +\infty$ (dans le cas où tous les pôles sont sur le cercle unité) ou rapport signal sur bruit (RSB) élevé (RSB $\to +\infty$). Il a ainsi été démontré qu'à N fixé, le choix optimal des paramètres n et l est $n = 2l$ ou $l = 2n$, et que pour ces valeurs les

deux méthodes atteignent une efficacité asymptotique voisine de 1, avec un léger avantage pour la méthode ESPRIT [Stoica et Söderström, 1991, Eriksson *et al.*, 1993].

Une fois que les pôles complexes ont été estimés, on peut en déduire une estimation des amplitudes complexes et de la variance du bruit. En effet, on démontre que sur l'horizon d'observation $[t, t+n-1]$ l'estimation des amplitudes complexes au sens du maximum de vraisemblance est équivalente à utiliser la méthode des moindres carrés [Badeau, 2005] : $\hat{\boldsymbol{\alpha}}(t) = \boldsymbol{V}^{n\dagger}\boldsymbol{x}(t)$, où $\hat{\boldsymbol{\alpha}}(t)$ est un estimateur du vecteur $\boldsymbol{D}_r^t\,[\alpha_1, \ldots, \alpha_r]^T$, dont les coefficients sont notés $\hat{\alpha}_1(t), \ldots, \hat{\alpha}_r(t)$. De même, la variance du bruit est estimée comme la puissance du résiduel : $\hat{\sigma}^2(t) = \frac{1}{n}\|\boldsymbol{x}(t) - \boldsymbol{V}^n\,\hat{\boldsymbol{\alpha}}(t)\|_2^2$.

Enfin, nous n'avons encore rien dit de l'estimation de l'ordre du modèle, que nous avons supposé connu jusqu'à présent. De nombreuses méthodes ont été proposées dans la littérature pour estimer r, dont les plus classiques sont issues de la théorie de l'information [Marcos, 1998]. En pratique, ces méthodes sont relativement satisfaisantes pour traiter des signaux qui vérifient bien le modèle de signal, mais leurs performances s'effondrent quand ce modèle est moins bien vérifié, en particulier quand le bruit est coloré. Dans la section III.2 page 8, nous présentons une nouvelle méthode d'estimation de l'ordre du modèle, plus robuste au bruit.

II.2 Analyse temps-fréquence (chapitre IV)

Dans la section II.1, nous avons défini une suite d'estimateurs $\hat{z}_k(t)$ des pôles complexes et une suite d'estimateurs $\hat{\alpha}_k(t)$ des amplitudes complexes du modèle ESM. Cette dépendance au temps t offre des perspectives intéressantes, soit pour affiner progressivement l'estimation (si les coefficients γ_τ définissent une fenêtre à mémoire longue), soit pour suivre d'éventuelles variations des paramètres du modèle (dans le cas d'une fenêtre à mémoire courte). Cependant, l'application directe de la méthode ESPRIT à chaque instant t induirait une complexité algorithmique prohibitive, même s'il est toujours possible d'optimiser les calculs (section III.3 page 9). Afin de réduire la complexité, on fait appel à des algorithmes adaptatifs, qui effectuent une poursuite rapide des paramètres. On utilise dans un premier temps un algorithme de **poursuite de sous-espace**, qui remplace le calcul exact de $\boldsymbol{W}_r(t)$ par un calcul récursif approché (section IV.1 page 11). Les étapes suivantes de la méthode ESPRIT sont ensuite implémentées de façon adaptative (sections IV.2 et IV.3 page 16).

II.3 Applications de l'analyse à haute résolution (chapitre V)

Diverses applications musicales des techniques présentées dans les chapitres III et IV ont été abordées dans le cadre de ma thèse, puis dans le cadre du projet DESAM de l'Agence Nationale de la Recherche (ANR) et de collaborations avec des laboratoires français et étrangers. La plupart de ces travaux se sont appuyés sur le développement d'un système complet d'analyse / synthèse, mettant en œuvre les algorithmes de poursuite présentés dans le chapitre IV et un certain nombre de prétraitements destinés à améliorer leurs performances (section V.1 page 17). Ce système a été par exemple utilisé pour calculer une représentation temps-fréquence à haute résolution du signal, pour séparer les composantes sinusoïdales de la composante de bruit ou pour créer des effets sonores. Nous avons aussi abordé des applications du domaine de l'extraction d'informations musicales - *Music Information Retrieval* (MIR), comme l'estimation de hauteurs de notes de piano (section V.2 page 18) ou l'estimation du tempo musical (section V.3 page 19). L'analyse à haute résolution a également été utilisée pour valider un modèle physique d'instrument de musique (section V.4 page 19) et dans le domaine du codage audio (section V.5 page 20). Enfin, nous avons abordé des applications dans des domaines autres que le traitement de signaux audio, comme l'estimation de canal en communications numériques (section V.6.1 page 21) ou la factorisation de tenseurs structurés (section V.6.2 page 21).

Chapitre III

Analyse spectrale

Dans ce chapitre, nous présentons trois variations autour de la méthode d'analyse spectrale à haute résolution résumée dans le chapitre II : une généralisation du modèle ESM (section III.1), une méthode d'estimation de l'ordre du modèle (section III.2) et une implémentation rapide de l'algorithme d'estimation (section III.3).

III.1 Modèle de mélange d'exponentielles complexes multipliées par des polynômes

Dans la section II.1 page 4, nous avons vu que si le signal $s(t)$ vérifie le modèle ESM à l'ordre $r < \min(n, t+1)$, alors la matrice $\boldsymbol{S}(t)$ de dimensions $n \times (t+1)$ est seulement de rang r. En revanche la réciproque n'est pas vraie : si la matrice $\boldsymbol{S}(t)$ est de rang r, le signal $s(t)$ ne vérifie pas nécessairement le modèle ESM. Nous avons démontré que $s(t)$ vérifie plus généralement un modèle que nous avons appelé *Polynomial Amplitude Complex Exponentials* (PACE) [Badeau *et al.*, 2006a] : $s(t) = \sum_{k=1}^{K} \alpha_k(t)\, {z_k}^t$, où $\alpha_k(t)$ est un polynôme de degré $M_k - 1$ avec $M_k \in \mathbb{N}^*$. L'ordre total du modèle PACE est alors défini comme $r = \sum_{k=1}^{K} M_k$. Tous les développements mathématiques précédents ont été généralisés au modèle PACE. Nous avons ainsi généralisé la factorisation (II.2) page 4 en utilisant des matrices que nous avons baptisées "Pascal-Vandermonde" et nous avons également généralisé la méthode ESPRIT à l'estimation du modèle PACE [Badeau *et al.*, 2006a]. En particulier nous avons démontré qu'en l'absence de bruit les pôles z_k sont des valeurs propres multiples de la matrice spectrale $\boldsymbol{\Phi}_r(t)$, de multiplicité M_k. Nous avons aussi démontré un résultat inattendu : en présence d'un bruit blanc, les pôles multiples sont éclatés en plusieurs pôles simples répartis de façon isotrope, formant les sommets d'un polygone régulier lorsque le RSB tend vers l'infini. Un estimateur naturel du pôle multiple z_k est donc obtenu en calculant la moyenne arithmétique des M_k pôles simples ainsi estimés. Nous avons également calculé analytiquement les bornes de Cramér-Rao pour les paramètres du modèle PACE [Badeau *et al.*, 2008a]. Cette étude a démontré qu'en pratique il n'était possible d'estimer ce modèle que pour des multiplicités M_k très faibles. Nous avons enfin étudié les performances asymptotiques de l'algorithme ESPRIT généralisé, sous une hypothèse de rapport signal sur bruit élevé [Badeau *et al.*, 2008c, Badeau *et al.*, 2007b]. Nous avons prouvé en particulier que les performances ne dépendent que de l'espace signal et pas du choix d'une base orthonormée particulière $\boldsymbol{W}_r(t)$ de ce sous-espace. De plus les estimateurs des pôles z_k et des coefficients des polynômes $\alpha_k(t)$ sont asymptotiquement non-biaisés et leurs variances asymptotiques ont été comparées aux bornes de Cramér-Rao. En particulier nous avons démontré que dans le cas d'une fenêtre rectangulaire de longueur l, sous une hypothèse de bruit blanc, si tous les pôles sont sur le cercle unité et si le RSB et les paramètres n et l tendent vers $+\infty$, alors l'efficacité des estimateurs des pôles simples est proche de 1. Plus précisément, l'efficacité asymptotique des estimateurs de tous les facteurs d'atténuation et toutes les fréquences est identique, indépendante des paramètres du modèle, et sa valeur optimale est égale à $9/8 = 1.125$, atteinte pour $n = 2l$ ou $l = 2n$. De même, l'efficacité asymptotique des estimateurs de toutes les amplitudes réelles et toutes les phases est identique, indépendante des paramètres du modèle, et sa valeur optimale est égale à $35/32 = 1.09375$, toujours atteinte pour $n = 2l$ ou $l = 2n$. Cependant, nos résultats de simulations numériques ont montré que les variances des estimateurs associés à un pôle multiple augmentent exponentiellement avec l'ordre du pôle. Cette observation confirme que l'estimation d'un modèle PACE est possible seulement pour des multiplicités très faibles. Par ailleurs, nous

avons cherché à savoir si le modèle PACE se prête à la représentation de signaux de musique [Badeau, 2005]. Pour cela, nous avons examiné trois exemples typiques de modulations d'amplitude ou de fréquence couramment rencontrées dans ce type de signaux : trémolo, vibrato et glissando. Dans les trois cas, l'algorithme ESPRIT a permis de représenter les signaux avec une grande précision, en disposant des pôles simples le long d'une courbe dans le plan complexe. Or comme nous l'avons démontré dans la référence [Badeau *et al.*, 2006a], les pôles multiples du modèle PACE ont plutôt tendance à se retrouver asymptotiquement dispersés de façon homogène et isotrope en présence de bruit, ce qui permet de les identifier comme les sommets d'un polygone régulier. Cette distribution spatiale ne correspond donc pas à la répartition curviligne observée dans les modulations musicales. Les tests que nous avons effectués confirment cet argument : nous avons pu observer qu'à nombre égal de paramètres, les signaux musicaux sont presque toujours mieux représentés avec des pôles simples qu'avec des pôles multiples (en terme d'erreur quadratique moyenne), sauf sur des horizons temporels très courts (de l'ordre de 4 millisecondes). Ainsi, dans le cadre du traitement de signaux de musique, nous avons préféré nous restreindre à l'utilisation du modèle ESM.

III.2 Méthode d'estimation de l'ordre du modèle

La détermination de l'ordre du modèle ESM est une étape essentielle du processus d'estimation, car elle conditionne toute la suite de l'analyse à haute résolution. Les performances des critères usuels de sélection de l'ordre s'avérant souvent décevantes dans le cas des signaux audio (*cf.* section II.1 page 6), nous avons proposé une nouvelle méthode, plus robuste, pour estimer l'ordre du modèle [Badeau *et al.*, 2004a, Badeau *et al.*, 2006b]. Cette méthode s'appuie sur l'analyse de l'erreur d'estimation des pôles que l'on commet lorsqu'on applique la méthode ESPRIT avec un ordre de modèle p erroné ($p \neq r$). On démontre ainsi qu'en l'absence de bruit, si $p > r$, les pôles z_k présents dans le signal font tous partie des valeurs propres de la matrice $\boldsymbol{\Phi}_p(t)$. En revanche, si $p < r$, pour toute valeur propre $\hat{z}(t)$ de $\boldsymbol{\Phi}_p(t)$, il existe un pôle z_k présent dans le signal tel que

$$|\hat{z}(t) - z_k| \leq \kappa_2(t) \, \|\boldsymbol{E}_p(t)\|_2$$

où $\kappa_2(t)$ est une constante indépendante de p, $\boldsymbol{E}_p(t) = \boldsymbol{W}_p(t)_{\uparrow} - \boldsymbol{W}_p(t)_{\downarrow} \, \boldsymbol{\Phi}_p(t)$ et $\|.\|_2$ désigne la norme spectrale d'une matrice, définie comme sa plus grande valeur singulière. Il est intéressant de constater que la borne d'erreur $\|\boldsymbol{E}_p(t)\|_2$ peut être calculée numériquement sans connaitre l'ordre exact du modèle.

Le critère *ESTimation ERror* (ESTER) est alors défini par la fonction $p \mapsto \frac{1}{\|\boldsymbol{E}_p(t)\|_2^2}$. On démontre qu'il est toujours supérieur ou égal à 1 et en l'absence de bruit il diverge pour $p = r$. En présence de bruit, nous avons montré que l'ordre du modèle peut être estimé comme la plus grande valeur p dans un intervalle $[0, \ldots, p_{\max}]$ pour laquelle le critère atteint un maximum local supérieur à une certaine fraction de son maximum global [Badeau *et al.*, 2006b]. Le calcul direct de ce critère étant assez coûteux (sa complexité[1] est égale à $\frac{2}{3} n p_{\max}^3$), nous avons proposé un algorithme rapide pour le calculer récursivement, dont la complexité est seulement de $3np_{\max}^2$. Nous avons également montré la pertinence de ce critère en tant que borne d'erreur et les performances de la méthode ESTER ont été illustrées sur des signaux synthétiques et sur des signaux de piano. Nous avons ainsi observé que cette méthode est plus robuste que certains critères usuels de théorie de l'information. De plus, les bornes d'erreur peuvent être utilisées pour quantifier l'adéquation d'un éventuel ordre de modélisation inférieur, ce qui offre des perspectives intéressantes pour une application de codage. Bien qu'il ait été spécifiquement conçu pour la méthode ESPRIT, le critère ESTER peut être utilisé avec n'importe quelle méthode HR. Il a également été généralisé au modèle PACE [Badeau, 2005]. Dans la littérature, le critère ESTER a d'abord été cité comme méthode de référence sur l'estimation de l'ordre du modèle ESM [Reissell et Pai, 2007a, Haddad et Noga, 2007, Nechaev et Zotov, 2008, Zhang *et al.*, 2008, Christensen et Jakobsson, 2009] et il a été utilisé dans diverses applications de l'analyse HR. Il a notamment été proposé pour une application de contrôle de l'état de systèmes d'embrayage [Ompusunggu *et al.*, 2009] ou pour blanchir du bruit en présence de sinusoïdes [Da Costa *et al.*, 2009a]. Il a également été utilisé pour une application de séparation harmoniques / bruit [Gillet et Richard, 2008], pour l'analyse modale de structures [Ege *et al.*, 2009b], pour l'estimation de hauteur de notes de piano [Emiya *et al.*, 2007b] et pour étudier les phénomènes de sympathie dans la harpe de concert [Le Carrou *et al.*, 2009]. Il a enfin inspiré de nouvelles méthodes d'estimation de l'ordre basées sur des principes similaires [Papy *et al.*, 2007, Jakobsson *et al.*, 2007, Christensen *et al.*, 2007, Christensen *et al.*, 2009].

[1] Dans ce mémoire les complexités algorithmiques sont quantifiées en *multiply / accumulate* (MAC), opération arithmétique composée d'une multiplication et d'une addition.

III.3 Algorithme rapide d'estimation du modèle

Un inconvénient de la méthode d'estimation présentée dans la section II.1 page 4 est son coût élevé. On montre en effet que la complexité de l'estimation de l'espace signal est $O(Nn^2)$ (dans le cas d'une fenêtre rectangulaire de longueur $l = N - n + 1$), celle du calcul de la matrice spectrale est $O(nr^2)$, celle du calcul des pôles complexes est $O(r^3)$ et celle de l'estimation des amplitudes complexes est $O(nr^2)$. Nous avons donc développé des algorithmes rapides pour diminuer ce coût [Badeau, 2005]. Toutes les étapes ont été optimisées, à commencer par l'estimation de l'espace signal, qui est l'étape la plus coûteuse, grâce à la méthode des puissances itérées (section III.3.1). Nous avons également proposé des méthodes originales de calcul de la matrice spectrale et d'estimation des amplitudes. La complexité globale de l'algorithme obtenu est $O(Nr\log_2(N)+nr^2)$, au lieu de $O(Nn^2)$ dans le cas d'une implémentation directe. Ce travail a constitué un premier pas vers le développement d'algorithmes adaptatifs, qui seront présentés dans le chapitre IV. Nous résumons ci-dessous le principe de la méthode des puissances itérées et de l'itération bi-orthogonale, dont le but est d'estimer l'espace signal.

III.3.1 Méthode des puissances itérées

La **méthode des puissances itérées** est une technique classique de calcul des r principaux vecteurs propres ou plus généralement d'une base $\boldsymbol{W}_r(t)$ du sous-espace principal d'une matrice (semi-)définie positive $\boldsymbol{C}_{xx}(t)$ de dimensions $n \times n$ [Golub et Van Loan, 1996]. Elle alterne les opérations suivantes (où i est l'indice d'itération) :

1) $\boldsymbol{C}_{xy}(i) = \boldsymbol{C}_{xx}(t)\boldsymbol{W}_r(i-1)$
2) $\boldsymbol{W}_r(i)\,\boldsymbol{R}(i) = \boldsymbol{C}_{xy}(i)$

La première étape est un produit matriciel de complexité n^2r et la seconde est une orthonormalisation de la matrice $\boldsymbol{C}_{xy}(i)$, de complexité nr^2. Si $\boldsymbol{C}_{xx}(t)$ satisfait la relation (II.3) page 5 avec $\gamma_\tau = 1_{[0,l-1]}(\tau)$ (fenêtre rectangulaire de longueur l) et si la matrice $\boldsymbol{X}(t)$ présente une structure Toeplitz ou Hankel, l'étape 1) peut être optimisée en utilisant des produits matriciels rapides, basés sur l'algorithme de transformée de Fourier rapide (sa complexité est alors réduite à $4Nr\log_2(N)$). Soient $\lambda_1(t) \geq \ldots \geq \lambda_n(t) \geq 0$ les valeurs propres de $\boldsymbol{C}_{xx}(t)$ rangées par ordre décroissant et $\mathcal{N}(t)$ le sous-espace mineur de dimension $n - r$ de $\boldsymbol{C}_{xx}(t)$. Les propriétés de convergence de cet algorithme sont bien connues : si $\lambda_r(t) > \lambda_{r+1}(t)$ et si la matrice initiale $\boldsymbol{W}_r(0)$ vérifie $\mathrm{Im}(\boldsymbol{W}_r(0)) \cap \mathcal{N}(t) = \{0\}$, alors l'espace image de la matrice $\boldsymbol{W}_r(i)$ converge vers le sous-espace principal de dimension r de la matrice $\boldsymbol{C}_{xx}(t)$. De plus la vitesse de convergence est linéaire, le taux de convergence étant égal à $\lambda_{r+1}(t)/\lambda_r(t)$. Dans tous les cas numériques usuels, les deux conditions assurant la convergence de l'algorithme sont vérifiées. Si de plus la seconde étape est implémentée par une factorisation orthogonale-triangulaire (la matrice $\boldsymbol{R}(t)$ est triangulaire supérieure), alors $\boldsymbol{W}_r(i)$ converge vers la base orthonormée des r principaux vecteurs propres de la matrice $\boldsymbol{C}_{xx}(t)$ et $\boldsymbol{R}(i)$ converge vers une matrice diagonale contenant les valeurs propres correspondantes. L'algorithme prend alors le nom d'**itération orthogonale**.

III.3.2 Itération bi-orthogonale

L'**itération bi-orthogonale** est une variante de la méthode d'itération orthogonale, permettant de calculer les r principaux vecteurs singuliers à gauche et à droite d'une matrice $\boldsymbol{X}'(t)$ quelconque de dimensions $n \times l$ [Golub et Van Loan, 1996]. Elle effectue récursivement les opérations suivantes (où i est l'indice d'itération) :

1) $\boldsymbol{Y}(i)^H = \boldsymbol{X}'(t)^H\boldsymbol{W}_r(i-1)$
2) $\boldsymbol{U}_r(i)\boldsymbol{R}_Y(i) = \boldsymbol{Y}(i)^H$
3) $\boldsymbol{Z}(i) = \boldsymbol{X}'(t)\,\boldsymbol{U}_r(i)$
4) $\boldsymbol{W}_r(i)\boldsymbol{R}_Z(i) = \boldsymbol{Z}(i)$

Les étapes 1) et 3) sont des produits matriciels de complexité nlr et les étapes 2) et 4) consistent à calculer la factorisation orthogonale-triangulaire des matrices $\boldsymbol{Y}(i)^H$ et $\boldsymbol{Z}(i)$, dont la complexité est respectivement lr^2 et nr^2. Si $\boldsymbol{X}'(t) = \boldsymbol{X}(t)\,\boldsymbol{\Gamma}(t)^{\frac{1}{2}}$, où $\boldsymbol{X}(t)$ présente une structure Toeplitz ou Hankel, les étapes 1) et 3) peuvent être optimisées avec des produits matriciels rapides (leur complexité est alors réduite à $4Nr\log_2(N)$). Soient $\sigma_1(t) \geq \ldots \geq \sigma_{\min(n,l)}(t) \geq 0$ les valeurs singulières de $\boldsymbol{X}'(t)$ rangées par ordre décroissant, et $\mathcal{N}(t)$ le sous-espace mineur de dimension $n - r$ de $\boldsymbol{C}_{xx}(t) = \boldsymbol{X}'(t)\,\boldsymbol{X}'(t)^H$. Les propriétés de convergence sont similaires à celles de l'itération orthogonale : si $\sigma_r(t) > \sigma_{r+1}(t)$ et si $\mathrm{Im}(\boldsymbol{W}_r(0)) \cap \mathcal{N}(t) = \{0\}$, alors $\boldsymbol{W}_r(i)$ et $\boldsymbol{U}_r(i)$ convergent vers les bases orthonormées des r principaux vecteurs singuliers à gauche et à droite de $\boldsymbol{X}'(t)$. De même, $\boldsymbol{R}_Y(i)$ et $\boldsymbol{R}_Z(i)$ convergent toutes deux vers une matrice diagonale contenant les valeurs singulières correspondantes. Enfin, la vitesse de convergence est linéaire, le taux de convergence étant égal à $\sigma^2_{r+1}(t)/\sigma^2_r(t)$.

Chapitre IV

Analyse temps-fréquence

Dans le chapitre précédent, nous nous sommes intéressés à l'estimation d'un mélange de composantes fréquentielles dont les paramètres de fréquence, d'atténuation, d'amplitude et de phase sont supposés constants, indépendants du temps. En réalité, un signal de musique ne peut être modélisé de la sorte que sur des échelles de temps suffisamment courtes (de l'ordre de quelques dizaines de millisecondes). Sur des échelles plus longues, il est nécessaire de tenir compte d'éventuelles variations de ces paramètres. Nous supposerons néanmoins que ces variations sont suffisamment lentes pour qu'à chaque instant t, les valeurs courantes des pôles $z_k(t)$ puissent être correctement estimées à partir de la matrice $\boldsymbol{C}_{xx}(t)$. L'application directe de la méthode ESPRIT à chaque instant t induirait cependant une complexité algorithmique prohibitive ($O(tn^2)$ pour le seul calcul de $\boldsymbol{C}_{xx}(t)$), même s'il est toujours possible d'optimiser les calculs, comme on l'a vu dans la section III.3. Afin de réduire la complexité, on fait appel à des algorithmes adaptatifs, qui effectuent une poursuite rapide des paramètres.

Pour mémoire, la première méthode de poursuite des pôles que nous avons publiée utilisait un algorithme de moindres carrés adaptatif de faible complexité ($O(r^3)$ à chaque instant t), qui consistait à mettre à jour un polynôme prédicteur et à extraire ses racines [David *et al.*, 2002]. Cet algorithme était initialisé une première fois à l'aide d'une variante de ESPRIT baptisée *Matrix Pencil* [Hua et Sarkar, 1990], coûteuse mais précise, puis il était réinitialisé dès que l'erreur de prédiction devenait trop élevée. Dans le cadre de l'application aux signaux de musique, la faible robustesse au bruit de l'algorithme de moindres carrés adaptatifs entraînait des réinitialisations trop fréquentes, ce qui en réduisait fortement l'intérêt pratique. Par la suite nous avons donc cherché à développer des algorithmes de poursuite rapides et plus robustes. L'approche suivie a consisté à optimiser tour à tour les trois étapes de la méthode ESPRIT.

Tout d'abord, on remarque que si les poids γ_τ sont bien choisis, la complexité peut être fortement réduite, en calculant récursivement la matrice $\boldsymbol{C}_{xx}(t)$ définie dans l'équation (II.3) page 5 :

- La **fenêtre rectangulaire** de longueur $l \in \mathbb{N}^*$ correspond à $\gamma_\tau = 1_{[0,l-1]}(\tau)$. Ce choix conduit à la récurrence $\boldsymbol{C}_{xx}(t) = \boldsymbol{C}_{xx}(t-1) + \boldsymbol{x}(t)\,\boldsymbol{x}(t)^H - \boldsymbol{x}(t-l)\,\boldsymbol{x}(t-l)^H$ (mise à jour de rang 2).
- La **fenêtre exponentielle** correspond à $\gamma_\tau = \beta^\tau$, où $\beta \in]0,1[$ est un facteur d'oubli. Ce choix conduit à la récurrence $\boldsymbol{C}_{xx}(t) = \beta\boldsymbol{C}_{xx}(t-1) + \boldsymbol{x}(t)\,\boldsymbol{x}(t)^H$ (mise à jour de rang 1).
- La **fenêtre hybride** correspond à $\gamma_\tau = 1_{[0,l-1]}(\tau)\beta^\tau$, avec les paramètres l et τ précédents. Ce choix conduit à la récurrence $\boldsymbol{C}_{xx}(t) = \beta\boldsymbol{C}_{xx}(t-1) + \boldsymbol{x}(t)\,\boldsymbol{x}(t)^H - \beta^l\boldsymbol{x}(t-l)\,\boldsymbol{x}(t-l)^H$ (mise à jour de rang 2).

La complexité du calcul de $\boldsymbol{C}_{xx}(t)$ est alors réduite à $2n^2$ (pour une mise à jour de rang 2) ou n^2 (pour une mise à jour de rang 1).

Il se trouve néanmoins que même si la matrice $\boldsymbol{C}_{xx}(t)$ satisfait l'une de ces trois récurrences, le coût du calcul de son sous-espace principal n'en est pas moins de $O(n^2r)$, ce qui reste prohibitif pour beaucoup d'applications, et il n'existe pas de méthode simple permettant de calculer récursivement $\boldsymbol{W}_r(t)$. Afin de réduire le coût de calcul, on utilise un algorithme de **poursuite de sous-espace**, qui remplace le calcul exact de $\boldsymbol{W}_r(t)$ par un calcul récursif en introduisant des approximations basées sur des hypothèses de régularité temporelle (section IV.1). Les étapes suivantes de la méthode ESPRIT sont ensuite également implémentées de façon adaptative (sections IV.2 et IV.3).

IV.1 Poursuite de sous-espace

La littérature portant sur la poursuite de sous-espace est extrêmement riche et les approches proposées très diverses [Badeau *et al.*, 2008b]. Les algorithmes existants peuvent être classés selon différents critères :

Poursuite du sous-espace principal ou mineur : Certaines méthodes d'estimation paramétrique comme ESPRIT requièrent la poursuite du sous-espace principal de $\boldsymbol{C}_{xx}(t)$, alors que d'autres comme MUSIC sont basées sur le sous-espace mineur. En pratique, la poursuite du sous-espace mineur s'avère plus difficile que celle du sous-espace principal (la vitesse de convergence est souvent beaucoup plus faible et il est fréquent de rencontrer des problèmes de stabilité numérique).

Poursuite des vecteurs propres : Certains algorithmes suivent les vecteurs propres de la matrice $\boldsymbol{C}_{xx}(t)$ (ou de façon équivalente les vecteurs singuliers à gauche de la matrice $\boldsymbol{X}(t)\boldsymbol{\Gamma}(t)^{\frac{1}{2}}$), alors que d'autres se contentent de suivre une base quelconque du sous-espace que ces vecteurs engendrent. En effet, les méthodes MUSIC et ESPRIT peuvent être utilisées indifféremment dans les deux cas. En revanche, la poursuite des vecteurs propres/singuliers permet de faire facilement varier l'ordre r du modèle (en cas d'apparition ou de disparition d'une ou plusieurs composantes fréquentielles).

Poursuite des vecteurs singuliers à droite : Dans le cas de la poursuite des vecteurs singuliers de la matrice $\boldsymbol{X}(t)\boldsymbol{\Gamma}(t)^{\frac{1}{2}}$, certains algorithmes calculent les vecteurs singuliers à droite en plus des vecteurs singuliers à gauche, de façon à effectuer le suivi de la SVD complète de cette matrice.

Orthonormalité : Dans le cas où $\boldsymbol{W}_r(t)$ n'est pas une base de vecteurs propres ou singuliers, certains algorithmes garantissent son orthonormalité, d'autres non. En effet, la propriété d'orthonormalité n'est pas non plus requise par les méthodes MUSIC et ESPRIT (même si leurs performances peuvent être affectées par l'utilisation d'une matrice non orthonormée). En pratique, l'orthonormalisation de $\boldsymbol{W}_r(t)$ est préférable, car elle a tendance à renforcer la stabilité numérique des algorithmes de poursuite de sous-espace.

Forme de la fenêtre : La fenêtre rectangulaire est beaucoup moins couramment utilisée que la fenêtre exponentielle, car cette dernière aboutit généralement à des expressions analytiques plus simples et à des algorithmes de plus faible complexité. En revanche, nous avons remarqué que la fenêtre rectangulaire offre une meilleure capacité de poursuite que la fenêtre exponentielle (réponse plus rapide à de brusques variations des paramètres) [Badeau, 2005].

Complexité algorithmique : La complexité d'un algorithme de poursuite de sous-espace correspond au coût d'une itération (pour chaque instant t). On distingue les algorithmes de complexité quadratique en n ($O(n^2r)$ ou $O(n^2)$) et les algorithmes de complexité linéaire en n ($O(nr^2)$ ou $O(nr)$). La complexité d'un algorithme dépend (entre autres) de la complexité de la tâche qu'il accomplit : poursuite du sous-espace principal ou mineur, poursuite des vecteurs propres/singuliers ou d'une base quelconque, orthonormalisation de $\boldsymbol{W}_r(t)$, utilisation d'une fenêtre exponentielle ou rectangulaire/hybride.

Stabilité numérique : Dans le sens le plus large, un algorithme de poursuite de sous-espace est numériquement stable si la matrice $\boldsymbol{W}_r(t)$ reste bornée quand les données en entrée sont bornées. Dans le cas particulier de la poursuite d'une base orthonormée du sous-espace, une mesure courante de la stabilité numérique est la variation temporelle de l'erreur d'orthonormalité : $\eta(t) = \|\boldsymbol{I}_r - \boldsymbol{W}_r(t)^H\boldsymbol{W}_r(t)\|_2$. Si l'algorithme est initialisé aléatoirement avec une matrice non orthonormée, cette erreur doit tendre vers 0 (ou du moins converger en moyenne vers une valeur strictement positive qui dépend du bruit de quantification). Dans ce sens, la plupart des algorithmes de poursuite du sous-espace mineur existant se sont avérés numériquement instables (*cf.* [Badeau *et al.*, 2008b]).

Par rapport à la littérature existante dans le domaine, les recherches que nous avons entreprises ont suivi quatre directions :

- **Améliorer la capacité de poursuite** d'algorithmes conçus pour une fenêtre exponentielle, en proposant des implémentations alternatives utilisant une fenêtre rectangulaire ou une fenêtre hybride ;
- **Réduire la complexité** d'algorithmes conçus pour effectuer le suivi des vecteurs propres de la matrice de corrélation, en se contentant de suivre une base quelconque du sous-espace propre ;
- Modifier des algorithmes conçus pour suivre le sous-espace principal, afin d'**effectuer la poursuite du sous-espace mineur** ;
- **Améliorer la stabilité numérique** des algorithmes vis-à-vis de l'orthonormalité de la matrice $\boldsymbol{W}_r(t)$.

Nos travaux ont ainsi porté sur l'amélioration de quatre techniques de poursuite, dérivées de la méthode d'itération bi-orthogonale [Strobach, 1997a] (section IV.1.1), d'une approximation par projection [Yang, 1995]

TAB. IV.1 – Récapitulatif des propriétés des algorithmes de poursuite de sous-espace

Section	Algorithme	Sous-espace	Vecteurs singuliers	Ortho-normalité	Fenêtre	Coût
IV.1.1	Bi-SVD2	principal	à gauche	oui	exponentielle	nr^2
	SWASVD2	principal	à gauche & **à droite**	oui	**rectangulaire**	$(n+l)r^2$
IV.1.2	PAST	principal	non	non	exponentielle	$3nr$
	SW-PAST	principal	non	non	**rectangulaire**	$5nr$
IV.1.2	OPAST	principal	non	oui	exponentielle	$4nr$
	SW-OPAST	principal	non	oui	**rectangulaire**	$9nr$
IV.1.3	LORAF2	principal	à gauche	oui	exponentielle	nr^2
	NP2	principal	non	oui	exponentielle	nr^2
	FAPI	principal	non	oui	exponentielle **ou hybride**	$\boldsymbol{3nr}$ ou $\boldsymbol{6nr + 4lr}$
IV.1.4	SP2	principal	à gauche	oui	exponentielle	$4nr^2$
	YAST	pr. **ou mineur**	non	oui	exp. **ou hybride**	$\boldsymbol{O(nr)}$

(section IV.1.2), de la méthode des puissances itérées [Strobach, 1996, Hua *et al.*, 1999] (section IV.1.3) ou encore de la maximisation du critère de Rayleigh [Davila, 2000] (section IV.1.4). Les propriétés des algorithmes de la littérature et des algorithmes que nous avons proposés sont résumées dans la table IV.1, où nos contributions apparaissent en gras.

IV.1.1 Méthode d'itération bi-orthogonale

Nos premiers travaux dans le domaine de la poursuite de sous-espace ont été basés sur la méthode d'itération bi-orthogonale présentée dans la section III.3.2 page 9, qui effectue le calcul de la SVD complète de la matrice de données. En effet cette méthode est facilement transposable à un traitement adaptatif : il suffit de remplacer l'index d'itération i par le temps t, de façon à n'effectuer qu'une itération à chaque instant. On obtient ainsi l'algorithme suivant :

1) $\boldsymbol{Y}(t)^H = \boldsymbol{X}'(t)^H \boldsymbol{W}_r(t-1)$
2) $\boldsymbol{U}_r(t)\boldsymbol{R}_Y(t) = \boldsymbol{Y}(t)^H$
3) $\boldsymbol{Z}(t) = \boldsymbol{X}'(t)\,\boldsymbol{U}_r(t)$
4) $\boldsymbol{W}_r(t)\boldsymbol{R}_Z(t) = \boldsymbol{Z}(t)$

Cet algorithme effectue la poursuite des vecteurs singuliers à gauche et à droite de la matrice $\boldsymbol{X}'(t) = \boldsymbol{X}(t)\,\boldsymbol{\Gamma}(t)^{\frac{1}{2}}$, rangés respectivement dans les matrices $\boldsymbol{W}_r(t)$ et $\boldsymbol{U}_r(t)$.

Comme dans la section III.3.2, si la matrice $\boldsymbol{X}(t)$ présente une structure Toeplitz ou Hankel et si on utilise une fenêtre rectangulaire de longueur $l = N - n + 1$, les étapes 1) et 3) peuvent être optimisées en utilisant des produits matriciels rapides (leur complexité est alors réduite à $4Nr\log_2(N)$). Si l'on souhaite réduire davantage le coût de calcul, il faut introduire une approximation. En remarquant que la matrice $\boldsymbol{X}'(t)$ satisfait une récurrence simple, une astuce permet ainsi d'éviter le calcul explicite des matrices $\boldsymbol{Y}(t)^H$ et $\boldsymbol{Z}(t)$, et d'approcher l'étape 4) par un calcul récursif de la matrice $\boldsymbol{W}_r(t)$. Cette approche a été initialement proposée par Peter Strobach, qui a développé un algorithme adaptatif baptisé *Bi-iteration SVD* (Bi-SVD). La complexité de son implémentation Bi-SVD2 était seulement de nr^2 [Strobach, 1997a]. En revanche cet algorithme ne pouvait être utilisé qu'avec une fenêtre exponentielle et n'effectuait pas la poursuite de la matrice $\boldsymbol{U}_r(t)$. Un second algorithme utilisant une fenêtre rectangulaire avait également été proposé dans [Strobach, 1997b], mais celui-ci ne pouvait être appliqué qu'à une matrice $\boldsymbol{X}(t)$ carrée et structurée Hankel. Nous avons donc développé une nouvelle implémentation à fenêtre rectangulaire, applicable à une matrice $\boldsymbol{X}(t)$ de dimensions quelconques et pas nécessairement structurée. Cette implémentation calcule récursivement les deux matrices $\boldsymbol{U}_r(t)$ et $\boldsymbol{W}_r(t)$, et effectue donc le suivi de la SVD complète de la matrice $\boldsymbol{X}'(t)$, c'est pourquoi nous l'avons baptisée *Sliding Window Adaptive SVD* (SWASVD) [Badeau *et al.*, 2004b]. L'algorithme SWASVD2 ainsi obtenu possède une complexité égale à $(n+l)r^2$. La complexité des algorithmes Bi-SVD2 et SWASVD2 est bien linéaire en n, cependant elle reste quadratique en r. Notons que des implémentations de complexité linéaire en

r ont bien été proposées, mais impliquent soit une approximation supplémentaire dans le cas de l'algorithme Bi-SVD3 [Strobach, 1997a], soit une complexité algorithmique de constante multiplicative élevée dans le cas de l'algorithme SWASVD3 [Badeau *et al.*, 2004b].

Nous avons depuis développé des algorithmes de complexité plus faible, présentant de meilleurs performances en terme d'estimation de l'espace signal (sections IV.1.2 à IV.1.4). L'algorithme SWASVD2 reste néanmoins le seul capable d'effectuer le suivi de la SVD complète de la matrice $\boldsymbol{X}'(t)$. Dans la littérature, il a été cité comme algorithme de référence pour la poursuite de la SVD complète [Hoegaerts *et al.*, 2007, Du et Swamy, 2006], pour l'utilisation d'une fenêtre rectangulaire en poursuite de sous-espace [Toolan et Tufts, 2006, Willink, 2008, Slavnicu et Ciochina, 2006, Christensen et Jakobsson, 2009] et il a notamment été utilisé en séparation aveugle de sources [Bertrand et Moonen, 2009, Zhang *et al.*, 2010] et pour développer une implémentation adaptative de la décomposition tensorielle *PARAllel FACtor analysis* (PARAFAC) [Nion et Sidiropoulos, 2009]. Il a également inspiré des algorithmes de poursuite basés sur des principes similaires [Ouyang et Hua, 2005, Wang et Ouyang, 2009, Strobach, 2010].

IV.1.2 Approximation par projection

Comme nous l'avons mentionné plus haut, la méthode ESPRIT ne nécessite en fait pas le calcul de la SVD complète de la matrice de données $\boldsymbol{X}'(t)$, mais seulement une estimation du sous-espace engendré par ses r principaux vecteurs singuliers à gauche. En théorie, n'importe quelle base $\boldsymbol{W}_r(t)$ de ce sous-espace peut convenir ; l'orthonormalité de $\boldsymbol{W}_r(t)$ n'est même pas nécessaire. Puisque l'estimation de $\boldsymbol{W}_r(t)$ peut être posée en des termes moins contraignants que le calcul de la SVD complète de $\boldsymbol{X}'(t)$, on peut s'attendre à trouver des algorithmes de complexité plus faible que ceux basés sur l'itération bi-orthogonale. En fait il a été montré dans la référence [Strobach, 2009b] qu'il existe une borne inférieure de complexité, égale à $3nr$.

Le premier algorithme de poursuite de sous-espace ayant atteint cette borne est l'algorithme *Projection Approximation Subspace Tracking* (PAST) [Yang, 1995], originellement conçu pour une fenêtre exponentielle. Cet algorithme est dérivé d'une descente de gradient et se résume en trois étapes :

1) $\boldsymbol{C}_{xy}(t) = \boldsymbol{C}_{xx}(t)\,\boldsymbol{W}_r(t-1)$
2) $\boldsymbol{C}_{yy}(t) = \boldsymbol{W}_r(t-1)^H \boldsymbol{C}_{xx}(t)\,\boldsymbol{W}_r(t-1)$
3) $\boldsymbol{W}_r(t) = \boldsymbol{C}_{xy}(t)\boldsymbol{C}_{yy}(t)^{-1}$

La matrice $\boldsymbol{C}_{xx}(t)$ étant définie récursivement, si on suppose que la base $\boldsymbol{W}_r(t)$ varie lentement au cours du temps (hypothèse baptisée approximation par projection, ou *projection approximation*), on montre que les étapes 1) et 2) peuvent être approchées par une récurrence sur les matrices $\boldsymbol{C}_{xy}(t)$ et $\boldsymbol{C}_{yy}(t)$. Le lemme d'inversion matricielle [Horn et Johnson, 1985] prouve alors que la matrice $\boldsymbol{C}_{yy}(t)^{-1}$ satisfait une relation de récurrence. Une simple substitution dans l'étape 3) aboutit enfin à un calcul récursif de la base $\boldsymbol{W}_r(t)$, dont la complexité est seulement de $3nr$. Notons qu'une version déflationniste de PAST a également été proposée dans [Yang, 1995], permettant l'estimation des vecteurs propres de $\boldsymbol{C}_{xx}(t)$ mais présentant une complexité accrue ($4nr$), et de moins bonnes performances en terme d'estimation de l'espace signal. Par ailleurs, une variante de PAST garantissant l'orthonormalité de la matrice $\boldsymbol{W}_r(t)$ a été présentée dans la référence [Abed-Meraim *et al.*, 2000]. L'algorithme *Orthonormal PAST* (OPAST) consiste simplement à orthonormaliser la matrice $\boldsymbol{W}_r(t)$ calculée par PAST, grâce à une astuce permettant un calcul rapide. Sa complexité est égale à $4nr$. Il a été montré dans [Abed-Meraim *et al.*, 2000] que l'étape d'orthonormalisation stabilise PAST et améliore sa vitesse de convergence.

Dans la référence [Badeau *et al.*, 2003a], nous avons introduit des implémentations à fenêtre rectangulaire des algorithmes PAST et OPAST, baptisées respectivement *Sliding Window PAST* (SW-PAST) et *Sliding Window OPAST* (SW-OPAST) et de complexités respectives $5nr$ et $9nr$. Ces algorithmes améliorent les capacités de poursuite par rapport à une fenêtre exponentielle, au prix d'un accroissement du coût de calcul.

IV.1.3 Méthode des puissances itérées

Comme nous l'avons mentionné dans la section précédente, PAST et ses variantes ont été les premiers algorithmes de poursuite de sous-espace permettant d'atteindre (ou d'approcher) la borne inférieure de complexité. Nous avons donc par la suite orienté nos recherches vers le développement d'algorithmes de complexité équivalente, mais présentant des performances accrues en terme de vitesse de convergence et de stabilité numérique. Nos premiers efforts se sont portés sur la méthode des puissances itérées (*power iteration*). En effet cette méthode, présentée dans la section III.3.1 page 9, est facilement transposable à un traitement adaptatif

des données. Comme pour l'algorithme d'itération bi-orthogonale, il suffit de remplacer l'index d'itération i par le temps t, de façon à n'effectuer qu'une itération à chaque instant.

On obtient ainsi l'algorithme suivant :

1) $\boldsymbol{C}_{xy}(t) = \boldsymbol{C}_{xx}(t)\boldsymbol{W}_r(t-1)$

2) $\boldsymbol{W}_r(t)\,\boldsymbol{R}(t) = \boldsymbol{C}_{xy}(t)$

Si la seconde étape est implémentée par une factorisation orthogonale-triangulaire (la matrice $\boldsymbol{R}(t)$ est triangulaire supérieure), l'algorithme effectue une poursuite des vecteurs propres de la matrice $\boldsymbol{C}_{xx}(t)$. Il prend alors le nom d'**itération séquentielle** (*sequential iteration*) [Badeau *et al.*, 2002, David *et al.*, 2003]. Dans ce cas, la matrice $\boldsymbol{W}_r(t)$ peut varier au cours du temps, même si les paramètres du modèle ESM restent constants [Badeau, 2005]. Si en revanche la seconde étape est implémentée par une factorisation "polaire" (la matrice $\boldsymbol{R}(t)$ est symétrique définie positive), l'algorithme prend le nom de **méthode des puissances naturelles** - *Natural Power method* (NP) [Hua *et al.*, 1999]. Dans ce cas, l'algorithme n'estime plus les vecteurs propres de $\boldsymbol{C}_{xx}(t)$, mais une base orthonormée quelconque de l'espace signal. En contrepartie, si les paramètres du modèle ESM sont constants, alors on montre que la matrice $\boldsymbol{W}_r(t)$ l'est également [Badeau *et al.*, 2005a]. L'algorithme d'itération séquentielle a notamment été utilisé en traitement du signal audio pour effectuer le suivi des paramètres du modèle ESM [Boyer et Abed-Meraim, 2004, Lu et Brown, 2007] et pour séparer les composantes sinusoïdales de la composante de bruit [Gillet et Richard, 2008, Alonso Arevalo *et al.*, 2007].

Comme dans la section III.3.1, la première étape peut être optimisée en utilisant des produits matriciels rapides si la matrice $\boldsymbol{X}(t)$ présente une structure Toeplitz ou Hankel et si on utilise une fenêtre rectangulaire de longueur $l = N-n+1$ (sa complexité est alors réduite à $4Nr\log_2(N)$). Si l'on souhaite réduire davantage le coût de calcul, il faut introduire une approximation. La matrice $\boldsymbol{C}_{xx}(t)$ étant définie récursivement, si l'on suppose que le projecteur sur l'espace signal varie lentement au cours du temps, alors l'étape 1) peut être approchée par une récurrence sur la matrice $\boldsymbol{C}_{xy}(t)$. La seconde étape peut être ensuite implémentée de façon récursive, soit en utilisant des rotations de Givens dans le cas d'une factorisation orthogonale-triangulaire (on obtient ainsi l'implémentation LORAF2 de l'algorithme *LOw Rank Adaptive Filters* (LORAF) [Strobach, 1996]), soit en utilisant l'approximation par projection [Yang, 1995] dans le cas d'une factorisation polaire (on obtient ainsi l'algorithme NP2 [Hua *et al.*, 1999]). Il apparaît donc que les algorithmes LORAF2 et NP2 sont deux implémentations adaptatives différentes de la méthode des puissances itérées ; ils diffèrent seulement par le choix d'une base particulière $\boldsymbol{W}_r(t)$ de ce sous-espace.

Ces deux algorithmes ont une complexité égale à nr^2. Remarquons que des algorithmes de complexité $O(nr)$, respectivement nommés LORAF3 et NP3, ont également été proposés dans les références [Strobach, 1996] et [Hua *et al.*, 1999], mais LORAF3 utilise une approximation supplémentaire qui dégrade la performance de l'estimation et NP3 s'avère numériquement instable. Nous avons donc développé une nouvelle implémentation adaptative de la méthode des puissances naturelles, que nous avons baptisée *Fast Approximated Power Iteration* (FAPI) [Badeau *et al.*, 2005a]. Cette implémentation exploite le lemme d'inversion matricielle pour atteindre une complexité égale à $3nr$ dans le cas d'une fenêtre exponentielle et $6nr + 4lr$ dans le cas d'une fenêtre rectangulaire. Par ailleurs, il est démontré dans [Badeau *et al.*, 2005a] que PAST et OPAST (de même que NP2) constituent des approximations de FAPI. En effet, ces trois algorithmes s'appuient sur une hypothèse plus forte : ils ne supposent pas seulement que le projecteur sur l'espace signal varie lentement au cours du temps, mais que la base $\boldsymbol{W}_r(t)$ elle-même varie lentement au cours du temps. Cela explique leurs moins bonnes performances en terme d'estimation.

En conclusion, l'algorithme FAPI réunit les avantages de ses principaux concurrents : il est aussi rapide que PAST et aussi précis que LORAF2 en terme d'estimation de sous-espace. De plus, comme OPAST, NP2 et LORAF2, il calcule une base $\boldsymbol{W}_r(t)$ orthonormée et s'avère parfaitement stable numériquement. Des implémentations de FAPI ont été développées pour divers types de fenêtre : exponentielle [Badeau *et al.*, 2003d], rectangulaire [Badeau *et al.*, 2003c] et hybride [Badeau *et al.*, 2005a]. Dans la littérature, il a été cité comme méthode de référence en poursuite de sous-espace [Bartelmaos et Abed-Meraim, 2008, Doukopoulos et Moustakides, 2008, Sun *et al.*, 2007, Hendriks *et al.*, 2009, Christensen et Jakobsson, 2009] et remarqué pour ses bonnes performances [Guo *et al.*, 2006, Belkacemi et Marcos, 2006, Beau et Marcos, 2010]. Il a notamment été utilisé dans diverses applications comme la localisation adaptative de sources [Guo *et al.*, 2006, Valizadeh *et al.*, 2007], la formation de voies adaptative [Foutz et Spanias, 2007], le traitement espace/temps adaptatif - *Space-Time Adaptive Processing* (STAP) de signaux radars aériens [Marcos, 2009, Belkacemi et Marcos, 2006, Beau et Marcos, 2010] ou la séparation sinusoïdes / bruit [David *et al.*, 2006b, Alonso Arevalo *et al.*, 2003a]. Il a depuis inspiré d'autres algorithmes de poursuite de sous-espace, basés sur des principes similaires [Strobach, 2009a, Strobach, 2009c,

Strobach, 2009b, Zhang *et al.*, 2010].

IV.1.4 Maximisation du critère de Rayleigh

Tous les algorithmes de poursuite de sous-espace mentionnés ci-dessus, bien que reposant *a priori* sur des principes différents, peuvent être interprétés selon un paradigme commun. En effet dans le cas d'une fenêtre exponentielle, le sous-espace estimé à l'instant t est obtenu en ajoutant la direction $\boldsymbol{x}(t)$ au sous-espace de dimension r engendré par $\boldsymbol{W}_r(t-1)$, puis en retranchant une direction $\boldsymbol{v}(t)$ pour revenir à la dimension r. En terme d'estimation de sous-espace, ces algorithmes diffèrent donc seulement par la direction retranchée. Partant de ce point de vue, il paraît naturel de rechercher une direction $\boldsymbol{v}(t)$ permettant d'obtenir des performances optimales. Le premier algorithme exploitant cette idée a été proposé par Carlos Davila [Davila, 2000]. Il repose sur la maximisation du critère de Rayleigh, défini par la relation $J(\boldsymbol{W}_r(t)) = \text{trace}\left(\left(\boldsymbol{W}_r(t)^H \boldsymbol{C}_{xx}(t) \boldsymbol{W}_r(t)\right)\left(\boldsymbol{W}_r(t)^H \boldsymbol{W}_r(t)\right)^{-1}\right)$, dont on montre qu'il atteint sa valeur maximale lorsque $\boldsymbol{W}_r(t)$ engendre le sous-espace principal de la matrice $\boldsymbol{C}_{xx}(t)$. Pour éviter l'optimisation globale de ce critère, trop coûteuse, Carlos Davila a proposé de restreindre la recherche au sous-espace de dimension $r+1$ engendré par $\boldsymbol{x}(t)$ et $\boldsymbol{W}_r(t-1)$. Dans le paradigme exposé ci-dessus, cela revient à rechercher la direction $\boldsymbol{v}(t)$ telle que la matrice $\boldsymbol{W}_r(t)$ ainsi obtenue maximise le critère de Rayleigh.

L'algorithme obtenu, baptisé *Subspace Projection* (SP), a permis d'atteindre des performances inattendues en terme de vitesse de convergence, surclassant largement tous les algorithmes précédents, dont ceux basés sur les méthodes de bi-itération (Bi-SVD, SWASVD), d'approximation par projection (PAST, OPAST) et des puissances itérées (LORAF, NP, FAPI). En revanche, sa complexité algorithmique est plus élevée : n^2 dans le cas de données non structurées et $4nr^2$ pour une matrice de données structurée Toeplitz ou Hankel. Par ailleurs, il est conçu exclusivement pour une fenêtre exponentielle et effectue le suivi des principaux vecteurs propres de la matrice $\boldsymbol{C}_{xx}(t)$. Enfin, nous avons observé qu'il est moins stable que FAPI vis-à-vis de la propriété d'orthonormalité de la matrice $\boldsymbol{W}_r(t)$. Notre travail a donc porté sur le développement d'un algorithme présentant la même vitesse de convergence que l'implémentation SP2, mais de complexité plus faible, compatible avec l'usage d'une fenêtre rectangulaire ou hybride, capable de suivre le sous-espace mineur aussi bien que le sous-espace principal de $\boldsymbol{C}_{xx}(t)$ (afin de pouvoir utiliser des méthodes d'estimation basées sur l'espace bruit, comme MUSIC) et plus stable vis-à-vis de la propriété d'orthonormalité. Afin de réduire le coût de calcul, nous avons utilisé la même technique que dans FAPI : en renonçant au suivi des vecteurs propres et en utilisant le lemme d'inversion matricielle, nous avons ainsi atteint une complexité de $4nr$ dans le cas d'une fenêtre exponentielle et de $7nr$ dans le cas d'une fenêtre rectangulaire ou hybride [Badeau *et al.*, 2005b]. Pour effectuer le suivi du sous-espace mineur, il s'agit non plus de maximiser mais de minimiser le critère de Rayleigh vis-à-vis du vecteur $\boldsymbol{v}(t)$ [Badeau *et al.*, 2006c]. En revanche nous avons observé que le vecteur $\boldsymbol{x}(t)$ n'est plus dans ce cas un choix judicieux pour augmenter le sous-espace engendré par $\boldsymbol{W}_r(t-1)$. De meilleures performances ont été obtenues en utilisant un autre vecteur, obtenu par la méthode du gradient conjugué [Badeau *et al.*, 2007a]. Enfin, l'algorithme obtenu présentait des problèmes de stabilité numérique. Nous avons pu nous en affranchir en remplaçant la méthode d'orthonormalisation de la matrice $\boldsymbol{W}_r(t)$, inspirée de FAPI, par une méthode issue d'un autre algorithme baptisé *Fast Data Projection Method* (FDPM) [Doukopoulos et Moustakides, 2008], légèrement plus coûteuse mais garantissant une parfaite stabilité.

L'algorithme finalement publié dans la référence [Badeau *et al.*, 2008b] présente toutes les propriétés souhaitées : la remarquable vitesse de convergence de SP2, la possibilité d'utiliser les trois types de fenêtres, une complexité voisine de FAPI ($5nr$ pour une fenêtre exponentielle), la capacité de suivre le sous-espace mineur comme le sous-espace principal et la stabilité numérique de FDPM. Il a également fait l'objet d'un brevet [Badeau *et al.*, 2005d]. Nous l'avons baptisé *Yet Another Subspace Tracker* (YAST), en clin d'œil à l'abondance de publications dans le domaine de la poursuite de sous-espace. Dans la littérature, cet algorithme a été cité comme méthode de référence pour la poursuite de sous-espace [Doukopoulos et Moustakides, 2008, Misono et Yamada, 2008, Perry et Wolfe, 2009] et il a également été remarqué pour ses excellentes performances [Kacha *et al.*, 2006, Bartelmaos et Abed-Meraim, 2007]. Il a notamment été utilisé dans le cadre de l'estimation minimax adaptative [Leibovitz *et al.*, 2008], l'estimation aveugle ou semi-aveugle de canaux [Altuna *et al.*, 2006, Kacha *et al.*, 2006, Lin et Yanfen, 2007], l'optimisation d'algorithmes évolutionnistes [Knight et Lunacek, 2007] et le traitement du signal audio pour séparer les composantes sinusoïdales du bruit [Alonso Arevalo *et al.*, 2007]. Il a également inspiré des algorithmes basés sur des principes similaires [Bartelmaos et Abed-Meraim, 2007, Crammer, 2007].

IV.2 Poursuite des pôles complexes

Après l'étape de poursuite de l'espace signal, nous présentons des techniques de poursuite de la matrice spectrale $\mathbf{\Phi}_r(t)$ introduite dans l'équation (II.4) page 5 et des pôles $\hat{z}_k(t)$, parfois appelées "algorithmes ESPRIT adaptatifs" dans la littérature. Chacune de ces techniques est généralement liée à un algorithme de poursuite de l'espace signal qui lui est spécifique [Badeau *et al.*, 2005c]. Nous présentons ci-dessous deux techniques adaptées à l'algorithme FAPI.

IV.2.1 Poursuite exacte de la matrice spectrale et des pôles

Dans les références [Badeau *et al.*, 2003b, Badeau *et al.*, 2005c], nous avons proposé une implémentation adaptative de la méthode ESPRIT, s'appuyant sur un algorithme de poursuite de sous-espace effectuant des mises à jour de rang 1, comme PAST ou FAPI. L'objectif était donc d'implémenter rapidement et sans approximation les deux étapes suivantes, à savoir la poursuite de la matrice spectrale $\mathbf{\Phi}_r(t)$ et de ses valeurs propres $\hat{z}_k(t)$. Nous avons tout d'abord remarqué que le lemme d'inversion matricielle permettait de calculer récursivement la matrice spectrale $\mathbf{\Phi}_r(t)$, avec des mises à jour de rang 3 : $\mathbf{\Phi}_r(t) = \mathbf{\Phi}_r(t-1) + \underline{\boldsymbol{a}}(t)\,\underline{\boldsymbol{b}}(t)^H$, où $\underline{\boldsymbol{a}}(t)$ et $\underline{\boldsymbol{b}}(t)$ sont deux matrices de dimensions $r \times 3$. La complexité du calcul de $\mathbf{\Phi}_r(t)$ est ainsi réduite de $2nr^2$ (pour un calcul direct) à $2nr$. Dans un deuxième temps, nous avons démontré que les pôles $\hat{z}_k(t)$ sont racines d'un polynôme dont les coefficients s'expriment en fonction des valeurs précédentes $z_k(t-1)$ et des matrices $\underline{\boldsymbol{a}}(t)$ et $\underline{\boldsymbol{b}}(t)$. Ce calcul se fait en $2r^3$ opérations (au lieu de $O(r^3)$ pour la diagonalisation directe de la matrice $\mathbf{\Phi}_r(t)$). Le coût de la poursuite du sous-espace étant seulement de $3nr$ avec PAST ou FAPI, la complexité totale de l'algorithme est seulement de $5nr + 2r^3$. Dans la littérature, cette méthode de poursuite a été proposée pour la modélisation de sons percussifs ou de contact [Lagrange *et al.*, 2008c, Lagrange *et al.*, 2008b, Lagrange et Scherrer, 2008] et a inspiré de nouvelles implémentations adaptatives de la méthode ESPRIT [Wong et Evans, 2008].

IV.2.2 Poursuite approchée des pôles

Si l'approche précédente permet d'effectuer rapidement et exactement la poursuite de l'ensemble des pôles, elle n'effectue cependant pas les appariements entre deux instants consécutifs (autrement dit elle ne fournit pas les trajectoires individuelles de chaque pôle). C'est pour cette raison que nous avons poursuivi notre effort dans la référence [David *et al.*, 2006a], où nous avons introduit une nouvelle méthode de poursuite des valeurs et des vecteurs propres de $\mathbf{\Phi}_r(t)$, basée sur une descente de gradient. En l'associant à l'algorithme FAPI et à la méthode de poursuite de $\mathbf{\Phi}_r(t)$ présentée dans la section IV.2.1, on obtient un algorithme qui fournit un suivi individuel des valeurs propres, dont la complexité est $5nr + O(r^3)$. On obtient ainsi, au prix d'une approximation, les trajectoires lissées de chaque pôle $\hat{z}_k(t)$.

IV.3 Poursuite des amplitudes complexes

Une méthode rapide de calcul des amplitudes complexes $\hat{\alpha}_k(t)$ définies dans la section II.1 page 6 a été proposée dans la référence [Badeau, 2005]. Elle exploite la structure particulière de la matrice de Vandermonde $\boldsymbol{V}^n$, mais ne tient pas compte d'une possible récurrence temporelle satisfaite par les pôles $\hat{z}_k(t)$. Sa complexité est de $nr + \frac{r^3}{6} + O(r^2)$ (au lieu de $nr^2 + nr + O(r^3)$ pour un calcul direct).

Pour réduire davantage le coût de calcul, nous avons dû nous placer dans le cas particulier où les pôles $z_k(t)$ sont constants. Nous avons alors démontré que les amplitudes $\hat{\alpha}_k(t)$ peuvent être calculées récursivement et de manière exacte en utilisant le lemme d'inversion matricielle. L'algorithme ainsi obtenu a une complexité de seulement $5r$ à chaque instant t (indépendante du nombre de lignes n de la matrice de Vandermonde). Il a notamment été appliqué avec succès à la modélisation de notes de piano [David et Badeau, 2007].

Chapitre V

Applications de l'analyse à haute résolution

Dans ce chapitre, nous présentons diverses applications des techniques présentées dans les chapitres III et IV. La plupart de ces travaux se sont appuyés sur le développement d'un système complet d'analyse / synthèse, mettant en œuvre la méthode ESPRIT et un certain nombre de prétraitements destinés à améliorer ses performances (section V.1).

V.1 Analyse / synthèse et représentation de signaux audio

Dans [Badeau, 2005], nous avons montré que le modèle ESM est particulièrement bien adapté à l'analyse spectrale de signaux de musique. Il n'en reste pas moins que la mise en œuvre des méthodes HR est délicate et nécessite de prendre certaines précautions. Tout d'abord, la performance de l'estimation est sensible à l'hypothèse de blancheur du bruit additif $b(t)$. Or les spectres de signaux audio présentent généralement une forte dynamique ; aussi le bruit présent dans ces signaux s'avère souvent très coloré. De plus, les composantes sinusoïdales de forte énergie (dans les basses fréquences) peuvent masquer celles de plus faible énergie (dans les hautes fréquences), qui ne sont pourtant pas négligeables d'un point de vue perceptif. Une autre difficulté est liée au nombre parfois élevé de composantes sinusoïdales présentes dans toute l'étendue du spectre audible (par exemple dans le cas de sons graves ou de musique polyphonique). Or pour des raisons de complexité algorithmique et de conditionnement numérique, il est souhaitable en pratique de ne pas dépasser un ordre de modélisation d'environ 25 composantes.

Nous avons donc développé un système complet d'analyse spectrale à haute résolution de signaux audio, qui met en œuvre un certain nombre de prétraitements, destinés à améliorer l'estimation et à réduire la complexité algorithmique [Badeau, 2005, Badeau et David, 2008a]. Le signal est tout d'abord filtré par un filtre de préaccentuation, destiné à rehausser les aigus par rapport aux graves. Il est ensuite découpé en sous-bandes à l'aide d'un banc de filtres. Cette opération présente plusieurs avantages : en effet la sélection d'une bande de fréquences suivie d'une décimation permet à la fois de réduire le nombre et la dynamique des composantes sinusoïdales, d'améliorer la résolution spectrale et de réduire la dynamique du bruit. Dans chaque sous-bande, le bruit est ensuite blanchi. La technique utilisée comprend un filtrage médian du spectre de Fourier, destiné à éliminer les composantes sinusoïdales pour estimer l'enveloppe spectrale et un modèle autorégressif (AR) du bruit. Le signal de sous-bande est enfin filtré à l'aide du filtre à réponse impulsionnelle finie inverse de l'AR, de façon à blanchir le bruit. Dans chaque sous-bande l'ordre de modélisation est estimé grâce à la méthode ESTER présentée dans la section III.2 page 8, les paramètres du modèle ESM sont ensuite estimés en utilisant l'une des méthodes de poursuite présentées dans le chapitre IV, puis redistribués en bande pleine et corrigés en tenant compte des filtres de préaccentuation, d'analyse et de blanchiment. A l'issue de ce traitement, on obtient une suite de pôles complexes $\hat{z}_k(t) = e^{-\hat{\delta}_k(t)+i2\pi\hat{\nu}_k(t)}$ et d'amplitudes complexes $\hat{\alpha}_k(t) = \hat{a}_k(t)\, e^{i\hat{\phi}_k(t)}$, dont on peut extraire les fréquences instantanées $\hat{\nu}_k(t)$ et les puissances instantanées $\hat{a}_k(t)^2$.

Ce système d'analyse permet de définir une représentation temps-fréquence du signal audio, constituée d'un nuage de points de coordonnées $(t, \hat{\nu}_k(t))$ auxquels sont associées les puissances instantanées $\hat{a}_k(t)^2$. Nous avons appelé cette représentation Spectrogramme à Haute Résolution (HR-ogramme), par analogie avec le spectro-

gramme [David *et al.*, 2003, David et Badeau, 2008]. Dans les figures V.1(b) et V.1(e) nous avons représenté les HR-ogrammes d'un son de piano et d'un son de violon, dont les spectrogrammes apparaissent dans les figures V.1(a) et V.1(d). On peut observer que les algorithmes de poursuite ont correctement estimé les fréquences et les puissances instantanées des composantes sinusoïdales. Cette technique de représentation temps-fréquence à été comparée à des méthodes classiques de réassignement [Emiya *et al.*, 2005, David *et al.*, 2006b]. En pratique, les méthodes de réassignement offrent une **précision** spectrale et temporelle accrues par rapport à un spectrogramme, mais ne bénéficient pas de la haute **résolution** spectrale du HR-ogramme.

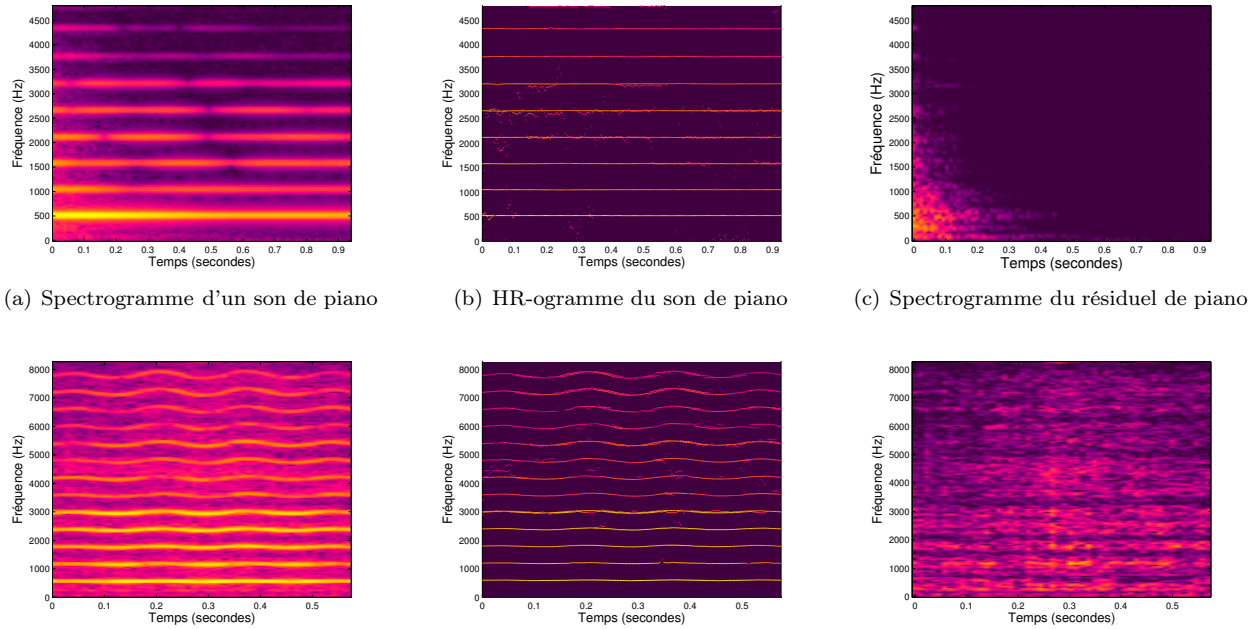

(a) Spectrogramme d'un son de piano (b) HR-ogramme du son de piano (c) Spectrogramme du résiduel de piano

(d) Spectrogramme d'un son de violon (e) HR-ogramme du son de violon (f) Spectrogramme du résiduel de violon

FIG. V.1 – Séparation des composantes sinusoïdales et bruitée de sons de piano et de violon

Ce système d'analyse a notamment été utilisé pour l'estimation de hauteurs de notes de piano (section V.2), le codage de signaux audio (section V.5) et l'analyse modale de structures, appliquée à la table d'harmonie du piano [Ege *et al.*, 2009b]. Il permet également de reconstruire la partie sinusoïdale du signal, soit par synthèse avec additions-recouvrements à partir des modèles ESM estimés dans chaque trame, soit par projection sur l'espace signal estimé dans chaque trame et dans chaque sous-bande, suivie d'un filtrage. La partie bruitée peut en être facilement déduite, soit en soustrayant du signal original la partie sinusoïdale synthétisée, soit par projection sur l'espace bruit [Badeau *et al.*, 2002, David *et al.*, 2003, Badeau, 2005, David *et al.*, 2006b]. Les spectrogrammes des parties bruitées des sons de piano et de violon analysés précédemment sont ainsi représentés sur les figures V.1(c) et V.1(f). On observe par rapport aux figures V.1(a) et V.1(d) que les sinusoïdes ont été complètement soustraites. Dans la littérature, cette méthode de séparation sinusoïdes / bruit a notamment été utilisée pour extraire des sons environnementaux transitoires de faible énergie [Grenier et David, 2003], pour séparer des pistes de batterie [Gillet et Richard, 2008] ou encore pour estimer le rythme musical [Alonso Arevalo *et al.*, 2007] (section V.3). Dans les références [Badeau *et al.*, 2002, David *et al.*, 2003] nous avons enfin proposé des méthodes de synthèse permettant de transformer la hauteur du son (en modifiant les composantes sinusoïdales sans modifier la partie bruitée) ou sa durée et d'inclure des effets comme la suppression ou l'ajout de vibrato ou de trémolo, la diminution ou l'augmentation de la puissance du bruit.

V.2 Estimation de hauteurs de notes de piano

Le problème de l'estimation de hauteurs [Christensen et Jakobsson, 2009] constitue une application naturelle de l'analyse spectrale à haute résolution, puisqu'il vise à modéliser le signal comme une somme de composantes

sinusoïdales, dont la distribution fréquentielle n'est pas quelconque, mais dépend de la fréquence fondamentale ν^0. Ainsi pour un son harmonique, la composante d'indice $h \in \mathbb{N}^*$ a une fréquence $\nu^h = h\nu^0$, alors que pour un son inharmonique comme une note de piano, les fréquences sont distribuées selon la loi $\nu^h = h\nu^0\sqrt{1+\beta h^2}$, où $\beta \geq 0$ est le coefficient d'inharmonicité [Fletcher et Rossing, 2008]. Une limitation connue de la plupart des méthodes d'estimation de hauteur est que leur performance décline quand la durée d'observation diminue ou lorsque la tessiture, c'est à dire l'intervalle de fréquences fondamentales possibles, s'élargit [Emiya *et al.*, 2007b]. Dans le cadre de la thèse de Valentin Emiya, nous avons proposé une nouvelle technique d'estimation de la hauteur de notes de piano visant à s'affranchir de ces limitations et basée sur l'algorithme d'analyse HR présenté dans la section V.1. La hauteur est ainsi estimée en recherchant le maximum d'une **fonction de détection** [Emiya *et al.*, 2007b]. Trois fonctions de détection ont été proposées : pour chaque trame, on calcule une fonction de détection de périodicité temporelle, une fonction de détection d'harmonicité fréquentielle ou une fonction de détection mixte à partir des paramètres du modèle ESM estimés, en tenant compte de l'inharmonicité des notes de piano. L'algorithme a été évalué sur une base de notes de piano isolées extraites des bases *Real World Computing* (RWC) [Goto *et al.*, 2002], PROSONUS [Prosonus, 1998] et *MIDI Aligned Piano Sounds* (MAPS) (section IX.1 page 36) et comparé à l'algorithme YIN [De Cheveigne et Kawahara, 2002]. Une note a été considérée correctement estimée quand sa fréquence fondamentale se trouve dans un intervalle d'un demi-ton centré autour de la référence *Musical Instrument Digital Interface* (MIDI). Des performances satisfaisantes ont ainsi été obtenues sur des trames de 60 millisecondes et sur un intervalle de 7,25 octaves (couvrant toute la tessiture du piano). Le taux d'erreur global de notre estimateur est voisin de 4.4%, ce qui est deux fois mieux que YIN. Cela est dû à un taux d'erreur très faible sur un large intervalle de fréquences fondamentales (1.1% dans l'intervalle 65-2000 Hz). En comparaison, les performances de YIN sont similaires dans les graves, un peu moins bonnes dans les médiums et plus faibles dans les aigus. Depuis la parution de ces résultats, d'autres travaux ont été entrepris sur l'application des méthodes HR au problème de l'estimation de hauteur, comme ceux menés au sein de l'Université d'Aalborg au Danemark [Christensen et Jakobsson, 2009] qui abordent le cas de fréquences fondamentales multiples.

V.3 Estimation du tempo musical

Malgré l'aisance que nous avons à battre la mesure quand nous écoutons de la musique, l'estimation automatique du tempo dans des œuvres musicales reste un problème difficile à résoudre [Alonso Arevalo, 2006]. Dans le cadre d'un travail de recherche mené avec Miguel Alonso, nous avons étudié le potentiel des méthodes HR pour estimer et suivre les variations du tempo, exprimé en terme de battements par minute (BPM). Le système proposé visait à traiter une grande variété de genres musicaux [Alonso Arevalo *et al.*, 2003a]. L'approche que nous avons suivie s'appuie sur l'observation suivante : lorsque l'on retire les composantes sinusoïdales d'un morceau de musique, le résiduel contient essentiellement des sons de nature transitoire : sons percussifs, attaques de notes, *etc.* Ainsi la battue doit être plus facilement détectable dans ce résiduel que dans l'enregistrement original. Nous avons donc utilisé la technique de séparation sinusoïdes / bruit présentée dans la section V.1 et effectué une détection de périodicité sur le résiduel extrait dans chaque sous-bande. Notre algorithme a été évalué sur une base de données composée de 54 morceaux annotés manuellement, contenant des extraits de rock, de pop, de jazz, de musique latine, classique et de chants traditionnels. Le tempo a été considéré correctement estimé lorsque l'estimation diffère de moins de 5% par rapport au tempo annoté, modulo un possible rapport 1/2 ou 2. Sur cette base de données, notre algorithme a surpassé trois autres méthodes de l'état de l'art [Scheirer, 1998, Paulus et Klapuri, 2002, Alonso Arevalo *et al.*, 2003b]. En utilisant un découpage en fenêtres de 6 secondes, il a ainsi atteint un taux de reconnaissance de 96%.

V.4 Analyse des modes couplés d'une harpe de concert

La harpe de concert est un instrument de musique composé de trois éléments acoustiques : la table d'harmonie, la caisse de résonance et les 47 cordes. Lorsqu'une corde est jouée, sa vibration se transmet à la table d'harmonie. Cette vibration donne lieu à un rayonnement acoustique qui, dans une certaine gamme de fréquences, est amplifiée par la caisse de résonance. Au final, le son émis par l'instrument s'avère être d'une grande complexité, allant même jusqu'à occasionner une gêne pour l'instrumentiste qui le décrit comme un "halo sonore". Il n'est pas rare en effet que la mise en vibration d'une corde fasse également vibrer de nombreuses

autres cordes, perturbant ainsi le jeu du musicien. La thèse de Jean-Loïc Le Carrou portant sur la "Vibro-acoustique de la harpe de concert" [Le Carrou, 2006] a permis d'analyser ce phénomène, appelé **vibration par sympathie**. Un modèle vibratoire d'une harpe de concert simplifiée a ainsi été développé. Ce modèle, constitué d'une poutre connectée à 35 cordes, a permis de calculer les modes propres de l'assemblage. Dans les déformées modales obtenues, les modes de cordes couplées, aussi appelés modes sympathiques, ont pu être identifiés. Afin de valider ce modèle théorique, l'identification expérimentale des modes sympathiques a fait l'objet d'une étude que nous avons menée en collaboration. La méthode ESPRIT décrite dans la section II.1 page 4 a ainsi été appliquée à des signaux vibratoires mesurés sur la table d'harmonie d'une harpe de concert, dont certaines cordes ont été amorties. L'erreur relative entre les fréquences modales expérimentales et théoriques s'est avérée très faible (inférieure à 0.11%), permettant de conclure positivement quant à la validité du modèle physique [Le Carrou *et al.*, 2006, Le Carrou *et al.*, 2007, Le Carrou *et al.*, 2009].

V.5 Codage audio

Les systèmes de codage audio actuels, comme le *MPEG-1/2 audio layer 3* (MP3) ou le *MPEG-2/4 AAC* (MP4), comportent généralement une transformée temps-fréquence, suivie d'un module de quantification des coefficients transformés puis d'un module de codage entropique [Spanias *et al.*, 2007]. Les transformées utilisées habituellement, comme la *Modified Discrete Cosine Transform* (MDCT), sont dérivées de l'analyse de Fourier. Elles correspondent donc à une décomposition du signal sur une famille de sinusoïdes d'amplitude constante. Jusqu'à récemment, les méthodes HR ont été rarement utilisées en codage audio alors que le modèle ESM permet de modéliser des sinusoïdes d'amplitude variable. Dans le cadre du projet DESAM de l'ANR et d'une collaboration avec Olivier Derrien du Laboratoire de Mécanique et Acoustique (LMA), nous avons montré qu'elles sont non seulement efficaces pour le codage à bas débit, déjà dominé par les codeurs paramétriques, mais aussi pour le codage à haut débit, dominé par les codeurs par transformées. Le système de codage que nous avons proposé dans [Derrien *et al.*, 2008] décompose d'abord le signal en 8 sous-bandes, à l'aide d'un banc de filtres *Pseudo-Quadrature Mirror Filters* (PQMF). Les signaux de sous-bandes sont ensuite segmentés en détectant les attaques et la méthode ESPRIT est appliquée à chaque segment. Pour chaque composante sinusoïdale estimée, les paramètres de fréquence, atténuation, amplitude et phase sont quantifiés. Le signal résiduel n'est pas codé. Ce système a été comparé à un codeur basé sur la MDCT, qui utilise la même transformée et la même quantification que MPEG-2/4 AAC, mais sans inclure le modèle perceptif et le codage entropique. Des tests objectifs et subjectifs préliminaires ont montré le potentiel de cette approche qui nécessite un espace de codage plus faible que le codeur basé sur la MDCT, à qualité sonore égale.

Par ailleurs, dans un codeur paramétrique les valeurs des paramètres, qui sont originellement continues, doivent être quantifiées avant de procéder au codage binaire. Nous nous sommes donc également intéressés à l'étape de quantification, qui a un impact majeur sur la performance globale du codeur. Comme les techniques modernes de communication utilisent couramment des codes binaires de longueur variable, la contrainte de débit doit être formulée en terme d'entropie des indices de quantification. Nous avons ainsi proposé une nouvelle méthode de quantification scalaire conjointe des paramètres d'atténuation, d'amplitude et de phase avec contrainte entropique [Derrien *et al.*, 2011, Derrien *et al.*, 2010], qui est une alternative aux techniques de quantification vectorielle [Chou *et al.*, 1989]. Ces deux approches ont été évaluées sur des données réelles obtenues par analyse HR de signaux de musique, en terme de rapport signal sur bruit de quantification (RSBQ). Deux bases de données contenant 8 extraits de divers genres musicaux échantillonnés à 44.1 kHz et d'une durée totale de 1 minute ont ainsi été utilisées pour l'apprentissage d'une part et pour l'évaluation d'autre part. La méthode de référence [Chou *et al.*, 1989] a présenté un meilleur RSBQ à résolution basse et moyenne (1-11 bits), aux dépends d'une complexité algorithmique très élevée et de pics de distorsion qui surviennent principalement dans les composantes à forte amplitude et forte atténuation, correspondant généralement à des attaques brusques. En comparaison, notre méthode ne s'est pas montrée aussi performante en terme de RSBQ, mais elle évite les pics de distorsion au niveau des attaques et sa complexité est beaucoup plus faible. Les deux approches ont également été évaluées selon un critère de qualité perceptive - *Objective Difference Grade* (ODG), calculé avec l'algorithme *PErception MOdel Quality* (PEMO-Q) [Huber et Kollmeier, 2006], qui compare les signaux originaux et re-synthétisés. Notre méthode a ainsi obtenu de meilleurs résultats que la méthode de référence, à résolutions moyenne et élevée (4-13 bits).

V.6 Autres applications

Outre l'analyse spectrale, les méthodes HR ont bien souvent été utilisées dans d'autres domaines, allant du traitement d'antenne aux communications numériques [Marcos, 1998]. Nous présentons ici deux exemples d'applications, que nous avons eu l'occasion de mener en collaboration avec des laboratoires français et étrangers.

V.6.1 Estimation de canal en communications numériques

La demande croissante de haute efficacité spectrale dans les communications sans fil à large bande, où la dispersion temporelle due aux trajets multiples est un problème majeur, a motivé le développement de techniques d'estimation de canal et d'égalisation (semi-)aveugles, qui requièrent peu ou pas de connaissance sur la séquence transmise. Un grand nombre de techniques (semi-)aveugles ont ainsi été proposées, s'appuyant sur des propriétés statistiques ou sur la structure algébrique du signal reçu [Tong et Perreau, 1998]. En particulier, plusieurs méthodes de type sous-espace ont été proposées pour l'identification de canal aveugle et semi-aveugle, après les travaux pionniers mené par l'équipe d'Eric Moulines à Télécom ParisTech [Cardoso *et al.*, 1995]. Ces algorithmes nécessitent dans un premier temps de déterminer le sous-espace signal ou le sous-espace bruit de la matrice de corrélation du signal reçu, puis de maximiser une forme quadratique pour obtenir la réponse impulsionnelle du canal à un facteur d'échelle près. En collaboration avec l'université de Mondragón en Espagne, nous avons ainsi proposé une nouvelle méthode adaptative rapide d'estimation aveugle de canal, basée sur l'estimation du sous-espace principal de la matrice de corrélation [Altuna *et al.*, 2006]. Cette méthode est pleinement adaptative, au sens où le sous-espace et la forme quadratique sont tous deux calculés récursivement. Elle est basée sur l'algorithme YAST (section IV.1.4 page 15), dont nous avons montré que les performances surpassent la plupart des autres algorithmes de poursuite de sous-espace, à la fois en terme de vitesse de convergence et de complexité algorithmique. Ces travaux ayant cependant été entrepris avant la mise au point d'une version numériquement stable de YAST, nous avons choisi d'utiliser YAST à l'initialisation pour tirer parti de sa remarquable vitesse de convergence, puis de le relayer par un autre algorithme de poursuite plus stable, tel que OPAST [Abed-Meraim *et al.*, 2000]. Des simulations numériques ont été effectuées pour comparer la performance de l'algorithme hybride ainsi obtenu avec l'estimateur basé sur la SVD exacte et avec un autre algorithme adaptatif utilisant uniquement OPAST. Les performances ont été mesurées en calculant l'erreur quadratique moyenne normalisée des réponses impulsionnelles estimées. Nous avons ainsi observé que l'algorithme hybride converge globalement vers la solution obtenue avec une SVD exacte, avec une vitesse supérieure à celui utilisant OPAST uniquement. La méthode proposée conserve donc une complexité linéaire et atteint des performances remarquables en termes de vitesse de convergence et de stabilité numérique.

V.6.2 SVD multilinéaire adaptative pour les tenseurs structurés

La décomposition en valeurs singulières d'ordre supérieur - *Higher-Order SVD* (HOSVD) est une généralisation de la SVD aux tenseurs d'ordre supérieur à 2, qui joue un rôle important dans de nombreux domaines [De Lathauwer, 1997]. Malheureusement, la complexité algorithmique de cette décomposition est très élevée, puisque son implémentation directe requiert le calcul des SVD de trois matrices très redondantes, appelées modes. Dans le cadre d'une collaboration avec Rémy Boyer du Laboratoire des Signaux et Systèmes (LSS), nous avons développé une méthode ultra-rapide de calcul de la HOSVD d'un tenseur structuré d'ordre 3 [Boyer et Badeau, 2006, Badeau et Boyer, 2008]. Le résultat clé de ce travail a été la possibilité de se ramener au calcul des SVD de trois matrices de Hankel non redondantes, dont les colonnes sont multipliées par des poids déterminés. Cette opération est une première étape qui a permis de réduire la complexité d'un ordre de grandeur. Pour réduire encore le coût de calcul, nous avons proposé d'effectuer des produits rapides entre matrices de Hankel et vecteurs, à l'aide de transformées de Fourier rapides. Enfin, dans un contexte adaptatif, nous avons proposé une implémentation séquentielle basée sur l'algorithme d'itération orthogonale présenté dans la section III.3.1 page 9. Dans le cas d'un tenseur cubique de dimension $I \times I \times I$, la complexité d'une implémentation directe de la HOSVD est $O(rI^3)$, où r est le rang de troncature de la factorisation, alors que notre algorithme optimisé le plus rapide présente une complexité quasi-linéaire $O(rI \log_2(I))$. Dans la littérature, cet algorithme a été proposé comme outil de factorisation de tenseurs structurés [Oseledets et Tyrtyshnikov, 2010, Ishteva *et al.*, 2009] et nous l'avons utilisé pour représenter des systèmes non-linéaires, en décomposant le noyau cubique d'un modèle de Volterra [Boyer *et al.*, 2011].

Deuxième partie

Décompositions non-négatives

Chapitre VI

Introduction de la Deuxième partie

Dans cette deuxième partie, nous modélisons le signal observé comme un processus aléatoire gaussien [Priestley, 1983]. Soit $x(t)$ (où $t \in \mathbb{Z}$) un processus aléatoire gaussien réel centré et stationnaire, dont la fonction d'autocovariance est sommable, et de densité spectrale de puissance $\widehat{v}(\nu)$ (où ν désigne la fréquence réduite). On considère une représentation de Fourier discrète du signal $x(t)$, observée sur un horizon de longueur $T \in \mathbb{N}^*$:

$$\forall f \in \{1 \dots F\}, x_f = \frac{1}{\sqrt{T}} \sum_{t \in \mathbb{Z}} x(t)\, g(t/T)\, e^{-2i\pi\nu_f t}$$

où la fenêtre g est une fonction de support $[-\frac{1}{2}, \frac{1}{2}]$, continue par morceaux et d'énergie unité ($\int_{-\frac{1}{2}}^{+\frac{1}{2}} g^2(u)\, du = 1$). Le signal étant réel, on considère des fréquences réduites $\nu_f \in [0, \frac{1}{2}]$, distinctes mais pas nécessairement uniformément réparties. Alors $\forall f \in \{1 \dots F\}$, x_f est une variable aléatoire gaussienne centrée, réelle si $\nu_f = 0$ ou $\nu_f = \frac{1}{2}$ et complexe sinon. De plus, si $\nu_f \in]0, \frac{1}{2}[$ et $T \to +\infty$, la variable aléatoire x_f est asymptotiquement circulaire. On note alors $\boldsymbol{x}$ et $\widehat{\boldsymbol{v}}$ les vecteurs de même dimension F, dont les coefficients sont respectivement x_f et $\widehat{v}_f = \widehat{v}(\nu_f)$. Lorsque $T \to +\infty$, $\boldsymbol{x}$ est un vecteur aléatoire gaussien centré, de matrice de covariance $\mathrm{diag}(\widehat{\boldsymbol{v}})$, ce que l'on note $\boldsymbol{x} \sim \mathcal{N}(\mathbf{0}, \mathrm{diag}(\widehat{\boldsymbol{v}}))$.

VI.1 Décompositions spectrales (chapitre VII)

Ce cadre théorique offre un formalisme commode pour estimer un modèle spectral paramétrique par la méthode du maximum de vraisemblance. En effet, si on considère une densité spectrale de puissance $\widehat{v}(\nu)$ paramétrée par un ensemble de paramètres $\boldsymbol{\Theta}$, la log-vraisemblance des observations x_f s'écrit sous la forme

$$L(\boldsymbol{\Theta}) \stackrel{\mathrm{c}}{=} -\sum_{f=1}^{F} d_{IS}(v_f | \widehat{v}_f(\boldsymbol{\Theta})) \tag{VI.1}$$

où le symbole $\stackrel{\mathrm{c}}{=}$ désigne l'égalité à une constante additive près, $v_f = |x_f|^2$ est la valeur du périodogramme de x à la fréquence ν_f et la fonction d_{IS} est la divergence d'Itakura-Saito (IS) : $d_{IS}(a|b) = \frac{a}{b} - \log\frac{a}{b} - 1$ [Itakura et Saito, 1968]. Si en particulier on souhaite faire apparaître la puissance du signal parmi les paramètres du modèle, on peut factoriser $\widehat{v}_f(\boldsymbol{\Theta})$ sous la forme $\widehat{v}_f(\boldsymbol{\Theta}) = \sigma^2\, w_f(\boldsymbol{\theta})$, où $\sigma^2 \geq 0$ est un paramètre scalaire de puissance et $w_f(\boldsymbol{\theta}) \geq 0$ représente la forme spectrale associée. L'ensemble des paramètres est alors $\boldsymbol{\Theta} = (\sigma^2, \boldsymbol{\theta})$. La maximisation de L par rapport à σ^2 conduit à la valeur $\hat{\sigma}^2 = \frac{1}{F}\sum_{f=1}^{F} \frac{v_f}{w_f(\boldsymbol{\theta})}$, Si on définit $L(\boldsymbol{\theta}) = \max_{\sigma^2} L(\boldsymbol{\Theta})$, on obtient par substitution dans l'équation (VI.1) la relation

$$\exp\left(\frac{L(\boldsymbol{\theta})}{F}\right) \propto \left(\prod_{f=1}^{F} \frac{v_f}{w_f(\boldsymbol{\theta})}\right)^{\frac{1}{F}} \Bigg/ \left(\frac{1}{F}\sum_{f=1}^{F} \frac{v_f}{w_f(\boldsymbol{\theta})}\right) \tag{VI.2}$$

où le symbole $\propto$ désigne l'égalité à un facteur multiplicatif près. La quantité dans le membre de droite de (VI.2), égale au quotient de la moyenne géométrique des valeurs $\frac{v_f}{w_f(\boldsymbol{\theta})}$ sur leur moyenne arithmétique, appartient

toujours à l'intervalle $[0, 1]$ et vaut 1 si et seulement si toutes ces valeurs sont égales à une même constante indépendante de f. Ainsi l'expression (VI.2) mesure la "platitude spectrale" (*spectral flatness*) [Johnston, 1988] du quotient du périodogramme v_f par la forme spectrale $w_f(\boldsymbol{\theta})$. Nous avons utilisé ce type de critère pour effectuer des estimations d'enveloppe spectrale (section VII.1.2 page 24) et pour définir des mesures de similarités spectrales (section VII.1.3 page 25).

Dans le cas où le signal observé n'est plus constitué d'une seule, mais de plusieurs composantes indépendantes, on modélise $x(t)$ comme la somme de K composantes gaussiennes réelles centrées et stationnaires $x^{(k)}(t)$, possédant chacune une densité spectrale de puissante $\widehat{v}^{(k)}(\nu)$, de sorte que $\boldsymbol{x}^{(k)} \sim \mathcal{N}\left(\mathbf{0}, \mathrm{diag}\left(\widehat{\boldsymbol{v}}^{(k)}\right)\right)$ et $\boldsymbol{x} \sim \mathcal{N}\left(\mathbf{0}, \mathrm{diag}\left(\sum_{k=1}^{K} \widehat{\boldsymbol{v}}^{(k)}\right)\right)$. Si les densités spectrales de puissance $\widehat{v}^{(k)}(\nu)$ suivent des modèles paramétriques différents, il est possible de les estimer conjointement en utilisant la même approche que précédemment. On obtient ainsi une log-vraisemblance de la forme (VI.1), où $\hat{v}_f(\boldsymbol{\Theta}) = \sum_{k=1}^{K} \widehat{v}_f^{(k)}(\boldsymbol{\Theta})$. Nous avons appliqué cette approche à l'estimation de mélanges de spectres harmoniques et de bruit coloré (section VII.2 page 26).

VI.2 Décompositions temps-fréquences (chapitre VIII)

Une autre façon de distinguer plusieurs composantes dans un mélange est de tenir compte de leurs variations temporelles différentes de puissance instantanée. On ne considère alors plus un unique vecteur aléatoire $\boldsymbol{x} = \sum_{k=1}^{K} \boldsymbol{x}^{(k)}$, mais une succession de N vecteurs aléatoires indépendants $\boldsymbol{x}_n = \sum_{k=1}^{K} \boldsymbol{x}_n^{(k)}$ (dont la juxtaposition constitue une représentation de Fourier à court terme du signal observé), où $\boldsymbol{x}_n^{(k)} \sim \mathcal{N}(0, \mathrm{diag}(\widehat{\boldsymbol{v}}_n^{(k)}))$. Comme dans la section VI.1, on peut factoriser $\widehat{v}_{fn}^{(k)}(\boldsymbol{\Theta})$ sous la forme $\widehat{v}_{fn}^{(k)}(\boldsymbol{\Theta}) = h_{kn}\, w_{fk}(\boldsymbol{\theta})$, où $h_{kn} \geq 0$ est un paramètre scalaire de puissance instantanée et $w_{fk}(\boldsymbol{\theta}) \geq 0$ représente la forme spectrale associée à la composante k. L'ensemble des paramètres est alors $\boldsymbol{\Theta} = (\boldsymbol{H}, \boldsymbol{\theta})$, où $\boldsymbol{H}$ est la matrice de dimensions $K \times N$ et de coefficients h_{kn}. La log-vraisemblance des observations x_{fn} s'écrit alors sous la forme

$$L(\boldsymbol{\Theta}) \stackrel{\mathrm{c}}{=} -\sum_{n=1}^{N} \sum_{f=1}^{F} d_{IS}(v_{fn} | \widehat{v}_{fn}(\boldsymbol{\Theta})) \tag{VI.3}$$

où $v_{fn} = |x_{fn}|^2$ est la valeur du spectrogramme de puissance de x à la fréquence ν_f et au temps n, et $\widehat{v}_{fn}(\boldsymbol{\Theta}) = \sum_{k=1}^{K} \widehat{v}_{fn}^{(k)}(\boldsymbol{\Theta}) = \sum_{k=1}^{K} h_{kn}\, w_{fk}(\boldsymbol{\theta})$. En définissant les matrices $\widehat{\boldsymbol{V}}(\boldsymbol{\Theta})$ et $\boldsymbol{W}(\boldsymbol{\theta})$, respectivement de dimensions $F \times N$ et $F \times K$ et de coefficients $\widehat{v}_{fn}(\boldsymbol{\Theta})$ et $w_{fk}(\boldsymbol{\theta})$, on aboutit à la factorisation matricielle $\widehat{\boldsymbol{V}}(\boldsymbol{\Theta}) = \boldsymbol{W}(\boldsymbol{\theta})\boldsymbol{H}$. Ainsi l'estimation du modèle de mélange au sens du maximum de vraisemblance revient à déterminer les matrices $\boldsymbol{W}(\boldsymbol{\theta})$ et $\boldsymbol{H}$ qui minimisent la divergence d'Itakura-Saito entre le spectrogramme de puissance $\boldsymbol{V}$ du signal observé et le modèle $\widehat{\boldsymbol{V}}(\boldsymbol{\Theta})$. Si la matrice $\boldsymbol{W}(\boldsymbol{\theta})$ n'est paramétrée que par ses propres coefficients, cette factorisation est connue sous le nom de NMF [Lee et Seung, 1999] et peut être calculée à l'aide d'algorithmes multiplicatifs (sections VIII.1 page 28 et VIII.2 page 30). En revanche l'utilisation d'un modèle plus contraint pour $\boldsymbol{W}(\boldsymbol{\theta})$ permet d'inclure divers critères comme l'harmonicité des composantes spectrales $w_{fk}(\boldsymbol{\theta})$ (section VIII.3.1.1 page 33).

Enfin, le cadre théorique présenté ci-dessus peut être facilement étendu à l'estimation bayésienne, qui substitue au critère du maximum de vraisemblance un critère de maximum *a posteriori* (MAP). Ce dernier intègre comme information *a priori* la distribution des paramètres $p(\boldsymbol{\Theta})$ et revient à ajouter à la log-vraisemblance $L(\boldsymbol{\Theta})$ le terme de pénalité $\ln(p(\boldsymbol{\Theta}))$. Nous avons utilisé cette approche pour intégrer des critères de régularité temporelle dans la NMF (section VIII.3.1.2 page 33).

VI.3 Applications des décompositions non-négatives (chapitre IX)

Les décompositions non-négatives ont été utilisées dans diverses applications musicales, dont l'estimation de hauteur et la transcription automatique (section IX.1 page 36) et la séparation de sources informée (section IX.2 page 38).

Chapitre VII

Décompositions spectrales

VII.1 Modélisation de spectres de sons stationnaires

Les notes de musique sont typiquement caractérisées par leur harmonicité (qui définit leur hauteur tonale) et leur enveloppe spectrale (qui est l'une des caractéristiques du timbre). Les travaux présentés ci-dessous ont porté sur la modélisation statistique conjointe de ces deux propriétés et sur le développement d'algorithmes d'estimation des paramètres. Ces travaux ont été appliqués à des problèmes d'estimation de hauteurs multiples et de transcription automatique de musique (chapitre IX). En utilisant le même cadre statistique, nous définissons également des mesures de similarité spectrales destinées à comparer les timbres de deux sons (section VII.1.3).

VII.1.1 Harmonicité

Dans la référence [Emiya *et al.*, 2010a], nous avons considéré la modélisation de notes de piano. Le modèle utilisé est un peu plus général que le modèle strictement harmonique, car il est capable de décrire l'inharmonicité due à la raideur des cordes. On montre en effet que les fréquences ν^h des partiels d'une note de piano (indexés par l'entier h) sont distribuées selon une loi qui s'exprime en fonction de la fréquence fondamentale ν^0 et d'un coefficient d'inharmonicité β : $\nu^h = h\nu^0\sqrt{1+\beta h^2}$ [Fletcher et Rossing, 2008]. On remarque en particulier que la valeur $\beta = 0$ correspond au modèle harmonique. On suppose par ailleurs que les amplitudes des partiels suivent le modèle spectral introduit dans la section VI.1 page 22 et qu'en dehors des fréquences ν^h, le spectre possède une énergie nulle. Pour estimer conjointement les paramètres ν^0 et β, nous avons proposé une méthode simple, qui s'appuie sur une recherche bidimensionnelle du maximum d'un produit spectral normalisé [Emiya *et al.*, 2010a].

VII.1.2 Enveloppe spectrale

La propriété d'harmonicité est une caractéristique essentielle pour estimer la hauteur de notes de musique. Elle n'est cependant pas suffisante, car sans hypothèse supplémentaire, tous les sous-multiples de la fréquence fondamentale sont des solutions admissibles. Pour supprimer cette ambiguïté, il est nécessaire de contraindre davantage le modèle. Une propriété très souvent utilisée dans la littérature est celle de **régularité spectrale** (*spectral smoothness*) [Klapuri, 2003]. Nous avons ainsi eu recours à des modèles d'enveloppe spectrale destinés à contraindre la régularité du spectre. Ces modèles ont également été utilisés pour représenter des bruits colorés.

VII.1.2.1 Modèles linéaires

Les modèles linéaires expriment la forme spectrale $w_f(\boldsymbol{\theta})$ introduite dans la section VI.1 comme une combinaison linéaire de motifs $P_m(f)$ appartenant à un dictionnaire prédéfini : $w_f(\boldsymbol{\theta}) = \sum_{m=1}^{M} e_m P_m(f)$, où les coefficients e_m et les motifs $P_m(f)$ sont à valeurs positives. L'ensemble de paramètres $\boldsymbol{\theta}$ est alors constitué des coefficients e_m. Ces modèles présentent l'avantage d'une grande souplesse puisque le dictionnaire peut être choisi librement. Les motifs $P_m(f)$ peuvent par exemple définir un découpage en sous-bandes, uniformes ou adaptées à une échelle psycho-acoustique comme l'échelle *Equivalent Rectangular Bandwidth* (ERB) [Moore et Glasberg, 1983], que nous avons utilisée dans un contexte d'estimation de hauteurs multiples (section VIII.3.1.1 page 33). Les coefficients e_m peuvent être estimés rapidement à l'aide d'algorithmes multiplicatifs appartenant à la famille décrite dans la section VIII.2.1 page 30.

VII.1.2.2 Modèles autorégressifs à moyenne ajustée

Par rapport aux modèles linéaires, l'intérêt du modèle autorégressif à moyenne ajustée (ARMA), outre sa justification physique (la propagation d'un son à travers un tube de section variable est modélisée par un filtre autorégressif), est sa capacité à représenter des formes complexes d'enveloppes spectrales avec un nombre réduit de coefficients. Son estimation s'avère en revanche plus délicate [Priestley, 1983]. Dans la référence [Badeau et David, 2008b], nous avons proposé de nouveaux algorithmes permettant d'estimer des modèles AR, à moyenne ajustée (MA) et ARMA à partir d'une observation partielle du spectre. L'approche utilisée consiste à optimiser un critère de maximum de vraisemblance pondéré et peut être vue comme une généralisation de la méthode *Discrete All-pole Modeling* (DAP), initialement proposée dans le contexte de la modélisation autorégressive de signaux de parole [El-Jaroudi et Makhoul, 1991]. Les modèles ARMA expriment la forme spectrale $w_f(\boldsymbol{\theta})$ introduite dans la section VI.1 page 22 comme une fraction rationnelle :

$$w_f(\boldsymbol{\theta}) = \left| \frac{1 + \sum_{q=1}^{Q} b^{(q)}\, e^{-i2\pi\nu_f q}}{1 + \sum_{p=1}^{P} a^{(p)}\, e^{-i2\pi\nu_f p}} \right|^2 .$$

Le vecteur de paramètres $\boldsymbol{\theta}$ est alors constitué des coefficients $b^{(q)}$ et $a^{(p)}$. Afin de pouvoir améliorer sélectivement la précision de l'estimation sur certaines fréquences par rapport à d'autres, nous avons introduit une pondération dans le critère (VI.1), en attribuant un poids $\gamma_f \geq 0$ à chaque fréquence f. Cette approche est par exemple intéressante pour estimer l'enveloppe spectrale d'un son (in)harmonique (section VII.1.1), dont le spectre ne contient qu'un sous-ensemble de fréquences pertinentes (il suffit de poser $\gamma_f = 0$ pour toute fréquence f non-harmonique). Pour estimer les paramètres de ce modèle, nous avons proposé des algorithmes multiplicatifs s'inspirant du paradigme présenté dans la section VIII.2.1 page 30. Dans le cas simple de l'estimation de modèles autorégressifs, nous avons démontré que si la fonction de vraisemblance admet un point critique dans le domaine des filtres stables, alors ce point critique est unique et correspond au maximum global. Nous avons aussi prouvé que le maximum global est un point fixe exponentiellement stable (ce qui prouve la convergence locale vers la solution optimale et une vitesse de convergence linéaire) et que la vitesse de convergence est améliorée en remplaçant le modèle AR estimé à chaque itération par le filtre stable équivalent. Enfin, la méthode classique de prédiction linéaire [Priestley, 1983] est un cas particulier de notre estimateur (obtenu en choisissant des poids γ_f égaux à toutes les fréquences). Dans le cas de l'estimation de modèles à moyenne ajustée, la convergence vers la solution optimale a été prouvée en introduisant une hypothèse supplémentaire. Dans le cas des modèles ARMA, la convergence n'a pas été démontrée, mais seulement observée. Dans tous les cas, la convergence vers la solution optimale nécessite que le nombre de paramètres du modèle soit inférieur au nombre de fréquences réduites dans la représentation spectrale. Cette approche a été utilisée pour l'estimation de modèles de mélanges (section VII.2) et pour généraliser la NMF de représentations temps-fréquences (section VIII.3.2.1 page 34).

VII.1.3 Mesures de similarités spectrales

Dans les références [Lagrange *et al.*, 2010c] et [Hennequin *et al.*, 2010c], nous avons proposé deux méthodes visant à évaluer quantitativement la similarité spectrale entre deux sons, afin de comparer leurs timbres. La mesure de la similarité est un problème central dans le domaine de l'extraction d'informations musicales - MIR, puisque plusieurs applications populaires reposent sur des comparaisons entre objets sonores, comme par exemple la recherche d'une séquence musicale dans une grande base de données, la classification automatique d'instruments ou encore la transcription automatique de musique. Les approches classiques consistent à comparer des descripteurs audio (*features*) comme les *Mel-Frequency Cepstral Coefficients* (MFCC) [Rabiner et Juang, 1993], en utilisant une fonction de base radiale - *Radial Basis Function* (RBF) [Buhmann, 2003] pour convertir une mesure de distance en mesure de similarité. Les deux méthodes que nous avons proposées s'expriment quant à elles dans le cadre statistique introduit dans la section VI.1.

VII.1.3.1 Mesures de similarités asymétriques

Dans la référence [Lagrange *et al.*, 2010c], nous avons proposé une nouvelle classe de mesures de similarité, qui comparent un spectre cible à un spectre de référence. L'application visée étant la recherche d'un son dans une grande base de données, nous avons considéré le cas où le spectre de référence (extrait de la base de données) est de bonne qualité acoustique, alors que le spectre cible (qui sert de requête) est dégradé, soit par l'environnement

acoustique, soit par des prétraitements, comme une compression ou une séparation de sources. Par ailleurs, la rapidité de la réponse du système dépend de la complexité de calcul des descripteurs du spectre cible, alors que les descripteurs du spectre de référence peuvent être calculés préalablement. Nous avons donc opté pour des mesures de similarité asymétriques, qui comparent des pics spectraux extraits de la requête (rapidement estimés) à l'enveloppe spectrale de la référence (estimée par la méthode True Envelope qui utilise un modèle cepstral [Röbel *et al.*, 2007]), à l'aide du critère de platitude spectrale (VI.2) page 22. Ces mesures se calculent rapidement et nous avons montré qu'elles s'avèrent robustes à divers types de dégradations acoustiques de la requête (ajout de bruit, filtrage passe-bas, compression, séparation de sources). Leur validation a été réalisée en étudiant leur capacité à déterminer si deux signaux audio ont été joués par le même instrument. Pour cette tâche, nous avons montré qu'elles surpassent les métriques basées sur des descripteurs spectraux standard comme les MFCC. Cette approche s'avère donc prometteuse pour la tâche de classification d'instruments de musique.

VII.1.3.2 Mesure invariante par transposition fréquentielle et changement d'amplitude

L'idée principale de la méthode proposée dans la référence [Hennequin *et al.*, 2010c] est de comparer un spectre cible à un spectre de référence, en utilisant la référence pour définir un modèle statistique $\hat{v}_f$, dont la cible est considérée comme une réalisation v_f. La vraisemblance L du spectre cible est alors calculée dans le but de mesurer la similarité entre les deux spectres. Afin de pouvoir comparer des sons possédant des hauteurs tonales et des intensités sonores différentes, le spectre de référence est adapté au spectre cible par transposition fréquentielle et multiplication par une constante, estimées de façon à maximiser la vraisemblance du modèle. La mesure ainsi obtenue étant asymétrique, nous avons également proposé un modèle conjoint, qui la rend symétrique. Nous avons utilisé cette mesure pour évaluer la similarité entre deux sons simples (c'est à dire des notes isolées, jouées par divers instruments). Nos résultats expérimentaux ont confirmé l'utilité de cette approche, dont avons étudié le potentiel pour la classification automatique d'instruments de musiques (en calculant une matrice de similarité) et l'estimation de hauteurs multiples (en exploitant la transposition fréquentielle estimée).

VII.2 Modélisation de mélanges spectraux

Nous abordons à présent le cas où plusieurs notes de musique sont jouées en même temps. L'identification des notes présentes dans le mélange nécessite d'être capable de séparer leurs différents spectres harmoniques. Cette tâche s'avère d'autant plus difficile que les recouvrements de partiels sont nombreux. L'hypothèse de régularité spectrale joue alors un rôle prépondérant pour estimer correctement l'amplitude de tous les partiels de chaque note. Nous avons ainsi choisi de modéliser le mélange spectral comme une somme de spectres (in)harmoniques (section VII.1.1) plus un spectre de bruit coloré, dont les enveloppes spectrales suivent un modèle paramétrique d'ordre faible afin de garantir la régularité (section VII.1.2). Ce modèle de mélange est estimé en recherchant le maximum de la fonction de vraisemblance (VI.1) ou d'une distribution *a posteriori* dans un cadre bayésien.

VII.2.1 Maximum *a posteriori* pondéré

La première méthode d'estimation d'un mélange de notes que nous avons proposée [Emiya *et al.*, 2007a, Emiya *et al.*, 2010a] a été développée pour une application d'estimation de hauteurs multiples de notes de piano (section IX.1.1.1 page 36). Nous avons choisi de modéliser les enveloppes des spectres harmoniques par des modèles AR, car ces modèles en l'absence de zéros ne peuvent pas décrire des atténuations fortes et localisées en fréquence. Ainsi un spectre possédant de nombreux harmoniques manquants ne peut pas être représenté fidèlement, ce qui le rend peu vraisemblable. On évite ainsi l'erreur classique d'estimation d'un sous-multiple d'une fréquence fondamentale. De même, nous avons choisi de modéliser l'enveloppe spectrale du bruit par un modèle MA, car ce modèle en l'absence de pôles ne peut pas décrire des résonances fortes et localisées en fréquences. Ainsi un spectre de bruit possédant plusieurs partiels résiduels ne peut pas être représenté fidèlement, ce qui le rend peu vraisemblable. On évite ainsi l'erreur classique d'estimation d'un multiple d'une fréquence fondamentale. La méthode d'estimation comprend une première étape de sélection d'un ensemble restreint de hauteurs candidates, obtenues par la méthode du produit spectral normalisé (section VII.1.1). Cette sélection préalable a pour but de réduire la complexité, en évitant d'explorer toutes les combinaisons possibles de fréquences fondamentales, dont la combinatoire serait gigantesque. Dans un deuxième temps on retient, parmi toutes les combinaisons de hauteurs sélectionnées, celle qui maximise un critère de maximum *a posteriori*

pondéré, incluant une approximation de la fonction de vraisemblance (VI.1) et des distributions *a priori* des paramètres. Pour chaque combinaison de hauteurs, le calcul du critère nécessite d'estimer tous les paramètres et en particulier de séparer les spectres harmoniques et bruité. On utilise pour cela un algorithme itératif qui alterne une estimation des enveloppes à partir des spectres séparés (section VII.1.2.2) et une séparation des spectres à partir des enveloppes estimées (on utilise l'estimateur linéaire optimal qui minimise l'erreur quadratique moyenne, conduisant à un filtrage de Wiener qui permet de prendre en compte les recouvrements de partiels).

VII.2.2 Algorithme Espérance-Maximisation

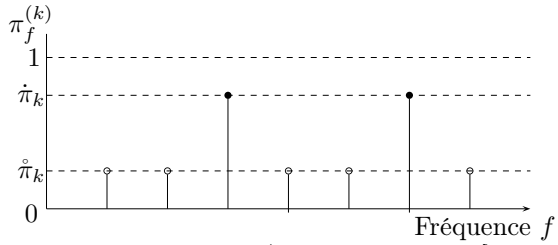

FIG. VII.1 – Modèle d'harmonicité (figure extraite de [Badeau *et al.*, 2009]).

L'inconvénient de l'approche précédente est que la convergence de l'algorithme itératif n'est pas démontrée mais seulement observée et rien ne prouve que la solution obtenue correspond à un maximum local du critère. Nous avons donc proposé dans [Badeau *et al.*, 2009] une seconde approche basée sur l'algorithme Espérance-Maximisation (EM), pour lequel il existe une preuve de convergence vers un maximum local, sous des hypothèses faibles [Jeff Wu, 1983]. L'utilisation de l'algorithme EM a cependant nécessité de modifier légèrement le modèle de mélange en introduisant des variables cachées. Ainsi la propriété d'harmonicité n'est plus modélisée de façon déterministe, mais par l'intermédiaire de variables aléatoires booléennes qui indiquent la présence ou l'absence d'un partiel à chaque fréquence. Pour chaque note indexée par un entier k et de fréquence fondamentale ν_k^0, on suppose ainsi que toutes les fréquences f harmoniques ont une même probabilité élevée de présence d'un partiel $\pi_f^{(k)} = \dot{\pi}_k$ et que toutes les autres fréquences f ont une même probabilité faible de présence d'un partiel $\pi_f^{(k)} = \mathring{\pi}_k$ (avec $0 < \mathring{\pi}_k < \dot{\pi}_k < 1$), comme on peut le voir sur la figure VII.1. Enfin, toutes les enveloppes spectrales sont modélisées par des modèles à moyenne ajustée (les modèles autorégressifs ayant présenté la propriété indésirable de capturer des partiels isolés en créant des résonances fortes et localisées en fréquences). Dans l'algorithme EM ainsi obtenu, l'étape E consiste à mettre à jour les probabilités *a posteriori* des variables cachées en utilisant le théorème de Bayes et l'étape M consiste à maximiser l'espérance conditionnelle de la log-vraisemblance des données complètes connaissant les observations. Cette deuxième étape comprend une mise à jour des probabilités *a priori* (qui revient à estimer les fréquences fondamentales indépendamment les unes des autres) et une mise à jour des enveloppes spectrales, qui comporte une règle multiplicative pour les coefficients de puissance et une technique semblable à celle présentée dans la section VII.1.2.2 pour les coefficients du filtre (avec des poids γ_f correspondant aux probabilités *a posteriori* des variables cachées). Cette approche s'est avérée particulièrement prometteuse, en raison de sa robustesse aux recouvrements de partiels et de sa capacité à simplifier la tâche d'estimation de hauteurs multiples en plusieurs estimations successives de hauteurs simples et d'enveloppes spectrales. En revanche, elle tend à rester piégée dans des maxima locaux de la fonction de vraisemblance. Elle nécessite donc une initialisation adaptée, utilisant par exemple un algorithme simple d'estimation de hauteurs multiples, et pourrait avantageusement tirer parti de règles heuristiques pour explorer plus efficacement le domaine des fréquences fondamentales. Son efficacité a été confirmée par des simulations numériques, réalisées sur des signaux synthétiques. Sa supériorité par rapport à la méthode présentée dans la section VII.2.1 tient à sa capacité à modifier automatiquement la combinaison de fréquences fondamentales testées à chaque itération (avec la possibilité d'effectuer des sauts de fréquences), ce qui permet de réduire proprement la combinatoire, en évitant d'estimer complètement tous les modèles de mélange correspondant à toutes les combinaisons de fréquences fondamentales possibles. Cependant, son application à des signaux réels nécessitera plusieurs améliorations qui avaient déjà été intégrées à la méthode présentée dans la section VII.2.1, comme la prise en compte de l'inharmonicité du piano et de la forme spectrale de la fenêtre temporelle utilisée, l'ajout d'informations *a priori* sur la distribution des paramètres dans un cadre bayésien et l'estimation automatique de la polyphonie (nombre de notes superposées) via l'ajout d'un terme de pénalité, similaire aux critères de théorie de l'information [Schwarz, 1978].

Chapitre VIII

Décompositions temps-fréquence

Dans ce chapitre, nous présentons différentes méthodes de décompositions inspirées de la factorisation en matrices non-négatives que nous avons mentionnée dans la section VI.2 page 23. Par rapport au chapitre précédent, il ne s'agit plus de décompositions de spectres de sons stationnaires, mais de décompositions bi-dimensionnelles de représentations temps-fréquences de signaux non-stationnaires. Nous commençons par une brève bibliographie sur la factorisation en matrices non-négatives (section VIII.1).

VIII.1 Factorisation en matrices non-négatives

VIII.1.1 Cadre déterministe

La NMF a été introduite par Lee et Sung pour décomposer des données bi-dimensionnelles non-négatives en une combinaison linéaire d'éléments dans un dictionnaire [Lee et Seung, 1999].

VIII.1.1.1 Définition et fonction objectif

Étant donnée une matrice de données $\boldsymbol{V}$ de dimensions $F \times N$ dont les coefficients sont non-négatifs, le problème de NMF consiste à calculer une approximation $\widehat{\boldsymbol{V}}$ de la matrice $\boldsymbol{V}$ tronquée au rang $K < \min(F, N)$, exprimée comme un produit $\widehat{\boldsymbol{V}} = \boldsymbol{W}\boldsymbol{H}$, où les deux matrices $\boldsymbol{W}$ de dimensions $F \times K$ et $\boldsymbol{H}$ de dimensions $K \times N$ ont des coefficients non-négatifs. Les colonnes de la matrice $\boldsymbol{W}$ constituent les éléments du dictionnaire et les lignes de $\boldsymbol{H}$ contiennent les coefficients de la décomposition. La dimension K est généralement choisie de telle sorte que $FK + KN \ll FN$, de façon à réduire la dimension des données. La NMF peut être considérée comme une technique d'apprentissage supervisée ou non-supervisée. Dans le cas de l'apprentissage supervisé, le dictionnaire $\boldsymbol{W}$ est préalablement estimé à partir de données d'apprentissage et la matrice $\boldsymbol{H}$ seulement doit être calculée à partir de la matrice $\boldsymbol{V}$. Dans le cas de l'apprentissage non-supervisé, les deux matrices $\boldsymbol{W}$ et $\boldsymbol{H}$ doivent être calculées conjointement à partir de $\boldsymbol{V}$. Dans les applications audio, $\boldsymbol{V}$ est souvent le spectrogramme d'amplitude ou de puissance, f désigne le canal fréquentiel et n la fenêtre temporelle. La figure VIII.1 représente la partition musicale, le spectrogramme et la NMF non supervisée de la mélodie d'"Au clair de la lune". Cette figure met clairement en évidence l'intérêt d'une telle décomposition : celle-ci fait apparaître les spectres des notes de musique dans la matrice $\boldsymbol{W}$ et leurs activations temporelles dans la matrice $\boldsymbol{H}$, ce qui permet d'envisager à la fois des applications de transcription et de séparation des notes de musique.

Le problème de NMF peut être formalisé comme la minimisation de la fonction objectif

$$D(\boldsymbol{V}|\boldsymbol{W}\boldsymbol{H}) = \sum_{n=1}^{N}\sum_{f=1}^{F} d\left(v_{fn}\,|\hat{v}_{fn}\right), \tag{VIII.1}$$

où $\hat{v}_{fn} = \sum_{k=1}^{K} w_{fk}h_{kn}$ et d est une divergence scalaire, c'est à dire une fonction telle que $\forall a, b \in \mathbb{R}_+, d(a|b) \geq 0$ et $d(a|b) = 0$ si et seulement si $b = a$. Si de plus la divergence $d(a|b)$ est convexe par rapport à b, alors la fonction objectif $D(\boldsymbol{V}|\boldsymbol{W}\boldsymbol{H})$ est elle-même convexe par rapport à $\boldsymbol{W}$ et par rapport à $\boldsymbol{H}$ (mais elle n'est pas convexe par rapport au couple $(\boldsymbol{W}, \boldsymbol{H})$). Les fonctions objectif les plus populaires pour la NMF sont la

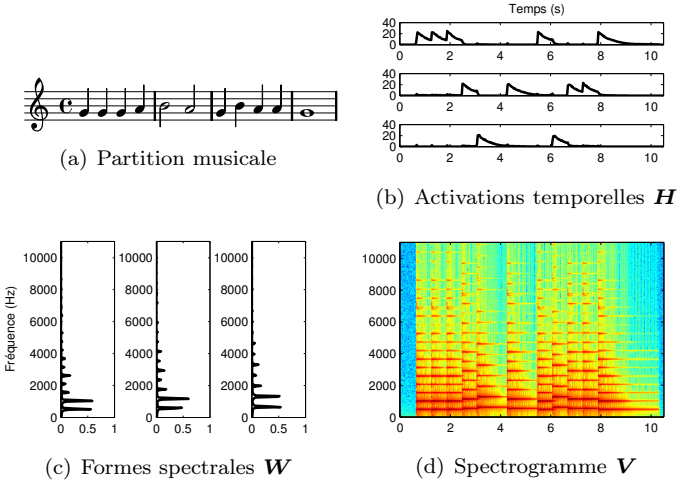

(a) Partition musicale

(b) Activations temporelles $\boldsymbol{H}$

(c) Formes spectrales $\boldsymbol{W}$

(d) Spectrogramme $\boldsymbol{V}$

FIG. VIII.1 – Décomposition du spectrogramme d'"Au clair de la Lune" (extraite de [Bertin, 2009, pp. 40–41])

distance euclidienne (EUC) et la divergence généralisée de Kullback-Leibler (KL), qui ont été particulièrement popularisées par Lee et Sung [Lee et Seung, 1999]. L'usage de la divergence d'Itakura-Saito dans les applications audio a été proposé dans [Févotte *et al.*, 2009]. Enfin, les **$\boldsymbol{\beta}$-divergences** [Eguchi et Kano, 2001] généralisent les trois divergences précédentes. Elles sont définies pour tout $\beta \in \mathbb{R}\backslash\{0,1\}$ comme

$$d_\beta(a|b) = \frac{1}{\beta(\beta-1)}\left(a^\beta + (\beta-1)b^\beta - \beta a b^{\beta-1}\right). \tag{VIII.2}$$

Dans la référence [Hennequin *et al.*, 2010e], nous avons présenté une preuve démontrant que les $\boldsymbol{\beta}$-divergences sont une sous-classe des divergences de Bregman [Bregman, 1967]. La distance euclidienne correspond à $\beta = 2$ et les divergences KL et IS sont respectivement obtenues par prolongement par continuité lorsque $\beta \to 1$ et $\beta \to 0$. Par ailleurs, on remarque que la fonction $d_\beta(a|b)$ est convexe par rapport à b si et seulement si $\beta \in [1,2]$. Ainsi les divergences EUC et KL sont convexes mais IS ne l'est pas.

VIII.1.1.2 Algorithmes

Diverses stratégies ont été proposées pour calculer une NMF. La plus populaire est l'algorithme de mises à jour multiplicatives proposé initialement par Lee et Seung [Lee et Seung, 2001] pour les divergences EUC et KL, qui a ensuite été généralisé à la β-divergence [Cichocki *et al.*, 2006, Kompass, 2007]. D'autres approches ont également été proposées, telles que la méthode du gradient projeté [Lin, 2007], les algorithmes de moindres carrés alternés [Finesso et Spreij, 2004] et l'algorithme quasi-Newton [Cichocki *et al.*, 2008]. Une bibliographie détaillée peut être consultée dans la référence [Cichocki *et al.*, 2008]. Nous nous sommes concentrés ici sur les règles multiplicatives, en raison de leur simplicité et de leur vitesse de convergence.

Les règles multiplicatives découlent de la même heuristique que les algorithmes multiplicatifs définis dans la section VIII.2.1 et conservent en particulier la non-négativité des facteurs $\boldsymbol{W}$ et $\boldsymbol{H}$. Dans le cas des β-divergences, elles s'expriment sous la forme suivante [Kompass, 2007] :

$$\boldsymbol{W} \leftarrow \boldsymbol{W} \otimes \frac{(\boldsymbol{V} \otimes (\boldsymbol{W}\boldsymbol{H})^{\beta-2})\boldsymbol{H}^T}{(\boldsymbol{W}\boldsymbol{H})^{\beta-1}\boldsymbol{H}^T} \tag{VIII.3}$$

$$\boldsymbol{H} \leftarrow \boldsymbol{H} \otimes \frac{\boldsymbol{W}^T(\boldsymbol{V} \otimes (\boldsymbol{W}\boldsymbol{H})^{\beta-2})}{\boldsymbol{W}^T(\boldsymbol{W}\boldsymbol{H})^{\beta-1}} \tag{VIII.4}$$

où le symbole $\otimes$ et la barre de fraction désignent respectivement le produit et la division matriciels terme à terme et les exponentiations sont également terme à terme. Dans le cas de la NMF non-supervisée, les règles (VIII.3) et (VIII.4) sont appliquées alternativement, alors que dans le cas de la NMF supervisée, la règle (VIII.4)

seulement est appliquée à chaque itération, la matrice $\boldsymbol{W}$ étant inchangée. Dans [Kompass, 2007], il est prouvé que si $\beta \in [1, 2]$, alors la fonction objectif décroît à chaque itération de (VIII.3) et (VIII.4). Cependant, cette propriété de décroissance ne garantit pas que la limite de la fonction objectif est un minimum local, ni que les valeurs successives de $\boldsymbol{W}$ et $\boldsymbol{H}$ convergent.

VIII.1.2 Cadres statistiques

De récents travaux ont conduit à formuler le problème de NMF comme un problème d'estimation paramétrique. Plus précisément, il a été prouvé que la minimisation de la fonction objectif (VIII.1) avec chacune des trois divergences les plus courantes (EUC, KL et IS) est équivalente à l'estimation au sens du maximum de vraisemblance de différents modèles statistiques. Ainsi la distance euclidienne correspond à l'estimation d'un modèle déterministe de NMF avec bruit additif gaussien [Schmidt et Laurberg, 2008], la divergence de Kullback-Leibler correspond à l'estimation d'un modèle de mélange de variables aléatoires suivant une loi de Poisson [Virtanen *et al.*, 2008] ou d'un modèle de mélange de composantes latentes indépendantes [Shashanka *et al.*, 2008], et la divergence d'Itakura-Saito correspond à l'estimation d'un modèle déterministe de NMF avec bruit multiplicatif de loi Gamma ou d'un modèle de mélange de variables aléatoires gaussiennes [Févotte *et al.*, 2009]. Ces cadres statistiques ont ouvert de nombreuses perspectives, puisqu'ils permettent d'incorporer aisément des informations sur les propriétés particulières des composantes à extraire, à l'aide de modèles paramétriques ou de distributions *a priori* dans le formalisme de l'inférence bayésienne. Ils ont également conduit au développement de nouveaux algorithmes de NMF, appartenant à la famille des algorithmes EM, comme l'algorithme *Space Alternating Generalized EM* (SAGE) [Fessler et Hero, 1994].

Le modèle de mélange de variables aléatoires gaussiennes [Févotte *et al.*, 2009] correspond à celui qui a été présenté dans le chapitre VI. Par rapport aux autres cadres statistiques mentionnés ci-dessus, ce modèle est le seul à tenir compte de la présence des phases dans la représentation temps-fréquence. Cette approche offre donc un formalisme commode pour séparer les composantes $x^{(k)}(t)$. On montre en effet que les estimateurs des composantes qui minimisent l'erreur quadratique moyenne sont obtenus par filtrage de Wiener. C'est notamment pour cette raison que nos propres travaux se sont concentrés sur ce modèle. Plus récemment, nous avons approfondi le lien entre les processus gaussiens et la NMF dans les références [Liutkus *et al.*, 2010a, Liutkus *et al.*, 2011], où les observations ne sont plus les coefficients d'une représentation temps-fréquence, mais directement les échantillons du signal $x(t)$.

VIII.2 Minimisation de la fonction objectif de la NMF

VIII.2.1 Algorithmes multiplicatifs

Les règles de mises à jour multiplicatives ont connu un grand succès pour résoudre des problèmes d'optimisation avec contraintes de positivité, tels que la NMF et ses nombreuses variantes. Néanmoins, malgré plusieurs années de recherche sur le sujet, la compréhension de leurs propriétés de convergence demeure imparfaite [Badeau *et al.*, 2010a]. Dans les références [Badeau *et al.*, 2010b, Badeau *et al.*, 2010c, Badeau *et al.*, 2010a], nous avons montré que la théorie de la stabilité de Lyapunov [Lasalle, 1986] fournit un point de vue très instructif sur le problème. Elle permet sous certaines hypothèses de prouver la **stabilité asymptotique** d'un point fixe $\boldsymbol{x}$ de l'algorithme (propriété selon laquelle il existe un voisinage de $\boldsymbol{x}$ dans lequel la convergence vers $\boldsymbol{x}$ est garantie) ou la **stabilité exponentielle** (qui garantit, en plus de la stabilité asymptotique, une vitesse de convergence linéaire).

VIII.2.1.1 Cadre théorique

On considère le problème de minimisation d'une fonction objectif $J : \Omega \subset \mathbb{R}_+^l \to \mathbb{R}$ de classe $\mathcal{C}^2$ (avec $l \in \mathbb{N}^*$). On suppose que le gradient de J peut être décomposé comme la différence de deux fonctions non-négatives : $\boldsymbol{\nabla} J(\boldsymbol{x}) = \boldsymbol{p}(\boldsymbol{x}) - \boldsymbol{m}(\boldsymbol{x})$, où les deux fonctions $\boldsymbol{p} : \Omega \to \mathbb{R}_+^l$ et $\boldsymbol{m} : \Omega \to \mathbb{R}_+^l$ sont de classe $\mathcal{C}^1$. Étant donné un pas $\eta \in \mathbb{R}$, on considère une fonction de récurrence $\boldsymbol{\phi}$ définie pour tout $\boldsymbol{x} \in \Omega$ tel que $\forall i \in \{1 \dots l\}$, $p_i(\boldsymbol{x}) > 0$ et $m_i(\boldsymbol{x}) > 0$: $\boldsymbol{\phi}(\boldsymbol{x}) = \operatorname{diag}\left(\frac{\boldsymbol{m}(\boldsymbol{x})}{\boldsymbol{p}(\boldsymbol{x})}\right)^{\eta} \boldsymbol{x}$. On remarque en particulier que pour tout $\boldsymbol{x} \in \Omega$ tel que $\forall i \in \{1 \dots l\}$, $p_i(\boldsymbol{x}) > 0$ et $m_i(\boldsymbol{x}) > 0$, la fonction $\boldsymbol{\phi}$ est définie et de classe $\mathcal{C}^1$ sur un voisinage de $\boldsymbol{x}$ dans Ω. Si de plus $\boldsymbol{x}$ est minimum local de J, alors $\boldsymbol{x}$ est un point fixe de $\boldsymbol{\phi}$.

Dans les références [Badeau *et al.*, 2010b, Badeau *et al.*, 2010c, Badeau *et al.*, 2010a], nous avons appelé **algorithme multiplicatif** la récurrence $\boldsymbol{x}^{(p+1)} = \boldsymbol{\phi}(\boldsymbol{x}^{(p)})$ et nous avons prouvé que cette récurrence constitue une méthode de descente (au même titre que peuvent l'être les méthodes de gradient). Nous avons également analysé ses propriétés de convergence. Nous avons ainsi appliqué la première et la seconde méthode de Lyapunov pour trouver des critères qui garantissent la stabilité exponentielle ou asymptotique des minima locaux de J, soit en analysant les valeurs propres de la matrice Jacobienne de $\boldsymbol{\phi}$, soit en introduisant une fonction de Lyapunov adéquate. Nous avons observé que la stabilité asymptotique conduit généralement à une vitesse de convergence sous-linéaire (selon l'initialisation). Nous avons également fourni l'expression analytique d'une borne supérieure $\eta^\star$ telle que la stabilité exponentielle ou asymptotique est garantie pour tout $\eta \in]0, \eta^\star[$. Cette étude a été directement appliquée à la NMF supervisée (section VIII.2.1.2).

VIII.2.1.2 Algorithmes multiplicatifs de NMF

Dans les références [Badeau *et al.*, 2010b, Badeau *et al.*, 2010c, Badeau *et al.*, 2010a, Badeau *et al.*, 2011], nous avons généralisé les règles (VIII.3) et (VIII.4) en appliquant le cadre théorique présenté dans la section VIII.2.1.1 au critère $J = D(\boldsymbol{V}|\boldsymbol{W}\boldsymbol{H})$ introduit dans l'équation (VIII.1) et utilisant la β-divergence définie dans l'équation (VIII.2). On introduit ainsi un pas d'exponentiation η :

$$\boldsymbol{W} \leftarrow \boldsymbol{W} \otimes \left(\frac{(\boldsymbol{V} \otimes (\boldsymbol{W}\boldsymbol{H})^{\beta-2})\boldsymbol{H}^T}{(\boldsymbol{W}\boldsymbol{H})^{\beta-1}\boldsymbol{H}^T} \right)^\eta \tag{VIII.5}$$

$$\boldsymbol{H} \leftarrow \boldsymbol{H} \otimes \left(\frac{\boldsymbol{W}^T(\boldsymbol{V} \otimes (\boldsymbol{W}\boldsymbol{H})^{\beta-2})}{\boldsymbol{W}^T(\boldsymbol{W}\boldsymbol{H})^{\beta-1}} \right)^\eta \tag{VIII.6}$$

Comme nous l'avons montré, l'intérêt du paramètre η est qu'il permet de contrôler le taux de convergence et en particulier de surpasser le cas standard $\eta = 1$ (qui correspond aux règles (VIII.3) et (VIII.4)). Nous avons démontré que ces règles satisfont la même propriété de décroissance de la fonction objectif que (VIII.3) et (VIII.4) si $\beta \in [1, 2]$ et pour tout $\eta \in]0, 1]$. Nous avons également appliqué la méthode générale d'analyse de stabilité des algorithmes multiplicatifs présentée dans la section VIII.2.1.1. Dans le cas de la NMF supervisée, nous avons ainsi présenté des critères qui garantissent la stabilité exponentielle ou asymptotique des règles multiplicatives pour tout $\eta \in]0, \eta^\star[$, où $\forall \beta \in \mathbb{R}$, $\eta^\star \in]0, 2]$ et si $\beta \in [1, 2]$, $\eta^\star = 2$. Analyser les propriétés de stabilité dans le cas non-supervisé s'est avéré particulièrement difficile, à cause des invariances de la factorisation, qui rendent les minima locaux de la fonction objectif non-isolés, donc non-asymptotiquement stables. Néanmoins, la première méthode de Lyapunov a fourni des résultats instructifs sur les propriétés de convergence de ces algorithmes. En particulier, nous avons prouvé l'instabilité des règles multiplicatives si $\eta \notin [0, 2]$. Finalement, les résultats théoriques ont été confirmés par des simulations numériques faisant intervenir la NMF supervisée et non-supervisée. Ces simulations ont montré que le taux de convergence dépend de la valeur du pas η et qu'il existe une valeur optimale de η (généralement différente de 1), qui fournit le taux de convergence le plus rapide (en particulier, plus rapide que celui des règles multiplicatives (VIII.3) et (VIII.4)).

VIII.2.2 Évitement des minima locaux

Dans la section précédente, nous avons examiné les propriétés de convergence des algorithmes multiplicatifs vers un minimum local de la fonction objectif. Nous nous intéressons à présent à la recherche d'un minimum global, qui nécessite de développer des stratégies pour éviter de rester piégé dans un minimum local indésirable.

VIII.2.2.1 Initialisation

Dans la référence [Bertin et Badeau, 2008], nous avons étudié d'une part la présence de minima locaux pour les trois divergences les plus populaires (EUC, KL et IS) et d'autre part les performances de différentes stratégies d'initialisations, dans une application de transcription de musique polyphonique. Nous avons ainsi observé que les trois fonctions objectif possèdent des minima locaux, même la distance euclidienne, dont nous supposions qu'elle pouvait en être exempte. Cette distance présente néanmoins le nombre le plus faible de minima locaux, alors que la divergence IS en présente le nombre le plus élevé (ce à quoi nous nous attendions puisque les divergences EUC et KL sont convexes alors que IS ne l'est pas). Deux techniques d'initialisation ont été testées :

- un algorithme de K-moyennes sphériques [Wild *et al.*, 2004] ;
- un algorithme de partitionnement de données hiérarchique, construit à partir d'une mesure de "proximité au rang 1" (*closeness to rank one*) [Kim et Choi, 2007].

Nous avons observé qu'aucune de ces deux techniques ne garantit la convergence vers un éventuel minimum global. Devant l'échec de l'approche par initialisation structurée, nous avons proposé une autre approche d'évitement des minima locaux, utilisant un algorithme de "refroidissement simulé" (section VIII.2.2.2).

VIII.2.2.2 Algorithme de "refroidissement simulé"

Dans la référence [Bertin *et al.*, 2009c], nous nous sommes concentrés sur l'algorithme standard de NMF ($\eta = 1$) utilisant la divergence d'Itakura-Saito ($\beta = 0$). Des résultats antérieurs [Févotte *et al.*, 2009] ont en effet montré les bonnes propriétés de cette divergence pour la décomposition de spectrogrammes de signaux audio. Cela est dû en particulier à sa propriété d'invariance par changement d'échelle, qui rend la fonction objectif robuste à la forte dynamique spectrale des signaux audio, propriété qui n'est pas partagée par EUC et KL. En raison des nombreux minima locaux néanmoins observés avec la divergence IS, nous avons proposé une nouvelle approche pour favoriser la convergence vers le minimum global. L'algorithme est basé sur l'utilisation de la β-divergence, où β agit comme un paramètre de température. Ce paramètre est ainsi fixé à une valeur comprise entre 1 et 2 lors des premières itérations (la fonction objectif est alors convexe par rapport à $\boldsymbol{W}$ et par rapport à $\boldsymbol{H}$), puis il décroit progressivement jusqu'à $\beta = 0$, valeur à laquelle il reste fixé pendant les dernières itérations. L'algorithme obtenu s'exprime sous la forme suivante :

$$\begin{aligned} \boldsymbol{W} &\leftarrow \boldsymbol{W} \otimes \frac{(\boldsymbol{V} \otimes (\boldsymbol{W}\boldsymbol{H})^{\beta(p)-2})\boldsymbol{H}^T}{(\boldsymbol{W}\boldsymbol{H})^{\beta(p)-1}\boldsymbol{H}^T} \\ \boldsymbol{H} &\leftarrow \boldsymbol{H} \otimes \frac{\boldsymbol{W}^T(\boldsymbol{V} \otimes (\boldsymbol{W}\boldsymbol{H})^{\beta(p)-2})}{\boldsymbol{W}^T(\boldsymbol{W}\boldsymbol{H})^{\beta(p)-1}} \end{aligned}$$

où β est fonction de l'indice d'itération p. Nos résultats expérimentaux sur des signaux synthétiques ont confirmé notre intuition que cette approche, en particulier l'exploitation de la propriété de convexité des β-divergences pour $1 \leq \beta \leq 2$ dans les premières itérations de l'algorithme, permet d'atteindre une valeur plus basse de la fonction objectif, en évitant donc de rester piégé dans des minima locaux. La valeur initiale optimale de β semble néanmoins dépendre du signal et son choix reste une question ouverte. Les bonnes propriétés de la divergence d'Itakura-Saito et de la stratégie de refroidissement simulé ont été confirmées sur une base de données d'enregistrements musicaux. Dans une application de transcription musicale, nous avons cependant remarqué que la meilleure performance de transcription ne correspond pas toujours à la valeur la plus faible de la fonction objectif, ce qui montre que la fonction objectif choisie n'est pas totalement pertinente pour cette application. Pour améliorer la performance de transcription, il est ainsi apparu que la stratégie à suivre n'est pas de chercher absolument à atteindre le minimum global de la fonction objectif, mais plutôt de contraindre la solution en lui imposant des propriétés telles que l'harmonicité ou la régularité temporelle (section VIII.3).

VIII.3 Ajout de contraintes

Dans la littérature, de nombreux auteurs ont constaté que la NMF ne faisait pas toujours ressortir les propriétés attendues dans les données analysées ; différentes stratégies ont donc été proposées pour contraindre la factorisation. Les premiers travaux en ce sens ont été effectués dans un cadre déterministe, en ajoutant un terme de pénalité à la fonction objectif usuelle de NMF destiné à favoriser les propriétés souhaitées dans la factorisation, comme la régularité temporelle [Virtanen, 2007, Chen *et al.*, 2006], la parcimonie [Virtanen, 2007, Raczyński *et al.*, 2007], ou la décorrélation [Chen *et al.*, 2006, Zhang et Fang, 2007, Raczyński *et al.*, 2007]). A chaque fois, des algorithmes multiplicatifs ont été proposés pour minimiser la nouvelle fonction objectif ainsi obtenue. Ces algorithmes ayant pour propriété de laisser inchangés les coefficients nuls, il a également été proposé d'initialiser à 0 certains coefficients dans la matrice $\boldsymbol{W}$ pour contraindre l'harmonicité spectrale [Raczyński *et al.*, 2007]. Des travaux plus récents ont été développés dans un cadre d'inférence bayésienne (section VIII.1.2), en introduisant des distributions *a priori* adaptées, destinées à renforcer dans $\boldsymbol{W}$ les propriétés d'harmonicité [Virtanen *et al.*, 2008], de parcimonie et de régularité [Schmidt et Laurberg, 2008] et dans $\boldsymbol{H}$ les propriétés de parcimonie [Schmidt et Laurberg, 2008, Smaragdis, 2009] et de régularité [Virtanen *et al.*, 2008, Schmidt et Laurberg, 2008, Févotte *et al.*, 2009]. Par ailleurs, plusieurs auteurs ont proposé d'intégrer des critères psycho-acoustiques dans la factorisation, par exemple en utilisant une échelle de fréquences adaptée à l'au-

dition humaine [Nikunen et Virtanen, 2010, Smaragdis, 2009] et/ou en introduisant des poids perceptifs dans la fonction objectif de la NMF, afin d'améliorer la qualité acoustique de la factorisation [Nikunen et Virtanen, 2010]. Enfin, notons que d'autres généralisations de la NMF ont été proposées en introduisant des produits de convolutions, comme l'utilisation d'un dictionnaire d'éléments temps-fréquences dans un cadre déterministe [Smaragdis, 2004] ou la modélisation de la transposition fréquentielle [Smaragdis, 2009] et de la réverbération [Singh *et al.*, 2010] dans un cadre d'analyse en composantes latentes indépendantes [Shashanka *et al.*, 2008].

Nos propres travaux ont reposé sur le modèle d'observations gaussiennes introduit dans le chapitre VI, dont nous avons proposé des extensions permettant par exemple de renforcer conjointement l'harmonicité spectrale et la régularité temporelle. Outre l'utilisation de distributions *a priori*, nous avons également défini des modèles paramétriques pour $\boldsymbol{W}$ et pour $\boldsymbol{H}$, destinés à contraindre la factorisation, voire à la généraliser (section VIII.3.2.1).

VIII.3.1 Harmonicité, régularité spectrale et régularité temporelle

VIII.3.1.1 Harmonicité et régularité spectrale

Dans les références [Vincent *et al.*, 2008, Vincent *et al.*, 2010], nous avons proposé un modèle destiné à contraindre conjointement l'harmonicité et la régularité des formes spectrales. Cette technique s'appuie sur un découpage en M sous-bandes du spectre, dans le cadre d'un modèle linéaire (section VII.1.2.1 page 24) qui permet d'adapter l'enveloppe spectrale à chaque composante. Les formes spectrales $w_{fk}(\boldsymbol{\theta})$ introduites dans la section VI.2 page 23 sont ainsi modélisées comme une combinaison linéaire à coefficients positifs e_{mk} de motifs spectraux harmoniques $P_{km}(f)$, dont le support fréquentiel correspond à la sous-bande m : $w_{fk}(\boldsymbol{\theta}) = \sum_{m=1}^{M} e_{mk} P_{km}(f)$. L'ensemble de paramètres $\boldsymbol{\theta}$ est alors constitué des coefficients e_{mk}. Les motifs $P_{km}(f)$ appartiennent à un dictionnaire prédéfini, alors que les coefficients e_{mk} sont estimés à l'aide d'un algorithme multiplicatif à partir du signal à décomposer, en même temps que les activations temporelles h_{kn}. Afin de garantir l'harmonicité de w_{fk}, pour toutes les sous-bandes m les motifs $P_{km}(f)$ partagent la même fréquence fondamentale ν_k^0. De plus la répartition des sous-bandes est uniforme sur l'échelle fréquentielle utilisée, qui peut être linéaire, logarithmique ou encore adaptée à l'échelle ERB, motivée par un critère perceptif [Moore et Glasberg, 1983]. Différentes formes de sous-bandes ont été testées, correspondant respectivement à des fenêtres rectangulaire, triangulaire, Hann et *gammatone* (également motivée par un critère perceptif [Johannesma, 1972]). Les motifs $P_{km}(f)$ sont soit appris à partir d'une base de données annotée, soit modélisés comme une somme de plusieurs partiels consécutifs à des fréquences harmoniques de ν_k^0, dont les amplitudes sont définies par la forme de la sous-bande m. Pour estimer ce modèle, deux fonctions objectif ont été testées : une distance euclidienne pondérée par des poids perceptifs [Vincent *et al.*, 2008] et la β-divergence définie dans l'équation (VIII.2) [Vincent *et al.*, 2010]. Dans une application d'estimation de hauteurs multiples (section IX.1.1.2 page 37), les meilleures performances ont été obtenues avec la β-divergence, en utilisant conjointement l'échelle ERB et les fenêtres gammatones. Par ailleurs l'apprentissage des motifs $P_{km}(f)$ n'a pas apporté d'amélioration par rapport au modèle harmonique. Enfin, nous avons montré que la valeur optimale de β est voisine de 0.5 et donc différente des valeurs entières généralement utilisées dans la littérature.

VIII.3.1.2 Harmonicité et régularités spectrale et temporelle

L'estimation au sens du maximum de vraisemblance des coefficients e_{mk} introduits dans la section VIII.3.1.1 et des activations temporelles h_{kn} a ouvert la possibilité de contraindre les solutions de la NMF en incluant des informations *a priori* sur les paramètres [Bertin *et al.*, 2010, Bertin *et al.*, 2009b]. Comme dans la référence [Févotte *et al.*, 2009], ce cadre théorique est exploité pour renforcer les régularités des lignes de la matrice $\boldsymbol{H}$. On fournit une information *a priori* sur l'ensemble de paramètres $\boldsymbol{\Theta}$ défini dans la section VI.2 page 23, exprimée comme une distribution *a priori* $p(\boldsymbol{\Theta})$. Grâce à la règle de Bayes, on obtient un estimateur MAP en maximisant la fonction objectif $C_{MAP}(\boldsymbol{\Theta}) = L(\boldsymbol{\Theta}) + \log(p(\boldsymbol{\Theta}))$, où l'expression de la vraisemblance $L(\boldsymbol{\Theta})$ a été donnée dans l'équation (VI.3). Pour la distribution *a priori*, on choisit une structure en chaîne de Markov similaire à celle proposée dans [Févotte *et al.*, 2009] : $p(\boldsymbol{H}) = \prod_{k=1}^{K} p(h_{k1}) \prod_{n=2}^{N} p(h_{kn}|h_{k(n-1)})$, où $p(h_{kn}|h_{k(n-1)})$ atteint son maximum en $h_{k(n-1)}$, ce qui favorise une variation lente de h_{kn} au cours du temps. Nous avons proposé la loi $p(h_{kn}|h_{k(n-1)}) = \mathcal{IG}(h_{kn}|\alpha_k, (\alpha_k+1)h_{k(n-1)})$, où $\mathcal{IG}(u|\alpha,\beta)$ est la distribution inverse-Gamma[1] de mode $\beta/(\alpha+1)$ et la distribution initiale $p(h_{k1})$ est la distribution *a priori* non-informative de Jeffrey :

[1] Distribution inverse-Gamma : $\mathcal{IG}(u|\alpha,\beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} u^{-(\alpha+1)} \exp(-\frac{\beta}{u})$, $u \geq 0$.

$p(h_{k1}) \propto 1/h_{k1}$. Les paramètres α_k sont fixés arbitrairement, selon le degré de régularité souhaité (plus α_k est élevé, plus h_{kn} est régulier), mais il serait envisageable de les apprendre. On ne considère en revanche pas d'information *a priori* sur l'ensemble $\boldsymbol{\theta}$ des paramètres e_{mk}. Pour calculer la NMF avec les contraintes d'harmonicité et de régularités spectrale et temporelle, nous avons d'abord proposé un algorithme de type SAGE pour maximiser la fonction objectif. Le principal avantage de SAGE est que ses propriétés de convergence sont bien connues [Fessler et Hero, 1994]. En revanche sa vitesse de convergence est plus lente que celle des algorithmes multiplicatifs. Le temps de calcul a donc été réduit en initialisant le système avec une variante originale de l'algorithme multiplicatif de NMF avec contrainte d'harmonicité (section VIII.3.1.1). Plus récemment, nous avons proposé une méthode d'estimation alternative, utilisant un algorithme multiplicatif qui prend en charge l'estimation de tous les paramètres de façon autonome [Bertin *et al.*, 2009b]. Ces deux algorithmes ont été utilisés dans le cadre d'une application de transcription automatique de musique de piano (section IX.1.2.3 page 38).

VIII.3.2 Instationnarités

Un autre inconvénient de la NMF est qu'elle n'est pas conçue pour représenter des éléments dont le contenu spectral varie au cours du temps (autrement que par un simple facteur multiplicatif). Si l'on souhaite représenter de tels éléments, une première solution consiste à surdimensionner la factorisation afin de consacrer plusieurs composantes à chaque élément. Cependant, quand la NMF n'est pas utilisée à des fins de synthèse sonore mais plutôt à des fins d'analyse sémantique, cette approche soulève le problème du regroupement de ces composantes en éléments distincts. Deux approches alternatives sont présentées ci-dessous, modélisant les variations temporelles des enveloppes spectrales (section VIII.3.2.1) et des fréquences fondamentales (section VIII.3.2.2).

VIII.3.2.1 Activations temps-fréquences

Certains sons naturels, comme les sons de clavecin ou les sons de parole dans une transition entre deux voyelles, possèdent une hauteur tonale constante mais une enveloppe spectrale variant au cours du temps. Pour modéliser cette variation, nous proposons dans [Hennequin *et al.*, 2010b, Hennequin *et al.*, 2010a] une extension de la NMF, où l'évolution spectrale est inclue dans l'activation temporelle, en rendant celle-ci dépendante de la fréquence. Le modèle $\hat{v}_{fn}$ introduit dans la section VI.2 page 23 s'exprime alors sous la forme $\hat{v}_{fn} = \sum_{k=1}^{K} w_{fk} h_{kn}(f)$. Cette dépendance fréquentielle est paramétrée afin d'éviter de trop accroître la dimension du modèle par rapport à celle des données. L'activation temps-fréquence est ainsi modélisée par un filtre ARMA $h_{kn}(f)$, qui permet de représenter une large classe de réponses en fréquences :

$$h_{kn}(f) = \sigma_{kn}^2 \frac{\left|1 + \sum_{q=1}^{Q} b_{kn}^{(q)} e^{-i2\pi\nu_f q}\right|^2}{\left|1 + \sum_{p=1}^{P} a_{kn}^{(p)} e^{-i2\pi\nu_f p}\right|^2}$$

où les coefficients $b_{kn}^{(q)}$ et $a_{kn}^{(p)}$ caractérisent respectivement les parties à moyenne ajustée et autorégressive du filtre, dont le gain global est σ_{kn}^2. Dans le cas particulier où $P = Q = 0$, $h_{kn}(f)$ ne dépend plus de f, ce qui correspond à une NMF standard avec des activations temporelles σ_{kn}^2.

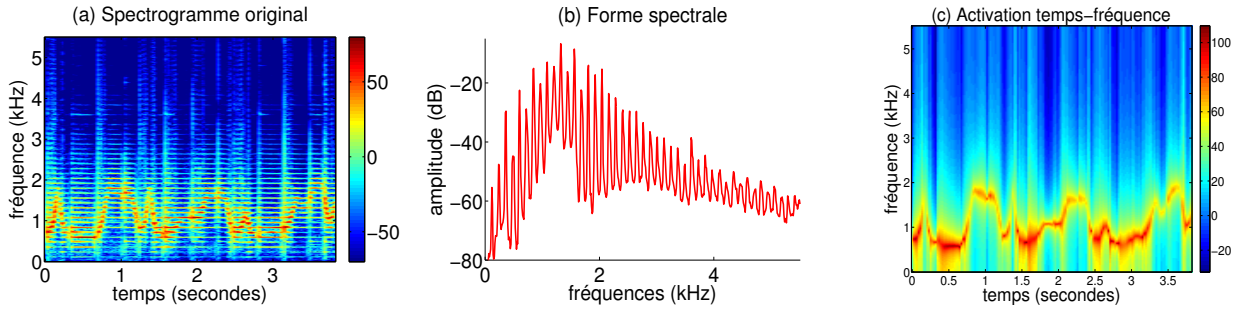

FIG. VIII.2 – Son de guimbarde décomposé avec une activation temps-fréquence paramétrée par un ARMA d'ordre (1,1) (figure extraite de [Hennequin, 2010])

Pour estimer ce modèle, nous avons proposé un algorithme multiplicatif dans [Hennequin *et al.*, 2010b, Hennequin *et al.*, 2010a] qui utilise les règles de mise à jour des filtres ARMA présentées dans la section VII.1.2.2 page 25. Cet algorithme a permis de représenter efficacement des sons non-stationnaires présentant de fortes variations spectrales, comme les sons de guimbarde. La guimbarde est un instrument constitué d'une tige métallique vibrante. Cette tige se place dans la bouche de l'instrumentiste qui module le son avec sa bouche. Il s'agit donc d'un son harmonique (de fréquence fondamentale fixe) présentant une forte résonance variant au cours du temps (voir le spectrogramme dans la figure VIII.2(a)). La décomposition obtenue avec notre algorithme (en utilisant une composante unique) est représentée dans les figures VIII.2(b) et VIII.2(c) : elle fait bien apparaître la forme harmonique du spectre d'une part et les variations temporelles de la résonance d'autre part.

VIII.3.2.2 Variations de fréquences fondamentales

La solution précédente se prête à la représentation d'éléments sonores dont l'enveloppe spectrale varie au cours du temps, mais ne permet pas de modéliser des variations de fréquence fondamentale, qui apparaissent dans certains effets sonores comme le vibrato. Dans la référence [Hennequin *et al.*, 2010d], nous proposons une méthode complémentaire qui utilise un dictionnaire de formes spectrales harmoniques correspondant aux notes de l'échelle MIDI, dont les fréquences fondamentales peuvent varier au cours du temps, dans un intervalle d'un demi-ton (mais dont les enveloppes spectrales sont constantes). Ces formes spectrales sont modélisées comme des peignes harmoniques paramétrés par leur fréquence fondamentale. Le modèle $\hat{v}_{fn}$ introduit dans la section VI.2 page 23 s'exprime alors sous la forme $\hat{v}_{fn} = \sum_{k=1}^{K} w_{fk}(\nu^0_{kn})\, h_{kn}$, où ν^0_{kn} est la fréquence fondamentale instantanée de la composante k à l'instant n. Le modèle paramétrique choisi pour représenter la forme spectrale w_{fk} s'écrit sous la forme $w_{fk}(\nu^0_{kn}) = \sum_{h=1}^{n_h(\nu^0_{kn})} a_h\, g(\nu_f - h\,\nu^0_{kn})$, où la fonction $g(\nu)$ est le module au carré de la transformée de Fourier de la fenêtre utilisée et $n_h(\nu^0_{kn})$ désigne le nombre d'harmoniques admissibles dans l'intervalle $[0, \frac{1}{2}]$ pour la fréquence fondamentale ν^0_{kn}. L'ensemble de paramètres $\boldsymbol{\theta}$ est alors constitué des fréquences fondamentales ν^0_{kn} et des amplitudes a_h (communes à toutes les formes spectrales du dictionnaire). Ces paramètres sont estimés récursivement, à l'aide d'un algorithme multiplicatif présenté dans [Hennequin *et al.*, 2010d], qui alterne leur mise à jour avec la mise à jour des activations temporelles h_{kn}.

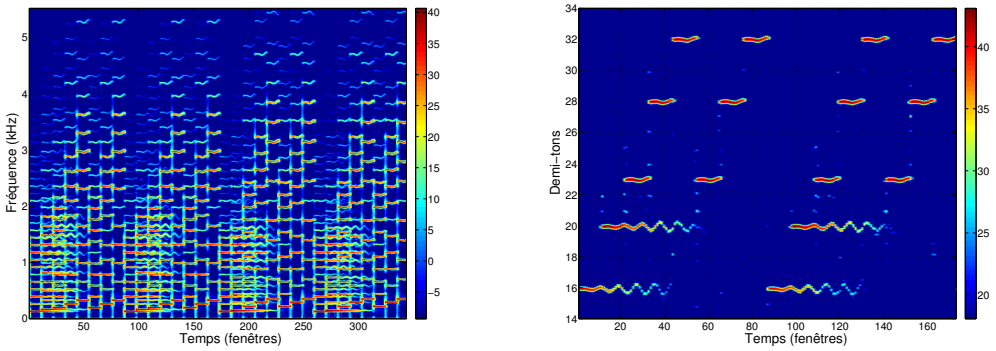

(a) Spectrogramme original

(b) Représentation des activations incluant les fréquences fondamentales (deux premières mesures de l'extrait). L'échelle de couleur est en dB.

FIG. VIII.3 – Décomposition d'un extrait du premier Prélude de Jean-Sébastien Bach (figure extraite de [Hennequin *et al.*, 2010d])

Nous avons utilisé ce modèle pour décomposer le spectrogramme d'un extrait (4 premières mesures) du premier prélude de Jean-Sébastien Bach joué par un synthétiseur (figure VIII.3(a)). Un léger vibrato a été ajouté dans les notes jouées afin de mettre en valeur l'estimation de fréquence fondamentale variable. La décomposition utilise 72 formes spectrales réparties tous les demi-tons. La figure VIII.3(b) représente les activations et les fréquences fondamentales ν^0_{kn} obtenues. Les notes du prélude apparaissent très clairement, avec l'effet de vibrato.

Dans la référence [Hennequin *et al.*, 2011], nous avons également utilisé ce modèle dans le cadre d'une application de séparation de sources informée par la partition musicale.

Chapitre IX

Applications des décompositions non-négatives

IX.1 Estimation de hauteurs et transcription automatique

La plupart des méthodes présentées dans les chapitres VII et VIII ont été conçues pour des applications d'estimation de hauteurs multiples et de transcription automatique de musique. L'estimation de hauteurs multiples consiste à déterminer la hauteur et parfois le nombre des notes jouées à un instant donné. La transcription est une tâche plus complexe, qui inclut la précédente et demande également de détecter les débuts et fins de notes, de façon à fournir une représentation symbolique de type MIDI. L'évaluation quantitative d'une transcription se fait sur la base d'une comparaison entre le fichier MIDI transcrit et un fichier MIDI de référence.

Base de données MAPS : afin d'évaluer les algorithmes, nous avons mis au point une base de données de sons de pianos annotés au format MIDI, que nous avons baptisée MAPS. Cette base se compose de notes isolées, d'accords aléatoires, d'accords usuels et de morceaux du répertoire de piano, enregistrés sur un piano droit DisKlavier de marque Yamaha ou synthétisés avec le logiciel Cubase SX de la société Steinberg, dans des conditions de jeu variées [Emiya *et al.*, 2010b]. Elle est diffusée sous licence *Creative Commons*.

Critères d'évaluation : la performance des algorithmes de transcription (ou d'estimation de hauteurs multiples) est évaluée quantitativement selon les critères généralement utilisés dans la communauté du MIR [Van Rijsbergen, 1979]. La **précision** (*precision*) est la proportion de notes correctement estimées parmi toutes les notes détectées. Le **rappel** (*recall*) est la proportion de notes de la référence MIDI qui sont correctement estimées. La **F-mesure** (*F-measure*) agrège les deux critères précédents en un critère unique (moyenne harmonique entre la précision et le rappel). En estimation de hauteurs multiples, une note est considérée correctement estimée si sa fréquence fondamentale se trouve dans un intervalle d'un demi-ton centré autour de la référence MIDI. En transcription automatique, une note est considérée correctement estimée si de plus son instant d'attaque est dans un certain voisinage de celui de la référence (typiquement 50 millisecondes). Notons que d'autres critères d'évaluation ont été proposés dans [Daniel *et al.*, 2008], incluant des poids destinés à mieux refléter l'importance perceptive des erreurs de transcription. Plusieurs algorithmes ont été évalués de cette manière, dont celui présenté dans la section IX.1.2.2 et une version antérieure de celui présenté dans la section IX.1.2.1.

Compétition MIREX : la compétition internationale *Music Information Retrieval Evaluation eXchange* (MIREX) est une campagne annuelle d'évaluation de techniques et d'algorithmes développés pour diverses tâches du domaine MIR (`http://www.music-ir.org/mirex/`). Ces algorithmes sont exécutés sur des bases de données standardisées et évalués selon des critères définis par la communauté. Les résultats de l'évaluation sont publiés chaque année et présentés lors de la conférence ISMIR (*International society for music information retrieval conference*). Certains de nos algorithmes d'estimation de hauteurs multiples et de transcription automatique ont ainsi été soumis et ont obtenu des résultats honorables lors des compétitions MIREX en 2007 et 2008.

IX.1.1 Estimation de hauteurs multiples

IX.1.1.1 Estimation d'un mélange de spectres harmoniques

Dans la référence [Emiya *et al.*, 2010a], nous avons évalué les performances de l'algorithme présenté dans la section VII.2.1 page 26 sur un sous-ensemble de la base MAPS et nous les avons comparées à deux algorithmes de l'état de l'art [Tolonen et Karjalainen, 2000, Klapuri, 2006]. Il est apparu que la méthode proposée donne des

résultats satisfaisants lorsque la polyphonie est inconnue et atteint des scores remarquablement élevés dans le cas de mélanges de notes présentant un fort recouvrement spectral. Les meilleurs résultats ont été obtenus pour les polyphonies 1 et 2, avec 93% de F-mesure, contre 88% en polyphonie 3 et 63% en polyphonie 6. Notons que notre algorithme a également été évalué dans la référence [David *et al.*, 2007] et comparé à d'autres méthodes, dont celle basée sur une analyse spectrale à haute résolution présentée dans la section V.2 page 18, ainsi que la méthode de transcription basée sur la NMF présentée dans la section IX.1.2.2. Cet algorithme a finalement été intégré dans un système de transcription automatique présenté dans la section IX.1.2.1.

IX.1.1.2 NMF avec contraintes d'harmonicité et de régularité spectrale

Dans [Vincent *et al.*, 2010], nous avons évalué l'algorithme présenté dans la section VIII.3.1.1 page 33 sur un sous-ensemble de la base MAPS et sur la base de données d'enregistrements d'instruments à vent proposée pour la compétition MIREX 2007. Quatre autres algorithmes d'estimation de hauteurs ont été testés à titre comparatif [Pertusa et Iñesta, 2008, Tolonen et Karjalainen, 2000, Klapuri, 2006, Emiya *et al.*, 2007a]. Sur les données de piano, notre algorithme a surpassé les autres approches en atteignant une F-mesure moyenne de 67%. Sur les données d'instruments à vent, notre algorithme a atteint des F-mesures allant de 63% à 68% en polyphonies 3 à 5, supérieures à celles de tous les autres algorithmes testés. Dans la compétition MIREX 2007, notre algorithme s'est hissé à la 5ème place (parmi 16 concurrents) en estimation de fréquences fondamentales multiples et à la 2ème place (parmi 11 concurrents) en transcription. Dans la compétition MIREX 2008, il a obtenu la 4ème place (parmi 13 concurrents) en transcription.

IX.1.2 Transcription automatique de musique

IX.1.2.1 Estimation de hauteurs et modèles de Markov cachés

Le système de transcription de piano que nous avons proposé dans [Emiya *et al.*, 2008, Emiya *et al.*, 2010a] est basé sur la méthode d'estimation de hauteurs multiples présentée dans la section IX.1.1.1. Il comprend une première phase de détection d'attaques, puis une phase de sélection d'un ensemble de hauteurs candidates après chaque attaque. Entre deux attaques consécutives, la poursuite des notes est effectuée en découpant le signal en trames et en estimant un modèle de Markov caché dont les états sont les combinaisons possibles de hauteurs candidates, obtenu en maximisant le critère MAP présenté dans la section VII.2.1 page 26. Enfin, un post-traitement corrige la transcription en détectant les répétitions ou les continuités de notes au niveau des attaques. Le système ainsi obtenu est capable d'estimer conjointement le niveau de polyphonie, dans la limite de 6 notes simultanées, et d'identifier les notes présentes, en estimant leur hauteur, leur instant d'attaque et leur durée. Des tests ont été effectués sur un ensemble de 90 morceaux de 30 secondes extraits de la base MAPS et des comparaisons avec des méthodes concurrentes du niveau de l'état de l'art [Marolt, 2004, Bertin *et al.*, 2007, Vincent *et al.*, 2008] ont montré des résultats prometteurs. On a ainsi obtenu une F-mesure de 65% et une bonne estimation des durées de notes, ce qui contribue à la qualité subjective de la transcription. Ce système de transcription a également été présenté à la compétition MIREX en 2008, où il s'est hissé au 3ème rang pour la F-mesure. Il a enfin été évalué selon des critères perceptifs dans la référence [Daniel *et al.*, 2008].

IX.1.2.2 NMF avec post-traitements

Le premier algorithme de transcription automatique par NMF que nous avons proposé [Bertin *et al.*, 2007] ne faisait pas encore usage des raffinements introduits dans la section VIII.3 page 32 (ajout de contraintes d'harmonicité et de régularités spectrale et temporelle). Il a donc fallu mettre en œuvre des post-traitements adéquats pour interpréter la factorisation obtenue, c'est à dire d'une part estimer les fréquences fondamentales des formes spectrales et d'autre part détecter les débuts et fins de notes sur les activations temporelles. Les fréquences fondamentales ont été estimées en recherchant le maximum d'un produit spectral. Les activations correspondant à des formes spectrales de mêmes hauteurs ont ensuite été sommées, puis seuillées pour déterminer les instants d'attaques et d'extinctions de notes. Enfin, les formes spectrales non-harmoniques et les notes de durées trop courtes ont été ignorées. Cet algorithme a été évalué sur une base de données de six morceaux de musique extraits du répertoire de piano classique, décrits dans [Bello *et al.*, 2006]. Ces morceaux ont été enregistrés sur un piano droit Yamaha Disklavier. Chaque morceau a été ensuite resynthétisé à partir du MIDI de façon à pouvoir tester les performances du système sur des données réelles et synthétiques. Les performances

ont été comparées à celles obtenues par un algorithme semblable, mais où la NMF a été remplacée par une *Non-negative K-SVD* (NKSVD) [Aharon *et al.*, 2005]. Les F-mesures moyennes obtenues par NMF ont été de 53% sur les données synthétiques et de 50% sur les données réelles. Celles obtenues par NKSVD ont été de 42% sur les données synthétiques et de 44% sur les données réelles. Outre ses performances légèrement supérieures, la NMF s'est également révélée plus robuste à l'initialisation et au rang de troncature choisis, et de complexité beaucoup plus faible que la NKSVD, c'est pourquoi nos travaux ultérieurs se sont concentrés sur la NMF uniquement. Cet algorithme a enfin été évalué selon des critères perceptifs dans la référence [Daniel *et al.*, 2008].

IX.1.2.3 NMF avec contraintes d'harmonicité et de régularités spectrale et temporelle

Dans [Bertin *et al.*, 2010, Bertin *et al.*, 2009b], nous avons appliqué les algorithmes de NMF avec contraintes d'harmonicité et de régularités spectrale et temporelle présentés dans la section VIII.3.1.2 page 33 à la transcription automatique de musique de piano. Ces algorithmes ont été testés sur des signaux synthétiques et sur des signaux enregistrés (30 morceaux de 30 secondes dans chaque cas) extraits de la base MAPS, et comparés à d'autres systèmes de transcription de l'état de l'art, dont certains sont également basés sur la NMF [Févotte *et al.*, 2009, Virtanen, 2007, Vincent *et al.*, 2008]. Nous avons ainsi obtenu une F-mesure moyenne de 67% sur les données synthétiques et de 45% sur les données réelles. Ces résultats se sont avérés supérieurs à ceux obtenus avec tous les autres algorithmes de NMF, mais inférieurs à ceux obtenus avec l'algorithme présenté dans la section IX.1.2.1 (respectivement 68% et 53%) et avec l'algorithme SONIC basé sur des réseaux de neurones [Marolt, 2004] (respectivement 76% et 58%). La contrainte de régularité temporelle n'apporte pas toutes les améliorations que nous aurions pu espérer par rapport à la robustesse au seuil de détection et l'estimation de la durée des notes, particulièrement dans le cas de sons transitoires. Néanmoins, elle est utile pour compenser la tendance de la NMF avec contrainte d'harmonicité à produire des activations temporelles irrégulières et conduit ainsi à une meilleure performance de transcription quand les deux contraintes sont utilisées conjointement.

IX.2 Séparation de sources informée

Par opposition au problème de séparation aveugle, la séparation informée - *Informed Source Separation* (ISS) [Parvaix *et al.*, 2010] considère l'apport d'une information annexe permettant d'améliorer la qualité de la séparation, comme le nombre et les caractéristiques timbrales des sources, leur distribution spatiale, leur fréquence fondamentale ou une représentation musicale symbolique de type MIDI. En utilisant la terminologie du codage audio, un système de séparation informée se compose d'un **codeur** et d'un **décodeur**. Au niveau du codeur, on dispose à la fois des sources et des mélanges et on utilise toute cette connaissance pour créer une information annexe, dont la taille est beaucoup plus faible que celle des sources. Au niveau du décodeur, les sources sont inconnues, mais on dispose des mélanges et de l'information transmise par le codeur pour estimer les sources. L'ISS vise ainsi à rendre la séparation de sources robuste et permet des applications comme l'écoute active de la musique, dont le but est de donner à l'auditeur la possibilité d'agir en temps-réel sur les sources séparées. Si l'information annexe est suffisamment petite, elle peut être directement embarquée dans les signaux de mélange par tatouage, ce qui permet l'écoute active sur le support CD audio stéréophonique [Parvaix *et al.*, 2010]. Dans ce contexte, nous avons proposé une nouvelle méthode de séparation informée dans le cas sous-déterminé [Liutkus *et al.*, 2010b], qui s'inscrit dans le cadre statistique introduit dans le chapitre VI. Notre méthode est basée sur une modélisation conjointe des sources du côté de l'encodeur, utilisant des variables latentes additives, supposées gaussiennes et indépendantes. Les opérations de mélange et de séparation sont modélisées comme des filtrages linéaires. L'information annexe est composée des coefficients spectraux et des activations de chaque composante latente, ainsi que des paramètres de mélange. L'espace nécessaire pour stocker cette information est principalement contrôlé par le nombre de variables latentes utilisées. En pratique, nous avons pu consacrer au codage des paramètres un poids total très faible comparé à celui des mélanges, typiquement inférieur à 5%. Même si cette information n'est ni quantifiée ni compressée, cela permet déjà de la masquer directement dans les signaux de mélange par tatouage. Lors de l'étape de décodage, la séparation est effectuée en utilisant un filtrage de Wiener généralisé des mélanges. La qualité de la séparation est directement liée au nombre de composantes latentes utilisées pour modéliser les sources et peut être connue de façon fiable par l'encodeur. La méthode proposée a ainsi permis d'atteindre d'excellentes performances et s'est avérée capable de séparer avec succès jusqu'à 11 sources dans des mélanges stéréophoniques.

Chapitre X

Conclusions et perspectives

X.1 Résumé des travaux effectués

Les travaux de recherche présentés dans ce mémoire ont porté sur deux techniques de traitement du signal : l'**analyse spectrale à haute résolution** et les **décompositions non-négatives**, dont le dénominateur commun est le principe de réduction de rang d'une représentation matricielle. La première partie a été consacrée à l'application des méthodes HR au traitement des signaux audio et musicaux. Nous avons d'abord généralisé le modèle ESM et la méthode ESPRIT en tenant compte de la possible multiplicité des pôles, puis proposé une nouvelle technique de sélection de l'ordre du modèle et une implémentation rapide de la méthode d'estimation. Par ailleurs, la non-stationnarité des signaux audio a nécessité de développer des algorithmes adaptatifs rapides effectuant la poursuite des variations temporelles de l'espace signal, des pôles et des amplitudes complexes. En outre, le nombre élevé de composantes fréquentielles et la forte dynamique spectrale ont nécessité de mettre en œuvre des prétraitements adéquats. Ces outils ont d'abord été appliqués à l'analyse/synthèse et la représentation temps-fréquence à haute résolution de signaux audio, ainsi qu'à la séparation des composantes tonales et bruitées et la synthèse avec effets sonores. Nous avons également abordé des applications d'estimation de hauteurs de notes de piano, d'estimation de tempo musical, d'analyse des modes couplés d'une harpe de concert, de codage audio, d'estimation de canal en communications numériques et de factorisation de tenseurs structurés. La deuxième partie du mémoire a porté sur des techniques de décompositions non-négatives de signaux audio à partir d'une représentation spectrale ou d'une représentation temps-fréquence, dans le but d'en extraire des éléments sonores possédant un sens musical, comme des notes de musique ou des accords. En utilisant une représentation spectrale, nous avons d'abord modélisé l'harmonicité et l'enveloppe spectrale caractéristiques de la hauteur et du timbre d'une note de musique, puis proposé des méthodes d'estimation de mélanges de notes. L'utilisation d'une représentation temps-fréquence permet en outre d'exploiter la dynamique temporelle des signaux audio, grâce à des techniques de factorisations matricielles comme la NMF, que nous avons formulée comme un problème d'estimation statistique. Nous avons ainsi proposé des algorithmes multiplicatifs de minimisation du critère de NMF dont nous avons également étudié et amélioré les propriétés de convergence, et nous avons contraint la décomposition afin de prendre en compte l'harmonicité, la régularité spectrale et les variations temporelles de chaque élément sonore, en introduisant des modèles paramétriques appropriés ou des distributions *a priori* des paramètres dans un cadre d'inférence bayésienne. Ces outils ont été appliqués à l'estimation de fréquences fondamentales multiples et à la transcription automatique de la musique, ainsi qu'à la séparation informée des différentes sources sonores mixées dans un enregistrement stéréophonique.

X.2 Projet de recherche à quatre ans

Les perspectives de ces travaux de recherche sont multiples. Une première question que l'on peut se poser est l'association des décompositions non-négatives et de l'analyse à haute résolution. Une voie à explorer serait l'estimation d'un mélange de composantes fréquentielles déterministes et de composantes aléatoires gaussiennes à l'aide d'un algorithme itératif, alternant des étapes d'analyse à haute résolution et de décompositions non-négatives. Une autre voie consisterait à appliquer les techniques de décompositions non-négatives à une représentation temps-fréquence à haute résolution du signal. Les décompositions non-négatives nécessitent habituellement un quadrillage (éventuellement irrégulier) du plan temps-fréquence, *a priori* incompatible avec la représentation obtenue par analyse HR (constituée d'un nuage de points arbitrairement distribués dans le plan), mais de nouvelles approches ont récemment été proposées pour s'affranchir de cette contrainte [Schmidt, 2009].

Techniques d'analyse à haute résolution

Nos travaux sur la poursuite de sous-espace sont arrivées à maturité avec les algorithmes FAPI et YAST. En revanche, les propriétés de convergence de ces deux algorithmes sont encore mal connues. Si la stabilité numérique de YAST vis-à-vis de l'orthonormalité a été démontrée, celle de FAPI a seulement été observée. Il serait aussi intéressant d'analyser et de comparer leurs vitesses de convergence d'un point de vue théorique. Par ailleurs, nos efforts se sont concentrés sur la poursuite des composantes fréquentielles, en supposant l'enveloppe spectrale du bruit et l'ordre du modèle constants et connus. Nous pourrions aussi envisager une poursuite conjointe des composantes fréquentielles, des paramètres de bruit et de l'ordre du modèle.

Applications de l'analyse à haute résolution

Des travaux pionniers menés à l'Université d'Aalborg ont conduit au développement de nouvelles méthodes d'estimation de hauteurs et de leur multiplicité [Christensen *et al.*, 2006]. Le potentiel de l'analyse HR dans ce domaine semble donc prometteur et nous envisageons de développer des algorithmes rapides de poursuite des fréquences fondamentales. Dans le domaine de l'analyse/synthèse, nous poursuivrons le développement de notre codeur audio à haute résolution basé sur le modèle ESM, auquel nous apporterons l'utilisation de critères psycho-acoustiques et la quantification conjointe des paramètres de fréquences avec contrainte entropique.

Techniques de décompositions non-négatives

Les récents travaux que nous avons menés pour approfondir le lien entre les processus gaussiens et la NMF ont offert de nouvelles perspectives pour modéliser des signaux non-stationnaires. Certaines questions concernant le calcul de la NMF restent ouvertes, comme l'estimation de l'ordre de la factorisation. Par ailleurs, les diverses contraintes que nous avons appliquées à la décomposition pourront être fusionnées dans un modèle unique, incluant des formes spectrales régulières et harmoniques et des activations temps-fréquences. La régularité des variations temporelles des enveloppes spectrales et des fréquences fondamentales pourra être induite par une distribution *a priori* structurée par une chaîne de Markov. Un tel modèle conviendrait pour représenter une grande variété de sons non-stationnaires, harmoniques ou bruités. Les sons transitoires seront quant à eux modélisés par une convolution temporelle entre une forme temps-fréquence de courte durée et une activation temporelle parcimonieuse. Pour identifier un son comportant simultanément des parties harmonique, bruitée et transitoire, les activations correspondantes pourront être alignées en introduisant des contraintes de corrélation. Nous projetons par ailleurs de généraliser ces décompositions non-négatives à l'analyse de signaux multi-canaux.

Applications des décompositions non-négatives

Notre algorithme EM d'estimation de hauteurs multiples s'est avéré très prometteur, mais n'a pas encore atteint le degré d'aboutissement de la méthode de maximum *a priori* pondéré et nécessitera certaines améliorations, comme la prise en compte de l'inharmonicité du piano et de la forme spectrale de la fenêtre temporelle utilisée, l'ajout d'informations *a priori* sur la distribution des paramètres et l'estimation automatique de la polyphonie. Par ailleurs, l'analyse en composantes latentes indépendantes s'est également avérée être un outil très puissant pour l'estimation et la poursuite de hauteurs multiples [Smaragdis, 2009], car elle permet de modéliser aisément des transpositions fréquentielles. Elle souffre en revanche d'une certaine rigidité, puisque chaque instrument est caractérisé par une forme spectrale unique, dont seules la puissance et la fréquence fondamentale varient au cours du temps. Nous envisageons donc de généraliser cette décomposition afin de prendre en compte de possibles variations d'enveloppe spectrale au cours du temps et sur l'échelle des hauteurs. Enfin, nous poursuivrons nos travaux dans le domaine de la séparation de sources informée, dans le cas de mélanges sous-déterminés (mono ou stéréophoniques), instantanés ou convolutifs. Les sources seront représentées par des modèles gaussiens. Nous considérerons tout d'abord le cas où l'information annexe est constituée d'une représentation musicale symbolique de type MIDI. L'approche que nous utiliserons reposera sur le modèle complet de signal incluant des parties harmoniques, bruitées et transitoires, dont l'estimation sera guidée par l'information MIDI. Concernant l'application de séparation informée pour l'écoute active de CD audio, où l'information annexe est directement embarquée dans les signaux de mélange par tatouage, la première approche que nous avons proposée comportait une modélisation des sources à l'aide d'une NMF multi-canaux. Les développements à venir utiliseront plus généralement une représentation compacte, éventuellement partielle, du contenu temps-fréquence de chaque source. Nous envisageons également la possibilité d'adapter nos méthodes de séparation informée au problème de séparation aveugle, en utilisant des données d'apprentissage et une étape d'estimation des modèles de sources. Nous utiliserons finalement nos techniques de décomposition et de séparation de sources pour appliquer diverses transformations à des signaux mono ou multicanaux, comme la modification de fréquence fondamentale, de durée, de timbre, de localisation spatiale, ou le remixage.

Bibliographie

[Abed-Meraim *et al.*, 2000] Karim Abed-Meraim, Ammar Chkeif, et Yingbo Hua. Fast orthonormal PAST algorithm. *IEEE Signal Proc. Letters*, 7(3):60–62, mars 2000.

[Aharon *et al.*, 2005] Michal Aharon, Michael Elad, et Alfred M. Bruckstein. K-SVD and its nonnegative variant for dictionary design. Dans *Proc. of the SPIE conference wavelets*, volume 5914, pages 327–339, San Diego, Californie, USA, juillet 2005.

[Allen et Rabiner, 1977] J. B. Allen et L. R. Rabiner. A unified approach to short-time Fourier analysis and synthesis. *Proceedings of the IEEE*, 65(11):1558–1564, novembre 1977.

[Alonso Arevalo, 2006] Miguel Alonso Arevalo. *Extraction d'information rythmique à partir d'enregistrements musicaux*. PhD thesis, École Nationale Supérieure des Télécommunications, Paris, France, novembre 2006.

[Alonso Arevalo *et al.*, 2003a] Miguel Alonso Arevalo, Roland Badeau, Bertrand David, et Gaël Richard. Musical tempo estimation using noise subspace projections. Dans *Proc. of IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pages 95–98, New Paltz, New York, USA, octobre 2003.

[Alonso Arevalo *et al.*, 2003b] Miguel Alonso Arevalo, Bertrand David, et Gaël Richard. A study of tempo tracking algorithms from polyphonic music signals. Dans *Proc. of 4th COST 276 Workshop*, pages 1–5, Bordeaux, France, mars 2003.

[Alonso Arevalo *et al.*, 2007] Miguel Alonso Arevalo, Gaël Richard, et Bertrand David. Accurate tempo estimation based on harmonic + noise decomposition. *EURASIP Journal on Advances in Signal Processing*, 2007(1):1–14, janvier 2007.

[Altuna *et al.*, 2006] Jon Altuna, Bernie Mulgrew, Roland Badeau, et Vicente Atxa. A fast adaptive method for subspace based blind channel estimation. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 4, pages 1121–1124, Toulouse, France, mai 2006.

[Auger et Flandrin, 1995] François Auger et Patrick Flandrin. Improving the readability of time-frequency and time-scale representations by the reassignment method. *IEEE Transactions on Signal Processing*, 43(5):1068–1089, mai 1995.

[Badeau, 2005] Roland Badeau. *Méthodes à haute résolution pour l'estimation et le suivi de sinusoïdes modulées. Application aux signaux de musique.* PhD thesis, École Nationale Supérieure des Télécommunications, ENST2005E007, Paris, France, avril 2005. Prix de thèse ParisTech 2006.

[Badeau *et al.*, 2003a] Roland Badeau, Karim Abed-Meraim, Gaël Richard, et Bertrand David. Sliding window orthonormal PAST algorithm. Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 5, pages 261–264, Hong Kong, Chine, avril 2003.

[Badeau *et al.*, 2010a] Roland Badeau, Nancy Bertin, et Emmanuel Vincent. Stability analysis of multiplicative update algorithms and application to non-negative matrix factorization. *IEEE Transactions on Neural Networks*, 21(12):1869–1881, décembre 2010.

[Badeau *et al.*, 2010b] Roland BADEAU, Nancy BERTIN, et Emmanuel VINCENT. Stability analysis of multiplicative update algorithms and application to non-negative matrix factorization. Rapport Technique 2010D018, Télécom ParisTech, Paris, France, septembre 2010.

[Badeau *et al.*, 2010c] Roland BADEAU, Nancy BERTIN, et Emmanuel VINCENT. Supporting document for the paper "stability analysis of multiplicative update algorithms and application to non-negative matrix factorization". Rapport Technique 2010D019, Télécom ParisTech, Paris, France, septembre 2010.

[Badeau *et al.*, 2011] Roland BADEAU, Nancy BERTIN, et Emmanuel VINCENT. Stability analysis of multiplicative update algorithms for non-negative matrix factorization. Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Prague, République Tchèque, mai 2011. Soumis pour publication.

[Badeau et Boyer, 2008] Roland BADEAU et Rémy BOYER. Fast multilinear singular value decomposition for structured tensors. *SIAM Journal on Matrix Analysis and Applications*, 30(3):1008–1021, septembre 2008.

[Badeau *et al.*, 2002] Roland BADEAU, Rémy BOYER, et Bertrand DAVID. EDS parametric modeling and tracking of audio signals. Dans *Proc. of the 5th International Conference on Digital Audio Effects (DAFx)*, pages 139–144, Hambourg, Allemagne, septembre 2002.

[Badeau et David, 2008a] Roland BADEAU et Bertrand DAVID. Adaptive subspace methods for high resolution analysis of music signals. Dans *Acoustics'08*, Paris, France, juillet 2008.

[Badeau et David, 2008b] Roland BADEAU et Bertrand DAVID. Weighted maximum likelihood autoregressive and moving average spectrum modeling. Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 3761–3764, Las Vegas, Nevada, USA, avril 2008.

[Badeau *et al.*, 2004a] Roland BADEAU, Bertrand DAVID, et Gaël RICHARD. Selecting the modeling order for the ESPRIT high resolution method: an alternative approach. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 2, pages 1025–1028, Montréal, Québec, Canada, mai 2004.

[Badeau *et al.*, 2005a] Roland BADEAU, Bertrand DAVID, et Gaël RICHARD. Fast approximated power iteration subspace tracking. *IEEE Transactions on Signal Processing*, 53(8):2931–2941, août 2005.

[Badeau *et al.*, 2005b] Roland BADEAU, Bertrand DAVID, et Gaël RICHARD. Yet another subspace tracker. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 4, pages 329–332, Philadelphie, Pennsylvanie, USA, mars 2005.

[Badeau *et al.*, 2006a] Roland BADEAU, Bertrand DAVID, et Gaël RICHARD. High resolution spectral analysis of mixtures of complex exponentials modulated by polynomials. *IEEE Transactions on Signal Processing*, 54(4):1341–1350, avril 2006.

[Badeau *et al.*, 2006b] Roland BADEAU, Bertrand DAVID, et Gaël RICHARD. A new perturbation analysis for signal enumeration in rotational invariance techniques. *IEEE Transactions on Signal Processing*, 54(2):450–458, février 2006.

[Badeau *et al.*, 2006c] Roland BADEAU, Bertrand DAVID, et Gaël RICHARD. YAST algorithm for minor subspace tracking. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 3, pages 552–555, Toulouse, France, mai 2006.

[Badeau *et al.*, 2007a] Roland BADEAU, Bertrand DAVID, et Gaël RICHARD. Conjugate gradient algorithms for minor subspace analysis. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 3, pages 1013–1016, Honolulu, Hawaii, USA, avril 2007.

[Badeau *et al.*, 2008a] Roland BADEAU, Bertrand DAVID, et Gaël RICHARD. Cramér-Rao bounds for multiple poles and coefficients of quasipolynomials in colored noise. *IEEE Transactions on Signal Processing*, 56(8):3458–3467, août 2008.

[Badeau *et al.*, 2009] Roland Badeau, Valentin Emiya, et Bertrand David. Expectation-maximization algorithm for multi-pitch estimation and separation of overlapping harmonic spectra. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3073–3076, Taipei, Taiwan, avril 2009.

[Badeau *et al.*, 2003b] Roland Badeau, Gaël Richard, et Bertrand David. Adaptive ESPRIT algorithm based on the PAST subspace tracker. Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 6, pages 229–232, Hong Kong, Chine, avril 2003.

[Badeau *et al.*, 2003c] Roland Badeau, Gaël Richard, et Bertrand David. Suivi d'espace dominant par la méthode des puissances itérées. Dans *Actes du colloque GRETSI*, volume 1, pages 137–140, Paris, France, septembre 2003.

[Badeau *et al.*, 2004b] Roland Badeau, Gaël Richard, et Bertrand David. Sliding window adaptive SVD algorithms. *IEEE Transactions on Signal Processing*, 52(1):1–10, janvier 2004.

[Badeau *et al.*, 2005c] Roland Badeau, Gaël Richard, et Bertrand David. Fast adaptive ESPRIT algorithm. Dans *Proc. of IEEE Workshop on Statistical Signal Processing (SSP)*, Bordeaux, France, juillet 2005.

[Badeau *et al.*, 2005d] Roland Badeau, Gaël Richard, et Bertrand David. Procédé de poursuite d'un sous-espace de dimension inférieure à celle des vecteurs de données, notamment audio. Brevet no. 05 50678, mars 2005. Télécom ParisTech, Paris, France.

[Badeau *et al.*, 2007b] Roland Badeau, Gaël Richard, et Bertrand David. Performance of ESPRIT for estimating mixtures of complex exponentials modulated by polynomials: Supporting document. Rapport Technique 2007D015, École Nationale Supérieure des Télécommunications, Paris, France, octobre 2007.

[Badeau *et al.*, 2008b] Roland Badeau, Gaël Richard, et Bertrand David. Fast and stable YAST algorithm for principal and minor subspace tracking. *IEEE Transactions on Signal Processing*, 56(8):3437–3446, août 2008.

[Badeau *et al.*, 2008c] Roland Badeau, Gaël Richard, et Bertrand David. Performance of ESPRIT for estimating mixtures of complex exponentials modulated by polynomials. *IEEE Transactions on Signal Processing*, 56(2):492–504, février 2008.

[Badeau *et al.*, 2003d] Roland Badeau, Gaël Richard, Bertrand David, et Karim Abed-Meraim. Approximated power iterations for fast subspace tracking. Dans *Proc. of the 7th International Symposium on Signal Processing and its Applications (ISSPA)*, volume 2, pages 583–586, Paris, France, juillet 2003.

[Bartelmaos et Abed-Meraim, 2007] Steve Bartelmaos et Karim Abed-Meraim. An efficient and stable algorithm for minor subspace tracking and stability analysis. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 3, pages 1301–1304, Honolulu, Hawaii, USA, avril 2007.

[Bartelmaos et Abed-Meraim, 2008] Steve Bartelmaos et Karim Abed-Meraim. Fast principal component extraction using givens rotations. *IEEE Signal Processing Letters*, 15:369–372, mars 2008.

[Beau et Marcos, 2010] Sophie Beau et Sylvie Marcos. Range dependent clutter rejection using range-recursive space-time adaptive processing (STAP) algorithms. *Signal Processing*, 90(1):57–68, janvier 2010.

[Belkacemi et Marcos, 2006] Hocine Belkacemi et Sylvie Marcos. Fast iterative subspace algorithms for airborne STAP radar. *EURASIP Journal on Applied Signal Processing*, 2006:1–8, 2006. Article ID 37296.

[Bello *et al.*, 2006] Juan Pablo Bello, Laurent Daudet, et Mark B. Sandler. Automatic piano transcription using frequency and time-domain information. *IEEE Transactions on Audio, Speech, and Language Processing*, 14(6):2242–2251, novembre 2006.

[Berry *et al.*, 2007] Michael W. Berry, Murray Browne, Amy N. Langville, V. Paul Pauca, et Robert J. Plemmons. Algorithms and applications for approximate nonnegative matrix factorization. *Computational Statistics and Data Analysis*, 52(1):155–173, septembre 2007.

[Bertin, 2009] Nancy Bertin. *Les factorisations en matrices non-négatives. Approches contraintes et probabilistes, application à la transcription automatique de musique polyphonique.* PhD thesis, École Nationale Supérieure des Télécommunications, Paris, France, octobre 2009. Prix de thèse 2010 du GdR-Isis/EEA/GRETSI.

[Bertin et Badeau, 2008] Nancy Bertin et Roland Badeau. Initialization, distances and local minima in audio applications of the non-negative matrix factorization. Dans *Acoustics'08*, Paris, France, juillet 2008.

[Bertin *et al.*, 2007] Nancy Bertin, Roland Badeau, et Gaël Richard. Blind signal decompositions for automatic transcription of polyphonic music: NMF and K-SVD on the benchmark. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 1, pages 65–68, Honolulu, Hawaii, USA, avril 2007.

[Bertin *et al.*, 2009b] Nancy Bertin, Roland Badeau, et Emmanuel Vincent. Fast Bayesian NMF algorithms enforcing harmonicity and temporal continuity in polyphonic music transcription. Dans *Proc. of IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pages 29–32, New Paltz, New York, USA, octobre 2009.

[Bertin *et al.*, 2010] Nancy Bertin, Roland Badeau, et Emmanuel Vincent. Enforcing harmonicity and smoothness in Bayesian non-negative matrix factorization applied to polyphonic music transcription. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(3):538–549, mars 2010.

[Bertin *et al.*, 2009c] Nancy Bertin, Cédric Févotte, et Roland Badeau. A tempering approach for Itakura-Saito non-negative matrix factorization. With application to music transcription. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1545–1548, Taipei, Taiwan, avril 2009.

[Bertrand et Moonen, 2009] Alexander Bertrand et Marc Moonen. Blind separation of non-negative sources using multiplicative updates and subspace projection. Rapport Technique, ESAT SCD-SISTA, Louvain, Belgique, septembre 2009.

[Boyer et Abed-Meraim, 2004] Rémy Boyer et Karim Abed-Meraim. Audio modeling based on delayed sinusoids. *IEEE Transactions on Speech and Audio Processing*, 12(2):110–120, mars 2004.

[Boyer et Badeau, 2006] Rémy Boyer et Roland Badeau. Adaptive multilinear SVD for structured tensors. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 3, pages 880–883, Toulouse, France, mai 2006.

[Boyer *et al.*, 2011] Rémy Boyer, Roland Badeau, et Gérard Favier. Fast orthogonal decomposition of Volterra cubic kernels using oblique unfolding. Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Prague, République Tchèque, mai 2011. Soumis pour publication.

[Bregman, 1967] L. M. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200–217, 1967.

[Buhmann, 2003] Martin D. Buhmann. *Radial Basis Functions: Theory and Implementations.* Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, juillet 2003.

[Cardoso *et al.*, 1995] Jean-François Cardoso, Éric Moulines, Pierre Duhamel, et Sylvie Mayrargue. Subspace methods for the blind identification of multichannel FIR filters. *IEEE Transactions on Signal Processing*, 43(2):516–525, février 1995.

[Chen *et al.*, 1998] Scott Shaobing Chen, David L. Donoho, et Michael A. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, août 1998.

[Chen *et al.*, 2006] Zhe Chen, Andrzej Cichocki, et Tomasz M. Rutkowski. Constrained non-negative matrix factorization method for EEG analysis in early detection of Alzheimer's disease. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 893–896, Toulouse, France, mai 2006.

[Chou *et al.*, 1989] Philip A. Chou, Tom Lookbaugh, et Robert M. Gray. Entropy-constrained vector quantization. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(1):31–42, janvier 1989.

[Christensen et Jakobsson, 2009] Mads Græsbøll Christensen et Andreas Jakobsson. *Multi-Pitch Estimation*. Synthesis Lectures on Speech and Audio Processing. Morgan and Claypool Publishers, 2009.

[Christensen *et al.*, 2006] Mads Græsbøll Christensen, Andreas Jakobsson, et Soren Holdt Jensen. Multi-pitch estimation using harmonic MUSIC. Dans *Record of Asilomar Conference on Signals, Systems, and Computers (ACSSC)*, pages 521–525, Pacific Grove, Californie, USA, octobre 2006.

[Christensen *et al.*, 2007] Mads Græsbøll Christensen, Andreas Jakobsson, et Soren Holdt Jensen. Sinusoidal order estimation using the subspace orthogonality and shift-invariance properties. Dans *Record of Asilomar Conference on Signals, Systems, and Computers (ACSSC)*, pages 651–655, Pacific Grove, Californie, USA, novembre 2007.

[Christensen *et al.*, 2009] Mads Græsbøll Christensen, Andreas Jakobsson, et Soren Holdt Jensen. Sinusoidal order estimation using angles between subspaces. *EURASIP Journal on Advances in Signal Processing*, 2009:1–11, 2009. Article ID 948756.

[Cichocki *et al.*, 2006] Andrzej Cichocki, Rafal Zdunek, et Shun-ichi Amari. Csiszar's divergences for non-negative matrix factorization: Family of new algorithms. Dans *Proc. of the 6th International Conference on Independent Component Analysis and Blind Signal Separation (ICA)*, pages 32–39, Charleston, Caroline du Sud, USA, mars 2006.

[Cichocki *et al.*, 2008] Andrzej Cichocki, Rafal Zdunek, et Shun-ichi Amari. Nonnegative matrix and tensor factorization. *IEEE Signal Processing Magazine*, 25(1):142–145, janvier 2008.

[Cichocki *et al.*, 2009] Andrzej Cichocki, Rafal Zdunek, Anh Huy Phan, et Shun-ichi Amari. *Nonnegative Matrix and Tensor Factorizations: Applications to Exploratory Multi-way Data Analysis and Blind Source Separation*. Wiley, novembre 2009.

[Cohen, 1995] Léon Cohen. *Time-Frequency Analysis*. Prentice-Hall, New York, NY, USA, 1995.

[Crammer, 2007] Koby Crammer. A conservative aggressive subspace tracker. Dans *Proc. of the 8th Annual Conference of the International Speech Communication Association (Interspeech)*, pages 498–501, Antwerp, Belgique, août 2007.

[Da Costa *et al.*, 2009a] João Paulo C. L. Da Costa, Florian Römer, et Martin Haardt. Deterministic prewhitening to improve subspace based parameter estimation techniques in severely colored noise environments. Dans *Proc. of the 54th International Scientific Colloquium (IWK)*, Ilmenau, Allemagne, septembre 2009.

[Daniel *et al.*, 2008] Adrien Daniel, Valentin Emiya, et Bertrand David. Perceptually-based evaluation of the errors usually made when automatically transcribing music. Dans *Proc. of the 9th International Society for Music Information Retrieval Conference (ISMIR)*, Philadelphie, Pennsylvanie, USA, septembre 2008.

[Daudet et Torrésani, 2006] Laurent Daudet et Bruno Torrésani. *Signal Processing Methods for Music Transcription*, Chapter "Sparse Adaptive Representations for Musical Signals", pages 65–98. Springer US, 2006.

[David et Badeau, 2007] Bertrand David et Roland Badeau. Fast sequential LS estimation for sinusoidal modeling and decomposition of audio signals. Dans *Proc. of IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pages 211–214, New Paltz, New York, USA, octobre 2007.

[David et Badeau, 2008] Bertrand David et Roland Badeau. Towards an adaptive subspace-based representation of musical spectral content. Dans *Acoustics'08*, Paris, France, juillet 2008.

[David *et al.*, 2007] Bertrand David, Roland Badeau, Nancy Bertin, Valentin Emiya, et Gaël Richard. Multipitch detection for piano music: Benchmarking a few approaches. *The Journal of the Acoustical Society of America (JASA)*, 122(5):2962, novembre 2007.

[David *et al.*, 2002] Bertrand David, Roland Badeau, et Gaël Richard. Sintrack analysis for tracking components of musical signals. Dans *Proc. of the Forum Acusticum*, Séville, Espagne, septembre 2002.

[David *et al.*, 2006a] Bertrand David, Roland Badeau, et Gaël Richard. HRHATRAC algorithm for spectral line tracking of musical signals. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 3, pages 45–48, Toulouse, France, mai 2006.

[David *et al.*, 2006b] Bertrand David, Valentin Emiya, Roland Badeau, et Yves Grenier. Harmonic plus noise decomposition: Time-frequency reassignment versus a subspace-based method. Dans *Proc. of 120th Convention of the Audio Engineering Society (AES)*, Paris, France, mai 2006.

[David *et al.*, 2003] Bertrand David, Gaël Richard, et Roland Badeau. An EDS modelling tool for tracking and modifying musical signals. Dans *Proc. of Stockholm Music Acoustics Conference (SMAC)*, volume 2, pages 715–718, Stockholm, Suède, août 2003.

[Davila, 2000] Carlos E. Davila. Efficient, high performance, subspace tracking for time-domain data. *IEEE Transactions on Signal Processing*, 48(12):3307–3315, décembre 2000.

[De Cheveigne et Kawahara, 2002] Alain De Cheveigne et Hideki Kawahara. YIN, a fundamental frequency estimator for speech and music. *The Journal of the Acoustical Society of America*, 111(4):1917–1930, avril 2002.

[De Lathauwer, 1997] Lieven De Lathauwer. *Signal Processing based on Multilinear Algebra*. PhD thesis, Université catholique de Louvain, Louvain, Belgique, septembre 1997.

[Derrien *et al.*, 2010] Olivier Derrien, Roland Badeau, et Gaël Richard. Calculation of an entropy-constrained quantizer for exponentially damped sinudoids parameters. Rapport Technique, Laboratoire de Mécanique et d'Acoustique, Marseille, France, juin 2010.

[Derrien *et al.*, 2011] Olivier Derrien, Roland Badeau, et Gaël Richard. Entropy-constrained quantization of exponentially damped sinusoids parameters. Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Prague, République Tchèque, mai 2011. Soumis pour publication.

[Derrien *et al.*, 2008] Olivier Derrien, Gaël Richard, et Roland Badeau. Damped sinusoids and subspace based approach for lossy audio coding. Dans *Acoustics'08*, Paris, France, juillet 2008.

[Donoho, 2006] David L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, avril 2006.

[Doukopoulos et Moustakides, 2008] Xenofon G. Doukopoulos et George V. Moustakides. Fast and stable subspace tracking. *IEEE Transactions on Signal Processing*, 56(4):1452–1465, avril 2008.

[Drakakis *et al.*, 2007] Konstantinos Drakakis, Scott Rickard, Ruairí De Fréin, et Andrzej Cichocki. Analysis of financial data using non-negative matrix factorization. *International Journal of Mathematical Sciences*, 6(2), juin 2007.

[Du et Swamy, 2006] K.-L. Du et M. N. S. Swamy. *Neural Networks in a Softcomputing Framework*. Springer, 2006.

[Durrieu *et al.*, 2008] Jean-Louis Durrieu, Gaël Richard, et Bertrand David. Singer melody extraction in polyphonic signals using source separation methods. Dans *Proc. of International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 169–172, Las Vegas, Nevada, USA, March 30-April 4 2008.

[Ege *et al.*, 2009b] Kerem Ege, Xavier Boutillon, et Bertrand David. High-resolution modal analysis. *Journal of Sound and Vibration*, 325(4-5):852–869, septembre 2009.

[Eguchi et Kano, 2001] Shinto Eguchi et Yutaka Kano. Robustifying maximum likelihood estimation. Rapport Technique, Tokyo Institute of Statistical Mathematics, Tokyo, Japon, 2001.

[El-Jaroudi et Makhoul, 1991] A. El-Jaroudi et J. Makhoul. Discrete all-pole modeling. *IEEE Transactions on Signal Processing*, 39(2):411–423, février 1991.

[Elad, 2010] Michael Elad. *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer, première édition, août 2010.

[Emiya *et al.*, 2007a] Valentin Emiya, Roland Badeau, et Bertrand David. Multipitch estimation of quasi-harmonic sounds in colored noise. Dans *Proc. of the 10th International Conference on Digital Audio Effects (DAFx)*, pages 93–98, Bordeaux, France, septembre 2007.

[Emiya *et al.*, 2008] Valentin Emiya, Roland Badeau, et Bertrand David. Automatic transcription of piano music based on HMM tracking of jointly-estimated pitches. Dans *Proc. of European Signal Processing Conference (EUSIPCO)*, Lausanne, Suisse, août 2008.

[Emiya *et al.*, 2010a] Valentin Emiya, Roland Badeau, et Bertrand David. Multipitch estimation of piano sounds using a new probabilistic spectral smoothness principle. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(6):1643–1654, août 2010.

[Emiya *et al.*, 2010b] Valentin Emiya, Nancy Bertin, Bertrand David, et Roland Badeau. MAPS - A piano database for multipitch estimation and automatic transcription of music. Rapport Technique 2010D017, Télécom ParisTech, Paris, France, juillet 2010.

[Emiya *et al.*, 2007b] Valentin Emiya, Bertrand David, et Roland Badeau. A parametric method for pitch estimation of piano tones. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 1, pages 249–252, Honolulu, Hawaii, USA, avril 2007.

[Emiya *et al.*, 2005] Valentin Emiya, Bertrand David, et Vincent Gibiat. Two representation tools to analyse non-stationary sounds in a perceptive context. Dans *Proc. of the Forum Acusticum*, Budapest, Hongrie, 2005.

[Eriksson *et al.*, 1993] Anders Eriksson, Petre Stoica, et Torsten Söderström. Second-order properties of MUSIC and ESPRIT estimates of sinusoidal frequencies in high SNR scenarios. *IEE Proceedings on Radar, Sonar and Navigation*, 140(4):266–272, août 1993.

[Fadili et Bullmore, 2005] Jalal M. Fadili et Ed Bullmore. Penalized partially linear models using sparse representations with an application to fMRI time series. *IEEE Transactions on Signal Processing*, 53(9):3436–3448, septembre 2005.

[Fessler et Hero, 1994] Jeffrey A. Fessler et Alfred O. Hero. Space-alternating generalized expectation-maximization algorithm. *IEEE Transactions on Signal Processing*, 42(10):2664–2677, octobre 1994.

[Finesso et Spreij, 2004] Lorenzo Finesso et Peter Spreij. Approximate nonnegative matrix factorization via alternating minimization. Dans *Proc. of the 16th International Symposium on Mathematical Theory of Networks and Systems (MTNS)*, Louvain, Belgique, juillet 2004.

[FitzGerald *et al.*, 2008] Derry FitzGerald, Matt Cranitch, et Eugene Coyle. Extended nonnegative tensor factorisation models for musical sound source separation. *Computational Intelligence and Neuroscience*, 2008:1–15, mai 2008. Article ID 872425.

[Fletcher et Rossing, 2008] Neville H. Fletcher et Thomas D. Rossing. *The Physics of Musical Instruments*. Springer-Verlag, New York, NY, USA, seconde édition, juin 2008.

[Foutz et Spanias, 2007] Jeffrey A. Foutz et Andreas S. Spanias. An adaptive low rank algorithm for semi-spherical antenna arrays. Dans *Proc. of IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 113–116, Nouvelle Orléans, Louisiane, USA, mai 2007.

[Févotte *et al.*, 2009] Cédric Févotte, Nancy Bertin, et Jean-Louis Durrieu. Nonnegative matrix factorization with the Itakura-Saito divergence. With application to music analysis. *Neural Computation*, 21(3):793–830, mars 2009.

[Gillet et Richard, 2008] Olivier Gillet et Gaël Richard. Transcription and separation of drum signals from polyphonic music. *IEEE Transactions on Audio, Speech, and Language Processing*, 16(3):529–540, mars 2008.

[Golub et Van Loan, 1996] Gene H. Golub et Charles F. Van Loan. *Matrix computations*. The Johns Hopkins University Press, Baltimore et Londres, troisième édition, 1996.

[Goto *et al.*, 2002] Masataka Goto, Hiroki Hashiguchi, Takuichi Nishimura, et Ryuichi Oka. RWC music database: Popular, classical, and jazz music databases. Dans *Proc. of the 3rd International Society for Music Information Retrieval Conference (ISMIR)*, pages 287–288, Paris, France, octobre 2002.

[Grenier et David, 2003] Yves Grenier et Bertrand David. Extraction of weak background transients from audio signals. Dans *Audio Engineering Society Convention Paper (AES)*, Amsterdam, Pays-Bas, mars 2003.

[Gribonval et Bacry, 2003] Rémi Gribonval et Emmanuel Bacry. Harmonic decomposition of audio signals with matching pursuit. *IEEE Transactions on Signal Processing*, 51(1):101–111, janvier 2003.

[Guo *et al.*, 2006] Xian Sheng Guo, Qun Wan, et Wan Lin Yang. Fast DOA tracking of coherently distributed sources based on subspace updating. Dans *Proc. of International Conference on Radar (CIE)*, pages 1–4, Shanghai, Chine, octobre 2006.

[Haddad et Noga, 2007] Darren H. Haddad et Andrew J. Noga. The matrix pencil and its applications to speech processing. Rapport Technique, Air Force Research Lab, New York, NY, USA, mars 2007.

[Hendriks *et al.*, 2009] Richard C. Hendriks, Richard Heusdens, Jesper Jensen, et Ulrik Kjems. Low complexity DFT-domain noise PSD tracking using high resolution periodograms. *EURASIP Journal on Advances in Signal Processing*, 2009:1–16, 2009. Article ID 925870.

[Hennequin, 2010] Romain Hennequin. Rapport à mi-parcours de travaux de thèse. Télécom ParisTech, avril 2010.

[Hennequin *et al.*, 2010a] Romain Hennequin, Roland Badeau, et Bertrand David. NMF with time-frequency activations to model non-stationary audio events. *IEEE Transactions on Audio, Speech, and Language Processing*, 2010. A paraître.

[Hennequin *et al.*, 2010b] Romain Hennequin, Roland Badeau, et Bertrand David. NMF with time-frequency activations to model non-stationary audio events. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 445–448, Dallas, Texas, USA, mars 2010.

[Hennequin *et al.*, 2010c] Romain Hennequin, Roland Badeau, et Bertrand David. Spectral similarity measure invariant to pitch shifting and amplitude scaling. Dans *Actes du 10ème Congrès Français d'Acoustique (CFA)*, Lyon, France, avril 2010.

[Hennequin *et al.*, 2010d] Romain Hennequin, Roland Badeau, et Bertrand David. Time-dependent parametric and harmonic templates in non-negative matrix factorization. Dans *Proc. of the 13th International Conference on Digital Audio Effects (DAFx)*, Graz, Autriche, septembre 2010.

[Hennequin *et al.*, 2010e] Romain Hennequin, Bertrand David, et Roland Badeau. Beta-divergence as a subclass of Bregman divergence. *IEEE Signal Processing Letters*, 18(2):83–86, février 2010.

[Hennequin *et al.*, 2011] Romain Hennequin, Bertrand David, et Roland Badeau. Score informed audio source separation using a parametric model of non-negative spectrogram. Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Prague, République Tchèque, mai 2011. Soumis pour publication.

[Hermus *et al.*, 2007] Kris Hermus, Patrick Wambacq, et Hugo Van hamme. A review of signal subspace speech enhancement and its application to noise robust speech recognition. *EURASIP Journal on Advances in Signal Processing*, 2007:1–15, 2007. Article ID 45821.

[Hoegaerts *et al.*, 2007] Luc Hoegaerts, Lieven De Lathauwer, Ivan Goethals, Johan A.K. Suykens, Joos Vandewalle, et Bart De Moor. Efficiently updating and tracking the dominant kernel principal components. *Neural Networks*, 20(2):220–229, mars 2007.

[Horn et Johnson, 1985] Roger A. Horn et Charles R. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, 1985.

[Hua et Sarkar, 1990] Yingbo Hua et Tapan K. Sarkar. Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38(5):814–824, mai 1990.

[Hua *et al.*, 1999] Yingbo Hua, Yong Xiang, Tianping Chen, Karim Abed-Meraim, et Yongfeng Miao. A new look at the power method for fast subspace tracking. *Digital Signal Processing*, 9(4):297–314, octobre 1999.

[Huang et Aviyente, 2007] Ke Huang et Selin Aviyente. Sparse representation for signal classification. Dans *Advances in Neural Information Processing Systems 19*, pages 609–616. MIT Press, Cambridge, MA, USA, 2007.

[Huber et Kollmeier, 2006] Rainer Huber et Birger Kollmeier. PEMO-Q – a new method for objective audio quality assessment using a model of auditory perception. *IEEE Transactions on Audio, Speech, and Language Processing*, 14(6):1902–1911, novembre 2006.

[Ishteva *et al.*, 2009] Mariya Ishteva, Lieven De Lathauwer, Pierre-Antoine Absil, et Sabine Van Huffel. Best low multilinear rank approximation of higher-order tensors, based on the Riemannian trust-region scheme. Rapport Technique 09-142, ESAT-SISTA, Université Catholique de Louvain, Louvain, Belgique, 2009.

[Itakura et Saito, 1968] Fumitada Itakura et Shuzo Saito. Analysis synthesis telephony based on the maximum likelihood method. Dans *Proc. of the 6th International Congress on Acoustics (ICA)*, pages C17–C20, Tokyo, Japon, août 1968.

[Jakobsson *et al.*, 2007] Andreas Jakobsson, Mads Græsbøll Christensen, et Soren H. Jensen. Frequency selective sinusoidal order estimation. *IEE Electronic Letters*, 43(21):1164–1165, octobre 2007.

[Jeff Wu, 1983] Chien-Fu Jeff Wu. On the convergence properties of the EM algorithm. *Annals of Statistics*, 11(1):95–103, mars 1983.

[Jensen *et al.*, 2004] Jesper Jensen, Richard Heusdens, et Søren Holdt Jensen. A perceptual subspace approach for modeling of speech and audio signals with damped sinusoids. *IEEE Transactions on Speech and Audio Processing*, 12(2):121–132, mars 2004.

[Johannesma, 1972] Peter I. M. Johannesma. The pre-response stimulus ensemble of neurons in the cochlear nucleus. Dans *Proc. of Symposium on Hearing Theory*, pages 58–69, Eindhoven, Pays-Bas, juin 1972.

[Johnston, 1988] James D. Johnston. Transform coding of audio signals using perceptual noise criteria. *IEEE Journal on Selected Areas in Communications*, 6(2):314–323, février 1988.

[Jolliffe, 2002] Ian T. Jolliffe. *Principal Component Analysis*. Statistical Theory and Methods. Springer, seconde édition, 2002.

[Kacha *et al.*, 2006] Ibrahim Kacha, Karim Abed-Meraim, et Adel Belouchrani. Fast adaptive blind MMSE equalizer for multichannel FIR systems. *EURASIP Journal on Applied Signal Processing*, 2006(20):1–17, 2006.

[Kim et Choi, 2007] Yong-Deok Kim et Seungjin Choi. A method of initialization for nonnegative matrix factorization. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 2, pages 537–540, Honolulu, Hawaii, USA, avril 2007.

[Klapuri, 2003] Anssi P. Klapuri. Multiple fundamental frequency estimation based on harmonicity and spectral smoothness. *IEEE Transactions on Speech and Audio Processing*, 11(6):804–816, novembre 2003.

[Klapuri, 2006] Anssi P. Klapuri. Multiple fundamental frequency estimation by summing harmonic amplitudes. Dans *Proc. of International Society for Music Information Retrieval Conference (ISMIR)*, pages 216–221, Victoria, Canada, octobre 2006.

[Knight et Lunacek, 2007] James N. Knight et Monte Lunacek. Reducing the space-time complexity of the CMA-ES. Dans *Proc. of the 9th Genetic and evolutionary computation conference (GECCO)*, pages 658–665, Londres, Royaume Uni, 2007.

[Kompass, 2007] Raul Kompass. A generalized divergence measure for nonnegative matrix factorization. *Neural Computation*, 19(3):780–791, mars 2007.

[Lagrange *et al.*, 2010b] Mathieu Lagrange, Roland Badeau, Bertrand David, Nancy Bertin, José Echeveste, Olivier Derrien, Sylvain Marchand, et Laurent Daudet. The DESAM toolbox: spectral analysis of musical audio. Dans *Proc. of the 13th International Conference on Digital Audio Effects (DAFx)*, Graz, Autriche, septembre 2010.

[Lagrange *et al.*, 2010c] Mathieu Lagrange, Roland Badeau, et Gaël Richard. Robust similarity metrics between audio signals based on asymmetrical spectral envelope matching. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 405–408, Dallas, Texas, USA, mars 2010.

[Lagrange *et al.*, 2008b] Mathieu Lagrange, Gary Scavone, et Philippe Depalle. Time-domain analysis / synthesis of the excitation signal in a source / filter model of contact sounds. Dans *Proc. of the 14th International Conference on Auditory Display (ICAD)*, Paris, France, juin 2008.

[Lagrange et Scherrer, 2008] Mathieu Lagrange et Bertrand Scherrer. Two-step modal identification for increased resolution analysis of percussive sounds. Dans *Proc. of International Conference on Digital Audio Effects (DAFx)*, Helsinki, Finlande, septembre 2008.

[Lagrange *et al.*, 2008c] Mathieu Lagrange, Nathan Whetsell, et Philippe Depalle. On the control of the phase of resonant filters with applications to percussive sound modelling. Dans *Proc. of the 11th Int. Conference on Digital Audio Effects (DAFx)*, Espoo, Finlande, septembre 2008.

[Lasalle, 1986] Joseph P. Lasalle. *The stability and control of discrete processes*. Springer-Verlag, New York, NY, USA, 1986.

[Laudadio, 2005] Teresa Laudadio. *Subspace-based quantification of Magnetic Resonance Spectroscopy data using biochemical prior knowledge*. PhD thesis, Université catholique de Louvain, Louvain, Belgique, 2005.

[Le Carrou, 2006] Jean-Loïc Le Carrou. *Vibro-acoustique de la harpe de concert*. PhD thesis, Laboratoire d'Acoustique de l'Université du Maine, Le Mans, France, décembre 2006.

[Le Carrou *et al.*, 2006] Jean-Loïc Le Carrou, François Gautier, et Roland Badeau. Analyse des modes de cordes couplées d'une harpe par une méthode à haute résolution. Dans *Actes du 8ème Congrès Français d'Acoustique (CFA)*, Tours, France, avril 2006.

[Le Carrou *et al.*, 2007] Jean-Loïc Le Carrou, François Gautier, et Roland Badeau. Theoretical and experimental investigations of harp's sympathetic modes. Dans *Proc. of International Congress on Acoustics (ICA)*, Madrid, Espagne, septembre 2007.

[Le Carrou *et al.*, 2009] Jean-Loïc Le Carrou, François Gautier, et Roland Badeau. Sympathetic string modes in the concert harp. *Acta Acustica united with Acustica*, 95(4):744–752, juillet/août 2009.

[Lee et Seung, 1999] Daniel D. Lee et H. Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, octobre 1999.

[Lee et Seung, 2001] Daniel D. Lee et H. Sebastian Seung. Algorithms for non-negative matrix factorization. Dans *Proc. of the Conference on Advances in Neural Information Processing Systems (NIPS)*, volume 13, pages 556–562, Vancouver, Colombie Britannique, Canada, décembre 2001. MIT Press.

[Leibovitz *et al.*, 2008] Guy Leibovitz, Assaf Elron, Zvi Ben-Haim, et Yonina C. Eldar. Recursive blind minimax estimation: improving MSE over recursive least squares. Dans *Proc. of the 25th IEEE Convention of Electrical and Electronics Engineers in Israel (IEEEI)*, pages 548–552, décembre 2008.

[Leveau *et al.*, 2007] Pierre LEVEAU, David SODOYER, et Laurent DAUDET. Automatic instrument recognition in a polyphonic mixture using sparse representations. Dans *Proc. of the 8th International Society for Music Information Retrieval Conference (ISMIR)*, Vienne, Autriche, septembre 2007.

[Lin, 2007] Chih-Jen LIN. Projected gradient methods for non-negative matrix factorization. *Neural Computation*, 19(10):2756–2779, octobre 2007.

[Lin et Yanfen, 2007] Wang LIN et Li YANFEN. A new semi-blind channel estimation for ZP-OFDM systems. Dans *Proc. of the 8th International Conference on Electronic Measurement and Instruments (ICEMI)*, volume 2, pages 167–171, Xi'an, Chine, août 2007.

[Liutkus *et al.*, 2010a] Antoine LIUTKUS, Roland BADEAU, et Gaël RICHARD. Gaussian processes for underdetermined source separation. *IEEE Transactions on Signal Processing*, 2010. Soumis pour publication.

[Liutkus *et al.*, 2010b] Antoine LIUTKUS, Roland BADEAU, et Gaël RICHARD. Informed source separation using latent components. Dans Vincent VIGNERON, Vicente ZARZOSO, Eric MOREAU, Rémi GRIBONVAL, et Emmanuel VINCENT, éditeurs, *Latent Variable Analysis and Signal Separation*, volume 6365 de *Lecture Notes in Computer Science*, pages 498–505, Saint Malo, France, septembre 2010. Springer.

[Liutkus *et al.*, 2011] Antoine LIUTKUS, Roland BADEAU, et Gaël RICHARD. Explaining IS-NMF as a particular case of Gaussian process regression. Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Prague, République Tchèque, mai 2011. Soumis pour publication.

[Lu et Brown, 2007] Jin LU et Lyndon J. BROWN. Control of exponentially damped sinusoidal signals. Dans *Proc. of the American Control Conference (ACC)*, New York, NY, USA, juillet 2007.

[Mallat, 2008] Stéphane MALLAT. *A Wavelet Tour of Signal Processing: The Sparse Way.* Academic Press, troisième édition, 2008.

[Mallat et Zhang, 1993] Stéphane MALLAT et Zhifeng ZHANG. Matching pursuit with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3415, décembre 1993.

[Marchand et Depalle, 2008] Sylvain MARCHAND et Philippe DEPALLE. Generalization of the derivative analysis method to non-stationary sinusoidal modeling. Dans *Proc. of the 11th International Conference on Digital Audio Effects (DAFx)*, Espoo, Finlande, septembre 2008.

[Marcos, 1998] Sylvie MARCOS, éditeur. *Les Méthodes à haute résolution: Traitement d'antenne et analyse spectrale.* Hermes Sciences Publications, 1998.

[Marcos, 2009] Sylvie MARCOS. Range recursive space-time adaptive processing (STAP) for MIMO airborne RADAR. Dans *Proc. of European Signal Processing Conference (EUSIPCO)*, Glasgow, Écosse, août 2009.

[Marolt, 2004] Matija MAROLT. A connectionist approach to automatic transcription of polyphonic piano music. *IEEE Transactions on Multimedia*, 6(3):439–449, juin 2004.

[Misono et Yamada, 2008] Masaki MISONO et Isao YAMADA. An efficient adaptive minor subspace extraction using exact nested orthogonal complement structure. *IEICE Transactions*, E91-A(8):1867–1874, août 2008.

[Moore et Glasberg, 1983] Brian C.J. MOORE et Brian R. GLASBERG. Suggested formulae for calculating auditory-filter bandwidths and excitation patterns. *Journal of the Acoustical Society of America (JASA)*, 74(3):750–753, septembre 1983.

[Nechaev et Zotov, 2008] Yu. B. NECHAEV et S. A. ZOTOV. Method for estimating the number of signals in problems of wireless direction finding with a small observation period. *Radioelectronics and Communications Systems*, 51(6):293–298, juin 2008.

[Nikunen et Virtanen, 2010] Joonas NIKUNEN et Tuomas VIRTANEN. Noise-to-mask ratio minimization by weighted non-negative matrix factorization. Dans *Proc. of International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 25–28, Dallas, Texas, USA, mars 2010.

[Nion et Sidiropoulos, 2009] Dimitri Nion et Nicholas D. Sidiropoulos. Adaptive algorithms to track the PARAFAC decomposition of a third-order tensor. *IEEE Transactions on Signal Processing*, 57(6):2299–2310, juin 2009.

[Ompusunggu *et al.*, 2009] A.P. Ompusunggu, J. Papy, S. Vandenplas, P. Sas, et H. Van Brussel. Exponential data fitting for features extraction in condition monitoring of paper-based wet clutches. Dans *Proc. of the 3rd International Operational Modal Analysis Conference (IOMAC)*, pages 323–330, Porto Novo, Ancona, Italie, mai 2009.

[Oseledets et Tyrtyshnikov, 2010] Ivan Oseledets et Eugene Tyrtyshnikov. TT-cross approximation for multidimensional arrays. *Linear Algebra and its Applications*, 432(1):70–88, janvier 2010.

[Ouyang et Hua, 2005] Shan Ouyang et Yingbo Hua. Bi-iterative least-square method for subspace tracking. *IEEE Transactions on Signal Processing*, 53(8):2984– 2996, août 2005.

[Papy *et al.*, 2007] Jean-Michel Papy, Lieven De Lathauwer, et Sabine Van Huffel. A shift invariance-based order-selection technique for exponential data modelling. *IEEE Signal Processing Letters*, 14(7):473–476, juillet 2007.

[Parvaix *et al.*, 2010] Mathieu Parvaix, Laurent Girin, et Jean-Marc Brossier. A watermarking-based method for informed source separation of audio signals with a single sensor. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(6):1464–1475, août 2010.

[Paulus et Klapuri, 2002] Jouni Paulus et Anssi Klapuri. Measuring the similarity of rhythmic patterns. Dans *Proc of the 3rd International Society for Music Information Retrieval Conference (ISMIR)*, Paris, France, octobre 2002.

[Perry et Wolfe, 2009] Patrick O. Perry et Patrick J. Wolfe. Minimax rank estimation for subspace tracking. *IEEE Journal of Selected Topics in Signal Processing*, 4(3):504–513, juin 2009.

[Pertusa et Iñesta, 2008] Antonio Pertusa et José M. Iñesta. Multiple fundamental frequency estimation using Gaussian smoothness. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 105–108, Las Vegas, Nevada, USA, avril 2008.

[Plumbley *et al.*, 2010] Mark D. Plumbley, Thomas Blumensath, Laurent Daudet, Rémi Gribonval, et Mike E. Davies. Sparse representations in audio and music: From coding to source separation. *Proceedings of the IEEE*, 98(6):995–1005, juin 2010.

[Priestley, 1983] Maurice B. Priestley. *Spectral Analysis and Time Series*. Probability and Mathematical Statistics. Academic Press, février 1983.

[Prosonus, 1998] Prosonus. Prosonus Grand Piano. Audio CD, 1998. Distribué par Big Fish Audio (`http://www.bigfishaudio.com/`).

[Rabiner et Juang, 1993] Lawrence Rabiner et Biing-Hwang Juang. *Fundamentals of Speech Recognition*. Signal Processing. Prentice Hall, USA, avril 1993.

[Raczyński *et al.*, 2007] Stanislaw Andrzej Raczyński, Nobutaka Ono, et Shigeki Sagayama. Multipitch analysis with harmonic nonnegative matrix approximation. Dans *Proc. of the 8th International Society for Music Information Retrieval Conference (ISMIR)*, Vienne, Autriche, septembre 2007.

[Ravelli *et al.*, 2008] Emmanuel Ravelli, Gaël Richard, et Laurent Daudet. Union of MDCT bases for audio coding. *IEEE Transactions on Audio, Speech, and Language Processing*, 16(8):1361–1372, novembre 2008.

[Röbel *et al.*, 2007] Axel Röbel, Fernando Villavicencio, et Xavier Rodet. On cepstral and all-pole based spectral envelope modeling with unknown model order. *Pattern Recognition Letters*, 28(11):1343–1350, août 2007.

[Reissell et Pai, 2007a] L.-M. Reissell et Dinesh K. Pai. High resolution analysis and resynthesis of environmental impact sounds. Dans *Proc. of the 13th International Conference on Auditory Display (ICAD)*, pages 189–196, Montréal, Québec, Canada, juin 2007.

[Riche de Prony, 1795] Gaspard-Marie Riche de Prony. Essai expérimental et analytique: sur les lois de la dilatabilité de fluides élastiques et sur celles de la force expansive de la vapeur de l'eau et de la vapeur de l'alcool à différentes températures. *Journal de l'école polytechnique*, 1(22):24–76, 1795.

[Roy et Kailath, 1989] Richard H. Roy et Thomas Kailath. ESPRIT-estimation of signal parameters via rotational invariance techniques. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(7):984–995, juillet 1989.

[Scheirer, 1998] Eric D. Scheirer. Tempo and beat analysis of acoustic music signals. *The Journal of the Acoustical Society of America (JASA)*, 103(1):588–601, janvier 1998.

[Schmidt, 2009] Mikkel N. Schmidt. Function factorization using warped Gaussian processes. Dans *Proc. of International Conference on Machine Learning (ICML)*, Montréal, Québec, Canada, juin 2009.

[Schmidt et Laurberg, 2008] Mikkel N. Schmidt et Hans Laurberg. Non-negative matrix factorization with Gaussian process priors. *Computational Intelligence and Neuroscience*, 2008:1–10, 2008. Article ID 361705.

[Schmidt, 1986] Ralph. O. Schmidt. Multiple emitter location and signal parameter estimation. *IEEE Transactions on Antennas and Propagation*, 34(3):276–280, mars 1986.

[Schwarz, 1978] Gideon E. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6(2):461–464, 1978.

[Shashanka *et al.*, 2008] Madhusudana Shashanka, Bhiksha Raj, et Paris Smaragdis. Probabilistic latent variable models as nonnegative factorizations. *Computational Intelligence and Neuroscience*, 2008:1–8, 2008. Article ID 947438.

[Singh *et al.*, 2010] Rita Singh, Bhiksha Raj, et Paris Smaragdis. Latent-variable decomposition based dereverberation of monaural andmulti-channel signals. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1914–1917, Dallas, Texas, USA, mars 2010.

[Slavnicu et Ciochina, 2006] Stefan Slavnicu et Silviu Ciochina. Gradient subspace method for tracking real-valued sinusoidal carriers. Dans *Proc. of International Conference on Wireless and Mobile Communications (ICWMC)*, pages 70–72, Bucarest, Roumanie, juillet 2006.

[Smaragdis, 2004] Paris Smaragdis. Non-negative matrix factor deconvolution; extraction of multiple sound sources from monophonic inputs. Dans *Proc. of International Conference on Independent Component Analysis and Blind Signal Separation (ICA)*, pages 494–499, Grenade, Espagne, septembre 2004.

[Smaragdis, 2009] Paris Smaragdis. Relative pitch tracking of multiple arbitrary sounds. *Journal of the Acoustical Society of America (JASA)*, 125(5):3406–3413, mai 2009.

[Smaragdis et Brown, 2003] Paris Smaragdis et Judith C. Brown. Non-negative matrix factorization for polyphonic music transcription. Dans *Proc. of IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pages 177–180, New Paltz, New York, USA, octobre 2003.

[Spanias *et al.*, 2007] Andreas Spanias, Ted Painter, et Venkatraman Atti. *Audio Signal Processing and Coding*. Wiley-Interscience, 2007.

[Stoica et Söderström, 1991] Petre Stoica et Torsten Söderström. Statistical analysis of MUSIC and subspace rotation estimates of sinusoidal frequencies. *IEEE Transactions on Signal Processing*, 39(8):1836–1847, août 1991.

[Strobach, 1996] Peter Strobach. Low-rank adaptive filters. *IEEE Transactions on Signal Processing*, 44(12):2932–2947, décembre 1996.

[Strobach, 1997a] Peter Strobach. Bi-iteration SVD subspace tracking algorithms. *IEEE Transactions on Signal Processing*, 45(5):1222–1240, mai 1997.

[Strobach, 1997b] Peter Strobach. Square hankel SVD subspace tracking algorithms. *Signal Processing*, 57(1):1–18, février 1997.

[Strobach, 2009a] Peter Strobach. The fast recursive row-Householder subspace tracking algorithm. *Signal Processing*, 89(12):2514–2528, décembre 2009.

[Strobach, 2009b] Peter Strobach. The Householder compressor theorem and its application in subspace tracking. *Signal Processing*, 89(5):857–875, mai 2009.

[Strobach, 2009c] Peter Strobach. Square-root Householder subspace tracking. *Numerische Mathematik*, 113(1):89–121, juillet 2009.

[Strobach, 2010] Peter Strobach. Sliding window adaptive SVD using the unsymmetric Householder partial compressor. *Signal Processing*, 90(1):352–362, janvier 2010.

[Sun *et al.*, 2007] Liping Sun, Guoan Bi, et Liren Zhang. Orthonormal subspace tracking algorithm for space-time multiuser detection in multipath CDMA channels. *IEEE Transactions on Vehicular Technology*, 56(6):3838–3845, novembre 2007.

[Tibshirani, 1996] Robert Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, janvier 1996.

[Tolonen et Karjalainen, 2000] Tero Tolonen et Matti Karjalainen. A computationally efficient multipitch analysis model. *IEEE Transactions on Speech and Audio Processing*, 8(6):708–716, juin 2000.

[Tong et Perreau, 1998] Lang Tong et Sylvie Perreau. Multichannel blind identification: From subspace to maximum likelihood methods. *Proceedings of IEEE*, 86(10):1951–1968, octobre 1998.

[Toolan et Tufts, 2006] Timothy M. Toolan et Donald W. Tufts. Efficient and accurate rectangular window subspace tracking. Dans *Proc. of the 4th IEEE Workshop on Sensor Array and Multichannel Processing (SAM)*, pages 60–64, Waltham, MA, USA, juillet 2006.

[Valizadeh *et al.*, 2007] Amir Valizadeh, Mahmood Karimi, et Omid Oliaei. An efficient algorithm for mutiple direction of arrival tracking based on the constrained projection approximation approach and Kalman filter. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 2, pages 1089–1092, Honolulu, Hawaii, USA, avril 2007.

[Van Rijsbergen, 1979] C.J Keith Van Rijsbergen. *Information Retrieval*. Butterworth-Heinemann, Londres, Royaume Uni, seconde édition, mars 1979.

[Vincent *et al.*, 2008] Emmanuel Vincent, Nancy Bertin, et Roland Badeau. Harmonic and inharmonic nonnegative matrix factorization for polyphonic pitch transcription. Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 109–112, Las Vegas, Nevada, USA, avril 2008.

[Vincent *et al.*, 2010] Emmanuel Vincent, Nancy Bertin, et Roland Badeau. Adaptive harmonic spectral decomposition for multiple pitch estimation. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(3):528–537, mars 2010.

[Virtanen, 2007] Tuomas Virtanen. Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria. *IEEE Transactions on Audio, Speech, and Language Processing*, 15(3):1066–1074, mars 2007.

[Virtanen *et al.*, 2008] Tuomas Virtanen, A. Taylan Cemgil, et Simon Godsill. Bayesian extensions to nonnegative matrix factorisation for audio signal modelling. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1825–1828, Las Vegas, Nevada, USA, avril 2008.

[Wang et Ouyang, 2009] Yajun Wang et Shan Ouyang. Inverse QR iterative method for subspace tracking. *Journal of Electronics*, 26(2):270–273, février 2009.

[Wild *et al.*, 2004] Stefan Wild, James Curry, et Anne Dougherty. Improving non-negative matrix factorizations through structured initialization. *Pattern Recognition*, 37(11):2217–2232, novembre 2004.

[Willink, 2008] Tricia J. Willink. Efficient adaptive SVD algorithm for MIMO applications. *IEEE Transactions on Signal Processing*, 56(2):615–622, février 2008.

[Wong et Evans, 2008] Ian C. Wong et Brian L. Evans. Sinusoidal modeling and adaptive channel prediction in mobile OFDM systems. *IEEE Transactions on Signal Processing*, 56(4):1601–1615., avril 2008.

[Wright *et al.*, 2009] John Wright, Allen Y. Yang, Arvind Ganesh, S. Shankar Sastry, et Yi Ma. Robust face recognition via sparse representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(2):210–227, février 2009.

[Wu *et al.*, 2009] Tong Tong Wu, Yi Fang Chen, Trevor Hastie, Eric Sobel, et Kenneth Lange. Genome-wide association analysis by Lasso penalized logistic regression. *Bioinformatics*, 25(6):714–721, janvier 2009.

[Yang, 1995] Bin Yang. Projection approximation subspace tracking. *IEEE Transactions on Signal Processing*, 44(1):95–107, janvier 1995.

[Young *et al.*, 2006] S. Stanley Young, Paul Fogel, et Douglas M. Hawkins. Clustering Scotch whiskies using non-negative matrix factorization. *Joint Newsletter for the Section on Physical and Engineering Sciences and the Quality and Productivity Section of the American Statistical Association*, 14(1):11–13, juin 2006.

[Zhang *et al.*, 2008] Johan Xi Zhang, Mads Græsbøll Christensen, Joachim Dahl, Søren Holdt Jensen, et Marc Moonen. Frequency-domain parameter estimations for binary masked signals. Dans *Proc. of the 9th annual conference of the international speech communication association (Interspeech)*, Brisbane, Australie, septembre 2008.

[Zhang *et al.*, 2010] Wei-Tao Zhang, Shun-Tian Lou, et Yan-Liang Zhang. Robust nonlinear power iteration algorithm for adaptive blind separation of independent signals. *Digital Signal Processing*, 20(2):541–551, mars 2010.

[Zhang et Fang, 2007] Ye Zhang et Yong Fang. A NMF algorithm for blind separation of uncorrelated signals. Dans *Proc. of IEEE International Conference on Wavelet Analysis and Pattern Recognition (ICWAPR)*, pages 999–1003, Beijing, Chine, novembre 2007.

Index

Troisième partie

Annexes

A. Curriculum Vitae détaillé

Dr. Roland BADEAU
Né le 28/08/1976 à Marseille
Nationalité française
Ingénieur en Chef du Corps des Mines
***Senior Member* de la société savante IEEE**

Maître de Conférences
Institut Télécom, Télécom ParisTech, CNRS LTCI
37/39 rue Dareau, 75014 Paris
Courriel : roland.badeau@telecom-paristech.fr
Page personnelle : `http://perso.telecom-paristech.fr/rbadeau/`

A.1 Formations et diplômes

Formation initiale

2001-2005 : Diplôme de Docteur de l'École Nationale Supérieure des Télécommunications (ENST), Paris, spécialité Signal et Images (thèse soutenue le 18 avril 2005). Sujet de thèse : "Méthodes à haute résolution pour l'estimation et le suivi de sinusoïdes modulées. Application aux signaux de musique" [1]. Mention Très Honorable. **Prix de thèse ParisTech 2006**.

2000-2001 : Diplôme d'Études Approfondies de l'École Normale Supérieure (ENS) de Cachan. DEA MVA (mathématiques / vision / apprentissage). Mention très bien.

1999-2001 : Diplôme d'ingénieur de l'ENST, Paris.
Dominante Signal et Images, option Audiovisuel et Multimédia.

1996-1999 : Diplôme d'ingénieur de l'École Polytechnique (X), Palaiseau.
Majeures de Mathématiques Appliquées et d'Informatique.

Formation continue

- **Gestion de projets et valorisation de la recherche**

  Octobre 2009 : Les règles de base de la propriété intellectuelle et de la valorisation (Télécom ParisTech).

  Juin 2006 : Montage et négociation des projets européens (Télécom ParisTech).

  Novembre 2005 : Les acteurs et contrats de la recherche (Télécom ParisTech).

- **Pédagogie et encadrement**

  Novembre 2008 et mars 2009 : Formation pédagogique ParisTech.

  Décembre 2008 et janvier 2009 : De l'encadrement du projet doctoral au devenir professionnel du docteur (Université Pierre et Marie Curie (UPMC), Paris).

- **Anglais**

  Depuis 2008 : Anglais avancé (Télécom ParisTech).

  2002-2003 : Cours d'anglais de professionnalisation (Télécom ParisTech).

A.2 Expériences professionnelles

Depuis déc. 2005 : Maître de Conférences au département Traitement du Signal et des Images (TSI) de Télécom ParisTech, Paris, groupe Audio, Acoustique et Ondes (AAO), projet structurant Audiosig.

Sep. 2001-nov. 2005 : Chargé d'Enseignement et de Recherche au département TSI de l'ENST, Paris.

Mars 2000-août 2000 : Stage de recherche et développement en traitement d'images ultrasonores effectué pour la formation du Corps Interministériel des Télécommunications, au sein des Laboratoires d'Électronique Philips, Limeil-Brévannes, en collaboration avec une société américaine spécialisée dans les ultrasons. Demande de brevet. Sujet : "Segmentation et suivi de la paroi cardiaque en imagerie ultrasonore 2D".

Avril 1999-juin 1999 : Stage de recherche en traitement des images radar effectué pour la formation d'ingénieur de l'X, au sein du département TSI de l'ENST. Félicitations du jury des stages d'option de l'X. Deux publications [59, 61]. Sujet : "Transformation de Mellin : aspects numériques et dérivation fractionnaire".

A.3 Production scientifique

Ma bibliographie complète est présentée dans l'annexe B page 66. Elle comprend les références suivantes :

Publications scientifiques :

- **Thèse de doctorat** [1] ;
- **17 articles de revues internationales** [2-18] ;
- **42 articles de conférences internationales** avec comité de lecture [19-60] ;
- **9 rapports techniques** [61-69].

Autres références :

- **Brevets**, logiciel, base de données [70-73] ;
- Rédaction de **4 supports de cours** [74-77] ;
- Participation à **5 conférences** internationales sans actes [78-82].

D'après le logiciel à télécharger du site *Publish or Perish* (`http://www.harzing.com/`), mon indice H (ou indice de Hirsch) est égal à 12 au 15/09/2010. Voici la liste des 12 publications les plus citées :

Citations	Réf.	Titre	Année	Publication
42	[4]	Fast approximated power iteration subspace tracking	2005	IEEE TSP [1]
38	[8]	Sliding window adaptive SVD algorithms	2004	IEEE TSP [1]
27	[60]	Harmonic and inharmonic nonnegative matrix factorization for polyphonic pitch transcription	2008	IEEE ICASSP [2]
25	[5]	High resolution spectral analysis of mixtures of complex exponentials modulated by polynomials	2006	IEEE TSP [1]
24	[26]	Yet another subspace tracker	2005	IEEE ICASSP [2]
23	[6]	A new perturbation analysis for signal enumeration in rotational invariance techniques	2006	IEEE TSP [1]
23	[23]	EDS parametric modeling and tracking of audio signals	2002	DAFx [3]
20	[34]	Blind signal decompositions for automatic transcription of polyphonic music : NMF and K-SVD on the benchmark	2007	IEEE ICASSP [2]
17	[32]	Fast adaptive ESPRIT algorithm	2005	IEEE SSP [4]
16	[25]	Selecting the modeling order for the ESPRIT high resolution method : an alternative approach	2004	IEEE ICASSP [2]
13	[10]	Performance of ESPRIT for estimating mixtures of complex exponentials modulated by polynomials	2008	IEEE TSP [1]
12	[19]	Musical tempo estimation using noise subspace projections	2003	IEEE WASPAA [5]

[1] *IEEE Transactions on Signal Processing.*
[2] *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing.*
[3] *Proc. of International Conference on Digital Audio Effects.*
[4] *Proc. of IEEE Workshop on Statistical Signal Processing.*
[5] *Proc. of IEEE Workshop on Applications of Signal Processing to Audio and Acoustics.*

A.4 Encadrements de travaux de recherche

Encadrements de thèses réalisés à Télécom ParisTech

Depuis janvier 2010 : Co-encadrement à 50% de la thèse d'Antoine Liutkus (co-directeur : Gaël Richard). Sujet : "Séparation de sources informée".

Depuis octobre 2009 : Co-encadrement à 50% de la thèse de Benoît Fuentes (co-directeur : Gaël Richard). Sujet : "Estimation de hauteurs multiples pour la transcription de musique polyphonique".

Depuis octobre 2008 : Co-encadrement à 50% de la thèse de Romain Hennequin, évaluée à mi-parcours le 5 mai 2010 (co-directeur : Bertrand David). Sujet : "Décomposition structurée de spectrogrammes musicaux".

2005-2009 : Co-encadrement à 90% de la thèse de Nancy Bertin, soutenue le 2 octobre 2009 (co-directeur : Gaël Richard). Sujet : "Les factorisations en matrices non-négatives. Approches contraintes et probabilistes, application à la transcription automatique de musique polyphonique" [83]. Mention Très Honorable. **Prix de thèse 2010 du GdR-Isis/EEA/GRETSI**.

2005-2008 : Co-encadrement à 40% de la thèse de Valentin Emiya, soutenue le 8 octobre 2008 (co-directeur : Bertrand David). Sujet : "Transcription automatique de la musique de piano" [84]. Mention Très Honorable.

Encadrements de stages réalisés à Télécom ParisTech

2007 : Co-encadrement à 50% du stage de Sylvain Streiff effectué pour le Master Sciences et Technologies de l'UPMC, parcours Acoustique, Traitement du signal, Informatique Appliqués à la Musique (ATIAM) (co-directeur : Bertrand David). Sujet : "Représentations à haute résolution pour l'estimation de fréquences fondamentales".

2006 : Encadrement à 100% du stage de Sylvain Rousselle effectué au titre du parcours ATIAM.
Sujet : "Séparation de la voix du locuteur et du fond musical dans des émissions radiodiffusées" [85].

2006 : Co-encadrement à 50% du stage de Sylvain Streiff effectué pour la formation de l'X (co-directeur : Bertrand David). Sujet : "Apprentissage automatique supervisé appliqué à la transcription automatique du piano".

Encadrement de projets d'élèves ingénieurs Télécom ParisTech

2008 : Projet de Al-Amin Alibay Premjee. Sujet : "Optimisation de portefeuille".

2008 : Projet de Romain Hennequin, co-encadré avec Bertrand David. Sujet : "Analyse de notes de piano".

2007 : Projet de Benoît Fuentes. Sujet : "Estimation de hauteur".

2006 : Projet de M. Garin et S. Kersulec. Sujet : "Segmentation Parole/Musique d'un document sonore".

2005 : Projet de Nicolas Durand.
Sujet : "Estimation de hauteurs multiples pour l'extraction de partitions à l'aide de signaux musicaux".

2004 : Projet de Nancy Bertin.
Sujet : "Poursuite de sous-espace, extraction du bruit d'attaque dans des notes de musique".

2003 : Projet de Nicolas Carayon, co-encadré avec Gaël Richard. Sujet : "Application des méthodes à haute résolution pour l'estimation de fréquences fondamentales multiples".

2002 : Projet de David Cournapeau. Sujet : "Estimation de hauteur".

A.5 Activités de recherche

A.5.1 Thèmes de recherche

- **Traitement de signaux audio et applications musicales** (modèles sinusoïdaux, modélisation de signaux non-stationnaires, analyse/synthèse, effets sonores, codage audio, estimation de hauteurs multiples, estimation du tempo, transcription automatique, séparation de sources)

- **Analyse spectrale à haute résolution** (estimation de l'ordre de modélisation, algorithmes adaptatifs rapides, poursuite de sous-espace, analyse de performances)
- **Décompositions non-négatives** (factorisation en matrices non-négatives, estimation bayésienne de mélanges de gaussiennes, contraintes temporelles et spectrales, convergence des algorithmes)

A.5.2 Participation à des projets de recherche

Depuis oct. 2009 : DReaM (le Disque Repensé pour l'écoute active de la Musique). L'objectif de ce projet est la réalisation d'un système d'écoute active de la musique. Il s'agit de permettre à un auditeur de modifier lui-même la musique pendant sa diffusion, en manipulant séparément les différentes sources présentes dans le mélange musical. Ce projet met en œuvre un concept nouveau, celui de la séparation de sources informée. Il est coordonné par le Laboratoire Bordelais de Recherche en Informatique (LaBRI) de l'Université Bordeaux 1 pour une durée de quatre ans et financé par le programme CONTenus et INTeractions (CONTINT) de l'ANR. Il regroupe quatre partenaires académiques, dont Télécom ParisTech, Paris, le laboratoire Grenoble-Images-Parole-Signal-Automatique (GIPSA-Lab), le laboratoire CNRS Institut Jean le Rond d'Alembert (IJLRD) de l'UPMC, Paris, et un industriel : iKlax Media, Bidart. **Rôle dans le projet :** j'ai co-écrit au titre de ce projet un article de revue [17] et un article de conférence [57].

Depuis mai 2008 : QUAERO. Il s'agit d'un programme fédérateur de recherche et d'innovation industrielle sur les technologies d'analyse automatique, de classification et d'utilisation de documents multimédia. Ce projet est coordonné par la société Technicolor pour une durée de cinq ans et financé par l'OSEO. Le consortium se compose de vingt-six partenaires industriels et organismes de recherche français et allemands des secteurs public et privé. **Rôle dans le projet :** j'ai co-publié au titre de ce projet quatre articles de conférences [35, 54, 49, 50].

Nov. 2006 - fév. 2010 : DESAM (Décompositions en Éléments Sonores et Applications Musicales). Ce projet consiste à mettre en œuvre des méthodes innovantes de traitement du signal, afin de décomposer des signaux de musique en éléments sonores. Ces décompositions sont appliquées à l'extraction de structures musicales (instruments, rythme, hauteurs de notes), à l'analyse physique et à la synthèse audio. Ce projet a été coordonné par le Laboratoire Traitement et Communication de l'Information (LTCI) de Télécom ParisTech, Paris, et financé par le programme Jeunes Chercheuses, Jeunes Chercheurs (JCJC) de l'ANR. Il a regroupé quatre partenaires académiques : le LaBRI de l'Université Bordeaux 1, l'équipe Lutheries-Acoustique-Musique (LAM) de l'Université Paris VI, le laboratoire Sciences et Technologies de l'Information et de la Communication (STIC) de l'Université de Toulon et du Var et le LMA, Marseille. **Rôle dans le projet :** j'ai co-publié au titre de ce projet six articles de revue [16, 15, 11, 18, 12, 2] et seize articles de conférences [47, 56, 45, 39, 24, 60, 78, 79, 82, 46, 29, 36, 35, 54, 49, 53]. **J'ai personnellement coordonné ce projet** et contribué au développement de la boîte à outils Matlab *DESAM Toolbox* [73] (*cf.* section A.7.1).

Jan. 2006 - mai 2009 : INFOMAGIC. L'objectif de ce projet est d'étudier et de proposer des prototypes logiciels de fonctions avancées d'analyse multimodale de données numériques à des échelles allant d'un seul ordinateur au réseau Internet, aussi bien pour des applications industrielles que pour des applications grand public. Il a été coordonné par la société Thalès, Neuilly-sur-Seine, et financé par le Pôle de compétitivité Cap Digital IMVN (Image, Multimédia et Vie Numérique). Il a regroupé plus de vingt partenaires académiques et industriels français. **Rôle dans le projet :** j'ai contribué dans ce projet à la production d'un livrable, lors de l'encadrement du stage de Sylvain Rousselle à Télécom ParisTech [85].

Jan. 2006 - jan. 2009 : KSPACE (*Knowledge space of semantic inference for automatic annotation and retrieval of multimedia content*). L'objectif de ce projet est de réduire la distance qui existe entre les descripteurs de contenus qui peuvent être extraits automatiquement et le niveau d'abstraction nécessaire à l'interprétation humaine. Ce réseau d'excellence a été coordonné par l'université Queen Mary, Londres, et financé par le 6ème programme cadre IST (*Information Society Technologies*) de la communauté européenne. Il regroupe quatorze partenaires académiques européens. **Rôle dans le projet :** j'ai co-publié au titre de ce projet trois articles de revues [9, 11, 12], six articles de conférences [27, 28, 34, 46, 29, 36] et rédigé un livrable [62].

Jan. 2007 - déc. 2007 : TAM-TAM (Transcription Automatique de la Musique : Traitements Avancés et Mise en œuvre). Ce projet s'intéresse à la transcription automatique des morceaux de piano. Il a été financé

par l'Institut Télécom sur Crédits Incitatifs et a regroupé deux écoles de l'Institut : Télécom ParisTech, Paris, et EURECOM, Sophia Antipolis. **Rôle dans le projet :** j'ai co-publié au titre de ce projet 4 articles de conférences [81, 39, 24, 80].

Nov. 2004 - nov. 2007 : Music Discover. Ce projet porte sur l'analyse, l'indexation, la représentation et la recherche d'informations dans les enregistrements musicaux, orientés vers le contenu et adaptés à l'utilisateur. Il a été coordonné par l'équipe Analyse/Synthèse de l'Institut de Recherche et Coordination Acoustique/Musique (IRCAM), Paris, et financé par le Ministère délégué à la Recherche et aux Nouvelles Technologies (MNRT) dans le cadre de l'Action Concertée Incitative (ACI) Masse de Données. Il a regroupé trois partenaires académiques, dont Télécom ParisTech, Paris, et le Laboratoire d'InfoRmatique en Image et Systèmes d'information (LIRIS), École Centrale Lyon. **Rôle dans le projet :** j'ai co-publié au titre de ce projet un article de revue [9] et quatre articles de conférences [27, 28, 34, 45].

A.5.3 Autres travaux collaboratifs hors projet de recherche

Depuis 2008 : Collaboration avec Emmanuel Vincent de l'Institut de Recherche en Informatique et Systèmes Aléatoires (IRISA), Rennes, sur le thème de la factorisation en matrices non-négatives et ses applications audio [2, 11, 18, 60, 35, 66, 69].

Depuis 2006 : Collaboration avec Rémy Boyer du LSS à Supélec, Gif-sur-Yvette, sur le développement de techniques rapides de traitement tensoriel du signal [3, 37, 38].

2006 - 2009 : Collaboration avec Jean-Loïc Le Carrou et François Gautier du Laboratoire d'Acoustique de l'Université du Maine (LAUM), Le Mans, sur l'analyse modale de la harpe à l'aide de méthodes à haute résolution [16, 55, 56].

2006 : Collaboration avec Jon Altuna et Vicente Atxa de l'université de Mondragón en Espagne sur le développement de méthodes sous-espace pour les communications numériques [20].

A.5.4 Travaux d'expertise

A.5.4.1 Relectures d'articles (*reviewing*)

Relecture de 38 articles de revues :

- 12 articles pour la revue *Transactions on Signal Processing* (IEEE),
- 8 articles pour la revue *Signal Processing Letters* (IEEE),
- 5 articles pour la revue *Signal Processing* (Elsevier),
- 7 articles pour la revue *Transactions on Audio, Speech and Language Processing* (IEEE),
- 1 article pour les revues *Transactions on Speech and Audio Processing* (IEEE), *Transactions on Neural Networks* (IEEE), *Signal Processing* (IET), *Proceedings on Vision, Image, and Signal Processing* (IEE), *Probability and statistics* (ESAIM) et *International Journal of Adaptive Control and Signal Processing* (John Wiley & Sons).

J'ai reçu en 2007 une lettre de remerciement de l'Éditeur en Chef de la revue *IEEE Transactions on Signal Processing*, pour mes travaux de relecture.

Relecture de 29 articles de conférences :

- 14 articles pour la conférence *Inf. Science, Signal Proc. and their Applications* (ISSPA 2003 et 2007),
- 6 articles pour la conférence *Int. Workshop on Signal Proc. and its Applications* (IEEE WoSPA, 2008),
- 3 articles pour la conférence *Int. Society for Music Information Retrieval conference* (ISMIR 2006),
- 2 articles pour la conférence *Global Communications Conference* (IEEE GlobeCom 2006 et 2008),
- 1 article pour les conférences *International Workshop on Acoustic Echo and Noise Control* (IWAENC 2006), *International Symposium on Information Theory* (IEEE ISIT 2007), *International Conference on Multimedia and Expo* (IEEE ICME 2010), *Multimedia Signal Processing Conference* (IEEE MMSP 2010).

A.5.4.2 Participations à des comités de programmes de conférences

- **MMSP 2010 :** *IEEE Multimedia Signal Processing Conference*
- **WoSPA 2008 :** *IEEE International Workshop on Signal Processing and its Applications*
- **ISSPA 2007 :** *Information Science, Signal Processing and their Applications*

A.5.4.3 Participations à des jurys de thèses

Sept. 2010 : Soutenance de thèse de Mathieu Parvaix au GIPSA-Lab, Grenoble (en tant qu'**examinateur**). Sujet : "La séparation de sources sonores informée par tatouage audio-numérique".

Oct. 2009 : Soutenance de thèse de Nancy Bertin à Télécom ParisTech, Paris, spécialité Signal et Images (en tant que **directeur de thèse**). Sujet : "Les factorisations en matrices non-négatives. Approches contraintes et probabilistes, application à la transcription automatique de musique polyphonique" [83].

Oct. 2008 : Soutenance de thèse de Valentin Emiya à Télécom ParisTech, Paris, spécialité Signal et Images (en tant que **directeur de thèse**). Sujet : "Transcription automatique de la musique de piano" [84].

Déc. 2007 : Soutenance de thèse de Guillaume Bouleux à Supélec, Gif-sur-Yvette, spécialité Traitement du Signal (en tant qu'**examinateur**). Sujet : "Traitement d'antenne adapté aux modèles linéaires intégrant une interférence structurée".

Déc. 2006 : Soutenance de thèse de Jean-Loïc Le Carrou au LAUM, Le Mans, spécialité Acoustique (en tant qu'**examinateur**). Sujet : "Vibro-acoustique de la harpe de concert".

A.6 Activités d'enseignement

Parallèlement à mes activités de recherche, j'effectue des activités d'enseignement depuis 2001. J'ai ainsi totalisé plus de 900 heures, soit **un peu plus d'une centaine d'heures d'enseignement par an en moyenne**, dont 60% de travaux pratiques, 25% de cours, 10% de travaux dirigés et 5% consacrées aux examens.

A.6.1 Enseignements en formation initiale

A.6.1.1 Activités à l'IRCAM

A l'IRCAM, je suis membre de l'équipe pédagogique du Master Sciences et Technologies de l'UPMC, parcours ATIAM. J'interviens plus particulièrement dans les deux unités d'enseignement suivantes (129 heures au total) :

Depuis 2006 : Traitement du Signal Audio (TSA)

Depuis 2004 : Élaboration et Transformation de Sons (ETS)

Par ailleurs, je participe aux jurys de stages et à l'audition des candidats au Master (depuis 2007).

A.6.1.2 Formation d'ingénieur Télécom ParisTech

A Télécom ParisTech, j'interviens dans la formation initiale des élèves ingénieurs, en 1ère, 2ème et 3ème années, et j'encadre régulièrement des projets d'élèves.

Enseignements de 1ère année (tronc commun)

Depuis 2007 : Outils et Applications pour le Signal, les Images et le Son (OASIS, 20 heures au total)

2005-2009 : Introduction au Traitement du Signal et des Images (ITSI, 50 heures au total)

2001-2007 : Bases de Traitement du Signal (BTS, 64 heures au total)

Enseignements de 2ème et 3ème année (Cycle Master)

- Dans le domaine **Traitement du signal et des images** (TSI)

  Depuis 2006 : Signal Audiofréquence et Parole (SAP, 15 heures au total)

  2008-2009 : Étude de Cas en Signal (ECS, 5 heures au total)

  2007-2008 : Traitement du Signal Adaptatif (TSA, 2 heures au total)

  2006-2008 : Traitement du Signal : Méthodes et Études de Cas (TSMEC, 29 heures au total)

  2006-2007 : Projets de Traitement du Signal (PTS)

  2001-2006 : Traitement du Signal : Méthodes et Applications (TSMAp, 286 heures au total)

  2001-2006 : Traitement du Signal : Études de Cas et Simulations (TSECS, 106 heures au total)

2001-2006 : Parole - Acoustique - Musique (PAMU, 92 heures au total)

2001-2006 : Traitement du Signal Adaptatif et Vectoriel (TSAV, 32 heures au total)

2001-2002 : Fréquence-Espace-Temps-Échelle (FETE, 10 heures au total)

- Dans le domaine **Mathématiques de l'ingénieur** (MDI)

  Depuis 2006 : Méthodes d'Optimisation Continue et Applications (MOCA, 69 heures au total)

  2006-2008 : Hilbert-Fourier (HF, 14 heures au total)

A.6.2 Enseignements en formation continue à Télécom ParisTech

Depuis 2006 : Traitement du Signal Audionumérique (3 heures au total).

A.6.3 Création de contenus pédagogiques

- **Rédaction de polycopiés de cours.**
  Dans le cadre de mes enseignements dans le Cycle Master de Télécom ParisTech, j'ai rédigé 4 polycopiés de cours : "Détection de hauteur" [74], "Estimation de sinusoïdes" [75], "Méthodes à haute résolution" [76] et "Transformées de Fourier à Temps Discret et à Temps Continu" [77].
- **Élaboration de "Sites Pédagogiques".**
  J'ai publié sur le site *ParisTech Graduate School* les contenus pédagogiques des unités d'enseignement "Méthodes d'Optimisation Continue et Applications" et "Traitement du Signal : Méthodes et Études de Cas" (`http://graduateschool.paristech.fr/`).

A.7 Responsabilités administratives et activités d'intérêt collectif

A.7.1 Coordination du projet de recherche DESAM

Le projet DESAM, réalisé de novembre 2006 à février 2010, est un projet du programme JCJC de l'ANR, auquel ont participé cinq laboratoires de recherche français. Les résultats du projet ont fait l'objet de sept articles de revues (dont trois pour le hors série *Signal Models and Representations of Musical and Environmental Sounds* du journal *IEEE Transactions on Acoustics, Speech and Language Processing*) et de vingt-deux articles de conférences. Le projet a été particulièrement mis en valeur lors du congrès américano-européen *Acoustics'08*, qui s'est tenu à Paris en juillet 2008, avec cinq interventions dans la session *Representations and models for musical sounds*. Enfin, les résultats obtenus ont été concrétisés par l'implémentation d'une boîte à outils Matlab, baptisée *DESAM Toolbox*. Il s'agit d'un ensemble de fonctions Matlab s'adressant à la communauté des chercheurs intéressés par la modélisation de signaux audio musicaux. Outre mes contributions scientifiques, mon rôle de coordinateur du projet a inclus un certain nombre de responsabilités administratives :

- Soumission du projet à l'ANR (2006),
- Supervision de l'accord de consortium (2007),
- Gestion du budget confié au LTCI (2006-2010),
- Recrutement en CDD de Mathieu Lagrange au LTCI pour une durée de 12 mois (octobre 2008- septembre 2009), pour une mission de recherche et de développement de la *DESAM Toolbox*,
- Organisation des réunions trimestrielles d'avancement du projet,
- Rédaction des rapports semestriels d'activité,
- Rédaction du rapport scientifique de fin de projet,
- Rédaction d'une fiche projet, pour la publication d'un recueil regroupant les projets financés au titre de l'édition 2006 des programmes Blanc et JCJC de l'ANR.

A.7.2 Responsabilités d'unités d'enseignement

Dans le cadre de mes activités d'enseignement à Télécom ParisTech et à l'IRCAM, j'ai eu la responsabilité de six unités d'enseignement, en formation initiale comme en formation continue.

Parcours ATIAM à l'IRCAM :

Depuis 2007 : Traitement du Signal Audio (TSA)

Formation initiale à Télécom ParisTech :

Depuis 2007 : Méthodes d'Optimisation Continue et Applications (MOCA)

2007-2008 : Hilbert-Fourier (HF)

2006-2007 : Traitement du Signal : Méthodes et Études de Cas (TSMEC)

2004-2006 : Traitement du Signal : Études de Cas et Simulations (TSECS)

Formation continue à Télécom ParisTech :

Depuis 2010 : Traitement du Signal Audionumérique (TSA)

Le rôle de responsable d'unité d'enseignement inclut les responsabilités administratives suivantes :
- Organisation des réunions d'équipes pédagogiques,
- Organisation des contenus et publication des emplois du temps,
- Réservation des salles de cours ou de travaux pratiques,
- Organisation des examens écrits ou oraux,
- Organisation des corrections et déclaration des notes,
- Déclaration des activités d'enseignement pour tous les intervenants.

A.7.3 Activités d'intérêt collectif et valorisation de la recherche

Depuis novembre 2010 : Membre suppléant du Conseil d'École de Télécom ParisTech.

Juin 2010 : Diffusion en *Creative Commons* de la **base de données** MAPS [71].

Mars 2010 : Diffusion sous les termes de la GPL et **déclaration d'invention de la *DESAM Toolbox*** [73].

Jan. 2010 : Participation au **Grand Colloque STIC 2010** de l'ANR, à Paris La Villette. Ce colloque a réuni 780 responsables de projets, autour du thème "Quelle recherche pour les STIC de demain ?"

2009 : Contribution au **Cahier Numéro 3** de l'ANR, "Calcul haute performance : une technologie clé pour de multiples applications" (présentation du projet DESAM à l'attention d'un large public).

Mai 2008 : Participation au **Colloque JCJC** de l'ANR à Hyères. Ce colloque a réuni les responsables de projets ANR dans des domaines scientifiques très variés.

2007 : Création du **site Web du projet DESAM**, présentant la thématique du projet, les dernières nouvelles, les équipes participants, les publications et un intranet destiné aux partenaires du projet (`http://perso.telecom-paristech.fr/rbadeau/desam/`).

Jan. 2006 : Participation à STEP 2006 (**Séminaire Télécom Paris**), réunissant des responsables et des enseignants-chercheurs de Télécom ParisTech avec des industriels, autour du thème de l'innovation.

Oct. 2005 : Coordination de la journée Portes Ouvertes du département TSI de Télécom ParisTech sur le site du 37/39 rue Dareau (Paris 14), à l'attention des personnels de l'école.

2005 : Création d'une **page Web pour la Société Française d'Acoustique (SFA)**, présentant des techniques de séparation de la partie tonale et de la partie bruitée d'un son.

B. Bibliographie personnelle

Thèse de Doctorat

[1] Roland Badeau. « *Méthodes à haute résolution pour l'estimation et le suivi de sinusoïdes modulées. Application aux signaux de musique.* ». PhD thesis, École Nationale Supérieure des Télécommunications, ENST2005E007, Paris, France, avril 2005. Prix de thèse ParisTech 2006.

Articles de revues

[2] Roland Badeau, Nancy Bertin et Emmanuel Vincent. « Stability analysis of multiplicative update algorithms and application to non-negative matrix factorization ». *IEEE Transactions on Neural Networks*, 21(12) :1869–1881, décembre 2010.

[3] Roland Badeau et Rémy Boyer. « Fast Multilinear Singular Value Decomposition For Structured Tensors ». *SIAM Journal on Matrix Analysis and Applications*, 30(3) :1008–1021, septembre 2008.

[4] Roland Badeau, Bertrand David et Gaël Richard. « Fast Approximated Power Iteration Subspace Tracking ». *IEEE Transactions on Signal Processing*, 53(8) :2931–2941, août 2005.

[5] Roland Badeau, Bertrand David et Gaël Richard. « High resolution spectral analysis of mixtures of complex exponentials modulated by polynomials ». *IEEE Transactions on Signal Processing*, 54(4) :1341–1350, avril 2006.

[6] Roland Badeau, Bertrand David et Gaël Richard. « A new perturbation analysis for signal enumeration in rotational invariance techniques ». *IEEE Transactions on Signal Processing*, 54(2) :450–458, février 2006.

[7] Roland Badeau, Bertrand David et Gaël Richard. « Cramér-Rao bounds for multiple poles and coefficients of quasipolynomials in colored noise ». *IEEE Transactions on Signal Processing*, 56(8) :3458–3467, août 2008.

[8] Roland Badeau, Gaël Richard et Bertrand David. « Sliding window adaptive SVD algorithms ». *IEEE Transactions on Signal Processing*, 52(1) :1–10, janvier 2004.

[9] Roland Badeau, Gaël Richard et Bertrand David. « Fast and stable YAST algorithm for principal and minor subspace tracking ». *IEEE Transactions on Signal Processing*, 56(8) :3437–3446, août 2008.

[10] Roland Badeau, Gaël Richard et Bertrand David. « Performance of ESPRIT for estimating mixtures of complex exponentials modulated by polynomials ». *IEEE Transactions on Signal Processing*, 56(2) :492–504, février 2008.

[11] Nancy Bertin, Roland Badeau et Emmanuel Vincent. « Enforcing Harmonicity and Smoothness in Bayesian Non-negative Matrix Factorization Applied to Polyphonic Music Transcription ». *IEEE Transactions on Audio, Speech, and Language Processing*, 18(3) :538–549, mars 2010.

[12] Valentin Emiya, Roland Badeau et Bertrand David. « Multipitch estimation of piano sounds using a new probabilistic spectral smoothness principle ». *IEEE Transactions on Audio, Speech, and Language Processing*, 18(6) :1643–1654, août 2010.

[13] Romain Hennequin, Roland Badeau et Bertrand David. « NMF with time-frequency activations to model non-stationary audio events ». *IEEE Transactions on Audio, Speech, and Language Processing*, 2010. A paraître.

[14] Romain Hennequin, Bertrand David et Roland Badeau. « Beta-divergence as a subclass of Bregman divergence ». *IEEE Signal Processing Letters*, 18(2) :83–86, février 2010.

[15] Mathieu Lagrange, Martin Raspaud, Roland Badeau et Gaël Richard. « Explicit Modeling of Temporal Dynamics within Musical Signals for Acoustical Unit Formation and Similarity ». *Pattern Recognition Letters (PRNSA)*, 31(12) :1498–1506, septembre 2010.

[16] Jean-Loïc Le Carrou, François Gautier et Roland Badeau. « Sympathetic string modes in the concert harp ». *Acta Acustica united with Acustica*, 95(4) :744–752, juillet/août 2009.

[17] Antoine Liutkus, Roland Badeau et Gaël Richard. « Gaussian Processes for Underdetermined Source Separation ». *IEEE Transactions on Signal Processing*, 2010. Soumis pour publication.

[18] Emmanuel Vincent, Nancy Bertin et Roland Badeau. « Adaptive harmonic spectral decomposition for multiple pitch estimation ». *IEEE Transactions on Audio, Speech, and Language Processing*, 18(3) :528–537, mars 2010.

Articles de conférences

[19] Miguel Alonso Arevalo, Roland Badeau, Bertrand David et Gaël Richard. « Musical tempo estimation using noise subspace projections ». Dans *Proc. of IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pages 95–98, New Paltz, New York, USA, octobre 2003.

[20] Jon Altuna, Bernie Mulgrew, Roland Badeau et Vicente Atxa. « A Fast Adaptive Method for Subspace Based Blind Channel Estimation ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 4, pages 1121–1124, Toulouse, France, mai 2006.

[21] Roland Badeau, Karim Abed-Meraim, Gaël Richard et Bertrand David. « Sliding Window Orthonormal PAST Algorithm ». Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 5, pages 261–264, Hong Kong, Chine, avril 2003.

[22] Roland Badeau, Nancy Bertin et Emmanuel Vincent. « Stability analysis of multiplicative update algorithms for non-negative matrix factorization ». Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Prague, République Tchèque, mai 2011. Soumis pour publication.

[23] Roland Badeau, Rémy Boyer et Bertrand David. « EDS parametric modeling and tracking of audio signals ». Dans *Proc. of the 5th International Conference on Digital Audio Effects (DAFx)*, pages 139–144, Hambourg, Allemagne, septembre 2002.

[24] Roland Badeau et Bertrand David. « Weighted maximum likelihood autoregressive and moving average spectrum modeling ». Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 3761–3764, Las Vegas, Nevada, USA, avril 2008.

[25] Roland Badeau, Bertrand David et Gaël Richard. « Selecting the modeling order for the ESPRIT high resolution method : an alternative approach ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 2, pages 1025–1028, Montréal, Québec, Canada, mai 2004.

[26] Roland Badeau, Bertrand David et Gaël Richard. « Yet Another Subspace Tracker ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 4, pages 329–332, Philadelphie, Pennsylvanie, USA, mars 2005.

[27] Roland Badeau, Bertrand David et Gaël Richard. « YAST Algorithm for Minor Subspace Tracking ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 3, pages 552–555, Toulouse, France, mai 2006.

[28] Roland Badeau, Bertrand David et Gaël Richard. « Conjugate gradient algorithms for minor subspace analysis ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 3, pages 1013–1016, Honolulu, Hawaii, USA, avril 2007.

[29] Roland Badeau, Valentin Emiya et Bertrand David. « Expectation-maximization algorithm for multipitch estimation and separation of overlapping harmonic spectra ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3073–3076, Taipei, Taiwan, avril 2009.

[30] Roland Badeau, Gaël Richard et Bertrand David. « Adaptive ESPRIT algorithm based on the PAST subspace tracker ». Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 6, pages 229–232, Hong Kong, Chine, avril 2003.

[31] Roland Badeau, Gaël Richard et Bertrand David. « Suivi d'espace dominant par la méthode des puissances itérées ». Dans *Actes du colloque GRETSI*, volume 1, pages 137–140, Paris, France, septembre 2003.

[32] Roland Badeau, Gaël Richard et Bertrand David. « Fast adaptive ESPRIT algorithm ». Dans *Proc. of IEEE Workshop on Statistical Signal Processing (SSP)*, Bordeaux, France, juillet 2005.

[33] Roland Badeau, Gaël Richard, Bertrand David et Karim Abed-Meraim. « Approximated power iterations for fast subspace tracking ». Dans *Proc. of the 7th International Symposium on Signal Processing and its Applications (ISSPA)*, volume 2, pages 583–586, Paris, France, juillet 2003.

[34] Nancy Bertin, Roland Badeau et Gaël Richard. « Blind signal decompositions for automatic transcription of polyphonic music : NMF and K-SVD on the benchmark ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 1, pages 65–68, Honolulu, Hawaii, USA, avril 2007.

[35] Nancy Bertin, Roland Badeau et Emmanuel Vincent. « Fast Bayesian NMF algorithms enforcing harmonicity and temporal continuity in polyphonic music transcription ». Dans *Proc. of IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pages 29–32, New Paltz, New York, USA, octobre 2009.

[36] Nancy Bertin, Cédric Févotte et Roland Badeau. « A tempering approach for Itakura-Saito non-negative matrix factorization. With application to music transcription ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1545–1548, Taipei, Taiwan, avril 2009.

[37] Rémy Boyer et Roland Badeau. « Adaptive Multilinear SVD for Structured Tensors ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 3, pages 880–883, Toulouse, France, mai 2006.

[38] Rémy Boyer, Roland Badeau et Gérard Favier. « Fast orthogonal decomposition of Volterra cubic kernels using oblique unfolding ». Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Prague, République Tchèque, mai 2011. Soumis pour publication.

[39] Bertrand David et Roland Badeau. « Fast sequential LS estimation for sinusoidal modeling and decomposition of audio signals ». Dans *Proc. of IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pages 211–214, New Paltz, New York, USA, octobre 2007.

[40] Bertrand David, Roland Badeau et Gaël Richard. « Sintrack analysis for tracking components of musical signals ». Dans *Proc. of the Forum Acusticum*, Séville, Espagne, septembre 2002.

[41] Bertrand David, Roland Badeau et Gaël Richard. « HRHATRAC Algorithm for Spectral Line Tracking of Musical Signals ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 3, pages 45–48, Toulouse, France, mai 2006.

[42] Bertrand David, Valentin Emiya, Roland Badeau et Yves Grenier. « Harmonic Plus Noise Decomposition : Time-Frequency Reassignment Versus a Subspace-Based Method ». Dans *Proc. of 120th Convention of the Audio Engineering Society (AES)*, Paris, France, mai 2006.

[43] Bertrand David, Gaël Richard et Roland Badeau. « An EDS modelling tool for tracking and modifying musical signals ». Dans *Proc. of Stockholm Music Acoustics Conference (SMAC)*, volume 2, pages 715–718, Stockholm, Suède, août 2003.

[44] Olivier Derrien, Roland Badeau et Gaël Richard. « Entropy-constrained quantization of exponentially damped sinusoids parameters ». Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Prague, République Tchèque, mai 2011. Soumis pour publication.

[45] Valentin Emiya, Roland Badeau et Bertrand David. « Multipitch estimation of quasi-harmonic sounds in colored noise ». Dans *Proc. of the 10th International Conference on Digital Audio Effects (DAFx)*, pages 93–98, Bordeaux, France, septembre 2007.

[46] Valentin Emiya, Roland Badeau et Bertrand David. « Automatic transcription of piano music based on HMM tracking of jointly-estimated pitches ». Dans *Proc. of European Signal Processing Conference (EUSIPCO)*, Lausanne, Suisse, août 2008.

[47] Valentin Emiya, Bertrand David et Roland Badeau. « A parametric method for pitch estimation of piano tones ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 1, pages 249–252, Honolulu, Hawaii, USA, avril 2007.

[48] Benoit Fuentes, Roland Badeau et Gaël Richard. « Adaptive harmonic decomposition using shift-invariant PLCA ». Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Prague, République Tchèque, mai 2011. Soumis pour publication.

[49] Romain Hennequin, Roland Badeau et Bertrand David. « NMF with time-frequency activations to model non-stationary audio events ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 445–448, Dallas, Texas, USA, mars 2010.

[50] Romain Hennequin, Roland Badeau et Bertrand David. « Spectral similarity measure invariant to pitch shifting and amplitude scaling ». Dans *Actes du 10ème Congrès Français d'Acoustique (CFA)*, Lyon, France, avril 2010.

[51] Romain Hennequin, Roland Badeau et Bertrand David. « Time-dependent parametric and harmonic templates in non-negative matrix factorization ». Dans *Proc. of the 13th International Conference on Digital Audio Effects (DAFx)*, Graz, Autriche, septembre 2010.

[52] Romain Hennequin, Bertrand David et Roland Badeau. « Score informed audio source separation using a parametric model of non-negative spectrogram ». Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Prague, République Tchèque, mai 2011. Soumis pour publication.

[53] Mathieu Lagrange, Roland Badeau, Bertrand David, Nancy Bertin, José Echeveste, Olivier Derrien, Sylvain Marchand et Laurent Daudet. « The DESAM Toolbox : spectral analysis of musical audio ». Dans *Proc. of the 13th International Conference on Digital Audio Effects (DAFx)*, Graz, Autriche, septembre 2010.

[54] Mathieu Lagrange, Roland Badeau et Gaël Richard. « Robust similarity metrics between audio signals based on asymmetrical spectral envelope matching ». Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 405–408, Dallas, Texas, USA, mars 2010.

[55] Jean-Loïc Le Carrou, François Gautier et Roland Badeau. « Analyse des modes de cordes couplées d'une harpe par une méthode à haute résolution ». Dans *Actes du 8ème Congrès Français d'Acoustique (CFA)*, Tours, France, avril 2006.

[56] Jean-Loïc Le Carrou, François Gautier et Roland Badeau. « Theoretical and experimental investigations of harp's sympathetic modes ». Dans *Proc. of International Congress on Acoustics (ICA)*, Madrid, Espagne, septembre 2007.

[57] Antoine Liutkus, Roland Badeau et Gaël Richard. « Informed Source Separation Using Latent Components ». Dans Vincent Vigneron, Vicente Zarzoso, Eric Moreau, Rémi Gribonval et Emmanuel Vincent, éditeurs, *Latent Variable Analysis and Signal Separation*, volume 6365 de *Lecture Notes in Computer Science*, pages 498–505, Saint Malo, France, septembre 2010. Springer.

[58] Antoine Liutkus, Roland Badeau et Gaël Richard. « Explaining IS-NMF as a particular case of Gaussian process regression ». Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Prague, République Tchèque, mai 2011. Soumis pour publication.

[59] Jean-Marie Nicolas, Alain Maruani et Roland Badeau. « Les moments d'ordre inférieur : Principes et application au filtrage des images RSO ». Dans *Actes du 12ème congrès francophone de Reconnaissance des Formes et Intelligence Artificielle (RFIA)*, volume 3, pages 27–36, Paris, France, février 2000.

[60] Emmanuel Vincent, Nancy Bertin et Roland Badeau. « Harmonic and inharmonic nonnegative matrix factorization for polyphonic pitch transcription ». Dans *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 109–112, Las Vegas, Nevada, USA, avril 2008.

Rapports techniques

[61] Roland Badeau. « Transformation de Mellin : aspects numériques et dérivation fractionnaire. Application aux lois des variables aléatoires positives. ». Rapport Technique, École Nationale Supérieure des Télécommunications, Paris, France, novembre 1999.

[62] Roland Badeau. « *Technical Report on Audio and Speech Processing* », Chapitre "High Resolution Methods". Groupe des Écoles des Télécommunications (GET), Paris, France, 2006. Livrable du réseau d'excellence européen K-Space.

[63] Roland Badeau, Nancy Bertin et Emmanuel Vincent. « Stability analysis of multiplicative update algorithms and application to non-negative matrix factorization ». Rapport Technique 2010D018, Télécom ParisTech, Paris, France, septembre 2010.

[64] Roland Badeau, Nancy Bertin et Emmanuel Vincent. « Supporting document for the paper "Stability analysis of multiplicative update algorithms and application to non-negative matrix factorization" ». Rapport Technique 2010D019, Télécom ParisTech, Paris, France, septembre 2010.

[65] Roland Badeau, Gaël Richard et Bertrand David. « Performance of ESPRIT for Estimating Mixtures of Complex Exponentials Modulated by Polynomials : Supporting Document ». Rapport Technique 2007D015, École Nationale Supérieure des Télécommunications, Paris, France, octobre 2007.

[66] Nancy Bertin, Roland Badeau et Emmanuel Vincent. « Enforcing harmonicity and smoothness in Bayesian non-negative matrix factorization applied to polyphonic music transcription ». Rapport Technique, École Nationale Supérieure des Télécommunications, 2009D006, janvier 2009.

[67] Olivier Derrien, Roland Badeau et Gaël Richard. « Calculation of an entropy-constrained quantizer for exponentially damped sinudoids parameters ». Rapport Technique, Laboratoire de Mécanique et d'Acoustique, Marseille, France, juin 2010.

[68] Valentin Emiya, Nancy Bertin, Bertrand David et Roland Badeau. « MAPS - A piano database for multipitch estimation and automatic transcription of music ». Rapport Technique 2010D017, Télécom ParisTech, Paris, France, juillet 2010.

[69] Emmanuel Vincent, Nancy Bertin et Roland Badeau. « Adaptive harmonic spectral decomposition for multiple pitch estimation ». Rapport Technique, IRISA, PI-1919, Rennes, France, janvier 2009.

Brevets, logiciel et base de données

[70] Roland Badeau, Gaël Richard et Bertrand David. « Procédé de poursuite d'un sous-espace de dimension inférieure à celle des vecteurs de données, notamment audio ». Brevet no. 05 50678, mars 2005. Télécom ParisTech, Paris, France.

[71] Valentin Emiya, Nancy Bertin, Bertrand David et Roland Badeau. « MIDI Aligned Piano Sounds (MAPS) », 2010. Télécom ParisTech, Paris, France. Diffusion en *Creative Commons*.

[72] Laurent Girin, Antoine Liutkus, Gaël Richard et Roland Badeau. « Procédé et dispositif de formation d'un signal mixé numérique audio, procédé et dispositif de séparation de signaux, et signal correspondant ». Demande de brevet no. B10/3035FR / GBO, octobre 2010. Institut Polytechnique de Grenoble et Institut Télécom, Télécom ParisTech.

[73] Mathieu Lagrange, Roland Badeau, Bertrand David, Nancy Bertin, José Echeveste, Olivier Derrien, Sylvain Marchand et Laurent Daudet. « DESAM Toolbox », 2010. Télécom ParisTech, Paris, France. Diffusion en *GNU General Public License* (GPL).

Supports de cours

[74] Roland Badeau. « *Musique et Signaux Audiofréquence (PAMU)* », Chapitre "Détection de hauteur". École Nationale Supérieure des Télécommunications, Paris, France, 2002.

[75] Roland Badeau. « *Séries Chronologiques (TSMAp)* », Chapitre "Estimation de sinusoïdes". École Nationale Supérieure des Télécommunications, Paris, France, 2004.

[76] Roland Badeau. « *Élaboration et transformations de sons (ATIAM)* », Chapitre "Méthodes à haute résolution". Institut de Recherche et Coordination Acoustique/Musique, Paris, France, 2006.

[77] Roland Badeau. « *Hilbert - Fourier (MDI)* », Chapitre "Transformées de Fourier à Temps Discret et à Temps Continu". École Nationale Supérieure des Télécommunications, Paris, France, 2008.

Conférences sans actes

[78] Roland Badeau et Bertrand David. « Adaptive subspace methods for high resolution analysis of music signals ». Dans *Acoustics'08*, Paris, France, juillet 2008.

[79] Nancy Bertin et Roland Badeau. « Initialization, distances and local minima in audio applications of the non-negative matrix factorization ». Dans *Acoustics'08*, Paris, France, juillet 2008.

[80] Bertrand David et Roland Badeau. « Towards an adaptive subspace-based representation of musical spectral content ». Dans *Acoustics'08*, Paris, France, juillet 2008.

[81] Bertrand David, Roland Badeau, Nancy Bertin, Valentin Emiya et Gaël Richard. « Multipitch detection for piano music : Benchmarking a few approaches ». *The Journal of the Acoustical Society of America (JASA)*, 122(5) :2962, novembre 2007.

[82] Olivier Derrien, Gaël Richard et Roland Badeau. « Damped sinusoids and subspace based approach for lossy audio coding ». Dans *Acoustics'08*, Paris, France, juillet 2008.

Travaux encadrés

[83] Nancy Bertin. « *Les factorisations en matrices non-négatives. Approches contraintes et probabilistes, application à la transcription automatique de musique polyphonique.* ». PhD thesis, École Nationale Supérieure des Télécommunications, Paris, France, octobre 2009. Prix de thèse 2010 du GdR-Isis/EEA/GRETSI.

[84] Valentin Emiya. « *Transcription automatique de la musique de piano* ». PhD thesis, École Nationale Supérieure des Télécommunications, Paris, France, octobre 2008.

[85] Sylvain Rousselle. « Séparation de la Voix du Locuteur et du Fond Musical dans des Emissions Radiodiffusées ». Master's thesis, École Nationale Supérieure des Télécommunications, Paris, France, juillet 2006. Délivrable du projet Infomagic, du pôle de compétitivité Cap Digital IMVN (Image, Multimédia et Vie Numérique).

C. Copies des dix publications les plus représentatives

Analyse spectrale à haute résolution

[Badeau *et al.*, 2008a] Roland Badeau, Gaël Richard, et Bertrand David. Performance of ESPRIT for estimating mixtures of complex exponentials modulated by polynomials. *IEEE Transactions on Signal Processing*, 56(2) :492–504, février 2008.

[Badeau *et al.*, 2006] Roland Badeau, Bertrand David, et Gaël Richard. A new perturbation analysis for signal enumeration in rotational invariance techniques. *IEEE Transactions on Signal Processing*, 54(2) :450–458, février 2006.

[Badeau *et al.*, 2005a] Roland Badeau, Bertrand David, et Gaël Richard. Fast approximated power iteration subspace tracking. *IEEE Transactions on Signal Processing*, 53(8) :2931–2941, août 2005.

[Badeau *et al.*, 2008b] Roland Badeau, Gaël Richard, et Bertrand David. Fast and stable YAST algorithm for principal and minor subspace tracking. *IEEE Transactions on Signal Processing*, 56(8) :3437–3446, août 2008.

[Badeau *et al.*, 2005b] Roland Badeau, Gaël Richard, et Bertrand David. Fast adaptive ESPRIT algorithm. Dans *Proc. of IEEE Workshop on Statistical Signal Processing (SSP)*, Bordeaux, France, juillet 2005.

Décompositions non négatives

[Emiya *et al.*, 2010] Valentin Emiya, Roland Badeau, et Bertrand David. Multipitch estimation of piano sounds using a new probabilistic spectral smoothness principle. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(6) :1643–1654, août 2010.

[Badeau *et al.*, 2009] Roland Badeau, Valentin Emiya, et Bertrand David. Expectation-maximization algorithm for multi-pitch estimation and separation of overlapping harmonic spectra. Dans *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3073–3076, Taipei, Taiwan, avril 2009.

[Badeau *et al.*, 2010] Roland Badeau, Nancy Bertin, et Emmanuel Vincent. Stability analysis of multiplicative update algorithms and application to non-negative matrix factorization. *IEEE Transactions on Neural Networks*, 21(12) :1869–1881, décembre 2010.

[Bertin *et al.*, 2010] Nancy Bertin, Roland Badeau, et Emmanuel Vincent. Enforcing harmonicity and smoothness in Bayesian non-negative matrix factorization applied to polyphonic music transcription. *IEEE Transactions on Audio, Speech and Language Processing*, 18(3) :538–549, mars 2010.

[Hennequin *et al.*, 2010] Romain Hennequin, Roland Badeau, et Bertrand David. NMF with time-frequency activations to model non-stationary audio events. *IEEE Transactions on Audio, Speech and Language Processing*, 2010. A paraître.

Performance of ESPRIT for estimating mixtures of complex exponentials modulated by polynomials

Roland Badeau, *Member, IEEE*, Gaël Richard, *Senior Member, IEEE*, and Bertrand David, *Member, IEEE*

***Abstract*—High Resolution (HR) methods are known to provide accurate frequency estimates for discrete spectra. The Polynomial Amplitude Complex Exponentials (PACE) model, also called quasipolynomial model in the literature, was presented as the most general model tractable by HR methods. A subspace-based estimation scheme was recently proposed, derived from the classical ESPRIT algorithm. In this paper, we focus on the performance of this estimator. We first present some asymptotic expansions of the estimated parameters, obtained at the first order under the assumption of a high signal-to-noise ratio. Then the performance of the generalized ESPRIT algorithm for estimating the parameters of this model is analyzed in terms of bias and variance, and compared to the Cramér-Rao bounds. This performance is studied in an asymptotic context, and it is proved that the efficiency of undamped single poles estimators is close to the optimality. Moreover, our results show that the best performance is obtained for a proper dimensioning of the data. To illustrate the practical capabilities of the generalized ESPRIT algorithm, we finally propose an application to ARMA filter synthesis, in the context of system conversion from continuous time to discrete time.**

***Index Terms*—performance analysis, perturbation theory, high resolution, ESPRIT, polynomial modulation, multiple eigenvalues.**

I. Introduction

HIGH Resolution (HR) methods, such as the well-known ESPRIT algorithm [1], are very classical techniques for estimating mixtures of complex exponentials in white noise. However the underlying Exponential Sinusoidal Model (ESM), although the most studied in the literature, is not the most general model tractable by HR methods. Indeed the ESM only accounts for systems with single poles, whereas one can find examples of systems involving multiple poles, which generate mixtures of complex exponentials modulated by polynomials, as shown in [2]. For instance, critically damped harmonic oscillators involve a double pole [3]. Laguerre functions are a special case of signals with multiple poles (the exponentials are modulated by Laguerre polynomials), often used in the estimation of time delays [4], [5], and in biomedical engineering, for modeling fluorescence decay [6]. Signals with multiple poles also appear in quantum physics, as solutions of the Schrödinger equation for hydrogen-like atoms [7], in laser physics, as transverse laser modes [8], and in finance, for modeling the evolution of interest rates [9].

In order to estimate the parameters of this more general Polynomial Amplitude Complex Exponentials (PACE) model, also referred to as the quasipolynomial model [10] in the literature, we proposed in [2] a generalization of the ESPRIT algorithm for estimating multiple poles. The polynomial amplitude parameters can then be recovered by means of a Least Squares (LS) method. The performance of an estimator is generally described in terms of bias and variance, the latter being generally compared to the Cramér-Rao bound (CRB), in terms of statistical efficiency [11]. An analysis of the Cramér-Rao bounds for the frequencies and damping factors of complex quasipolynomials in white noise was proposed in [10]. In [12], we derive analytic expressions of the Cramér-Rao bounds for the frequencies, damping factors, amplitudes and phases of quasipolynomials in colored noise, and these expressions are simplified in an asymptotic context. In particular, it is shown that the CRB for the parameters associated to a multiple pole present an exponential increase with the order of the pole, which suggests that the practical estimation of the PACE model is only possible if the exponentials are modulated by polynomials of low order.

Unfortunately, in the case of HR methods, the bias and variance cannot be calculated analytically, because the extraction of polynomial roots, or matrix eigenvalues, induces a complex relationship between the statistics of the signal and those of the estimators. In the case of the ESM however, asymptotic results were obtained with the perturbation theory. These results rely either on the hypothesis of a high window length ($N \rightarrow +\infty$, in the case where all the poles are on the unit circle), or on the hypothesis of a high signal-to-noise ratio (SNR$\rightarrow +\infty$). For instance, it was established in [13], [14] that the Prony [15] and Pisarenko [16] methods are very inefficient: their variances are much greater than the Cramér-Rao bounds. Conversely, the Minimum Norm method [17], MUSIC [18], ESPRIT [1] and Matrix Pencil [19] have an asymptotic efficiency close to 1, as shown in [20]–[24]. More precisely, it was proved in [22], [23] in the case of undamped sinusoids that MUSIC and ESPRIT perform similarly, with a slight advantage for ESPRIT. This was confirmed in [24] in the more general case of the ESM: ESPRIT and Matrix Pencil are less sensitive to noise than MUSIC.

In the case of the PACE model, it was shown in [2] that the presence of noise scatters the multiple poles into several single poles, forming the vertices of a regular polygon as a first order approximation. However the original multiple pole can be recovered by computing the arithmetic mean of the scattered poles. Below, we analyze the performance of this approach in presence of *colored* noise and under the high SNR hypothesis, in terms of first order perturbations. These developments are utilized to show that the estimators proposed in [2] for the PACE model are unbiased, and their variances are calculated and compared to the Cramér-Rao bounds. Additionally, by

Roland Badeau, Gaël Richard and Bertrand David are with the Department of Signal and Image Processing, École Nationale Supérieure des Télécommunications (ENST), Paris, France. E-mail: [roland.badeau, bertrand.david, gael.richard]@enst.fr.

considering a high observation length and a white noise, it is shown that the efficiency of the estimators is close to 1. In particular, we generalize a result presented in [19], which provides the ideal dimensioning of the data matrix in order to improve the efficiency of single poles estimators. However, our simulation results confirm that the practical estimation of the parameters is only possible for poles of low order. To illustrate the capabilities of our estimation method, we finally propose an application to ARMA filter synthesis, in the context of system conversion from continuous time to discrete time.

The paper is organized as follows. Section II describes the PACE model and the estimation method introduced in [2]. Then the influence of an additive perturbation onto the estimated frequencies, damping factors, amplitudes and phases is studied in section III. Section IV analyzes the performance of the estimators: their first order bias and variances are calculated in section IV-A, then their asymptotic expansions are derived in section IV-B (in the case of undamped single poles). These developments are illustrated in section V-A, where the variances of the estimators are compared to the Cramér-Rao bounds, and the generalized ESPRIT algorithm is applied to ARMA filter synthesis in section V-B. The main conclusions of this paper are summarized in section VI. Finally, the mathematical developments for the perturbation analysis are provided in the appendix.

II. General framework

In sections II-A and II-B, we summarize the basics of the PACE model, also called quasipolynomial model, and the generalized ESPRIT algorithm, which were presented in [2].

A. Polynomial Amplitude Complex Exponentials

Definition 1. Let $K \in \mathbb{N}^*$. For all $k \in \{0 \dots K-1\}$, define the partial order $M_k \in \mathbb{N}^*$, the frequency $f_k \in \left]-\frac{1}{2}, \frac{1}{2}\right]$, the damping (or amplifying) factor $\delta_k \in \mathbb{R}$, and the complex pole $z_k = e^{\delta_k + i2\pi f_k}$. Suppose that the complex poles are distinct from one another. Then a discrete signal $s(t)$ satisfies the PACE model of order $r \triangleq \sum_{k=0}^{K-1} M_k$ if and only if it can be written in the form

$$\boxed{s(t) = \sum_{k=0}^{K-1} \alpha_k[t]\, {z_k}^t} \tag{1}$$

where $\forall k \in \{0, \dots, K-1\}$, $\alpha_k[t]$ is a complex polynomial of order $M_k - 1$.

The polynomial $\alpha_k[t]$ can be decomposed onto the polynomial basis of falling factorials:

Definition 2 (Falling factorial)**.** For all $m \in \mathbb{Z}$, the falling factorial of order m is the polynomial[1]

$$F_m[X] = \begin{cases} 0 & \text{if } m < 0 \\ 1 & \text{if } m = 0 \\ \frac{1}{m!} \prod_{m'=0}^{m-1} (X - m') & \text{if } m > 0. \end{cases}$$

[1]Note that this definition does not match the classical definition of the falling factorial [25], [26], from which the multiplicative factor $\frac{1}{m!}$ is missing.

In this basis, equation (1) can be rewritten in the form

$$s(t) = \sum_{k=0}^{K-1} \sum_{m=0}^{M_k-1} \alpha_k^{(m)} F_m[t]\, {z_k}^{t-m} \tag{2}$$

where $\forall k \in \{0 \dots K-1\}$, $\forall m \in \{0, M_k - 1\}$, $\alpha_k^{(m)}$ is a complex amplitude. Define the real amplitude $a_k^{(m)} = \left|\alpha_k^{(m)}\right|$, and the phase[2] $\phi_k^{(m)} = \Im\left(\ln\left(\alpha_k^{(m)}\right)\right)$.

The PACE model can be characterized by means of matrix analysis. Indeed, the samples of the discrete signal $s(t)$ can be arranged into a Hankel data matrix with $n > r$ rows and $l \geq r$ columns:

$$\boldsymbol{S} = \begin{bmatrix} s(t-l+1) & \cdots & s(t-1) & s(t) \\ s(t-l+2) & \cdots & s(t) & s(t+1) \\ \vdots & \cdots & \vdots & \vdots \\ s(t-l+n) & \cdots & s(t+n-2) & s(t+n-1) \end{bmatrix}. \tag{3}$$

In particular, the range space of $\boldsymbol{S}$ can be characterized by the *generalized Pascal* and *Pascal-Vandermonde* matrices.

Definition 3 (Generalized Pascal matrices)**.** Let $z \in \mathbb{C}$ and $M \in \mathbb{N}^*$. The generalized Pascal matrix denoted $\boldsymbol{C}_M^n(z)$ is an $n \times M$ matrix whose coefficients are $\boldsymbol{C}_M^n(z)_{(i,j)} = F_j[i]\, z^{i-j}$ for all $i \in \{0 \dots n-1\}$ and $j \in \{0 \dots M-1\}$.

Example 4. If $M = 3$ and $n = 6$,

$$\boldsymbol{C}_3^6(z) = \begin{bmatrix} 1 & 0 & 0 \\ z & 1 & 0 \\ z^2 & 2\,z & 1 \\ z^3 & 3\,z^2 & 3\,z \\ z^4 & 4\,z^3 & 6\,z^2 \\ z^5 & 5\,z^4 & 10\,z^3 \end{bmatrix}.$$

Definition 5 (Pascal-Vandermonde matrices)**.** The $n \times r$ Pascal-Vandermonde matrix is the matrix formed by concatenating the generalized Pascal matrices $\boldsymbol{C}_k^n = \boldsymbol{C}_{M_k}^n(z_k)$:

$$\boldsymbol{V}^n = \left[\boldsymbol{C}_0^n, \dots, \boldsymbol{C}_{(K-1)}^n\right].$$

Based on the above definitions, the following proposition, which is proved in [2], shows a factorization of the Hankel data matrix[3]:

Proposition 1 (Factorization of the data matrix)**.** *An $n \times l$ Hankel matrix $\boldsymbol{S}$ of the form* (3) *where $s(t)$ is the signal defined in equation* (2) *can be factorized in the form*

$$\boldsymbol{S} = \boldsymbol{V}^n\, \boldsymbol{D}\, {\boldsymbol{V}^l}^T \tag{4}$$

where $\boldsymbol{V}^n$ and $\boldsymbol{V}^l$ are the $n \times r$ and $l \times r$ Pascal-Vandermonde matrices, and $\boldsymbol{D}$ is an $r \times r$ block-diagonal matrix

$$\boldsymbol{D} = \operatorname{diag}(\boldsymbol{H}_0 \dots \boldsymbol{H}_{(K-1)})$$

whose k^{th} block $\boldsymbol{H}_k$ is an $M_k \times M_k$ upper anti-triangular Hankel matrix (in the particular case $M_k = 1$, $\boldsymbol{H}_k = z_k^{t-l+1}\alpha_k^{(0)}$).

[2]In the whole paper, the notation $\ln(.)$ denotes the determination of the complex logarithm which corresponds to an angle lying in $]-\pi, \pi[$.

[3]In linear systems realization theory, state space representations also lead to low-rank factorizations of Hankel matrices [27].

B. Estimation of the model parameters

Proposition 1 shows that the matrix $\boldsymbol{S}$ has rank r, and that its range space, called *signal subspace*, is also spanned by the Pascal-Vandermonde matrix $\boldsymbol{V}^n$.

1) Rotational Invariance Property: The ESPRIT method relies on a particular property of Vandermonde matrices known as *rotational invariance* [1], which reflects the invariance of the signal subspace to time shifts. Theorem 2, shown in [2], generalizes this property to Pascal-Vandermonde matrices.

Theorem 2 (Rotational Invariance Property of Pascal-Vandermonde matrices)**.** *Let $\boldsymbol{V}^n_{\downarrow}$ be the matrix extracted from $\boldsymbol{V}^n$ by deleting the last row. Similarly, let $\boldsymbol{V}^n_{\uparrow}$ be the matrix extracted from $\boldsymbol{V}^n$ by deleting the first row. Then $\boldsymbol{V}^n_{\downarrow}$ and $\boldsymbol{V}^n_{\uparrow}$ span the same subspace, and*

$$\boldsymbol{V}^n_{\uparrow} = \boldsymbol{V}^n_{\downarrow}\, \boldsymbol{J} \tag{5}$$

where $\boldsymbol{J}$ is the $r \times r$ block-diagonal matrix

$$\boldsymbol{J} = \operatorname{diag}\left(\boldsymbol{J}_0, \ldots, \boldsymbol{J}_{(K-1)}\right) \tag{6}$$

whose k^{th} block $\boldsymbol{J}_k$ is the $M_k \times M_k$ Jordan block whose diagonal coefficients are equal to z_k.

The interesting fact in theorem 2 is that equation (5) involves a Jordan matrix[4] $\boldsymbol{J}$, which characterizes the poles z_k and their multiplicity M_k. As shown below, the generalized ESPRIT algorithm consists in computing $\boldsymbol{J}$ as a by-product of the Jordan canonical decomposition of a so-called *spectral matrix*.

2) The generalized ESPRIT method: In practice, the Pascal-Vandermonde matrix $\boldsymbol{V}^n$ is unknown. Nevertheless, it was shown in [2] that in presence of white noise an $n \times r$ orthonormal matrix $\boldsymbol{W}$ spanning the signal subspace can be estimated by computing the left dominant r-dimensional singular subspace of the noisy data matrix, or by using subspace tracking methods [29]–[31]. Since $\boldsymbol{W}$ and $\boldsymbol{V}^n$ span the same subspace, there is an $r \times r$ invertible matrix $\boldsymbol{G}$ such that

$$\boldsymbol{V}^n = \boldsymbol{W}\, \boldsymbol{G}. \tag{7}$$

Substituting equation (7) into equation (5) shows that $\boldsymbol{W}$ satisfies an equation similar to equation (5): $\boldsymbol{W}_{\uparrow} = \boldsymbol{W}_{\downarrow}\, \boldsymbol{\Phi}$ where $\boldsymbol{\Phi}$, herein called the *spectral matrix*, is defined by its Jordan canonical decomposition:

$$\boldsymbol{\Phi} \triangleq \boldsymbol{G}\, \boldsymbol{J}\, \boldsymbol{G}^{-1}. \tag{8}$$

It can be noticed that $\forall t \in \mathbb{Z}$, the spectral matrix $\boldsymbol{\Phi}$, which depends on the observation window $\{t-l+1, \ldots, t+n-1\}$, is similar to the unique Jordan matrix $\boldsymbol{J}$. Finally, the generalized ESPRIT algorithm consists in:

- estimating a basis $\boldsymbol{W}$ of the signal subspace, via singular value decomposition for instance[5],
- computing the spectral matrix[6] $\boldsymbol{\Phi} = \boldsymbol{W}^{\dagger}_{\downarrow}\boldsymbol{W}_{\uparrow}$.
- computing the eigenvalues of $\boldsymbol{\Phi}$ from which the estimated poles and their multiplicities can be extracted.

Note that in a noisy context, the estimated spectral matrix does not have multiple eigenvalues in practice, and the generalized ESPRIT algorithm cannot be applied as it is. This problem will be discussed in section A.

3) Estimation of the complex amplitudes: The complex amplitudes are estimated by means of the LS method. Let $\widehat{\boldsymbol{V}}^N$ be the $N \times r$ Pascal-Vandermonde matrix defined from the estimated poles, and $\widehat{\boldsymbol{s}}$ the N-dimensional vector containing the successive samples of the observed signal. Then the LS-estimate of the vector $\boldsymbol{\alpha} = \left[\boldsymbol{\alpha}_0, \ldots, \boldsymbol{\alpha}_{K-1}\right]^T$ (with $\boldsymbol{\alpha}_k = \left[\alpha_k^{(0)}, \ldots, \alpha_k^{(M_k-1)}\right]^T$) containing the complex amplitudes is

$$\widehat{\boldsymbol{\alpha}} = \widehat{\boldsymbol{V}}^{N\dagger}\widehat{\boldsymbol{s}}. \tag{9}$$

III. Perturbation analysis

The objective here is to measure the performance of the estimators presented above in terms of bias and dispersion. Unfortunately, it is not possible to establish analytic formulae in the general case, because of the eigen or singular value decompositions. However, asymptotic results could be obtained by using the perturbation theory in the case of the sinusoidal model [23] and in the case of the ESM model [24], under the hypothesis of a high signal-to-noise ratio. We propose below to apply the perturbation theory in the more general framework of the PACE model, in order to finally derive the first and second moments of the estimators. First, we analyze the perturbation induced onto the frequencies and damping factors, from which we derive the perturbation induced onto the amplitudes and phases. The detailed mathematical developments can be found in the appendix.

Suppose that the PACE signal $s(t)$ is corrupted by a perturbation $\varepsilon\Delta s(t)$ (where $\varepsilon << 1$), so that the observed signal is $s(t) + \varepsilon\Delta s(t)$. In other terms, the $n \times l$ Hankel data matrix $\boldsymbol{S}$ is corrupted by an additive perturbation $\varepsilon\Delta\boldsymbol{S}$, where $\Delta\boldsymbol{S}$ is the $n \times l$ Hankel matrix containing the samples of $\Delta s(t)$, so that the observed matrix is

$$\boldsymbol{S}(\varepsilon) = \boldsymbol{S} + \varepsilon\Delta\boldsymbol{S}. \tag{10}$$

Then suppose that the generalized ESPRIT algorithm is applied to the perturbed matrix $\boldsymbol{S}(\varepsilon)$ instead of the exact matrix $\boldsymbol{S}$. In section A in the appendix, it is shown that the perturbed subspace weighting matrix $\boldsymbol{W}(\varepsilon)$ (section A1), spectral matrix $\boldsymbol{\Phi}(\varepsilon)$ (section A2), poles $z_k(\varepsilon)$, frequencies $f_k(\varepsilon)$ and damping factors $\delta_k(\varepsilon)$ (section A3) are $\mathcal{C}^\infty$ functions of ε in the neighborhood of $\varepsilon = 0$, leading to the first order expansions in proposition 3.

Proposition 3 (Perturbation of the frequencies and damping factors)**.** *Let*[7] $\delta_k(\varepsilon) = \Re\left(\ln(z_k(\varepsilon))\right)$ *and* $f_k(\varepsilon) = \frac{1}{2\pi}\Im\left(\ln(z_k(\varepsilon))\right)$*. Then the functions $\varepsilon \mapsto \delta_k(\varepsilon)$ and $\varepsilon \mapsto f_k(\varepsilon)$ are $\mathcal{C}^\infty$ and admit the first order expansions:*

$$\left\{\begin{array}{rcl} \delta_k(\varepsilon) & = & \delta_k + \varepsilon\,\Delta\delta_k + O\left(\varepsilon^2\right) \\ f_k(\varepsilon) & = & f_k + \varepsilon\,\Delta f_k + O\left(\varepsilon^2\right). \end{array}\right. \tag{11}$$

[4] See [28, pp. 121–142] for a definition of Jordan canonical decomposition.

[5] In linear systems realization theory, Ho's algorithm is a well-known method for identifying a state space representation [27]. The use the singular value decomposition in this context was early proposed in [32].

[6] In the whole paper, the symbol † denotes the Moore-Penrose pseudo-inverse.

[7] It is supposed here that all frequencies lie between $-\frac{1}{2}$ and $\frac{1}{2}$.

The first order perturbations $\Delta\delta_k$ and Δf_k are of the form

$$\left\{\begin{array}{rcl}\Delta\delta_k &=& \frac{1}{M_k}\,\Re\left(\frac{{\boldsymbol{u}_k}^H\Delta\boldsymbol{s}}{z_k\,\alpha_k^{(M_k-1)}}\right)\\ \Delta f_k &=& \frac{1}{2\pi\,M_k}\,\Im\left(\frac{{\boldsymbol{u}_k}^H\Delta\boldsymbol{s}}{z_k\,\alpha_k^{(M_k-1)}}\right),\end{array}\right. \tag{12}$$

where the vectors $\Delta\boldsymbol{s}=[\Delta s(t-l+1),\ldots,\Delta s(t+n-1)]^T$ and $\boldsymbol{u}_k$ (whose coefficients, defined in equation (45) *in the appendix, depend on the model parameters) have dimension $N=n+l-1$.*

Knowing the influence of a perturbation of the data onto the estimated poles, it is then possible to analyze the perturbation induced onto the amplitudes and phases, obtained from these poles by the least squares method. More precisely, it is shown in section B2 in the appendix that the complex amplitudes $\alpha_k^{(m)}(\varepsilon)$, real amplitudes $a_k^{(m)}(\varepsilon)$ and phases $\phi_k^{(m)}(\varepsilon)$ are $\mathcal{C}^\infty$ functions of ε in the neighborhood of $\varepsilon=0$, leading to the first order expansions in proposition 4.

Proposition 4 (Perturbation of the amplitudes and phases)**.** *Let $a_k^{(m)}(\varepsilon)=|\alpha_k^{(m)}(\varepsilon)|$ and $\phi_k^{(m)}(\varepsilon)=\frac{1}{2\pi}\,\Im\left(\ln(\alpha_k^{(m)}(\varepsilon))\right)$. If $a_k^{(m)}\neq 0$, then the functions $\varepsilon\mapsto a_k^{(m)}(\varepsilon)$ and $\varepsilon\mapsto\phi_k^{(m)}(\varepsilon)$ are $\mathcal{C}^\infty$ in the neighborhood of $\varepsilon=0$, and admit the first order expansion:*

$$\left\{\begin{array}{rcl}a_k^{(m)}(\varepsilon) &=& a_k^{(m)}+\varepsilon\,\Delta a_k^{(m)}+O\left(\varepsilon^2\right)\\ \phi_k^{(m)}(\varepsilon) &=& \phi_k^{(m)}+\varepsilon\,\Delta\phi_k^{(m)}+O\left(\varepsilon^2\right).\end{array}\right. \tag{13}$$

The first order perturbations $\Delta a_k^{(m)}$ and $\Delta\phi_k^{(m)}$ are of the form

$$\left\{\begin{array}{rcl}\Delta a_k^{(m)} &=& a_k^{(m)}\,\Re\left(\frac{{\boldsymbol{b}_k^{(m)}}^H\Delta\boldsymbol{s}}{\alpha_k^{(m)}}\right)\\ \Delta\phi_k^{(m)} &=& \Im\left(\frac{{\boldsymbol{b}_k^{(m)}}^H\Delta\boldsymbol{s}}{\alpha_k^{(m)}}\right),\end{array}\right. \tag{14}$$

where the vectors $\boldsymbol{b}_k^{(m)}$ (whose coefficients, defined in section B in the appendix, depend on the model parameters) have dimension N.

IV. Performance of the estimators

Here we aim at exploiting the results of the perturbation analysis in section III to derive the first and second moments of the generalized ESPRIT estimators [2] in the case of a high signal-to-noise ratio. The most remarkable property shown below is that the asymptotic efficiency of the estimators of all parameters associated to single undamped poles is minimum if and only if the number of rows n and the number of columns l of the data matrix satisfy either $n=2l$ or $l=2n$. In practice, this result allows to properly dimension the data matrix when the length of the observation window $N=n+l-1$ is fixed.

From now on, we suppose that $\Delta s(t)$ is a circular complex stationary process of variance 1. Then the stationary process $\varepsilon\Delta s(t)$ can be viewed as a complex noise of variance $\sigma^2=\varepsilon^2<<1$. Using vector notations, we can write $\boldsymbol{s}(\varepsilon)=\boldsymbol{s}+\varepsilon\Delta\boldsymbol{s}$, where the vectors $\boldsymbol{s}(\varepsilon)$, $\boldsymbol{s}$ and $\Delta\boldsymbol{s}$ contain the N successive samples of the corresponding signals in the observation window $\{t-l+1\ldots t+n-1\}$. The covariance matrix $\boldsymbol{\Gamma}\triangleq\mathbb{E}\left[\Delta\boldsymbol{s}\,\Delta\boldsymbol{s}^H\right]$ of the circular complex random vector $\Delta\boldsymbol{s}$ is an $N\times N$ Toeplitz matrix, whose diagonal coefficients are equal to 1. In the particular case of white noise, $\boldsymbol{\Gamma}$ is equal to the identity.

Section IV-A presents the first order performance analysis, which is then simplified in section IV-B in an asymptotic context. The proofs of the results presented in this section can be found in the appendix.

A. First order performance

The following proposition gives the variances of the estimated poles $\widehat{z}_k$, damping (or amplifying) factors $\widehat{\delta}_k$ and frequencies $\widehat{f}_k$, obtained by applying the generalized ESPRIT algorithm [2] to the perturbed signal $\widehat{\boldsymbol{s}}=\boldsymbol{s}(\varepsilon)$.

Proposition 5. *The estimator $\widehat{z}_k=z_k(\varepsilon)$ of the pole z_k is unbiased at the first order, and its variance is of the form*

$$\mathrm{var}\,(\widehat{z}_k)\sim\frac{\sigma^2}{\left(M_k\,a_k^{(M_k-1)}\right)^2}\,{\boldsymbol{u}_k}^H\boldsymbol{\Gamma}\,\boldsymbol{u}_k. \tag{15}$$

where the vector $\boldsymbol{u}_k$ is defined in equation (45)*. In the same way, the estimators $\widehat{\delta}_k=\delta_k(\varepsilon)$ and $\widehat{f}_k=f_k(\varepsilon)$ of the damping factor δ_k and of the frequency f_k are unbiased at the first order. Moreover, their respective variances are of the form*

$$\mathrm{var}(\widehat{\delta}_k)\sim\frac{\sigma^2\,e^{-2\delta_k}}{2\left(M_k\,a_k^{(M_k-1)}\right)^2}\,{\boldsymbol{u}_k}^H\boldsymbol{\Gamma}\,\boldsymbol{u}_k \tag{16}$$

$$\mathrm{var}(\widehat{f}_k)\sim\frac{1}{4\pi^2}\,\mathrm{var}(\widehat{\delta}_k). \tag{17}$$

The variances of $\widehat{\delta}_k$ and $\widehat{f}_k$ are derived from their first order expansions, presented in section A3. A remarkable similitude can be noticed between their expressions and the Cramér-Rao bounds in equations (9) and (10). However equations (16) and (17) are not self-explanatory, because they involve many auxiliary variables, via the vectors $\boldsymbol{u}_k$. It will be shown in section IV-B that under some additional hypotheses, they can be simplified.

These variances satisfy the following properties[8]:

- if the noise is white ($\boldsymbol{\Gamma}=\boldsymbol{I}_N$), the variances of $\widehat{\delta}_k$ and $\widehat{f}_k$ depend on the frequencies only by their differences[9],
- if z_k is a single pole, the variances of $\widehat{\delta}_k$ and $\widehat{f}_k$ do not depend on any phase, and they are inversely proportional to ${a_k^{(0)}}^2$, but they do not depend on any other amplitude[10].

Proposition 6 below gives the variances of the estimated complex amplitudes $\widehat{\boldsymbol{\alpha}}$, real amplitudes $\widehat{a}_k^{(m)}$ and phases $\widehat{\phi}_k^{(m)}$, obtained by means of the Least Squares (LS) method [2].

Proposition 6. *The estimator $\widehat{\boldsymbol{\alpha}}=\boldsymbol{\alpha}(\varepsilon)$ of the vector of complex amplitudes $\boldsymbol{\alpha}$ is unbiased at first order, and its covariance matrix is of the form*

$$\mathrm{cov}\,(\widehat{\boldsymbol{\alpha}})\sim\sigma^2\,\boldsymbol{B}^H\boldsymbol{\Gamma}\,\boldsymbol{B}, \tag{18}$$

[8]These properties can be proved by a thorough analysis of the particular structure of the vector $\boldsymbol{u}_k$, defined in equation (45).

[9]This property is also satisfied by the Cramér-Rao bounds [12].

[10]In the case of the Cramér Rao bounds, this property also applies to multiple poles [12].

where the $N \times r$ matrix $\boldsymbol{B}$ is defined in equation (53)*. In the same way, the estimators $\widehat{a}_k^{(m)} = a_k^{(m)}(\varepsilon)$ and $\widehat{\phi}_k^{(m)} = \phi_k^{(m)}(\varepsilon)$ of the real-valued amplitude $a_k^{(m)}$ and the phase $\phi_k^{(m)}$ are unbiased at first order. Moreover, their respective variances are of the form*

$$\mathrm{var}(\widehat{a}_k^{(m)}) \sim \frac{\sigma^2}{2}\, \boldsymbol{b}_k^{(m)H} \boldsymbol{\Gamma}\, \boldsymbol{b}_k^{(m)} \quad (19)$$

$$\mathrm{var}(\widehat{\phi}_k^{(m)}) \sim \frac{1}{{a_k^{(m)}}^2}\, \mathrm{var}(\widehat{a}_k^{(m)}), \quad (20)$$

where the vectors $\boldsymbol{b}_k^{(m)}$, introduced in proposition 4, are the columns of the matrix $\boldsymbol{B}$.

As for proposition 5, the variances of $\widehat{\alpha}_k$, $\widehat{a}_k$ and $\widehat{\phi}_k$ are obtained from their first order expansions, presented in section B2.

B. Asymptotic performance

Let us suppose that the noise is white ($\boldsymbol{\Gamma} = \boldsymbol{I}_N$) and that all poles are on the unit circle ($\forall k \in \{0 \dots K-1\}$, $\delta_k \to 0$). We present below some asymptotic expansions of the estimator variances with respect to n, l and N.

Corollary 7. *If z_k is a single pole, the variances of the estimators $\widehat{\delta}_k$ and $\widehat{f}_k$ admit the expansions*

$$\mathrm{var}(\widehat{\delta}_k) \sim \frac{\sigma^2}{\max(n,l)^2 \min(n,l) {a_k^{(0)}}^2} + O\left(\frac{1}{N^4}\right) \quad (21)$$

$$\mathrm{var}(\widehat{f}_k) \sim \frac{\sigma^2}{4\pi^2 \max(n,l)^2 \min(n,l) {a_k^{(0)}}^2} + O\left(\frac{1}{N^4}\right) \quad (22)$$

Both of them are minimal for $n = 2l = \frac{2}{3}(N+1)$ or for $l = 2n = \frac{2}{3}(N+1)$ (if N equals 2 modulo 3), and these minima admit the asymptotic expansions

$$\mathrm{var}(\widehat{\delta}_k) \sim \frac{27}{4} \frac{\sigma^2}{N^3 {a_k^{(0)}}^2}$$

$$\mathrm{var}(\widehat{f}_k) \sim \frac{27}{4} \frac{\sigma^2}{4\pi^2 N^3 {a_k^{(0)}}^2}.$$

The proof of corollary 7 relies on the first order expansion of the vector $\boldsymbol{u}_k$ involved in equations (16) and (17), which admits a simple closed form. Attention must be paid to the fact that expressions (21) and (22) are only valid for a single pole. If z_k is a multiple pole, these variances cannot be formulated in such a simple way and are function of the complex amplitudes associated to z_k for all indices $m \in \{0 \dots M_k - 1\}$ (the optimal values of n and l are also function of these amplitudes in this case). These variances can be compared to the asymptotic Cramér-Rao bounds given in equations (13) and (15):

Under the above hypotheses, the asymptotic efficiency of the estimators of all damping factors and all frequencies associated to the single poles is the same one, independent from the model parameters, and equal to $9/8 = 1.125$ if $n = 2l$ or $l = 2n$.

In this way, the results obtained in [19] about the Matrix Pencil method (in the particular case of a single complex sinusoid) are recovered. Figure 1-a represents the ratio between the Cramér-Rao bound and the variance of the estimators in a logarithmic scale, as a function of the ratio $\frac{n}{N+1}$. Thus it can be verified that the maximum is reached at $n = \frac{N+1}{3}$ and $n = \frac{2(N+1)}{3}$ as expected. Besides, the performance collapses when n becomes too high or too small.

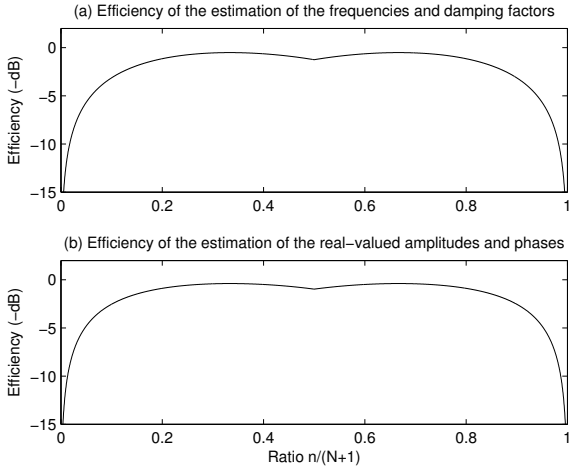

Fig. 1. Efficiency of the estimators

The following corollary is the analogue of corollary 7 for the real-valued amplitudes and phases.

Corollary 8. *If z_k is a single pole, the variances of the estimators $\widehat{a}_k^{(0)}$ and $\widehat{\phi}_k^{(0)}$ admit the expansions*

$$\mathrm{var}(\widehat{a}_k^{(0)}) \sim \frac{\sigma^2}{2}\left(\frac{1}{N} + \frac{N^2}{2\max(n,l)^2 \min(n,l)}\right) + O\left(\frac{1}{N^2}\right) \quad (23)$$

$$\mathrm{var}(\widehat{\phi}_k^{(0)}) \sim \frac{\sigma^2}{2{a_k^{(0)}}^2}\left(\frac{1}{N} + \frac{N^2}{2\max(n,l)^2 \min(n,l)}\right) + O\left(\frac{1}{N^2}\right). \quad (24)$$

Both of them are minimal for $n = 2l = \frac{2}{3}(N+1)$ or for $l = 2n = \frac{2}{3}(N+1)$ (if N equals 2 modulo 3), and these minima admit the asymptotic expansions

$$\mathrm{var}(\widehat{a}_k^{(0)}) \sim \frac{35}{16} \frac{\sigma^2}{N}$$

$$\mathrm{var}(\widehat{\phi}_k^{(0)}) \sim \frac{35}{16} \frac{\sigma^2}{N {a_k^{(0)}}^2}.$$

The proof of corollary 8 relies on the first order expansions of the vectors $\boldsymbol{b}_k^{(m)}$ involved in equations (19) and (20), which admit a simple closed form. Again, attention must be paid to the fact that expressions (23) and (24) are only valid for a single pole. These variances can be compared to the asymptotic Cramér-Rao bounds given in equations (16) and (17):

Under the above hypotheses, the asymptotic efficiency of the estimators of all the real-valued amplitudes and phases associated to single poles is the same one, independent from the model parameters, and equal to $35/32 = 1.09375$ if $n = 2l$ or $l = 2n$.

This efficiency is even better than that of the estimators $\widehat{\delta}_k$ and $\widehat{f}_k$. It can also be noticed that the optimum is obtained for the same values of n and l as in the previous case. Figure 1-b represents the ratio between the Cramér-Rao bound and the variance of the estimators in a logarithmic scale, as a function of the ratio $\frac{n}{N+1}$. Again, the maximum is reached at $n = \frac{N+1}{3}$ and $n = \frac{2(N+1)}{3}$, and the performance collapses when n becomes too high or too small. The similitude between the curves represented in figures 1-a and 1-b is noticeable. This could be explained by the fact that the estimation of the amplitudes and the phases directly relies on the estimation of the frequencies and the damping factors.

V. Simulation results

A. Dependence of the variances with respect to the PACE parameters

This section illustrates the variations of the estimators variances with respect to the parameters of the PACE model. Note that propositions 5, 6 and corollaries 7, 8 show a rather simple dependency on the amplitudes and the variance σ^2. Therefore we focus here on the dependency on the frequency gap between two components (section V-A1), the damping factor (section V-A2), the spectral flatness of the noise (section V-A3), and the order of a pole (section V-A4). For these simulations, the same synthetic signals as those introduced in [12] are used. In the figures below, the solid lines represent the theoretical variance of the frequency estimators or that of the damping factor estimators, which are equivalent according to (17). In the same way, the dashed lines represent the theoretical *relative* variance of the amplitude estimators, which is equal to that of the phase estimators, according to (20).

1) Variation of the variances with respect to frequency gaps: We consider a signal of length $N = 200$, composed of two undamped components ($K = 2$) of same order $M_0 = M_1 = 1$, in white noise ($\mathbf{\Gamma} = \boldsymbol{I}_N$ and $\sigma^2 = 1$). These components have zero phases, and same amplitudes, such that $\mathrm{SNR}_0^{(0)} = \mathrm{SNR}_1^{(0)} = 50$ dB. Figures 2-a and 2-b show the variations of the variances of the estimators obtained with $n = \frac{2}{3}(N+1)$, with respect to the frequency gap $\Delta f = |f_1 - f_0| \in]0, 0.5]$ (f_0 was set to 0). The diamonds represent the empirical variance, obtained by averaging 100 runs of the ESPRIT algorithm. They match the theoretical variance, which confirms the validity of our perturbation analysis for this SNR. The variation rate of the variances is similar to that of the Cramér Rao bounds [12]: it is broken at $\Delta f = \frac{1}{N} = 5\,10^{-3}$, which corresponds to the resolution limit of Fourier analysis. At this limit point, the relative variance of the amplitude estimate is still lower than -60 dB, which shows the good resolution of the ESPRIT algorithm. The efficiencies of both estimators are represented in figure 2-c. It can be noticed that they remain close to 1, even when the frequency gap tends to zero.

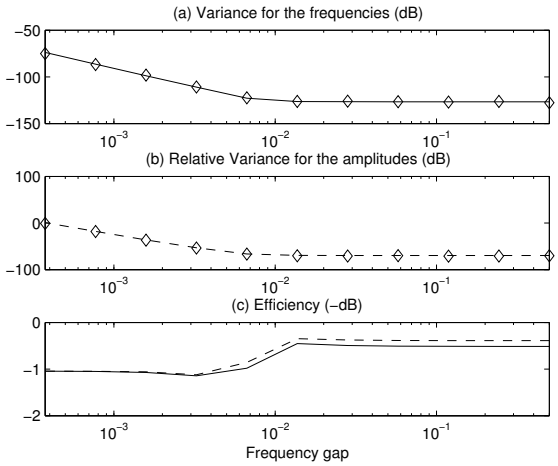

Fig. 2. Variation of the variances with respect to the frequency gap

2) Variation of the variances with respect to the damping factor: We consider a signal of length $N = 100$, composed of one component ($K = 1$) of order $M_0 = 1$, in white noise ($\mathbf{\Gamma} = \boldsymbol{I}_N$ and $\sigma^2 = 1$). This component has zero frequency and phase, and an amplitude such that $\mathrm{SNR}_0^{(0)} = 50$ dB. Figures 3-a and 3-b show the variations of the variances of the estimators obtained with $n = \frac{2}{3}(N+1)$, with respect to the damping factor δ_0. The diamonds represent the empirical variance, obtained by averaging 500 runs of the ESPRIT algorithm. They no longer exactly match the theoretical variance for negative values of delta, which shows the validity limit of our perturbation analysis with respect to the high SNR hypothesis. Again, these variations are very similar to those of the Cramér-Rao bounds illustrated in [12]. The efficiencies of both estimators are represented in figure 3-c. They remain close to 1 whatever the value of the damping factor is.

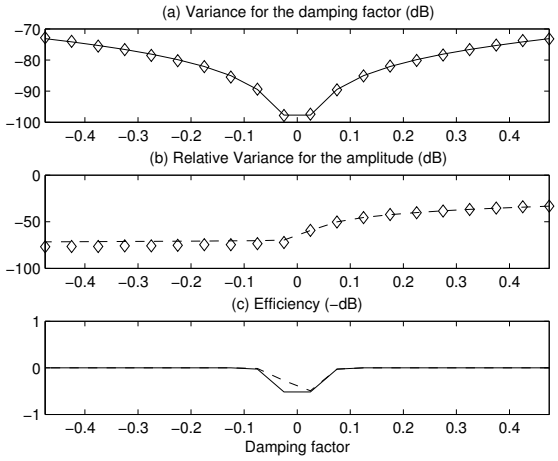

Fig. 3. Variation of the variances with respect to the damping factor

3) Variation of the variances with respect to the spectral flatness of the noise: We consider a signal of length $N = 100$, composed of one undamped component ($K = 1$) of order

$M_0 = 1$, in colored noise. This component has a zero phase, a normalized frequency equal to 0.05, and an amplitude such that $\mathrm{SNR}_0^{(0)}$ = 50 dB. The noise is obtained by filtering a white noise by the filter of transfer function $H_a(z) = \frac{1}{1-a\,z^{-1}}$ (where $0 \leq a < 1$), such that $\mathbf{\Gamma} = \mathrm{Toeplitz}(1, a, a^2 \ldots a^{N-1})$. The Spectral Flatness (SF) measure of the noise is defined as

$$\mathrm{SF}(a) = \frac{\exp\left(\int_0^1 \ln\left(|H_a(e^{i2\pi f})|^2\right) \mathrm{d}f\right)}{\int_0^1 |H_a(e^{i2\pi f})|^2 \mathrm{d}f}.$$

By tuning the parameter a, it is possible to make the spectral flatness map the range $]0, 1]$ (the case SF= 1 corresponds to white noise). Figure 4 illustrates the variations of the variances of the estimators obtained with $n = \frac{2}{3}(N+1)$, with respect to the spectral flatness of the noise. As expected, figure 4-c shows that the efficiency degrades when the spectral flatness decreases (note that the ESPRIT algorithm explicitly relies on the white noise assumption). In other respects, figures 4-a and 4-b show that the variances admit a maximum when SF $\simeq$ 0.5. A this point, the theoretical and the empirical variances (obtained by averaging 1000 runs) no longer match in figure 4-b, which shows the validity limit of our perturbation analysis for this SNR. In the range SF $\in [0.5, 1]$, the variances decrease when the spectral flatness increases, as expected. In the range SF $\in]0.02, 0.5]$, we observe the inverse phenomenon. Indeed, as mentioned in [12], the power spectral density of the noise becomes a sharp peak when SF becomes low, and converges to a spectral line when SF$\rightarrow$ 0. Therefore the problem of estimating the single undamped component in colored noise becomes close to the problem of estimating two undamped components *without noise*. However, contrary to what is observed in [12] in the case of the Cramér-Rao bounds, we note that the variances are stationary in the interval SF $\in]0, 0.02]$. This is because the ESPRIT algorithm is applied with an "erroneous" model order ($K = 1$ instead of 2)[11].

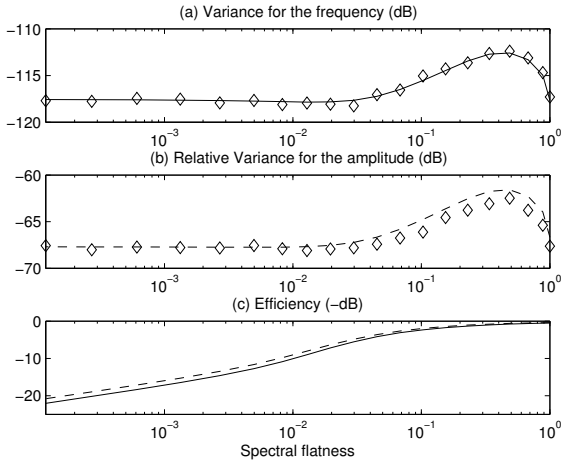

Fig. 4. Variation of the variances with respect to the spectral flatness of the noise

[11]The impact of an erroneous modeling order on the estimated parameters was studied in [33].

4) Variation of the variances with respect to the pole order: We consider a signal of length $N = 20$, composed of one undamped component ($K = 1$) of order $M_0 \in \{1 \ldots 4\}$, in white noise ($\mathbf{\Gamma} = \boldsymbol{I}_N$ and $\sigma^2 = 1$). This component has zero phases, and amplitudes such that $\mathrm{SNR}_0^{(M_0-1)}$ = 50 dB, and $\forall m < M_0 - 1$, $\mathrm{SNR}_0^{(m)}$ = 0. The corresponding pole is $z_0 = 1$. Figures 5-a and 5-b show the variations of the variances of the estimators obtained with $n = \frac{2}{3}(N+1)$, with respect to the pole order M_0. They confirm the results obtained for the Cramér-Rao bounds [12]: estimating multiple poles is all the more difficult as their order is high. Actually this estimation is no longer possible in this case if $M_0 > 4$. Besides, the empirical variance (obtained by averaging 100 runs) and the theoretical variance no longer exactly match for $M_0 \geq 4$, which shows the validity limit of our perturbation analysis for this SNR (we observed that they match again if the SNR becomes greater than 80 dB). In other respects, figure 5-c shows that the efficiency rapidly degrades when M_0 increases. We may infer that the arithmetic mean of the scattered eigenvalues is not a sufficiently reliable estimator for a pole of high multiplicity. Some ideas to improve this estimator are suggested in [34].

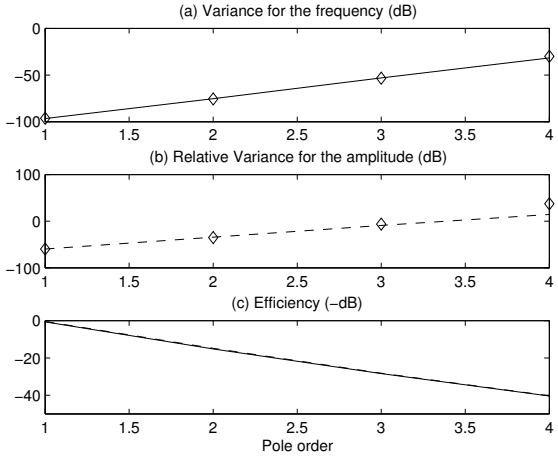

Fig. 5. Variation of the variances with respect to the pole order

B. Application to ARMA filter synthesis

As shown in section V-A4, estimating multiple poles is a difficult task. In order to illustrate the practical capabilities of the generalized ESPRIT algorithm in presence of multiple poles, we propose below an application to ARMA filter synthesis, in the context of system conversion from continuous time to discrete time.

1) Principle: We consider continuous time systems defined by an ordinary differential equation (ODE) with real-valued coefficients (below τ denotes the continuous time, in seconds)

$$a_0\widetilde{y}(\tau) + a_1\widetilde{y}'(\tau) + \ldots + a_p\widetilde{y}^{(p)}(\tau) = b_0\widetilde{x}(\tau) + \ldots + b_q\widetilde{x}^{(q)}(\tau),$$

whose transfer function can be written in the form

$$\widetilde{H}(\rho) = \frac{\sum_{k=0}^{q} b_k \rho^k}{\sum_{k=0}^{p} a_k \rho^k},$$

where $\rho = i2\pi\nu$ and ν denotes the frequency, in Hertz. Moreover, we suppose that $\widetilde{x}(\tau)$ and $\widetilde{y}(\tau)$ satisfy the conditions of the Nyquist-Shannon sampling theorem: $\exists T > 0$ such that the Fourier transforms $\widetilde{X}(\nu)$ and $\widetilde{Y}(\nu)$ are zero outside the range $[-\frac{1}{2T}, \frac{1}{2T}]$, and we consider the discrete signals $x(t) = \widetilde{x}(tT)$ and $y(t) = \widetilde{y}((t-t_0)T)$ for all $t \in \mathbb{Z}$ (where $t_0 \in \mathbb{R}$), obtained by sampling $x(\tau)$ and $y(\tau)$ at the frequency $1/T$. Then it is well known that $y(t)$ can be obtained from $x(t)$ by applying the discrete filter of frequency response

$$H(e^{i2\pi f}) = \widetilde{H}\left(i2\pi\frac{f}{T}\right) e^{-i2\pi t_0 f} = \frac{\sum_{k=0}^{q} b_k \left(i2\pi\frac{f}{T}\right)^k}{\sum_{k=0}^{p} a_k \left(i2\pi\frac{f}{T}\right)^k} e^{-i2\pi t_0 f}$$

where $f \in \left[-\frac{1}{2}, \frac{1}{2}\right]$ is the normalized frequency. It can be noticed that the time delay t_0 can be chosen such that $H(e^{i2\pi\frac{1}{2}}) = H(e^{-i2\pi\frac{1}{2}})$. In this case, if the denominator is never zero, the 1-periodic function $f \mapsto H(e^{i2\pi f})$ is continuous and piecewise continuously differentiable, which proves that $h(t) = O(\frac{1}{t^2})$, thus the discrete filter h is stable. The impulse response $h(t)$ can then be obtained by numerically computing the inverse discrete time Fourier transform of $H(e^{i2\pi f})$. However this impulse response $h(t)$ is generally infinite, and we aim at approximating it by an AutoRegressive Moving Average (ARMA) filter $g(t)$.

It is well known that the general rational transfer function of a stable ARMA filter can be decomposed in the form

$$G(z) = \sum_{m=0}^{M_0-1} \alpha_0^{(m)} z^{-m} + \sum_{0<|z_k|<1} \sum_{m=0}^{M_k-1} \frac{\alpha_k^{(m)}}{(1-z_k z^{-1})^{(1+m)}} + \sum_{1<|z_k|<+\infty} \sum_{m=0}^{M_k-1} \frac{\beta_k^{(m)} (-z/z_k)^{(1+m)}}{(1-z/z_k)^{(1+m)}} \quad (25)$$

where for all $k \in \{0 \dots K-1\}$ z_k is a pole of multiplicity M_k (here we assume $z_0 = 0$). This transfer function corresponds to the impulse response

$$g(t) = \sum_{0\le|z_k|<1} \sum_{m=0}^{M_k-1} \alpha_k^{(m)} F_m(t) z_k^{t-m}$$

for all $t \ge 0$, and

$$g(t) = \sum_{1<|z_k|<+\infty} \sum_{m=0}^{M_k-1} (-1)^{(1+m)} \beta_k^{(m)} F_m(|t|-1) z_k^{-|t|}$$

for all $t < 0$.

Thus both the causal and anticausal parts of $g(t)$ satisfy a PACE model[12]. Following this remark, we can find an ARMA filter g which approximates the discrete filter h by applying the generalized ESPRIT algorithm to the impulse response $h(t)$ on two appropriately chosen finite intervals.

[12]Note that the causal part generally contains a multiple pole $z_0 = 0$, whereas the anticausal part never contains a pole at $z = \infty$.

2) Example: The ARMA filter synthesis method could be successfully used for designing differentiator or integrator filters. Here it is applied to the continuous time system

$$\widetilde{y}(\tau) + 3\widetilde{y}'(\tau) + 3\widetilde{y}''(\tau) + \widetilde{y}'''(\tau) = \widetilde{x}(\tau) - \widetilde{x}'(\tau), \quad (26)$$

with parameters $T = 1$ and[13] $t_0 \simeq 2.3924$. The impulse and the frequency response of the corresponding discrete filter h are represented in figure 6.

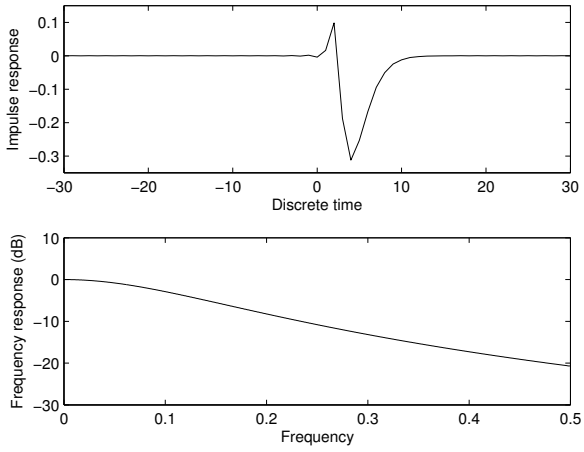

Fig. 6. Impulse and frequency response of the discrete filter

Equation (26) shows that the continuous time filter contains a triple pole at $\rho = -1$. When synthesizing the corresponding ARMA filter, we thus expect to find a triple pole at $z = e^{\frac{\rho}{T}} = 1/e$. Figure 7 represents the estimated poles of the causal part of g in the complex plane, obtained by applying the ESPRIT algorithm with $N = 128$, $n = (N+1)/3 = 43$, and[14] $r = 15$. As expected, we observe a triple pole in the neighborhood of $1/e \simeq 0.3679$ (which is scattered into three single eigenvalues forming the vertices of an equilateral triangle). We also observe a double pole in the neighborhood of 0, which corresponds to the polynomial part of the transfer function in equation (25). Finally, we obtained an ARMA filter g with 26 poles and 27 zeros. The approximation error for the impulse response was $\max_{n\in[-128,128]} |h(n) - g(n)| < 5\,10^{-8}$, and the approximation error for the frequency response was $\max_{f\in[-\frac{1}{2},\frac{1}{2}]} |20\log_{10} H(e^{i2\pi f}) - 20\log_{10} G(e^{i2\pi f})| < 10^{-5}$.

VI. Conclusions

In this paper, the performance of the generalized ESPRIT algorithm for estimating the parameters of the PACE model, also called quasipolynomial model, was investigated in the context of high SNRs. This work was based on the analysis of the first order perturbations induced by an additive noise. In particular, it was shown that the perturbation of the estimated

[13]The fractional part of the delay t_0 was chosen in order to make the frequency response $H(e^{i2\pi f})$ continuous, and the integer part of t_0 was chosen as the smallest integer such that the anticausal part of the estimated filter g has no pole at $z = \infty$.

[14]The order r was selected by means of Information Theoretic Criteria [35].

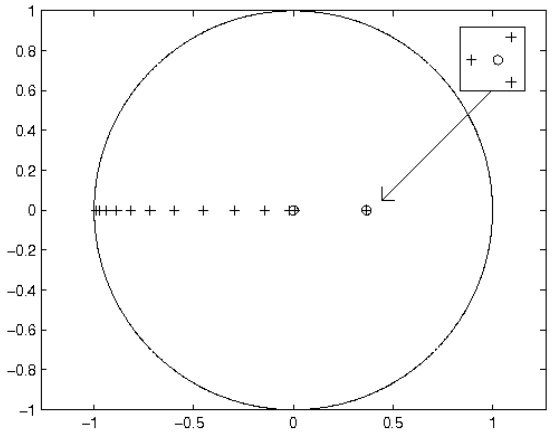

Fig. 7. Poles of the causal part of the ARMA filter

poles is not sensitive to the particular choice of the orthonormal subspace weighting matrix. In other respects, the presence of noise scatters multiple poles into several single eigenvalues, forming the vertices of a regular polygon. However it was proved that the estimation of multiple poles can be improved by calculating the arithmetic mean of the scattered eigenvalues.

Then it was shown that the estimators of all the parameters obtained in this way are unbiased, and their variances were calculated and compared to the Cramér-Rao bounds. By supposing that the noise is white, that all poles are on the unit circle, and that the SNR, n and $l \to +\infty$, it was shown that the efficiency of single poles estimators is close to 1. More precisely, the asymptotic efficiency of the estimators of all damping factors and frequencies is the same one, independent from the model parameters, and equal to $9/8 = 1.125$ if $n = 2l$ or $l = 2n$. In other respects, the asymptotic efficiency of the estimators of all the real-valued amplitudes and phases is the same one, independent from the model parameters, and equal to $35/32 = 1.09375$ if $n = 2l$ or $l = 2n$.

However, our simulation results showed that the variances of the estimators associated to a multiple pole present an exponential increase with the order of the pole. Thus the practical estimation of the PACE model parameters is only possible if the exponentials are modulated by polynomials of low order. Nevertheless, some recent advances in linear algebra computations, such as the techniques proposed in [34], offer interesting outlooks for improving the estimation of multiple poles.

APPENDIX

The following developments lead to propositions 3 and 4 presented in section III. The exhaustive proofs of these results are presented in a supporting document which is available at `http://www.tsi.enst.fr/~rbadeau/`.

A. Perturbation of the frequencies and damping factors

Below, we successively derive the following first order expansions:

- $\boldsymbol{W}(\varepsilon) = \boldsymbol{W} + \varepsilon\Delta\boldsymbol{W} + O(\varepsilon^2)$, where the expression of $\Delta\boldsymbol{W}$ is a linear function of $\Delta\boldsymbol{S}$ (proposition 9);
- $\boldsymbol{\Phi}(\varepsilon) = \boldsymbol{\Phi} + \varepsilon\Delta\boldsymbol{\Phi} + O(\varepsilon^2)$, where the expression of $\Delta\boldsymbol{\Phi}$ is a linear function of $\Delta\boldsymbol{W}$ (proposition 10);
- $z_k(\varepsilon) = z_k + \varepsilon\Delta z_k + O(\varepsilon^2)$, where the expression of Δz_k is a linear function of $\Delta\boldsymbol{\Phi}$ (proposition 12);
- $\begin{cases} \delta_k(\varepsilon) = \delta_k + \varepsilon\Delta\delta_k + O(\varepsilon^2) \\ f_k(\varepsilon) = f_k + \varepsilon\Delta f_k + O(\varepsilon^2) \end{cases}$ where the expressions of $\Delta\delta_k$ and Δf_k are functions of Δz_k (corollary 3).

Finally, successive substitutions lead to the expressions of $\Delta\delta_k$ and Δf_k as functions of the additive perturbation Δs.

1) Perturbation of the signal subspace: Here we analyze the influence of a perturbation of the data onto the signal subspace. For all ε, let $\boldsymbol{\Pi}(\varepsilon)$ be the $n \times n$ projector onto the r-dimensional dominant subspace of the $n \times n$ positive semidefinite matrix $\boldsymbol{S}(\varepsilon)\,\boldsymbol{S}(\varepsilon)^H$. If $\boldsymbol{W}$ is orthonormal, $\boldsymbol{\Pi}(0) = \boldsymbol{W}\boldsymbol{W}^H$. Then the perturbation theory shows that the function $\varepsilon \mapsto \boldsymbol{\Pi}(\varepsilon)$ is $\mathcal{C}^\infty$ in a neighborhood of $\varepsilon = 0$. In the literature, the asymptotic performance analysis of some subspace-based algorithms was performed by investigating the perturbation of this projector at the first order [36]. However we are interested here in the perturbation of the subspace weighting matrix $\boldsymbol{W}$, which is analyzed in the following proposition.

Proposition 9 (Perturbation of the signal subspace)**.** *There exists an infinity*[15] *of $\mathcal{C}^\infty$ functions $\varepsilon \mapsto \boldsymbol{W}(\varepsilon)$ defined in a neighborhood of $\varepsilon = 0$ and with values in the group of complex orthonormal matrices $\mathbb{U}^{n\times r}$, which span the r-dimensional dominant subspace of $\boldsymbol{S}(\varepsilon)$ (*i.e. *such that $\boldsymbol{W}(0) = \boldsymbol{W}$ and $\boldsymbol{W}(\varepsilon)\,\boldsymbol{W}(\varepsilon)^H = \boldsymbol{\Pi}(\varepsilon)$). Each $\mathcal{C}^\infty$ function $\varepsilon \mapsto \boldsymbol{W}(\varepsilon)$ admits a first order expansion*

$$\boxed{\boldsymbol{W}(\varepsilon) = \boldsymbol{W} + \varepsilon\Delta\boldsymbol{W} + O(\varepsilon^2).} \tag{27}$$

The first order perturbation $\Delta\boldsymbol{W}$ can be decomposed as

$$\Delta\boldsymbol{W} = \Delta\boldsymbol{W}^{\perp} - \boldsymbol{W}\boldsymbol{A}, \tag{28}$$

where the $r \times r$ matrix $\boldsymbol{A} \triangleq \Delta\boldsymbol{W}^H\boldsymbol{W}$ is antihermitian and the $n \times r$ matrix $\Delta\boldsymbol{W}^{\perp}$ is orthogonal to $\mathrm{span}(\boldsymbol{W})$*:*

$$\boxed{\Delta\boldsymbol{W}^{\perp} = \left(\boldsymbol{I}_n - \boldsymbol{W}\,\boldsymbol{W}^H\right)\Delta\boldsymbol{S}\,\boldsymbol{S}^{\dagger}\,\boldsymbol{W}.} \tag{29}$$

The proof of proposition 9 relies on the implicit definition of $\boldsymbol{W}(\varepsilon)$ as the unique minimum point of a cost function $\mathcal{J}$. Equations (27) to (29) are thus derived by zeroing the first derivative of $\mathcal{J}$.

Proof of proposition 9:

It can be verified that the function

$$\boldsymbol{W}(\varepsilon) \triangleq \boldsymbol{\Pi}(\varepsilon)\boldsymbol{W}\left(\boldsymbol{W}^H\boldsymbol{\Pi}(\varepsilon)\boldsymbol{W}\right)^{-\frac{1}{2}}$$

satisfies all the properties mentioned in proposition 9. It is also clear that any function of the form $\boldsymbol{W}(\varepsilon)\,\boldsymbol{\Theta}(\varepsilon)$, where $\varepsilon \mapsto \boldsymbol{\Theta}(\varepsilon)$ is a $\mathcal{C}^\infty$ function, whose values belong to the group of orthonormal matrices $\mathcal{O}_r(\mathbb{C})$, and which reaches the value $\boldsymbol{I}_r$ at $\varepsilon = 0$, also satisfies these properties. Lastly, if $\varepsilon \mapsto \boldsymbol{W}'(\varepsilon)$ is an other function satisfying all these properties, then $\boldsymbol{\Pi}(\varepsilon) = \boldsymbol{W}(\varepsilon)\,\boldsymbol{W}(\varepsilon)^H = \boldsymbol{W}'(\varepsilon)\,\boldsymbol{W}'(\varepsilon)^H$. Therefore $\boldsymbol{W}'(\varepsilon) = \boldsymbol{W}(\varepsilon)\,\boldsymbol{\Theta}(\varepsilon)$, where $\boldsymbol{\Theta}(\varepsilon) \triangleq \boldsymbol{W}(\varepsilon)^H\boldsymbol{W}'(\varepsilon)$ is a $\mathcal{C}^\infty$ function, whose values belong to the group of orthonormal matrices $\mathcal{O}_r(\mathbb{C})$ since $\boldsymbol{W}(\varepsilon)$ and $\boldsymbol{W}'(\varepsilon)$ are two orthonormal bases of the same subspace, and which reach the value $\boldsymbol{W}^H\boldsymbol{W} = \boldsymbol{I}_r$ at $\varepsilon = 0$. Then note that according to [37], any orthonormal matrix $\boldsymbol{W}(\varepsilon)$ spanning the principal subspace of the matrix $\boldsymbol{S}(\varepsilon)\boldsymbol{S}(\varepsilon)^H$ minimizes the function

$$\boxed{\begin{array}{rccl} \mathcal{J}: & \mathbb{C}^{n\times r} & \to & \mathbb{R} \\ & \boldsymbol{U} & \mapsto & \frac{1}{2}\left\|\boldsymbol{S}(\varepsilon) - \boldsymbol{U}\,\boldsymbol{U}^H\,\boldsymbol{S}(\varepsilon)\right\|_F^2. \end{array}}$$

[15] All these functions are obtained by right-multiplying any of them by a $\mathcal{C}^\infty$ function, with values in the group of complex orthonormal matrices $\mathbb{U}^{r\times r}$, and reaching the value $\boldsymbol{I}_r$ at $\varepsilon = 0$.

Consequently, the gradient $\boldsymbol{\nabla}\mathcal{J}(\boldsymbol{U})$ is zero at $\boldsymbol{U}=\boldsymbol{W}(\varepsilon)$. However it can be verified[16] that $\boldsymbol{\nabla}\mathcal{J}(\boldsymbol{U})=$

$$\left(-2\boldsymbol{S}(\varepsilon)\boldsymbol{S}(\varepsilon)^H+\boldsymbol{S}(\varepsilon)\boldsymbol{S}(\varepsilon)^H\boldsymbol{U}\boldsymbol{U}^H+\boldsymbol{U}\boldsymbol{U}^H\boldsymbol{S}(\varepsilon)\boldsymbol{S}(\varepsilon)^H\right)\boldsymbol{U}.$$

Let $\boldsymbol{W}(\varepsilon)=\boldsymbol{W}+\varepsilon\Delta\boldsymbol{W}+O(\varepsilon^2)$ be the first order expansion of the function $\varepsilon\mapsto\boldsymbol{W}(\varepsilon)$. Then

$$\boldsymbol{\nabla}\mathcal{J}(\boldsymbol{W}(\varepsilon))=\varepsilon\boldsymbol{\nabla}\mathcal{J}^{(1)}+O(\varepsilon^2) \tag{30}$$

where

$$\begin{aligned}\boldsymbol{\nabla}\mathcal{J}^{(1)} &= -\left(\boldsymbol{I}_n-\boldsymbol{W}\boldsymbol{W}^H\right)\Delta\boldsymbol{S}\,\boldsymbol{S}^H\boldsymbol{W}\\ &\quad+\Delta\boldsymbol{W}\left(\boldsymbol{W}^H\boldsymbol{S}\,\boldsymbol{S}^H\boldsymbol{W}\right)+\boldsymbol{W}\,\boldsymbol{N}\end{aligned} \tag{31}$$

and

$$\begin{aligned}\boldsymbol{N} &\triangleq \left(\boldsymbol{W}^H\boldsymbol{S}\,\boldsymbol{S}^H\boldsymbol{W}\right)\left(\boldsymbol{W}^H\Delta\boldsymbol{W}+\Delta\boldsymbol{W}^H\boldsymbol{W}\right)\\ &\quad+\Delta\boldsymbol{W}^H\boldsymbol{W}\left(\boldsymbol{W}^H\boldsymbol{S}\,\boldsymbol{S}^H\boldsymbol{W}\right).\end{aligned} \tag{32}$$

However, the first order expansion of the orthonormality condition $\boldsymbol{W}(\varepsilon)^H\boldsymbol{W}(\varepsilon)=\boldsymbol{I}_r$ shows that $\boldsymbol{W}^H\Delta\boldsymbol{W}+\Delta\boldsymbol{W}^H\boldsymbol{W}=\boldsymbol{0}_{(r\times r)}$, which means that the matrix $\boldsymbol{A}\triangleq\Delta\boldsymbol{W}^H\boldsymbol{W}$ is antihermitian. Thus $\boldsymbol{N}=\boldsymbol{A}\left(\boldsymbol{W}^H\boldsymbol{S}\,\boldsymbol{S}^H\boldsymbol{W}\right)$. Since $\boldsymbol{\nabla}\mathcal{J}(\boldsymbol{W}(\varepsilon))=\boldsymbol{0}_{(n\times r)}$, equations (30) to (32) yield

$$\Delta\boldsymbol{W}=\left(\boldsymbol{I}_n-\boldsymbol{W}\boldsymbol{W}^H\right)\Delta\boldsymbol{S}\boldsymbol{S}^H\boldsymbol{W}\left(\boldsymbol{W}^H\boldsymbol{S}\boldsymbol{S}^H\boldsymbol{W}\right)^{-1}-\boldsymbol{W}\boldsymbol{A}.$$

Finally, by noting that $\boldsymbol{S}^H\,\boldsymbol{W}\left(\boldsymbol{W}^H\boldsymbol{S}\,\boldsymbol{S}^H\boldsymbol{W}\right)^{-1}=\boldsymbol{S}^\dagger\,\boldsymbol{W}$, equations (27) to (29) can be derived. ∎

2) Perturbation of the spectral matrix: The following proposition complements the result of proposition 9 by showing how the spectral matrix is perturbed.

Proposition 10 (Perturbation of the spectral matrix)**.** *Suppose that the matrix $\boldsymbol{W}_\downarrow$ is full-rank. Then in the neighborhood of $\varepsilon=0$, $\boldsymbol{W}_\downarrow(\varepsilon)$ is also full-rank. Moreover, the function*

$$\boldsymbol{\Phi}(\varepsilon)\triangleq\boldsymbol{W}(\varepsilon)_\downarrow^\dagger\,\boldsymbol{W}(\varepsilon)_\uparrow \tag{33}$$

is $\mathcal{C}^\infty$, and admits the first order expansion:

$$\boxed{\boldsymbol{\Phi}(\varepsilon)=\boldsymbol{\Phi}+\varepsilon\Delta\boldsymbol{\Phi}+O(\varepsilon^2).} \tag{34}$$

The first order perturbation $\Delta\boldsymbol{\Phi}$ can be written in the form

$$\Delta\boldsymbol{\Phi}=\Delta\boldsymbol{\Phi}^\perp+\boldsymbol{A}\boldsymbol{\Phi}-\boldsymbol{\Phi}\boldsymbol{A}, \tag{35}$$

where the $r\times r$ matrix $\Delta\boldsymbol{\Phi}^\perp$ is defined as

$$\boxed{\Delta\boldsymbol{\Phi}^\perp=-\boldsymbol{W}_\downarrow^\dagger\Delta\boldsymbol{W}_\downarrow^\perp\,\boldsymbol{\Phi}+\boldsymbol{\Phi}\,\boldsymbol{W}_\uparrow^\dagger\Delta\boldsymbol{W}_\uparrow^\perp.} \tag{36}$$

Equations (34) to (36) are obtained by substituting equations (27) to (29) into the first order expansion of equation (33) which defines $\boldsymbol{\Phi}(\varepsilon)$.

In the following step, the estimated poles are defined as the eigenvalues of the perturbed spectral matrix $\boldsymbol{\Phi}(\varepsilon)$. In order to compute their first order expansion, we first need to introduce the matrix $\boldsymbol{J}(\varepsilon)\triangleq\boldsymbol{G}^{-1}\boldsymbol{\Phi}(\varepsilon)\boldsymbol{G}$, which is similar to $\boldsymbol{\Phi}(\varepsilon)$; thus the estimated poles can also be viewed as the eigenvalues of $\boldsymbol{J}(\varepsilon)$. Note however that $\boldsymbol{J}(\varepsilon)$ is generally no longer a Jordan matrix when $\varepsilon>0$. The following corollary provides the first order expansion of this matrix. Let us define the vectors $\boldsymbol{v}_\uparrow$ and $\boldsymbol{v}_\downarrow$ of same dimension r as the conjugate transpose of the first and last row of the matrix $\boldsymbol{V}^n$ respectively, and consider the $r\times r$ positive definite matrix $\boldsymbol{Z}=\boldsymbol{V}^{n\,H}\boldsymbol{V}^n$.

[16]To compute $\boldsymbol{\nabla}\mathcal{J}(\boldsymbol{U})$, the following derivation rule has to be applied: if $\boldsymbol{M}$ is a constant $n\times r$ matrix, $\boldsymbol{\nabla}\,\mathrm{trace}\left(\Re\left(\boldsymbol{U}^H\boldsymbol{M}\right)\right)=\boldsymbol{M}$. As a consequence, if $\boldsymbol{C}$ is a constant $n\times n$ hermitian matrix, $\boldsymbol{\nabla}\,\mathrm{trace}\left(\boldsymbol{U}^H\boldsymbol{C}\boldsymbol{U}\right)=2\boldsymbol{C}\boldsymbol{U}$.

Corollary 11. *Let $\boldsymbol{J}(\varepsilon)=\boldsymbol{G}^{-1}\boldsymbol{\Phi}(\varepsilon)\boldsymbol{G}$. The function $\varepsilon\mapsto\boldsymbol{J}(\varepsilon)$ is $\mathcal{C}^\infty$ in a neighborhood of $\varepsilon=0$, and admits the expansion*

$$\boxed{\boldsymbol{J}(\varepsilon)=\boldsymbol{J}+\varepsilon\Delta\boldsymbol{J}+O(\varepsilon^2)} \tag{37}$$

where the first order perturbation $\Delta\boldsymbol{J}=\boldsymbol{G}^{-1}\Delta\boldsymbol{\Phi}\,\boldsymbol{G}$ can be written in the following form (where $\boldsymbol{A}'=\boldsymbol{G}^{-1}\boldsymbol{A}\,\boldsymbol{G}$):

$$\boxed{\Delta\boldsymbol{J}=\Delta\boldsymbol{J}^\perp+\boldsymbol{A}'\boldsymbol{J}-\boldsymbol{J}\boldsymbol{A}'.} \tag{38}$$

Moreover, the matrix $\Delta\boldsymbol{J}^\perp=\boldsymbol{G}^{-1}\Delta\boldsymbol{\Phi}^\perp\,\boldsymbol{G}$ has rank 2:

$$\Delta\boldsymbol{J}^\perp=\boldsymbol{v}'_\downarrow{\boldsymbol{e}_\downarrow}^H\Delta\boldsymbol{S}\,{\boldsymbol{V}^l}^{\dagger T}\boldsymbol{D}^{-1}\boldsymbol{J}-\boldsymbol{J}\,\boldsymbol{v}'_\uparrow{\boldsymbol{e}_\uparrow}^H\Delta\boldsymbol{S}\,{\boldsymbol{V}^l}^{\dagger T}\boldsymbol{D}^{-1} \tag{39}$$

where the $l\times r$ Pascal-Vandermonde matrix $\boldsymbol{V}^l$ and the $r\times r$ block-diagonal matrix $\boldsymbol{D}$ were introduced in proposition 1, the r-dimensional vectors $\boldsymbol{v}'_\downarrow$ and $\boldsymbol{v}'_\uparrow$ are defined as

$$\begin{aligned}\boldsymbol{v}'_\downarrow &= \frac{\boldsymbol{Z}^{-1}\boldsymbol{v}_\downarrow}{1-\boldsymbol{v}_\downarrow^H\boldsymbol{Z}^{-1}\boldsymbol{v}_\downarrow}\\ \boldsymbol{v}'_\uparrow &= \frac{\boldsymbol{Z}^{-1}\boldsymbol{v}_\uparrow}{1-\boldsymbol{v}_\uparrow^H\boldsymbol{Z}^{-1}\boldsymbol{v}_\uparrow},\end{aligned}$$

and the n-dimensional vectors $\boldsymbol{e}_\downarrow$ and $\boldsymbol{e}_\uparrow$ are defined as

$$\begin{aligned}\boldsymbol{e}_\downarrow &= [0\dots0,1]^T-\boldsymbol{V}^n\boldsymbol{Z}^{-1}\boldsymbol{v}_\downarrow\\ \boldsymbol{e}_\uparrow &= [1,0\dots0]^T-\boldsymbol{V}^n\boldsymbol{Z}^{-1}\boldsymbol{v}_\uparrow.\end{aligned}$$

Corollary 11 is derived from proposition 10 by means of the basis change (7).

3) Perturbation of the poles: Finally, we focus in this section on the perturbation of the poles. Theoretically, they are obtained by computing the Jordan form of the spectral matrix $\boldsymbol{\Phi}$. In practice, contrary to $\boldsymbol{\Phi}$, the perturbed spectral matrix does not have multiple eigenvalues: multiple poles are scattered into several single eigenvalues.

More precisely, it was shown in [2] that if $M_k>1$:

- the first order perturbation of the pole z_k is homogeneous and isotropic, so that the M_k perturbed eigenvalues $z_k^{(m)}(\varepsilon)$ form the vertices of a regular polynomial of order M_k in the complex plan;
- the perturbation of the scattered eigenvalues is of order $\varepsilon^{\frac{1}{M_k}}$, which suggests that multiple poles are more sensitive to perturbations than single poles.

In fact it is possible to overcome this problem by no longer considering the eigenvalues $z_k^{(m)}(\varepsilon)$ as M_k distinct estimators of the same pole z_k, but rather by forming a single estimator of this pole by averaging the $z_k^{(m)}(\varepsilon)$. The following proposition shows that the arithmetic mean $z_k(\varepsilon)$ of the scattered eigenvalues admits a series expansion.

Proposition 12 (Perturbation of the poles)**.** *Let $z_k(\varepsilon)=\frac{1}{M_k}\sum_{m=0}^{M_k-1}z_k^{(m)}(\varepsilon)\ \forall k\in\{0\dots K-1\}$. Then the function $\varepsilon\mapsto z_k(\varepsilon)$ is $\mathcal{C}^\infty$ and admits the first order expansion*

$$\boxed{z_k(\varepsilon)=z_k+\varepsilon\Delta z_k+O\left(\varepsilon^2\right)} \tag{40}$$

where Δz_k is the complex number

$$\Delta z_k=\frac{1}{M_k}\,\mathrm{trace}\left(\Delta\boldsymbol{J}_k^\perp\right). \tag{41}$$

Here, $\Delta\boldsymbol{J}_k^\perp$ is the sub-block of dimension $M_k\times M_k$ extracted from the matrix $\Delta\boldsymbol{J}^\perp$, which is associated to z_k[17]. The complex number Δz_k can also be written in the form

$$\boxed{\Delta z_k=\frac{1}{M_k\,\alpha_k^{(M_k-1)}}\left({\boldsymbol{e}_\downarrow}^H\Delta\boldsymbol{S}\,\boldsymbol{f}_{\downarrow k}-{\boldsymbol{e}_\uparrow}^H\Delta\boldsymbol{S}\,\boldsymbol{f}_{\uparrow k}\right).} \tag{42}$$

[17]This corresponds to the rows and columns of indices $\sum_{k'=0}^{k-1}M_{k'}$ to $\sum_{k'=0}^{k}M_{k'}-1$.

The l-dimensional vectors $\boldsymbol{f}_{\downarrow k}$ and $\boldsymbol{f}_{\uparrow k}$ are defined[18] *as*

$$\left\{ \begin{array}{ll} \boldsymbol{f}_{\downarrow k} & = \alpha_k^{(M_k-1)}\, \boldsymbol{V}^{l\,\dagger T} \boldsymbol{D}^{-1} \boldsymbol{J}\, \boldsymbol{v}'_{\downarrow k} \\ \boldsymbol{f}_{\uparrow k} & = \alpha_k^{(M_k-1)}\, \boldsymbol{V}^{l\,\dagger T} \boldsymbol{D}^{-1} \boldsymbol{J}\, \boldsymbol{v}'_{\uparrow k}, \end{array} \right. \tag{43}$$

where $\boldsymbol{v}'_{\downarrow k}$ and $\boldsymbol{v}'_{\uparrow k}$ are the r-dimensional vectors whose coefficients are equal to those of $\boldsymbol{v}'_{\downarrow}$ and $\boldsymbol{v}'_{\uparrow}$ inside the k^{th} sub-block[19]*, and zero outside this sub-block.*

Equation (42) is derived by substituting equation (39) into equation (41). It can be noticed that the antihermitian matrix $\boldsymbol{A}$, introduced in proposition 9, is no longer involved in proposition 12. We can conclude that the performance of the generalized ESPRIT algorithm is not sensitive to the particular choice of the orthonormal basis $\boldsymbol{W}(\varepsilon)$ (at the first order).

Since the matrix $\Delta\boldsymbol{S}$ is Hankel and contains the samples of the PACE signal, the right member of equation (42) contains linear combinations of $\Delta s(\tau)$ for $\tau \in \{t-l+1 \dots t+n-1\}$. Therefore equation (42) can also be written as a scalar product:

$$\Delta z_k = \frac{1}{M_k\, \alpha_k^{(M_k-1)}}\, \boldsymbol{u}_k{}^H \Delta\boldsymbol{s} \tag{44}$$

where for all $\tau \in [0, \dots, n+l-2]$, the coefficient of index τ in $\boldsymbol{u}_k$ is[20]

$$u_k(\tau) = \sum_{\upsilon=0}^{n-1} \left(e_\downarrow(\upsilon) f^*_{\downarrow k}(\tau-\upsilon) - e_\uparrow(\upsilon) f^*_{\uparrow k}(\tau-\upsilon) \right) \mathbf{1}_{\{\tau-l<\upsilon\leq\tau\}}. \tag{45}$$

Equation (45) involves two convolution products, which are due to the Hankel matrix / vector products in equation (42).

Finally, proposition 3 in section III shows how the perturbation of the poles affects the frequencies and damping factors.

B. Perturbation of the amplitudes and phases

Below, we successively derive the following first order expansions:

- $\boldsymbol{V}^N(\varepsilon) = \boldsymbol{V}^N + \varepsilon\Delta\boldsymbol{V}^N + O(\varepsilon^2)$, where the expression of $\Delta\boldsymbol{V}^N$ is a linear function of the Δz_k's (lemma 13);
- $\boldsymbol{\alpha}(\varepsilon) = \boldsymbol{\alpha} + \varepsilon\Delta\boldsymbol{\alpha} + O(\varepsilon^2)$, where the expression of $\Delta\boldsymbol{\alpha}$ is a linear function of Δs (proposition 14);
- $\left\{ \begin{array}{l} a_k(\varepsilon) = a_k + \varepsilon\Delta a_k + O(\varepsilon^2) \\ \phi_k(\varepsilon) = \phi_k + \varepsilon\Delta\phi_k + O(\varepsilon^2) \end{array} \right.$ where the expressions of Δa_k and $\Delta\phi_k$ are functions of $\Delta\boldsymbol{\alpha}$ (corollary 4).

1) Perturbation of the Pascal-Vandermonde matrix:

Lemma 13 (Perturbation of the Pascal-Vandermonde matrix)**.** *Let $\boldsymbol{V}^N(\varepsilon)$ be the $N \times r$ Pascal-Vandermonde matrix associated to the estimated poles $\{z_0(\varepsilon), \dots, z_{K-1}(\varepsilon)\}$ defined in proposition 12. Then the function $\varepsilon \mapsto \boldsymbol{V}^N(\varepsilon)$ is $\mathcal{C}^\infty$ in the neighborhood of $\varepsilon = 0$, and admits the first order expansion:*

$$\boxed{\boldsymbol{V}^N(\varepsilon) = \boldsymbol{V}^N + \varepsilon\, \Delta\boldsymbol{V}^N + O(\varepsilon^2)} \tag{46}$$

where the matrix $\Delta\boldsymbol{V}^N$ can be written in the form

$$\boxed{\Delta\boldsymbol{V}^N = \overline{\boldsymbol{V}}^N \Delta\boldsymbol{Z}} \tag{47}$$

where $\overline{\boldsymbol{V}}^N$ is the $N \times (r+K)$ Pascal-Vandermonde matrix obtained by concatenating the generalized Pascal matrices $\overline{\boldsymbol{C}}_k^N = \boldsymbol{C}_{M_k+1}^N(z_k)$ of dimension $N \times (M_k+1)$, and

$$\Delta\boldsymbol{Z} = \mathrm{diag}(\Delta\boldsymbol{Z}_0, \dots, \Delta\boldsymbol{Z}_{K-1}) \tag{48}$$

is a $(r+K) \times r$ matrix whose diagonal blocks

$$\Delta\boldsymbol{Z}_k = \Delta z_k \times \left[\begin{array}{c} \mathbf{0}_{(1\times M_k)} \\ \mathrm{diag}(1, 2, \dots, M_k) \end{array} \right] \tag{49}$$

have dimension $(M_k+1) \times M_k$.

Lemma 13 is proved by substituting the first order expansion of the poles (40) into the expression of the Pascal-Vandermonde matrix, and by extracting the terms of order 1.

2) Perturbation of the amplitudes and phases: The perturbation of the complex-valued amplitudes can be derived from lemma 13. Let

$$\boldsymbol{s} = [s(t-l+1), \dots, s(t+n-1)]^T$$

be the N-dimensional vector containing the samples of the PACE signal. For all $\varepsilon \in \mathbb{R}$, the N-dimensional observed vector is $\boldsymbol{s}(\varepsilon) = \boldsymbol{s} + \varepsilon\, \Delta\boldsymbol{s}$.

Proposition 14 (Perturbation of the complex-valued amplitudes)**.** *The perturbed LS estimate defined in equation* (9) *can be written in the form*

$$\boldsymbol{\alpha}(\varepsilon) = \boldsymbol{V}^N(\varepsilon)^\dagger \boldsymbol{s}(\varepsilon). \tag{50}$$

Then the function $\varepsilon \mapsto \boldsymbol{\alpha}(\varepsilon)$ is $\mathcal{C}^\infty$ in the neighborhood of $\varepsilon = 0$, and admits the first order expansion:

$$\boxed{\boldsymbol{\alpha}(\varepsilon) = \boldsymbol{\alpha} + \varepsilon\, \Delta\boldsymbol{\alpha} + O(\varepsilon^2).} \tag{51}$$

The r-dimensional vector $\Delta\boldsymbol{\alpha}$ satisfies

$$\boxed{\Delta\boldsymbol{\alpha} = \boldsymbol{B}^H \Delta\boldsymbol{s}} \tag{52}$$

where the $r \times N$ matrix

$$\boldsymbol{B}^H = \boldsymbol{V}^{N\,\dagger} \left(\boldsymbol{I}_N - \overline{\boldsymbol{V}}^N \left[\begin{array}{c} \boldsymbol{B}_0 \\ \vdots \\ \boldsymbol{B}_{K-1} \end{array} \right] \right) \tag{53}$$

involves the $(M_k+1) \times N$ matrices $\boldsymbol{B}_k$, of rank 1:

$$\boldsymbol{B}_k = \left[0, \frac{1}{M_k} \frac{\alpha_k^{(0)}}{\alpha_k^{(M_k-1)}}, \dots, \frac{M_k-1}{M_k} \frac{\alpha_k^{(M_k-2)}}{\alpha_k^{(M_k-1)}}, 1 \right]^T \boldsymbol{u}_k^H. \tag{54}$$

Equations (51) to (54) are obtained by substituting equations (46) to (49) into the first order expansion of equation (50).

Finally, proposition 4 in section III shows how the perturbation of the complex-valued amplitudes influences the real-valued amplitudes and phases. The vector $\boldsymbol{b}_k^{(m)}$ is the column of $\boldsymbol{B}$ associated to the pole z_k at index m, *i.e.* the column of index $m + \sum_{k'=0}^{k-1} M_{k'}$. The derivation of equation (14) from equation (52) is straightforward.

REFERENCES

[1] R. Roy, A. Paulraj, and T. Kailath, "ESPRIT–A subspace rotation approach to estimation of parameters of cisoids in noise," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 34, no. 5, pp. 1340–1342, Oct. 1986.

[2] R. Badeau, B. David, and G. Richard, "High resolution spectral analysis of mixtures of complex exponentials modulated by polynomials," *IEEE Trans. Signal Processing*, vol. 54, no. 4, pp. 1341–1350, Apr. 2006.

[3] R. A. Serway and J. W. Jewett, *Physics for Scientists and Engineers*. Brooks/Cole, 2003.

[4] B. Fischer and A. Medvedev, "$\mathbf{L}^2$ Time Delay Estimation by Means of Laguerre Functions," in *Proceedings of the American Control Conference*, San Diego, California, USA, June 1999.

[5] A. M. Sabatini, "Correlation Receivers Using Laguerre Filter Banks for Modelling Narrowband Ultrasonic Echoes and Estimating Their Time-of-Flights," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 44, no. 6, pp. 1253–1263, Nov. 1997.

[18] It can be verified that if $M_k = 1$, the vectors $\boldsymbol{f}_{\downarrow k}$ and $\boldsymbol{f}_{\uparrow k}$ do not depend on any complex amplitude.

[19] This corresponds to the coefficients of indices $\sum_{k'=0}^{k-1} M_{k'}$ to $\sum_{k'=0}^{k} M_{k'} - 1$.

[20] Here the function $\mathbf{1}_{(.)}$ is one if its argument is true and zero otherwise.

[6] O. V. Ivanova, L. Marcu, and M. C. K. Khoo, "A Nonparametric Method for Analysis of Fluorescence Emission in Combined Time and Wavelength Dimensions," *Annals of Biomedical Engineering*, vol. 33, no. 4, pp. 531–544, Apr. 2005.

[7] T. W. Hänsch, A. L. Schawlow, and G. W. Series, "The Spectrum of Atomic Hydrogen," *Sci. Amer.*, vol. 240, no. 94, pp. 531–544, Mar. 1979.

[8] P. W. Milonni and J. H. Eberly, *Lasers*, ser. Wiley Series in Pure and Applied Optics. New York, USA: Wiley-Interscience, 1988.

[9] D. Filipovic, "Exponential-polynomial families and the term structure of interest rates," *Bernoulli*, vol. 6, no. 6, pp. 1081–1107, 2000.

[10] V. Slivinskas, M. Radavicius, and V. Simonyte, "Cramér-rao bound for the estimates of frequencies and damping factors of quasipolynomials in noise," Report UPTEC 92022R, Institute of Technology, Uppsala University, Sweden, Tech. Rep., Feb. 1992.

[11] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1993.

[12] R. Badeau, B. David, and G. Richard, "Cramér-Rao bounds for multiple poles and coefficients of quasipolynomials in colored noise," submitted for publication.

[13] A. Kot, S. Parthasarathy, D. Tufts, and R. Vaccaro, "The statistical performance of state-variable balancing and Prony's method in parameter estimation," in *Proc. of ICASSP'87*, vol. 12, Apr. 1987, pp. 1549–1552.

[14] P. Stoica and A. Nehorai, "Study of the statistical performance of the Pisarenko harmonic decomposition method," *IEE Proceedings Radar and Signal Processing*, vol. 135, no. 2, pp. 161–168, Apr. 1988.

[15] G. M. Riche de Prony, "Essai expérimental et analytique sur les lois de la dilatabilité de fluides élastiques et sur celles de la force expansive de la vapeur de l'eau et de la vapeur de l'alcool à différentes températures," *Journal de l'école polytechnique*, vol. 1, no. 22, pp. 24–76, 1795, in French.

[16] V. F. Pisarenko, "The retrieval of harmonics from a covariance function," *Geophysical J. Royal Astron. Soc.*, vol. 33, pp. 347–366, 1973.

[17] R. Kumaresan and D. W. Tufts, "Estimating the parameters of exponentially damped sinusoids and pole-zero modeling in noise," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 30, no. 6, pp. 833–840, Dec. 1982.

[18] R. O. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. Antennas Propagat.*, vol. 34, no. 3, pp. 276–280, Mar. 1986.

[19] Y. Hua and T. K. Sarkar, "Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 38, no. 5, pp. 814–824, May 1990.

[20] B. Porat and B. Friedlander, "On the accuracy of the Kumaresan-Tufts method for estimating complex damped exponentials," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 35, no. 2, pp. 231–235, Feb. 1987.

[21] Y. Hua and T. K. Sarkar, "Perturbation analysis of TK method for harmonic retrieval problems," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 36, pp. 228–240, Feb. 1988.

[22] P. Stoica and T. Söderström, "Statistical Analysis of MUSIC and Subspace Rotation Estimates of Sinusoidal Frequencies," *IEEE Trans. Signal Processing*, vol. 39, pp. 1836–1847, Aug. 1991.

[23] A. Eriksson, P. Stoica, and T. Soderstrom, "Second-order properties of MUSIC and ESPRIT estimates of sinusoidal frequencies in high SNR scenarios," *IEE Proceedings on Radar, Sonar and Navigation*, vol. 140, no. 4, pp. 266–272, Aug. 1993.

[24] Y. Hua and T. K. Sarkar, "On SVD for estimating generalized eigenvalues of singular matrix pencil in noise," *IEEE Trans. Signal Processing*, vol. 39, no. 4, pp. 892–900, Apr. 1991.

[25] S. Roman, *The Umbral Calculus*. New York: Academic Press, 1984, §1.2: The Lower Factorial Polynomial.

[26] R. L. Graham, D. E. Knuth, and O. Patashnik, *Concrete Mathematics: A Foundation for Computer Science*, 2nd ed. Reading, MA: Addison-Wesley, 1994.

[27] T. Kailath, *Linear systems*. Englewood Cliffs, New Jersey, USA: Prenctice-Hall, 1980.

[28] R. A. Horn and C. R. Johnson, *Matrix analysis*. Cambridge University Press, 1985.

[29] R. Badeau, G. Richard, and B. David, "Sliding window adaptive SVD algorithms," *IEEE Trans. Signal Processing*, vol. 52, no. 1, pp. 1–10, Jan. 2004.

[30] R. Badeau, B. David, and G. Richard, "Fast Approximated Power Iteration Subspace Tracking," *IEEE Trans. Signal Processing*, vol. 53, no. 8, pp. 2931–2941, Aug. 2005.

[31] ——, "Yet Another Subspace Tracker," in *Proc. of ICASSP'05*, vol. 4. IEEE, Mar. 2005, pp. 329–332.

[32] H. Zeiger and A. McEwen, "Approximate linear realizations of given dimension via Ho's algorithm," *IEEE Trans. Automat. Contr.*, vol. 19, no. 2, p. 153, Apr. 1974.

[33] R. Badeau, B. David, and G. Richard, "A new perturbation analysis for signal enumeration in rotational invariance techniques," *IEEE Trans. Signal Processing*, vol. 54, no. 2, pp. 450–458, Feb. 2006.

[34] Z. Zeng, "On ill-conditioned eigenvalues, multiple roots of polynomials, and their accurate computations," MSRI Preprint No. 1998-048, Northeastern Illinois University, Chicago, USA, Tech. Rep., 1998.

[35] M. Wax and T. Kailath, "Detection of signals by information theoretic criteria," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 33, no. 2, pp. 387–392, Apr. 1985.

[36] K. Abed-Meraim, P. Loubaton, and E. Moulines, "A Subspace Algorithm for Certain Blind Identification Problems," *IEEE Trans. Inform. Theory*, vol. 43, no. 2, pp. 499–511, Mar. 1997.

[37] B. Yang, "Projection Approximation Subspace Tracking," *IEEE Trans. Signal Processing*, vol. 44, no. 1, pp. 95–107, Jan. 1995.

A new perturbation analysis for signal enumeration in rotational invariance techniques

Roland Badeau, *Member, IEEE,* Bertrand David, and Gaël Richard, *Member, IEEE*

***Abstract*—The ESPRIT algorithm is a subspace-based high resolution method used in source localization and spectral analysis, which provides very accurate estimates of the signal parameters. However, the underlying theory assumes a known model order, which is usually not the case in many applications. In particular, it is well known that under-evaluating the model order biases the estimation. In this paper, we analyze the perturbation induced by an erroneous model order, and we present an error bound for the estimated parameters. Based on this theoretical framework, we propose a new method for selecting an appropriate modeling order, which consists in minimizing the error bound. This approach is applied to both synthetic and musical signals and its performance is compared to that of existing methods, such as the Information Theoretic Criteria (ITC).**

***Index Terms*—ESPRIT, rotational invariance, model order selection, signal enumeration, perturbation theory.**

I. Introduction

ESTIMATING a line spectrum is an important task for many applications, such as speech signal analysis and synthesis [1] and musical signal modification [2]. Although the Fourier transform remains a prominent tool for frequency estimation, the ESPRIT algorithm [3] overcomes the resolution limit of the Fourier analysis and provides straight estimates of the model parameters. This method relies on the rotational invariance property of the signal subspace spanned by the sinusoids. Its drawback is that the model order is supposed to be known, which is not the case in practice.

Many methods were proposed in the literature for estimating the number of sinusoids in white noise. The most classical ones are the maximum likelihood method [4] and the Information Theoretic Criteria (ITC) [5], among which the Akaike Information Criterion (AIC) [6] and the Maximum Description Length (MDL) by Schwartz [7] and Rissanen [8]. Another consistent procedure in the framework of the ITC is the Efficient Detection Criteria (EDC) [9], which proves to be robust to non-additive white noise [10]. The various ITC rely on the similarity of the eigenvalues within the noise subspace, and not on the existence of a gap between the signal and noise subspaces [11]. A criterion for model order selection based on this gap, which looks for a *maximally stable* decomposition, has been developed in [12]. Other methods proposed for model order selection include the Wishart matrices [13] and the cross-validation [14] approaches. However in the presence of a correlated noise, these methods tend to overestimate the model order. Consequently, specific methods have been designed to address the colored noise case, including new information theoretic criteria [15], [16], a technique based on a band noise covariance matrix model [17], and a maximum a posteriori criterion [18].

In other respects, we show in this paper how applying the ESPRIT high resolution method with an erroneous model order perturbs the estimation of the sinusoids. Note that in the literature, most papers rather focus on the perturbations induced by the additive noise. For example, the asymptotic second-order properties of ESPRIT were studied in the Direction of Arrival (DOA) [19] and in the frequency estimation [20] context, for a finite Signal to Noise Ratio (SNR). Reciprocally, a similar study was carried out for a finite number of data samples under a large SNR hypothesis [21]. In [22], a class of modeling errors were analyzed. However, to the best of our knowledge, no perturbation analysis of the ESPRIT estimates in the case of erroneous modeling order has ever been published (in the case of the MUSIC algorithm, a study is available in [23]).

Note that all the above mentioned performance analyses of the ESPRIT algorithm, as well as the subspace perturbation approach in [24], rely on first order approximations. Conversely, we present in this paper error bounds for the frequency estimates, which are derived without approximation and which can be easily computed. Furthermore, they are more precise than those presented in [25]. Based on this result, we propose a new model order selection method which consists in minimizing the perturbation. Contrary to the other methods proposed in the literature, which select the model order by analyzing the spectral properties of the additive noise, our approach focuses on the signal itself. Although it relies on a noiseless model, we observed that it outperforms the classical ITC, even in low SNR scenarios.

The paper is organized as follows. Section II summarizes the principles of the ESPRIT high resolution method. In section III, the perturbation of the poles induced by an erroneous model order is analyzed. Then our new model order selection method, referred to as the ESTER method, is introduced in section IV, where a fast implementation is proposed. In section V, the relevance of our criterion as an error bound is examined, and the performance of the ESTER method is compared to that of some existing methods. Finally, the main conclusions of this paper are summarized in section VI.

II. The ESPRIT method for spectral analysis

The noiseless Exponential Sinusoidal Model (ESM), also known as the Exponentially Damped Sinusoidal (EDS)

Roland Badeau, Bertrand David, and Gaël Richard are with the Department of Signal and Image Processing, Ecole Nationale Supérieure des Télécommunications (ENST), Paris, France. E-mail: [roland.badeau, bertrand.david, gael.richard]@enst.fr.

model [26], defines the discrete signal as

$$x(t) = \sum_{k=1}^{r} \alpha_k z_k^t$$

where r is the order of the model, $\alpha_k \in \mathbb{C} - \{0\}$ are the complex amplitudes, and $z_k \in \mathbb{C} - \{0\}$ are the complex poles.

Let $n > r$, $l > r$, and define the $n \times l$ Hankel data matrix

$$\boldsymbol{X} = \begin{bmatrix} x(0) & x(1) & \cdots & x(l-1) \\ x(1) & x(2) & \cdots & x(l) \\ \vdots & \vdots & \cdots & \vdots \\ x(n-1) & x(n) & \cdots & x(n+l-2) \end{bmatrix}.$$

which involves $N = n + l - 1$ samples of the signal. This matrix can be factorized in the form $\boldsymbol{X} = \boldsymbol{V}^n \, \boldsymbol{A} \, \boldsymbol{V}^{l^T}$, where $\boldsymbol{V}^n$ is a $n \times r$ Vandermonde matrix

$$\boldsymbol{V}^n = \begin{bmatrix} 1 & \dots & 1 \\ z_1 & \dots & z_r \\ \vdots & \vdots & \vdots \\ z_1^{n-1} & \dots & z_r^{n-1} \end{bmatrix},$$

$\boldsymbol{A} = \mathrm{diag}(\alpha_1 \dots \alpha_r)$, and $\boldsymbol{V}^l$ is a $l \times r$ Vandermonde matrix [27]. If the r poles $\{z_1, \dots, z_r\}$ are distinct, $\boldsymbol{X}$ has a r-dimensional range space, spanned by the full-rank matrix $\boldsymbol{V}^n$. This range space fully characterizes the signal poles. It is thus referred to as the *signal subspace*.

Let $\boldsymbol{V}^n_{\downarrow}$ be the matrix extracted from $\boldsymbol{V}^n$ by deleting the last row. Similarly, let $\boldsymbol{V}^n_{\uparrow}$ be the matrix extracted from $\boldsymbol{V}^n$ by deleting the first row. Then the Vandermonde matrix $\boldsymbol{V}^n$ satisfies the so-called *rotational invariance property*:

$$\boldsymbol{V}^n_{\uparrow} = \boldsymbol{V}^n_{\downarrow} \, \boldsymbol{D} \tag{1}$$

where $\boldsymbol{D} = \mathrm{diag}(z_1, \dots, z_r)$.

In practice, $\boldsymbol{V}^n$ is unknown, but the signal subspace can be obtained by computing the singular value decomposition (SVD) of $\boldsymbol{X}$ (or via subspace tracking techniques [28]–[30] in an adaptive context). Indeed, if $\{\boldsymbol{w}(1), \dots, \boldsymbol{w}(n)\}$ are the left singular vectors associated to the singular values $\sigma_1 \geq \dots \geq \sigma_n \geq 0$ sorted in decreasing order, then the signal subspace is spanned by the $n \times r$ orthonormal matrix $\boldsymbol{W}(r) = [\boldsymbol{w}(1), \dots, \boldsymbol{w}(r)]$ (the $n - r$ last singular values being equal to 0). Since $\boldsymbol{W}(r)$ and $\boldsymbol{V}^n$ span the same subspace, there is a non-singular matrix $\boldsymbol{G}$ of dimension $r \times r$ such that

$$\boldsymbol{V}^n = \boldsymbol{W}(r) \, \boldsymbol{G}. \tag{2}$$

By deleting the last row in equation (2) we obtain

$$\boldsymbol{W}_{\downarrow}(r) = \boldsymbol{V}^n_{\downarrow} \, \boldsymbol{G}^{-1}. \tag{3}$$

Similarly, deleting the first row in equation (2) and substituting equation (1) yields

$$\boldsymbol{W}_{\uparrow}(r) = \boldsymbol{V}^n_{\downarrow} \, \boldsymbol{D} \, \boldsymbol{G}^{-1}. \tag{4}$$

Substituting equation (3) into equation (4) leads to the rotational invariance property of the matrix $\boldsymbol{W}(r)$:

$$\boldsymbol{W}_{\uparrow}(r) = \boldsymbol{W}_{\downarrow}(r) \, \boldsymbol{\Phi}(r)$$

where $\boldsymbol{\Phi}(r)$ is defined by its eigenvalue decomposition: $\boldsymbol{\Phi}(r) \triangleq \boldsymbol{G} \, \boldsymbol{D} \, \boldsymbol{G}^{-1}$. Finally, the ESPRIT algorithm [3] consists of the following steps:

- computing $\boldsymbol{W}(r)$,
- computing $\boldsymbol{\Phi}(r) = \boldsymbol{W}_{\downarrow}(r)^{\dagger} \boldsymbol{W}_{\uparrow}(r)$ (where the symbol $\dagger$ denotes the Moore-Penrose pseudo-inverse),
- extracting the poles z_k as the eigenvalues of $\boldsymbol{\Phi}(r)$.

III. Impact of an erroneous model order

In practice, the model order r is unknown. We assume below that the ESPRIT algorithm is applied with an erroneous model order p and we focus on how the estimation of the poles is affected. For all $p \in \{1 \dots n\}$, let $\boldsymbol{W}(p) \triangleq [\boldsymbol{w}(1), \dots, \boldsymbol{w}(p)]$ and

$$\boldsymbol{\Phi}(p) \triangleq \boldsymbol{W}_{\downarrow}(p)^{\dagger} \boldsymbol{W}_{\uparrow}(p). \tag{5}$$

The estimated poles are defined as the eigenvalues of $\boldsymbol{\Phi}(p)$.

A. Over-estimation of the model order

If $p \geq r$, the following proposition shows that the r true poles belong to the whole set of eigenvalues of $\boldsymbol{\Phi}(r)$.

Proposition III.1. *Suppose that $r \leq p < n$ and $\boldsymbol{W}_{\downarrow}(p)$ is full rank. Then $\forall k \in \{1, \dots, r\}$, z_k is an eigenvalue of $\boldsymbol{\Phi}(p)$.*

Proof: Let $\boldsymbol{v}_k$ be the right eigenvector of $\boldsymbol{\Phi}(r)$ associated to the eigenvalue z_k and consider the p-dimensional vector $\overline{\boldsymbol{v}}_k \triangleq \begin{bmatrix} \boldsymbol{v}_k \\ \boldsymbol{0} \end{bmatrix}$. Note that $\boldsymbol{W}(p) \, \overline{\boldsymbol{v}}_k = \boldsymbol{W}(r) \, \boldsymbol{v}_k$. Consequently, $\boldsymbol{W}_{\uparrow}(p) \overline{\boldsymbol{v}}_k = \boldsymbol{W}_{\uparrow}(r) \, \boldsymbol{v}_k = \boldsymbol{W}_{\downarrow}(r) \, \boldsymbol{\Phi}(r) \, \boldsymbol{v}_k = z_k \, \boldsymbol{W}_{\downarrow}(r) \, \boldsymbol{v}_k = z_k \, \boldsymbol{W}_{\downarrow}(p) \overline{\boldsymbol{v}}_k$. Since $\boldsymbol{W}_{\downarrow}(p)$ is full rank, left multiplying the previous equality by $\boldsymbol{W}_{\downarrow}(p)^{\dagger}$ yields $\boldsymbol{\Phi}(p) \, \overline{\boldsymbol{v}}_k = z_k \, \overline{\boldsymbol{v}}_k$. ∎

B. Under-estimation of the model order

If $p < r$, it is well known that the eigenvalues of $\boldsymbol{\Phi}(p)$ do not match the poles in the general case. More precisely, let $\widehat{z}$ be an eigenvalue of $\boldsymbol{\Phi}(p)$. In this section, it will be shown that $\widehat{z}$ approximates one of the eigenvalues of $\boldsymbol{\Phi}(r)$, and that an error bound can be easily computed. First, we need to define the *upper condition number* of the signal subspace[1]:

$$\kappa_2 = \inf_{\boldsymbol{\Lambda} \in \mathrm{diag}(\mathbb{R}^{+r})} \frac{\sigma_{\max}(\boldsymbol{V}^n \, \boldsymbol{\Lambda})}{\sigma_{\min}(\boldsymbol{V}^n_{\downarrow} \, \boldsymbol{\Lambda})} \tag{6}$$

where $\sigma_{\max}(.)$ denotes the largest singular value of a matrix, and $\sigma_{\min}(.)$ denotes the smallest one. This condition number characterizes the noiseless signal itself, and does not depend on p. It is an unknown constant for our problem, which does not need to be calculated. It is involved in the following theorem, whose proof can be found in the appendix.

[1] In [25], κ_2 was defined as the upper condition number of the *Vandermonde matrix* $\boldsymbol{V}^n$, equal to $\frac{\sigma_{\max}(\boldsymbol{V}^n)}{\sigma_{\min}(\boldsymbol{V}^n_{\downarrow})}$. The new definition of κ_2 in equation (6) yields better error bounds, due to the presence of the infimum.

Notation: $\mathbb{R}^+$ denotes the set of all positive real numbers.

Theorem III.2 (a priori error bound[2])**.** *For all $\widehat{z} \in \mathbb{C}$, there is an eigenvalue z_k of $\boldsymbol{\Phi}(r)$ for which*

$$|\widehat{z} - z_k| \leq \kappa_2 \; \sigma_{\min}\left(\boldsymbol{W}_{\uparrow}(p) - \widehat{z}\,\boldsymbol{W}_{\downarrow}(p)\right). \tag{7}$$

Note that the bound $\sigma_{\min}\left(\boldsymbol{W}_{\uparrow}(p) - \widehat{z}\,\boldsymbol{W}_{\downarrow}(p)\right)$ can be computed without knowing the exact value of r. Corollary III.3, which follows from theorem III.2, has a certain similarity with the well known Bauer-Fike theorem [31, pp. 365], [32, pp. 321]. It gives an error bound valid for all the eigenvalues of $\boldsymbol{\Phi}(p)$. Again, this bound can be computed without knowing the exact value of r. It involves the spectral norm of a matrix (or 2-norm), also denoted $\|.\|_2$, defined as $\|\boldsymbol{M}\|_2 \triangleq \max_{\|\boldsymbol{u}\|_2=1} \|\boldsymbol{M}\,\boldsymbol{u}\|_2 = \sigma_{\max}(\boldsymbol{M})$.

Corollary III.3 (a posteriori error bound[3])**.** *For each eigenvalue $\widehat{z}$ of $\boldsymbol{\Phi}(p)$, there is an eigenvalue z_k of $\boldsymbol{\Phi}(r)$ for which*

$$|\widehat{z} - z_k| \leq \kappa_2 \; \|\boldsymbol{E}(p)\|_2 \tag{8}$$

where

$$\boldsymbol{E}(p) = \boldsymbol{W}_{\uparrow}(p) - \boldsymbol{W}_{\downarrow}(p)\,\boldsymbol{\Phi}(p). \tag{9}$$

Proof: Let $\widehat{z}$ be an eigenvalue of $\boldsymbol{\Phi}(p)$ and $\widehat{\boldsymbol{v}}$ a unitary eigenvector associated with $\widehat{z}$. Let $\boldsymbol{e}(p) \triangleq \left(\boldsymbol{W}_{\uparrow}(p) - \widehat{z}\,\boldsymbol{W}_{\downarrow}(p)\right)\widehat{\boldsymbol{v}}$. Since $\widehat{\boldsymbol{v}}$ is unitary, $\sigma_{\min}\left(\boldsymbol{W}_{\uparrow}(p) - \widehat{z}\,\boldsymbol{W}_{\downarrow}(p)\right) \leq \|\boldsymbol{e}(p)\|_2$. In other respects, $\boldsymbol{e}(p) = \left(\boldsymbol{W}_{\uparrow}(p) - \boldsymbol{W}_{\downarrow}(p)\,\boldsymbol{\Phi}(p)\right)\widehat{\boldsymbol{v}}$, thus $\|\boldsymbol{e}(p)\|_2 \leq \|\boldsymbol{E}(p)\|_2$. Consequently,

$$\sigma_{\min}\left(\boldsymbol{W}_{\uparrow}(p) - \widehat{z}\,\boldsymbol{W}_{\downarrow}(p)\right) \leq \|\boldsymbol{E}(p)\|_2. \tag{10}$$

Finally, substituting equation (10) into (7) yields equation (8). ■

Remark. *Let $p < n-1$. We know that if $p = r$, $\boldsymbol{E}(p) = 0$. Conversely, if $\boldsymbol{E}(p) = 0$, then the matrices $\boldsymbol{W}_{\downarrow}(p)$ and $\boldsymbol{W}_{\uparrow}(p)$ span the same subspace, which means that the rotational invariance property is satisfied at order p. Thus p complex exponentials can be extracted from the observed signal, and the corresponding complex poles can be estimated by means of the ESPRIT algorithm. Since the signal does not contain more than r complex exponentials, we expect that $p \leq r$. The case $p < r$ can happen if the signal parameters satisfy some particular relationships. In practice, r is always the greatest value of p for which $\boldsymbol{E}(p) = 0$.*

IV. Selection of an appropriate modeling order based on the ESTimation ERror

The practical interest of corollary III.3 is that $\|\boldsymbol{E}(p)\|_2$ (which will be referred to as the *a posteriori error bound*) can be computed for all $p \in \{1 \ldots p_{\max}\}$, where $1 \leq p_{\max} < n-1$. If $p_{\max}$ happens to be lower than r, the a posteriori error bound gives a quantitative criterion for selecting an appropriate modeling order, such that the estimation error bound is minimum. If $p_{\max}$ happens to be greater than r, then r is the greatest value of $p \in \{1 \ldots p_{\max}\}$ for which the a posteriori error bound is zero. In any case, detecting the maxima of the inverse error function $J : p \mapsto \frac{1}{\|\boldsymbol{E}(p)\|_2{}^2}$ in the range $\{1 \ldots p_{\max}\}$ is a relevant approach for selecting the modeling order. Below, this function will be referred to as the ESTimation ERror (ESTER) criterion. In presence of noise, we observed that a robust way of selecting the modeling order consists in detecting the greatest value of p for which the function $J(p)$ reaches a local maximum which is greater than a fraction of its global maximum (typically one tenth of the global maximum). Examples of the function J are represented in section V. Proposition IV.1 shows that its values are in the interval $[1, +\infty[$ (the proof is given in the appendix).

Proposition IV.1. *For all $p \in \{1, \ldots, n\}$, $\|\boldsymbol{E}(p)\|_2 \leq 1$.*

Note that this ESTER criterion measures the rotational invariance of $\boldsymbol{W}(p)$, since by definition the rotational invariance property is satisfied exactly if $\|\boldsymbol{E}(p)\|_2 = 0$. The drawback of the ESTER method is that a direct implementation would lead to a very computationally demanding algorithm. First, the singular vectors $\boldsymbol{w}(p)$ have to be computed for all $p \in \{1 \ldots p_{\max}\}$, which requires $O(N \log_2(N) p_{\max} + n p_{\max}^2)$ operations, by means of a variant of the orthogonal iteration EVD algorithm presented in [26]. Then the matrix $\boldsymbol{E}(p)$ must be calculated for all $p \in \{1, \ldots, p_{\max}\}$. Such a computation would involve $3np^2$ MAC for each p, so that the overall complexity would be $np_{\max}^3$ MAC[4]. This computational cost is to be compared to that of the ITC criteria illustrated in section V. In particular, the complexity of the AIC, MDL [5] and EDC [9] criteria is linear in $p_{\max}$. However, contrary to the ESTER method, these ITC require the full SVD of the data matrix, whose complexity is $O(N^3)$. Besides, the complexity of the criteria proposed in [16] for addressing the colored noise case also have a complexity equal to $O(N^3)$. Consequently, the relative complexities of the ESTER method and the various ITC depend on $p_{\max}$.

To make the ESTER method faster, we developed a recursive implementation, presented in table I, which only involves $6np + O(p^2)$ MAC for each p, so that its overall complexity is $3np_{\max}^2 + O(p_{\max}^3)$ (plus the computation of the singular vectors). In particular, it can be noticed that computing the matrices $\boldsymbol{E}(p)$ for all $p \in \{1 \ldots p_{\max}\}$ in this way is not more computationally demanding than computing $\boldsymbol{E}(p_{\max})$ directly. Sections IV-A and IV-B present fast methods for computing $\boldsymbol{\Phi}(p)$ and $\boldsymbol{E}(p)$ recursively.

A. Recursive computation of $\boldsymbol{\Phi}(p)$

A direct calculation of $\boldsymbol{\Phi}(p)$ for all $p \in \{1 \ldots p_{\max}\}$ from equation (5) would involve $2np^2 + O(p^3)$ MAC for each p, and the overall complexity would be $\frac{2}{3}\,np_{\max}^3 + O(p_{\max}^4)$. This section aims at computing the $p \times p$ matrix $\boldsymbol{\Phi}(p)$ recursively, in

[2] In comparison, the a priori error bound proposed in [25] was equal to $\frac{\sigma_{\max}(\boldsymbol{V}^n)}{\sigma_{\min}(\boldsymbol{V}^n_{\downarrow})} \left\|\boldsymbol{W}_{\uparrow}(p)\,\widehat{\boldsymbol{v}} - \widehat{z}\,\boldsymbol{W}_{\downarrow}(p)\,\widehat{\boldsymbol{v}}\right\|_2$, where $\widehat{\boldsymbol{v}}$ was an arbitrary unitary vector. Note that the condition number κ_2 defined in equation (6) is lower than $\frac{\sigma_{\max}(\boldsymbol{V}^n)}{\sigma_{\min}(\boldsymbol{V}^n_{\downarrow})}$. Moreover $\sigma_{\min}\left(\boldsymbol{W}_{\uparrow}(p) - \widehat{z}\,\boldsymbol{W}_{\downarrow}(p)\right) \leq \left\|\boldsymbol{W}_{\uparrow}(p)\,\widehat{\boldsymbol{v}} - \widehat{z}\,\boldsymbol{W}_{\downarrow}(p)\,\widehat{\boldsymbol{v}}\right\|_2$ for all unitary vector $\widehat{\boldsymbol{v}}$. Thus the a posteriori error bound in equation (7) is lower than that proposed in [25].

[3] This a posteriori error bound is lower than that proposed in [25], because of the lower value of κ_2.

[4] In this paper, operations counts are expressed in terms of multiply / accumulate (MAC) operations. Note that $p_{\max}$ is supposed to be much lower than n.

order to reduce the complexity. Equation (5) can be rewritten in the form

$$\boldsymbol{\Phi}(p) = \boldsymbol{\Omega}(p)\,\boldsymbol{\Psi}(p) \tag{11}$$

where $\boldsymbol{\Omega}(p)$ and $\boldsymbol{\Psi}(p)$ are the $p \times p$ matrices

$$\boldsymbol{\Omega}(p) \triangleq \left(\boldsymbol{W}_{\downarrow}(p)^H\,\boldsymbol{W}_{\downarrow}(p)\right)^{-1} \tag{12}$$

$$\boldsymbol{\Psi}(p) \triangleq \boldsymbol{W}_{\downarrow}(p)^H\,\boldsymbol{W}_{\uparrow}(p). \tag{13}$$

The matrix $\boldsymbol{\Omega}(p)$ can be easily calculated. Indeed, since $\boldsymbol{W}(p)$ is orthonormal, $\boldsymbol{W}(p)^H\,\boldsymbol{W}(p) = \boldsymbol{I}_p$. In particular, this equation yields $\boldsymbol{W}_{\downarrow}(p)^H\,\boldsymbol{W}_{\downarrow}(p) = \boldsymbol{I}_p - \boldsymbol{\nu}(p)\,\boldsymbol{\nu}(p)^H$, where $\boldsymbol{\nu}(p)$ is the p-dimensional vector such that $\boldsymbol{\nu}(p)^H$ is the last row of $\boldsymbol{W}(p)$. Finally, the matrix inversion lemma [31, pp. 18-19] shows that

$$\boldsymbol{\Omega}(p) = \boldsymbol{I}_p + \frac{1}{1-\|\boldsymbol{\nu}(p)\|^2}\,\boldsymbol{\nu}(p)\,\boldsymbol{\nu}(p)^H. \tag{14}$$

Moreover, $\boldsymbol{\Psi}(p)$ can be recursively updated. Indeed, equation (13) yields

$$\boldsymbol{\Psi}(p) = \left[\begin{array}{c|c} \boldsymbol{\Psi}(p-1) & \boldsymbol{\psi}_r(p) \\ \hline \boldsymbol{\psi}_l(p)^H & \psi_{lr}(p) \end{array}\right] \tag{15}$$

where $\boldsymbol{\psi}_r(p) \triangleq \boldsymbol{W}_{\downarrow}(p-1)^H\,\boldsymbol{w}_{\uparrow}(p)$, $\boldsymbol{\psi}_l(p) \triangleq \boldsymbol{W}_{\uparrow}(p-1)^H\,\boldsymbol{w}_{\downarrow}(p)$ and $\psi_{lr}(p) \triangleq \boldsymbol{w}_{\downarrow}(p)^H\,\boldsymbol{w}_{\uparrow}(p)$. It can be noticed that the computation of $\boldsymbol{\Psi}(p)$ from $\boldsymbol{\Psi}(p-1)$ requires only $2np$ MAC.

Finally, $\boldsymbol{\Phi}(p)$ can be computed from $\boldsymbol{\Psi}(p)$. Indeed, substituting equation (14) into equation (11) yields

$$\boldsymbol{\Phi}(p) = \boldsymbol{\Psi}(p) + \frac{1}{1-\|\boldsymbol{\nu}(p)\|^2}\,\boldsymbol{\nu}(p)\,\boldsymbol{\varphi}(p)^H \tag{16}$$

where

$$\boldsymbol{\varphi}(p) \triangleq \boldsymbol{\Psi}(p)^H\boldsymbol{\nu}(p). \tag{17}$$

It can be noticed that the computation of $\boldsymbol{\Phi}(p)$ from $\boldsymbol{\Psi}(p)$ requires only $p^2 + O(p)$ MAC, plus the computation of $\boldsymbol{\varphi}(p)$. This last operation normally requires p^2 MAC, but lemma IV.2 suggests a recursive implementation, which involves only $O(p)$ MAC.

Lemma IV.2. *Let $\mu(p)$ be the complex number such that*

$$\boldsymbol{\nu}(p) = \left[\begin{array}{c} \boldsymbol{\nu}(p-1) \\ \hline \mu(p) \end{array}\right] \tag{18}$$

Then $\boldsymbol{\varphi}(p)$ satisfies the recursion

$$\boldsymbol{\varphi}(p) = \left[\begin{array}{c} \boldsymbol{\varphi}(p-1) + \mu(p)\,\boldsymbol{\psi}_l(p) \\ \hline \boldsymbol{\psi}_r(p)^H\boldsymbol{\nu}(p-1) + \mu(p)\,\psi_{lr}(p)^* \end{array}\right] \tag{19}$$

Proof: The assertion can be shown by substituting equations (15) an (18) into equation (17). ∎

Finally, the recursive computation of $\boldsymbol{\Phi}(p)$ consists in computing $\boldsymbol{\Psi}(p)$ from $\boldsymbol{\Psi}(p-1)$ with equation (15), then computing $\boldsymbol{\varphi}(p)$ from $\boldsymbol{\varphi}(p-1)$ with equation (19), then computing $\boldsymbol{\Phi}(p)$ from $\boldsymbol{\Psi}(p)$ with equation (16). This method requires $2np + O(p^2)$ MAC at each step. Therefore, its overall computational cost is $np^2_{\max} + O(p^3_{\max})$ MAC.

B. Recursive computation of $\boldsymbol{E}(p)$

Here we suppose that all the $\boldsymbol{\Phi}(p)$ have been computed. A direct calculation of $\boldsymbol{E}(p)$ for all $p \in \{1 \dots p_{\max}\}$ from equation (9) would involve np^2 MAC for each p, and the overall complexity would be $\frac{1}{3}\,np^3_{\max}$. This section aims at computing $\boldsymbol{E}(p)$ recursively, in order to reduce the complexity.

Substituting equation (16) into equation (9) shows that

$$\boldsymbol{E}(p) = \boldsymbol{\Xi}(p) - \frac{1}{1-\|\boldsymbol{\nu}(p)\|^2}\,\left(\boldsymbol{W}_{\downarrow}(p)\,\boldsymbol{\nu}(p)\right)\boldsymbol{\varphi}(p)^H \tag{20}$$

where

$$\boldsymbol{\Xi}(p) \triangleq \boldsymbol{W}_{\uparrow}(p) - \boldsymbol{W}_{\downarrow}(p)\,\boldsymbol{\Psi}(p). \tag{21}$$

Note that the computation of $\boldsymbol{E}(p)$ from $\boldsymbol{\Xi}(p)$ requires $2np$ MAC. Then substituting equation (15) into equation (21) yields a recursion for the $(n-1) \times p$ matrix $\boldsymbol{\Xi}(p)$:

$$\boldsymbol{\Xi}(p) = \left[\begin{array}{c|c} \boldsymbol{\Xi}(p-1) - \boldsymbol{w}_{\downarrow}(p)\,\boldsymbol{\psi}_l(p)^H & \boldsymbol{\xi}(p) \end{array}\right] \tag{22}$$

where $\boldsymbol{\xi}(p) \triangleq \boldsymbol{w}_{\uparrow}(p) - \boldsymbol{W}_{\downarrow}(p-1)\,\boldsymbol{\psi}_r(p) - \boldsymbol{w}_{\downarrow}(p)\,\psi_{lr}(p)$. The computation of $\boldsymbol{\Xi}(p)$ from $\boldsymbol{\Xi}(p-1)$ involves $2np$ MAC. Finally, the recursive computation of $\boldsymbol{E}(p)$ consists in computing $\boldsymbol{\Xi}(p)$ from $\boldsymbol{\Xi}(p-1)$ with equation (22), then computing $\boldsymbol{E}(p)$ from $\boldsymbol{\Xi}(p)$ with equation (20). This method requires $4np$ MAC at each step. Thus its overall computational cost is $2np^2_{\max}$ MAC. As a result, it can be noticed that computing the matrices $\boldsymbol{\Phi}(p)$ and $\boldsymbol{E}(p)$ for all $p \in \{1 \dots p_{\max}\}$ is not more computationally demanding than just computing them for $p = p_{\max}$. In both cases, the overall complexity is $3np^2_{\max} + O(p^3_{\max})$.

The complete pseudo code for computing $\boldsymbol{E}(p)$ for all $p \in \{1 \dots p_{\max}\}$ is presented in table I. Note that the calculation of the matrices $\boldsymbol{\Psi}(p)$ and $\boldsymbol{\Phi}(p)$ is not even required.

TABLE I
RECURSIVE COMPUTATION OF $\boldsymbol{E}(p)$

Initialization
compute $\boldsymbol{w}(p)$ for all $p = 1 \dots p_{\max}$, $\boldsymbol{\varphi}(0) = []$, $\boldsymbol{\Xi}(0) = []$

For $p = 1$ to $p_{\max}$

Update of the auxiliary matrix $\boldsymbol{\Psi}(p)$	**Cost**
$\boldsymbol{\psi}_r(p) = \boldsymbol{W}_{\downarrow}(p-1)^H\,\boldsymbol{w}_{\uparrow}(p)$	np
$\boldsymbol{\psi}_l(p) = \boldsymbol{W}_{\uparrow}(p-1)^H\,\boldsymbol{w}_{\downarrow}(p)$	np
$\psi_{lr}(p) = \boldsymbol{w}_{\downarrow}(p)^H\,\boldsymbol{w}_{\uparrow}(p)$	n
Update of the auxiliary matrix $\boldsymbol{\Xi}(p)$	
$\boldsymbol{\xi}(p) = \boldsymbol{w}_{\uparrow}(p) - \boldsymbol{W}_{\downarrow}(p-1)\,\boldsymbol{\psi}_r(p) - \boldsymbol{w}_{\downarrow}(p)\,\psi_{lr}(p)$	np
$\boldsymbol{\Xi}(p) = \left[\begin{array}{c\|c} \boldsymbol{\Xi}(p-1) - \boldsymbol{w}_{\downarrow}(p)\,\boldsymbol{\psi}_l(p)^H & \boldsymbol{\xi}(p) \end{array}\right]$	np
Computation of $\boldsymbol{E}(p)$ from $\boldsymbol{\Xi}(p)$	
$\boldsymbol{\varphi}(p) = \left[\begin{array}{c} \boldsymbol{\varphi}(p-1) + \mu(p)\,\boldsymbol{\psi}_l(p) \\ \hline \boldsymbol{\psi}_r(p)^H\boldsymbol{\nu}(p-1) + \mu(p)\,\psi_{lr}(p)^* \end{array}\right]$	$2p$
$\boldsymbol{E}(p) = \boldsymbol{\Xi}(p) - \frac{1}{1-\|\boldsymbol{\nu}(p)\|^2}\,\left(\boldsymbol{W}_{\downarrow}(p)\,\boldsymbol{\nu}(p)\right)\boldsymbol{\varphi}(p)^H$	$2np$

V. SIMULATION RESULTS

Section V-A illustrates the relevance of our error bounds. Then the ESTER method is applied to synthetic signals (sections V-B, V-C) and to a musical signal (section V-D).

A. Relevance of the a priori and a posteriori error bounds

In this section, the relevance of the a priori and a posteriori error bounds is illustrated. The test signal is a sum of $r = 20$ undamped complex exponentials of the same amplitude $\alpha_k = 1$, whose frequencies are randomly distributed in $[-\frac{1}{2}, \frac{1}{2}]$. The upper condition number of the signal subspace satisfies $1 \leq \kappa_2 \leq 1.05$. The singular vectors forming the matrices $\boldsymbol{W}(p)$ have been obtained by computing the SVD of a Hankel data matrix of $n = 512$ rows and $l = 512$ columns, containing the $n + l - 1 = 1023$ samples of the whole signal. For all $p \in \{1 \dots p_{\max} = r\}$, the eigenvalues $\left\{\widehat{z}_{(p,m)}\right\}_{m\in\{1\dots p\}}$ of the matrix $\boldsymbol{\Phi}(p)$ have been computed.

In figure 1-a, the solid line represents the eigenvalue errors obtained for $p = 7 < r$, *i.e.*

$$\left\{\min_{k\in\{1\dots r\}} |\widehat{z}_{(7,m)} - z_k|\right\}_{m\in\{1\dots 7\}}$$

sorted in increasing order. The dotted line represents the corresponding a priori error bounds, *i.e.*

$$\left\{\kappa_2\ \sigma_{\min}\left(\boldsymbol{W}_{\uparrow}(7) - \widehat{z}_{(7,m)}\,\boldsymbol{W}_{\downarrow}(7)\right)\right\}_{m\in\{1\dots 7\}}$$

It can be noticed that the dotted line is above the solid line (as expected), and more importantly that the variations of the dotted line follow those of the solid line, which suggests that the a priori error bound is relevant.

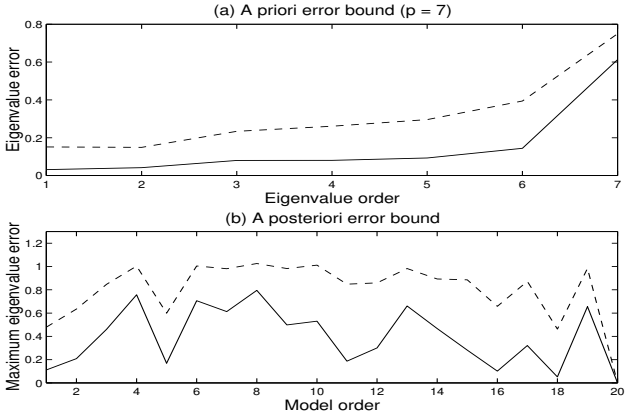

Fig. 1. A priori and a posteriori error bounds
(a) a priori error bound *vs.* eigenvalue error at order $p = 7$
(b) a posteriori error bound *vs.* maximum eigenvalue error as a function of p

In figure 1-b, the solid line represents the maximum eigenvalue error obtained for all modeling orders, *i.e.*

$$\max_{m\in\{1\dots p\}}\ \min_{k\in\{1\dots r\}} |\widehat{z}_{(p,m)} - z_k|$$

as a function of p. The dotted line represents the corresponding a posteriori error bounds, *i.e.* $\kappa_2\ \|\boldsymbol{E}(p)\|_2$ as a function of p. As for the a priori error bound, it can be noticed that the dotted line is above the solid line (as expected), and that the variations of the dotted line follow those of the solid line. This suggests that the a posteriori error bound is a relevant criterion for minimizing the bias of the estimated eigenvalues.

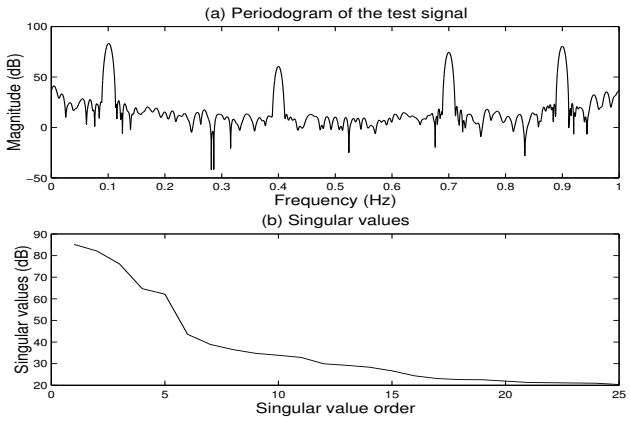

Fig. 2. Periodogram and singular values of a synthetic signal
(a) periodogram of the synthetic signal in dB
(b) singular values of the synthetic signal sorted in decreasing order

B. Model order selection for a synthetic signal

The test signal is a sum of $r = 5$ undamped complex exponentials plus a complex colored noise. For each $k \in \{1, \dots, r\}$, the k^{th} exponential is characterized by its amplitude α_k and its pole $z_k = e^{i2\pi f_k}$, where $f_k \in \mathbb{R}$ is its frequency. The values of the parameters are given in table II. The additive noise has been obtained by applying the filter

TABLE II
PARAMETERS OF THE SYNTHETIC SIGNAL

Frequency	0.1 Hz	0.102 Hz	0.4 Hz	0.7 Hz	0.9 Hz
Amplitude	100	100	10	50	100

$H(z) = \frac{1}{1-0.95\,z^{-1}}$ to a complex white gaussian noise, whose variance has been chosen so that the resulting Signal Noise Ratio (SNR) is 40 dB. The periodogram of the resulting test signal is represented in figure 2-a. It was computed from a signal of length 255, multiplied by a Blackman window, chosen for its high leakage rejection (-57 dB), and zero-padded to obtain 65536 points in the frequency domain. In particular, it can be noticed that the two complex exponentials of lowest frequency are not resolved by the periodogram.

The ESTER method is compared to several other signal enumeration techniques, among which three Information Theoretic Criteria (ITC), known as the *Akaike Information Criterion* (AIC) [5], the *Minimum Description Length* (MDL) [5], and the *Efficient Detection Criterion* (EDC) [9] which is known to be a robust generalization of AIC and MDL. These methods consist in minimizing a cost function which involves the singular values $\{\sigma_1, \dots, \sigma_n\}$:

$$\text{ITC}(p) = -(n-p)\,l\,\ln\left(\frac{\left(\prod_{q=p+1}^{n}\sigma_q^2\right)^{\frac{1}{n-p}}}{\frac{1}{n-p}\sum_{q=p+1}^{n}\sigma_q^2}\right) + p\,(2n-p)\,C(l) \quad (23)$$

where $C(l)$ is a function of l. The AIC criterion is defined by choosing $C(l) = 1$ and the MDL criterion is defined by choosing $C(l) = \frac{1}{2}\ln(l)$. The EDC criteria are obtained for all

functions of l such that $\lim_{l\to+\infty} \frac{C(l)}{l} = 0$ and $\lim_{l\to+\infty} \frac{C(l)}{\ln(\ln(l))} = +\infty$. We chose $C(l) = \sqrt{l\ln(\ln(l))}$, for which we obtained the best results. The singular values have been obtained by computing the SVD of a Hankel data matrix containing $n = 128$ rows and $l = 128$ columns, involving the $n+l-1 = 255$ samples of the whole signal. Figure 2-b displays the highest $p_{\max} = 25$ singular values. Note that the singular values do not present a significant decrease beyond $p = 5$.

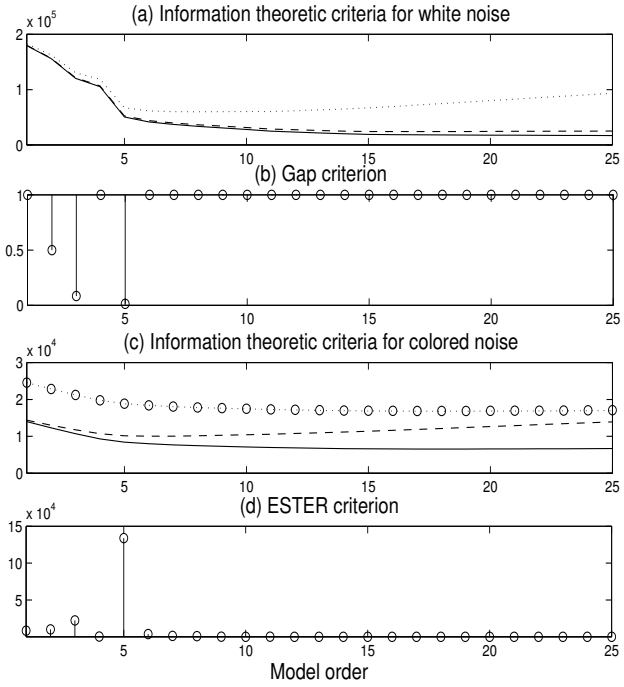

Fig. 3. Model order selection for the synthetic signal
(a) AIC, MDL, and EDC criteria
(b) Gap criterion
(c) C_1, C_2, C_{m1}, and C_{m2} criteria
(d) ESTER criterion

Figure 3-a displays the values of the AIC (solid line), MDL (dashed line) and EDC (dotted line) criteria, such as formulated in the above equation, for $p \in \{1 \dots p_{\max}\}$. None of them reaches a minimum at $p = r = 5$. This failure might be explained by the presence of the surrounding noise, whose power spectral density is not uniform, contrary to the additive white noise hypothesis on which these estimators basically rely. As expected, the EDC criterion is more robust than AIC and MDL, but its minimum is obtained for $p = 8$. Figure 3-b represents the criterion proposed in [12] for detecting the gap in the singular values decrease. It can be noticed that this criterion selects the right value $p = r = 5$, but the value $p = 3$ is almost as much emphasized. Figure 3-c displays new ITC criteria proposed in [16] to address the colored noise case[5]. The best results were obtained with C_2, which reaches a minimum at $p = 6 \simeq r$.

Finally, figure 3-d displays the ESTER criterion $J(p)$ for $p \in \{1, \dots, p_{\max}\}$. It can be noticed that the global maximum is reached at $p = r = 5$, despite the surrounding noise, which was not included in the model.

[5]These new criteria are referred to as C_1 (solid line), C_2 (dashed line), C_{m1} (dotted line) and C_{m2} (circles line). The common value of the parameters M_1 and M_2 defined in [16] was set to $\frac{n}{2} - 1$.

TABLE III
RATES OF SUCCESSFUL SIGNAL ENUMERATION FOR VARIOUS N

N	125	250	500
AIC	8 %	2 %	0 %
MDL	48 %	61 %	70 %
EDC	38 %	58 %	77 %
gap criterion	47 %	63 %	76 %
C2	35 %	44 %	44 %
ESTER	48 %	63 %	76 %

TABLE IV
RATES OF SUCCESSFUL SIGNAL ENUMERATION FOR VARIOUS SNR

SNR	10 dB	20 dB	30 dB
AIC	3 %	2 %	2 %
MDL	45 %	61 %	65 %
EDC	18 %	58 %	77 %
gap criterion	28 %	63 %	75 %
C2	13 %	44 %	61 %
ESTER	36 %	63 %	76 %

C. Statistical performance comparison

Below, the ESTER method and the above-mentioned signal enumeration techniques are applied to various synthetic signals. These signals consist of a sum of real-valued and undamped sinusoids, plus a colored noise. The number of sinusoids is uniformly distributed between 1 and 10, so that the model order r belongs to $\{2 \dots 20\}$. Their amplitudes, phases and frequencies are randomly distributed in the intervals $[1, 10]$, $]-\pi, \pi]$ and $]-\frac{1}{2}, \frac{1}{2}]$. The additive noise is obtained by filtering a white gaussian noise by the high-pass filter $1 - 0.5\,z^{-1}$ (whose rejection is lower than 10 dB).

As proposed in section IV, the robustness of the ESTER method is improved here by detecting the greatest value of p for which the ESTER criterion $J(p)$ reaches a local maximum which is greater than one tenth of the global maximum. Tables III and IV show the rates of successful signal enumeration, averaged over 10000 independent runs, for various values of the window length N and the SNR (other analysis parameters are $n = \lfloor N/2 \rfloor$ and $p_{\max} = 22$). In table III, the SNR is fixed to 20 dB, and the experiment is run for $N = 125$, $N = 250$ and $N = 500$. In table IV, N is fixed to 250, and the experiment is run for a low SNR (10 dB), a moderate SNR (20 dB) and a high SNR (30 dB). The obtained percentages have been rounded towards the nearest integer, since the number of independent runs (10000) guarantees that the confidence interval lies between $\pm 1\%$ around the estimated rate, for a 95% confidence level.

It can be noticed that the AIC is ineffective for processing these synthetic data. Besides, the successful rate of the C_2 criterion [16] is always lower than that of all the other methods. The MDL criterion seems to be more robust than ESTER to low SNRs. However, ESTER outperforms MDL both for high values of N and high SNRs. Compared to the EDC criterion, ESTER presents similar performance for high N and high SNRs, but it is more robust to low values of N and low SNRs. Finally, ESTER behaves similarly to the gap criterion [12] for all values of N. However, it is more robust to

low SNRs. It can be noticed that although the ESTER method relies on a noise-free signal model, its performance does not collapse at low SNRs.

D. Model order selection for a musical signal

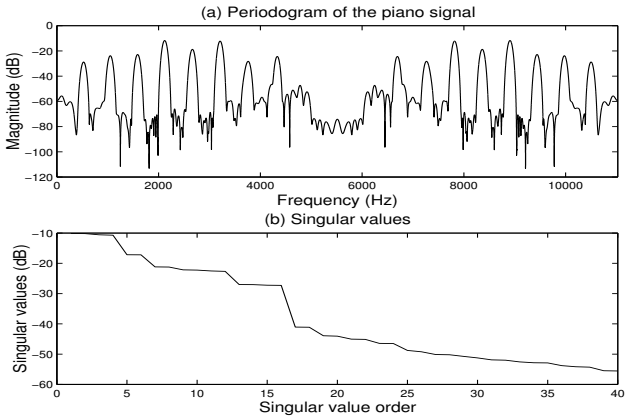

Fig. 4. Periodogram and singular values of a piano signal
(a) periodogram of the piano signal in dB
(b) singular values of the piano signal sorted in decreasing order

This section illustrates the application of the ESTER method to a musical signal. The study deals with a piano tone, C5, sampled at 11025 Hz, from which a segment of 255 samples (23 ms) has been extracted. Note that audio signals often require some pre-processing before applying the ESPRIT algorithm. For example, signals with a high number of sinusoids (typically low-pitched sounds) may first be decomposed into several sub-band signals (via filtering/decimating, as proposed in [33]). In this example, this pre-processing is not used, since the chosen piano tone has few sinusoidal components. In other respects, it is well known that the energy of audio signals is not evenly distributed over frequencies. Therefore we used a pre-emphasis filter obtained by linear prediction at order 7 to compensate the energy decrease.

The periodogram of the filtered piano signal is displayed in figure 4-a. In this figure, sixteen sharp spectral peaks clearly rise above the surrounding noise level. The $p_{\max} = 40$ highest singular values of the data matrix[6] are represented in figure 4-b. Clearly, these singular values collapse beyond $p = 16$, which suggests a modeling order equal to 16.

Figure 5-a displays the AIC (solid line), MDL (dashed line) and EDC (dotted line) criteria. Only EDC reaches a minimum at $p = 16$. However this minimum is not substantially lower than the neighboring values. Figure 5-b represents the gap criterion proposed in [12]. Contrary to the above-mentioned ITC, this criterion here selects the right value $p = 16$. Figure 5-c displays the ITC criteria proposed in [16], with the same parameters as in section V-B. None of them reaches a minimum at $p = 16$. Finally, figure 5-d displays the ESTER criterion $J(p)$ for all $p \in \{1, \ldots, p_{\max}\}$. The global maximum is reached at $p = 16$. It can also be noticed that the error bounds obtained for lower values of p are relevant. Indeed, high values are reached at $p = 4$, 6, 12, which in fact correspond to small jumps in the decrease of the singular values (represented in figure 4-b). Therefore, the ESTER method gives the expected model order, and moreover the error bounds can be used to quantify the adequacy of a possible lower modeling order. In particular, it can be noticed that odd model orders do not fit the signal. Indeed, since this signal is real-valued and centered, its spectrum is hermitian symmetric with no constant component, which underlies an even model order.

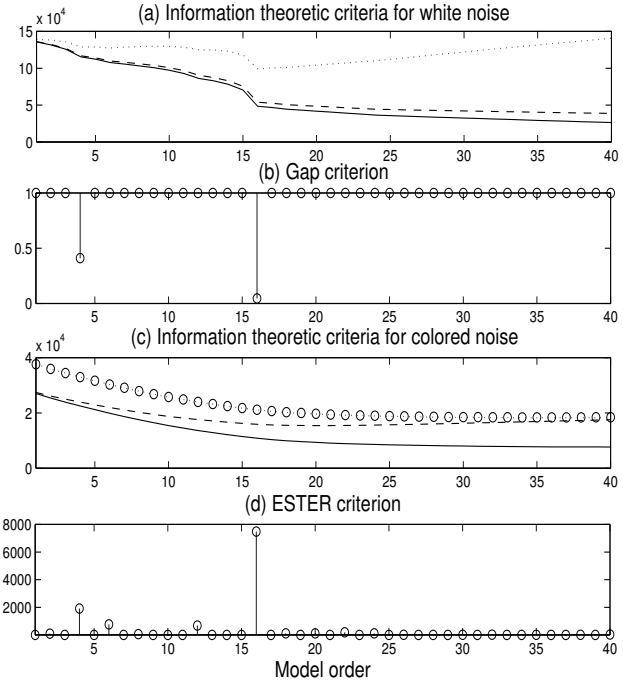

Fig. 5. Model order selection for the piano signal
(a) AIC, MDL, and EDC criteria
(b) Gap criterion
(c) C_1, C_2, C_{m1}, and C_{m2} criteria
(d) ESTER criterion

VI. Conclusions

In this paper, we described how the estimation of a noiseless ESM model is affected by applying the ESPRIT algorithm with an erroneous model order. If the model order is overestimated, the true poles are among those estimated. On the contrary, if the order is underestimated, the estimated poles can be seen as approximations of some of the true ones. In this last case, an *a posteriori error bound* was given, which can be computed without knowing the exact model order. Following from this observation, we introduced the ESTER criterion for selecting an appropriate model order. Since the initial method was computationally expensive, we proposed a fast algorithm for recursively computing the a posteriori error bounds. Then, we showed the relevance of our criterion as an error bound, and we illustrated the performance of the ESTER criterion on synthetic signals and on a piano signal. In addition, we noted that the error bounds could be used to quantify the adequacy

[6]The singular values have been obtained by computing the SVD of a Hankel data matrix containing $n = 128$ rows and $l = 128$ columns, as in section V-B.

of a possible lower modeling order. Finally, although it was designed for the ESPRIT algorithm, the ESTER criterion can be used with any High Resolution method. It can also be adapted to estimate the order of the more general *Polynomial Amplitude Complex Exponentials* (PACE) model [34].

In presence of noise, we mentioned in section IV that a robust way of selecting the modeling order consists in detecting the greatest value of p for which the ESTER criterion $J(p)$ reaches a local maximum greater than a fraction of its global maximum. Future work will be dedicated to an analysis of the effect of noise, in order to better exploit the information provided by the ESTER criterion.

APPENDIX

Proof of theorem III.2:

If $\widehat{z} = z_k$ for some $k \in \{1 \dots r\}$, the assertion is trivial, so we may assume that $\forall k \in \{1 \dots r\}$, $\widehat{z} \neq z_k$. Let $\widehat{\boldsymbol{v}}$ be a given unitary vector, and define the residual

$$\boldsymbol{e}(p) \triangleq (\boldsymbol{W}_{\uparrow}(p) - \widehat{z}\,\boldsymbol{W}_{\downarrow}(p))\,\widehat{\boldsymbol{v}}. \tag{24}$$

Since $\boldsymbol{W}(p)\,\widehat{\boldsymbol{v}} = \boldsymbol{W}(r) \begin{bmatrix} \widehat{\boldsymbol{v}} \\ \boldsymbol{0} \end{bmatrix}$, equation (24) yields

$$\boldsymbol{e}(p) = (\boldsymbol{W}_{\uparrow}(r) - \widehat{z}\,\boldsymbol{W}_{\downarrow}(r)) \begin{bmatrix} \widehat{\boldsymbol{v}} \\ \boldsymbol{0} \end{bmatrix}. \tag{25}$$

Substituting equations (3) and (4) into equation (25) yields

$$\boldsymbol{e}(p) = \boldsymbol{V}_{\downarrow}^{n}\,(\boldsymbol{D} - \widehat{z}\,\boldsymbol{I}_r)\,\boldsymbol{G}^{-1} \begin{bmatrix} \widehat{\boldsymbol{v}} \\ \boldsymbol{0} \end{bmatrix}. \tag{26}$$

Since $\forall k \in \{1 \dots r\}$, $\widehat{z} \neq z_k$, $\boldsymbol{D} - \widehat{z}\,\boldsymbol{I}_r$ is non-singular. Therefore, equation (26) yields

$$\begin{bmatrix} \widehat{\boldsymbol{v}} \\ \boldsymbol{0} \end{bmatrix} = \boldsymbol{G}\,(\boldsymbol{D} - \widehat{z}\,\boldsymbol{I}_r)^{-1}\,\boldsymbol{V}_{\downarrow}^{n\dagger}\,\boldsymbol{e}(p). \tag{27}$$

Consider a given diagonal matrix $\boldsymbol{\Lambda}$ whose diagonal coefficients are positive. Then equation (27) is equivalent to

$$\begin{bmatrix} \widehat{\boldsymbol{v}} \\ \boldsymbol{0} \end{bmatrix} = \boldsymbol{G}\,\boldsymbol{\Lambda}\,(\boldsymbol{D} - \widehat{z}\,\boldsymbol{I}_r)^{-1}\,\boldsymbol{\Lambda}^{-1}\,\boldsymbol{V}_{\downarrow}^{n\dagger}\,\boldsymbol{e}(p). \tag{28}$$

Applying the 2-norm into equation (28) yields

$$\|\widehat{\boldsymbol{v}}\|_2 \leq \|\boldsymbol{G}\,\boldsymbol{\Lambda}\|_2 \left\|(\boldsymbol{D} - \widehat{z}\,\boldsymbol{I}_r)^{-1}\right\|_2 \left\|\boldsymbol{\Lambda}^{-1}\boldsymbol{V}_{\downarrow}^{n\dagger}\right\|_2 \|\boldsymbol{e}(p)\|_2. \tag{29}$$

Since $\boldsymbol{W}(r)$ is orthonormal and $\boldsymbol{V}^n\,\boldsymbol{\Lambda} = \boldsymbol{W}(r)\,\boldsymbol{G}\,\boldsymbol{\Lambda}$, $\|\boldsymbol{G}\,\boldsymbol{\Lambda}\|_2 = \|\boldsymbol{V}^n\,\boldsymbol{\Lambda}\|_2 = \sigma_{\max}(\boldsymbol{V}^n\,\boldsymbol{\Lambda})$. Moreover, $(\boldsymbol{D} - \widehat{z}\,\boldsymbol{I}_r)^{-1}$ is diagonal with diagonal entries $\frac{1}{z_k - \widehat{z}}$, thus $\left\|(\boldsymbol{D} - \widehat{z}\,\boldsymbol{I}_r)^{-1}\right\|_2 = \frac{1}{\min\limits_{k\in\{1\dots r\}} |\widehat{z} - z_k|}$. Since $\boldsymbol{V}_{\downarrow}^{n}$ is full-rank, the singular values of $\boldsymbol{\Lambda}^{-1}\boldsymbol{V}_{\downarrow}^{n\dagger}$ are the inverses of those of $\boldsymbol{V}_{\downarrow}^{n}\,\boldsymbol{\Lambda}$, so that $\left\|\boldsymbol{\Lambda}^{-1}\boldsymbol{V}_{\downarrow}^{n\dagger}\right\|_2 = \frac{1}{\sigma_{\min}(\boldsymbol{V}_{\downarrow}^{n}\,\boldsymbol{\Lambda})}$. Finally, since $\widehat{\boldsymbol{v}}$ is unitary, equation (29) yields

$$\min_{k\in\{1\dots r\}} |\widehat{z} - z_k| \leq \frac{\sigma_{\max}(\boldsymbol{V}^n\,\boldsymbol{\Lambda})}{\sigma_{\min}(\boldsymbol{V}_{\downarrow}^{n}\,\boldsymbol{\Lambda})}\,\|\boldsymbol{e}(p)\|_2. \tag{30}$$

Note that equation (30) is satisfied for all unitary vector $\widehat{\boldsymbol{v}}$ and all matrix $\boldsymbol{\Lambda} \in \mathrm{diag}(\mathbb{R}_+^{*\,r})$. Consequently, equation (7) follows from equation (30). ■

Proof of proposition IV.1: It can be noticed that

$$\boldsymbol{E}(p) = \left(\boldsymbol{W}_{\uparrow}(p)\,\boldsymbol{W}_{\uparrow}(p)^{\dagger} - \boldsymbol{W}_{\downarrow}(p)\,\boldsymbol{W}_{\downarrow}(p)^{\dagger}\right)\boldsymbol{W}_{\uparrow}(p).$$

Applying the 2-norm yields

$$\|\boldsymbol{E}(p)\|_2 \leq \mathrm{dist}\,(\boldsymbol{W}_{\uparrow}(p), \boldsymbol{W}_{\downarrow}(p))\;\|\boldsymbol{W}_{\uparrow}(p)\|_2 \tag{31}$$

where

$$\mathrm{dist}\,(\boldsymbol{W}_{\uparrow}(p), \boldsymbol{W}_{\downarrow}(p)) \triangleq \left\|\boldsymbol{W}_{\uparrow}(p)\,\boldsymbol{W}_{\uparrow}(p)^{\dagger} - \boldsymbol{W}_{\downarrow}(p)\,\boldsymbol{W}_{\downarrow}(p)^{\dagger}\right\|_2$$

is the distance between the subspaces $\mathrm{Span}(\boldsymbol{W}_{\uparrow}(p))$ and $\mathrm{Span}(\boldsymbol{W}_{\downarrow}(p))$, which satisfies $\mathrm{dist}\,(\boldsymbol{W}_{\uparrow}(p), \boldsymbol{W}_{\downarrow}(p)) \leq 1$, as shown in [32, pp. 76-77]. Since $\|\boldsymbol{W}_{\uparrow}(p)\|_2 \leq \|\boldsymbol{W}(p)\|_2 = 1$, the result follows from equation (31). ■

REFERENCES

[1] R. J. McAulay and T. F. Quatieri, "Speech analysis and synthesis based on a sinusoidal representation," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 34, no. 4, Aug. 1986.

[2] B. David, G. Richard, and R. Badeau, "An EDS modeling tool for tracking and modifying musical signals," in *Proc. of SMAC 03*, vol. 2, Stockholm, Sweden, Aug. 2003, pp. 715–718.

[3] R. Roy, A. Paulraj, and T. Kailath, "ESPRIT–A subspace rotation approach to estimation of parameters of cisoids in noise," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 34, no. 5, pp. 1340–1342, Oct. 1986.

[4] G. Bienvenu and L. Kopp, "Optimality of high-resolution array processing using the eigensystem method," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 31, no. 5, pp. 1235–1245, Oct. 1983.

[5] M. Wax and T. Kailath, "Detection of signals by information theoretic criteria," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 33, no. 2, pp. 387–392, Apr. 1985.

[6] H. Akaike, "Information theory and an extension of the maximum likelihood principle," in *Proc. of the 2nd International Symposium on Information Theory*, B. N. Petrov and F. Csaki, Eds. Budapest: Akademia Kiado, 1973, pp. 267–281.

[7] G. Schwarz, "Estimating the dimension of a model," *The Annals of Statistics*, vol. 6, no. 2, pp. 461–464, 1978.

[8] J. Rissanen, "Modeling by shortest data description," *Automatica*, vol. 14, pp. 465–471, 1978.

[9] L. C. Zhao, P. R. Krishnaiah, and Z. D. Bai, "On detection of the number of signals in presence of white noise," *Journal of Multivariate Analysis*, vol. 20, no. 1, pp. 1–25, 1986.

[10] F. Gini and F. Bordoni, "On the behavior of information theoretic criteria for model order selection of InSAR signals corrupted by multiplicative noise," *Signal Processing*, vol. 83, pp. 1047–1063, 2003.

[11] A. P. Liavas and P. A. Regalia, "On the behavior of Information Theoretic Criteria for model order selection," *IEEE Trans. Signal Processing*, vol. 49, no. 8, pp. 1689–1695, Aug. 2001.

[12] A. P. Liavas, P. A. Regalia, and J.-P. Delmas, "Blind channel approximation: effective channel order determination," *IEEE Trans. Signal Processing*, vol. 47, no. 12, pp. 3336–3344, Dec. 1999.

[13] J. Grouffaud, P. Larzabal, and H. Clergeot, "Some properties of ordered eigenvalues of a Wishart matrix: application in detection test and model order selection," in *Proc. of ICASSP'96*, vol. 5. IEEE, 1996, pp. 2465–2468.

[14] D. Kundu and A. Mitra, "Detecting the number of signals for an undamped exponential model using cross-validation approach," *Signal Processing*, vol. 80, no. 3, pp. 525–534, 2000.

[15] L. C. Zhao, P. R. Krishnaiah, and Z. D. Bai, "On detection of the number of signals when the noise covariance matrix is arbitrary," *Journal of Multivariate Analysis*, vol. 20, no. 1, pp. 26–49, 1986.

[16] Q. T. Zhang and K. M. Wong, "Information theoretic criteria for the determination of the number of signals in spatially correlated noise," *IEEE Trans. Signal Processing*, vol. 41, no. 4, pp. 1652–1663, Apr. 1993.

[17] J. J. Fuchs, "Estimation of the number of signals in the presence of unknown correlated sensor noise," *IEEE Trans. Signal Processing*, vol. 40, no. 5, pp. 1053–1061, May 1992.

[18] W. B. Bishop and P. M. Djuric, "Model order selection of damped sinusoids in noise by predictive densities," *IEEE Trans. Signal Processing*, vol. 44, no. 3, pp. 611–619, Mar. 1996.

[19] B. Rao and K. Hari, "Performance analysis of ESPRIT and TAM in determining the direction of arrival of plane waves in noise," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 37, no. 12, pp. 1990–1995, Dec. 1989.

[20] P. Stoica and T. Söderström, "Statistical Analysis of MUSIC and Subspace Rotation Estimates of Sinusoidal Frequencies," *IEEE Trans. Signal Processing*, vol. 39, pp. 1836–1847, Aug. 1991.

[21] A. Eriksson, P. Stoica, and T. Soderstrom, "Second-order properties of MUSIC and ESPRIT estimates of sinusoidal frequencies in high SNR scenarios," in *IEE Proceedings on Radar, Sonar and Navigation*, vol. 140, no. 4, Aug. 1993, pp. 266–272.

[22] A. Kangas, P. Stoica, and T. Soderstrom, "Finite sample and modelling error effects on ESPRIT and MUSIC direction estimators," in *IEE Proceedings on Radar, Sonar and Navigation*, vol. 141, no. 5, Oct. 1994, pp. 249–255.

[23] H. Saarnisaari, "Robustness of the MUSIC algorithm to Errors in Estimation the Dimensions of the Subspaces: Delay Estimation in DS/SS in the Presence of Interference," in *Proc. of the Military Communications Conference MILCOM'99*, Atlantic City, USA, 1999.

[24] F. Li, R. J. Vaccaro, and D. W. Tufts, "Performance analysis of the state-space realization (TAM) and ESPRIT algorithms for DOA estimation," *IEEE Trans. Antennas Propagat.*, vol. 39, no. 3, pp. 418–423, Mar. 1991.

[25] R. Badeau, B. David, and G. Richard, "Selecting the modeling order for the ESPRIT high resolution method: an alternative approach," in *Proc. of ICASSP'04*, vol. 2. Montreal, Quebec: IEEE, May 2004, pp. 1025–1028.

[26] R. Badeau, R. Boyer, and B. David, "EDS parametric modeling and tracking of audio signals," in *Proc. of Int. Conf. on Digital Audio Effects DAFx-02*, Hamburg, Sept. 2002, pp. 139–144.

[27] Y. Hua and T. K. Sarkar, "Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 38, no. 5, pp. 814–824, May 1990.

[28] R. Badeau, G. Richard, and B. David, "Sliding window adaptive SVD algorithms," *IEEE Trans. Signal Processing*, vol. 52, no. 1, pp. 1–10, Jan. 2004.

[29] R. Badeau, B. David, and G. Richard, "Fast Approximated Power Iteration Subspace Tracking," *IEEE Trans. Signal Processing*, Aug. 2005.

[30] ——, "Yet Another Subspace Tracker," in *Proc. of ICASSP'05*, vol. 4. Philadelphia, PA, USA: IEEE, Mar. 2005, pp. 329–332.

[31] R. A. Horn and C. R. Johnson, *Matrix analysis*. Cambridge: Cambridge University Press, 1985.

[32] G. H. Golub and C. F. V. Loan, *Matrix computations*, 3rd ed. Baltimore and London: The Johns Hopkins University Press, 1996.

[33] J. Laroche, "The use of the Matrix Pencil method for the spectrum analysis of musical signals," *Journal of the Acoustical Society of America*, vol. 94, no. 4, pp. 1958–1965, Oct. 1993.

[34] R. Badeau, B. David, and G. Richard, "High resolution spectral analysis of mixtures of complex exponentials modulated by polynomials," *IEEE Trans. Signal Processing*, (accepted for publication).

Fast Approximated Power Iteration Subspace Tracking

Roland Badeau, *Member, IEEE,* Bertrand David, and Gaël Richard, *Member, IEEE*

***Abstract*—This paper introduces a fast implementation of the power iteration method for subspace tracking, based on an approximation less restrictive than the well known *projection approximation*. This algorithm, referred to as the fast API method, guarantees the orthonormality of the subspace weighting matrix at each iteration. Moreover, it outperforms many subspace trackers related to the power iteration method, such as PAST, NIC, NP3 and OPAST, while having the same computational complexity. The API method is designed for both exponential windows and sliding windows. Our numerical simulations show that sliding windows offer a faster tracking response to abrupt signal variations.**

***Index Terms*—Subspace tracking, projection approximation, power iteration.**

I. Introduction

THE interest in subspace-based methods stems from the fact that they consist in splitting the observations into a set of desired and a set of disturbing components, which can be viewed in terms of signal and noise subspaces. These methods have applications in numerous domains including the fields of adaptive filtering, source localization, or parameter estimation [1]. The estimation of the signal subspace is commonly based on the traditional eigenvalue decomposition (EVD) or singular value decomposition (SVD). However, the main drawback of these decompositions is their inherent complexity. Therefore, there is a real need for fast subspace tracking techniques in the context of adaptive signal processing.

Due to this interest, a large number of approaches have already been introduced. A reference method is I. Karasalo's algorithm [2], which involves the full SVD of a small matrix. A fast tracking method based on Givens rotations (the FST algorithm) is proposed in [3]. Other approaches consist in interlacing a recursive update of the estimated correlation matrix or the data matrix with one or a few steps of a standard SVD or power iteration algorithm. This is the case of the Jacobi SVD method [4], the transposed QR-iteration [5], the orthogonal / bi-orthogonal iteration [6], [7], and the power method [8]. Other matrix decompositions have also successfully been used in subspace tracking (for example the rank-revealing QR factorization [9], the rank-revealing URV decomposition [10], and the Lankzos (bi)-diagonalization [11]). Other techniques rely on the noise and signal subspace averaging method [12], the maximum likelihood principle [13], the operator restriction analysis [14], or the perturbation theory [15].

Roland Badeau, Bertrand David and Gaël Richard are with the Department of Signal and Image Processing, Ecole Nationale Supérieure des Télécommunications (ENST), Paris, France. E-mail: [roland.badeau, bertrand.david, gael.richard]@enst.fr.

The estimation of the signal subspace can also be viewed as a constrained or unconstrained optimization problem [16]–[21], for which the introduction of a projection approximation hypothesis lead to fast subspace tracking methods (see *e.g.* the PAST [22] and NIC [23] algorithms). In [8], it is shown that these subspace trackers are closely linked to the classical power iteration method [24]. Several implementations of this method based on QR factorizations are proposed in [6], among which the Loraf2 and Loraf3 algorithms. However, compared to PAST and NIC, Loraf2 is more computationally demanding and the performance of Loraf3 is degraded. Another fast implementation of the power iteration method, the NP3 algorithm which relies on rank-one matrix updates, is proposed in [8], but our numerical simulations showed that this algorithm does not converge in many situations. An orthonormal version of the PAST algorithm, proposed in [25], can be seen as a fast implementation of the power method and outperforms PAST, NIC and NP3. Concurrently, the recent API method [26], based on the power iteration method and on a new projection approximation, has the same computational complexity as the above mentioned algorithms but provides a better estimation of the dominant subspace.

All these adaptive techniques are designed for exponential windows. Indeed, this choice tends to smooth the variations of the signal parameters, and thus allows a low-complexity update at each time step. However, it is only suitable for slowly changing signals. Conversely, a few subspace trackers are based on sliding windows, which generally require more computations, but offer a faster tracking response to sudden signal changes [22], [27]–[30]. In particular, a sliding window version of the API algorithm is proposed in [31].

This paper presents several fast implementations of the API method. These algorithms present several advantages:

- they can be applied either on an infinite *exponential window* or on a *truncated window* (*e.g.* a sliding window which may have an exponential decrease),
- an orthonormal subspace basis is computed *at each time step*, which is required for some subspace-based estimation methods, such as MUSIC [32],
- they rely on a new projection approximation, less restrictive than the classical one, which leads to better tracking results. In particular, it is shown that the PAST and OPAST subspace trackers can be viewed as approximations of the fast API method.

The paper is organized as follows. In section II, we introduce the various window shapes applied to the data. In section III, the classical power iteration method is reviewed, then the projection approximation is discussed in section IV.

Our approximated power iteration (API) method is introduced in section V, and a fast implementation of this algorithm is proposed in section VI. In section VII, it is shown that both PAST and OPAST can be seen as approximations of the fast API algorithm. In section VIII, the performance of this method is compared to that of several subspace trackers, among which PAST and OPAST. Finally, the main conclusions of this paper are summarized in section IX.

II. Data windowing

Let $\{\boldsymbol{x}(t)\}_{t\in\mathbb{Z}}$ be a sequence of n-dimensional data vectors. We are interested in computing the dominant subspace spanned by its correlation matrix. This matrix can be estimated according to the nature of the data window.

A. Exponential window

The estimated $n \times n$ correlation matrix is defined as

$$\boldsymbol{C}_{xx}(t) = \sum_{u=-\infty}^{t} \beta^{t-u}\, \boldsymbol{x}(u)\, \boldsymbol{x}(u)^H$$

where $0 < \beta < 1$ is the forgetting factor. It can be recursively updated according to the following scheme:

$$\boldsymbol{C}_{xx}(t) = \beta\, \boldsymbol{C}_{xx}(t-1) + \boldsymbol{x}(t)\, \boldsymbol{x}(t)^H. \tag{1}$$

B. Truncated window

The $n \times n$ correlation matrix $\boldsymbol{C}_{xx}(t)$ is estimated on a window of length $l \in \mathbb{N}^*$:

$$\boldsymbol{C}_{xx}(t) = \sum_{u=t-l+1}^{t} \beta^{t-u}\, \boldsymbol{x}(u)\, \boldsymbol{x}(u)^H. \tag{2}$$

where $0 < \beta \leq 1$. The case $\beta = 1$ corresponds to a rectangular (or sliding) window. This matrix can be recursively updated according to the following scheme:

$$\boldsymbol{C}_{xx}(t) = \beta\, \boldsymbol{C}_{xx}(t-1) + \boldsymbol{x}(t)\, \boldsymbol{x}(t)^H - \beta^l \boldsymbol{x}(t-l)\, \boldsymbol{x}(t-l)^H. \tag{3}$$

C. Unified formalization

Both equations (1) and (3) can be written in the form

$$\boldsymbol{C}_{xx}(t) = \beta\, \boldsymbol{C}_{xx}(t-1) + \underline{\boldsymbol{x}}(t)\, \underline{J}\, \underline{\boldsymbol{x}}(t)^H \tag{4}$$

where $\underline{\boldsymbol{x}}(t)$ and $\underline{J}$ are defined according to the window shape:

- in the exponential window case:

$$\underline{J} = 1 \tag{5}$$
$$\underline{\boldsymbol{x}}(t) = \boldsymbol{x}(t) \tag{6}$$

- in the truncated window case:

$$\underline{J} = \begin{bmatrix} 1 & 0 \\ 0 & -\beta^l \end{bmatrix} \tag{7}$$
$$\underline{\boldsymbol{x}}(t) = \left[\, \boldsymbol{x}(t) \;\middle|\; \boldsymbol{x}(t-l) \,\right]. \tag{8}$$

Let p be the rank of the update involved in equation (4). Since $p = 1$ in the exponential window case and $p = 2$ in the truncated window case, p characterizes the window shape. In particular, $\underline{\boldsymbol{x}}(t)$ is a $n \times p$ matrix and $\underline{J}$ is a $p \times p$ matrix.

III. The classical power iteration method

The power iteration method [8] tracks the dominant subspace[1] of dimension $r \leq n$ spanned by the $n \times n$ matrix $\boldsymbol{C}_{xx}(t)$. At each time step, a basis of this subspace is computed, represented by an orthonormal matrix $\boldsymbol{W}(t)$ of dimension $n \times r$. The computation of $\boldsymbol{W}(t)$ consists of a data compression step (9) and an orthonormalization step (10) of the compressed matrix at each iteration:

$$\boldsymbol{C}_{xy}(t) = \boldsymbol{C}_{xx}(t)\, \boldsymbol{W}(t-1) \tag{9}$$

$$\boldsymbol{W}(t)\, \boldsymbol{R}(t) = \boldsymbol{C}_{xy}(t). \tag{10}$$

where $\boldsymbol{C}_{xy}(t)$ can be considered as a $n \times r$ correlation matrix between the n-dimensional data vectors $\boldsymbol{x}(t)$ and the r-dimensional compressed data vectors

$$\boldsymbol{y}(t) = \boldsymbol{W}(t-1)^H \boldsymbol{x}(t). \tag{11}$$

The orthonormalization step (10) involves a $r \times r$ matrix $\boldsymbol{R}(t)$, such that $\boldsymbol{R}(t)^H \boldsymbol{R}(t) = \boldsymbol{\Phi}(t)$, where $\boldsymbol{\Phi}(t)$ is the $r \times r$ positive definite matrix $\boldsymbol{C}_{xy}(t)^H \boldsymbol{C}_{xy}(t)$. Consequently, $\boldsymbol{R}(t)^H$ is a square root of $\boldsymbol{\Phi}(t)$. In particular, $\boldsymbol{R}(t)^H$ is equal to the positive definite square root of $\boldsymbol{\Phi}(t)$, right multiplied by a unitary matrix[2]. For example, $\boldsymbol{R}(t)$ can be triangular [6], or positive definite [8].

If $\boldsymbol{C}_{xx}(t)$ remains constant and if its first r eigenvalues are strictly larger than the $(n-r)^{\text{th}}$ others, the power iteration method converges globally and exponentially to the principal subspace [8] [24, pp. 410-411]. Note that the multiplication in step (9) involves $n^2 r$ operations, and the orthonormalization step (10) requires $O(nr^2)$ operations[3]. Because of its high computational cost, this algorithm is not suitable for real-time processing.

IV. The projection approximation

We are now looking for an approximation that will allow us to reduce the complexity. Suppose that $\boldsymbol{W}(t-1)$ *exactly* spans the r-dimensional dominant subspace of $\boldsymbol{C}_{xx}(t)$. Then equation (9) yields

$$\boldsymbol{C}_{xy}(t) = \boldsymbol{W}(t-1)\, \boldsymbol{C}_{yy}(t) \tag{12}$$

where the matrix $\boldsymbol{C}_{yy}(t) \triangleq \boldsymbol{W}(t-1)^H \boldsymbol{C}_{xx}(t)\, \boldsymbol{W}(t-1)$ can be seen as the correlation matrix of the compressed data

[1] The r-dimensional dominant subspace of the positive semidefinite matrix $\boldsymbol{C}_{xx}(t)$ is the subspace spanned by the r eigenvectors of $\boldsymbol{C}_{xx}(t)$ associated to the r eigenvalues of highest magnitude (which are supposed to be strictly greater than the $n-r$ others).

[2] If $\boldsymbol{T}$ is a positive definite matrix, a *square root* of $\boldsymbol{T}$ is any matrix $\boldsymbol{S}$ of the same dimension such that $\boldsymbol{S}\, \boldsymbol{S}^H = \boldsymbol{T}$. Such a matrix is denoted $\boldsymbol{S} = \boldsymbol{T}^{\frac{1}{2}}$. There is only one positive definite square root of $\boldsymbol{T}$. The other square roots are obtained by right multiplying this positive definite square root by any unitary transform. The notation $\boldsymbol{S}^{\frac{1}{2}}$ can denote any of them.

[3] In this paper, operations counts are expressed in terms of multiply / accumulate (MAC) operations, herein referred to as *flops*. Whenever a specific matrix function is used, such as orthonormalization, inversion or square rooting, only the order of the operations count is presented, since the exact operations count depends on the way this function is implemented. Nevertheless, r is supposed to be much lower than n, so that the dominant cost of the power iteration method is that of the first step, whose exact operations count is known ($n^2 r$).

vectors. In this case, $\boldsymbol{W}(t)$ and $\boldsymbol{W}(t-1)$ are two orthonormal matrices spanning the range space of $\boldsymbol{C}_{xy}(t)$, thus

$$\boldsymbol{W}(t) = \boldsymbol{W}(t-1)\,\boldsymbol{\Theta}(t) \tag{13}$$

where $\boldsymbol{\Theta}(t) \triangleq \boldsymbol{W}(t-1)^H \boldsymbol{W}(t)$ is a $r \times r$ orthonormal matrix. Substituting equation (12) into equation (10) and left multiplying by $\boldsymbol{W}(t)^H$ yields the polar decomposition of $\boldsymbol{R}(t)^H$:

$$\boldsymbol{R}(t)^H = \boldsymbol{C}_{yy}(t)\,\boldsymbol{\Theta}(t) \tag{14}$$

where $\boldsymbol{C}_{yy}(t)$ is the positive definite factor and $\boldsymbol{\Theta}(t)$ is the orthonormal factor. Now suppose that $\boldsymbol{W}(t-1)$ *approximately* spans the dominant subspace of $\boldsymbol{C}_{xx}(t)$. Then equations (13) and (14) become approximations:

$$\boldsymbol{W}(t) \simeq \boldsymbol{W}(t-1)\,\boldsymbol{\Theta}(t) \tag{15}$$

$$\boldsymbol{R}(t)^H \simeq \boldsymbol{C}_{yy}(t)\,\boldsymbol{\Theta}(t) \tag{16}$$

where the $r \times r$ matrix $\boldsymbol{\Theta}(t)$ is *nearly* orthonormal.

Compared to equation (15), the classical projection approximation [22] is equivalent to $\boldsymbol{W}(t) \simeq \boldsymbol{W}(t-1)$ at each time step[4]. The validity of this approximation additionally requires that $\boldsymbol{\Theta}(t)$ is close to the $r \times r$ identity matrix (herein denoted $\boldsymbol{I}_r$). In this case, equation (16) shows that $\boldsymbol{R}(t)^H$ must be nearly positive definite[5]. Consequently, the choice of the square root $\boldsymbol{R}(t)^H$ of $\boldsymbol{\Phi}(t)$ is restricted (*e.g.* $\boldsymbol{R}(t)$ can no longer be upper triangular, as it was in [6]).

The NP3 implementation of the power method [8] is based on this approximation, but this algorithm relies on a matrix $\boldsymbol{R}(t)$ which deviates from the positive definite structure constraint. Therefore, the classical projection approximation does not stand, and this subspace tracker is not guaranteed to converge.

Concurrently, the algorithms presented in section V do not have to face this limitation, since they rely on the less restrictive approximation (15). Also note that (15) is the best approximation of $\boldsymbol{W}(t)$ in terms of mean square error, since the solution to the minimization problem

$$\arg \min_{\boldsymbol{\Theta} \in \mathbb{C}^{r \times r}} \|\boldsymbol{W}(t) - \boldsymbol{W}(t-1)\,\boldsymbol{\Theta}\|_F^2$$

is $\boldsymbol{\Theta}(t) = \boldsymbol{W}(t-1)^H \boldsymbol{W}(t)$ (where $\boldsymbol{W}(t-1)$ is supposed to be orthonormal).

V. Approximated Power Iteration

The complexity of the power iteration method can be reduced by introducing approximation (15) at time $t-1$ in step (9). Then the $n \times r$ matrix $\boldsymbol{C}_{xy}(t)$ can be computed recursively, as shown in section V-A, and factorization (10) can be updated, as shown in section V-C. This fast update requires the introduction of a $r \times r$ auxiliary matrix $\boldsymbol{Z}(t)$, introduced in section V-B.

[4]In fact, the projection approximation in [22] is defined as $\boldsymbol{W}(t')^H \boldsymbol{x}(t) \approx \boldsymbol{W}(t-1)^H \boldsymbol{x}(t) \triangleq \boldsymbol{y}(t)\ \forall t' \geq t$. It was shown in [8, pp. 301] that this approximation is equivalent to $\boldsymbol{W}(t) \simeq \boldsymbol{W}(t-1)$ at each time step.

[5]Conversely, if $\boldsymbol{R}(t)^H$ is chosen close to the only positive definite square root of $\boldsymbol{\Phi}(t)$, the approximate polar decomposition (16) shows that $\boldsymbol{\Theta}(t) \simeq \boldsymbol{I}_r$, so that equation (15) yields $\boldsymbol{W}(t) \simeq \boldsymbol{W}(t-1)$.

A. Recursion for the matrix $\boldsymbol{C}_{xy}(t)$

It is shown in this section that the $n \times r$ matrix $\boldsymbol{C}_{xy}(t)$ can be updated in the same way as the $n \times n$ matrix $\boldsymbol{C}_{xx}(t)$ in equation (4):

$$\boldsymbol{C}_{xy}(t) = \beta\,\boldsymbol{C}_{xy}(t-1)\,\boldsymbol{\Theta}(t-1) + \underline{\boldsymbol{x}}(t)\,\underline{J}\,\underline{\widehat{\boldsymbol{y}}}(t)^H. \tag{17}$$

In the exponential window case, equation (17) involves a rank-one update ($\underline{\boldsymbol{x}}(t)$ and $\underline{\widehat{\boldsymbol{y}}}(t)$ are vectors and $\underline{J}$ is a scalar), whereas in the truncated window case it involves a rank-two update ($\underline{\boldsymbol{x}}(t)$ and $\underline{\widehat{\boldsymbol{y}}}(t)$ are two-column matrices and $\underline{J}$ is a 2×2 matrix).

1) Truncated window: first, equation (2) can be written

$$\boldsymbol{C}_{xx}(t) = \boldsymbol{X}(t)\,\boldsymbol{D}\,\boldsymbol{X}(t)^H \tag{18}$$

where $\boldsymbol{X}(t) \triangleq [\boldsymbol{x}(t-l+1), \boldsymbol{x}(t-l+2), \ldots, \boldsymbol{x}(t)]$ is the $n \times l$ data matrix and $\boldsymbol{D}$ is the $l \times l$ diagonal matrix $\mathrm{diag}(\beta^{l-1}, \beta^{l-2}, \ldots, \beta, 1)$.

Substituting equation (18) into equation (9) yields

$$\boldsymbol{C}_{xy}(t) = \boldsymbol{X}(t)\,\boldsymbol{D}\,\boldsymbol{Y}(t)^H \tag{19}$$

where $\boldsymbol{Y}(t) \triangleq \boldsymbol{W}(t-1)^H \boldsymbol{X}(t)$ is the $r \times l$ compressed data matrix. Now let us show recursions for matrices $\boldsymbol{X}(t)$ and $\boldsymbol{Y}(t)$. The first one is straightforward:

$$\left[\ \boldsymbol{x}(t-l)\ \middle|\ \boldsymbol{X}(t)\ \right] = \left[\ \boldsymbol{X}(t-1)\ \middle|\ \boldsymbol{x}(t)\ \right]. \tag{20}$$

Then left multiplying equation (20) by $\boldsymbol{W}(t-1)^H$ yields

$$\left[\ \boldsymbol{v}(t-l)\ \middle|\ \boldsymbol{Y}(t)\ \right] = \left[\ \boldsymbol{W}(t-1)^H \boldsymbol{X}(t-1)\ \middle|\ \boldsymbol{y}(t)\ \right] \tag{21}$$

where $\boldsymbol{y}(t)$, defined in equation (11), and

$$\boldsymbol{v}(t-l) \triangleq \boldsymbol{W}(t-1)^H \boldsymbol{x}(t-l) \tag{22}$$

are r-dimensional compressed data vectors. Applying approximation (15) at time $t-1$ to equation (21) yields the recursion $\left[\ \boldsymbol{v}(t-l)\ \middle|\ \boldsymbol{Y}(t)\ \right] \simeq \left[\ \widehat{\boldsymbol{V}}(t-1)\ \middle|\ \boldsymbol{y}(t)\ \right]$, where $\widehat{\boldsymbol{V}}(t-1)$ is the $r \times l$ compressed data matrix

$$\widehat{\boldsymbol{V}}(t-1) \triangleq \boldsymbol{\Theta}(t-1)^H \boldsymbol{Y}(t-1). \tag{23}$$

From now on, the exact definition of $\boldsymbol{Y}(t)$ is therefore replaced by

$$\left[\ \widehat{\boldsymbol{v}}(t-l)\ \middle|\ \boldsymbol{Y}(t)\ \right] \triangleq \left[\ \widehat{\boldsymbol{V}}(t-1)\ \middle|\ \boldsymbol{y}(t)\ \right] \tag{24}$$

where the r-dimensional vector $\widehat{\boldsymbol{v}}(t-l)$, defined by the first column in the left side of equation (24), is an approximation of the vector $\boldsymbol{v}(t-l)$. Equations (19), (20), (23) and (24) finally yield

$$\begin{aligned} \boldsymbol{C}_{xy}(t) = \ & \beta\,\boldsymbol{C}_{xy}(t-1)\,\boldsymbol{\Theta}(t-1) \\ & + \boldsymbol{x}(t)\,\boldsymbol{y}(t)^H - \beta^l\,\boldsymbol{x}(t-l)\,\widehat{\boldsymbol{v}}(t-l)^H \end{aligned} \tag{25}$$

This recursion can be seen as a particular case of equation (17), where $\underline{J}$ and $\underline{\boldsymbol{x}}(t)$ are defined in equations (7) and (8) and the $r \times p$ (with $p = 2$) matrix

$$\underline{\widehat{\boldsymbol{y}}}(t) \triangleq \left[\ \boldsymbol{y}(t)\ \middle|\ \widehat{\boldsymbol{v}}(t-l)\ \right] \tag{26}$$

is an approximation of

$$\underline{\boldsymbol{y}}(t) \triangleq \boldsymbol{W}(t-1)^H \underline{\boldsymbol{x}}(t) = \left[\ \boldsymbol{y}(t)\ \middle|\ \boldsymbol{v}(t-l)\ \right]. \tag{27}$$

2) Exponential window: substituting equation (1) into equation (9) yields

$$\boldsymbol{C}_{xy}(t) = \beta\, \boldsymbol{C}_{xx}(t-1)\boldsymbol{W}(t-1) + \boldsymbol{x}(t)\, \boldsymbol{y}(t)^H. \quad (28)$$

Applying the projection approximation (15) at time $t-1$, equation (28) can be replaced by the following recursion:

$$\boldsymbol{C}_{xy}(t) = \beta\, \boldsymbol{C}_{xy}(t-1)\, \boldsymbol{\Theta}(t-1) + \boldsymbol{x}(t)\, \boldsymbol{y}(t)^H. \quad (29)$$

This recursion can be seen as a particular case of equation (17), where $\underline{J}$ and $\underline{\boldsymbol{x}}(t)$ are defined in equations (5) and (6) and the $r \times p$ (with $p = 1$) matrix $\widehat{\underline{\boldsymbol{y}}}(t) \triangleq \boldsymbol{y}(t)$ is now equal to the vector $\underline{\boldsymbol{y}}(t) \triangleq \boldsymbol{W}(t-1)^H \underline{\boldsymbol{x}}(t) = \boldsymbol{y}(t)$.

B. Recursion for the matrix $\boldsymbol{Z}(t)$

Now, we aim at updating factorization (10) by means of equation (17). This calculation requires the introduction of an auxiliary matrix, denoted $\boldsymbol{Z}(t)$. Let $\boldsymbol{S}(t-1) \triangleq (\boldsymbol{R}(t-1)\, \boldsymbol{\Theta}(t-1))^H$ and suppose that the $r \times r$ matrix $\boldsymbol{S}(t-1)$ is non-singular. Then let

$$\boldsymbol{Z}(t-1) \triangleq \boldsymbol{S}(t-1)^{-1}. \quad (30)$$

Proposition 5.1: The $r \times r$ matrix

$$\boldsymbol{S}(t) \triangleq (\boldsymbol{R}(t)\, \boldsymbol{\Theta}(t))^H \quad (31)$$

is non-singular if and only if the $p \times p$ matrix $\beta \underline{J}^{-1} + \underline{\boldsymbol{y}}(t)^H \underline{\boldsymbol{h}}(t)$ is non-singular, where

$$\underline{\boldsymbol{h}}(t) \triangleq \boldsymbol{Z}(t-1)\, \widehat{\underline{\boldsymbol{y}}}(t). \quad (32)$$

has dimension $r \times p$. In this case, the $r \times r$ matrix

$$\boldsymbol{Z}(t) \triangleq \boldsymbol{S}(t)^{-1} \quad (33)$$

satisfies the recursion

$$\boldsymbol{Z}(t) = \frac{1}{\beta} \boldsymbol{\Theta}(t)^H \left(\boldsymbol{I}_r - \underline{\boldsymbol{g}}(t)\, \underline{\boldsymbol{y}}(t)^H\right) \boldsymbol{Z}(t-1)\, \boldsymbol{\Theta}(t)^{-H} \quad (34)$$

where $\underline{\boldsymbol{g}}(t)$ is the $r \times p$ matrix

$$\underline{\boldsymbol{g}}(t) \triangleq \underline{\boldsymbol{h}}(t)\, \left(\beta \underline{J}^{-1} + \underline{\boldsymbol{y}}(t)^H \underline{\boldsymbol{h}}(t)\right)^{-1}. \quad (35)$$

Proof:

Substituting equation (10) into equation (17) and left multiplying by $\boldsymbol{W}(t-1)^H$ leads to

$$\boldsymbol{\Theta}(t)\, \boldsymbol{R}(t) = \beta\, \boldsymbol{S}(t-1)^H + \underline{\boldsymbol{y}}(t)\, \underline{J}\, \widehat{\underline{\boldsymbol{y}}}(t)^H. \quad (36)$$

Next, the following matrix inversion lemma [33, pp. 18-19] will be applied to invert the right member of this equality. The interest of this approach is that the $r \times r$ matrix inversion problem is converted into a smaller $p \times p$ matrix inversion (with $p = 1$ or 2).

Lemma 5.2: Let $\boldsymbol{A}$ be a $r \times r$ non-singular complex matrix. Consider the $r \times r$ matrix $\boldsymbol{B} = \boldsymbol{A} + \boldsymbol{P}\, \boldsymbol{J}\, \boldsymbol{Q}$, where $\boldsymbol{P}$, $\boldsymbol{J}$ and $\boldsymbol{Q}$ have dimensions $r \times m$, $m \times m$ and $m \times r$, and $\boldsymbol{J}$ is supposed to be non-singular. Then $\boldsymbol{B}$ is non-singular if and only if $\boldsymbol{J}^{-1} + \boldsymbol{Q}\, \boldsymbol{A}^{-1}\, \boldsymbol{P}$ is non-singular, and in this case

$$\boldsymbol{B}^{-1} = \boldsymbol{A}^{-1} - \boldsymbol{A}^{-1}\, \boldsymbol{P} \left(\boldsymbol{J}^{-1} + \boldsymbol{Q}\, \boldsymbol{A}^{-1}\, \boldsymbol{P}\right)^{-1} \boldsymbol{Q}\, \boldsymbol{A}^{-1}.$$

Lemma 5.2 applied to equation (36) shows that the $r \times r$ matrix $\boldsymbol{\Theta}(t)\, \boldsymbol{R}(t)$ is non-singular if and only if the $p \times p$ matrix $\beta \underline{J}^{-1} + \underline{\boldsymbol{y}}(t)^H \underline{\boldsymbol{h}}(t)$ is non-singular (which provides a fast way of detecting the singularity of $\boldsymbol{R}(t)$ or $\boldsymbol{\Theta}(t)$). In the non-singular case, lemma 5.2 leads to the equation

$$(\boldsymbol{\Theta}(t)\, \boldsymbol{R}(t))^{-1} = \frac{1}{\beta} \boldsymbol{Z}(t-1)^H \left(\boldsymbol{I}_r - \underline{\boldsymbol{y}}(t)\, \underline{\boldsymbol{g}}(t)^H\right).$$

Finally, left multiplying the complex conjugate transpose of this last equation by $\boldsymbol{\Theta}(t)^H$ and right multiplying it by $\boldsymbol{\Theta}(t)^{-H}$ yields recursion (34). ■

C. Recursion for the matrix $\boldsymbol{W}(t)$

Next, proposition 5.3 introduces a fast update for the subspace weighting matrix.

Proposition 5.3: If the $p \times p$ matrix $\beta \underline{J}^{-1} + \underline{\boldsymbol{y}}(t)^H \underline{\boldsymbol{h}}(t)$ is non-singular, $\boldsymbol{W}(t)$ satisfies the recursion

$$\boldsymbol{W}(t) = \left(\boldsymbol{W}(t-1) + \underline{\boldsymbol{e}}(t)\, \underline{\boldsymbol{g}}(t)^H\right) \boldsymbol{\Theta}(t) \quad (37)$$

where $\underline{\boldsymbol{e}}(t)$ is the $n \times p$ matrix

$$\underline{\boldsymbol{e}}(t) \triangleq \underline{\boldsymbol{x}}(t) - \boldsymbol{W}(t-1)\, \underline{\boldsymbol{y}}(t). \quad (38)$$

Proof:

Substituting equation (10) into equation (17) and right multiplying by $\boldsymbol{\Theta}(t)$ shows that $\boldsymbol{W}(t)$ satisfies the recursion

$$\boldsymbol{W}(t)\boldsymbol{S}(t)^H = \left(\beta \boldsymbol{W}(t-1)\boldsymbol{S}(t-1)^H + \underline{\boldsymbol{x}}(t)\, \underline{J}\, \widehat{\underline{\boldsymbol{y}}}(t)^H\right) \boldsymbol{\Theta}(t)$$

Substituting equations (36) and (38) into the above equation yields

$$\boldsymbol{W}(t)\boldsymbol{S}(t)^H = \boldsymbol{W}(t-1)\boldsymbol{\Theta}(t)\boldsymbol{S}(t)^H + \underline{\boldsymbol{e}}(t)\, \underline{J}\, \widehat{\underline{\boldsymbol{y}}}(t)^H \boldsymbol{\Theta}(t). \quad (39)$$

However, left multiplying (36) by $\underline{\boldsymbol{g}}(t)^H$ and replacing $\underline{\boldsymbol{g}}(t)$ by its definition in equation (35) leads to

$$\begin{aligned} \underline{\boldsymbol{g}}(t)^H \boldsymbol{\Theta}(t) \boldsymbol{R}(t) = {} & \left(\beta \underline{J}^{-1} + \underline{\boldsymbol{y}}(t)^H \underline{\boldsymbol{h}}(t)\right)^{-H} \\ & \left(\left(\beta \boldsymbol{S}(t-1) \underline{\boldsymbol{h}}(t)\right)^H + \left(\underline{\boldsymbol{y}}(t)^H \underline{\boldsymbol{h}}(t)\right)^H \underline{J}\, \widehat{\underline{\boldsymbol{y}}}(t)^H\right). \end{aligned} \quad (40)$$

Then equations (32) and (30) show that

$$\left(\beta \boldsymbol{S}(t-1) \underline{\boldsymbol{h}}(t)\right)^H = \beta\, \widehat{\underline{\boldsymbol{y}}}(t)^H = \beta \underline{J}^{-1}\, \underline{J}\, \widehat{\underline{\boldsymbol{y}}}(t)^H. \quad (41)$$

Substituting equation (41) into equation (40) yields

$$\underline{\boldsymbol{g}}(t)^H\, \boldsymbol{\Theta}(t)\, \boldsymbol{R}(t) = \underline{J}\, \widehat{\underline{\boldsymbol{y}}}(t)^H. \quad (42)$$

Finally, substituting equation (42) into equation (39) and right multiplying by $\boldsymbol{S}(t)^{-H} = \boldsymbol{Z}(t)^H$ yields equation (37). ■

Note that if $\beta\, \underline{J}^{-1} + \underline{\boldsymbol{y}}(t)^H \underline{\boldsymbol{h}}(t)$ is singular, $\boldsymbol{Z}(t)$ and $\boldsymbol{W}(t)$ can no longer be updated with equations (34) and (37). In practice, we never encountered this rank deficiency case in our numerical simulations[6].

[6] A solution consists in computing $\boldsymbol{W}(t)$ and $\boldsymbol{R}(t)$ by means of a SVD or a QR factorization of $\boldsymbol{C}_{xy}(t)$. Then $\boldsymbol{\Theta}(t) = \boldsymbol{W}(t-1)^H \boldsymbol{W}(t)$ can be deduced. Note that the whole processing requires $O(nr^2)$ operations; this technique must be used while $\boldsymbol{R}(t)$ or $\boldsymbol{\Theta}(t)$ remains singular. When both $\boldsymbol{R}(t)$ and $\boldsymbol{\Theta}(t)$ become non-singular again, then $\boldsymbol{Z}(t)$ can be computed, and the algorithm can switch back to the fully adaptive processing.

TABLE I
EXPONENTIAL WINDOW API ALGORITHM

Initialization :[8]

$$\boldsymbol{W}(0) = \begin{bmatrix} \boldsymbol{I}_r \\ \boldsymbol{0}_{(n-r)\times r} \end{bmatrix},\ \boldsymbol{Z}(0) = \boldsymbol{I}_r$$

For each time step do

Input vector : $\boldsymbol{x}(t)$

		Cost
PAST main section		
$\boldsymbol{y}(t) = \boldsymbol{W}(t-1)^H \boldsymbol{x}(t)$	(11)	nr
$\boldsymbol{h}(t) = \boldsymbol{Z}(t-1)\,\boldsymbol{y}(t)$	(32)	r^2
$\boldsymbol{g}(t) = \frac{\boldsymbol{h}(t)}{\beta + \boldsymbol{y}(t)^H \boldsymbol{h}(t)}$	(35)	$2r$
API main section		
$\boldsymbol{e}(t) = \boldsymbol{x}(t) - \boldsymbol{W}(t-1)\,\boldsymbol{y}(t)$	(38)	nr
$\boldsymbol{\Theta}(t) = \left(\boldsymbol{I}_r + \|\boldsymbol{e}(t)\|^2 \boldsymbol{g}(t)\,\boldsymbol{g}(t)^H\right)^{-\frac{1}{2}}$	(43)	$n + O(r^3)$
$\boldsymbol{Z}(t) = \begin{matrix} \frac{1}{\beta}\boldsymbol{\Theta}(t)^H \left(\boldsymbol{I}_r - \boldsymbol{g}(t)\,\boldsymbol{y}(t)^H\right) \\ \boldsymbol{Z}(t-1)\,\boldsymbol{\Theta}(t)^{-H} \end{matrix}$	(34)	$O(r^3)$
$\boldsymbol{W}(t) = \left(\boldsymbol{W}(t-1) + \boldsymbol{e}(t)\,\boldsymbol{g}(t)^H\right)\boldsymbol{\Theta}(t)$	(37)	$nr^2 + nr$

Since $\boldsymbol{W}(t-1)$ is orthonormal, $\underline{\boldsymbol{e}}(t)$ is orthogonal to $\boldsymbol{W}(t-1)$. Moreover, the orthonormality of $\boldsymbol{W}(t)$, associated to equation (37), yields

$$\boldsymbol{\Theta}(t)\,\boldsymbol{\Theta}(t)^H = \left(\boldsymbol{I}_r + \underline{\boldsymbol{g}}(t)\,\left(\underline{\boldsymbol{e}}(t)^H \underline{\boldsymbol{e}}(t)\right)\,\underline{\boldsymbol{g}}(t)^H\right)^{-1}. \quad (43)$$

Therefore, $\boldsymbol{\Theta}(t)$ is an inverse square root of the $r \times r$ positive definite matrix $\boldsymbol{I}_r + \underline{\boldsymbol{g}}(t)\,\left(\underline{\boldsymbol{e}}(t)^H \underline{\boldsymbol{e}}(t)\right)\,\underline{\boldsymbol{g}}(t)^H$. The choice of this inverse square root does not affect the subspace tracking performance[7].

The pseudo-code of the exponential window API algorithm is presented in table I, and that of the truncated window API algorithm (TW-API) is presented in table II. It can be noted that the first section of API is exactly the same as that of the PAST subspace tracker [22]; it requires only $nr + r^2 + O(r)$ operations per time step, while the rest of the algorithm has a $nr^2 + o(nr^2)$ computational complexity. In the same way, the first section of TW-API is similar to that of the sliding window version of PAST [29]; it requires only $2nr + 2r^2 + O(r)$, while the rest of the algorithm has a $(n+l)r^2 + o(nr^2)$ computational complexity. Note that the implementations of API and TW-API presented in tables I and II are of limited interest, since a number of faster subspace trackers have already been proposed in the literature, which have a $O(nr)$ complexity (among which [3], [22], [23], [25], [29], [34] are illustrated in section VIII). A faster implementation of API and TW-API is proposed in section VI.

[7]Let $\boldsymbol{\Theta}^P(t)$ be the only positive definite inverse square root. Then $\boldsymbol{\Theta}(t)$ can be written in the form

$$\boldsymbol{\Theta}(t) = \boldsymbol{\Theta}^P(t)\,\boldsymbol{U}(t) \quad (44)$$

where $\boldsymbol{U}(t)$ is a $r \times r$ orthonormal matrix. Substituting equation (44) into equation (37) yields

$$\boldsymbol{W}(t) = \left\{\left(\boldsymbol{W}(t-1) + \boldsymbol{e}(t)\,\boldsymbol{g}(t)^H\right)\boldsymbol{\Theta}^P(t)\right\}\boldsymbol{U}(t).$$

It can be readily seen in this last equation that $\boldsymbol{U}(t)$ does not affect the subspace spanned by $\boldsymbol{W}(t)$; it only affects the particular orthonormal basis $\boldsymbol{W}(t)$ of this subspace. Consequently, the choice of a particular inverse square root $\boldsymbol{\Theta}(t)$ has no impact on the subspace tracking performance.

[8]The initial values $\boldsymbol{W}(0)$ and $\boldsymbol{Z}(0)$ have to be chosen suitably:

- $\boldsymbol{W}(0)$ should be a $n \times r$ orthonormal matrix,
- $\boldsymbol{Z}(0)$ should be a $r \times r$ positive definite matrix.

Both matrices can be calculated from an initial block of data or from arbitrary initial data. The simplest way, however, is to set $\boldsymbol{W}(0) = \begin{bmatrix} \boldsymbol{I}_r \\ \boldsymbol{0}_{(n-r)\times r} \end{bmatrix}$ and $\boldsymbol{Z}(0) = \boldsymbol{I}_r$. The choice of these initial values affects the transient behavior but not the steady state performance of the algorithm.

TABLE II
TRUNCATED WINDOW API (TW-API) ALGORITHM

Initialization :

$$\boldsymbol{W}(0) = \begin{bmatrix} \boldsymbol{I}_r \\ \boldsymbol{0}_{(n-r)\times r} \end{bmatrix},\ \boldsymbol{Z}(0) = \boldsymbol{I}_r,\ \boldsymbol{X}(0) = \boldsymbol{0}_{n\times l},\ \widehat{\boldsymbol{V}}(0) = \boldsymbol{0}_{r\times l}$$

For each time step do

Input vector : $\boldsymbol{x}(t)$

		Cost
Section similar to SW − PAST		
$\left[\ \boldsymbol{x}(t-l)\ \middle|\ \boldsymbol{X}(t)\ \right] = \left[\ \boldsymbol{X}(t-1)\ \middle|\ \boldsymbol{x}(t)\ \right]$	(20)	
$\boldsymbol{y}(t) = \boldsymbol{W}(t-1)^H \boldsymbol{x}(t)$	(11)	nr
$\left[\ \widehat{\boldsymbol{v}}(t-l)\ \middle|\ \widehat{\boldsymbol{Y}}(t)\ \right] = \left[\ \widehat{\boldsymbol{V}}(t-1)\ \middle|\ \boldsymbol{y}(t)\ \right]$	(24)	
$\boldsymbol{v}(t-l) = \boldsymbol{W}(t-1)^H \boldsymbol{x}(t-l)$	(22)	nr
$\underline{\boldsymbol{x}}(t) = \left[\ \boldsymbol{x}(t)\ \middle|\ \boldsymbol{x}(t-l)\ \right]$	(8)	
$\underline{\widehat{\boldsymbol{y}}}(t) = \left[\ \boldsymbol{y}(t)\ \middle|\ \widehat{\boldsymbol{v}}(t-l)\ \right]$	(26)	
$\underline{\boldsymbol{y}}(t) = \left[\ \boldsymbol{y}(t)\ \middle|\ \boldsymbol{v}(t-l)\ \right]$	(27)	
$\underline{\boldsymbol{h}}(t) = \boldsymbol{Z}(t-1)\,\underline{\widehat{\boldsymbol{y}}}(t)$	(32)	$2r^2$
$\underline{\boldsymbol{g}}(t) = \underline{\boldsymbol{h}}(t)\,\left(\beta\,\underline{J}^{-1} + \underline{\boldsymbol{y}}(t)^H \underline{\boldsymbol{h}}(t)\right)^{-1}$	(35)	$8r$
TW − API main section		
$\underline{\boldsymbol{e}}(t) = \underline{\boldsymbol{x}}(t) - \boldsymbol{W}(t-1)\,\underline{\boldsymbol{y}}(t)$	(38)	$2nr$
$\boldsymbol{\Theta}(t) = \left(\boldsymbol{I}_r + \underline{\boldsymbol{g}}(t)\,\left(\underline{\boldsymbol{e}}(t)^H \underline{\boldsymbol{e}}(t)\right)\,\underline{\boldsymbol{g}}(t)^H\right)^{-\frac{1}{2}}$	(43)	$4n + O(r^3)$
$\boldsymbol{Z}(t) = \begin{matrix} \frac{1}{\beta}\boldsymbol{\Theta}(t)^H \left(\boldsymbol{I}_r - \underline{\boldsymbol{g}}(t)\,\underline{\boldsymbol{y}}(t)^H\right) \\ \boldsymbol{Z}(t-1)\,\boldsymbol{\Theta}(t)^{-H} \end{matrix}$	(34)	$O(r^3)$
$\boldsymbol{W}(t) = \left(\boldsymbol{W}(t-1) + \underline{\boldsymbol{e}}(t)\,\underline{\boldsymbol{g}}(t)^H\right)\boldsymbol{\Theta}(t)$	(37)	$nr^2 + 2nr$
$\widehat{\boldsymbol{V}}(t) = \boldsymbol{\Theta}(t)^H \widehat{\boldsymbol{Y}}(t)$	(23)	lr^2

VI. FAST API METHOD

In this section, a fast implementation of the API method is proposed, based on a particular choice of the matrix $\boldsymbol{\Theta}(t)$. It is supposed that $\beta\,\underline{J}^{-1} + \underline{\boldsymbol{y}}(t)^H \underline{\boldsymbol{h}}(t)$ is non-singular, so that $\boldsymbol{\Theta}(t)$ is also non-singular. Below, the $p \times p$ identity matrix is denoted $\underline{I}_p$.

A. A particular solution to equation (43)

Let $\underline{\varepsilon}(t)$ be a square root of the $p \times p$ matrix $\underline{\boldsymbol{e}}(t)^H \underline{\boldsymbol{e}}(t)$:

$$\underline{\varepsilon}(t)\,\underline{\varepsilon}(t)^H = \left(\underline{\boldsymbol{e}}(t)^H \underline{\boldsymbol{e}}(t)\right). \quad (45)$$

Substituting equation (45) into equation (43) and applying the matrix inversion lemma shows that[9]

$$\boldsymbol{\Theta}(t)\,\boldsymbol{\Theta}(t)^H = \boldsymbol{I}_r - \underline{\boldsymbol{g}}(t)\,\underline{\varepsilon}(t)\,\underline{\rho}(t)^{-1}\underline{\varepsilon}(t)^H \underline{\boldsymbol{g}}(t)^H \quad (46)$$

where $\underline{\rho}(t)$ is the $p \times p$ positive definite matrix

$$\underline{\rho}(t) = \underline{I}_p + \underline{\varepsilon}(t)^H\,\left(\underline{\boldsymbol{g}}(t)^H \underline{\boldsymbol{g}}(t)\right)\underline{\varepsilon}(t). \quad (47)$$

Considering equation (46), we are looking for a special solution of the form

$$\boldsymbol{\Theta}(t) = \boldsymbol{I}_r - \underline{\boldsymbol{g}}(t)\,\underline{\varepsilon}(t)\,\underline{\sigma}(t)^{-1}\underline{\varepsilon}(t)^H \underline{\boldsymbol{g}}(t)^H \quad (48)$$

[9]Lemma 5.2 is applied with $\boldsymbol{A} = \boldsymbol{I}_r$, $\boldsymbol{P} = \underline{\boldsymbol{g}}(t)\,\underline{\varepsilon}(t)$, $\boldsymbol{J} = \underline{I}_p$ and $\boldsymbol{Q} = \underline{\varepsilon}(t)^H \underline{\boldsymbol{g}}(t)^H$. In particular, the non-singularity of $\boldsymbol{\Theta}(t)$ is equivalent to the non-singularity of $\underline{\rho}(t)$.

where $\underline{\sigma}(t)$ is a $p \times p$ non-singular matrix. The interest of this approach is that the $r \times r$ matrix square rooting problem in equation (46) is converted into a smaller $p \times p$ matrix square rooting. Indeed, substituting equation (48) into equation (46) yields a sufficient condition :

$$\underline{\sigma}(t)^{-1} + \underline{\sigma}(t)^{-H} + \underline{\sigma}(t)^{-1} \left(\underline{I}_p - \underline{\rho}(t)\right) \underline{\sigma}(t)^{-H} = \underline{\rho}(t)^{-1}.$$

Left multiplying the two members of this last equation by $\underline{\sigma}(t)$ and right multiplying them by $\underline{\sigma}(t)^H$ yields the equation[10]

$$\left(\underline{\sigma}(t) - \underline{\rho}(t)\right) \underline{\rho}(t)^{-1} \left(\underline{\sigma}(t) - \underline{\rho}(t)\right)^H = \underline{I}_p$$

whose solution is

$$\underline{\sigma}(t) = \underline{\rho}(t) + \underline{\rho}(t)^{\frac{1}{2}^H}. \tag{49}$$

Even if other choices would be possible, from now on we suppose that the square root of $\underline{\rho}(t)$ which is involved in the above equation is the only positive definite square root. This condition guarantees that $\underline{\sigma}(t)$ is positive definite, so that $\boldsymbol{\Theta}(t)$ is hermitian[11]. Then define the $p \times p$ positive definite matrix

$$\underline{\tau}(t) = \underline{\varepsilon}(t)\, \sigma(t)^{-1} \underline{\varepsilon}(t)^H. \tag{50}$$

Substituting equation (50) into equation (48) yields

$$\boldsymbol{\Theta}(t) = \boldsymbol{I}_r - \underline{\boldsymbol{g}}(t)\, \underline{\tau}(t)\, \underline{\boldsymbol{g}}(t)^H. \tag{51}$$

B. Fast implementation of the particular solution

Based on the low-rank matrix update of $\boldsymbol{\Theta}(t)$ in equation (51), it is shown below that the matrices $\boldsymbol{Z}(t)$, $\boldsymbol{W}(t)$ and $\widehat{\boldsymbol{V}}(t)$ can also be efficiently updated. Consider the $p \times p$ matrix

$$\underline{\eta}(t) = \underline{I}_p - \left(\underline{\boldsymbol{g}}(t)^H \underline{\boldsymbol{g}}(t)\right) \underline{\tau}(t). \tag{52}$$

Since $\boldsymbol{\Theta}(t)$ is non-singular, the matrix inversion lemma shows that $\underline{\eta}(t)$ is also non-singular[12]. Then substituting equation (51) into equation (34) yields

$$\boldsymbol{Z}(t) = \frac{1}{\beta} \left(\boldsymbol{Z}(t-1) - \underline{\boldsymbol{g}}(t)\, \underline{\boldsymbol{h}}'(t)^H + \underline{\boldsymbol{\epsilon}}(t)\, \underline{\boldsymbol{g}}(t)^H\right) \tag{53}$$

where the $r \times p$ matrices $\underline{\boldsymbol{h}}'(t)$ and $\underline{\boldsymbol{\epsilon}}(t)$ are defined by

$$\begin{aligned}
\underline{\boldsymbol{y}}'(t) &= \underline{\boldsymbol{y}}(t)\, \underline{\eta}(t) + \underline{\boldsymbol{g}}(t)\, \underline{\tau}(t) && (54)\\
\underline{\boldsymbol{h}}'(t) &= \boldsymbol{Z}(t-1)^H \underline{\boldsymbol{y}}'(t) && (55)\\
\underline{\boldsymbol{\epsilon}}(t) &= \left(\boldsymbol{Z}(t-1)\underline{\boldsymbol{g}}(t) - \underline{\boldsymbol{g}}(t) \left(\underline{\boldsymbol{h}}'(t)^H \underline{\boldsymbol{g}}(t)\right)\right) \\
&\quad \left(\underline{\tau}(t)\, \underline{\eta}(t)^{-1}\right)^H && (56)
\end{aligned}$$

Then substituting equation (51) into equation (37) yields

$$\boldsymbol{W}(t) = \boldsymbol{W}(t-1) + \underline{\boldsymbol{e}}'(t)\, \underline{\boldsymbol{g}}(t)^H \tag{57}$$

where $\underline{\boldsymbol{e}}'(t)$ is the $n \times p$ matrix

$$\underline{\boldsymbol{e}}'(t) = \underline{\boldsymbol{e}}(t)\, \underline{\eta}(t) - \boldsymbol{W}(t-1)\, \underline{\boldsymbol{g}}(t)\, \underline{\tau}(t). \tag{58}$$

However, substituting equations (38) and (54) into equation (58) yields

$$\underline{\boldsymbol{e}}'(t) = \underline{\boldsymbol{x}}(t)\, \underline{\eta}(t) - \boldsymbol{W}(t-1)\, \underline{\boldsymbol{y}}'(t). \tag{59}$$

Finally, substituting equation (51) into equation (23) yields

$$\widehat{\boldsymbol{V}}(t) = \boldsymbol{Y}(t) - \underline{\boldsymbol{g}}(t) \left(\underline{\boldsymbol{g}}(t)\, \underline{\tau}(t)\right)^H \boldsymbol{Y}(t). \tag{60}$$

The pseudo-code of the exponential window fast API algorithm (FAPI) is presented in table III, and that of the truncated window fast API algorithm (TW-FAPI) is presented in table IV. The overall computational cost of FAPI is $n(3r+2) + 5r^2 + O(r)$ flops per iteration[13] (whereas the complexities of PAST [22] and OPAST [25] are respectively $3nr + 2r^2 + O(r)$ and $n(4r+1) + 2r^2 + O(r)$). The overall computational cost of TW-FAPI is $n(6r+8) + 4lr + O(r^2)$ flops per iteration[14] (whereas the complexities of SW-PAST and SW-OPAST [29] are respectively $5nr + 4r^2 + O(r)$ and $n(15r+28) + 12r^2 + O(r)$). Note that the presence of a $4lr$ term in the complexity of TW-FAPI may make this algorithm more computationally demanding in applications for which l is much larger than n. However, in the context of frequency estimation, it has been proved that optimal Cramer-Rao bounds were obtained for $\frac{1}{2}n \leq l \leq 2n$ [35], and in section VIII-A, TW-FAPI is tested with $l = \frac{3}{2}n$.

TABLE III
EXPONENTIAL WINDOW FAST API (FAPI) ALGORITHM

Initialization (*cf.* table I)

For each time step do

Input vector : $\boldsymbol{x}(t)$
PAST main section (*cf.* table I)

FAPI main section :		**Cost**
$\varepsilon^2(t) = \|\boldsymbol{x}(t)\|^2 - \|\boldsymbol{y}(t)\|^2$	(45)	$n + r$
$\tau(t) = \frac{\varepsilon^2(t)}{1+\varepsilon^2(t)\|\boldsymbol{g}(t)\|^2 + \sqrt{1+\varepsilon^2(t)\|\boldsymbol{g}(t)\|^2}}$	(50)	r
$\eta(t) = 1 - \tau(t)\, \|\boldsymbol{g}(t)\|^2$	(52)	1
$\boldsymbol{y}'(t) = \eta(t)\, \boldsymbol{y}(t) + \tau(t)\, \boldsymbol{g}(t)$	(54)	$2r$
$\boldsymbol{h}'(t) = \boldsymbol{Z}(t-1)^H \boldsymbol{y}'(t)$	(55)	r^2
$\boldsymbol{\epsilon}(t) = \frac{\tau(t)}{\eta(t)} \left(\boldsymbol{Z}(t-1)\boldsymbol{g}(t) - \left(\boldsymbol{h}'(t)^H \boldsymbol{g}(t)\right) \boldsymbol{g}(t)\right)$	(56)	$r^2 + 3r$
$\boldsymbol{Z}(t) = \frac{1}{\beta} \left(\boldsymbol{Z}(t-1) - \boldsymbol{g}(t)\, \boldsymbol{h}'(t)^H + \boldsymbol{\epsilon}(t)\, \boldsymbol{g}(t)^H\right)$	(53)	$2r^2$
$\boldsymbol{e}'(t) = \eta(t)\, \boldsymbol{x}(t) - \boldsymbol{W}(t-1)\, \boldsymbol{y}'(t)$	(59)	$nr + n$
$\boldsymbol{W}(t) = \boldsymbol{W}(t-1) + \boldsymbol{e}'(t)\, \boldsymbol{g}(t)^H$	(57)	nr

VII. LINK WITH THE PAST AND OPAST ALGORITHMS

In this section, it is shown that the classical exponential window PAST algorithm can be seen as a first order approximation of the FAPI algorithm. Indeed, the error $\boldsymbol{e}(t)$ is the component of $\boldsymbol{x}(t)$ that does not belong to the signal subspace spanned by $\boldsymbol{W}(t-1)$. Thus, if this subspace slowly varies upon time, and if the Signal to Noise Ratio (SNR) is high, $\boldsymbol{e}(t) \simeq \boldsymbol{0}$. If the second order term $\|\boldsymbol{e}(t)\|^2$ is disregarded in

[10] Remember that $\underline{\rho}(t)$ is an hermitian matrix.

[11] More precisely, $\boldsymbol{\Theta}(t)$ is positive definite. Indeed, equation (49) shows that $\underline{\sigma}(t)$ and $\underline{\rho}(t)$ are simultaneously diagonalizable, and the eigenvalues of $\underline{\sigma}(t)$ are strictly greater than those of $\underline{\rho}(t)$. Therefore, $\underline{\rho}(t)^{-1} - \underline{\sigma}(t)^{-1}$ is a positive definite matrix. Then subtracting equation (46) from equation (48) shows that $\boldsymbol{\Theta}(t)$ is positive definite.

[12] Lemma 5.2 is applied to equation (51), with $\boldsymbol{A} = \boldsymbol{I}_r$, $\boldsymbol{P} = \underline{\boldsymbol{g}}(t)\, \underline{\tau}(t)$, $\boldsymbol{J} = \underline{I}_p$ and $\boldsymbol{Q} = \underline{\boldsymbol{g}}(t)^H$.

[13] Note that this implementation of FAPI is faster than that proposed in [26], whose global cost was $n(4r+2) + 5r^2 + O(r)$.

[14] This implementation of TW-FAPI is also faster than that proposed in [31], whose global cost was $n(8r+8) + 4lr + O(r^2)$.

TABLE IV
TRUNCATED WINDOW FAST API (TW-FAPI) ALGORITHM

Initialization (*cf.* table II)

For each time step do

Section similar to SW − PAST (*cf.* table II)

TW − FAPI main section		**Cost**
$\underline{\varepsilon}(t) = \left(\underline{\boldsymbol{x}}(t)^H \underline{\boldsymbol{x}}(t) - \underline{\boldsymbol{y}}(t)^H \underline{\boldsymbol{y}}(t)\right)^{\frac{1}{2}}$	(45)	$4n + 4r$
$\underline{\rho}(t) = \underline{I}_p + \underline{\varepsilon}(t)^H \left(\underline{\boldsymbol{g}}(t)^H \underline{\boldsymbol{g}}(t)\right) \underline{\varepsilon}(t)$	(47)	$4r$
$\underline{\tau}(t) = \underline{\varepsilon}(t) \left(\underline{\rho}(t) + \underline{\rho}(t)^{\frac{1}{2}H}\right)^{-1} \underline{\varepsilon}(t)^H$	(50)	$O(1)$
$\underline{\eta}(t) = \underline{I}_p - \left(\underline{\boldsymbol{g}}(t)^H \underline{\boldsymbol{g}}(t)\right) \underline{\tau}(t)$	(52)	$O(1)$
$\underline{\boldsymbol{y}}'(t) = \underline{\boldsymbol{y}}(t)\, \underline{\eta}(t) + \underline{\boldsymbol{g}}(t)\, \underline{\tau}(t)$	(54)	$8r$
$\underline{\boldsymbol{h}}'(t) = \underline{\boldsymbol{Z}}(t-1)^H \underline{\boldsymbol{y}}'(t)$	(55)	$2r^2$
$\underline{\boldsymbol{\varepsilon}}(t) = \left(\underline{\boldsymbol{Z}}(t-1)\underline{\boldsymbol{g}}(t) - \underline{\boldsymbol{g}}(t)\left(\underline{\boldsymbol{h}}'(t)^H \underline{\boldsymbol{g}}(t)\right)\right) \left(\underline{\tau}(t)\, \underline{\eta}(t)^{-1}\right)^H$	(56)	$2r^2 + 12r$
$\underline{\boldsymbol{Z}}(t) = \frac{1}{\beta}\left(\underline{\boldsymbol{Z}}(t-1) - \underline{\boldsymbol{g}}(t)\, \underline{\boldsymbol{h}}'(t)^H + \underline{\boldsymbol{\varepsilon}}(t)\, \underline{\boldsymbol{g}}(t)^H\right)$	(53)	$4r^2$
$\underline{\boldsymbol{e}}'(t) = \underline{\boldsymbol{x}}(t)\, \underline{\eta}(t) - \boldsymbol{W}(t-1)\, \underline{\boldsymbol{y}}'(t)$	(59)	$2nr + 4n$
$\boldsymbol{W}(t) = \boldsymbol{W}(t-1) + \underline{\boldsymbol{e}}'(t)\, \underline{\boldsymbol{g}}(t)^H$	(57)	$2nr$
$\widehat{\boldsymbol{V}}(t) = \boldsymbol{Y}(t) - \underline{\boldsymbol{g}}(t) \left(\underline{\boldsymbol{g}}(t)\, \underline{\tau}(t)\right)^H \boldsymbol{Y}(t)$	(60)	$4lr$

table III, $\tau(t) = 0$, $\eta(t) = 1$ and $\boldsymbol{\Theta}(t)$ becomes the $r \times r$ identity matrix. Then equations (57) and (53) become

$$\boldsymbol{W}(t) = \boldsymbol{W}(t-1) + \boldsymbol{e}(t)\, \boldsymbol{g}(t)^H \tag{61}$$

$$\boldsymbol{Z}(t) = \frac{1}{\beta}\left(\boldsymbol{Z}(t-1) - \boldsymbol{g}(t)\, \boldsymbol{h}(t)^H\right) \tag{62}$$

(in particular, it can be recursively shown that $\boldsymbol{Z}(t)$ is always hermitian). Consequently, this first order approximation of the fast API method is an exact implementation of the classical PAST subspace tracker [22], which only provides a *nearly* orthonormal subspace weighting matrix. In other respects, a thorough examination of the OPAST algorithm presented in [25] shows that $\boldsymbol{W}(t)$ is updated as in equation (57) (which guarantees the orthonormality, contrary to equation (61)). However, $\boldsymbol{Z}(t)$ is updated as in equation (62). Consequently, OPAST can be seen as an intermediary between PAST and FAPI.

VIII. SIMULATION RESULTS

In this section, the performance of the subspace estimation is analyzed in the context of frequency estimation, in terms of the maximum principal angle between the true dominant subspace of the correlation matrix $\boldsymbol{C}_{xx}(t)$ (obtained via an exact eigenvalue decomposition), and the estimated dominant subspace of the same correlation matrix (obtained with the subspace tracker). This error criterion was initially proposed by P. Comon and G.H. Golub as a measure of the distance between equidimensional subspaces [24, pp. 603-604]). In section VIII-A, the FAPI and TW-FAPI algorithms are compared to other existing subspace trackers. In section VIII-B, the behavior of the API method regarding the SNR and the parameters n and r is investigated.

A. Comparison of FAPI and TW-FAPI with other existing subspace trackers

In this section, the test signal is a sum of $r = 4$ complex sinusoidal sources plus a complex white gaussian noise (the SNR is 5.7 dB). The frequencies of the sinusoids vary according to a jump scenario originally proposed by P. Strobach in the context of Direction Of Arrival estimation [36]: their values abruptly change at different time instants, between which they remain constant. Their variations are represented on Figure 1-a. This signal is processed in section VIII-A1 by means of an exponential window whose forgetting factor is $\beta \approx 0.99$, and in section VIII-A2 by means of a sliding window of length $l = 120$. This parameters were chosen so that the effective window length is the same in both cases, *i.e.* $\beta = \frac{1}{1-1/l}$. Section VIII-A3 focuses on the orthonormality of the subspace weighting matrix. The complexities of the various subspace trackers illustrated in this section are given in table V.

Algorithm	Complexity (flops)	Window	Figure
FAPI	$n(3r+2) + 5r^2 + O(r)$	exponential	Fig. 1
PAST	$3nr + 2r^2 + O(r)$		
NIC	$4nr + 2r^2 + O(r)$		
OPAST	$n(4r+1) + 2r^2 + O(r)$		
Karasalo	$nr^2 + n(3r+2) + O(r^3)$	exponential	Fig. 2
FST	$n(6r+2) + 12r^2 + O(r)$		
Householder PAST	$n(4r+1) + 2r^2 + O(r)$		
Loraf2	$nr^2 + n(3r+2) + O(r^3)$		
SP1	$4nr^2 + n(4r+2) + O(r^3)$		
TW-FAPI	$n(6r+8) + 4lr + O(r^2)$	sliding	Fig. 3
SW-PAST	$5nr + 4r^2 + O(r)$		
SW-NIC	$6nr + 4r^2 + O(r)$		
SW-OPAST	$n(15r+28) + 12r^2 + O(r)$		

TABLE V
COMPARISON OF THE COMPLEXITIES

1) Exponential window case: figure 1-b shows the maximum principal angle error trajectory $\theta_{\mathrm{FAPI}}(t)$, obtained with the FAPI method with parameters $n = 80$ and $\beta \approx 0.99$. Then this result is compared to that obtained with the PAST subspace tracker: figure 1-c shows the ratio in dB of the trajectories obtained with FAPI and PAST, *i.e.*

$$20 \log_{10}\left(\frac{\theta_{\mathrm{FAPI}}(t)}{\theta_{\mathrm{PAST}}(t)}\right).$$

At initialization, it can be noticed that FAPI converges faster than PAST. Moreover, PAST does not provide an orthonormal subspace weighting matrix. Figure 1-d shows the ratio in dB of the trajectories obtained with FAPI and the NIC subspace tracker[15], which is a robust generalization of PAST [23]. It can be seen that the subspace estimation error is always smaller with FAPI. As PAST, NIC does not guarantee the orthonormality of the subspace weighting matrix. Figure 1-e shows the ratio of the trajectories obtained with FAPI and OPAST. The two algorithms reach the same performance, except at initialization, where FAPI converges faster. In fact, the difference is much more distinct with the sliding window versions of these algorithms (see section VIII-A2).

In figure 2, the FAPI algorithm is compared to five other well-known subspace trackers :

[15]The learning step η is equal to 0.7.

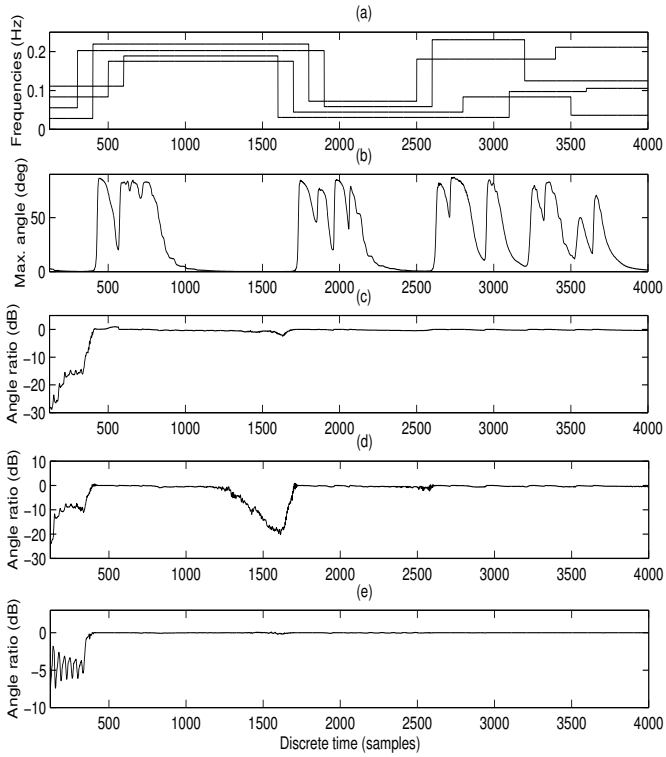

Fig. 1. Subspace tracking based on an exponential window
(a) Normalized frequencies of the sinusoids
(b) Maximum principal angle trajectory obtained with FAPI
(c) Ratio of the trajectories obtained with FAPI and PAST
(d) Ratio of the trajectories obtained with FAPI and NIC
(e) Ratio of the trajectories obtained with FAPI and OPAST

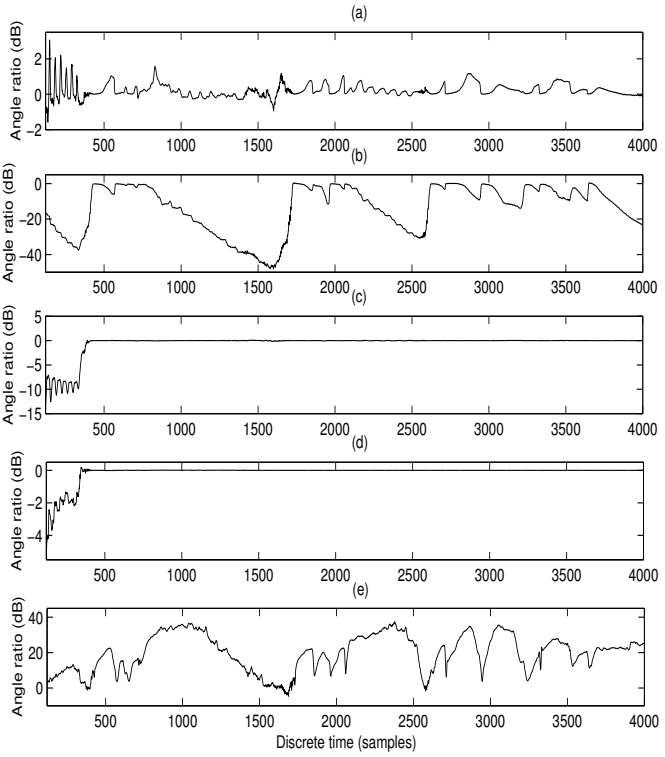

Fig. 2. Subspace tracking based on an exponential window
(a) Ratio of the trajectories obtained with FAPI and Karasalo
(b) Ratio of the trajectories obtained with FAPI and FST
(c) Ratio of the trajectories obtained with FAPI and Householder PAST
(d) Ratio of the trajectories obtained with FAPI and Loraf2
(e) Ratio of the trajectories obtained with FAPI and SP1

- I. Karasalo's algorithm [2],
- the Fast Subspace Tracking (FST) algorithm [3],
- the novel PAST algorithm employing Householder transformations, herein called Householder PAST [34],
- the Low-Rank Adaptive Filter (Loraf2) algorithm [7],
- and the Subspace Projection (SP1) algorithm [37].

Figure 2-a shows that the behaviors of FAPI and Karasalo's algorithm are very similar. However the dominant cost of the latter is nr^2 (see table V). Figure 2-b shows that FAPI converges to the signal subspace much more precisely than FST. Moreover, FST is more computationally demanding than FAPI. Figure 2-c shows that FAPI and Householder PAST reach the same performance, except at initialization, where FAPI converges faster. Figure 2-d shows that the same remark can be made about FAPI and Loraf2. Besides, the dominant complexity of Loraf2 is nr^2.

Among the various subspace trackers that we have tested, SP1 is the only one which really outperformed FAPI (see figure 2-e). However, table V shows that SP1 is the most computationally demanding algorithm. In other respects, it is only suitable for time series data analysis, and was only designed for exponential windows.

2) Sliding window case: figure 3-a shows the maximum principal angle error trajectory $\theta_{\mathrm{TW-FAPI}}(t)$, obtained with the TW-FAPI method with parameters $\beta = 1$ (which turns the truncated window into a sliding window), $n = 80$ and $l = 120$. It can be noticed that this algorithm has a fast convergence rate after each frequency jump. This result can be compared to that of figure 1-b, obtained with the exponential window FAPI method, for which the response to frequency jumps is slower, because of the nature of the window which tends to smooth the signal variations. Figure 3-b shows the ratio in dB of the trajectories obtained with TW-FAPI and the sliding window version of PAST, herein called SW-PAST [22], [29]. It can be seen that TW-FAPI converges faster than SW-PAST at initialization. Note that as PAST, SW-PAST does not provide an orthonormal subspace weighting matrix. Figure 3-c shows the ratio in dB of the trajectories obtained with TW-FAPI and a sliding window version of the NIC algorithm, herein called SW-NIC[16]. Finally, figure 3-d shows the ratio in dB of the trajectories obtained with TW-FAPI and the sliding window OPAST algorithm [29]. It can be noticed that the maximum principal angle error trajectory obtained with TW-FAPI is about 20 dB lower than those obtained with SW-NIC and SW-OPAST in regions where the frequencies are constant.

3) Orthonormality error: the orthonormality of the subspace weighting matrix $\boldsymbol{W}(t)$ can be measured by means of the following error criterion:

$$20 \log_{10} \left(\| \boldsymbol{W}(t)^H \boldsymbol{W}(t) - \boldsymbol{I}_r \|_F \right).$$

Table VI shows the maximum orthonormality error reached by the above mentioned algorithms while tracking the test

[16]SW-NIC is also implemented with $\eta = 0.7$.

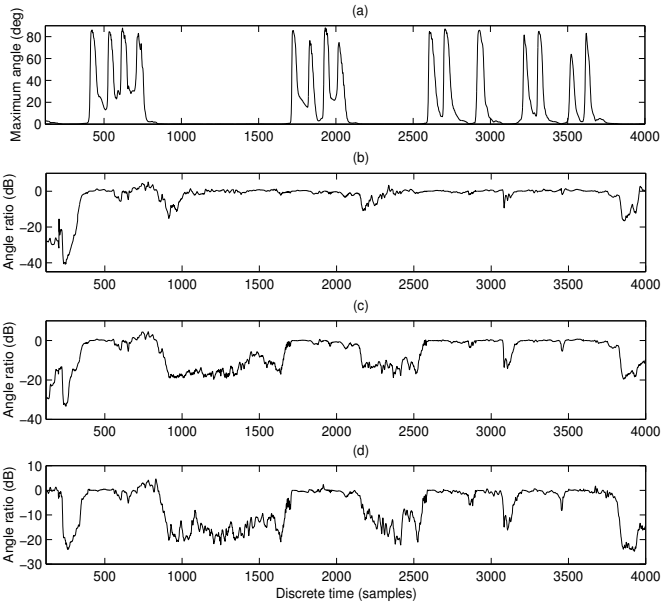

Fig. 3. Subspace tracking based on a sliding window
(a) Maximum principal angle trajectory obtained with TW-FAPI
(b) Ratio of the trajectories obtained with TW-FAPI and SW-PAST
(c) Ratio of the trajectories obtained with TW-FAPI and SW-NIC
(d) Ratio of the trajectories obtained with TW-FAPI and SW-OPAST

Algorithms	Orthonormality error
FAPI, TW-FAPI, OPAST, Householder PAST	about -300 dB
Karasalo, FST, Loraf2	about -280 dB
SP1, SW-OPAST	about -240 dB
PAST, NIC	about -25 dB
SW-PAST, SW-NIC	about -5 dB

TABLE VI
MAXIMUM ORTHONORMALITY ERROR

signal variations. We observed that FAPI, TW-FAPI, OPAST and Householder PAST outperformed all the other algorithms, whereas PAST, NIC, and their sliding window versions do not guarantee the orthonormality of the subspace weighting matrix.

B. Behavior of the API method regarding the SNR and the parameters n and r

In this section, the test signal is still a sum of $r = 4$ complex sinusoidal sources plus a complex white gaussian noise. However, the frequencies of the sinusoids are constant, equal to the initial values given in figure 1-a.

1) Influence of the SNR: in this section, the effect of the SNR onto the subspace estimation is investigated. To this end, the noise part of the test signal was synthesized so that the SNR varies linearly from +30 dB to -30 dB (see figure 4-a).

Figure 4-b shows the maximum principal angle error trajectory obtained with the FAPI method with parameters $n = 80$ and $\beta \approx 0.99$. It can be seen that the performance of the subspace estimation collapses beyond $n \simeq 2600$. Figure 4-a shows that from this time instant the SNR is lower than -10 dB. Figure 4-c shows the maximum principal angle error trajectory obtained with the TW-FAPI method with parameters $\beta = 1$, $n = 80$ and $l = 120$. Again, the performance of the subspace estimation collapses beyond $n \simeq 2600$. Although they are not illustrated here, we observed that the performance of all the above mentioned subspace trackers similarly collapse beyond the same SNR limit (-10 dB).

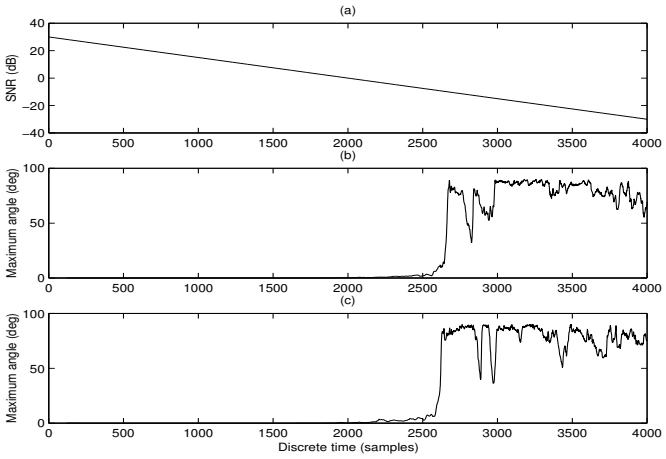

Fig. 4. Influence of the Signal to Noise Ratio
(a) Signal to Noise Ratio in dB
(b) Maximum principal angle trajectory obtained with FAPI
(c) Maximum principal angle trajectory obtained with TW-FAPI

2) Influence of the ratio n/r: in this section, we focus on the influence of the ratio n/r onto the subspace estimation. The SNR is constant, equal to 5.7 dB.

Figure 5-a shows the mean of $\theta_{\mathrm{FAPI}}(t)$, as a function of the ratio n/r, for all $n \in \{r+1, \ldots, 80\}$ (with $\beta \approx 0.99$). It can be seen that the subspace estimation becomes reliable as soon as $n/r \geq 7$. Figure 5-b shows the mean of $\theta_{\mathrm{TW-FAPI}}(t)$, as a function of the ratio n/r, for all $n \in \{r+1, \ldots, 80\}$ (with $\beta = 1$ and $l = 120$). Again, it can be seen that the subspace estimation becomes reliable as soon as $n/r \geq 7$. Although they are not illustrated here, we observed that the same remark is valid for all the above mentioned subspace trackers.

3) Tracking a subspace of wrong dimension: since the dimension r of the signal subspace is unknown in many applications, we investigate in this section the performance of the FAPI and TW-FAPI algorithms when applied with a wrong subspace dimension r. The SNR is constant, equal to 5.7 dB. The performance of the subspace estimation is analyzed in terms of the maximum principal angle between the true 4-dimensional signal subspace and the estimated r-dimensional subspace.

Figure 5-c shows the mean of $\theta_{\mathrm{FAPI}}(t)$, as a function of r, for all $r \in \{1, \ldots, 20\}$ (with parameters $\beta \approx 0.99$ and $n = 80$). Similarly, figure 5-d shows the mean of $\theta_{\mathrm{TW-FAPI}}(t)$, as a function of r, for all $r \in \{1, \ldots, 20\}$ (with parameters $l = 120$ and $n = 80$). It can be seen that the subspace estimation is reliable in all cases:

- if $r = 4$, the maximum principal angle is very low (as expected),
- if $r < 4$, the maximum principal angle remains low, which means that the estimated lower-dimensional subspace is nearly included in the true signal subspace,
- if $r > 4$, the maximum principal angle is even lower than in the case $r = 4$, which means that the true

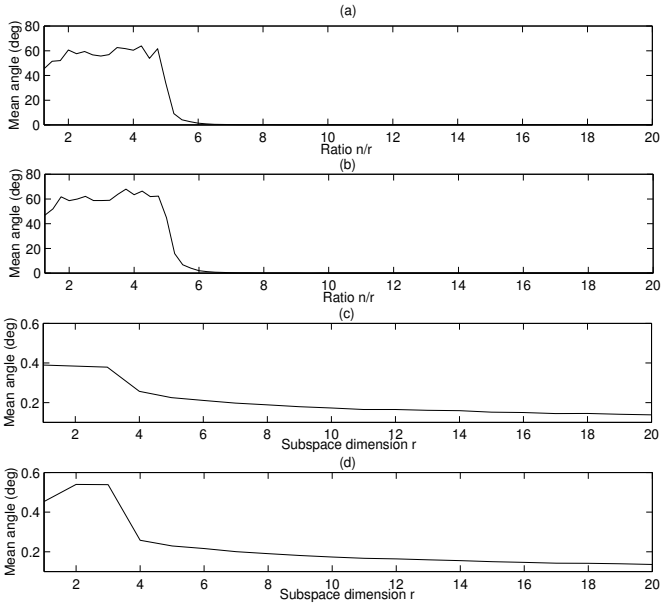

Fig. 5. Influence of the parameters n and r
(a) Average max. angle obtained with FAPI as a function of n/r
(b) Average max. angle obtained with TW-FAPI as a function of n/r
(c) Average max. angle obtained with FAPI as a function of r
(d) Average max. angle obtained with TW-FAPI as a function of r

signal subspace is nearly included in the estimated upper-dimensional subspace. Moreover, it can be noticed that the maximum principal angle decreases as the dimension of the estimated subspace increases.

We can conclude that FAPI and TW-FAPI are robust to erroneous subspace dimension r.

IX. Conclusions

In this paper, several implementations of the API algorithm for subspace tracking were presented, based either on exponential windows or on truncated windows. These algorithms reach a linear complexity and guarantee the orthonormality of the subspace weighting matrix at each time step. In the context of frequency estimation, the method proves able to track abrupt frequency variations robustly, and outperforms many subspace trackers, both in terms of subspace estimation and computational complexity. Finally, these subspace tracking algorithms can be considered as the starting point of a real-time frequency tracker, whose full implementation can involve our adaptive version of the ESPRIT algorithm [38].

References

[1] P. Comon and G. H. Golub, "Tracking a few extreme singular values and vectors in signal processing," in *Proc. IEEE*, vol. 78, Aug. 1990, pp. 1327–1343.

[2] I. Karasalo, "Estimating the covariance matrix by signal subspace averaging," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 34, pp. 8–12, Feb. 1986.

[3] D. J. Rabideau, "Fast, rank adaptive subspace tracking and applications," *IEEE Trans. Signal Processing*, vol. 44, no. 9, pp. 2229–2244, Sept. 1996.

[4] M. Moonen, P. V. Dooren, and J. Vandewalle, "An SVD updating algorithm for subspace tracking," *SIAM J. Matrix Ana. Appl.*, vol. 13, no. 4, pp. 1015–1038, 1992.

[5] E. M. Dowling, L. P. Ammann, and R. D. DeGroat, "A TQR-iteration based adaptive SVD for real time angle and frequency tracking," *IEEE Trans. Signal Processing*, vol. 42, no. 4, pp. 914–926, Apr. 1994.

[6] P. Strobach, "Low-rank adaptive filters," *IEEE Trans. Signal Processing*, vol. 44, no. 12, pp. 2932–2947, Dec. 1996.

[7] ——, "Bi-iteration SVD subspace tracking algorithms," *IEEE Trans. Signal Processing*, vol. 45, no. 5, pp. 1222–1240, May 1997.

[8] Y. Hua, Y. Xiang, T. Chen, K. Abed-Meraim, and Y. Miao, "A new look at the power method for fast subspace tracking," *Digital Signal Processing*, Oct. 1999.

[9] C. H. Bischof and G. M. Shroff, "On updating signal subspaces," *IEEE Trans. Signal Processing*, vol. 40, pp. 96–105, 1992.

[10] G. W. Stewart, "An updating algorithm for subspace tracking," *IEEE Trans. Signal Processing*, vol. 40, pp. 1535–1541, June 1992.

[11] G. Xu, H. Zha, G. H. Golub, and T. Kailath, "Fast algorithms for updating signal subspaces," *IEEE Trans. Circuits Syst.*, vol. 41, no. 8, pp. 537–549, Aug. 1994.

[12] R. D. DeGroat, "Noniterative subspace tracking," *IEEE Trans. Signal Processing*, vol. 40, no. 3, pp. 571–577, Mar. 1992.

[13] T. Chonavel, B. Champagne, and C. Riou, "Fast adaptive eigenvalue decomposition: a maximum likelihood approach," *Signal processing*, vol. 83, no. 2, pp. 307–324, Feb. 2003.

[14] C. S. MacInnes, "Fast, accurate subspace tracking using operator restriction analysis," in *Proc. of IEEE Int. Conf. on Acoustic, Speech and Signal Processing*, 1998, pp. 1357–1360.

[15] B. Champagne, "SVD-updating via constrained perturbations with application to subspace tracking," *Signals, Systems and Computers*, vol. 2, pp. 1379–1385, 1996.

[16] E. Oja, "Neural networks, principal components and subspaces," *Int. journal of neural systems*, vol. 1, no. 1, pp. 61–68, 1989.

[17] L. Xu, "Least mean square error reconstruction principle for selg-organizing neural nets," *Neural Networks*, vol. 6, pp. 627–648, 1993.

[18] T. Chen and S. Amari, "Unified stabilization approach to principal and minor components extraction algorithms," *Neural Networks*, vol. 14, no. 10, pp. 1377–1387, 2001.

[19] S. Y. Kung, K. I. Diamantaras, and J. S. Taur, "Adaptive principal component extraction (apex) and applications," *IEEE Trans. Signal Processing*, vol. 43, no. 1, pp. 1202–1217, Jan. 1995.

[20] G. Mathew and V. U. Reddy, "Adaptive estimation of eigensubspace," *IEEE Trans. Signal Processing*, vol. 43, no. 2, pp. 401–411, Feb. 1995.

[21] Z. Fu and E. M. Dowling, "Conjugate gradient eigenstructure tracking for adaptive spectral estimation," *IEEE Trans. Signal Processing*, vol. 43, no. 5, pp. 1151–1160, May 1995.

[22] B. Yang, "Projection Approximation Subspace Tracking," *IEEE Trans. Signal Processing*, vol. 44, no. 1, pp. 95–107, Jan. 1995.

[23] Y. Miao and Y. Hua, "Fast subspace tracking and neural network learning by a novel information criterion," *IEEE Trans. Signal Processing*, vol. 46, no. 7, pp. 1967–1979, July 1998.

[24] G. H. Golub and C. F. V. Loan, *Matrix computations*, 3rd ed. Baltimore and London: The Johns Hopkins University Press, 1996.

[25] K. Abed-Meraim, A. Chkeif, and Y. Hua, "Fast orthonormal PAST algorithm," *IEEE Signal Proc. Letters*, vol. 7, no. 3, pp. 60–62, Mar. 2000.

[26] R. Badeau, G. Richard, and B. David, "Approximated power iterations for fast subspace tracking," in *Proc. of 7th Int. Symp. on Signal Proc. and its Applications*, vol. 2, Paris, France, July 2003, pp. 583–586.

[27] P. Strobach, "Square hankel SVD subspace tracking algorithms," *Signal Processing*, vol. 57, no. 1, pp. 1–18, Feb. 1997.

[28] E. C. Real, D. W. Tufts, and J. W. Cooley, "Two algorithms for fast approximate subspace tracking," *IEEE Trans. Signal Processing*, vol. 47, no. 7, pp. 1936–1945, July 1999.

[29] R. Badeau, K. Abed-Meraim, G. Richard, and B. David, "Sliding Window Orthonormal PAST Algorithm," in *Proc. of IEEE Int. Conf. on Acoustic, Speech and Signal Processing*, vol. 5, Apr. 2003, pp. 261–264.

[30] R. Badeau, G. Richard, and B. David, "Sliding window adaptive SVD algorithms," *IEEE Trans. Signal Processing*, vol. 52, no. 1, pp. 1–10, Jan. 2004.

[31] ——, "Suivi d'espace dominant par la méthode des puissances itérées," in *Proc. of 19ème colloque GRETSI sur le traitement du signal et des images*, vol. 1, Sept. 2003, pp. 137–140.

[32] R. O. Schmidt, "A signal subspace approach to multiple emitter location and spectral estimation," Ph.D. dissertation, Stanford University, Nov. 1981.

[33] R. A. Horn and C. R. Johnson, *Matrix analysis*. Cambridge: Cambridge University Press, 1985.

[34] S. C. Douglas, "Numerically-robust adaptive subspace tracking using Householder transformations," in *Proc. of IEEE Sensor Array and Multichannel Signal Proc. Workshop*, 2000, pp. 499 –503.

[35] Y. Hua and T. K. Sarkar, "Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 38, no. 5, pp. 814–824, May 1990.

[36] P. Strobach, "Fast recursive subspace adaptive ESPRIT algorithms," *IEEE Trans. Signal Processing*, vol. 46, no. 9, pp. 2413–2430, Sept. 1998.

[37] C. E. Davila, "Efficient, high performance, subspace tracking for time-domain data," *IEEE Trans. Signal Processing*, vol. 48, no. 12, pp. 3307–3315, Dec. 2000.

[38] R. Badeau, G. Richard, and B. David, "Adaptive ESPRIT algorithm based on the PAST subspace tracker," in *Proc. of IEEE ICASSP'03*, vol. 6, Apr. 2003, pp. 229–232.

Fast and stable YAST algorithm for principal and minor subspace tracking

Roland Badeau, *Member, IEEE,* Gaël Richard, *Senior Member, IEEE*, and Bertrand David, *Member, IEEE*

***Abstract*—This paper presents a new implementation of the YAST algorithm for principal and minor subspace tracking. YAST was initially derived from the Subspace Projection (SP) algorithm by C.E. Davila, which was known for its exceptional convergence rate, compared to other classical principal subspace trackers. The novelty in the YAST algorithm was the lower computational cost (linear if the data correlation matrix satisfies a so-called shift-invariance property), and the extension to minor subspace tracking. However, the original implementation of the YAST algorithm suffered from a numerical stability problem (the subspace weighting matrix slowly loses its orthonormality). We thus propose in this paper a new implementation of YAST, whose stability is established theoretically and tested via numerical simulations. This algorithm combines all the desired properties for a subspace tracker: remarkably high convergence rate, lowest steady state error, linear complexity, and numerical stability regarding the orthonormality of the subspace weighting matrix.**

***Index Terms*—Adaptive estimation, Time series, Algorithms, Numerical stability, Convergence of numerical methods, Complexity theory.**

I. Introduction

FAST estimation and tracking of the principal or minor subspace of a sequence of random vectors is a major problem in many applications, such as adaptive filtering, system identification, source localization, and spectral analysis for instance [1]. This problem can be stated as follows: considering a sequence of n-dimensional random data vectors $\{\boldsymbol{x}(t)\}_{t\in\mathbb{Z}}$, whose time-varying correlation matrix is denoted $\boldsymbol{C}_{xx}(t)$, we aim at estimating and tracking the r-dimensional principal or minor eigensubspace[1] of $\boldsymbol{C}_{xx}(t)$ (with $r < n$), which can be represented by its orthogonal projector $\boldsymbol{\Pi}(t)$, of rank r. The exact solution of this problem is well known: the principal (resp. minor) subspace of $\boldsymbol{C}_{xx}(t)$ is that which maximizes (resp. minimizes) the *generalized Rayleigh quotient*

$$\mathcal{J}\left(\boldsymbol{\Pi}(t)\right) = \operatorname{trace}\left(\boldsymbol{C}_{xx}(t)\,\boldsymbol{\Pi}(t)\right). \tag{1}$$

The orthogonal projector $\boldsymbol{\Pi}(t)$ is generally parameterized by an $n \times r$ orthonormal matrix $\boldsymbol{W}(t)$, such that

$$\boldsymbol{\Pi}(t) = \boldsymbol{W}(t)\,\boldsymbol{W}(t)^H, \tag{2}$$

where the symbol ^{H} denotes the conjugate transpose of a matrix (or a vector). In practice, a straightforward calculation of this so-called *subspace weighting matrix* $\boldsymbol{W}(t)$ is computationally demanding: its complexity is $O(n^2 r)$, which is inappropriate for real-time applications. Therefore the art of subspace tracking consists in recursively updating a matrix $\boldsymbol{W}(t)$, as close as possible to this exact solution, with a minimum of computations.

A large number of algorithms have been proposed for performing this task. Amidst this abundant literature[2], the PAST [3], LORAF3 [4], OPAST [5], and FAPI [6] algorithms belong to the family of Principal Subspace Trackers (PST), whereas the QRI [7] and HFRANS [8] algorithms belong to the family of Minor Subspace Trackers (MST). Besides, some algorithms such as ODKA [9], NOOJA [10], and FDPM [2], [11] can handle both principal and minor subspace tracking. The performance of these various algorithms is generally evaluated according to several criteria : their computational complexity, which is quantified as the number of operations performed at each iteration; the steady-state error, which is the mean estimation error when the algorithm has converged; and the convergence rate, which corresponds to the number of required iterations before reaching the steady state. Another desirable property for a subspace tracker is to guarantee the orthonormality of the subspace weighting matrix $\boldsymbol{W}(t)$ at each iteration. It is then worth evaluating the stability of the algorithm with respect to the orthonormality of $\boldsymbol{W}(t)$.

Recently, a new PST, referred to as the Subspace Projection (SP) algorithm was proposed by C.E. Davila [12]. We observed in [6] that the convergence rate of this algorithm is much higher than that of other classical PST's. However, this remarkable subspace tracker is computationally demanding. Indeed, the fastest implementation of SP relies on a so-called *shift invariance property* of the correlation matrix $\boldsymbol{C}_{xx}(t)$ (which typically holds in the case of time series analysis for instance), and has a complexity of $O(nr^2)$, whereas a number of existing subspace trackers only require $O(nr)$ operations per iteration (which is the case of all the above mentioned PST algorithms). Nevertheless, we found out that this drawback could be circumvented, and we proposed in [13] and [14] a new algorithm, referred to as YAST, which computes the same signal subspace as the SP algorithm, but only requires $O(nr)$ operations (when the correlation matrix $\boldsymbol{C}_{xx}(t)$ satisfies the shift invariance property). A version of YAST dedicated to MST was then proposed in [15].

Despite the great diversity of approaches proposed in the literature for performing PST and MST, most of the algorithms

Roland Badeau, Gaël Richard, and Bertrand David are with the Department of Signal and Image Processing, TELECOM ParisTech (ENST) - LTCI CNRS, Paris, France. E-mail: [roland.badeau, bertrand.david, gael.richard]@telecom-paristech.fr.

The research leading to this paper was supported by the ACI *Masse de données* Music Discover, and by the European Commission under contract FP6-027026, Knowledge Space of semantic inference for automatic annotation and retrieval of multimedia content - K-Space.

[1]The principal (resp. minor) eigensubspace of $\boldsymbol{C}_{xx}(t)$ is defined as the subspace spanned by the r eigenvectors of $\boldsymbol{C}_{xx}(t)$ associated to r greatest (resp. lowest) eigenvalues.

[2]A recent literature review can be found in [2] for instance.

can be interpreted in the same unified framework: at each iteration t, the range space of $\boldsymbol{W}(t)$ belongs to a subspace of dimension $r+1$, spanned by $\boldsymbol{W}(t-1)$, plus one additional direction, which is often chosen as the data vector $\boldsymbol{x}(t)$ (this is the case of all the above mentioned algorithms for instance, although some of them involve several additional directions at each iteration). From this perspective, the exceptional convergence rate of SP and YAST can be easily explained: at each iteration, those two algorithms compute the "best" r-dimensional subspace inside this augmented subspace, *i.e.* the subspace which explicitly maximizes (or minimizes) the generalized Rayleigh quotient. Several versions of YAST have already been presented in the literature:

- Various additional search directions have been proposed, such as $\boldsymbol{x}(t)$ and $\boldsymbol{C}_{xx}(t-1)\boldsymbol{x}(t)$ in [13], [15]. A conjugate gradient method was also presented in [16], for computing an appropriate search direction in the case of MST. This method can easily be adapted to PST.
- Several window shapes were proposed for estimating the correlation matrix $\boldsymbol{C}_{xx}(t)$: an exponential window in [13], [15], and a truncated (or sliding) window in [14].

Unfortunately, it was observed in [17] and [18] that the implementations of the YAST PST presented in [13] and [15] suffer from a numerical stability problem (the subspace weighting matrix slowly loses its orthonormality). In order to solve this problem, we propose in this paper to apply a new orthogonalization procedure to the YAST algorithm, which is similar to that proposed in [2] for the FDPM subspace tracker. The stability of the resulting algorithm is established theoretically and tested via numerical simulations. To keep the mathematical developments as concise as possible, the new implementation of YAST introduced in this paper thus involves the following changes with respect to earlier versions:

- this implementation stands for both principal and minor subspace tracking;
- there is only one additional search direction, which is the data vector $\boldsymbol{x}(t)$;
- the update of the correlation matrix is based on the simplest window shape (exponential window);
- this implementation involves an orthogonalization procedure similar to that proposed in [2].

Nevertheless, this implementation can be easily generalized to truncated windows and other search directions (*cf.* [13]–[16]). We also present a convergence analysis of the YAST algorithm, and we prove the numerical stability of the new implementation. The paper is organized as follows: in section II, the basic principle of the YAST algorithm is summarized. Then a fast implementation of YAST is proposed in section III. The performance of this subspace tracker is illustrated in section IV. The main conclusions of this paper are summarized in section V. Finally, the convergence and the numerical stability of the YAST algorithm are investigated in the Appendix.

II. Overview of the YAST algorithm

A. Principle

Consider a sequence of independent n-dimensional random data vectors $\{\boldsymbol{x}(t)\}_{t\in\mathbb{Z}}$, whose $n\times n$ correlation matrix $\boldsymbol{C}_{xx}(t)$ is estimated by applying an exponential window to the data, leading to the update

$$\boldsymbol{C}_{xx}(t) = \beta\,\boldsymbol{C}_{xx}(t-1) + \boldsymbol{x}(t)\,\boldsymbol{x}(t)^H \tag{3}$$

where $0<\beta<1$ is the exponential forgetting factor. As mentioned in the introduction, the generalized Rayleigh quotient is maximized (resp. minimized) when the subspace weighting matrix $\boldsymbol{W}(t)$ spans the principal subspace (resp. the minor subspace) of $\boldsymbol{C}_{xx}(t)$. Unfortunately, implementing this optimization over all orthonormal matrices is computationally demanding, and does not lead to a simple recursion between $\boldsymbol{W}(t)$ and $\boldsymbol{W}(t-1)$. In order to reduce the complexity, the implementation of YAST proposed below recursively performs this search within the $(r+1)$-dimensional subspace spanned by $\boldsymbol{W}(t-1)$ and $\boldsymbol{x}(t)$.

Let $\underline{\boldsymbol{\Pi}}(t)$ be the orthogonal projector on the augmented subspace, and $\underline{\boldsymbol{W}}(t)$ an $n\times(r+1)$ orthonormal matrix such that

$$\underline{\boldsymbol{\Pi}}(t) = \underline{\boldsymbol{W}}(t)\underline{\boldsymbol{W}}(t)^H. \tag{4}$$

Then any r-dimensional subspace of $\mathrm{span}(\underline{\boldsymbol{\Pi}}(t))$ can be represented by the orthogonal projector

$$\boldsymbol{\Pi}(t) = \underline{\boldsymbol{\Pi}}(t) - \underline{\boldsymbol{v}}(t)\,\underline{\boldsymbol{v}}(t)^H, \tag{5}$$

where the unitary vector $\underline{\boldsymbol{v}}(t)$ belongs to the range space of $\underline{\boldsymbol{\Pi}}(t)$. In particular, $\underline{\boldsymbol{v}}(t)$ can be written in the form

$$\underline{\boldsymbol{v}}(t) = \underline{\boldsymbol{W}}(t)\,\underline{\boldsymbol{\phi}}(t), \tag{6}$$

where $\underline{\boldsymbol{\phi}}(t)$ is an $(r+1)$-dimensional unitary vector. Substituting equations (4) to (6) into equation (1), the criterion to be optimized can be rewritten as

$$\mathcal{J}\left(\boldsymbol{\Pi}(t)\right) = \mathrm{trace}\left(\underline{\boldsymbol{C}}_{yy}(t)\right) - \underline{\boldsymbol{\phi}}(t)^H\underline{\boldsymbol{C}}_{yy}(t)\,\underline{\boldsymbol{\phi}}(t) \tag{7}$$

where $\underline{\boldsymbol{C}}_{yy}(t)$ is the $(r+1)\times(r+1)$ matrix

$$\underline{\boldsymbol{C}}_{yy}(t) = \underline{\boldsymbol{W}}(t)^H\,\boldsymbol{C}_{xx}(t)\,\underline{\boldsymbol{W}}(t). \tag{8}$$

In this manner, the optimization of $\mathcal{J}$ with respects to the $n\times r$ orthonormal matrix $\boldsymbol{W}(t)$ is replaced by the optimization of $\mathcal{J}$ with respects to the $(r+1)$-dimensional unitary vector $\underline{\boldsymbol{\phi}}(t)$. The result of this optimization is well-known: according to equation (7), $\mathcal{J}$ is maximized (resp. minimized) when $\underline{\boldsymbol{\phi}}(t)$ is the minor (resp. principal) eigenvector of $\underline{\boldsymbol{C}}_{yy}(t)$. Finally, given the new data vector $\boldsymbol{x}(t)$, the YAST algorithm updates the previous subspace weighting matrix $\boldsymbol{W}(t-1)$ by successively computing

1) an orthonormal basis $\underline{\boldsymbol{W}}(t)$ of the augmented subspace;
2) the compressed matrix $\underline{\boldsymbol{C}}_{yy}(t)$ defined in equation (8);
3) the principal or minor eigenvector $\underline{\phi}(t)$ of $\underline{\boldsymbol{C}}_{yy}(t)$ (for MST or PST resp.), and a basis $\boldsymbol{W}(t)$ of the range space of the projector $\boldsymbol{\Pi}(t)$ defined in equation (5);
4) the $r\times r$ compressed matrix $\boldsymbol{C}_{yy}(t)$, defined as

$$\boldsymbol{C}_{yy}(t) = \boldsymbol{W}(t)^H\boldsymbol{C}_{xx}(t)\boldsymbol{W}(t), \tag{9}$$

which will permit a fast calculation of $\underline{\boldsymbol{C}}_{yy}(t)$ in step 2).

The convergence of this YAST algorithm is proved in Appendix B page 8, and a fast implementation is proposed in section III.

B. Numerical stability

As mentioned in the introduction, it was shown in [17] that the implementation of the YAST PST presented in [13] suffers from a numerical stability problem. Indeed, although YAST was originally designed to compute an orthonormal subspace weighting matrix $\boldsymbol{W}(t)$ at each iteration, when the estimated subspace gets too close to the exact subspace (which happens in the case of stationary data with a high signal-to-noise ratio), a sudden loss of orthogonality makes the algorithm diverge. The solution proposed in [17] to cope with this problem consisted in switching to another subspace tracker such as FAPI [6] when this critical point was reached. In the context of MST [15], we addressed the problem in a different way, by modifying our implementation of YAST in order to restore the orthogonality of the subspace weighting matrix at each iteration. Although this modification prevents the sudden divergence of YAST, it was proved in [18] that the accumulation of rounding errors still leads to a loss of orthogonality of the subspace weighting matrix, at a much lower rate however. In order to alleviate this effect, it was proposed in [18] to use a partial orthogonalization scheme called Pairwise Gram-Schmidt (PGS), which improves the numerical stability of YAST without increasing its computational cost. However, we will show in section IV an example of instable behavior of YAST-PGS.

In fact, the loss of orthogonality is a shortcoming common to most of the existing MST's of linear complexity. In particular, the classical Oja algorithm [19] is known to diverge. A very classical approach for improving the numerical stability of MST consists in using Householder transforms [8]–[10]. However this technique does not completely prevent the slow accumulation of rounding errors, as observed in [2], and the convergence rate of these algorithms is much lower than that of the classical PST techniques. To the best of our knowledge, the first MST of linear complexity which guaranteed a perfect stability with respect to the orthogonality was the Fast Data Projection Method (FDPM) by G. Doukopoulos and V. Moustakides [2], [11]. This algorithm relies on a specific orthogonalization procedure, whose numerical stability was theoretically proved. Considering the exceptional stability of FDPM, we propose in the following section to apply a similar orthogonalization procedure to the YAST algorithm, in order to solve the problems observed in our previous implementations. The numerical stability of this new implementation is proved in Appendix C page 9.

III. Fast implementation of YAST

Below, a fast implementation of the YAST algorithm is proposed, whose global cost is only $5nr$ flops[3]. Following the structure proposed in section II-A, it is composed of four steps: computation of $\underline{\boldsymbol{W}}(t)$ (section III-A), computation of $\underline{\boldsymbol{C}}_{yy}(t)$ (section III-B), update of $\boldsymbol{W}(t)$ (section III-C), and update of $\boldsymbol{C}_{yy}(t)$ (section III-D). The computational complexity of this implementation is addressed in section III-E.

[3]In this paper, a flop is a multiply / accumulate (MAC) operation.

A. Computation of $\underline{\boldsymbol{W}}(t)$

Define the r-dimensional compressed data vector

$$\boldsymbol{y}(t) = \boldsymbol{W}(t-1)^H \boldsymbol{x}(t). \tag{10}$$

Then the n-dimensional vector

$$\boldsymbol{e}(t) = \boldsymbol{x}(t) - \boldsymbol{W}(t-1)\,\boldsymbol{y}(t) \tag{11}$$

is orthogonal to the range space of $\boldsymbol{W}(t-1)$. Let $\sigma(t)$ be its Hermitian norm:

$$\sigma(t) = \|\boldsymbol{e}(t)\|_2. \tag{12}$$

If $\sigma(t) = 0$, then the new search direction $\boldsymbol{x}(t)$ is actually included in the old subspace spanned by $\boldsymbol{W}(t-1)$. This means that the new subspace is equal to the old one, thus we can simply choose $\boldsymbol{W}(t) = \boldsymbol{W}(t-1)$, and the iteration is over. Otherwise $\sigma(t) \neq 0$, and we can define the normalized vector

$$\boldsymbol{u}(t) = \boldsymbol{e}(t) \,/\, \sigma(t). \tag{13}$$

Then the $n \times (r+1)$ augmented matrix

$$\underline{\boldsymbol{W}}(t) = \left[\, \boldsymbol{W}(t-1) \,\middle|\, \boldsymbol{u}(t) \,\right] \tag{14}$$

defines an orthonormal basis of $\mathrm{span}([\boldsymbol{W}(t-1), \boldsymbol{x}(t)])$.

B. Computation of $\underline{\boldsymbol{C}}_{yy}(t)$

Substituting equation (3) and equations (9) to (14) into equation (8) yields

$$\underline{\boldsymbol{C}}_{yy}(t) = \left[\begin{array}{c|c} \boldsymbol{C}'_{yy}(t) & \boldsymbol{z}(t) \\ \hline \boldsymbol{z}(t)^H & \gamma(t) \end{array}\right], \tag{15}$$

where

$$\boldsymbol{C}'_{yy}(t) = \beta\,\boldsymbol{C}_{yy}(t-1) + \boldsymbol{y}(t)\,\boldsymbol{y}(t)^H \tag{16}$$

$$\boldsymbol{z}(t) = \beta \boldsymbol{W}(t-1)^H \boldsymbol{C}_{xx}(t-1)\boldsymbol{u}(t) + \sigma(t)\boldsymbol{y}(t) \tag{17}$$

$$\gamma(t) = \beta \boldsymbol{u}(t)^H \boldsymbol{C}_{xx}(t-1)\boldsymbol{u}(t) + \sigma(t)^2. \tag{18}$$

In order to reduce the computational cost, it will be useful to remark that $\boldsymbol{z}(t)$ and $\gamma(t)$ can be efficiently computed as

$$\boldsymbol{z}(t) = \beta \frac{\boldsymbol{y}''(t) - \boldsymbol{y}'(t)}{\sigma(t)} + \sigma(t)\boldsymbol{y}(t), \tag{19}$$

$$\gamma(t) = \beta \frac{\boldsymbol{x}(t)^H \boldsymbol{x}'(t) - 2\Re(\boldsymbol{y}(t)^H \boldsymbol{y}''(t)) + \boldsymbol{y}(t)^H \boldsymbol{y}'(t)}{\sigma(t)^2} + \sigma(t)^2 \tag{20}$$

where we have introduced the auxiliary vectors

$$\boldsymbol{x}'(t) = \boldsymbol{C}_{xx}(t-1)\,\boldsymbol{x}(t), \tag{21}$$

$$\boldsymbol{y}'(t) = \boldsymbol{C}_{yy}(t-1)\,\boldsymbol{y}(t), \tag{22}$$

$$\boldsymbol{y}''(t) = \boldsymbol{W}(t-1)^H \boldsymbol{x}'(t). \tag{23}$$

C. Update of $\boldsymbol{W}(t)$

Let $\underline{\boldsymbol{\phi}}(t)$ be the unitary eigenvector associated to the greatest (or lowest) eigenvalue of the matrix $\underline{\boldsymbol{C}}_{yy}(t)$ of dimension $(r+1)\times(r+1)$. The YAST algorithm consists in estimating the new subspace as the orthogonal complement of the vector $\underline{\boldsymbol{v}}(t) = \underline{\boldsymbol{W}}(t)\,\underline{\boldsymbol{\phi}}(t)$ (defined in equation (6)) in the range space of $\underline{\boldsymbol{W}}(t)$. To this end, we will first introduce a non-orthonormal matrix $\boldsymbol{T}(t)$ which spans this new subspace, then we will describe how this matrix can be orthonormalized.

1) Computation of the non-orthonormal subspace weighting matrix $\boldsymbol{T}(t)$: Vector $\underline{\boldsymbol{\phi}}(t)$ is now decomposed in the form

$$\underline{\boldsymbol{\phi}}(t) = \left[\begin{array}{c} \varepsilon(t)\boldsymbol{\phi}(t) \\ \hdashline \varphi(t) \end{array}\right] \theta(t), \tag{24}$$

where $\varphi(t)\theta(t)$ is the polar form of the last coefficient of $\underline{\boldsymbol{\phi}}(t)$ (*i.e.* $\varphi(t) \in \mathbb{R}$ is non-negative, and $\theta(t)$ is a complex number such that $|\theta(t)| = 1$), $\varepsilon(t) \geq 0$, and $\boldsymbol{\phi}(t)$ is an r-dimensional unitary vector. In particular, $\varepsilon(t)$ and $\varphi(t)$ are such that $\varepsilon(t)^2 + \varphi(t)^2 = 1$. The following lemma which involves $\varepsilon(t)$ is proved in Appendix A page 8.

Lemma III-C.1: $\varepsilon(t)$ is bounded by the distance between the estimated subspaces at times t and $t-1$, in the sense that

$$\varepsilon(t)^2 \leq \|\boldsymbol{\Pi}(t) - \boldsymbol{\Pi}(t-1)\|_2$$

where $\|.\|_2$ denotes the matrix spectral norm, or 2-norm.

Lemma III-C.1 shows that when the algorithm has reached its steady state, *i.e.* when the subspace admits slow variations ($\|\boldsymbol{\Pi}(t) - \boldsymbol{\Pi}(t-1)\|_2 \ll 1$), then $\varepsilon(t)$ is much smaller than 1.

Next, the following lemma introduces the non-orthonormal subspace weighting matrix $\boldsymbol{T}(t)$.

Lemma III-C.2: The range space of the $n \times r$ matrix

$$\boldsymbol{T}(t) = \boldsymbol{W}(t-1) - \varepsilon(t)\boldsymbol{u}(t)\boldsymbol{\phi}(t)^H \tag{25}$$

is a subspace of $\mathrm{span}(\underline{\boldsymbol{W}}(t))$, such that

$$\left\|\underline{\boldsymbol{v}}(t)^H \boldsymbol{T}(t)\right\|_2 \leq (\varepsilon(t))^3 . \tag{26}$$

Lemma III-C.2 shows that when $\varepsilon(t) \to 0$, the range space of $\boldsymbol{T}(t)$ asymptotically matches the orthogonal complement of $\underline{\boldsymbol{v}}(t)$ in the augmented subspace $\mathrm{span}(\underline{\boldsymbol{W}}(t))$. Lemma III-C.2 is also proved in Appendix A page 8.

2) Computation of the orthonormal subspace weighting matrix $\boldsymbol{W}(t)$: The following developments aim at orthonormalizing $\boldsymbol{T}(t)$ to obtain the new subspace weighting matrix $\boldsymbol{W}(t)$ [4]. In order to ensure the numerical stability of our algorithm, the proposed orthonormalization method is similar to that presented in [2]. First, equation (25) yields

$$\boldsymbol{T}(t)^H \boldsymbol{T}(t) = \boldsymbol{I}_r + \varepsilon(t)^2\, \boldsymbol{\phi}(t)\, \boldsymbol{\phi}(t)^H . \tag{27}$$

(here $\boldsymbol{I}_r$ denotes the $r \times r$ identity matrix). Define the r-dimensional unitary vector

$$\boldsymbol{e}_1(t) = -e^{i\,\mathrm{angle}(\phi_1(t))}[1, 0 \ldots 0]^T, \tag{28}$$

where the symbol ^{T} denotes the transpose of a vector (or a matrix), and $\phi_1(t)$ is the first coefficient of $\boldsymbol{\phi}(t)$. Then define the r-dimensional unitary vector

$$\boldsymbol{a}(t) = \frac{\boldsymbol{\phi}(t) - \boldsymbol{e}_1(t)}{\|\boldsymbol{\phi}(t) - \boldsymbol{e}_1(t)\|_2}, \tag{29}$$

and consider the Householder matrix

$$\boldsymbol{S}(t) = \boldsymbol{I}_r - 2\boldsymbol{a}(t)\boldsymbol{a}(t)^H, \tag{30}$$

which transforms the vector $\boldsymbol{\phi}(t)$ into $\boldsymbol{e}_1(t)$. Note that the denominator in the right member of equation (29) can never be zero (more precisely, it belongs to the interval $[1, 2]$, which prevents a division by zero). Let

$$\boldsymbol{Q}(t) = \boldsymbol{T}(t)\, \boldsymbol{S}(t). \tag{31}$$

Then equations (27) and (31) yield

$$\boldsymbol{Q}(t)^H \boldsymbol{Q}(t) = \boldsymbol{I}_r + \varepsilon(t)^2 \boldsymbol{e}_1(t)\, \boldsymbol{e}_1(t)^H,$$

which shows that the columns of the matrix $\boldsymbol{Q}(t)$ are orthogonal. Therefore an orthonormal matrix $\boldsymbol{W}(t)$ can be obtained by normalizing the first column of $\boldsymbol{Q}(t)$:

$$\boldsymbol{W}(t) = \boldsymbol{Q}(t)\, \boldsymbol{D}(t) \tag{32}$$

where $\boldsymbol{D}(t)$ is the $r \times r$ diagonal matrix

$$\boldsymbol{D}(t) = \mathrm{diag}(1/\|\boldsymbol{q}_1(t)\|_2, 1 \ldots 1), \tag{33}$$

$\boldsymbol{q}_1(t)$ being the first column of $\boldsymbol{Q}(t)$. Finally, substituting equations (25) and (30) into equation (31) yields

$$\boldsymbol{Q}(t) = \boldsymbol{W}(t-1) - 2\boldsymbol{b}(t)\boldsymbol{a}(t)^H - \varepsilon(t)\boldsymbol{u}(t)\, \boldsymbol{e}_1(t)^H \tag{34}$$

where

$$\boldsymbol{b}(t) = \boldsymbol{W}(t-1)\, \boldsymbol{a}(t). \tag{35}$$

D. Update of $\boldsymbol{C}_{yy}(t)$

The auxiliary matrix $\boldsymbol{C}_{yy}(t) = \boldsymbol{W}(t)^H \boldsymbol{C}_{xx}(t) \boldsymbol{W}(t)$ defined in equation (9) can also be efficiently updated. Indeed, substituting equations (31), (30), (25), and (14) into equation (32) yields

$$\boldsymbol{W}(t) = \underline{\boldsymbol{W}}(t)\, \boldsymbol{U}(t)\, \boldsymbol{D}(t) \tag{36}$$

where the $(r+1) \times r$ matrix $\boldsymbol{U}(t)$ is defined as

$$\boldsymbol{U}(t) = \left[\begin{array}{c} \boldsymbol{I}_r - 2\boldsymbol{a}(t)\boldsymbol{a}(t)^H \\ \hdashline -\varepsilon(t)\boldsymbol{e}_1(t)^H \end{array}\right]. \tag{37}$$

Then substituting equations (8) and (36) into equation (9) yields

$$\boldsymbol{C}_{yy}(t) = \boldsymbol{D}(t)\, \boldsymbol{C}''_{yy}(t)\, \boldsymbol{D}(t). \tag{38}$$

where

$$\boldsymbol{C}''_{yy}(t) = \boldsymbol{U}(t)^H \underline{\boldsymbol{C}}_{yy}(t) \boldsymbol{U}(t). \tag{39}$$

Finally, substituting equations (15) and (37) into equation (39) yields

$$\boldsymbol{C}''_{yy}(t) = \mathrm{Herm}\left(\boldsymbol{C}'_{yy}(t) - \boldsymbol{a}'(t)\boldsymbol{a}(t)^H - \varepsilon(t)\boldsymbol{z}'(t)\boldsymbol{e}_1(t)^H\right) \tag{40}$$

where $\mathrm{Herm}(.)$ denotes the Hermitian part of a matrix ($\mathrm{Herm}(\boldsymbol{M}) = \frac{\boldsymbol{M}+\boldsymbol{M}^H}{2}$), and

$$\boldsymbol{a}'(t) = 4\boldsymbol{C}'_{yy}(t)\boldsymbol{a}(t) - 4(\boldsymbol{a}(t)^H \boldsymbol{C}'_{yy}(t)\boldsymbol{a}(t))\, \boldsymbol{a}(t), \tag{41}$$

$$\boldsymbol{z}'(t) = 2\boldsymbol{z}(t) - 4(\boldsymbol{a}(t)^H \boldsymbol{z}(t))\boldsymbol{a}(t) - \varepsilon(t)\gamma(t)\boldsymbol{e}_1(t). \tag{42}$$

[4] Actually, an exact implementation of YAST would require that $\boldsymbol{T}(t)$ be exactly orthogonal to $\underline{\boldsymbol{v}}(t)$. However we observed in practice that this (negligible) approximation improves both the stability and the convergence rate of YAST, as illustrated in section IV-B2.

TABLE I
PSEUDO-CODE OF THE YAST ALGORITHM

	eq.:	flops:
Input: $\boldsymbol{x}(t)$		
$\boldsymbol{y}(t) = \boldsymbol{W}(t-1)^H \boldsymbol{x}(t)$	(10)	nr
$\boldsymbol{e}(t) = \boldsymbol{x}(t) - \boldsymbol{W}(t-1)\boldsymbol{y}(t)$	(11)	nr
$\sigma(t) = \|\boldsymbol{e}(t)\|_2$	(12)	n
if $\sigma(t) \neq 0$,		
$\boldsymbol{u}(t) = \boldsymbol{e}(t) \,/\, \sigma(t)$	(13)	
$\boldsymbol{x}'(t) = \boldsymbol{C}_{xx}(t-1)\,\boldsymbol{x}(t)$	(21)	$9n$
$\boldsymbol{y}'(t) = \boldsymbol{C}_{yy}(t-1)\,\boldsymbol{y}(t)$	(22)	r^2
$\boldsymbol{y}''(t) = \boldsymbol{W}(t-1)^H \boldsymbol{x}'(t)$	(23)	nr
$\boldsymbol{C}'_{yy}(t) = \beta\,\boldsymbol{C}_{yy}(t-1) + \boldsymbol{y}(t)\,\boldsymbol{y}(t)^H$	(16)	r^2
$\boldsymbol{z}(t) = \beta\,(\boldsymbol{y}''(t) - \boldsymbol{y}'(t))/\sigma(t) + \sigma(t)\boldsymbol{y}(t)$	(19)	r
$\gamma(t) = \beta\frac{\boldsymbol{x}(t)^H\boldsymbol{x}'(t) - 2\Re(\boldsymbol{y}(t)^H\boldsymbol{y}''(t)) + \boldsymbol{y}(t)^H\boldsymbol{y}'(t)}{\sigma(t)^2} + \sigma(t)^2$	(20)	n
$\underline{\boldsymbol{C}}_{yy}(t) = \left[\begin{array}{c\|c} \boldsymbol{C}'_{yy}(t) & \boldsymbol{z}(t) \\ \hline \boldsymbol{z}(t)^H & \gamma(t) \end{array}\right]$	(15)	
if Principal Subspace Tracking		
$\left(\underline{\boldsymbol{\phi}}(t), \lambda(t)\right) = \min\left\{\text{eig}\left(\underline{\boldsymbol{C}}_{yy}(t)\right)\right\}$		$O(r^2)$
else if Minor Subspace Tracking		
$\left(\underline{\boldsymbol{\phi}}(t), \lambda(t)\right) = \max\left\{\text{eig}\left(\underline{\boldsymbol{C}}_{yy}(t)\right)\right\}$		$O(r^2)$
end if		
$\left[\varepsilon(t)\boldsymbol{\phi}(t)^T, \varphi(t)\right]\,\theta(t) = \underline{\boldsymbol{\phi}}(t)^T$	(24)	$2r$
$\boldsymbol{e}_1(t) = -e^{i\,\text{angle}(\phi_1(t))}[1, 0 \ldots 0]^T$	(28)	
$\boldsymbol{a}(t) = \frac{\boldsymbol{\phi}(t) - \boldsymbol{e}_1(t)}{\|\boldsymbol{\phi}(t) - \boldsymbol{e}_1(t)\|_2}$	(29)	r
$\boldsymbol{b}(t) = \boldsymbol{W}(t-1)\,\boldsymbol{a}(t)$	(35)	nr
$\boldsymbol{Q}(t) = \boldsymbol{W}(t-1) - 2\boldsymbol{b}(t)\boldsymbol{a}(t)^H - \varepsilon(t)\boldsymbol{u}(t)\,\boldsymbol{e}_1(t)^H$	(34)	nr
$\boldsymbol{D}(t) = \text{diag}(1/\|\boldsymbol{q}_1(t)\|_2, 1 \ldots 1)$	(33)	n
$\boldsymbol{W}(t) = \boldsymbol{Q}(t)\,\boldsymbol{D}(t)$	(32)	
$\boldsymbol{a}'(t) = 4\,\boldsymbol{C}'_{yy}(t)\boldsymbol{a}(t) - 4\,(\boldsymbol{a}(t)^H\boldsymbol{C}'_{yy}(t)\boldsymbol{a}(t))\,\boldsymbol{a}(t)$	(41)	r^2
$\boldsymbol{z}'(t) = 2\,\boldsymbol{z}(t) - 4(\boldsymbol{a}(t)^H\boldsymbol{z}(t))\boldsymbol{a}(t) - \varepsilon(t)\gamma(t)\boldsymbol{e}_1(t)$	(42)	$2r$
$\boldsymbol{C}''_{yy}(t) = \text{Herm}(\boldsymbol{C}'_{yy}(t) - \boldsymbol{a}'(t)\boldsymbol{a}(t)^H - \varepsilon(t)\boldsymbol{z}'(t)\boldsymbol{e}_1(t)^H)$	(39)	$2r^2$
$\boldsymbol{C}_{yy}(t) = \boldsymbol{D}(t)\,\boldsymbol{C}''_{yy}(t)\,\boldsymbol{D}(t)$	(40)	
end if		
Total:		$5nr + O(r^2 + n)$

E. Computational complexity

The complete pseudo-code of YAST is summarized in table I[5]. Note that this implementation requires computing the principal or minor[6] eigenvector $\underline{\boldsymbol{\phi}}(t)$ of the compressed matrix $\underline{\boldsymbol{C}}_{yy}(t)$. This operation can be performed by means of a recursive algorithm, such as the *shifted* power iteration method [20], whose dominant cost is r^2 flops per iteration. This leads to an overall complexity of $O(r^2)$ if we assume that the number of performed iterations is much lower than r, or $O(r^3)$ if we consider that this number is proportional to r. In order to reduce the number of performed iterations, since $\varepsilon(t)$ is small when the algorithm has reached its steady state (as mentioned in section III-C), we propose to choose the vector $[0 \ldots 0, 1]^T$ as an initial guess for starting this recursion. Alternately, we introduced in [16] a conjugate gradient (CG) algorithm for computing the minor eigenvector of a positive semidefinite matrix, which presents a much faster convergence rate than the power iteration method, for the same computational cost. The pseudo-code of this CG algorithm applied to the computation of the minor or principal eigenvector of $\underline{\boldsymbol{C}}_{yy}(t)$ is presented in table II, with a proper initialization. In practice, we observed that the performance of the YAST algorithm was unaltered if one iteration only of the CG algorithm is performed[7], which permits to compute $\underline{\boldsymbol{\phi}}(t)$ with an overall complexity of $O(r^2)$ flops only.

In other respects, the calculation cost of the vector $\boldsymbol{x}'(t) = \boldsymbol{C}_{xx}(t-1)\,\boldsymbol{x}(t)$ is normally n^2 flops. However, if the data correlation matrix satisfies a shift invariance property, this cost can be reduced to $9n$ flops by means of the technique described in [12]. Therefore the global cost of the YAST algorithm is reduced to $5nr$ flops. This cost can then be compared to that of other existing subspace trackers. In the PST category for instance, the complexity of PAST [3] and FAPI [6] is $3nr$, that of OPAST [5] is $4nr$, that of LORAF3 [21] is $18nr$, and that of SP1 [12] is $4nr^2$. In the MST category, the complexity of HFRANS [8] and NOOJA [10] is $4nr$, that of FDPM [2], [11] and YAST-PGS [18] is $5nr$, that of ODKA [9] is $7nr$, and that of QRI [7] is n^2r.

TABLE II
PSEUDO-CODE OF THE CONJUGATE GRADIENT ALGORITHM

	flops:
Initialization of the CG algorithm at time t	
$\boldsymbol{g} = \boldsymbol{z}(t) \,/\, \|\boldsymbol{z}(t)\|$	r
$\boldsymbol{p} = \boldsymbol{C}'_{yy}(t)\,\boldsymbol{g} - (\boldsymbol{g}^H\boldsymbol{C}'_{yy}(t)\boldsymbol{g})\,\boldsymbol{g}$	r^2
$\boldsymbol{S}(0) = \left[\begin{array}{ccc} \boldsymbol{0}_r & \boldsymbol{p}\,/\,\|\boldsymbol{p}\| & \boldsymbol{g} \\ 1 & 0 & 0 \end{array}\right]$	r
$\underline{\boldsymbol{C}}_{ys}(0) = \underline{\boldsymbol{C}}_{yy}(t)\boldsymbol{S}(0)$	r^2
$\boldsymbol{C}_{ss}(0) = \boldsymbol{S}(0)^H\underline{\boldsymbol{C}}_{ys}(0)$	$6r$
Repeat at each iteration $k \geq 1$	
if minor eigenvector computation	
$(\boldsymbol{\theta}(k), \lambda(k)) = \min\{\text{eig}\,(\boldsymbol{C}_{ss}(k-1))\}$	$O(1)$
else if principal eigenvector computation	
$(\boldsymbol{\theta}(k), \lambda(k)) = \max\{\text{eig}\,(\boldsymbol{C}_{ss}(k-1))\}$	$O(1)$
end if	
$\boldsymbol{\vartheta}(k) = \left[\begin{array}{c} -\sqrt{\|\theta_2(k)\|^2 + \|\theta_3(k)\|^2} \\ \theta_1^*(k)\frac{\theta_2(k)}{\|\theta_2(k)\|}\Big/\sqrt{1 + \left\|\frac{\theta_3(k)}{\theta_2(k)}\right\|^2} \\ \theta_1^*(k)\frac{\theta_3(k)}{\|\theta_3(k)\|}\Big/\sqrt{1 + \left\|\frac{\theta_2(k)}{\theta_3(k)}\right\|^2} \end{array}\right]$	$O(1)$
$\boldsymbol{\Theta}(k) = [\boldsymbol{\theta}(k), \boldsymbol{\vartheta}(k)]$	
$\boldsymbol{S}(k)_{(:,1:2)} = \boldsymbol{S}(k-1)\,\boldsymbol{\Theta}(k)$	$6r$
$\underline{\boldsymbol{C}}_{ys}(k)_{(:,1:2)} = \underline{\boldsymbol{C}}_{ys}(k-1)\boldsymbol{\Theta}(k)$	$6r$
$\boldsymbol{C}_{ss}(k)_{(1:2,1:2)} = \boldsymbol{\Theta}(k)^H\boldsymbol{C}_{ss}(k-1)\boldsymbol{\Theta}(k)$	12
$\frac{1}{2}\boldsymbol{\nabla}J(k) = \underline{\boldsymbol{C}}_{ys}(k)_{(:,1)} - \lambda(k)\boldsymbol{S}(k)_{(:,1)}$	r
$\boldsymbol{g}(k) = \boldsymbol{\nabla}J(k) \,/\, \|\boldsymbol{\nabla}J(k)\|$	r
$\boldsymbol{g}(k) \leftarrow \boldsymbol{g}(k) - \boldsymbol{S}(k)_{(:,1:2)}\left(\boldsymbol{S}(k)^H_{(:,1:2)}\boldsymbol{g}(k)\right)$	$4r$
$\boldsymbol{g}(k) \leftarrow \boldsymbol{g}(k) \,/\, \|\boldsymbol{g}(k)\|$	r
$\boldsymbol{S}(k)_{(:,3)} = \boldsymbol{g}(k)$	
$\underline{\boldsymbol{C}}_{ys}(k)_{(:,3)} = \underline{\boldsymbol{C}}_{yy}(t)\,\boldsymbol{S}(k)_{(:,3)}$	r^2
$\boldsymbol{C}_{ss}(k)_{(:,3)} = \boldsymbol{S}(k)^H\underline{\boldsymbol{C}}_{ys}(k)_{(:,3)}$	$3r$
$\boldsymbol{C}_{ss}(k)_{(3,:)} = \boldsymbol{C}_{ss}(k)^H_{(:,3)}$	
Until $\|\boldsymbol{\nabla}J(k)\| <$ **threshold**	
$\lambda(t) \leftarrow \lambda(k)$	
$\underline{\boldsymbol{\phi}}(t) \leftarrow \boldsymbol{S}(k)_{(:,1)}$	

IV. SIMULATION RESULTS

In this section, the performance of YAST is illustrated and compared to that of other existing algorithms: the FAPI, LORAF3, and SP1 PST's, and the FDPM, HFRANS, and YAST-PGS MST's. Other performance comparisons, involving the same test signals, can be found in references [13], [15], [16],

[5] In table I, the computational costs of multiplications and divisions of complex numbers by positive numbers have been neglected.

[6] In the case of PST, this pseudo-code involves the computation of the minor eigenvector of $\underline{\boldsymbol{C}}_{yy}(t)$. Since this operation is generally more complex than computing the principal eigenvector, an other possible implementation of the YAST PST algorithm consists in recursively updating the inverse of $\underline{\boldsymbol{C}}_{yy}(t)$, and computing $\underline{\boldsymbol{\phi}}(t)$ as the principal eigenvector of this matrix, as initially proposed in [13].

[7] In particular, the simulation results presented in section IV are unchanged.

where the OPAST, QRI, NOOJA, and ODKA algorithms were compared to YAST. In the following simulations, the subspace trackers are initialized with the subspace weighting matrix formed by the r first columns of the $n \times n$ identity matrix.

A. Principal subspace tracking

In this section, the performance of the subspace estimation is analyzed in terms of the maximum principal angle between the true principal subspace of the correlation matrix $\boldsymbol{C}_{xx}(t)$ (obtained via an exact eigenvalue decomposition), and the estimated principal subspace of the same correlation matrix (obtained with the subspace tracker). This error criterion was initially proposed by P. Comon and G.H. Golub as a measure of the distance between equidimensional subspaces [20], and used as a PST performance factor in [6], [13], [21], [22]. The test signal is a sum of $r = 4$ complex sinusoidal sources plus a complex white Gaussian noise (the SNR is 5.7 dB). The frequencies of the sinusoids vary according to a jump scenario originally proposed by P. Strobach [21]: their values abruptly change at different time instants, between which they remain constant. Their variations are represented in Figure 1-a.

Figure 1-b shows the maximum principal angle error trajectory (averaged over $K = 50$ independent runs) obtained with several PST algorithms: FAPI [6], LORAF3 [21], SP1 [12], and YAST. Those algorithms were applied with the same vector length $n = 80$ and the same forgetting factor $\beta \approx 0.99$. It can be noticed that FAPI and LORAF3 behave similarly[8]. Besides, the results obtained with YAST and SP1 cannot be distinguished, which is not surprising, since those two algorithms compute the same subspace estimate (the computational complexity of YAST, however, is lower than that of SP1). It can be noticed that those two algorithms converge much faster than the other ones, and their principal angle error is always much lower. To the best of our knowledge, YAST and SP1 are the only PST algorithms which present such an exceptional convergence rate.

In order to measure the orthonormality of the estimated subspace weighting matrix, we calculate the same average performance factor as in [7]–[10]:

$$\eta(t) = \frac{1}{K} \sum_{k=1}^{K} \left\| \boldsymbol{W}_k(t)^H \boldsymbol{W}_k(t) - \boldsymbol{I}_r \right\|_F^2 \qquad (43)$$

where the number of algorithm runs is $K = 50$, k indicates that the associated variable depends on the particular run, and $\|.\|_F$ denotes the Frobenius norm. This orthogonality error is plotted in Figure 1-c. Computations were performed in double precision with Matlab® (64 bits precision). Aside the somewhat erratic behavior of SP1, the four algorithms prove to be stable. Besides, it can be noticed that YAST provides an improved orthonormality compared to SP1. The lowest orthonormality error is reached by FAPI.

[8]Contrary to what we erroneously suggested in the introduction of reference [6], the performance of LORAF3 in terms of subspace estimation also proves to be comparable to that of other classical implementations of the power iteration method, such as PAST [3], LORAF2 [4], OPAST [5], and NIC [23].

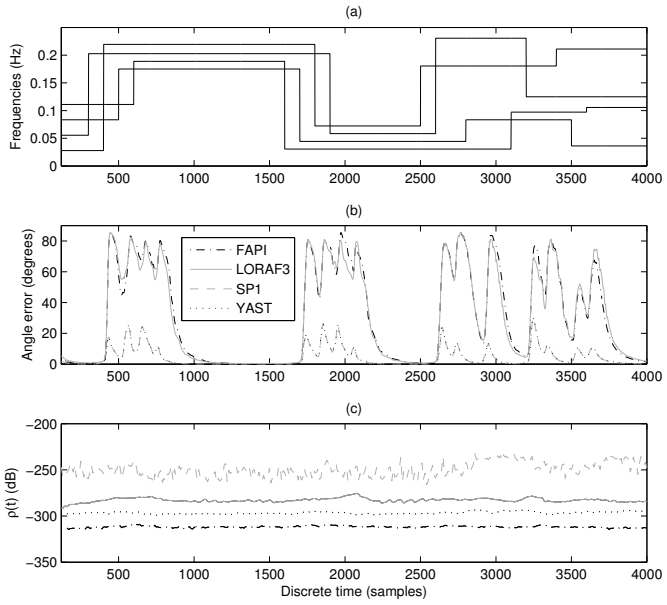

Fig. 1. Simulation results
(a) Normalized frequencies of the sinusoids,
(b) Angle error trajectories,
(c) Departure from orthonormality.

B. Minor subspace tracking

1) A classical example: In the following, $\boldsymbol{x}(t)$ is a sequence of $n = 4$ dimensional independent jointly-Gaussian random vectors, with zero mean and covariance matrix

$$\boldsymbol{C} = \begin{bmatrix} 0.9 & 0.4 & 0.7 & 0.3 \\ 0.4 & 0.3 & 0.5 & 0.4 \\ 0.7 & 0.5 & 1.0 & 0.6 \\ 0.3 & 0.4 & 0.6 & 0.9 \end{bmatrix}.$$

This is a classical example, which served as a reference for testing several MST algorithms [7]–[10], with parameter $r = 2$. In order to measure the performance of the algorithms in terms of subspace estimation, we calculate the ensemble averages of the following performance factor, as in [7]–[10]:

$$\rho(t) = \frac{1}{K} \sum_{k=1}^{K} \frac{\operatorname{trace}\left(\boldsymbol{W}_k(t)^H \boldsymbol{E}_1 \, \boldsymbol{E}_1^H \boldsymbol{W}_k(t) \right)}{\operatorname{trace}\left(\boldsymbol{W}_k(t)^H \boldsymbol{E}_2 \, \boldsymbol{E}_2^H \boldsymbol{W}_k(t) \right)},$$

where the number of algorithm runs is $K = 50$, $\boldsymbol{E}_1$ is the exact $(n - r)$-dimensional principal subspace, and $\boldsymbol{E}_2$ is the exact r-dimensional minor subspace of $\boldsymbol{C}$.

Figure 2-a shows the tracking results obtained with four MST's: our new implementation of YAST, the Pairwise Gram-Schmidt (PGS) implementation of YAST called YAST-PGS [18], and the HFRANS [8] and FDPM [2], [11] algorithms, which were derived from the Data Projection Method (DPM) of Yang and Kaveh [24] [9]. It can be noticed that the results obtained with HFRANS and FDPM cannot be

[9]The HFRANS and FDPM algorithms were implemented with parameter $\mu = 0.13$, and YAST with $\beta = 0.99$.

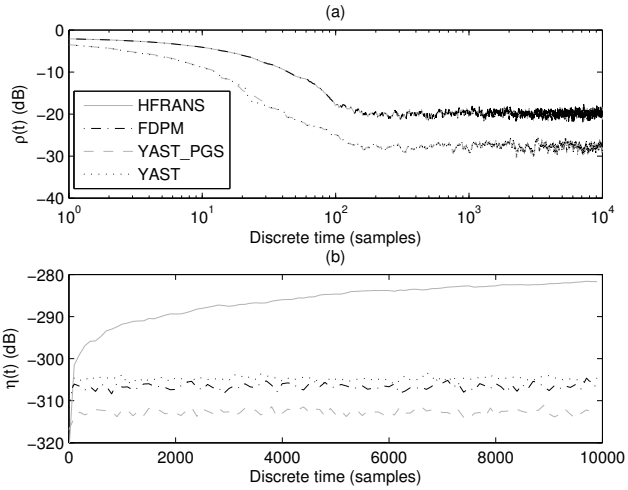

Fig. 2. A classical example
(a) Error trajectories,
(b) Departure from orthonormality.

distinguished, and that those obtained with YAST and YAST-PGS are almost identical. Moreover, both versions of YAST converge much faster than the algorithms derived from DPM, and the steady-state error obtained with YAST is lower. Figure 2-b shows the departures from orthogonality of the four above-mentioned subspace trackers, as defined in equation (43). It can be noticed that YAST, YAST-PGS and FDPM are stable, whereas the orthonormality error obtained with HFRANS keeps growing.

Note that this example, usually encountered in the literature, does not belong to the field of times series analysis. Therefore the optimization mentioned in section III-E cannot be applied, and the global complexity of YAST becomes n^2 in this case[10].

2) Time series analysis: Here the YAST algorithm is applied to frequency estimation, so that its complexity is only $5nr$. The test signal is a sum of 4 complex sinusoidal sources, of same unitary amplitude, random phases, and normalized frequencies equal to 0.2, 0.4, 0.5 and 0.8. This signal is perturbed by an additive white Gaussian noise, so that the SNR is 30 dB. The data vectors are composed of $n = 12$ successive samples of the noisy signal, so that the dimension of the noise subspace is $r = 8$. However, the $n \times n$ matrix $\boldsymbol{C}$ is defined as the covariance matrix of the noiseless data vectors. The performance of the various subspace trackers is measured by the functions $\rho(t)$ and $\eta(t)$ defined above.

Figure 3-a shows the tracking results obtained with the four MST's, with the same parameters as above. The same observations can be made regarding the convergence properties of the various algorithms in terms of subspace estimation, except that our new implementation of YAST converges even faster than YAST-PGS. This is due to the negligible approximation introduced in section III-C, which surprisingly improves the convergence rate of YAST. Regarding the departures from orthogonality, represented in Figure 3-b, it can be noticed that the YAST-PGS algorithm slowly loses the orthogonality, whereas FDPM and our new implementation of YAST remain perfectly stable.

[10]Note that even in this unfavorable case, YAST remains competitive in comparison with the subspace trackers of linear complexity, since its much higher convergence rate reduces the total number of operations required for reaching a given error level.

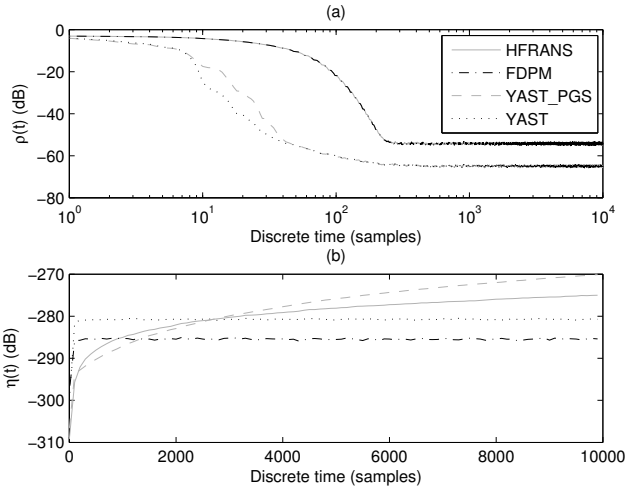

Fig. 3. Time series analysis
(a) Error trajectories,
(b) Departure from orthonormality.

V. Conclusions

In this paper, we presented a new implementation of the YAST algorithm for principal and minor subspace tracking, whose convergence and numerical stability are proved in Appendix B and Appendix C. The exceptional performance of YAST observed in our simulation results, both in terms of convergence rate and steady state error, is inherited from the SP1 algorithm by C.E. Davila [12]. Besides, the remarkable stability of this new implementation was obtained by applying an orthonormalization procedure similar to that of the FDPM algorithm by X.G. Doukopoulos and G.V. Moustakides [2], [11]. Regarding the computational complexity, if the data correlation matrix satisfies a shift invariance property, which typically holds in the case of time series analysis for instance, YAST involves only $5nr$ flops (instead of n^2 flops otherwise). This is to be compared to the $3nr$ complexity of the fastest PST algorithms, PAST [3] and FAPI [6], and the $4nr$ complexity of the fastest MST algorithms, among which NOOJA [10] and HFRANS [8].

For the clarity of the presentation, the proposed implementation involves a simple window shape for updating the correlation matrix, and a simple additional search direction at each iteration. Nevertheless, this implementation can be easily generalized to various window shapes and search directions, as proposed in [13]–[16].

APPENDIX

A. Asymptotic analysis

Proof of lemma III-C.1 (page 4): Substituting equations (2), (4), and (14) into equation (5) yields

$$\boldsymbol{\Pi}(t) - \boldsymbol{\Pi}(t-1) = \boldsymbol{u}(t)\boldsymbol{u}(t)^H - \underline{\boldsymbol{v}}(t)\underline{\boldsymbol{v}}(t)^H.$$

This equation can be interpreted as follows: the new subspace, represented by the projector $\boldsymbol{\Pi}(t)$, is obtained by adding the search direction $\boldsymbol{u}(t)$, and then removing the direction $\underline{\boldsymbol{v}}(t)$ to the old subspace represented by $\boldsymbol{\Pi}(t-1)$. Then left and right multiplying this equation by $\boldsymbol{u}(t)$ and substituting equations (6), (14) and (24) yields

$$\boldsymbol{u}(t)^H \left(\boldsymbol{\Pi}(t) - \boldsymbol{\Pi}(t-1)\right) \boldsymbol{u}(t) = \varepsilon(t)^2.$$

Finally, since the matrix 2-norm is the norm induced by the Hermitian vector norm, we obtain $\varepsilon(t)^2 \leq \|\boldsymbol{\Pi}(t) - \boldsymbol{\Pi}(t-1)\|_2$. ∎

Proof of lemma III-C.2 (page 4): Equations (6) and (25) yield

$$\underline{\boldsymbol{v}}(t)^H \boldsymbol{T}(t) = \underline{\boldsymbol{\Phi}}(t)^H \underline{\boldsymbol{W}}(t)^H \left(\boldsymbol{W}(t-1) - \varepsilon(t)\boldsymbol{u}(t)\boldsymbol{\phi}(t)^H\right). \tag{44}$$

Substituting equations (14) and (24) into equation (44) yields

$$\underline{\boldsymbol{v}}(t)^H \boldsymbol{T}(t) = \varepsilon(t)(1-\varphi(t))\theta(t)^* \boldsymbol{\phi}(t)^H.$$

Noting that $1-\varphi(t) = \frac{\varepsilon(t)^2}{1+\varphi(t)}$, we finally obtain equation (26). ∎

B. Convergence analysis

In this section we show that, assuming stationary data vectors of constant correlation matrix $\boldsymbol{C}_{xx}$, the YAST algorithm converges to the true principal or minor subspace of $\boldsymbol{C}_{xx}$. More precisely, assuming that $\forall t$, $\boldsymbol{C}_{xx}(t) = \boldsymbol{C}_{xx}$, and that the r-dimensional principal or minor subspace of $\boldsymbol{C}_{xx}$ is unique[11], we show that the range space of the orthogonal projector $\boldsymbol{\Pi}(t) = \boldsymbol{W}(t)\boldsymbol{W}(t)^H$ converges to this subspace. In order to simplify the discussion, the following proof focuses on PST, but the same proof stands for MST. This proof is decomposed into 5 steps. We thus successively prove that:

1) there is a subsequence $\{\boldsymbol{\Pi}(t_n)\}$ of $\{\boldsymbol{\Pi}(t)\}$, which converges toward a rank-r orthogonal projector $\boldsymbol{\Pi}^*$;
2) $\boldsymbol{\Pi}^*$ is a critical point of the generalized Rayleigh quotient on the manifold $\mathbb{P}$ of rank-r orthogonal projectors;
3) $\mathcal{J}(\boldsymbol{\Pi}(t))$ is an increasing convergent sequence;
4) $\boldsymbol{\Pi}^*$ is the global maximum of the generalized Rayleigh quotient on $\mathbb{P}$;
5) $\mathcal{J}(\boldsymbol{\Pi}(t))$ converges to $\mathcal{J}(\boldsymbol{\Pi}^*)$;
6) the whole sequence $\boldsymbol{\Pi}(t)$ converges to $\boldsymbol{\Pi}^*$.

We first need to introduce the *natural gradient* $\nabla_{\mathbb{P}}\mathcal{J}$ of function $\mathcal{J}$ on the manifold $\mathbb{P}$. By definition, $\nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi})$ is the orthogonal projection of the gradient of function $\mathcal{J}$ onto the tangent space of $\mathbb{P}$ at point $\boldsymbol{\Pi}$. More precisely, it can be proved that the natural gradient of the generalized Rayleigh quotient is[12]

$$\nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi}) = \boldsymbol{\Pi}\boldsymbol{C}_{xx} + \boldsymbol{C}_{xx}\boldsymbol{\Pi} - 2\boldsymbol{\Pi}\boldsymbol{C}_{xx}\boldsymbol{\Pi}. \tag{45}$$

This expression shows that the critical points of the generalized Rayleigh quotient (for which $\nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi}) = \boldsymbol{0}$) are the rank-$r$ projectors which span an *invariant subspace* of $\boldsymbol{C}_{xx}$.

Step 1): Since $\mathbb{P}$ is a compact set, the Bolzano-Weierstrass theorem proves that there is a subsequence $\{\boldsymbol{\Pi}(t_n)\}$ of $\{\boldsymbol{\Pi}(t)\}$, which converges toward a rank-r orthogonal projector $\boldsymbol{\Pi}^*$.

Step 2): Since $\boldsymbol{\Pi}(t)$ is the global maximum of the generalized Rayleigh quotient $\mathcal{J}$ on the subset of orthogonal projectors whose range space is included in $\mathrm{span}(\underline{\boldsymbol{\Pi}}(t))$, $\nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi}(t)) \perp \underline{\boldsymbol{\Pi}}(t)$. In particular, since $\boldsymbol{x}(t) \in \mathrm{span}(\underline{\boldsymbol{\Pi}}(t))$,

$$\forall t,\ \boldsymbol{x}(t) \perp \mathrm{span}(\nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi}(t))). \tag{46}$$

Let $\varepsilon > 0$. Since $\mathcal{J}$ is $\mathcal{C}^1$, $\exists n_0$ such that $\forall n \geq n_0$ $\|\nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi}^*) - \nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi}(t_n))\|_2 < \varepsilon$. Therefore $\forall n \geq n_0$, equation (46) yields

$$\begin{aligned} & \left\| \frac{\boldsymbol{x}(t_n)}{\|\boldsymbol{x}(t_n)\|_2}^H \nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi}^*) \right\|_2 \\ = & \left\| \frac{\boldsymbol{x}(t_n)}{\|\boldsymbol{x}(t_n)\|_2}^H \left(\nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi}^*) - \nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi}(t_n))\right) \right\|_2 \\ \leq & \|\nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi}^*) - \nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi}(t_n))\|_2 < \varepsilon. \end{aligned}$$

However $\boldsymbol{x}(t_n)$ is a sequence of independent random vectors of full-rank covariance matrix; thus any family of n vectors in this sequence almost surely spans the whole space $\mathbb{C}^n$. This finally proves that $\nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi}^*) = 0$, *i.e.* $\boldsymbol{\Pi}^*$ is a critical point of the generalized Rayleigh quotient on $\mathbb{P}$.

Step 3): Substituting equations (5), (4), (2), and (14) into equation (1) yields

$$\mathcal{J}\left(\boldsymbol{\Pi}(t)\right) = \mathcal{J}\left(\boldsymbol{\Pi}(t\text{-}1)\right) + \boldsymbol{u}(t)^H \boldsymbol{C}_{xx}\boldsymbol{u}(t) - \underline{\boldsymbol{v}}(t)^H \boldsymbol{C}_{xx}\underline{\boldsymbol{v}}(t).$$

In this equation, both unitary vectors $\boldsymbol{u}(t)$ and $\underline{\boldsymbol{v}}(t)$ belong to the augmented subspace $\mathrm{span}(\underline{\boldsymbol{\Pi}}(t))$. However, $\underline{\boldsymbol{v}}(t)$ corresponds the global minimum or the function $\boldsymbol{w} \mapsto \boldsymbol{w}^H\boldsymbol{C}_{xx}\boldsymbol{w}/\|\boldsymbol{w}\|_2^2$ within this subspace. Thus $\mathcal{J}(\boldsymbol{\Pi}(t))$ increases at each iteration. Moreover, since $\mathcal{J}(\boldsymbol{\Pi}(t))$ is bounded, $\mathcal{J}(\boldsymbol{\Pi}(t))$ is an increasing convergent sequence.

Step 4): Therefore the sequence $\boldsymbol{\Pi}(t_n)$, which converges to a critical point $\boldsymbol{\Pi}^*$ of $\mathcal{J}$ on $\mathbb{P}$, is such that $\mathcal{J}(\boldsymbol{\Pi}(t_n))$ is increasing. However it is well known that the generalized Rayleigh quotient admits a unique local maximum on $\mathbb{P}$, which is the global maximum (other critical points are either unstable

[11] In the case of PST, the r-dimensional principal subspace is uniquely defined if and only if $\lambda_r > \lambda_{r+1}$, where $\lambda_1 \geq \ldots \geq \lambda_n \geq 0$ are the eigenvalues of $\boldsymbol{C}_{xx}$ sorted in decreasing order.

[12] Let $\mathbb{T}_{\boldsymbol{\Pi}}$ denote the tangent space of the manifold $\mathbb{P}$ at point $\boldsymbol{\Pi}$, and note that the standard gradient of function $\mathcal{J}$ is the constant matrix $\boldsymbol{C}_{xx}$. By definition, $\nabla_{\mathbb{P}}\mathcal{J}(\boldsymbol{\Pi})$ is the orthogonal projection of $\boldsymbol{C}_{xx}$ onto $\mathbb{T}_{\boldsymbol{\Pi}}$. Equation (45) thus follows from the following observations:

- $\mathbb{T}_{\boldsymbol{\Pi}}$ is the $(nr-r^2)$-dimensional space of Hermitian matrices $\boldsymbol{X}$ such that $\boldsymbol{\Pi}\boldsymbol{X} + \boldsymbol{X}\boldsymbol{\Pi} = \boldsymbol{X}$.
- $\mathbb{T}_{\boldsymbol{\Pi}}$ is also the linear space of all matrices of the form $\boldsymbol{\Pi}\boldsymbol{H} + \boldsymbol{H}\boldsymbol{\Pi} - 2\boldsymbol{\Pi}\boldsymbol{H}\boldsymbol{\Pi}$, where $\boldsymbol{H}$ is Hermitian.
- the orthogonal complement $\mathbb{T}_{\boldsymbol{\Pi}}^{\perp}$ of $\mathbb{T}_{\boldsymbol{\Pi}}$ is the (n^2+r^2-nr)-dimensional space of all matrices of the form $\boldsymbol{A} + \boldsymbol{H} - \boldsymbol{\Pi}\boldsymbol{H} - \boldsymbol{H}\boldsymbol{\Pi} + 2\boldsymbol{\Pi}\boldsymbol{H}\boldsymbol{\Pi}$, where $\boldsymbol{A}$ is skew-Hermitian and $\boldsymbol{H}$ is Hermitian.
- Let $\boldsymbol{M}$ be any $n \times n$ matrix, and let $\boldsymbol{A}$ denote its skew-Hermitian part and $\boldsymbol{H}$ denote its Hermitian part. Then the orthogonal projection of $\boldsymbol{M}$ onto $\mathbb{T}_{\boldsymbol{\Pi}}$ is $\boldsymbol{\Pi}\boldsymbol{H} + \boldsymbol{H}\boldsymbol{\Pi} - 2\boldsymbol{\Pi}\boldsymbol{H}\boldsymbol{\Pi}$, and the orthogonal projection of $\boldsymbol{M}$ onto $\mathbb{T}_{\boldsymbol{\Pi}}^{\perp}$ is $\boldsymbol{A} + \boldsymbol{H} - \boldsymbol{\Pi}\boldsymbol{H} - \boldsymbol{H}\boldsymbol{\Pi} + 2\boldsymbol{\Pi}\boldsymbol{H}\boldsymbol{\Pi}$.

saddle points, or the global minimum). Thus $\mathbf{\Pi}^*$ is the unique global maximum of the generalized Rayleigh quotient on $\mathbb{P}$.

Step 5): Since $\mathcal{J}(\mathbf{\Pi}(t_n))$ converges to the maximum value $\mathcal{J}^*$ of the generalized Rayleigh quotient and the whole sequence $\mathcal{J}(\mathbf{\Pi}(t))$ is increasing, $\mathcal{J}(\mathbf{\Pi}(t))$ also converges to $\mathcal{J}^*$.

Step 6): Finally, let prove by contradiction that the whole sequence $\mathbf{\Pi}(t)$ converges to $\mathbf{\Pi}^*$. Suppose that $\exists\ \varepsilon > 0$ and a subsequence $\mathbf{\Pi}(t_k)$ of $\mathbf{\Pi}(t)$ such that $\forall k$, $\|\mathbf{\Pi}(t_k) - \mathbf{\Pi}^*\|_2 \geq \varepsilon$. Let $\mathbb{P}_\varepsilon$ be the set of all rank-r orthogonal projectors $\mathbf{\Pi}$ such that $\|\mathbf{\Pi} - \mathbf{\Pi}^*\|_2 \geq \varepsilon$, and let $\mathcal{J}_\varepsilon^*$ be the supremum of $\mathcal{J}$ on $\mathbb{P}_\varepsilon$. Since $\mathcal{J}$ is a continuous function, and $\mathbb{P}_\varepsilon$ is a compact set, the extreme value theorem proves that $\mathcal{J}$ attains its supremum $\mathcal{J}_\varepsilon^*$ on $\mathbb{P}_\varepsilon$. However $\mathbf{\Pi}^*$ is the unique global maximum of $\mathcal{J}$ on $\mathbb{P}$; thus $\mathcal{J}^* > \mathcal{J}_\varepsilon^*$. Consequently, the subsequence $\mathbf{\Pi}(t_k)$ satisfies $\forall k$, $\mathcal{J}(\mathbf{\Pi}(t_k)) \leq \mathcal{J}_\varepsilon^* < \mathcal{J}^*$. This contradicts the fact that $\mathcal{J}(\mathbf{\Pi}(t))$ converges to $\mathcal{J}^*$. As a conclusion, the whole sequence $\mathbf{\Pi}(t)$ converges to $\mathbf{\Pi}^*$.

C. Numerical stability

As mentioned in the introduction, our previous implementations of YAST [13], [15] suffered from numerical stability problems when the algorithm had converged, as observed in [17] in the case of PST, and as proved in [18] in the case of MST. In this section, we prove that our new implementation is stable both for PST and MST. Indeed, suppose that, because of rounding errors, $\boldsymbol{W}(t-1)^H\boldsymbol{W}(t-1) = \boldsymbol{I}_r + \boldsymbol{E}(t-1)$ where $\boldsymbol{E}(t-1) \neq \mathbf{0}$, and let show that the orthogonality error $\boldsymbol{E}(t)$ induced at time t is smaller than that at time $t-1$. In the following developments, we first show that matrix $\boldsymbol{E}(t)$ satisfies a simple recursion. Then, noting that this recursion fits into the framework of the proof presented in [2], we refer to this proof to conclude about the numerical stability of YAST.

First, let $\boldsymbol{e}(t) = -\boldsymbol{W}(t-1)^H\boldsymbol{u}(t)$. Equations (10) and (13) show that $\boldsymbol{e}(t) = \frac{\boldsymbol{E}(t-1)\boldsymbol{y}(t)}{\sigma(t)}$. Then matrix $\boldsymbol{T}(t)$ defined in equation (25) satisfies

$$\begin{aligned}\boldsymbol{T}(t)^H\boldsymbol{T}(t) =& \boldsymbol{I}_r + \varepsilon(t)^2\boldsymbol{\phi}(t)\boldsymbol{\phi}(t)^H + \boldsymbol{E}(t-1)\\ &+\varepsilon(t)(\boldsymbol{e}(t)\boldsymbol{\phi}(t)^H + \boldsymbol{\phi}(t)\boldsymbol{e}(t)^H)\\ =& \boldsymbol{I}_r + \mathcal{O}(\varepsilon(t)^2) + \boldsymbol{E}(t-1) + \mathcal{O}(\varepsilon(t)\|\boldsymbol{E}(t-1)\|_2).\end{aligned}$$

As mentioned in section III-C, when the algorithm has reached its steady state, $\varepsilon(t) \ll 1$. Thus the previous equation can be simplified in the form $\boldsymbol{T}(t)^H\boldsymbol{T}(t) \approx \boldsymbol{I}_r + \boldsymbol{E}(t-1)$, where $\approx$ stands for the approximation $\varepsilon(t) \ll 1$. Then equation (31) yields $\boldsymbol{Q}(t)^H\boldsymbol{Q}(t) \approx \boldsymbol{I}_r + \boldsymbol{S}(t)\boldsymbol{E}(t-1)\boldsymbol{S}(t)$. Since matrix $\boldsymbol{W}(t)$ is obtained by normalizing the first column of $\boldsymbol{Q}(t)$, this last equation shows that matrix $\boldsymbol{D}(t)$ in equation (32) is of the form $\boldsymbol{D}(t) = \boldsymbol{I}_r - \frac{\epsilon(t)}{2}\boldsymbol{e}_1\,\boldsymbol{e}_1^T$, where $\epsilon(t) = \mathcal{O}(\|\boldsymbol{E}(t-1)\|_2)$ and $\boldsymbol{e}_1 = [1, 0 \dots 0]^T$. Finally, equation (32) yields $\boldsymbol{W}(t)^H\boldsymbol{W}(t) \approx \boldsymbol{I}_r + \boldsymbol{E}(t) + \mathcal{O}(\|\boldsymbol{E}(t)\|_2^2)$, where

$$\boldsymbol{E}(t) = \boldsymbol{S}(t)\boldsymbol{E}(t-1)\boldsymbol{S}(t) - \epsilon(t)\boldsymbol{e}_1\boldsymbol{e}_1^T. \tag{47}$$

However it can be noticed that the top left coefficient of $\boldsymbol{E}(t)$ is necessary zero since the first column of $\boldsymbol{W}(t)$ is explicitly normalized. Thus $\epsilon(t)$ is also the top left coefficient of matrix $\boldsymbol{S}(t)\boldsymbol{E}(t-1)\boldsymbol{S}(t)$, *i.e.* $\epsilon(t) = \boldsymbol{e}_1^T\boldsymbol{S}(t)\boldsymbol{E}(t-1)\boldsymbol{S}(t)\boldsymbol{e}_1$.

By considering the column-wise version $\mathrm{col}\{\boldsymbol{E}(t)\}$ of the error $\boldsymbol{E}(t)$ (*i.e.* where the matrix elements are read in a column by column manner), equation (47) can be rewritten in the form

$$\mathrm{col}\{\boldsymbol{E}(t)\} = \boldsymbol{P}_{r^2}\ (\boldsymbol{S}(t) \otimes \boldsymbol{S}(t))\ \mathrm{col}\{\boldsymbol{E}(t\text{-}1)\} \tag{48}$$

where $\otimes$ denotes the Kronecker product, and $\boldsymbol{P}_{r^2} = \boldsymbol{I}_{r^2} - (\boldsymbol{e}_1 \otimes \boldsymbol{e}_1)(\boldsymbol{e}_1^T \otimes \boldsymbol{e}_1^T)$. Based on an equation similar to equation (48), detailed mathematical developments proving the numerical stability of FDPM were provided in Appendix III in [2], in the framework of stochastic approximation theory. These developments can also be applied here, to finally prove the stability of the YAST algorithm. The basic idea is the following: the algorithm is stable if and only if the eigenvalues of $\boldsymbol{P}_{r^2}\ (\boldsymbol{S}(t) \otimes \boldsymbol{S}(t))$ are strictly smaller than one. This is almost surely the case because $\boldsymbol{P}_{r^2}$ is a rank deficient orthogonal projector (which generally lowers the vector norm), and $\boldsymbol{S}(t) \otimes \boldsymbol{S}(t)$ is a random unitary matrix (which maintains the vector norm).

REFERENCES

[1] A.-J. Van der Veen, E. F. Deprettere, and A. L. Swindlehurst, "Subspace based signal analysis using singular value decomposition," *Proc. of IEEE*, vol. 81, no. 9, pp. 1277–1308, Sept. 1993.

[2] X. Doukopoulos and G. Moustakides, "Fast and Stable Subspace Tracking," *IEEE Trans. Signal Processing*, vol. 56, no. 4, pp. 1452–1465, Apr. 2008.

[3] B. Yang, "Projection Approximation Subspace Tracking," *IEEE Trans. Signal Processing*, vol. 44, no. 1, pp. 95–107, Jan. 1995.

[4] P. Strobach, "Low-rank adaptive filters," *IEEE Transactions on Signal Processing*, vol. 44, no. 12, pp. 2932–2947, Dec. 1996.

[5] K. Abed-Meraim, A. Chkeif, and Y. Hua, "Fast orthonormal PAST algorithm," *IEEE Signal Proc. Letters*, vol. 7, no. 3, pp. 60–62, Mar. 2000.

[6] R. Badeau, B. David, and G. Richard, "Fast Approximated Power Iteration Subspace Tracking," *IEEE Trans. Signal Processing*, vol. 53, no. 8, pp. 2931–294, Aug. 2005.

[7] P. Strobach, "Square root QR inverse iteration for tracking the minor subspace," *IEEE Trans. Signal Processing*, vol. 48, no. 11, pp. 2994–2999, 2000.

[8] S. Attallah, "The Generalized Rayleigh's Quotient Adaptive Noise Subspace Algorithm: A Householder Transformation-Based Implementation," *IEEE Trans. Circuits Syst. II*, vol. 53, no. 1, pp. 3–7, Jan. 2006.

[9] S. C. Douglas, "Numerically-robust adaptive subspace tracking using Householder transformations," in *Proc. of Sensor Array and Multichannel Signal Proc. Workshop*. IEEE, 2000, pp. 499–503.

[10] S. Attallah and K. Abed-Meraim, "Fast algorithms for subspace tracking," *IEEE Signal Proc. Letters*, vol. 8, no. 7, pp. 203–206, 2001.

[11] X. Doukopoulos and G. Moustakides, "Blind Adaptive Channel Estimation in OFDM Systems," *IEEE Trans. Wireless Commun.*, vol. 5, no. 7, pp. 1716–1725, July 2006.

[12] C. E. Davila, "Efficient, high performance, subspace tracking for time-domain data," *IEEE Trans. Signal Processing*, vol. 48, no. 12, pp. 3307–3315, Dec. 2000.

[13] R. Badeau, B. David, and G. Richard, "Yet Another Subspace Tracker," in *Proc. of 2005 International Conference on Acoustics, Speech, and Signal Processing ICASSP'05*, vol. 4. IEEE, Mar. 2005, pp. 329–332.

[14] R. Badeau, "Méthodes à haute résolution pour l'estimation et le suivi de sinusoïdes modulées. Application aux signaux de musique." Ph.D. dissertation, École Nationale Supérieure des Télécommunications, Paris, Apr. 2005, in French.

[15] R. Badeau, B. David, and G. Richard, "YAST Algorithm for Minor Subspace Tracking," in *Proc. of 2006 International Conference on Acoustics, Speech, and Signal Processing ICASSP'06*, vol. 3. Toulouse, France: IEEE, May 2006, pp. 552–555.

[16] ——, "Conjugate gradient algorithms for minor subspace analysis," in *2007 International Conference on Acoustics, Speech, and Signal Processing ICASSP'07*, vol. 3, Honolulu, Hawaii, USA, Apr. 2007, pp. 1013–1016.

[17] J. Altuna, B. Mulgrew, R. Badeau, and V. Atxa, "A fast adaptive method for subspace based blind channel estimation," in *Proc. of 2006 International Conference on Acoustics, Speech, and Signal Processing ICASSP'06*, vol. 4. Toulouse, France: IEEE, May 2006, pp. 1121–1124.

[18] S. Bartelmaos and K. Abed-Meraim, "An efficient & stable algorithm for minor subspace tracking and stability analysis," in *Proc. of ICASSP'07*, vol. III. Honolulu, Hawaii, USA: IEEE, Apr. 2007, pp. 1301–1304.

[19] E. Oja, "Principal components, minor components, and linear neural networks," *Neural Networks*, vol. 5, pp. 927–935, Nov./Dec. 1992.

[20] G. H. Golub and C. F. Van Loan, *Matrix computations*, 3rd ed. Baltimore and London: The Johns Hopkins University Press, 1996.

[21] P. Strobach, "Fast recursive subspace adaptive ESPRIT algorithms," *IEEE Trans. Signal Processing*, vol. 46, no. 9, pp. 2413–2430, Sept. 1998.

[22] R. Badeau, G. Richard, and B. David, "Sliding window adaptive SVD algorithms," *IEEE Trans. Signal Processing*, vol. 52, no. 1, pp. 1–10, Jan. 2004.

[23] Y. Miao and Y. Hua, "Fast subspace tracking and neural network learning by a novel information criterion," *IEEE Trans. Signal Processing*, vol. 46, no. 7, pp. 1967–1979, July 1998.

[24] J. Yang and M. Kaveh, "Adaptive Eigensubspace Algorithms for Direction or Frequency Estimation and Tracking," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 36, no. 2, pp. 241–251, Feb. 1988.

FAST ADAPTIVE ESPRIT ALGORITHM

Roland Badeau, Gaël Richard and Bertrand David

École Nationale Supérieure des Télécommunications
Département de Traitement du Signal et des Images
46 rue Barrault, 75634 Paris Cedex 13 France
roland.badeau, gael.richard, bertrand.david@enst.fr

ABSTRACT

The ESPRIT algorithm[1] is a subspace-based high resolution method used in source localization and spectral analysis. It relies on the rotational invariance property of the signal subspace, and provides accurate estimates of the signal parameters. However, its main drawback is a high computational cost.

In an adaptive context, some very fast algorithms were proposed to robustly track the variations of the signal subspace. Based on these subspace trackers, we propose a new adaptive implementation of the ESPRIT algorithm, faster than the existing methods.

1. INTRODUCTION

The ESPRIT algorithm [1] first consists in estimating the signal subspace. Then a so-called *spectral matrix* is derived from this subspace, whose eigenvalues characterize the signal parameters.

Adaptive implementations of the ESPRIT algorithm for tracking the spectral matrix were already proposed in the literature, based on specific subspace trackers. For example, the technique presented in [2] relies on a *Singular Value Decomposition* (SVD) updating algorithm [3], and that proposed in [4] is based on the rank-revealing URV decomposition [5]. The complexity of these methods is proportional to $O(n^2)$ operations for each time update, where n is the dimension of the data vectors. In [6], other adaptive algorithms were proposed for the LORAF [7] and the Bi-SVD [8] subspace trackers. In comparison with [2] and [4], the complexity of these algorithms is reduced to $O(nr^2)$ or $O(nr)$, where r is the dimension of the signal subspace (in practice $r << n$). However we observed that only the $O(nr^2)$ algorithms in [6] are of practical interest, due to the degraded performance of the $O(nr)$ subspace trackers. In [9], we proposed a new adaptive algorithm based on the PAST [10] and related subspace trackers, whose complexity is $O(nr)$. Its performance is the same as that of the $O(nr^2)$ algorithms in [6]. Below, we propose a more efficient implementation of the approach introduced in [9].

The second step of the ESPRIT algorithm consists in computing the eigenvalues of the $r \times r$ estimated spectral matrix. They can be obtained via an *Eigen-Value Decomposition* (EVD), whose complexity is $O(r^3)$, as proposed in [4]. However they can also be recursively updated. A first approach proposed in [11] relies on perturbation theory, but we observed that it fails to robustly track frequency variations. Another technique proposed in [6] relies on the *split-Schur* theorem, but proves to be less precise than an EVD, for a similar computational cost. We propose in this paper a new method wich exactly and efficiently updates the eigenvalues of the spectral matrix at each time step.

This paper is organized as follows. Section 2 briefly describes the ESPRIT method. Section 3 introduces a new algorithm for tracking the spectral matrix. Then our adaptive EVD algorithm is presented in section 4. The performance of this fast adaptive ESPRIT method is compared to that of existing methods in section 5. Section 6 presents the analysis of an audio signal. Finally, section 7 summarizes the main conclusions of this paper.

2. ESTIMATION OF SIGNAL PARAMETERS VIA ROTATIONAL INVARIANCE TECHNIQUES

In source localization and spectral analysis, the observed data can be arranged into a sequence of n-dimensional vectors $\{\boldsymbol{x}(t)\}$, where $t \in \mathbb{Z}$. If there is no noise, it can be shown that these vectors belong to the *signal subspace* spanned by the Vandermonde matrix

$$\boldsymbol{V} = \begin{bmatrix} 1 & \cdots & 1 \\ z_1 & \cdots & z_r \\ \vdots & \cdots & \vdots \\ {z_1}^{n-1} & \cdots & {z_r}^{n-1} \end{bmatrix},$$

where r is the model order and $\forall k \in \{1 \dots r\}$, $z_k \in \mathbb{C} - \{0\}$ is a complex pole of the signal. In practice, the matrix $\boldsymbol{V}$ is unknown. However, in presence of white noise, a $n \times r$ matrix $\boldsymbol{W}(t)$ spanning the signal subspace can be estimated from the vectors observed in the analysis window. This operation involves the SVD of a data matrix, or the EVD of a correlation matrix, which requires $O(n^3)$ operations. In an adaptive context, the signal variations can be tracked by time-shifting the analysis window. The complexity of the subspace estimation can be reduced by using a fast subspace tracker which efficiently updates $\boldsymbol{W}(t)$.

The ESPRIT algorithm relies on the rotational invariance property of the Vandermonde matrix $\boldsymbol{V}$. Let $\boldsymbol{V}_\downarrow$ be the matrix extracted from $\boldsymbol{V}$ by deleting the last row. Similarly, let $\boldsymbol{V}_\uparrow$ be the matrix extracted from $\boldsymbol{V}$ by deleting the first row. Then it can be shown that $\boldsymbol{V}_\uparrow = \boldsymbol{V}_\downarrow \boldsymbol{D}$, where $\boldsymbol{D} = \mathrm{diag}(z_1 \dots z_r)$. The estimated matrix $\boldsymbol{W}(t)$ satisfies a similar invariance property: $\boldsymbol{W}(t)_\uparrow = \boldsymbol{W}(t)_\downarrow \boldsymbol{\Phi}(t)$, where the eigenvalues of the $r \times r$ matrix $\boldsymbol{\Phi}(t)$, herein called *spectral matrix*, are the complex poles z_k.

[1] ESPRIT stands for *Estimation of Signal Parameters via Rotational Invariance Techniques*.

Consequently, the LS-ESPRIT algorithm [1] consists in:

1. estimating the subspace weighting matrix $\boldsymbol{W}(t)$,
2. computing $\boldsymbol{\Phi}(t)$ by means of the *Least Squares* (LS) method:
$$\boldsymbol{\Phi}(t) = \boldsymbol{W}(t)_{\downarrow}{}^{\dagger} \boldsymbol{W}(t)_{\uparrow} \quad (1)$$
(where $\dagger$ denotes the Moore-Penrose pseudo-inverse),
3. estimating the complex poles z_k as the eigenvalues of $\boldsymbol{\Phi}(t)$.

3. SPECTRAL MATRIX TRACKING

As mentioned in section 2, the first step of the ESPRIT algorithm consists in estimating the subspace weighting matrix $\boldsymbol{W}(t)$. Although no assumption is made about the structure of this matrix, many subspace estimation techniques produce an orthonormal matrix $\boldsymbol{W}(t)$. In section 3.1, we show that this property can be taken into account in the second step of the ESPRIT method, in order to efficiently compute the spectral matrix. Besides, in an adaptive context, the complexity can be further reduced by exploiting the structure of the update involved in the subspace tracker, as shown in section 3.2. We thus obtain a fast algorithm for tracking the spectral matrix. Its cost is lower than that of the algorithm that we proposed in [9], which did not exploit the orthonormality of $\boldsymbol{W}(t)$.

3.1. Fast computation of the spectral matrix

If the matrix $\boldsymbol{W}(t)_{\downarrow}$ is full-rank (which is always the case in practice), equation (1) can be written in the form

$$\boldsymbol{\Phi}(t) = \boldsymbol{\Omega}(t)\, \boldsymbol{\Psi}(t) \quad (2)$$

where $\boldsymbol{\Omega}(t)$ and $\boldsymbol{\Psi}(t)$ are the $r \times r$ matrices

$$\boldsymbol{\Omega}(t) \triangleq \left(\boldsymbol{W}_{\downarrow}(t)^H \boldsymbol{W}_{\downarrow}(t) \right)^{-1}, \quad (3)$$

$$\boldsymbol{\Psi}(t) \triangleq \boldsymbol{W}_{\downarrow}(t)^H \boldsymbol{W}_{\uparrow}(t). \quad (4)$$

The computation of $\boldsymbol{\Psi}(t)$ involves nr^2 MAC[2], and that of $\boldsymbol{\Omega}(t)$ involves $nr^2 + O(r^3)$ MAC. The overall computation of $\boldsymbol{\Phi}(t)$ as defined in equation (2) therefore involves $2nr^2 + O(r^3)$ MAC. It will be shown below that this cost can be reduced by assuming that the matrix $\boldsymbol{W}(t)$ is orthonormal.

Indeed, if $\boldsymbol{W}(t)$ is orthonormal, $\boldsymbol{\Omega}(t)$ can be easily calculated. The equation $\boldsymbol{W}(t)^H \boldsymbol{W}(t) = \boldsymbol{I}_r$ shows that $\boldsymbol{W}_{\downarrow}(t)^H \boldsymbol{W}_{\downarrow}(t) = \boldsymbol{I}_r - \boldsymbol{\nu}(t)\, \boldsymbol{\nu}(t)^H$, where $\boldsymbol{\nu}(t)^H$ is the last row of $\boldsymbol{W}(t)$. Thus $\boldsymbol{W}_{\downarrow}(t)^H \boldsymbol{W}_{\downarrow}(t)$ is a rank-one modification of the identity matrix. Finally, the matrix inversion lemma [12, pp. 18-19] shows that

$$\boldsymbol{\Omega}(t) = \boldsymbol{I}_r + \frac{1}{1 - \|\boldsymbol{\nu}(t)\|^2}\, \boldsymbol{\nu}(t)\, \boldsymbol{\nu}(t)^H.$$

Substituting this last equation into equation (2) yields

$$\boxed{\boldsymbol{\Phi}(t) = \boldsymbol{\Psi}(t) + \frac{1}{1-||\boldsymbol{\nu}(t)||^2}\, \boldsymbol{\nu}(t)\, \boldsymbol{\varphi}(t)^H} \quad (5)$$

where the r-dimensional vector $\boldsymbol{\varphi}(t)$ is defined as

$$\boldsymbol{\varphi}(t) = \boldsymbol{\Psi}(t)^H \boldsymbol{\nu}(t). \quad (6)$$

Thus the matrix $\boldsymbol{\Phi}(t)$ is a rank-one modification of $\boldsymbol{\Psi}(t)$. By successively computing equations (4), (6) and (5), the calculation of the spectral matrix involves only $nr^2 + O(r^2)$ MAC.

[2] In this paper, the computational costs are expressed in terms of Multiply-ACcumulate (MAC) operations.

Table 1. Adaptive computation of the spectral matrix

	Eq.	Cost
$\boldsymbol{e}_-(t) = \boldsymbol{W}_{\downarrow}(t-1)^H \boldsymbol{e}_{\uparrow}(t)$	(8)	nr
$\boldsymbol{e}_+(t) = \boldsymbol{W}_{\uparrow}(t-1)^H \boldsymbol{e}_{\downarrow}(t)$	(9)	nr
$\boldsymbol{e}'_+(t) = \boldsymbol{e}_+(t) + \boldsymbol{g}(t)\, \left(\boldsymbol{e}_{\uparrow}(t)^H \boldsymbol{e}_{\downarrow}(t)\right)$	(10)	n
$\boldsymbol{\Psi}(t) = \boldsymbol{\Psi}(t-1) + \boldsymbol{e}_-(t)\, \boldsymbol{g}(t)^H + \boldsymbol{g}(t)\, \boldsymbol{e}'_+(t)^H$	(11)	$2r^2$
$\boldsymbol{\varphi}(t) = \boldsymbol{\Psi}(t)^H \boldsymbol{\nu}(t)$	(6)	r^2
$\boldsymbol{\Phi}(t) = \boldsymbol{\Psi}(t) + \frac{1}{1-||\boldsymbol{\nu}(t)||^2}\, \boldsymbol{\nu}(t)\, \boldsymbol{\varphi}(t)^H$	(5)	r^2

3.2. Adaptive computation of the spectral matrix

Several subspace trackers such as PAST [10] and NIC [13] provide a rank-one update for $\boldsymbol{W}(t)$:

$$\boldsymbol{W}(t) = \boldsymbol{W}(t-1) + \boldsymbol{e}(t)\, \boldsymbol{g}(t)^H \quad (7)$$

where $\boldsymbol{e}(t)$ and $\boldsymbol{g}(t)$ are n and r-dimensional vectors. Here we focus on those which additionally guarantee the orthonormality of $\boldsymbol{W}(t)$, such as in [14–17], in order to use the results in section 3.1. In this case, it will be shown below that $\boldsymbol{\Psi}(t)$ satisfies a similar recursion. Indeed, let

$$\boldsymbol{e}_-(t) = \boldsymbol{W}_{\downarrow}(t-1)^H \boldsymbol{e}_{\uparrow}(t), \quad (8)$$

$$\boldsymbol{e}_+(t) = \boldsymbol{W}_{\uparrow}(t-1)^H \boldsymbol{e}_{\downarrow}(t), \quad (9)$$

$$\boldsymbol{e}'_+(t) = \boldsymbol{e}_+(t) + \boldsymbol{g}(t)\, \left(\boldsymbol{e}_{\uparrow}(t)^H \boldsymbol{e}_{\downarrow}(t)\right). \quad (10)$$

Substituting equations (7), (8), (9) and (10) into equation (4) shows that $\boldsymbol{\Psi}(t)$ satisfies a rank-two recursion:

$$\boxed{\boldsymbol{\Psi}(t) = \boldsymbol{\Psi}(t-1) + \boldsymbol{e}_-(t)\, \boldsymbol{g}(t)^H + \boldsymbol{g}(t)\, \boldsymbol{e}'_+(t)^H.} \quad (11)$$

This recursion leads to an adaptive computation of the spectral matrix, as shown in table 1. The global cost is only $n(2r+1) + 4r^2$ MAC (instead of in $8nr + O(r^2)$ in [9]).

Besides, the last row of equation (7) can be written in the form:

$$\boldsymbol{\nu}(t) = \boldsymbol{\nu}(t-1) + \boldsymbol{g}(t)\, e_n(t)^* \quad (12)$$

where $e_n(t)$ is the last coefficient of $\boldsymbol{e}(t)$. Substituting equations (11) and (12) into (5) yields a rank-three recursion for $\boldsymbol{\Phi}(t)$:

$$\begin{aligned} \boldsymbol{\Phi}(t) = {} & \boldsymbol{\Phi}(t-1) + \boldsymbol{g}(t)\, \boldsymbol{e}''_+(t)^H \\ & + \boldsymbol{e}_-(t)\, \boldsymbol{g}(t)^H + \boldsymbol{\nu}(t-1)\, \Delta\boldsymbol{\varphi}(t)^H, \end{aligned} \quad (13)$$

where

$$\boldsymbol{e}''_+(t) = \boldsymbol{e}'_+(t) + \frac{e_n(t)}{1 - \|\boldsymbol{\nu}(t)\|^2}\, \boldsymbol{\varphi}(t),$$

$$\Delta\boldsymbol{\varphi}(t) = \frac{\boldsymbol{\varphi}(t)}{1 - \|\boldsymbol{\nu}(t)\|^2} - \frac{\boldsymbol{\varphi}(t-1)}{1 - \|\boldsymbol{\nu}(t-1)\|^2}.$$

Below, equation (13) will be compacted in the form

$$\boldsymbol{\Phi}(t) = \boldsymbol{\Phi}(t-1) + \underline{\boldsymbol{a}}(t)\, \underline{\boldsymbol{b}}(t)^H \quad (14)$$

where $\underline{\boldsymbol{a}}(t)$ and $\underline{\boldsymbol{b}}(t)$ are $r \times 3$ matrices. Note that the additional calculation of $\boldsymbol{e}''_+(t)$ and $\Delta\boldsymbol{\varphi}(t)$ only involves $O(r)$ MAC.

4. EIGENVALUES TRACKING

The eigenvalues of $\boldsymbol{\Phi}(t)$ can be obtained by computing its EVD, which involves $O(r^3)$ MAC. However, in an adaptive context, we propose in this section a fast algorithm for diagonalizing $\boldsymbol{\Phi}(t)$, which exploits the structure of the update in equation (14).

Suppose that an EVD of the $r \times r$ matrix $\boldsymbol{\Phi}(t-1)$ is known:

$$\boldsymbol{D}(t-1) = \boldsymbol{G}'(t-1)^H \, \boldsymbol{\Phi}(t-1) \, \boldsymbol{G}(t-1) \tag{15}$$

where $\boldsymbol{D}(t-1) = \mathrm{diag}\{z_1(t-1) \dots z_r(t-1)\}$, $\boldsymbol{G}(t-1)$ is non-singular, and $\boldsymbol{G}'(t-1) = \boldsymbol{G}(t-1)^{-H}$. Then define the $r \times r$ matrix

$$\widetilde{\boldsymbol{D}}(t) \triangleq \boldsymbol{G}'(t-1)^H \, \boldsymbol{\Phi}(t) \, \boldsymbol{G}(t-1). \tag{16}$$

Note that the eigenvalues of $\widetilde{\boldsymbol{D}}(t)$ are the same as those of $\boldsymbol{\Phi}(t)$. Substituting equations (14) and (15) into equation (16) yields

$$\widetilde{\boldsymbol{D}}(t) = \boldsymbol{D}(t-1) + \underline{\widetilde{\boldsymbol{a}}}(t) \, \underline{\widetilde{\boldsymbol{b}}}(t)^H, \tag{17}$$

where

$$\underline{\widetilde{\boldsymbol{a}}}(t) = \boldsymbol{G}'(t-1)^H \underline{\boldsymbol{a}}(t) \tag{18}$$

$$\underline{\widetilde{\boldsymbol{b}}}(t) = \boldsymbol{G}(t-1)^H \underline{\boldsymbol{b}}(t). \tag{19}$$

Thus the matrix $\widetilde{\boldsymbol{D}}(t)$ is a low-rank modification of $\boldsymbol{D}(t-1)$. Now suppose that $\widetilde{\boldsymbol{D}}(t)$ is diagonalizable in the form

$$\boldsymbol{D}(t) = \widetilde{\boldsymbol{G}}'(t)^H \, \widetilde{\boldsymbol{D}}(t) \, \widetilde{\boldsymbol{G}}(t), \tag{20}$$

where $\boldsymbol{D}(t) = \mathrm{diag}\{z_1(t) \dots z_r(t)\}$, $\widetilde{\boldsymbol{G}}(t)$ is non-singular, and $\widetilde{\boldsymbol{G}}'(t) = \widetilde{\boldsymbol{G}}(t)^{-H}$. Then an EVD of $\boldsymbol{\Phi}(t)$ can be derived from this last EVD. Indeed, substituting equation (20) into (16) yields

$$\boldsymbol{D}(t) = \boldsymbol{G}'(t)^H \, \boldsymbol{\Phi}(t) \, \boldsymbol{G}(t),$$

where the $r \times r$ matrices $\boldsymbol{G}(t)$ and $\boldsymbol{G}(t)$ are non-singular:

$$\boldsymbol{G}(t) = \boldsymbol{G}(t-1) \, \widetilde{\boldsymbol{G}}(t), \tag{21}$$

$$\boldsymbol{G}'(t) = \boldsymbol{G}'(t-1) \, \widetilde{\boldsymbol{G}}'(t). \tag{22}$$

Note that the column vectors $\{\widetilde{\boldsymbol{g}}_1(t), \dots, \widetilde{\boldsymbol{g}}_r(t)\}$ of the matrix $\widetilde{\boldsymbol{G}}(t)$ are the right eigenvectors of the matrix $\widetilde{\boldsymbol{D}}(t)$. In the same way, the column vectors $\{\widetilde{\boldsymbol{g}}'_0(t), \dots, \widetilde{\boldsymbol{g}}'_{r-1}(t)\}$ of the matrix $\widetilde{\boldsymbol{G}}'(t)$ are the left eigenvectors of the matrix $\widetilde{\boldsymbol{D}}(t)$. Below, proposition 1 describes the eigenstructure of the matrix $\widetilde{\boldsymbol{D}}(t)$. Its proof can be found in appendix A.

Proposition 1. *Suppose that all eigenvalues of $\boldsymbol{\Phi}(t-1)$ and $\boldsymbol{\Phi}(t)$ are simple, and that all eigenvalues of $\boldsymbol{\Phi}(t)$ differ from those of $\boldsymbol{\Phi}(t-1)$. Consider the 3×3 matrix*

$$\underline{\boldsymbol{\phi}}(z;t) \triangleq \underline{\boldsymbol{I}} - \underline{\widetilde{\boldsymbol{b}}}(t)^H \, (z\boldsymbol{I}_r - \boldsymbol{D}(t-1))^{-1} \, \underline{\widetilde{\boldsymbol{a}}}(t), \tag{23}$$

defined for all $z \in \mathbb{C} - \{\mathrm{eig}(\boldsymbol{\Phi}(t-1))\}$.

Then the r eigenvalues of $\boldsymbol{\Phi}(t)$ match the roots of the equation

$$\det(\underline{\boldsymbol{\phi}}(z;t)) = 0. \tag{24}$$

If $z_k(t)$ is such an eigenvalue, the corresponding right eigenvector $\widetilde{\boldsymbol{g}}_k(t)$ and left eigenvector $\widetilde{\boldsymbol{g}}'_k(t)$ of $\widetilde{\boldsymbol{D}}(t)$ can be written in the form

$$\widetilde{\boldsymbol{g}}_k(t) = (z_k(t) \, \boldsymbol{I}_r - \boldsymbol{D}(t-1))^{-1} \, \underline{\widetilde{\boldsymbol{a}}}(t) \, \underline{v}_k(t), \tag{25}$$

$$\widetilde{\boldsymbol{g}}'_k(t) = (z_k(t) \, \boldsymbol{I}_r - \boldsymbol{D}(t-1))^{-1*} \, \underline{\widetilde{\boldsymbol{b}}}(t) \, \underline{v}'_k(t), \tag{26}$$

where

$$\underline{v}_k(t) \in \ker\left\{\underline{\boldsymbol{\phi}}(z_k(t);t)\right\}, \tag{27}$$

$$\underline{v}'_k(t) \in \ker\left\{\underline{\boldsymbol{\phi}}(z_k(t);t)^H\right\}. \tag{28}$$

Proposition 1 suggests a fast way of diagonalizing $\widetilde{\boldsymbol{D}}(t)$, from which the EVD of $\boldsymbol{\Phi}(t)$ can be deduced. The pseudo-code of the complete algorithm is summarized in table 2. Its complexity is $2r^3 + O(r^2)$, instead of $O(r^3)$ for a full EVD[3].

Table 2. Adaptive computation of the eigenvalues

	Eq.	Cost
$\underline{\widetilde{\boldsymbol{a}}}(t) = \boldsymbol{G}'(t-1)^H \, \underline{\boldsymbol{a}}(t)$	(18)	$3r^2$
$\underline{\widetilde{\boldsymbol{b}}}(t) = \boldsymbol{G}(t-1)^H \underline{\boldsymbol{b}}(t)$	(19)	$3r^2$
Compute[4] the roots $\{z_k(t)\}$ of $\det\left(\underline{\boldsymbol{\phi}}(z;t)\right) = 0$	(24)	$O(r^2)$
For $k = 1$ to r do		$r\times$
$\underline{\boldsymbol{\phi}}_k(t) = \underline{\boldsymbol{\phi}}(z_k(t);t)$	(23)	$O(1)$
Extract $\underline{v}_k(t)$ from $\ker\left\{\underline{\boldsymbol{\phi}}_k(t)\right\}$	(27)	$O(1)$
$\widetilde{\boldsymbol{g}}_k(t) \propto (z_k(t) \, \boldsymbol{I}_r - \boldsymbol{D}(t-1))^{-1} \, \underline{\widetilde{\boldsymbol{a}}}(t) \, \underline{v}_k(t)$	(25)	$3r$
Extract $\underline{v}'_k(t)$ from $\ker\left\{\underline{\boldsymbol{\phi}}_k(t)^H\right\}$	(28)	$O(1)$
$\widetilde{\boldsymbol{g}}'_k(t) \propto (z_k(t) \, \boldsymbol{I}_r - \boldsymbol{D}(t-1))^{-1*} \, \underline{\widetilde{\boldsymbol{b}}}(t) \, \underline{v}'_k(t)$	(26)	$3r$
End for		
Re-normalize the columns of $\widetilde{\boldsymbol{G}}(t)$ and $\widetilde{\boldsymbol{G}}'(t)$		
in order to obtain $\widetilde{\boldsymbol{G}}'(t)^H \widetilde{\boldsymbol{G}}(t) = \boldsymbol{I}_r$		$4r^2$
$\boldsymbol{G}(t) = \boldsymbol{G}(t-1) \, \widetilde{\boldsymbol{G}}(t)$	(21)	r^3
$\boldsymbol{G}'(t) = \boldsymbol{G}'(t-1) \, \widetilde{\boldsymbol{G}}'(t)$	(22)	r^3

5. SIMULATION RESULTS

In this section, the performance of the tracking is illustrated in the context of frequency estimation. The test signal is a sum of $r = 4$ complex sinusoidal sources plus a complex white gaussian noise (the SNR is 5.7 dB). The frequencies of the sinusoids vary according to a scenario illustrated in figure 1-a: two of them are constant and have close values (below the Fourier resolution limit), the third one is sinusoidally modulated, and the last one is a discontinuous, piecewise linear function. The tracking of these frequencies is performed by means of three algorithms, whose computational costs are summarized in table 3: a non-adaptive estimation scheme based on the EVD of the signal correlation matrix, the adaptive ESPRIT algorithm proposed in [6], and our new methods for tracking the spectral matrix and its eigenvalues, associated to the YAST subspace tracker [16].

Figures 1-b to 1-d show the frequency tracking results, obtained with parameters[5] $n = 80$ and $\beta = 0.99$. The dotted lines

[3] The computation of an EVD is an iterative process whose exact operation count depends on the particular matrix to be decomposed. However, the number of MAC involved is much greater than $2r^3$ in usual cases.

[4] It can be shown that $\det(\underline{\boldsymbol{\phi}}(z;t))$ is proportional to the characteristic polynomial of $\widetilde{\boldsymbol{D}}(t)$. The roots of this equation can be computed via standard numerical methods.

[5] β is the forgetting factor of the exponential window used for estimating the signal correlation matrix.

Fig. 1. Frequency estimation and tracking

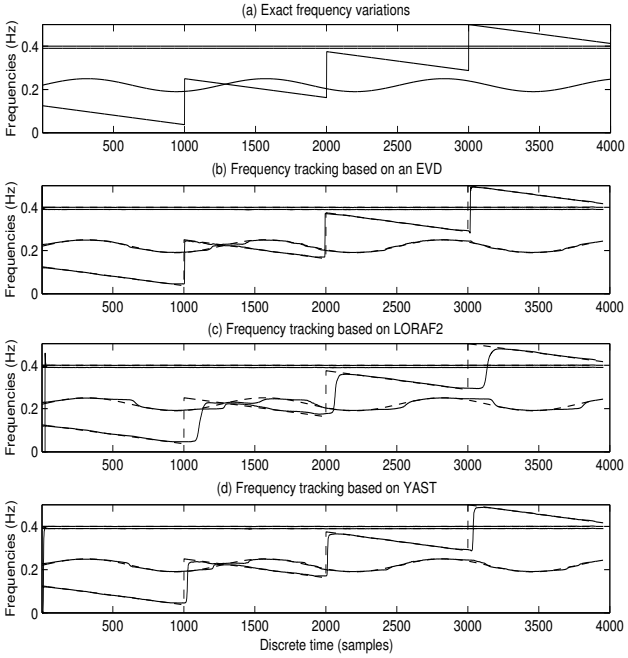

Table 3. Complexities of various adaptive ESPRIT algorithms

Step	*Algo. 1*	*Algo. 2*	*Algo. 3*
1. Subspace estimation	EVD $O(n^3)$	LORAF2 [7] nr^2	YAST [16] $4nr$
2. Spectral matrix	eq. (1) $2nr^2$	ref. [6] nr^2	Table 1 $2nr$
3. Complex poles	EVD $O(r^3)$	EVD $O(r^3)$	Table 2 $2r^3$

indicate the true frequency parameters, while the solid lines indicate the estimated frequencies. The frequency trajectories obtained with the EVD-based algorithm are displayed in figure 1-b. This result is to be compared to that obtained with the adaptive ESPRIT algorithm [6] associated with the LORAF2 subspace tracker [7], represented in figure 1-c. The response of this algorithm to frequency variations is slower. Finally, figure 1-d shows the trajectories estimated by the tracking methods presented in this paper, associated to the YAST subspace tracker [16]. The performance of this technique is closer to that of the EVD-based algorithm. Besides, it is the less computationally demanding approach.

6. APPLICATION TO AUDIO

This section illustrates the application of our adaptive ESPRIT algorithm to a musical signal. The study deals with two piano tones, C5 and G5, sampled at 11025 Hz.

First, it must be mentioned that audio signals often require some pre-processing before estimating and tracking the sinusoids. For example, signals with a high number of sinusoids (typically low-pitched sounds) may first be decomposed into several sub-band signals (via filtering/decimating, as proposed in [18]). In the example proposed below, this pre-processing is not used, since the chosen piano tones have few sinusoidal components. In addition, it is well known that the energy of audio signals is not evenly distributed over frequencies. To compensate the energy decrease, we use a pre-emphasis filter obtained by linear prediction at order 4.

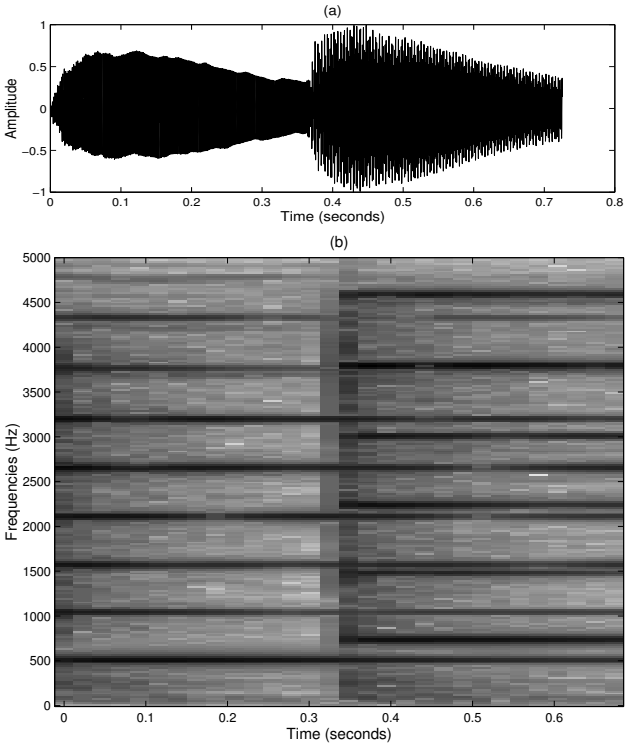

Fig. 2. *Time waveform (a) and spectrogram (b) of the piano signal for the frequency tracking test.*

The frequency tracking method presented in sections 3 and 4 is applied to the following signal: the C5 tone is played at time $t = 0s$, then the G5 is played at time $t = 0.36s$, while the C5 is maintained (the time waveform and the spectrogram of this signal are represented in figure 2). Figure 3 shows the result of the tracking. The analysis parameters are $r = 28$ (*i.e.* 14 real sinusoids), $n = 280$ (*i.e.* 25ms), and $\beta = 0.99$, and the sinusoids energies (estimated by means of the LS method) are represented on a logarithmic scale using gray levels for the plot. Since the number of sinusoids is over-estimated in the first half of the sound, some spurious poles are detected. However, the algorithm proves to be robust to the attack of the second tone.

7. CONCLUSIONS

In this paper, we proposed a fully adaptive implementation of the ESPRIT algorithm. This new approach can be used in conjunction with any subspace tracker which involves a rank-one update and guarantees the orthonormality of the subspace weighting matrix, such as [14–17]. These two conditions are necessary for efficiently updating the spectral matrix $\mathbf{\Phi}(t)$, by means of a rank-three update. The resulting tracking method is faster than the algorithm that we proposed in [9], and than any other existing method. Be-

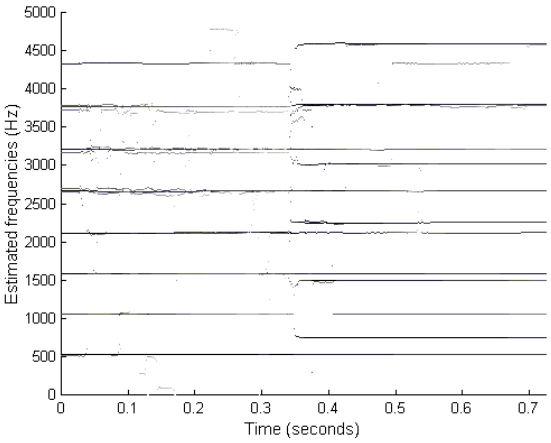

Fig. 3. *Frequency tracking of the piano signal.*

sides, the EVD of the spectral matrix can be updated by taking the structure of the recursion for $\mathbf{\Phi}(t)$ into account. The overall complexity of our new adaptive ESPRIT algorithm is only $2nr + 2r^3$ MAC. Our simulation results show that this algorithm robustly tracks abrupt frequency variations, and can be successfully applied to audio signals. It reaches the best performance when used in conjunction with the YAST subspace tracker [16].

A. APPENDIX

Proof of proposition 1. Let $z_k(t)$ be an eigenvalue of $\widetilde{\boldsymbol{D}}(t)$, and $\widetilde{\boldsymbol{g}}_k(t)$ and $\widetilde{\boldsymbol{g}}'_k(t)$ the corresponding right and left eigenvectors. Then

$$(z\boldsymbol{I}_r - \boldsymbol{D}(t-1))\,\widetilde{\boldsymbol{g}} = \widetilde{\underline{\boldsymbol{a}}}(t)\left(\widetilde{\underline{\boldsymbol{b}}}(t)^H\widetilde{\boldsymbol{g}}\right), \quad (29)$$

$$(z\boldsymbol{I}_r - \boldsymbol{D}(t-1))^*\,\widetilde{\boldsymbol{g}}' = \widetilde{\underline{\boldsymbol{b}}}(t)\left(\widetilde{\underline{\boldsymbol{a}}}(t)^H\widetilde{\boldsymbol{g}}'\right). \quad (30)$$

Define

$$\underline{v}_k(t) \triangleq \widetilde{\underline{\boldsymbol{b}}}(t)^H\widetilde{\boldsymbol{g}}_k(t), \quad (31)$$

$$\underline{v}'_k(t) \triangleq \widetilde{\underline{\boldsymbol{a}}}(t)^H\widetilde{\boldsymbol{g}}'_k. \quad (32)$$

If $z_k(t)$ is not an eigenvalue of $\mathbf{\Phi}(t-1)$, the matrix $z_k(t)\boldsymbol{I}_r - \boldsymbol{D}(t-1)$ is non-singular. Then equation (25) is obtained by left-multiplying equation (29) by $(z\boldsymbol{I}_r - \boldsymbol{D}(t-1))^{-1}$. Substituting equation (25) into equation (31) yields equation (27).

Since the vector $\widetilde{\boldsymbol{g}}_k(t)$ is non-zero, $\underline{v}_k(t)$ is also non-zero, thus the matrix $\phi(z_k(t);t)$ is singular and its determinant is zero. Then equation (26) is obtained by left-multiplying equation (30) by $(z\boldsymbol{I}_r - \boldsymbol{D}(t-1))^{-1*}$. Substituting equation (26) into equation (32) finally yields equation (28). Conversely, if $z_k(t)$ is defined as a root of the equation (24), then the vectors $\underline{v}_k(t)$ and $\underline{v}'_k(t)$ can be defined as in equations (27) and (28), and the vectors $\widetilde{\boldsymbol{g}}_k(t)$ and $\widetilde{\boldsymbol{g}}_k(t)$ can be defined as in equations (25) and (26). Then it can be readily verified that $\widetilde{\boldsymbol{g}}_k(t)$ and $\widetilde{\boldsymbol{g}}_k(t)$ are the right and left eigenvectors of $\widetilde{\boldsymbol{D}}(t)$ associated to the eigenvalue $z_k(t)$. Thus the eigenvalues of $\widetilde{\boldsymbol{D}}(t)$ match the roots of equation (24). □

B. REFERENCES

[1] R. Roy and T. Kailath, "ESPRIT-Estimation of Signal Parameters via Rotational Invariance Techniques," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 37, no. 7, pp. 984–995, July 1989.

[2] M. Moonen, F. J. Vanpoucke, and E. F. Deprettere, "Parallel and adaptive high-resolution direction finding," *IEEE Trans. Signal Processing*, vol. 42, no. 9, pp. 2439–2448, Sept. 1994.

[3] M. Moonen, P. Van Dooren, and J. Vandewalle, "An SVD updating algorithm for subspace tracking," *SIAM J. Matrix Ana. Appl.*, vol. 13, no. 4, pp. 1015–1038, 1992.

[4] K. J. R. Liu, D. P. O'Leary, G. W. Stewart, and Y-J. J. Wu, "URV ESPRIT for tracking time-varying signals," *IEEE Trans. Signal Processing*, vol. 42, no. 12, pp. 3443–3448, Dec. 1994.

[5] G. W. Stewart, "An updating algorithm for subspace tracking," *IEEE Trans. Signal Processing*, vol. 40, pp. 1535–1541, June 1992.

[6] P. Strobach, "Fast recursive subspace adaptive ESPRIT algorithms," *IEEE Trans. Sig. Proc.*, vol. 46, no. 9, pp. 2413–2430, Sept. 1998.

[7] P. Strobach, "Low-rank adaptive filters," *IEEE Trans. Signal Processing*, vol. 44, no. 12, pp. 2932–2947, Dec. 1996.

[8] P. Strobach, "Bi-iteration SVD subspace tracking algorithms," *IEEE Trans. Signal Proc.*, vol. 45, no. 5, pp. 1222–1240, May 1997.

[9] R. Badeau, G. Richard, and B. David, "Adaptive ESPRIT algorithm based on the PAST subspace tracker," in *Proc. of ICASSP'03*, Hong Kong, China, Apr. 2003, vol. 6, pp. 229–232, IEEE.

[10] B. Yang, "Projection Approximation Subspace Tracking," *IEEE Trans. Signal Processing*, vol. 44, no. 1, pp. 95–107, Jan. 1995.

[11] Q. G. Liu and B. Champagne, "An adaptive ESPRIT algorithm based on perturbation of unsymmetrical matrices," in *Proc. 8th EUSIPCO*, Trieste, Italy, Sept. 1996, pp. 539–542.

[12] R. A. Horn and C. R. Johnson, *Matrix analysis*, Cambridge University Press, Cambridge, 1985.

[13] Y. Miao and Y. Hua, "Fast subspace tracking and neural network learning by a novel information criterion," *IEEE Trans. Signal Processing*, vol. 46, no. 7, pp. 1967–1979, July 1998.

[14] K. Abed, A. Chkeif, and Y. Hua, "Fast orthonormal PAST algorithm," *IEEE Sig. Proc. Letters*, vol. 7, no. 3, pp. 60–62, Mar. 2000.

[15] S. C. Douglas, "Numerically-robust adaptive subspace tracking using Householder transformations," in *SAM'00*. 2000, pp. 499–503, IEEE.

[16] R. Badeau, B. David, and G. Richard, "Yet Another Subspace Tracker," in *Proc. of ICASSP'05*. Mar. 2005, vol. 4, pp. 329–332, IEEE.

[17] R. Badeau, B. David, and G. Richard, "Fast Approximated Power Iteration Subspace Tracking," *IEEE Trans. Signal Processing*, vol. 53, no. 8, Aug. 2005, to be published.

[18] J. Laroche, "The use of the Matrix Pencil method for the spectrum analysis of musical signals," *Journal of the Acoustical Society of America*, vol. 94, no. 4, pp. 1958–1965, Oct. 1993.

Multipitch estimation of piano sounds using a new probabilistic spectral smoothness principle

Valentin Emiya, *Member, IEEE,* Roland Badeau, *Member, IEEE,* Bertrand David, *Member, IEEE*

***Abstract*—A new method for the estimation of multiple concurrent pitches in piano recordings is presented. It addresses the issue of overlapping overtones by modeling the spectral envelope of the overtones of each note with a smooth autoregressive model. For the background noise, a moving-average model is used and the combination of both tends to eliminate harmonic and sub-harmonic erroneous pitch estimations. This leads to a complete generative spectral model for simultaneous piano notes, which also explicitly includes the typical deviation from exact harmonicity in a piano overtone series. The pitch set which maximizes an approximate likelihood is selected from among a restricted number of possible pitch combinations as the one. Tests have been conducted on a large homemade database called MAPS, composed of piano recordings from a real upright piano and from high-quality samples.**

***Index Terms*—Acoustic signal analysis, audio processing, multipitch estimation, piano, transcription, spectral smoothness.**

I. Introduction

The issue of monopitch estimation has been addressed for decades by different approaches such as retrieving a periodic pattern in a waveform [1; 2] or matching a regularly spaced pattern to an observed spectrum [3–5], or even by combining both spectral and temporal cues [6; 7]. Conversely the multipitch estimation (MPE) problem has become a rather active research area in the last decade [8–11] and is mostly handled by processing spectral or time-frequency representations. This MPE task has also become a central tool in musical scene analysis [11], particularly when the targeted application is the automatic transcription of music (ATM) [12–16]. While recent works consider pitch and time dimensions jointly to perform this task, MPE on single frames has been historically used as a prior processing to pitch tracking over time and musical note detection. Indeed, in a signal processing perspective, a period or a harmonic series can be estimated from a short signal snapshot. In terms of perception, only a few cycles are needed to identify a pitched note [17].

In this polyphonic context, two issues have proved difficult and interesting for computational MPE: the overlap between the overtones of different notes and the unknown number of such notes occurring simultaneously. As the superposition of two sounds in octave relationship leads to a spectral ambiguity and thus to an ill-posed problem, some knowledge is often used for the further spectral modeling of the underlying sources. For instance, when iterative spectral subtraction is employed, several authors have adopted a smoothness assumption for the spectral envelope of the notes [8–10].

Several preceding works [12–14; 16] specifically address the case of the piano. One of the interests in studying a single instrument is to incorporate in the algorithms some results established from physical considerations [18; 19], in the hope that a more specific model will lead to a better performance, with the obvious drawback of narrowing the scope of the application. Indeed, while a large number of musical pieces are composed for piano solo, the performance of ATM systems is relatively poor for this instrument in comparison to others [6].

Two deviations from a simple harmonic spectral model are specific to free vibrating string instruments and are particularly salient for the piano: a small departure from exact harmonicity (the overtone series is slightly stretched) and the beatings between very close frequency components (for a single string two different vibrating polarizations occur and the different strings in a doublet or triplet are slightly detuned on purpose [20]).

In this work, we propose a new spectral model in which the inharmonic distribution is taken into account and adjusted for each possible note. The spectral envelope of the overtones is modeled by a smooth autoregressive (AR) model, the smoothness resulting from a low model order. A smooth spectral model is similarly introduced for the residual noise by using a low-order moving-average (MA) process, which is particular efficient against residual sinusoids. The proposed method follows a preliminary work [21] in which the AR/MA approach was already used and MPE was addressed as an extension of a monopitch estimation approach. In addition to the opportunity to deal with higher polyphony levels, the current paper introduces a new signal model for simultaneous piano notes and a corresponding estimation scheme. The major advance from the previous study is thus to model and estimate the spectral overlap between note spectra. An iterative estimation framework is proposed which leads to a maximum likelihood estimation for each possible set of F_0s.

This paper is structured as follows. In section II, the overall MPE approach is described. The sound model is first detailed, followed by the principle of the algorithm: statistical background, adaptive choice of the search space and detection function. The estimation of the model parameters is then explained in section III. After specifying the implementation details, the algorithm is tested in section IV using a database of piano sounds and the results are analyzed and compared with those obtained with some state-of-the-art approaches.

V. Emiya is with the Metiss team at INRIA, Centre Inria Rennes - Bretagne Atlantique, Rennes, France, and, together with R. Badeau and B. David, with Institut Télécom; Télécom ParisTech; CNRS LTCI, Paris, France.

The research leading to this paper was supported by the French GIP ANR under contract ANR-06-JCJC-0027-01, *Décomposition en Éléments Sonores et Applications Musicales - DESAM* and by the European Commission under contract FP6-027026-K-SPACE. The authors would also like to thanks Peter Weyer-Brown from Telecom ParisTech and Nathalie Berthet for their useful comments to improve the English usage.

Conclusions are finally drawn in section V.

Note that in the following, * denotes the complex conjugate, $^{\mathrm{T}}$ the transpose operator, † the conjugate transpose, $\lfloor . \rfloor$ the floor function and $|.|$ the number of elements in a set, the absolute value of a real number or the modulus of a complex number. In addition, the term *polyphony* will be applied to any possible mixture of notes, including silence (*polyphony 0*) and single notes (*polyphony 1*).

II. Multipitch estimation algorithm

A. Generative sound model

At the frame level, a mixture of piano sounds is modeled as a sum of sinusoids and background noise. The amplitude of each sinusoid is considered as a random variable. The spectral envelope for the overtone series of a note is introduced as second order statistical properties of this random variable, making it possible to adjust its smoothness. The noise is considered as a moving-average (MA) process.

More precisely, a mixture of $P \in \mathbb{N}$ simultaneous piano notes, observed in a discrete-time N-length frame, is modeled as a sum $x(t) \triangleq \sum_{p=1}^{P} x_p(t) + x_b(t)$ of the note signals x_p and of noise x_b. The sinusoidal model for x_p is

$$x_p(t) \triangleq \sum_{h=1}^{H_p} \left(\alpha_{hp} e^{2i\pi f_{hp} t} + \alpha_{hp}^* e^{-2i\pi f_{hp} t}\right) \quad (1)$$

where α_{hp} are the complex amplitudes and f_{hp} the frequencies of the H_p overtones (here, H_p is set to the maximum number of overtones below the Nyquist frequency). Note p is parameterized by $\mathcal{C}_p = (f_{0p}, \beta_p)$, f_{0p} being the *fundamental frequency* (F_0) and $\beta_p(f_{0p})$ being the so-called *inharmonicity coefficient* of the piano note [20], such that

$$f_{hp} \triangleq h f_{0p} \sqrt{1 + \beta_p(f_{0p})\, h^2} \quad (2)$$

We introduce a *spectral envelope* for note p as an autoregressive (AR) model of order Q_p, parameterized by $\theta_p \triangleq \left(\sigma_p^2, A_p(z)\right)$, where σ_p^2 is the power and $\frac{1}{A_p(z)}$ is the transfer function of the related AR filter. Order Q_p should be low and proportional to H_p, in order to obtain a smooth envelope and to avoid overfitting issues for high pitches. The *amplitude* α_{hp} of overtone h of note p is the outcome of a zero-mean complex Gaussian random variable[1], with variance equal to the spectral envelope power density at the frequency f_{hp} of the overtone:

$$\alpha_{hp} \sim \mathcal{N}\left(0, \frac{\sigma_p^2}{\left|A_p\left(e^{2i\pi f_{hp}}\right)\right|^2}\right) \quad (3)$$

In order to model a smooth spectral envelope without residual peaks, the *noise* x_b is modeled as an MA process with a low order Q_b, parameterized by $\theta_b \triangleq \left(\sigma_b^2, B(z)\right)$, σ_b^2 being the power and $B(z)$ the finite impulse response (FIR) filter of the process.

Furthermore, high note powers and low noise powers are favored by choosing an inverse gamma prior for σ_p^2: $\sigma_p^2 \sim \mathcal{IG}\left(k_{\sigma_p^2}, E_{\sigma_p^2}\right)$ and a gamma prior for σ_b^2: $\sigma_b^2 \sim \Gamma\left(k_{\sigma_b^2}, E_{\sigma_b^2}\right)$. A non-informative prior[2] is assumed for $A_p(z)$ and $B(z)$.

[1] Using these centered Gaussian variables, x_p is thus a harmonic process, in the statistical sense.

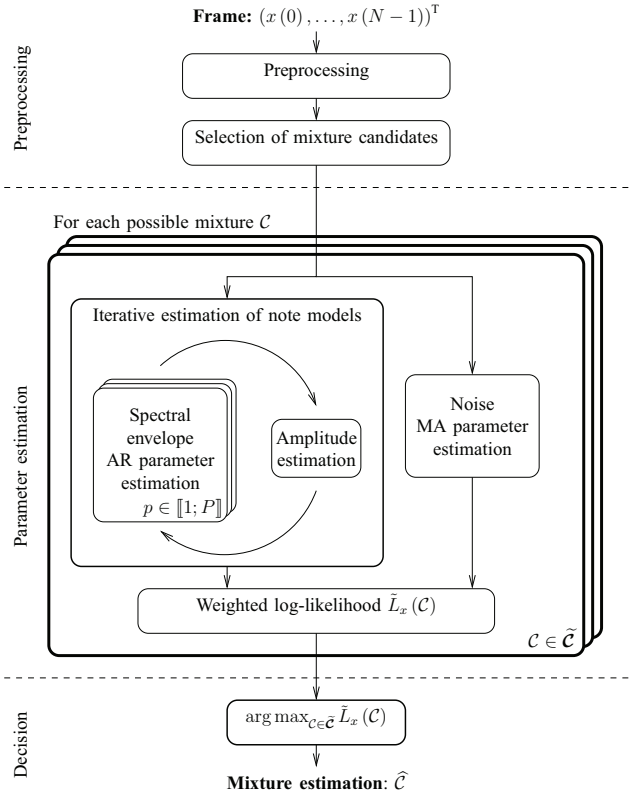

Fig. 1. Multipitch estimation block diagram.

With a user-defined weighting window w, we can equivalently consider the Discrete Time Fourier Transforms (DTFT) $X_p(f)$ of $x_p(t)\,w(t)$, $X_b(f)$ of $x_b(t)\,w(t)$ and $X(f) \triangleq \sum_{p=1}^{P} X_p(f) + X_b(f)$ of $x(t)\,w(t)$. W being the DTFT of w, we thus have

$$X_p(f) = \sum_{h=1}^{H_p} \left(\alpha_{hp} W(f - f_{hp}) + \alpha_{hp}^* W^*(f + f_{hp})\right) \quad (4)$$

In addition, the following synthetic notations will be used: a mixture of notes is denoted by $\mathcal{C} \triangleq (\mathcal{C}_1, \ldots, \mathcal{C}_P)$; the set of parameters for spectral envelope models by $\theta \triangleq (\theta_1, \ldots, \theta_P)$; the set of amplitudes of overtones of note p by $\alpha_p \triangleq \left(\alpha_{1p}, \ldots, \alpha_{H_p p}\right)$; and the set of amplitudes of overtones of all notes by $\alpha \triangleq (\alpha_1, \ldots, \alpha_P)$.

B. Principle of the algorithm

The algorithm is illustrated in Fig. 1. The main principles are described in the current section, except for the model parameter estimation, which is detailed in Section III.

[2] *e.g.* an improper, constant prior density, which will not affect the detection function.

1) Statistical background: Let us consider the observed frame x and the set $\boldsymbol{\mathcal{C}}$ of all the possible mixtures of piano notes. Ideally, the multipitch detection function can be expressed as the maximum *a posteriori* (MAP) estimator $\hat{\mathcal{C}}$, which is equivalent to the ML estimator if no *a priori* information on mixtures is specified:

$$\begin{aligned}\hat{\mathcal{C}} &= \arg\max_{\mathcal{C}\in\boldsymbol{\mathcal{C}}} p(\mathcal{C}|x) = \arg\max_{\mathcal{C}\in\boldsymbol{\mathcal{C}}} \frac{p(x|\mathcal{C})\,p(\mathcal{C})}{p(x)} \\ &= \arg\max_{\mathcal{C}\in\boldsymbol{\mathcal{C}}} p(x|\mathcal{C}) \end{aligned} \tag{5}$$

2) Search space: The number of combinations among Q possible notes being $\binom{Q}{P}$ for polyphony P, the size of the search set $\boldsymbol{\mathcal{C}}$ is $\sum_{P=0}^{Q}\binom{Q}{P} = 2^Q$, *i.e.* around $3{\cdot}10^{26}$ for a typical piano with $Q = 88$ keys. Even when the polyphony is limited to $P_{\max} = 6$ and the number of possible notes to $Q = 60$ (*e.g.* by ignoring the lowest and highest-pitched keys), the number of combinations reaches $\sum_{P=0}^{P_{\max}}\binom{Q}{P} \approx 56{\cdot}10^6$, which remains a too huge search set for a realistic implementation. In order to reduce the size of this set, we propose to select a given number N_c of note candidates. We use the normalized product spectrum function defined in dB as

$$\Pi_X(f_0,\beta) \triangleq \frac{1}{H(f_0,\beta)^{\nu}} 10\log \prod_{h=1}^{H(f_0,\beta)} |X(f_h)|^2 \tag{6}$$

where X is the observed spectrum, $H(f_0,\beta)$ is the number of overtones for the note with fundamental frequency f_0 and inharmonicity β, $f_h \triangleq hf_0\sqrt{1+\beta h^2}$, and ν is a parameter adjusted to balance the values of the function between bass and treble notes. As the true notes generate local peaks in Π_X, selecting an oversized set of N_c greatest peaks is an efficient way for adaptively reducing the number of note candidates. In addition, for candidate $n_c \in [\![1;N_c]\!]$, accurate values of its fundamental frequency f_{0n_c} and inharmonicity β_{n_c} are found by a local two-dimensional maximization[3] of $\Pi_X(f_{0n_c},\beta_{n_c})$. They are used in the subsequent estimation stages to locate the frequencies of the overtones. In the current implementation, the number of note candidates is set to $N_c = 9$. This choice results from the balance between increasing the number of candidates and limiting the computational time. The size of the set $\widetilde{\boldsymbol{\mathcal{C}}}$ of possible mixtures thus equals $\sum_{P=0}^{P_{\max}}\binom{N_c}{P} = 466$, for $P_{\max} = 6$.

3) Multipitch detection function: The multipitch detection function aims at finding the correct mixture among all the possible candidates. For each candidate $\mathcal{C}\in\widetilde{\boldsymbol{\mathcal{C}}}$, the parameters $\widehat{\alpha}$, $\widehat{\theta}$ and $\widehat{\theta_b}$ of the related model are estimated as explained in section III. The detection function is then computed, and the multipitch estimation is defined as the mixture with the greatest value.

Since equation (5) is intractable, we define an alternate detection function which involves the following terms:

- $L_p(\theta_p) \triangleq \ln p(\alpha_p|\theta_p,\mathcal{C}_p)$ is the log-likelihood of the amplitudes α_p of overtones of note p, related to the spectral envelope models θ_p;
- $L_b(\theta_b) \triangleq \ln p(x|\alpha,\theta_b,\mathcal{C})$ is equal[4] to the log-likelihood $\ln p_{x_b}$ related to the noise model θ_b;
- the priors $\ln p(\theta_p)$ on the spectral envelope of the note p and $\ln p(\theta_b)$ on noise parameters.

We empirically define the detection function as a weighted sum of these log-densities, computed with the estimated parameters (see Appendix A for a discussion):

$$\begin{aligned}\tilde{L}_x(\mathcal{C}) \triangleq\ & w_1\sum_{p=1}^{P}\tilde{L}_p\left(\widehat{\theta_p}\right)/P + w_2\tilde{L}_b\left(\widehat{\theta_b}\right) \\ & + w_3\sum_{p=1}^{P}\ln p\left(\widehat{\sigma_p^2}\right)/P + w_4\ln p\left(\widehat{\sigma_b^2}\right) - \mu_{\text{pol}}\,P \end{aligned} \tag{7}$$

where

$$\begin{cases} \tilde{L}_p\left(\widehat{\theta_p}\right) \triangleq \frac{1}{H_p}L_p\left(\widehat{\sigma_p^2},\widehat{A_p}\right) - \mu_{\text{env}}\,H_p \\ \tilde{L}_b\left(\widehat{\theta_b}\right) \triangleq \frac{1}{|\mathcal{F}_b|}L_b\left(\widehat{\sigma_b^2},\widehat{B}\right) - \mu_b\,|\mathcal{F}_b| \\ w_1,\ldots,w_4,\mu_{\text{pol}},\mu_{\text{env}},\mu_b \text{ are user-defined coefficients} \\ |\mathcal{F}_b| \text{ is the number of noisy bins (see eq. (14)).} \end{cases}$$

III. Model parameter estimation

A. Iterative estimation of spectral envelope parameters and amplitudes of notes

We consider a possible mixture $\mathcal{C} = (\mathcal{C}_1,\ldots,\mathcal{C}_P)$. The estimation of the unknown spectral envelope parameters θ and of the unknown amplitudes of the overtones α is performed by iteratively estimating the former and the latter, as described in this section.

1) Spectral envelope parameter estimation: let us assume the amplitudes α are known in order to estimate the spectral envelope parameters θ in the maximum likelihood (ML) sense. Given that note models are independent, the likelihood of α is $p(\alpha|\theta,\mathcal{C}) = \prod_{p=1}^{P}p(\alpha_p|\theta_p,\mathcal{C}_p)$. The optimization w.r.t. θ thus consists in maximizing the log-likelihood $L_p\left(\sigma_p^2,A_p\right) \triangleq \ln p\left(\alpha_p|\sigma_p^2,A_p,\mathcal{C}_p\right)$ w.r.t. $\theta_p = \left(\sigma_p^2,A_p\right)$, independently for each note p. As proved in the Appendix B, the maximization w.r.t. σ_p^2 leads to the expression

$$L_p\left(\widehat{\sigma_p^2},A_p\right) = c + \frac{H_p}{2}\ln\rho(A_p) \tag{8}$$

with

$$c \triangleq -\frac{H_p}{2}\ln(2\pi e) - \frac{1}{2}\sum_{h=1}^{H_p}\ln|\alpha_{hp}|^2 \tag{9}$$

$$\rho(A_p) \triangleq \frac{\left(\prod_{h=1}^{H_p}|\alpha_{hp}|^2\left|A_p\left(e^{2i\pi f_{hp}}\right)\right|^2\right)^{\frac{1}{H_p}}}{\frac{1}{H_p}\sum_{h=1}^{H_p}|\alpha_{hp}|^2\left|A_p\left(e^{2i\pi f_{hp}}\right)\right|^2} \tag{10}$$

$$\widehat{\sigma_p^2} = \frac{1}{H_p}\sum_{h=1}^{H_p}|\alpha_{hp}|^2\left|A_p\left(e^{2i\pi f_{hp}}\right)\right|^2 \tag{11}$$

[3] The `fminsearch` Matlab function was used for this optimization.

[4] by using the substitution $x \mapsto x - 2\text{Re}\left(\sum_{p=1}^{P}\sum_{h=1}^{H_p}\alpha_{hp}e^{2i\pi f_{hp}\mathbf{t}}\right)$, $\mathbf{t}$ being the time instants related to the frame, this term is equal to $\ln p_{x_b}\left(x - 2\text{Re}\left(\sum_{p=1}^{P}\sum_{h=1}^{H_p}\alpha_{hp}e^{2i\pi f_{hp}\mathbf{t}}\right)|\theta_b\right)$

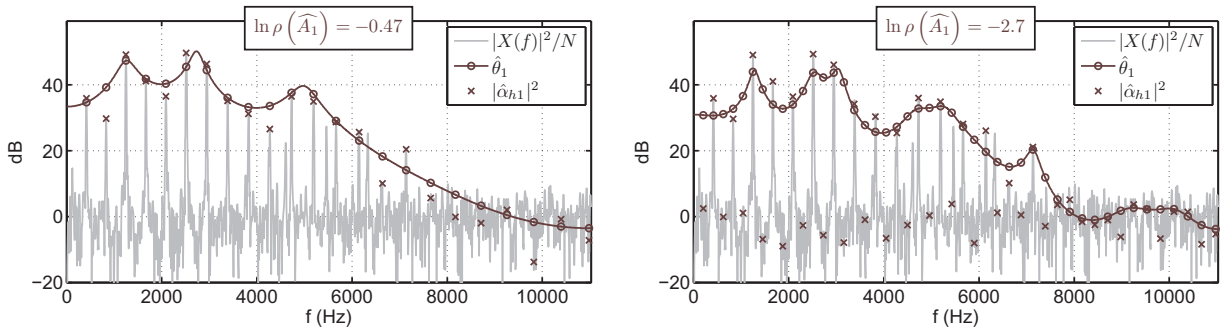

(a) True note: A♭$_4$. Selected model: A♭$_4$.

(b) True note: A♭$_4$. Selected model: A♭$_3$.

Fig. 2. Estimation of spectral envelopes: the AR estimation performs a better fitting of the amplitudes for the correct model (left) than for the sub-octave model (right).

As c is a constant w.r.t. A_p, the optimization consists in maximizing $\rho(A_p)$, which measures the spectral flatness (*i.e.* the ratio between the geometrical and arithmetical means) of $\left\{|\alpha_{hp}|^2 \left|A_p\left(e^{2i\pi f_{hp}}\right)\right|^2\right\}_{1\leq h\leq H_p}$. This quantity lies in $[0;1]$ and is maximum when the latter coefficients are constant, *i.e.* when the filter A_p perfectly whitens the amplitudes α_p. The estimation of A_p from the discrete data α_p thanks to the Digital All-Pole (DAP) method [22] leads to a solution that actually maximizes $\rho(A_p)$, as proved in [23]. It is applied here to the amplitudes α_p, for each note p.

The estimation of the spectral envelope of a note is illustrated in Fig. 2. A note is analyzed by considering two possible models: the model related to the true note (Fig. 2(a)) and the model related to its sub-octave (Fig. 2(b)). In the former case, the estimate of the spectral envelope is close to the amplitudes of the overtones, whereas in the latter case, the AR spectral envelope model is not adapted to amplitudes that are alternatively high and low: the low values obtained for the spectral flatness ρ (and, consequently, for the likelihood) are here a good criterion to reject wrong models like the sub-octave (see Fig. 6(a) for the whole spectral flatness curve).

2) Estimation of the amplitudes of the overtones: let us now assume that spectral envelope parameters θ are known, that the frequencies of the overtones may overlap, and that the amplitudes α are unknown. In all that follows, we assume that at a given overtone frequency f_{hp}, the power spectrum of the noise is not significant in comparison with the power spectrum of the amplitudes of the overtone.

When the overtone h of note p is not overlapping with any other overtone, the amplitude α_{hp} is directly given by the spectrum value $X(f_{hp})$. In the alternative case of overlapping overtones, the observed spectrum results from the contribution of overtones with close frequencies. We propose an estimate of the hidden random variable α, given the observation X and the parameters θ that control the second-order statistics of α.

As defined by eq. (3), the amplitude α_{hp} is a random variable such that $\alpha_{hp} \sim \mathcal{N}(0, v_{hp})$ with $v_{hp} \triangleq \frac{\sigma_p^2}{\left|A_p\left(e^{2i\pi f_{hp}}\right)\right|^2}$. One can rewrite the sound model x as a sum of K sinusoids and noise $x(t) = \sum_{k=1}^{K} \alpha_k e^{2i\pi f_k t} + x_b(t)$ with $\alpha_k \sim \mathcal{N}(0, v_k)$ and $K = 2\sum_{p=1}^{P} H_p$. In this formula, the couples of indexes (h,p) related to overtones and notes have been replaced by a single index k. In the spectral domain, the DFT X of the N-length frame x using a weighting window w of DFT W is $X(f) = \sum_{k=1}^{K} \alpha_k W(f - f_k)$, the noise spectrum being removed since we are only interested in the values of X at frequencies f_{hp}, where the noise term is insignificant.

In these conditions, for $1 \leq k_0 \leq K$, the optimal linear estimator $\hat{\alpha}_{k_0}$ of α_{k_0} as a function of $X(f_{k_0})$, obtained by minimizing the mean squared error $\epsilon_{k_0} = \mathbb{E}\left[\left|\alpha_{k_0} - \hat{\alpha}_{k_0}\right|^2\right]$ is

$$\hat{\alpha}_{k_0} = \frac{W^*(0)\, v_{k_0}}{\sum_{k=1}^{K} |W(f_{k_0} - f_k)|^2 v_k} X(f_{k_0}) \qquad (12)$$

The proof of eq. (12) is given in the Appendix B. By ignoring non-overlapping overtones, the expression of the estimator of α_{hp} can be simplified as

$$\widehat{\alpha_{hp}} \triangleq \frac{W^*(0)\, v_{hp}^2}{\sum\limits_{\left|f_{h'p'} - f_{hp}\right| < \Delta_w} \left|W\left(f_{h'p'} - f_{hp}\right)\right|^2 v_{h'p'}^2} X(f_{hp}) \qquad (13)$$

where Δ_w is the width of the main lobe of W.

The amplitude estimation is illustrated in Fig. 3 on a synthetic signal, composed of two octave-related notes, *i.e.* with overlapping spectra. As expected, the amplitudes are perfectly estimated when there is no overlap (see odd-order peaks of note 1). When the overtones overlap, predominant amplitudes are well estimated (see note 1 for $f = 0.04$ or note 2 for $f = 0.4$), as well as amplitudes with the same order of magnitude (see amplitudes at $f = 0.2$). The estimation may be less accurate when an amplitude is much weaker than the other (see note 1 at $f = 0.36$).

3) Iterative algorithm for the estimation of note parameters: in the last two parts, we have successively seen how to estimate the spectral envelope models from known amplitudes and the amplitudes from known spectral envelope models. None of these quantities are actually known, and both must be estimated. We propose to use the two methods above in the Algorithm 1 to iteratively estimate both the spectral envelope and the amplitudes of the overtones. It alternates the estimation of the former and of the latter, with an additional thresholding of the weak amplitudes in order to avoid the amplitude estimates becoming much lower than the observed

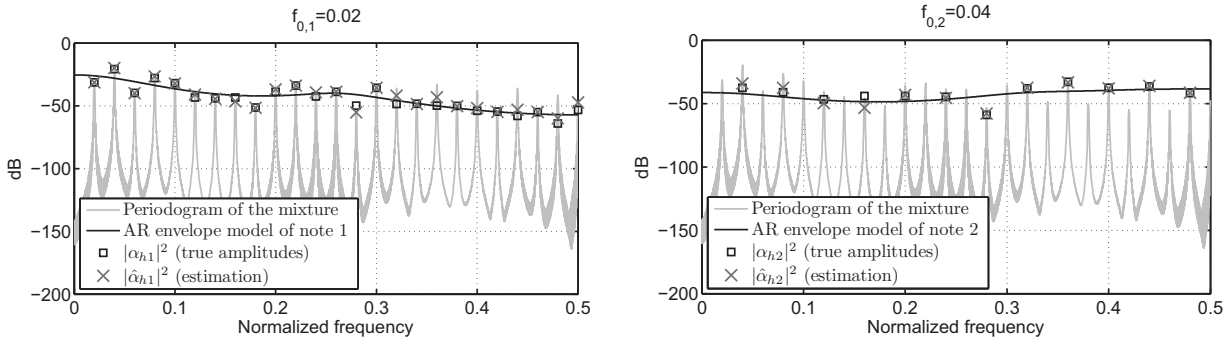

(a) Note 1. (b) Note 2.

Fig. 3. Estimation of amplitudes from overlapping spectra: example on a synthetic mixture of 2 notes. The F_0s $f_{0,1}$ and $f_{0,2}$ of the two notes are in an octave relation and the estimation is performed from a $N = 2048$-length observation. For each note, the true amplitudes (squares) have been generated from a spectral envelope AR model (black line) using eq. (3). The amplitudes are estimated (crosses) from the observation of the mixture (grey line) using the AR envelope information (eq. (12)).

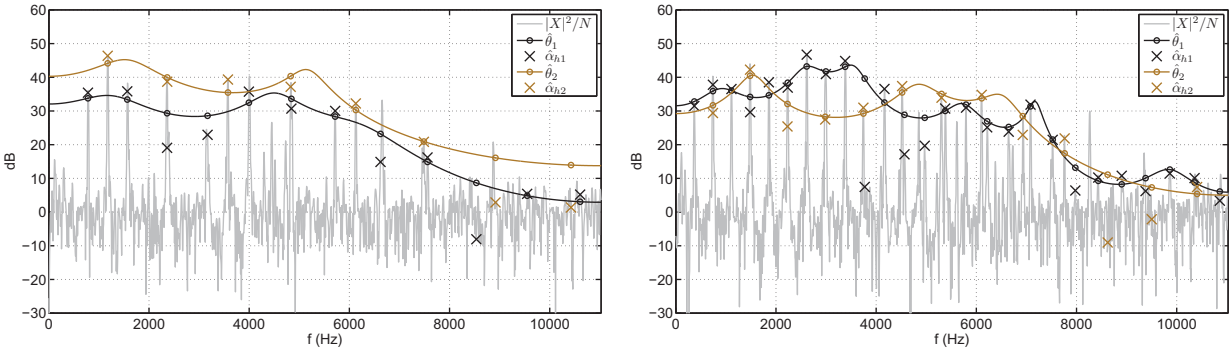

(a) True notes: G_5+D_6. Selected model: G_5+D_6. (b) True notes: $F\#_4$+$F\#_5$. Selected model: $F\#_4$+$F\#_5$.

Fig. 4. Iterative estimation of amplitudes $\widehat{\alpha_p}$ and spectral envelopes $\widehat{\theta_p}$: example of a fifth (Fig. 4(a)) and of an octave (Fig. 4(b)).

spectral coefficients[5]. For $p \in [\![1; P]\!]$ and $h \in [\![1; H_p]\!]$, the estimate $\widehat{\alpha_{hp}}$ of α_{hp} and $\widehat{\theta_p}$ of θ_p are respectively the values of $\left(\alpha_{hp}^{(i)}\right)$ and $\left(\theta_p^{(i)}\right)$ after a number of iterations. Note that the ML estimates of the envelopes are directly related to the MAP estimation of the multipitch contents while amplitudes are estimated in the least mean square sense, which is a different objective function. Consequently, the convergence of the iterative algorithm is not proved or guaranteed but it is observed after about $N_{\text{it}} \triangleq 20$ iterations.

The estimation is illustrated in Fig. 4 with two typical cases. While the energy is split between the overlapping overtones, non-overlapping overtones help estimating the spectral envelopes, resulting in a good estimation in the case of a fifth (Fig. 4(a)), the octave being a more difficult – but successful – case (Fig. 4(b)).

Algorithm 1 Iterative estimation of note parameters

Require: spectrum X, notes $\mathcal{C}_1, \ldots, \mathcal{C}_P$.
  Initialize the amplitudes $\alpha_{hp}^{(0)}$ of overtones to spectral values $X(f_{hp})$ for $p \in [\![1; P]\!]$ and $h \in [\![1; H_p]\!]$.
  for each iteration i **do**
    for each note p **do**
      Estimate $\theta_p^{(i)}$ from $\alpha_p^{(i-1)}$. {AR estimation}
    end for
    Jointly estimate all $\alpha_{hp}^{(i)}$ from $X(f_{hp})$ and $\theta_p^{(i)}$. {eq. (13)}
    Threshold $\alpha_{hp}^{(i)}$ at a minimum value.
  end for
Ensure: estimation of α_{hp} and θ_p for $p \in [\![1; P]\!]$ and $h \in [\![1; H_p]\!]$.

B. Estimation of noise parameters

We assume that the noise signal results from the circular filtering[6] of a white centered Gaussian noise with variance σ_b^2 by a Q_b-order FIR filter parameterized by its transfer function $B(z) = \sum_{k=0}^{Q_b} b_k z^{-k}$, with $b_0 = 1$.

As we assumed that in a bin close to the frequency of an overtone, the spectral coefficients of noise are negligible w.r.t. the spectral coefficients of notes, the information related to noise is mainly observable in the remaining bins, *i.e.* on the

[5]In our implementation, this threshold was set to the minimum observed amplitude $\min_f \frac{|X(f)|}{\sqrt{N}}$.

[6]The circularity assumption, which is commonly used and asymptotically valid, leads to a simplified ML solution.

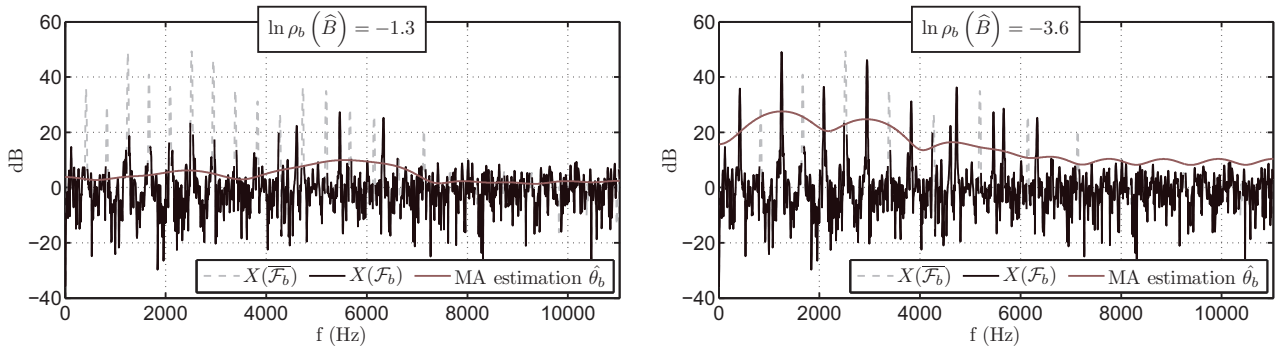

(a) True note: $E\flat_4$. Selected model: $E\flat_4$.

(b) True note: $E\flat_4$. Selected model: $E\flat_5$.

Fig. 5. Noise parameter estimation when selecting the true model (left) and the octave model (right). $\overline{\mathcal{F}_b}$ denotes the complement of $\mathcal{F}_b$.

frequency support defined by

$$\mathcal{F}_b \triangleq \left\{ \frac{k}{N_{\text{fft}}} \middle/ \forall p \in [\![1; P]\!], \forall h \in [\![1; H_p]\!], \left| \frac{k}{N_{\text{fft}}} - f_{hp} \right| > \frac{\Delta_w}{2} \right\} \tag{14}$$

where Δ_w denotes the width of the main lobe of w ($\Delta_w = \frac{4}{N}$ for a Hann window) and N_{fft} is the size of the DFT. The noise parameter estimation is then performed using this subset of bins, which gives satisfying results asymptotically (*i.e.* when $N \rightarrow +\infty$, the number of removed bins being much smaller than the total number of bins). As proved in the Appendix B, it results in the maximization, w.r.t. B, of the expression

$$L_b\left(\widehat{\sigma_b^2}, B\right) = c_b + \frac{|\mathcal{F}_b|}{2} \ln \rho_b\left(B\right) \tag{15}$$

where

$$c_b \triangleq -\frac{|\mathcal{F}_b|}{2} \ln 2\pi e - \frac{1}{2} \sum_{f \in \mathcal{F}_b} \ln \frac{|X_b\left(f\right)|^2}{N} \tag{16}$$

$$\rho_b\left(B\right) \triangleq \frac{\left(\prod_{f \in \mathcal{F}_b} \left| \frac{X_b(f)}{B(e^{2i\pi f})} \right|^2 \right)^{\frac{1}{|\mathcal{F}_b|}}}{\frac{1}{|\mathcal{F}_b|} \sum_{f \in \mathcal{F}_b} \left| \frac{X_b(f)}{B(e^{2i\pi f})} \right|^2} \tag{17}$$

$$\widehat{\sigma_b^2} = \frac{1}{|\mathcal{F}_b|} \sum_{f \in \mathcal{F}_b} \frac{1}{N} \left| \frac{X_b\left(f\right)}{B\left(e^{2i\pi f}\right)} \right|^2 \tag{18}$$

As for the spectral envelope AR parameters (eq. (11) and (8)), the ML estimation of noise parameters consists in the maximization of a spectral flatness $\rho_b\left(B\right)$. It can be achieved thanks to the MA estimation approach proposed in [23]. However, in order to speed up the computations, we use the Algorithm 3 described in the Appendix C, which is simpler and gives satisfying results for the current application.

The noise parameter estimation is illustrated in Fig. 5. When the true model is selected (Fig. 5(a)), most of primary lobes of the sinusoidal components are removed from $\mathcal{F}_b$ and the resulting spectral coefficients are well-fitted by an MA envelope. In the case a wrong model is selected and a lot of $\mathcal{F}_b$ bins are related to sinusoids (Fig. 5(b)), the MA model is not adapted to the observations which are not well-fitted. The resulting ρ_b value and likelihood will be low so that the wrong model will be rejected. As shown in Fig. 6, the spectral flatness criterion for the spectral envelope and noise estimations are complementary cues for rejecting wrong models and selecting the right one. The criterion on spectral envelopes (Fig. 6(a)) shows many high values but also large minima for sub-harmonic errors (*e.g.* notes $D\flat_3$, $A\flat_3$ and $D\flat_3$ in this example). Conversely, the criterion for the noise model (Fig. 6(b)) has a few high peaks located at the sub-harmonics F_0s of the true note – *i.e.* for every model all the sinusoids have been removed from the noise observation –, but is efficient for discriminating the other errors. Thus, if taken separately, these criteria are not good pitch estimators, but they efficiently combine into the detection function.

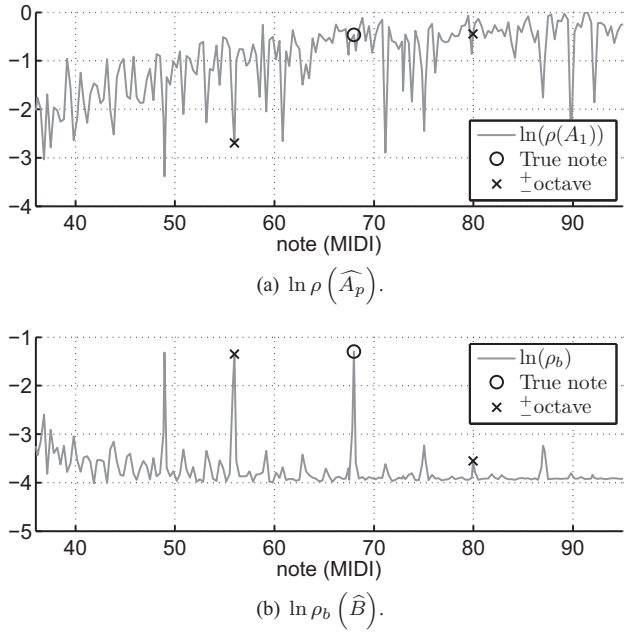

(a) $\ln \rho\left(\widehat{A_p}\right)$.

(b) $\ln \rho_b\left(\widehat{B}\right)$.

Fig. 6. $\ln \rho\left(\widehat{A_1}\right)$ (top) and $\ln \rho_b\left(\widehat{B}\right)$ (bottom) as a function of a single-note model (true note: $A\flat_4$).

IV. Experimental results

A. Implementation and tuning

The whole method is summarized by Algorithm 2. The MPE algorithm is implemented in Matlab and C. It is designed to analyze a 93ms frame and to estimate its polyphonic content. A preprocessing stage aims at reducing the spectral dynamics and consists in flattening the global spectral decrease by means of a median filtering of the spectrum. The values of the parameters of the MPE algorithm are given in Table I. The order Q_p of the spectral envelope filter $A_p(z)$ of a note p is set to $Q_p = H_p/2$ in order to be large enough to fit the data and small enough to obtain a smooth spectral envelope[7]. Parameters $\left(k_{\sigma_p^2}, E_{\sigma_p^2}\right)$, $\left(k_{\sigma_b^2}, E_{\sigma_b^2}\right)$, $\mu_{\text{env}}, \mu_{\text{b}}, \mu_{\text{pol}}$ and $(w_1, \ldots, w_4)$ have been estimated on a development database composed of about 380 mixtures from two different pianos, with polyphony levels from 1 to 6. They were jointly adjusted by optimizing the F-measure using a grid search, the bounds of the grid being set by hand.

Algorithm 2 Multipitch estimation.

Preprocessing
Computation of periodogram $\frac{1}{N}|X|^2$
Selection of N_c note candidates
for each possible note combination $\mathcal{C}$ **do**
  Iterative estimation of the amplitudes of overtones and of spectral envelope of notes (algorithm 1)
  For each note p, derivation of L_p (eq. (8)) and of prior $p\left(\widehat{\sigma_p^2}\right)$
  Noise parameter estimation (algorithm 3)
  Computation of L_b (eq. (15)) and of prior $p\left(\widehat{\sigma_b^2}\right)$
  Derivation of the detection function related to the note combination (eq. (7))
end for
Maximization of the detection function

Parameter	value	Parameter	value
$P_{\max}$	6	$\left(k_{\sigma_p^2}, E_{\sigma_p^2}\right)$	$(1, 10^{-4})$
N_c	9	$\left(k_{\sigma_b^2}, E_{\sigma_b^2}\right)$	$(2, 10^{-3})$
ν	0.38	μ_{env}	$-8.9{\cdot}10^{-3}$
f_s	22050Hz	μ_b	$-2.2{\cdot}10^{-4}$
N	2048	μ_{pol}	25
w	Hann	w_1	$8.1{\cdot}10^{-1}$
Q_p	$H_p/2$	w_2	$1.4{\cdot}10^{4}$
Q_b	20	w_3	$6.2{\cdot}10^{2}$
N_{it}	20	w_4	5.8

TABLE I
Parameters of the algorithm.

B. Evaluation

The proposed MPE algorithm has been tested on a database called MAPS[8] and composed of around 10000 piano sounds either recorded by using an upright Disklavier piano or generated by several virtual piano software products based on sampled sounds. The development set and the test set are disjointed. Two pianos are used in the former while the latter comprises sounds from other five pianos. In total, two upright pianos and five grand pianos were used. Recording conditions and tunings vary from one instrument to the other. Polyphony levels lie between 1 and 6 and notes are uniformly distributed between C_2 (65Hz) and B_6 (1976Hz). One part of the polyphonic mixtures is composed of randomly related pitches whereas the other part comprises usual chords from western music (major, minor, etc.). For each sound, a single 93ms-frame located 10ms after the onset time is extracted and analyzed. Two additional algorithms have been tested on the same database for comparison purposes. The first one is Tolonen's multipitch estimator [24] for which the implementation from the MIR Toolbox [25] was used. As the performance of the algorithm decreases when F_0s are greater than 500Hz (C_5), the system was additionally tested in restricted conditions – denoted Tolonen-500 – by selecting from the database the subset of sounds composed of notes between C_2 and B_4 only. The second one is Klapuri's system [26]. The code of the latter was provided by its author.

Our algorithm was implemented in Matlab and C. The computational cost on a recent PC is about $150\times$real time. It thus requires more computations than other algorithms like [26] but the proposed algorithm is computationally tractable, particularly when comparing it to the greedy and intractable joint-estimation approach where no note candidate selection is performed, as discussed in section II-B.

General results are presented in Fig. 7. Relevant items are defined as correct notes after rounding each F_0 to the nearest half-tone. Typical metrics are used: the recall is the ratio between the number of relevant items and of original items; the precision is the ratio between the number of relevant items and of detected items; and the F-measure is the harmonic mean between the precision and the recall. In this context, our system performs the best results for polyphony 1 and 2: 93% F-measure vs 85% (polyphony 1) and 91% (polyphony 2) for Klapuri's system. The trend then reverses between Klapuri's system and the proposed one, the F-measure being respectively 91% and 88% for polyphony 3, and 72% and 63% for polyphony 6. Moreover, the precision is high for all polyphony levels whereas the recall is decreasing when polyphony increases. Results from Tolonen's system are weaker for the overall tests, even with the restricted F_0-range configuration.

The ability of Klapuri's system and ours to detect polyphony levels (independently of the pitches) is presented in Tab. II. For polyphony levels from 1 to 5, both systems succeed in detecting the correct polyphony level more often than any other level. It should also be noted that the proposed system was tested in more difficult conditions since polyphony 0 (*i.e.* silence) may be detected, which is the case for a few sounds[9]. The proposed system tends to underestimate the polyphony

[7]Note that the number of degrees of freedom for the AR models is $H_p/4 = Q_p$ since half the poles are affected to the positive frequencies, the remaining poles being their complex conjugates.

[8]MAPS stands for MIDI Aligned Piano Sounds and is available on request.

[9]Note that silence detection could be improved by using an activity detector as a preprocessing stage, since silence is a very specific case of signal that can be detected using more specific methods than the proposed one.

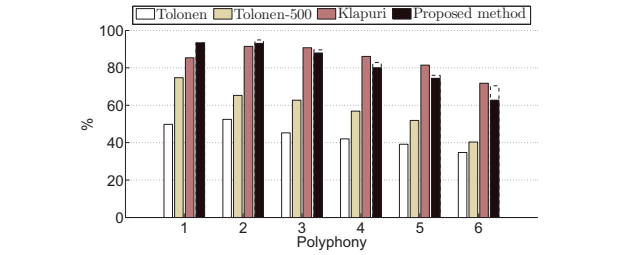

(a) F-measure

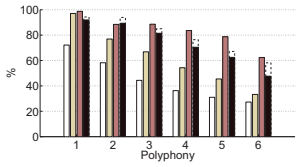

(b) Recall

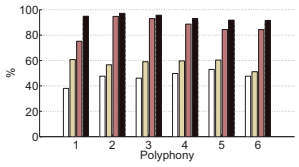

(c) Precision

Fig. 7. Multipitch estimation results with unknown polyphony: for each algorithm, the F-measure, the precision and the recall are rendered as a function of the true polyphony. For the proposed method, the performance on the training set is plotted using a dashed line.

	P					
P_{est}	1	2	3	4	5	6
0	7	2	4	6	9	0
1	**90**	14	8	8	8	12
2	3	**83**	18	12	8	19
3	0	1	**68**	27	16	25
4	0	0	2	**47**	27	32
5	0	0	0	1	**31**	11
6	0	0	0	0	1	**1**

Proposed method

	P					
P_{est}	1	2	3	4	5	6
0	0	0	0	0	0	0
1	**92**	21	5	2	1	3
2	1	**74**	18	7	4	7
3	1	3	**65**	24	10	12
4	0	1	9	**49**	24	27
5	1	0	1	14	**36**	23
6	5	1	1	4	25	**28**

Klapuri's method

TABLE II
POLYPHONY ESTIMATION: DETECTION RATE (%) W.R.T. TO THE NUMBER OF SOUNDS WITH TRUE POLYPHONY P, AS A FUNCTION OF THE TRUE POLYPHONY P AND OF THE ESTIMATED POLYPHONY P_{EST}.

level since the parameter tuning consists in optimizing the F-measure on the development set. This objective function could have been changed to take the polyphony level balance into account. This would result in reducing the polyphony underestimation trend. However, the overall F-measure would decrease. From a perceptive point of view, it has been shown that a missing note is generally less annoying than an added note when listening to a resynthesized transcription [27]. Thus, underestimating the polyphony may be preferred to overestimating it. Still, this trend turns out to be the main shortcoming of the proposed method, and should be fixed in the future for efficiently addressing sounds with polyphony higher than 5 notes.

The case of octave detection has been analyzed and results are presented in Fig. 8. Octave detection is one of the most difficult cases of multipitch estimation and the results are lower than the global results obtained for polyphony 2 (see Fig. 7). The proposed method reaches the highest results here, the F-measure being 81%, versus 77% for Klapuri's system and $77\%/66\%$ for Tolonen's (depending on the two F_0 ranges). This would suggest that the proposed models for spectral envelopes and overlapping spectra are significatively efficient.

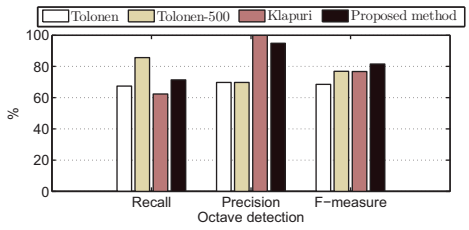

Fig. 8. Octave detection: results for the 97 octave sounds extracted from the database (45 sounds for the Tolonen-500 system).

Finally, we report additional results which are not depicted. First, if polyphony is known, the performance of our system reaches the second rank after Klapuri's system, with correct detection rates from 95% (polyphony 1) down to 55% (polyphony 6). Second, F-measure values are between 5% and 10% better for usual chords than for random-pitch chords. This has been observed with all the tested methods for polyphony levels higher than two. While the algorithms face more harmonically-related notes in usual chords – *i.e.* more spectral overlap –, it seems that simultaneous notes with widely-spread F_0s in random chords are a bigger difficulty. Third, the results obtained on the development set and the test set are comparable, with only a few $\%$ deviation. This shows that the parameter learning is not overfitting and that the system has good generalization abilities for other models of piano and different recording conditions. Fourth, the whole database includes sounds from seven different pianos. The results are comparable from one piano to the other, with about 3%-standard deviation. Results do not significantly depend on the upright/grand piano differences, on recording conditions or on whether a real piano or a software-based one is used, which suggests robustness for varying production conditions. Fifth, the proposed method uses a candidate selection stage which may fail, causing the subsequent multipitch estimation to fail. The performance of the candidate selection stage has thus been checked: satisfying results are obtained for low to medium polyphony levels, 99% of true notes being selected for polyphony 1 and 2, 94% for polyphony 3 and 86% for polyphony 4. Performance is slightly decreasing in polyphony 5 and 6, the scores being 78% and 71% respectively. Along the original-pitch dimension, these errors are mainly located below A_2 (110Hz), and above A_6 (1760Hz). Selecting a set of likely F_0s is a challenging issue for future works since a lot of MPE algorithms would benefit from it. Indeed, they often consist in optimizing an objective criterion that presents a lot of local optima along the pitch dimension. Some basic solutions like exhaustive or grid search [27] or more elaborated ones like Markov chain Monte Carlo methods [28] have already been proposed for searching the global optimum, and an efficient candidate selection stage would be very helpful to reach the result more quickly.

C. Integration in an automatic transcription system

The proposed MPE method has been introduced for a single-frame analysis. Beyond the multipitch issue, the automatic transcription task is addressed by integrating this method in

a system that performs pitch tracking over time and outputs not only pitches but note events. In the current case, the MPE method can be used in the transcription framework proposed in [19]. The automatic transcription of piano music is thus obtained by:

1) using an onset detector to localize new events;
2) detecting a set of pitch candidates after each onset;
3) between two consecutive onsets, decoding the multipitch contents thanks to a frame-based Hidden Markov Model (HMM), the states being the possible combinations of simultaneous pitches, the likelihood being given by the proposed MPE method (eq. (7));
4) postprocessing the decoded pitches to detect repetitions or continuations of pitches when an onset occurs.

This transcription system was presented at the *Multiple Fundamental Frequency Estimation & Tracking* task of the MIREX[10] contest in 2008 (system called EBD2). For the piano-only test set, the evaluation is performed in two different ways, depending on how a correct note is defined. In the first, more constraining evaluation, where a correct note implies a correct onset (up to a 50-ms deviation), a correct offset (up to a 20%-duration or 50-ms deviation) and a correct pitch (up to a quartertone deviation), the proposed system reaches the 3^{rd} rank with a 33.1% F-measure, out of 13 systems scoring from 6.1% to 36.8%. For that evaluation, the duration estimation is evaluated by the average overlap between correct estimations and reference notes: the proposed method reaches the 7^{th} position with 80%, the values being in a narrow interval ($[77.4\%; 84\%]$). In the second evaluation, where correctness only requires the onset and pitch criteria above, the proposed system reaches the 7^{th} rank with a 56.9% F-measure, out of 13 systems scoring from 24.5% to 75.7%. The best rank is obtained for average overlap with a 61% value, the lowest value being 40.1%.

Thus, when the proposed MPE method is integrated in a full transcription system for piano music, state-of-the art results are reached in terms of global performance. In addition, good average overlap scores show that the proposed method is suitable for pitch tracking over time and note extraction.

Note that one can wonder whether the reported performance is significant when comparing MPE and ATM quantitative results (with roughly 20% F-measure deviation, no matter the method is). It actually depends not only on the database, but on the different testing protocols. While isolated frames composed of one chord are used in the proposed MPE evaluation, ATM evaluation implies other difficulties like having asynchronous notes overlaping in time; detecting onsets; estimating the end of damping notes; dealing with reverberation queues and so on. More, polyphony levels in musical recordings are often high and are not uniformly distributed at all, with a 4.5 average polyphony and a 3.1 standard deviation reported for a number of classical music pieces [29, p.114]. Hence, F-measure for MPE and ATM should be evaluated separately.

[10]http://www.music-ir.org/mirex/2008/

V. Conclusions

In this paper, a sound model and a method were proposed for the multipitch estimation problem in the case of piano sounds. The approach was based on an adaptive scheme, including the adaptive matching of the inharmonic distribution of the frequencies of the overtones and an autoregressive spectral envelope. We have shown the advantage of using a moving average model for the residual noise. The estimation of the parameters was performed by taking the possible overlap between the spectra of the notes into account. Finally, a weighted likelihood criterion was introduced to determine the estimated mixture of notes in the analyzed frame among a set of possible mixtures.

The performance of the method was measured on a large database of piano sounds and compared to the state-of-the-art. The proposed method provides satisfying results when polyphony is unknown, and reaches particularly good scores for mixtures of harmonically-related notes.

This approach has been successively integrated in a full transcription system [19] and alternative investigations [30] are ongoing. Future works may deal with the use of different spectral envelopes in a similar modeling and estimation framework. Indeed, while the proposed AR envelope model may be too generic to characterize the spectral envelope smoothness of musical instruments, this information can be easily replaced, as the variance of the amplitude Gaussian models, by any parametric envelope model. The method may also be extended to other instruments or to voice. Keeping the proposed sound model, it could be useful to investigate other estimation strategies in order to reduce the computational cost. This may be achieved by means of a hybrid approach with an iterative selection of the notes and a joint estimation criterion. Furthermore, the log-likelihood function related to the model for spectral envelopes is not very selective. Hence, some investigations on more appropriate functions could improve the performance of decision step. Finally, one could conduct others investigations about how to prune the set of all possible note combinations more efficiently. They may include iterative candidate selection or MCMC methods.

Appendix A
Discussion on the detection function

According to eq. (5), the multipitch estimation is theoretically obtained by maximizing, w.r.t. $\mathcal{C}$, the likelihood $p(x|\mathcal{C})$, which is expressed as a function of the parameters of the proposed model as:

$$p(x|\mathcal{C}) = \iint p(x, \alpha, \theta_b|\mathcal{C})\, d\alpha d\theta_b \tag{19}$$

$$= \iint p(x|\alpha, \theta_b, \mathcal{C})\, p(\alpha|\mathcal{C})\, p(\theta_b)\, d\alpha d\theta_b \tag{20}$$

$$= \iiint p(x|\alpha, \theta_b, \mathcal{C})\, p(\alpha|\theta, \mathcal{C})\, p(\theta)\, p(\theta_b)\, d\theta d\alpha d\theta_b \tag{21}$$

Since the derivation of eq. (21) is not achievable, a simpler detection function must be built. The proposed function can be interpreted as a weighted log-likelihood and is obtained from the integrand of eq. (21) by:

- considering the estimated parameters $\widehat{\theta_p}, \widehat{\alpha_p}, \widehat{\theta_b}$;
- taking the logarithm, thus turning the product into a sum;
- normalizing the sums over the number of notes by P, and log-likelihoods L_p and L_b by the size H_p and $|\mathcal{F}_b|$ of their respective random variable;
- introducing penalization terms using the $\mathcal{C}$-dependent sizes P, H_p, $|\mathcal{F}_b|$ with coefficients $\mu_{\text{pol}}, \mu_{\text{env}}, \mu_b$;
- weighting each quantity by the coefficients $w_1, \ldots, w_4$.

The normalization operation aims at obtaining data that can be compared from one mixture $\mathcal{C}$ to another, since the dimension of the random variables depends on the mixture. The penalization is inspired by model selection approaches [31], in which an integral like (21) is replaced by a likelihood penalized by a linear function of the model order. The weighting by $w_1, \ldots, w_4$ is finally used to find an optimal balance between the various probability density functions. Note that in eq. (7), priors related to filters A_p and B have been removed since they are non-informative, and thus can be considered as constant terms.

Appendix B
Proofs

Proof of eq. (8)*:* Using eq. (3), the log-likelihood to be maximized is

$$L_p\left(\sigma_p^2, A_p\right) = -\frac{H_p}{2}\ln\left(2\pi\sigma_p^2\right) + \frac{1}{2}\sum_{h=1}^{H_p}\ln\left|A_p\left(e^{2i\pi f_{hp}}\right)\right|^2 - \frac{1}{2\sigma_p^2}\sum_{h=1}^{H_p}\left|\alpha_{hp}\right|^2\left|A_p\left(e^{2i\pi f_{hp}}\right)\right|^2 \tag{22}$$

The maximization w.r.t. σ_p^2 leads to the estimate $\widehat{\sigma_p^2}$ given by eq. (11), which is used in eq. (22) to obtain eq. (8). ■

Proof of eq. (12) *and* (13)*:* The expected linear estimator is expressed as $\hat{\alpha}_{k_0} \triangleq \eta X\left(f_{k_0}\right)$ with $\eta \in \mathbb{C}$. The mean squared error is then $\epsilon_{k_0}(\eta) = \mathbb{E}\left[\left|\alpha_{k_0} - \eta X\left(f_{k_0}\right)\right|^2\right]$. The optimal value $\hat{\eta}$ is such that $\frac{d\epsilon_{k_0}}{d\eta}(\hat{\eta}) = 0$, which is equivalent to the decorrelation between the error $(\alpha_{k_0} - \hat{\eta}X\left(f_{k_0}\right))$ and the data $X\left(f_{k_0}\right)$:

$$\begin{aligned} 0 &= \mathbb{E}\left[\left(\alpha_{k_0}^* - \hat{\eta}^* X^*\left(f_{k_0}\right)\right) X\left(f_{k_0}\right)\right] \\ &= \mathbb{E}\left[\alpha_{k_0}^* X\left(f_{k_0}\right)\right] - \hat{\eta}^*\mathbb{E}\left[\left|X\left(f_{k_0}\right)\right|^2\right] \end{aligned} \tag{23}$$

which leads to $\hat{\eta} = \frac{\mathbb{E}\left[\alpha_{k_0} X^*\left(f_{k_0}\right)\right]}{\mathbb{E}\left[\left|X\left(f_{k_0}\right)\right|^2\right]}$, where

$$\begin{aligned} \mathbb{E}\left[\alpha_{k_0} X^*\left(f_{k_0}\right)\right] &= \sum_{k=1}^{K} W^*\left(f_{k_0} - f_k\right)\mathbb{E}\left[\alpha_{k_0}\alpha_k^*\right] \\ &= W^*(0)\, v_{k_0} \end{aligned} \tag{24}$$

$$\begin{aligned} \text{and } \mathbb{E}\left[\left|X\left(f_{k_0}\right)\right|^2\right] &= \sum_{k=1}^{K}\sum_{l=1}^{K} W\left(f_{k_0} - f_k\right) \\ &\quad \times W^*\left(f_{k_0} - f_l\right)\mathbb{E}\left[\alpha_k\alpha_l^*\right] \\ &= \sum_{k=1}^{K}\left|W\left(f_{k_0} - f_k\right)\right|^2 v_k \end{aligned} \tag{25}$$

The related error is

$$\begin{aligned} \epsilon_{k_0}(\hat{\eta}) &= \mathbb{E}\left[\left|\alpha_{k_0}\right|^2\right] + \left|\hat{\eta}\right|^2\mathbb{E}\left[\left|X\left(f_{k_0}\right)\right|^2\right] \\ &\quad - \hat{\eta}\mathbb{E}\left[\alpha_{k_0}^* X\left(f_{k_0}\right)\right] - \hat{\eta}^*\mathbb{E}\left[\alpha_{k_0} X^*\left(f_{k_0}\right)\right] \\ &= \mathbb{E}\left[\left|\alpha_{k_0}\right|^2\right] - \frac{\left|\mathbb{E}\left[\alpha_{k_0} X^*\left(f_{k_0}\right)\right]\right|^2}{\mathbb{E}\left[\left|X\left(f_{k_0}\right)\right|^2\right]} \\ &= \left(1 - \frac{\left|W(0)\right|^2 v_{k_0}}{\sum_{k=1}^{K}\left|W\left(f_{k_0} - f_k\right)\right|^2 v_k}\right) v_{k_0} \end{aligned} \tag{26}$$

Note that this approach has some connections with the Wiener filtering technique [32] widely used in the field of audio source separation. In both cases, the aim is to estimate a hidden variable thanks to its second-order statistics and to the observations. In the considered short-term analysis context, the proposed approach explicitly models the overlap phenomenon and takes the resulting frequency leakage into account. As expected, the estimation error is all the larger as the estimated overtone is masked by another component, which justifies the approximation of eq. (12) in eq. (13). ■

Proof of eq. (15)*:* In a matrix form, the filtering process is expressed as $\underline{x_b} = B_{\text{circ}}\underline{w_b}$, where $\underline{x_b}$ is an N-length frame of the noise process x_b, $\underline{w_b} \sim \mathcal{N}\left(0, \boldsymbol{I}_N\sigma_b^2\right)$ and B_{circ} is the $N \times N$ circulant matrix with first column $(b_0, \ldots, b_{Q_b}, 0, \ldots, 0)$. Thus, we have $\underline{x_b} \sim \mathcal{N}\left(0, B_{\text{circ}}\left(B_{\text{circ}}\right)^\dagger\sigma_b^2\right)$. B_{circ} is circulant, so $\det\left(B_{\text{circ}}B_{\text{circ}}^\dagger\right) = \prod_{k=0}^{N-1}\left|B\left(e^{2i\pi\frac{k}{N}}\right)\right|^2$. The likelihood of $\underline{x_b}$ is then

$$p\left(\underline{x_b}\right) = \frac{e^{-\frac{1}{2}\underline{x_b}^\dagger\left(B_{\text{circ}}B_{\text{circ}}^\dagger\sigma_b^2\right)^{-1}\underline{x_b}}}{\sqrt{(2\pi)^N\det\left(B_{\text{circ}}B_{\text{circ}}^\dagger\sigma_b^2\right)}} \tag{27}$$

$$= \frac{e^{-\frac{\left\|B_{\text{circ}}^{-1}\underline{x_b}\right\|^2}{2\sigma_b^2}}}{\sqrt{\left(2\pi\sigma_b^2\right)^N\prod_{k=0}^{N-1}\left|B\left(e^{2i\pi\frac{k}{n}}\right)\right|^2}} \tag{28}$$

In the spectral domain, the log-likelihood of the noise signal

is thus obtained using the Parseval identity:

$$L_b\left(\sigma_b^2, B\right) = -\frac{N}{2}\ln 2\pi\sigma_b^2 - \frac{1}{2}\sum_{k=0}^{N-1}\ln\left|B\left(e^{2i\pi\frac{k}{N}}\right)\right|^2 - \frac{1}{2\sigma_b^2}\sum_{k=0}^{N-1}\frac{1}{N}\left|\frac{X_b\left(\frac{k}{N}\right)}{B\left(e^{2i\pi\frac{k}{N}}\right)}\right|^2 \quad (29)$$

When only considering the observations on the spectral bins $\mathcal{F}_b$, the log-likelihood (29) becomes

$$L_b\left(\sigma_b^2, B\right) = -\frac{|\mathcal{F}_b|}{2}\ln 2\pi\sigma_b^2 - \frac{1}{2}\sum_{f\in\mathcal{F}_b}\ln\left|B\left(e^{2i\pi f}\right)\right|^2 - \frac{1}{2\sigma_b^2}\sum_{f\in\mathcal{F}_b}\frac{1}{N}\left|\frac{X_b\left(f\right)}{B\left(e^{2i\pi f}\right)}\right|^2 \quad (30)$$

The maximization w.r.t. σ_b^2 leads to the estimate given by eq. (18), which is used in eq. (30) to obtain eq. (15).

■

Appendix C
MA estimation

The algorithm is based on the decomposition $\underline{r_b} = \mathbf{B}\underline{b}$ of the first $(Q_b + 1)$ terms $\underline{r_b}$ of the autocorrelation function of the process, as a product of the coefficient vector $\underline{b} \triangleq (b_0, \ldots, b_{Q_b})^{\mathrm{T}}$ by the matrix

$$\mathbf{B} \triangleq \begin{pmatrix} b_0 & b_1 & \ldots & b_{Q_b} \\ 0 & b_0 & \ldots & b_{Q_b-1} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \ldots & 0 & b_0 \end{pmatrix} \quad (31)$$

The estimate $\widehat{\underline{b}}$ of $\underline{b}$ is obtained using Algorithm 3. As $\widehat{\mathbf{B}}$ is upper triangular, the estimation of $\widehat{\underline{b}}$ in $\widehat{r_b} = \widehat{\mathbf{B}}\widehat{\underline{b}}$ is fast performed by back substitution, instead of inverting a full matrix. The algorithm convergence was observed after about 20 iterations.

Algorithm 3 Iterative estimation of noise parameters

estimate the autocorrelation vector $\underline{r_b}$ by the empiric correlation coefficients $\widehat{r_b}$ obtained from the spectral observations $\{X(f)\}_{f\in\mathcal{F}_b}$;
initialize $\widehat{\underline{b}} \leftarrow (1, 0, \ldots, 0)^{\mathrm{T}}$;
for each iteration **do**
  update the estimate $\widehat{\mathbf{B}}$ of $\mathbf{B}$ from $\widehat{\underline{b}}$;
  re-estimate $\widehat{\underline{b}}$ by solving $\widehat{r_b} = \widehat{\mathbf{B}}\widehat{\underline{b}}$;
  normalize $\widehat{\underline{b}}$ by its first coefficient;
end for

References

[1] L. Rabiner, "On the use of autocorrelation analysis for pitch detection," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 25, no. 1, pp. 24–33, 1977.

[2] A. de Cheveigné and H. Kawahara, "YIN, a fundamental frequency estimator for speech and music," *J. Acous. Soc. Amer.*, vol. 111, no. 4, pp. 1917–1930, 2002.

[3] M. R. Schroeder, "Period histogram and product spectrum: New methods for fundamental-frequency measurement," *J. Acous. Soc. Amer.*, vol. 43, no. 4, pp. 829–834, 1968.

[4] J. C. Brown, "Musical fundamental frequency tracking using a pattern recognition method," *J. Acous. Soc. Amer.*, vol. 92, no. 3, pp. 1394–1402, 1992.

[5] B. Doval and X. Rodet, "Fundamental frequency estimation and tracking using maximum likelihood harmonic matching and HMMs," in *Proc. Int. Conf. Audio Speech and Sig. Proces. (ICASSP)*, vol. 1, Apr. 1993, pp. 221–224.

[6] G. Peeters, "Music pitch representation by periodicity measures based on combined temporal and spectral representations," in *Proc. Int. Conf. Audio Speech and Sig. Proces. (ICASSP)*, Toulouse, France, May 2006, pp. 53–56.

[7] V. Emiya, B. David, and R. Badeau, "A parametric method for pitch estimation of piano tones," in *Proc. Int. Conf. Audio Speech and Sig. Proces. (ICASSP)*, Honolulu, HI, USA, Apr. 2007, pp. 249–252.

[8] A. Klapuri, "Multiple fundamental frequency estimation based on harmonicity and spectral smoothness," *IEEE Trans. Speech Audio Process.*, vol. 11, no. 6, pp. 804–816, 2003.

[9] C. Yeh, A. Robel, and X. Rodet, "Multiple fundamental frequency estimation of polyphonic music signals," in *Proc. Int. Conf. Audio Speech and Sig. Proces. (ICASSP)*, Philadelphia, PA, USA, Mar. 2005, pp. 225–228.

[10] H. Kameoka, T. Nishimoto, and S. Sagayama, "A Multi-pitch Analyzer Based on Harmonic Temporal Structured Clustering," *IEEE Trans. Audio, Speech and Language Processing*, vol. 15, no. 3, pp. 982–994, 2007.

[11] M. Christensen and A. Jakobsson, *Multi-Pitch Estimation*, ser. Synthesis lectures on speech and audio processing, B. Juang, Ed. Morgan and Claypool Publishers, 2009.

[12] C. Raphael, "Automatic transcription of piano music," in *Proc. Int. Conf. Music Information Retrieval (ISMIR)*, Paris, France, Oct. 2002, pp. 15–19.

[13] M. Marolt, "A connectionist approach to automatic transcription of polyphonic piano music," *IEEE Trans. Multimedia*, vol. 6, no. 3, pp. 439–449, 2004.

[14] J. Bello, L. Daudet, and M. Sandler, "Automatic piano transcription using frequency and time-domain information," *IEEE Trans. Audio, Speech and Language Processing*, vol. 14, no. 6, pp. 2242–2251, Nov. 2006.

[15] A. Klapuri and M. Davy, *Signal Processing Methods for Music Transcription*. Springer, 2006.

[16] G. Poliner and D. Ellis, "A discriminative model for polyphonic piano transcription," *EURASIP Journal on Advances in Signal Processing*, vol. 8, pp. 1–9, 2007.

[17] K. Robinson and R. Patterson, "The duration required to identify the instrument, the octave, or the pitch chroma of a musical note," *Music Perception*, vol. 13, no. 1, pp. 1–15, 1995.

[18] L. I. Ortiz-Berenguer, F. J. Casajús-Quirós, and M. Torres-Guijarro, "Non-linear effects modeling for

polyphonic piano transcription," in *Proc. Int. Conf. Digital Audio Effects (DAFx)*, London, UK, Sep. 2003.

[19] V. Emiya, R. Badeau, and B. David, "Automatic transcription of piano music based on HMM tracking of jointly-estimated pitches," in *Proc. Eur. Conf. Sig. Proces. (EUSIPCO)*, Lausanne, Switzerland, Aug. 2008.

[20] N. H. Fletcher and T. D. Rossing, *The Physics of Musical Instruments*. Springer, 1998.

[21] V. Emiya, R. Badeau, and B. David, "Multipitch estimation of inharmonic sounds in colored noise," in *Proc. Int. Conf. Digital Audio Effects (DAFx)*, Bordeaux, France, Sep. 2007, pp. 93–98.

[22] A. El-Jaroudi and J. Makhoul, "Discrete all-pole modeling," *IEEE Trans. Signal Process.*, vol. 39, no. 2, pp. 411–423, 1991.

[23] R. Badeau and B. David, "Weighted maximum likelihood autoregressive and moving average spectrum modeling," in *Proc. Int. Conf. Audio Speech and Sig. Proces. (ICASSP)*, Las Vegas, NV, USA, Mar.-Apr. 2008, pp. 3761–3764.

[24] T. Tolonen and M. Karjalainen, "A computationally efficient multipitch analysis model," *IEEE Trans. Speech Audio Process.*, vol. 8, no. 6, pp. 708–716, 2000.

[25] O. Lartillot and P. Toiviainen, "A Matlab toolbox for musical feature extraction from audio," in *Proc. Int. Conf. Digital Audio Effects (DAFx)*, Bordeaux, France, Sep. 2007, pp. 237–244.

[26] A. Klapuri, "Multiple fundamental frequency estimation by summing harmonic amplitudes," in *Proc. Int. Conf. Music Information Retrieval (ISMIR)*, Victoria, Canada, Oct. 2006, pp. 216–221.

[27] A. Cemgil, H. Kappen, and D. Barber, "A generative model for music transcription," *IEEE Trans. Audio, Speech and Language Processing*, vol. 14, no. 2, pp. 679–694, 2006.

[28] P. Walmsley, S. Godsill, and P. Rayner, "Polyphonic pitch tracking using joint Bayesian estimation of multiple frame parameters," in *Proc. IEEE Work. Appli. Sig. Proces. Audio and Acous. (WASPAA)*, New Paltz, NY, USA, Oct. 1999, pp. 119–122.

[29] V. Emiya, "Transcription automatique de la musique de piano," Ph.D. dissertation, École Nationale Supérieure des Télécommunications, France, 2008.

[30] R. Badeau, V. Emiya, and B. David, "Expectation-maximization algorithm for multi-pitch estimation and separation of overlapping harmonic spectra," in *Proc. Int. Conf. Audio Speech and Sig. Proces. (ICASSP)*, Taipei, Taiwan, Apr. 2009.

[31] P. Stoica and Y. Selen, "Model-order selection: a review of information criterion rules," *IEEE Signal Process. Mag.*, vol. 21, no. 4, pp. 36–47, 2004.

[32] N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. The MIT Press, 1964.

EXPECTATION-MAXIMIZATION ALGORITHM FOR MULTI-PITCH ESTIMATION AND SEPARATION OF OVERLAPPING HARMONIC SPECTRA

Roland BADEAU, Valentin EMIYA, Bertrand DAVID

CNRS LTCI, TELECOM ParisTech (ENST)
roland.badeau@telecom-paristech.fr

ABSTRACT

This paper addresses the problem of multi-pitch estimation, which consists in estimating the fundamental frequencies of multiple harmonic sources, with possibly overlapping partials, from their mixture. The proposed approach is based on the expectation-maximization algorithm, which aims at maximizing the likelihood of the observed spectrum, by performing successive single-pitch and spectral envelope estimations. This algorithm is illustrated in the context of musical chord identification.

Index Terms— Spectral analysis, Maximum likelihood estimation, Algorithms, Harmonic analysis, Envelope detection.

1. INTRODUCTION

Estimating the fundamental frequency (or pitch) of a harmonic signal is a prominent task in audio signal processing. Its main difficulty lies in the intrinsic ambiguity of the harmonic model, which typically leads to octave errors. However this problem can be circumvented by taking the smoothness of the spectral envelope into account [1]. In presence of multiple sources, the multi-pitch estimation task is even more difficult, because of the spectral overlap between the harmonic components (see [2, 3] for a review). There are two categories of methods for performing this task:

- iterative approaches, which recursively estimate the dominant pitch, and remove its harmonics from the mixture;
- joint approaches, which simultaneously estimate all the fundamental frequencies by optimizing a joint criterion.

Methods in the first category are fast, but they tend to accumulate errors at each iteration, because of the spectral overlap between partials belonging to different sources. Methods in the second category do not accumulate errors, but they are more computationally demanding, since they require to explore a space whose dimension is the number of harmonic components in the mixture (which can be circumvented by reducing this search space beforehand, like in [4]).

In this paper, we propose a new approach for addressing the multi-pitch estimation task, which combines the advantages of both categories: we optimize a joint criterion by means of a recursive algorithm which performs successive single-pitch estimations. This is achieved by applying the expectation-maximization (EM) approach [5] to an appropriate spectrum model, which provides a proper statistical framework to our method. The paper is organized as follows: the spectral mixture model is introduced in Section 2, and the EM algorithm is described in Section 3. Numerical simulations are presented in Section 4, and the conclusions are summarized in Section 5. The following notations will be used throughout the paper:

- we use normal symbols for scalars, underlined symbols for vectors, and doubly underlined symbols for matrices;
- for any integer variable N, $\overline{N}$ denotes the set $\{0 \dots N-1\}$.

The research leading to this paper was supported by the French GIP ANR under contract ANR-06-JCJC-0027-01, Décompositions en Éléments Sonores et Applications Musicales - DESAM, and by the European Commission under contract FP6-027026, Knowledge Space of semantic inference for automatic annotation and retrieval of multimedia content - KSPACE.

2. SPECTRAL MIXTURE MODEL

We denote $\{Y_i\}_{i\in\overline{I}}$ the samples of the outcome power spectrum (where $\overline{I} = \{0 \dots I-1\}$ is the set of all frequency bins), and $\underline{Y}$ the random vector $[Y_0 \dots Y_{I-1}]^T$. This power spectrum is modeled as the squared magnitude of a sum of J complex spectra $\underline{X}_j = [X_{0,j} \dots X_{I-1,j}]^T$, whose presence or absence at frequency i is indicated by a Boolean variable $B_{i,j}$, plus a complex spectrum $\underline{N} = [N_0 \dots N_{I-1}]^T$ corresponding to an additive noise present at all frequencies, so that $Y_i = |N_i + \sum_{j\in\overline{J}} B_{i,j} X_{i,j}|^2$. This model should be understood as follows: since we aim at modeling a mixture of harmonic spectra, $\underline{X}_j$ is the j^{th} complete spectrum, and the Booleans $B_{i,j} \in \mathcal{B} = \{0,1\}$ act as a selector for the harmonic frequencies (vector $\underline{B}_j = [B_{0,j} \dots B_{I-1,j}]^T$ is typically shaped as a comb).

2.1. Spectral envelope model

As usually assumed in the literature, for all $j \in \overline{J}$, we suppose that $\underline{X}_j$ is a complex Gaussian random vector, of zero mean, and covariance matrix $\mathrm{diag}(\underline{s}_j)$, where $\underline{s}_j = [s_{0,j} \dots s_{I-1,j}]^T$ (the outcome values at all frequency bins are assumed independent). In the same way, $\underline{N}$ is a complex Gaussian random vector, of zero mean, and covariance matrix $\mathrm{diag}(\underline{s}_J)$. All those spectral components are assumed mutually independent. Moreover, $\forall j \in \overline{J}\cup\{J\}$ ($j = J$ referring to the noise component), the smooth spectral envelope $\underline{s}_j \in \mathbb{R}_+^I$ is parameterized by a moving average (MA) model of order K:

$$s_{i,j} = \sigma_j^2 \left| \sum_{k=0}^{K} \alpha_{k,j} e^{-2i\pi k \frac{i}{I}} \right|^2 \tag{1}$$

where $\forall j \in \overline{J} \cup \{J\}$, $\alpha_{0,j} = 1$.

2.2. Harmonicity model

All Booleans $B_{i,j}$ are assumed independent, and we denote $\pi_{i,j} \in]0,1[$ the probability that $B_{i,j} = 1$. The harmonicity of the spectral components is expressed as follows: we consider L subsets $\mathcal{H}_l$ of $\overline{I}$, which locate the harmonic frequency bins for L pitch candidates. We further assume that for all the spectral components, *i.e.* $\forall j \in \overline{J}$, there is a unique pitch identified by an index $l_j \in \overline{L}$, and two

constants $\dot{\pi}_j$ and $\mathring{\pi}_j$ in $]0,1[$, such that $\forall i \in \mathcal{H}_{l_j}, \pi_{i,j} = \dot{\pi}_j$, and $\forall i \notin \mathcal{H}_l, \pi_{i,j} = \mathring{\pi}_j$ (all harmonic frequency bins in $\mathcal{H}_{l_j}$ have the same high probability of presence $\dot{\pi}_j$, and the other ones have the same low probability of presence $\mathring{\pi}_j$, as illustrated in Fig. 1). For all $i \in \overline{I}$, let $\underline{B}^i = [B_{i,0} \dots B_{i,J-1}]^T$. Since all Booleans $B_{i,j}$ are independent, for any vector $\underline{b} \in \mathcal{B}^J$, the probability that $\underline{B}^i = \underline{b}$ is

$$\pi_{i,\underline{b}} = \prod_{j \in \overline{J}} (\pi_{i,j})^{b_j} (1 - \pi_{i,j})^{1-b_j}. \tag{2}$$

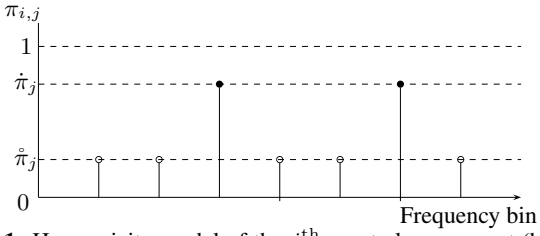

Fig. 1. Harmonicity model of the j^{th} spectral component (here the fundamental frequency equals 3 Fourier bins; • stands for frequencies in $\mathcal{H}_{l_j}$, ∘ stands for frequencies not in $\mathcal{H}_{l_j}$).

2.3. Resulting mixture model

For any vector $\underline{b} \in \mathcal{B}^J$, the mixture $N_i + \sum_{j \in \overline{J}} b_j X_{i,j}$ is a complex Gaussian random variable, of zero mean, and variance

$$s_{i,\underline{b}} = s_{i,J} + \sum_{j \in \overline{J}} b_j s_{i,j}, \tag{3}$$

where $s_{i,J}$ refers to the noise component (see section 2.1). The conditional law of Y_i given that $\underline{B}^i = \underline{b}$ is a chi-square distribution with two degrees of freedom, of probability density function

$$p(Y_i | \underline{B}^i = \underline{b}) = \frac{1}{s_{i,\underline{b}}} e^{-\frac{Y_i}{s_{i,\underline{b}}}}. \tag{4}$$

3. EXPECTATION-MAXIMIZATION ALGORITHM

In this section, we propose an EM algorithm to estimate the spectral mixture model. Here the observations are the power spectrum samples Y_i, and the hidden states are the Booleans $B_{i,j}$, arranged in the $I \times J$ matrix $\underline{\underline{B}}$. Their joint probability is written in the form

$$p(\underline{Y}, \underline{\underline{B}}) = \prod_{i \in \overline{I}} \sum_{\underline{b} \in \mathcal{B}^J} \mathbf{1}_{\{\underline{B}^i = \underline{b}\}} \pi_{i,\underline{b}}\, p(Y_i | \underline{B}_i = \underline{b}),$$

where $\pi_{i,\underline{b}}$ was defined in equation (2). Substituting equation (4) into the above equation leads to the following expression of the joint log-likelihood $L(\underline{Y}, \underline{\underline{B}}) = \ln(p(\underline{Y}, \underline{\underline{B}}))$:

$$L(\underline{Y}, \underline{\underline{B}}) = \sum_{i \in \overline{I}} \sum_{\underline{b} \in \mathcal{B}^J} \mathbf{1}_{\{\underline{B}^i = \underline{b}\}} \left(\ln \left(\frac{\pi_{i,\underline{b}}}{s_{i,\underline{b}}} \right) - \frac{Y_i}{s_{i,\underline{b}}} \right)$$

To estimate the set of prior probabilities $\pi_{i,j}$, arranged in the $I \times J$ matrix $\underline{\underline{\pi}}$, and the set of envelope coefficients $s_{i,j}$, arranged in the $I \times (J+1)$ matrix $\underline{\underline{s}}$, the EM algorithm generates a sequence $(\underline{\underline{\pi}}^n, \underline{\underline{s}}^n)_{n \geq 0}$ by recursively maximizing the conditional expectation $Q^n_{\underline{\underline{\pi}}, \underline{\underline{s}}} = \mathbb{E}\left[L(\underline{Y}, \underline{\underline{B}}) | \underline{Y}; \underline{\underline{\pi}}^n, \underline{\underline{s}}^n \right]$ with respect to (w.r.t.) the model parameters: $(\underline{\underline{\pi}}^{n+1}, \underline{\underline{s}}^{n+1}) = \underset{\underline{\underline{\pi}}, \underline{\underline{s}}}{\operatorname{argmax}}\, Q^n_{\underline{\underline{\pi}}, \underline{\underline{s}}}$. An important result about the EM algorithm is that the log-likelihood of the sequence $(\underline{\underline{\pi}}^n, \underline{\underline{s}}^n)_{n \geq 0}$ generated in this way is non-decreasing.

Here $Q^n_{\underline{\underline{\pi}}, \underline{\underline{s}}}$ is the sum of two terms $Q^n_{\underline{\underline{\pi}}}$ and $Q^n_{\underline{\underline{s}}}$, defined as

$$Q^n_{\underline{\underline{\pi}}} = \sum_{i \in \overline{I}} \sum_{\underline{b} \in \mathcal{B}^J} \gamma^n_{i,\underline{b}} \ln(\pi_{i,\underline{b}}) \tag{5}$$

$$Q^n_{\underline{\underline{s}}} = \sum_{i \in \overline{I}} \sum_{\underline{b} \in \mathcal{B}^J} \gamma^n_{i,\underline{b}} \left(\ln \left(\frac{1}{s_{i,\underline{b}}} \right) - \frac{Y_i}{s_{i,\underline{b}}} \right) \tag{6}$$

where $\gamma^n_{i,\underline{b}}$ is the posterior probability that $\underline{B}^i = \underline{b}$ given Y_i.

3.1. E-Step: updating the posterior probabilities

The Bayes theorem proves that

$$\gamma^n_{i,\underline{b}} = \frac{\pi^n_{i,\underline{b}}\, p(Y_i | \underline{B}^i = \underline{b}; \underline{\underline{s}}^n)}{p(Y_i; \underline{\underline{\pi}}^n, \underline{\underline{s}}^n)} \propto \frac{\pi^n_{i,\underline{b}}}{s^n_{i,\underline{b}}} e^{-\frac{Y_i}{s^n_{i,\underline{b}}}}.$$

according to equation (4). These posterior probabilities are thus calculated by using the property $\sum_{\underline{b} \in \mathcal{B}^J} \gamma^n_{i,\underline{b}} = 1$. The posterior marginal probability $\gamma^n_{i,j}$ that $B_{i,j} = 1$ given Y_i is then obtained according to

$$\gamma^n_{i,j} = \sum_{\underline{b} \in \mathcal{B}^J \,\&\, b_j = 1} \gamma^n_{i,\underline{b}}. \tag{7}$$

3.2. M-step: updating the model parameters

3.2.1. Updating the prior probabilities

We are now looking for the prior probabilities $\underline{\underline{\pi}}^{n+1}$ which maximize $Q^n_{\underline{\underline{\pi}}}$ defined in equation (5). First, equations (2) and (7) prove that $Q^n_{\underline{\underline{\pi}}}$ can be decomposed in the form $Q^n_{\underline{\underline{\pi}}} = \sum_{j \in \overline{J}} Q^n_j$, where $Q^n_j = \sum_{i \in \overline{I}} \gamma^n_{i,j} \ln(\pi_{i,j}) + (1 - \gamma^n_{i,j}) \ln(1 - \pi_{i,j})$. The prior probabilities are thus estimated by independently maximizing each Q^n_j w.r.t. the I variables $\pi_{0,j} \dots \pi_{I-1,j}$, which are entirely defined by the three parameters $(l_j, \dot{\pi}_j, \mathring{\pi}_j)$, as described in section 2.2. In particular, this maximization involves estimating the pitch the j^{th} component, via parameter l_j. We thus write Q^n_j as a function of the triplet $(l, \dot{\pi}, \mathring{\pi}) \in \overline{L} \times]0,1[\times]0,1[$, which will have to be maximized to obtain the estimates $(l^{n+1}_j, \dot{\pi}^{n+1}_j, \mathring{\pi}^{n+1}_j)$:

$$Q^n_j(l, \dot{\pi}, \mathring{\pi}) = |\mathcal{H}_l|\, \dot{Q}^n_{l,j}(\dot{\pi}) + (I - |\mathcal{H}_l|)\, \mathring{Q}^n_{l,j}(\mathring{\pi}) \tag{8}$$

where $|\mathcal{H}_l|$ denotes the cardinality of set $\mathcal{H}_l$ and

$$\begin{cases} \dot{Q}^n_{l,j}(\dot{\pi}) = \frac{1}{|\mathcal{H}_l|} \sum_{i \in \mathcal{H}_l} \gamma^n_{i,j} \ln(\dot{\pi}) + (1 - \gamma^n_{i,j}) \ln(1 - \dot{\pi}) \\ \mathring{Q}^n_{l,j}(\mathring{\pi}) = \frac{1}{I - |\mathcal{H}_l|} \sum_{i \notin \mathcal{H}_l} \gamma^n_{i,j} \ln(\mathring{\pi}) + (1 - \gamma^n_{i,j}) \ln(1 - \mathring{\pi}) \end{cases} \tag{9}$$

For all $l \in \overline{L}$, we independently maximize functions $\dot{Q}^n_{l,j}$ and $\mathring{Q}^n_{l,j}$ w.r.t. $\dot{\pi}$ and $\mathring{\pi}$, resulting in the optimal values

$$\begin{cases} \dot{\pi}^n_{l,j} = \frac{1}{|\mathcal{H}_l|} \sum_{i \in \mathcal{H}_l} \gamma^n_{i,j} \\ \mathring{\pi}^n_{l,j} = \frac{1}{I - |\mathcal{H}_l|} \sum_{i \notin \mathcal{H}_l} \gamma^n_{i,j}. \end{cases} \tag{10}$$

Substituting equation (10) into equation (9), we obtain the maximal values of $\dot{Q}^n_{l,j}(\dot{\pi})$ and $\mathring{Q}^n_{l,j}(\mathring{\pi})$ w.r.t. $\dot{\pi}$ and $\mathring{\pi}$:

$$\begin{cases} \dot{Q}^n_{l,j} = \dot{\pi}^n_{l,j} \ln(\dot{\pi}^n_{l,j}) + (1 - \dot{\pi}^n_{l,j}) \ln(1 - \dot{\pi}^n_{l,j}) \\ \mathring{Q}^n_{l,j} = \mathring{\pi}^n_{l,j} \ln(\mathring{\pi}^n_{l,j}) + (1 - \mathring{\pi}^n_{l,j}) \ln(1 - \mathring{\pi}^n_{l,j}) \end{cases}$$

Equation (8) becomes $Q^n_{l,j} = |\mathcal{H}_l|\,\dot{Q}^n_{l,j} + (I - |\mathcal{H}_l|)\,\mathring{Q}^n_{l,j}$. Maximizing this expression w.r.t. l yields $l_j^{n+1} = \underset{l}{\operatorname{argmax}}\, Q^n_{l,j}$, which identifies the estimates $\dot{\pi}_j^{n+1} = \dot{\pi}^n_{(l_j^{n+1},j)}$ and $\mathring{\pi}_j^{n+1} = \mathring{\pi}^n_{(l_j^{n+1},j)}$. The joint probabilities $\pi^{n+1}_{i,\underline{b}}$ are then calculated via equation (2).

3.2.2. Updating the envelope coefficients

Now we are looking for the coefficients σ_j^{n+1} and $\alpha_{k,j}^{n+1}$ which maximize $Q^n_{\underline{\underline{s}}}$ defined in equation (6). Note that in this equation, $s_{i,\underline{b}}$ depends on σ_j and $\alpha_{k,j}$ via equations (3) and (1) (condition $b_j = 1$ is relevant for $j \in \overline{J}$; for $j = J$, this condition is assumed always true in the following developments).

3.2.2.1. Estimation of the variances

Differentiating equation (6) w.r.t. σ_j^2 shows that $\forall j \in \overline{J} \cup \{J\}$, $\sigma_j^2 \frac{\partial Q^n_{\underline{\underline{s}}}}{\partial \sigma_j^2} = \rho_j^+ - \rho_j^-$, where $\begin{cases} \rho_j^- &= \sum\limits_{i\in\overline{I}} \sum\limits_{\underline{b}\in\mathcal{B}^J \& b_j=1} \gamma^n_{i,\underline{b}} \frac{s_{i,j}}{s_{i,\underline{b}}} \\ \rho_j^+ &= \sum\limits_{i\in\overline{I}} \sum\limits_{\underline{b}\in\mathcal{B}^J \& b_j=1} \gamma^n_{i,\underline{b}} \frac{s_{i,j}}{s_{i,\underline{b}}} \frac{Y_i}{s_{i,\underline{b}}} \end{cases}$

Zeroing this derivative does not admit a closed-form solution, but it can be proved that the multiplicative update rule $\sigma_j^2 \leftarrow \frac{\rho_j^+}{\rho_j^-}\,\sigma_j^2$ forms an ascent method which converges to the maximum of $Q^n_{\underline{\underline{s}}}$. In practice, we perform one such update per iteration of the EM algorithm, which still guarantees that the log-likelihood $Q^n_{\underline{\underline{s}}}$ is non-decreasing.

3.2.2.2. Estimation of the MA parameters

Differentiating (6) w.r.t. $\alpha_{k,j}$ yields $\forall k \in \{0 \dots K\}$, $\forall j \in \overline{J} \cup \{J\}$, $\frac{1}{\sigma_j^2} \frac{\partial Q^n_{\underline{\underline{s}}}}{\partial \alpha_{k,j}} = \sum\limits_{\kappa=0}^{K} \alpha_{\kappa,j}\, r_j^+(k-\kappa) - \sum\limits_{\kappa=0}^{K} \alpha_{\kappa,j}\, r_j^-(k-\kappa)$, where

$$\begin{cases} r_j^-(k) &= \sum\limits_{i\in\overline{I}} 2\cos\left(2\pi k \frac{i}{I}\right) \sum\limits_{\underline{b}\in\mathcal{B}^J \& b_j=1} \frac{\gamma^n_{i,\underline{b}}}{s_{i,\underline{b}}} \\ r_j^+(k) &= \sum\limits_{i\in\overline{I}} 2\cos\left(2\pi k \frac{i}{I}\right) \sum\limits_{\underline{b}\in\mathcal{B}^J \& b_j=1} \frac{\gamma^n_{i,\underline{b}}}{s_{i,\underline{b}}} \frac{Y_i}{s_{i,\underline{b}}} \end{cases}$$

Let $\underline{\alpha}_j = [\alpha_{0,j} \dots \alpha_{K,j}]^T$, and $\underline{\underline{R}}_j^-$ and $\underline{\underline{R}}_j^+$ be the $(K+1) \times (K+1)$ Toeplitz matrices whose coefficients are $r_j^-(k)$ and $r_j^+(k)$, respectively. Again, the maximum does not admit a closed-form expression, but the multiplicative update rule $\underline{\alpha}_j \leftarrow (\underline{\underline{R}}_j^-)^{-1}(\underline{\underline{R}}_j^+)\underline{\alpha}_j$ guarantees that the log-likelihood $Q^n_{\underline{\underline{s}}}$ is non-decreasing. The convergence of such multiplicative rules for estimating MA models was studied in [6]. Note that at each iteration, vector $\underline{\alpha}_j^n$ should be remapped into the set of minimum phase filters, following the approach proposed in [6].

3.3. Discussion

The two main advantages of this EM algorithm are its low complexity (multi-pitch estimation is performed by means of J successive single-pitch estimations, instead of exploring a vector space of dimension J), and its ability to handle spectral overlap between the harmonic components (by taking the smoothness of the spectral envelopes into account). However, the log-likelihood function $Q^n_{\underline{\pi},\underline{\underline{s}}}$ to be optimized is generally not smooth, and the algorithm tends to stay trapped in local maxima[1]. Therefore it is necessary to develop strategies to escape from these maxima, for instance by testing multiples or sub-multiples of the current estimated frequencies. Initialization is also an important point. An approach (illustrated in section 4) would be to use the EM algorithm as a refinement, after a first stage of basic multi-pitch estimation (any fast multi-pitch estimator can be used, such as [3]). Indeed, we observed that this algorithm is capable of correcting some wrongly estimated pitches when initialized properly. To summarize, the algorithm requires the following inputs:

- the initial values of the multiple pitches via the indexes l_j. We then propose to initialize the a priori probabilities $\pi_{i,j}$ as harmonic combs, as illustrated in Fig. 1, with $\dot{\pi}_j = 1 - \varepsilon$, $\mathring{\pi}_j = \varepsilon$, and $\varepsilon \ll 1$. Finally, the posterior probabilities $\gamma_{i,j}$ are taken equal to $\pi_{i,j}$.
- the initial spectral envelopes. We propose to take flat MA envelopes ($\alpha_{k,j} = 0\ \forall k \geq 1$), with variances of different orders, and to let the multiplicative updates presented in section 3.2.2 converge (contrary to the EM algorithm itself, where only one update is performed per iteration).

Note that in practice, like any EM-based approach, this algorithm has to be implemented very carefully to avoid numerical errors, due to possible divisions by numbers much smaller than one.

4. SIMULATION RESULTS: MUSICAL CHORD IDENTIFICATION

In this section, our EM algorithm is applied to an audio-like synthetic signal. Music signals indeed provide an interesting, challenging field of application to our algorithm, since they often contain mixtures of harmonic components with a substantial spectral overlap (which is typically the case of consonant chords). Here we use a sampling frequency of 22 kHz, and we consider a chord formed of $J = 4$ notes, of fundamental frequencies 330, 440, 550, and 660 Hz (corresponding to musical notes E4 - A4 - C♯5 - E5, forming a major triad involving an octave, which generates a large spectral overlap). For modeling the spectral envelopes of the three notes, we use MA models of order 5, whose coefficients are normalized so that the maximal values of the envelopes of the four notes are -30, -10, -20, and 0 dB, respectively. The amplitudes $X_{i,j}$ of the partials of the four notes are then obtained by sampling their MA envelope at the harmonic frequencies (their phases being generated as independent, uniformly distributed random variables). The additive noise N_i is then synthesized as a white Gaussian noise, whose variance is chosen so that its constant power density function (PSD) is -60 dB. We compute $I = 1000$ samples of this mixture signal, corresponding to a 45 ms-long frame, which is a typical analysis length for multi-pitch estimation. The digital Fourier transform of the signal is then computed with the same number of samples I, without applying a tapering window, in order to avoid spectral leakage, and correlation between adjacent Fourier bins. The periodogram of the resulting signal is represented in Fig. 3-(a).

The EM algorithm is applied with $L = 88$ pitch candidates, distributed according to the piano MIDI scale which ranges from note A0 to C8, including the $J = 4$ true pitches. To simulate a first stage of basic multi-pitch estimation, the algorithm is initialized with the following pitches : 330, 880, 748, and 660 Hz, corresponding to notes E4, A5, F♯5, and E5 (two erroneous pitches out of four). The MA envelopes are initialized as constant PSD of magnitude 0 dB, and -40 dB for the noise part. Fig. 2 represents the 4 detection functions $l \mapsto Q^n_{l,j}$ obtained after $n = 25$ iterations of the algorithm. It is worth noticing that the multi-pitch estimation problem is reduced to four independent single-pitch estimation problems, and that the detection functions admit a strong maximum at the true pitch values.

[1] Local maxima are inherent to the pitch estimation problem. Such maxima appear for instance in the pitch detection functions represented in Fig. 2.

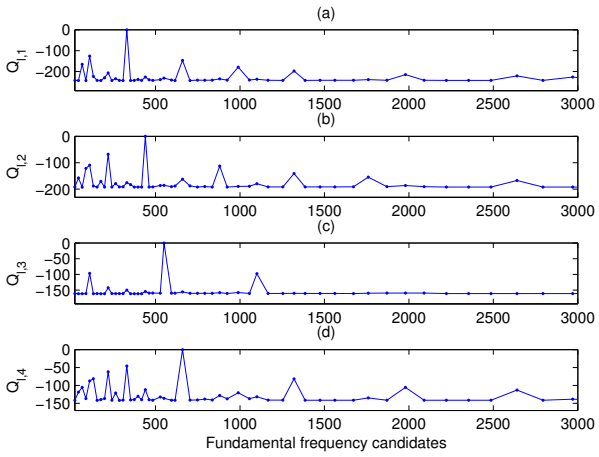

Fig. 2. Pitch detection functions (dB)
(a) E4, 330 Hz, (b) A4, 440 Hz, (c) C♯5, 550 Hz, (d) E5, 660 Hz

The output of the algorithm also permits to separate the magnitude spectra of the four notes. Indeed, for all $j \in \overline{J} \cup \{J\}$, the squared magnitude of the j^{th} spectrum component $|X_{i,j}|^2$ (if $j < J$) or $|N_i|^2$ (if $j = J$) can be estimated as a weighted sum of products between the outcome spectrum Y_i and the Wiener-like filters $\frac{s_{i,j}^2}{s_{i,\underline{b}}^2}$: $\widehat{|X|}_{i,j}^2 = \sum_{\underline{b} \in \mathcal{B}^J \,\&\, b_j = 1} \pi_{i,\underline{b}} \frac{s_{i,j}^2}{s_{i,\underline{b}}^2} Y_i$. The separated spectra are represented in Fig. 3-(b-f) (solid lines), superimposed with the original MA envelopes (dashed lines). Note that the spectral components are efficiently estimated. However in the case of overlapping partials, the amplitude of the weakest one tends to be under-estimated (the estimation of the even harmonics of note E4 in Fig. 3-(b) is perturbed by the partials of note E5, represented in Fig. 3-(e)).

5. CONCLUSIONS

In this paper, we introduced a novel approach for multi-pitch estimation, based on the statistical framework of the EM algorithm. The proposed method is particularly promising, due to its robustness to overlapping partials, and its capacity to simplify the multi-pitch estimation task into successive single-pitch estimations. It requires a proper initialization, involving a first stage of basic multi-pitch estimation for instance, and could advantageously make use of heuristics, in order to avoid to stay trapped in local maxima. The effectiveness of this approach is confirmed by our simulations, performed on audio-like synthetic signals. In order to apply this algorithm to real audio signals, and compare its performance to that of competing methods, some additional improvements will be helpful:

- The design of the sets $\mathcal{H}_l$ of harmonic frequencies may be refined by means of peak-picking in the Fourier spectrum. In the case of inharmonic instruments such as the piano, the inharmonicity coefficient may also be included in the model.
- Since the estimation of a PSD usually involves windows such as Hann or Hamming, the spectral leakage and correlation between adjacent Fourier bins should be integrated in the model.
- The number J of harmonic components, which is supposed known herein this paper, could be estimated by incorporating a statistical criterion, such as BIC [7].

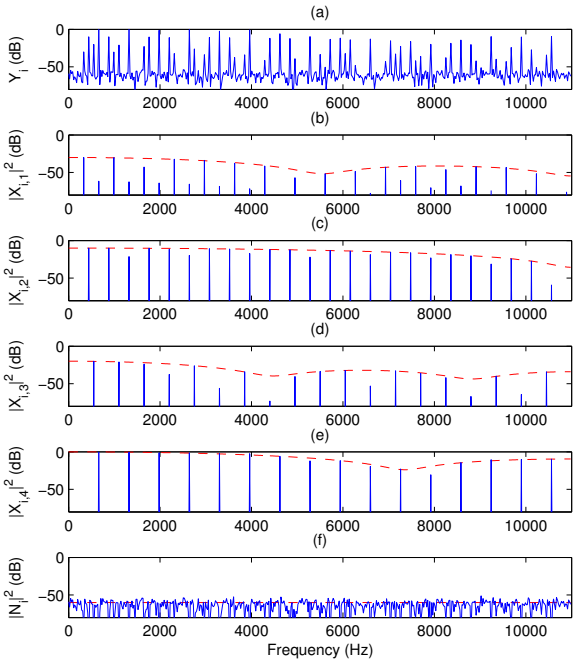

Fig. 3. Observed (a) and estimated (b-f) spectra

Besides, the proposed framework permits to chose various models for the spectral envelopes, which could also be parameterized by using cepstral representations for instance, or mixtures of template spectra, either defined according to a psychoacoustic frequency scale, such as Mel, Bark, and ERB, or learned from a database of harmonic signals.

6. REFERENCES

[1] A. Klapuri, "Multiple fundamental frequency estimation by harmonicity and spectral smoothness," *IEEE Trans. Speech Audio Processing*, vol. 11, no. 6, pp. 804–816, Nov. 2003.

[2] A. Klapuri and M. Davy, Eds., *Signal Processing Methods for Music Transcription*, Springer, New York, 2006.

[3] A. de Cheveigné, D. Wang, and G.J. Brown, Eds., *Computational Auditory Scene Analysis: Principles, Algorithms and Applications*, chapter Multiple F0 estimation, Wiley-IEEE Press, Piscataway, NJ, 2006.

[4] V. Emiya, R. Badeau, and B. David, "Multipitch estimation of inharmonic sounds in colored noise," in *Proc. of 10th International Conference on Digital Audio Effects DAFx-07*, Bordeaux, France, Sept. 2007, pp. 93–98.

[5] A. Dempster, N. Laird, and D. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society, Series B*, vol. 39, no. 1, pp. 1–38, 1977.

[6] R. Badeau and B. David, "Weighted maximum likelihood autoregressive and moving average spectrum modeling," in *Proc. of 2008 International Conference on Acoustics, Speech, and Signal Processing ICASSP'08*, Las Vegas, Nevada, USA, Apr. 2008, pp. 3761–3764, IEEE.

[7] G. Schwarz, "Estimating the dimension of a model," *Annals of Statistics*, vol. 6, no. 2, pp. 461–464, 1978.

Stability analysis of multiplicative update algorithms and application to non-negative matrix factorization

Roland Badeau, *Senior Member, IEEE,* Nancy Bertin, and Emmanuel Vincent, *Senior Member, IEEE*

***Abstract*—Multiplicative update algorithms have encountered a great success to solve optimization problems with non-negativity constraints, such as the famous non-negative matrix factorization (NMF) and its many variants. However, despite several years of research on the topic, the understanding of their convergence properties is still to be improved. In this paper, we show that Lyapunov's stability theory provides a very enlightening viewpoint on the problem. We prove the exponential or asymptotic stability of the solutions to general optimization problems with non-negative constraints, including the particular case of supervised NMF, and finally study the more difficult case of unsupervised NMF. The theoretical results presented in the paper are confirmed by numerical simulations involving both supervised and unsupervised NMF, and the convergence speed of NMF multiplicative updates is investigated.**

***Index Terms*—Optimization methods, non-negative matrix factorization, multiplicative update algorithms, convergence of numerical methods, stability, Lyapunov methods.**

I. Introduction

NON-NEGATIVE matrix factorization (NMF) is a powerful decomposition technique allowing the decomposition of two-dimensional non-negative data as a linear combination of meaningful elements in a dictionary [24]. Given an $F \times T$ data matrix $\boldsymbol{V}$ having non-negative coefficients, NMF consists in computing a rank-K truncated approximation $\widehat{\boldsymbol{V}}$ of matrix $\boldsymbol{V}$ (with $K < \min(F,T)$) as a product $\widehat{\boldsymbol{V}} = \boldsymbol{W}\boldsymbol{H}$, where both the $F \times K$ matrix $\boldsymbol{W}$ and the $K \times T$ matrix $\boldsymbol{H}$ have non-negative coefficients. The columns of matrix $\boldsymbol{W}$ form the elements of the dictionary, and the rows of $\boldsymbol{H}$ contain the coefficients of the decomposition. The computation of this factorization is generally formalized as a constrained optimization problem. The objective functions most often encountered in NMF literature rely on the Euclidean (EUC) distance, the Kullback-Leibler (KL) divergence [24], [25], or the Itakura-Saito (IS) divergence [19]. The three of them are enclosed in the general framework of β-divergences [10], [14].

NMF can be considered either as a *supervised*, or as an *unsupervised* learning tool. In the case of supervised learning [12], [18], [30], [32], the dictionary $\boldsymbol{W}$ is estimated from training data in a preprocessing stage, and matrix $\boldsymbol{H}$ only has to be computed given the data in matrix $\boldsymbol{V}$. In the case of unsupervised learning [24], both matrices $\boldsymbol{W}$ and $\boldsymbol{H}$ have to be computed given $\boldsymbol{V}$. Several algorithms have been proposed in order to compute an NMF. The most popular is the multiplicative update algorithm initially proposed by Lee and Seung [25] for the EUC and KL divergences, which has then been generalized to the β-divergence [10], [20]. This algorithm can be applied both to supervised and to unsupervised NMF. For the interested reader, other approaches have also been proposed, such as the projected gradient method [28], alternating least squares (ALS) algorithms [16], the quasi-Newton algorithm [11], a multilayer technique [9], or a space-alternating expectation-maximization (SAGE) algorithm derived in a statistical framework [15]. See for instance [11] for a recent survey on the topic. Theoretical and numerical comparisons between Lee and Sung's multiplicative updates and other algorithms are already available in the literature, see *e.g.* [2], [8], [9], [11], [17], [28].

After Lee and Seung's paper, the multiplicative update philosophy has been applied to various optimization problems involving non-negativity constraints, such as some variants of the NMF. These algorithms generally aim at enhancing (or enforcing) a particular property in the decomposition, depending on the application. In the context of image representation and recognition for instance, various properties have been investigated, such as orthogonality [7], spatial localization [26], or transformation-invariance [13]. In the context of multipitch and music transcription, some desired properties are spectral harmonicity [4], [5], [33] and temporal continuity [4], [5]. In source separation, classical constraints include sparseness and temporal continuity [34], decorrelation [36], and shift-invariance [31]. Note that multiplicative updates have also been applied to non-negative tensor factorization *via* unfolding [11].

A curious point is that to the best of our knowledge, despite many years of research and several papers on the topic, the convergence properties of multiplicative update algorithms for unsupervised NMF have not been clearly identified:

- Lee and Seung proved that the objective function based on EUC and KL decreases at each iteration [25] (and the proof was later generalized to β-divergences [20], for $\beta \in [1,2]$). However, this proves neither that the limit value of the objective function is a local minimum, nor that the successive iterates converge to a limit point.
- In constrained optimization, all the local minima of the objective function are proved to be stationary points, defined as the solutions to Karush, Kuhn and Tucker's (KKT) optimality conditions. Stationary points of NMF were studied in [2], [8], [27]. Some numerical examples have been presented in [17], where the KKT conditions

The research leading to this paper was supported by the French GIP ANR under contract ANR-06-JCJC-0027-01, *Décompositions en Éléments Sonores et Applications Musicales* - DESAM.

Roland Badeau is with Institut Télécom; Télécom ParisTech; CNRS LTCI, 75634 Paris cedex 13, France (e-mail: roland.badeau@telecom-paristech.fr). Emmanuel Vincent and Nancy Bertin are with INRIA, Centre Inria Rennes - Bretagne Atlantique, 35042 Rennes cedex, France (e-mail: emmanuel.vincent@inria.fr; nancy.bertin@inria.fr).

are not fulfilled after a high (but finite) number of iterations, but this does not contradict the possible asymptotic convergence to a local minimum.
- Since the multiplicative updates involve ratios, numerical problems could be encountered if the denominator becomes arbitrarily small. In order to circumvent this problem, it is proposed in [27] to add a small positive quantity to the denominator, and it is proved that any accumulation point of the sequence of the iterates computed in this way is a stationary point[1]. However there is no guarantee that such a stationary point is a local minimum, nor that the algorithm converges to this accumulation point.

Analyzing the convergence properties of unsupervised NMF multiplicative updates is difficult because at each iteration, these algorithms usually switch between two different updates: one for the left factor $\boldsymbol{W}$, and one for the right factor $\boldsymbol{H}$. Nevertheless, the convergence analysis happens to be simpler in the case of supervised NMF, where only one of the two factors is updated, the other one being kept unchanged throughout the iterations. In this paper, we intend to analyze the convergence of general multiplicative update algorithms, where all variables are updated at once (which includes the particular case of supervised NMF), before studying the case of unsupervised NMF. Two important aspects of these algorithms have to be taken into account:

- *Local convergence:* Since the objective function generally admits several local minima [2], [9], there is no guarantee that the algorithm converges to the global minimum. So the best result we can prove is the local convergence to a local minimum[2]. This means that if the algorithm is initialized in a given neighborhood of a local minimum called *basin of attraction*, then the algorithm will converge to this local minimum.
- *Stability:* Because of the multiplicative form of the algorithm, a zero entry remains zero in all subsequent iterations. Zeroing may happen because of a bad initialization or because of the finite machine precision for instance. This prevents the convergence to a local minimum whose corresponding coefficient would be non-zero. However, this problem can be very easily circumvented. Indeed, in this case, numerical simulations show that the limit point of the algorithm is generally unstable: replacing the zero entry by any arbitrarily small quantity will make the algorithm escape from this trap and finally converge to a stable limit point. Other well-known examples of unstable stationary points of the algorithm (with non-zero entries) are saddle points [2], [8], [16].

These remarks show that an appropriate notion for analyzing the convergence of multiplicative update algorithms is the *asymptotic stability* in the sense of Lyapunov's theory [22], which implies both local and stable convergence. In this paper, we analyze the convergence properties of general multiplicative update algorithms, designed to solve optimization problems with non-negativity constraints. We thus apply Lyapunov's first and second methods to find some criteria which guarantee the exponential or asymptotic stability of the local minima of the objective function. This analysis is then applied to prove the stability of supervised NMF multiplicative updates, and we finally show how Lyapunov's first method provides some interesting insights into the convergence properties of unsupervised NMF multiplicative updates. The numerical simulations illustrate those theoretical results, and the convergence speed of NMF multiplicative updates is investigated. The paper is organized as follows: in section II, we present some elementary results about NMF, general optimization problems with non-negativity constraints, and multiplicative update algorithms. The convergence of these algorithms is analyzed by means of Lyapunov's stability theory in section III, and the case of non-negative matrix factorization is studied in section IV. Some numerical simulations are presented in section V, and the main conclusions are summarized in section VI. Finally, the mathematical proofs of the main results presented in this paper are included in the Appendix (due to the lack of space, some proofs have been moved to a separate document [1]).

II. Theoretical background

A. Multiplicative update algorithms and NMF

Given a matrix $\boldsymbol{V} \in \mathbb{R}_+^{F\times T}$ and an integer $K < \min(F,T)$, unsupervised NMF consists in computing a reduced-rank approximation of $\boldsymbol{V}$ as a product $\widehat{\boldsymbol{V}} = \boldsymbol{W}\boldsymbol{H}$, where $\boldsymbol{W} \in \mathbb{R}_+^{F\times K}$ and $\boldsymbol{H} \in \mathbb{R}_+^{K\times T}$. This problem can be formalized as the minimization of an objective function

$$D(\boldsymbol{V}|\boldsymbol{W}\boldsymbol{H}) = \sum_{f=1}^{F}\sum_{t=1}^{T} d\left(v_{ft}\middle|\sum_{k=1}^{K} w_{fk}h_{kt}\right), \qquad (1)$$

where d is a scalar divergence (*i.e.* a function such that $\forall x,y \in \mathbb{R}_+, d(x|y) \geq 0$, and $d(x|y) = 0$ if and only if $y = x$).

β-divergences [10], [14] are defined for all $\beta \in \mathbb{R}\backslash\{0,1\}$ as

$$d_\beta(x|y) = \frac{1}{\beta(\beta-1)}\left(x^\beta + (\beta-1)y^\beta - \beta x y^{\beta-1}\right). \qquad (2)$$

The Euclidean distance corresponds to $\beta = 2$, and KL and IS divergences are obtained when $\beta \to 1$ and $\beta \to 0$, respectively: $d_{KL}(x|y) = x\ln\left(\frac{x}{y}\right) - x + y$, and $d_{IS}(x|y) = \frac{x}{y} - \ln\left(\frac{x}{y}\right) - 1$.

The generalization of Lee and Seung's multiplicative updates to the β-divergence takes the following form [10], [20]:

$$\boldsymbol{W} \leftarrow \boldsymbol{W} \otimes \frac{(\boldsymbol{V}\otimes(\boldsymbol{W}\boldsymbol{H})^{\beta-2})\boldsymbol{H}^T}{(\boldsymbol{W}\boldsymbol{H})^{\beta-1}\boldsymbol{H}^T} \qquad (3)$$

$$\boldsymbol{H} \leftarrow \boldsymbol{H} \otimes \frac{\boldsymbol{W}^T(\boldsymbol{V}\otimes(\boldsymbol{W}\boldsymbol{H})^{\beta-2})}{\boldsymbol{W}^T(\boldsymbol{W}\boldsymbol{H})^{\beta-1}} \qquad (4)$$

where the symbol $\otimes$ and the fraction bar denote entrywise matrix product and division respectively, and the exponentiations must also be understood entrywise. In the case of unsupervised NMF, updates (3) and (4) are computed alternately, whereas in

[1]Note that this proof only stands for the Euclidean distance, and does not apply when the added quantity is zero.

[2]Avoiding to be trapped in a local minimum distinct from the global minimum is out of the scope of this paper (the interested reader can have a look at [3], [6]).

the case of supervised NMF, the update (4) only is computed at each iteration, matrix $\boldsymbol{W}$ being kept unchanged.

In [20], it is proved that if $\beta \in [1,2]$, then the objective function is non-increasing at each iteration of (3) and (4). As in [25], this algorithm can be interpreted as a gradient descent with an adaptive step size for each entry, defined as a function of both matrices $\boldsymbol{W}$ and $\boldsymbol{H}$, chosen so that the successive iterates remain non-negative. Alternatively, we consider the approach used in [4], [7], [34]: the recursion for $\boldsymbol{W}$ in (3) can be written $\forall f,k,\ w_{fk} \leftarrow w_{fk}\frac{m^w_{fk}}{p^w_{fk}}$, where $p^w_{fk} \geq 0$ and $m^w_{fk} \geq 0$ are such that the partial derivative of the objective function w.r.t. w_{fk} is equal to $p^w_{fk} - m^w_{fk}$:

$$\frac{\partial D}{\partial w_{fk}} = \underbrace{\sum_{t=1}^{T} \hat{v}_{ft}^{\beta-1} h_{kt}}_{p^w_{fk}} - \underbrace{\sum_{t=1}^{T} v_{ft}\, \hat{v}_{ft}^{\beta-2} h_{kt}}_{m^w_{fk}}, \tag{5}$$

where

$$\hat{v}_{ft} = \sum_{k=1}^{K} w_{fk} h_{kt}. \tag{6}$$

Thus if $\frac{\partial D}{\partial w_{fk}} > 0$, $\frac{m^w_{fk}}{p^w_{fk}} < 1$ so that w_{fk} decreases, and conversely if $\frac{\partial D}{\partial w_{fk}} < 0$, $\frac{m^w_{fk}}{p^w_{fk}} > 1$ so that w_{fk} increases. The same remark stands for matrix $\boldsymbol{H}$. This confirms that the updates (3) and (4) form a descent method.

We focus in this paper on a generalization of this approach which involves an exponent step size $\eta > 0$:

$$\boldsymbol{W} \leftarrow \boldsymbol{W} \otimes \left(\frac{(\boldsymbol{V} \otimes (\boldsymbol{W}\boldsymbol{H})^{\beta-2})\boldsymbol{H}^T}{(\boldsymbol{W}\boldsymbol{H})^{\beta-1}\boldsymbol{H}^T} \right)^{\eta} \tag{7}$$

$$\boldsymbol{H} \leftarrow \boldsymbol{H} \otimes \left(\frac{\boldsymbol{W}^T(\boldsymbol{V} \otimes (\boldsymbol{W}\boldsymbol{H})^{\beta-2})}{\boldsymbol{W}^T(\boldsymbol{W}\boldsymbol{H})^{\beta-1}} \right)^{\eta} \tag{8}$$

Note that standard multiplicative updates (3) and (4) correspond to the particular case $\eta = 1$. As will be shown in section V, η actually permits to control the convergence rate, and in particular to outperform the standard case $\eta = 1$. This approach was first introduced in [4].

The convergence properties of these generalized updates will be analyzed in section IV. The following proposition proves that (7) and (8) satisfy the same decrease property as (3) and (4) for all $\eta \in]0,1]$ (*i.e.* $0 < \eta \leq 1$)[3].

Proposition 1. *Consider the objective function $D(\boldsymbol{V}|\boldsymbol{W}\boldsymbol{H})$ defined in equation* (1)*, involving the β-divergence* (2)*, with $\beta \in [1,2]$. If $\eta \in]0,1]$, if all coefficients in the numerator and denominator in recursion* (7) *are non-zero and if $(\boldsymbol{W},\boldsymbol{H})$ is not a fixed point of* (7) *, then* (7) *makes the objective function strictly decrease. Similarly, if $\eta \in]0,1]$, if all coefficients in the numerator and denominator in recursion* (8) *are non-zero and if $(\boldsymbol{W},\boldsymbol{H})$ is not a fixed point of* (8) *, then* (8) *makes the objective function strictly decrease.*

Proposition 1 is proved in Appendix A. Note that this decrease property does not guarantee that that the limit value of the objective function is a local minimum, nor that the successive values of $\boldsymbol{W}$ and $\boldsymbol{H}$ converge to a limit point. From now on and until section IV, we will introduce a general framework for multiplicative update algorithms, where all variables are updated at once (which includes the particular case of supervised NMF).

[3]In this paper we use the ISO notation for intervals, which uses inwards pointing brackets to indicate inclusion of the endpoint, and outwards pointing brackets for exclusion.

B. General optimization problems with non-negativity constraints

We consider the minimization in the first orthant $\mathbb{R}^n_+$ of an objective function $J : \mathbb{R}^n_+ \to \mathbb{R}$, which is twice continuously differentiable in its domain. For any vector $\boldsymbol{x} = [x_1 \dots x_n]^T \in \mathbb{R}^n_+$, the constraint $x_i \geq 0$ is said to be *active* if $x_i = 0$, or *inactive* if $x_i > 0$. The following two propositions are classical results in constrained optimization theory [29].

Proposition 2 (First order KKT optimality conditions)**.** *Let $\boldsymbol{\nabla} J(\boldsymbol{x})$ denote the gradient vector of the objective function $J : \mathbb{R}^n_+ \to \mathbb{R}$, which is twice continuously differentiable in its domain. Then for any local minimum $\boldsymbol{x}$ of J in $\mathbb{R}^n_+$,*

- *$\forall i \in \{1 \dots n\}$ such that $x_i > 0$, $\nabla_i J = 0$;*
- *$\forall i \in \{1 \dots n\}$ such that $x_i = 0$, $\nabla_i J \geq 0$.*

If $x_i = 0$ and $\nabla_i J > 0$, the constraint is said to be *strictly active*. Following these considerations, we introduce the following notation for denoting the extraction of particular sub-vectors or sub-matrices:

Notation 1.

- $[.]_0$ is obtained by selecting the coefficients (of a vector) or the rows and columns (of a matrix) corresponding to strictly active constraints, *i.e.* whose index i is such that $\nabla_i J(\boldsymbol{x}) > 0$ (and $x_i = 0$).
- $[.]_+$ is obtained by selecting the coefficients, or the rows and columns, whose index i is such that $\nabla_i J(\boldsymbol{x}) = 0$ (and either $x_i > 0$ or $x_i = 0$).
- $[.]^\star_+$ is obtained by selecting the coefficients, or the rows and columns, corresponding to inactive constraints, *i.e.* whose index i is such that $x_i > 0$ (and $\nabla_i J(\boldsymbol{x}) = 0$).

We then have the following optimality condition at order 2:

Proposition 3 (Second order optimality condition)**.** *Let $\boldsymbol{\nabla}^2 J(\boldsymbol{x})$ denote the $n \times n$ Hessian matrix of the objective function $J : \mathbb{R}^n_+ \to \mathbb{R}$, which is twice continuously differentiable in its domain. Then for any local minimum $\boldsymbol{x}$ of J in $\mathbb{R}^n_+$, the sub-matrix $[\boldsymbol{\nabla}^2 J(\boldsymbol{x})]_+$ is positive semi-definite.*

C. Multiplicative update algorithms

In order to introduce general multiplicative update algorithms, we first have to assume that function J satisfies certain conditions.

Assumption 1 (Decomposability of the objective function)**.** Let $J : \mathbb{R}^n_+ \to \mathbb{R}$ be an objective function which is twice continuously differentiable in its domain. We assume that the gradient of J can be decomposed as the difference of two non-negative functions:

$$\boldsymbol{\nabla} J(\boldsymbol{x}) = \boldsymbol{p}(\boldsymbol{x}) - \boldsymbol{m}(\boldsymbol{x}), \tag{9}$$

where both functions $\boldsymbol{p} : \mathbb{R}_+^n \to \mathbb{R}_+^n$ and $\boldsymbol{m} : \mathbb{R}_+^n \to \mathbb{R}_+^n$ are continuously differentiable in the domain of J.

Note that in Assumption 1, functions $\boldsymbol{p}$ and $\boldsymbol{m}$ are not unique. Indeed, any non-negative constant (or continuously differentiable function) can be added to both $\boldsymbol{p}$ and $\boldsymbol{m}$, without changing their difference.

Example. An example of function J satisfying Assumption 1 will be provided in equation (17), at the beginning of section IV. It corresponds to the NMF objective function introduced in equation (1), whose partial derivatives can be written as the difference of two continuously differentiable non-negative functions, defined in equations (5) and (32).

Assumption 2. Let $\boldsymbol{x} \in \mathbb{R}_+^n$. Given an objective function J satisfying Assumption 1, we assume that $\forall i \in \{1 \dots n\}$, $p_i(\boldsymbol{x}) > 0$ and $m_i(\boldsymbol{x}) > 0$.

Note that if Assumption 1 stands, there always exist functions $p_i(\boldsymbol{x})$ and $m_i(\boldsymbol{x})$ such that Assumption 2 also stands $\forall \boldsymbol{x} \in \mathbb{R}_+^n$. Indeed, any positive constant (or continuously differentiable function) can be added to both functions $p_i(\boldsymbol{x})$ and $m_i(\boldsymbol{x})$, without changing their difference.

Definition 1 (Multiplicative mapping)**.** Consider the minimization of an objective function J satisfying Assumption 1. For any step size $\eta \in \mathbb{R}$, the multiplicative mapping $\boldsymbol{\phi}$ is defined in the domain of all $\boldsymbol{x} \in \mathbb{R}_+^n$ satisfying Assumption 2, as

$$\boldsymbol{\phi}(\boldsymbol{x}) = \boldsymbol{\Lambda}(\boldsymbol{x})^\eta \, \boldsymbol{x} \tag{10}$$

where $\boldsymbol{\Lambda}(\boldsymbol{x}) = \operatorname{diag}\left(\frac{\boldsymbol{m}(\boldsymbol{x})}{\boldsymbol{p}(\boldsymbol{x})}\right)$, $\boldsymbol{m}(\boldsymbol{x})$ and $\boldsymbol{p}(\boldsymbol{x})$ have been defined in Assumption 1, and $\operatorname{diag}(.)$ denotes the diagonal matrix whose diagonal coefficients are those of the vector argument.

The proof of the following lemma is straightforward.

Lemma 4 (Regularity of the mapping)**.** *Let J be an objective function satisfying Assumption 1, and let $\boldsymbol{x} \in \mathbb{R}_+^n$ satisfying Assumption 2. Then $\forall \eta \in \mathbb{R}$, the mapping $\boldsymbol{\phi}$ introduced in Definition 1 is defined and continuously differentiable in a neighborhood of $\boldsymbol{x}$.*

The following lemma is a corollary of Proposition 2.

Lemma 5. *Let $\boldsymbol{x} \in \mathbb{R}_n^+$ be a local minimum of a function J satisfying Assumption 1. If Assumption 2 holds, then $\boldsymbol{x}$ is a fixed point of the mapping $\boldsymbol{\phi}$ introduced in Definition 1 (*i.e. *$\boldsymbol{\phi}(\boldsymbol{x}) = \boldsymbol{x}$).*

Proof: Obviously, if $x_i = 0$, $\phi_i(\boldsymbol{x}) = \boldsymbol{x}_i$. Otherwise Proposition 2 proves that $\forall i \in \{1 \dots n\}$ such that $x_i > 0$, $p_i(\boldsymbol{x}) = m_i(\boldsymbol{x})$, thus $\phi_i(\boldsymbol{x}) = \boldsymbol{x}_i$. ∎

Definition 2 (Multiplicative update algorithm)**.** Consider the minimization of an objective function J satisfying Assumption 1. A multiplicative update algorithm is defined by a recursion of the form $\boldsymbol{x}^{(p+1)} = \boldsymbol{\phi}(\boldsymbol{x}^{(p)})$, where the mapping $\boldsymbol{\phi} : \mathbb{R}_+^n \to \mathbb{R}_+^n$ was defined in Definition 1.

Note that if there exists p such that $\boldsymbol{x}^{(p)}$ does not satisfy Assumption 2, then $\boldsymbol{x}^{(p+1)}$ may be undefined, and the algorithm must stop[4]. Also note that recursion $\boldsymbol{x}^{(p+1)} = \boldsymbol{\phi}(\boldsymbol{x}^{(p)})$ can be seen as a descent method. Indeed, if $\boldsymbol{x} \in \mathbb{R}_+^n$ satisfies Assumption 2, the first order expansion of function $\eta \mapsto J\left(\boldsymbol{\Lambda}(\boldsymbol{x})^\eta \, \boldsymbol{x}\right)$ in a neighborhood of $\eta = 0$ yields

$$\begin{aligned} &J(\boldsymbol{\phi}(\boldsymbol{x})) - J(\boldsymbol{x}) = \eta \boldsymbol{\nabla} J(\boldsymbol{x})^T \left(\ln\left(\boldsymbol{\Lambda}(\boldsymbol{x})\right) \boldsymbol{\phi}(\boldsymbol{x})\right) + O(\eta^2) \\ &= -\eta \sum_{i=1}^{n} \phi_i(\boldsymbol{x})(m_i(\boldsymbol{x}) - p_i(\boldsymbol{x})) \ln\left(\tfrac{m_i(\boldsymbol{x})}{p_i(\boldsymbol{x})}\right) + O(\eta^2) \end{aligned}$$

This equation shows that if $\boldsymbol{\phi}(\boldsymbol{x}) \neq \boldsymbol{x}$ and if $\eta > 0$ is small enough, then $J(\boldsymbol{\phi}(\boldsymbol{x})) - J(\boldsymbol{x}) < 0$, which means that the objective function decreases.

III. Stability analysis of multiplicative updates

We analyze the convergence of multiplicative update algorithms by means of Lyapunov's stability theory. In neural networks' litterature, this theory has been used for various problems, such as analyzing the global exponential stability of discrete recurrent neural networks with time-varying delays [35], [37] or analyzing the discrete-time dynamics of a class of self-stabilizing minor component analysis (MCA) extraction algorithms [21].

A. Stability definitions

Let us recall a few classical definitions in Lyapunov's stability theory of discrete dynamical systems [22]. Notation $\|.\|$ denotes any vector norm.

Definition 3 (Lyapunov stability)**.** A fixed point $\boldsymbol{x} \in \mathbb{R}_+^n$ of the recursion $\boldsymbol{x}^{(p+1)} = \boldsymbol{\phi}(\boldsymbol{x}^{(p)})$, where mapping $\boldsymbol{\phi} : \mathbb{R}_+^n \to \mathbb{R}_+^n$ is continuous in a neighborhood of $\boldsymbol{x}$, is said to be *Lyapunov stable* if $\forall \varepsilon > 0$, $\exists \delta > 0$ such that $\forall \boldsymbol{x}^{(0)} \in \mathbb{R}_+^n$, $\|\boldsymbol{x}^{(0)} - \boldsymbol{x}\| < \delta \;\Rightarrow \|\boldsymbol{x}^{(p)} - \boldsymbol{x}\| < \varepsilon \;\forall p \in \mathbb{N}$.

This property means that initializing the recursion close enough to $\boldsymbol{x}$ guarantees that the subsequent iterates remain in a given bounded domain around $\boldsymbol{x}$. However, it does not guarantee local convergence. A fixed point which is not Lyapunov stable is called *unstable*.

Definition 4 (Asymptotic stability)**.** A fixed point $\boldsymbol{x} \in \mathbb{R}_+^n$ of the recursion $\boldsymbol{x}^{(p+1)} = \boldsymbol{\phi}(\boldsymbol{x}^{(p)})$, where mapping $\boldsymbol{\phi} : \mathbb{R}_+^n \to \mathbb{R}_+^n$ is continuous in a neighborhood of $\boldsymbol{x}$, is said to be *asymptotically stable* if it is Lyapunov stable and there exists $\delta > 0$ such that $\forall \boldsymbol{x}^{(0)} \in \mathbb{R}_+^n$, $\|\boldsymbol{x}^{(0)} - \boldsymbol{x}\| < \delta \;\Rightarrow \boldsymbol{x}^{(p)} \underset{p \to +\infty}{\longrightarrow} \boldsymbol{x}$.

This property means that initializing the recursion close enough to $\boldsymbol{x}$ guarantees the convergence to $\boldsymbol{x}$. A fixed point which is Lyapunov stable, but not asymptotically stable, is sometimes called *marginally stable*.

Definition 5 (Exponential stability and rate of convergence)**.** A fixed point $\boldsymbol{x} \in \mathbb{R}_+^n$ of the recursion $\boldsymbol{x}^{(p+1)} = \boldsymbol{\phi}(\boldsymbol{x}^{(p)})$, where mapping $\boldsymbol{\phi} : \mathbb{R}_+^n \to \mathbb{R}_+^n$ is continuous in a neighborhood of $\boldsymbol{x}$, is said to be *exponentially stable* if there exists $\delta, \alpha > 0$ and $\rho \in]0, 1[$ such that $\forall \boldsymbol{x}^{(0)} \in \mathbb{R}_+^n$,

$$\|\boldsymbol{x}^{(0)} - \boldsymbol{x}\| < \delta \;\Rightarrow \|\boldsymbol{x}^{(p)} - \boldsymbol{x}\| \leq \alpha \|\boldsymbol{x}^{(0)} - \boldsymbol{x}\| \; \rho^p \;\forall p \in \mathbb{N}. \tag{11}$$

[4]However this singular case is never observed in practical NMF problems.

In this case, the minimum value of ρ such that equation (11) stands is called the *rate of convergence* of the recursion.

This property ensures a *linear* speed of convergence; it also implies asymptotic stability. A fixed point which is asymptotically stable, but not exponentially stable, is generally characterized by a *sub-linear* speed of convergence (depending on the initialization). Note that all the stability properties defined above are *local*, which means that those properties hold in a neighborhood of the fixed point $\boldsymbol{x}$.

B. Lyapunov's first (or indirect) method

Lyapunov's first (or indirect) method permits to characterize the exponential stability and the corresponding convergence rate of a dynamical system. Let us first recall its principle, that we apply in the domain $\mathbb{R}^n_+$.

Theorem 6 (Lyapunov's first stability theorem)**.** *Let $\boldsymbol{x} \in \mathbb{R}^n_+$ be a fixed point of the recursion $\boldsymbol{x}^{(p+1)} = \boldsymbol{\phi}(\boldsymbol{x}^{(p)})$, where mapping $\boldsymbol{\phi} : \mathbb{R}^n_+ \to \mathbb{R}^n_+$ is continuously differentiable in a neighborhood of $\boldsymbol{x}$. Let $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ be the Jacobian matrix[5] of mapping $\boldsymbol{\phi}$ at point $\boldsymbol{x}$. Then the exponential stability (or unstability) of $\boldsymbol{x}$ is characterized by the eigenvalues of $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$:*

- *$\boldsymbol{x}$ is an exponentially stable fixed point if and only if all the eigenvalues of $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ have a magnitude lower than 1. In this case, the rate of convergence of the recursion is equal to the spectral radius[6] of matrix $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$, which is denoted $\rho\left(\nabla\boldsymbol{\phi}^T(\boldsymbol{x})\right) < 1$.*
- *If at least one eigenvalue of $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ has a magnitude greater than 1, then $\boldsymbol{x}$ is unstable.*

Note that if all eigenvalues of $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ have a magnitude lower than or equal to 1, and at least one of them has magnitude 1, then Theorem 6 does not permit to conclude: the fixed point can be Lyapunov stable or unstable. In order to apply Theorem 6 to the mapping $\boldsymbol{\phi}$ defined in equation (10), we characterize the eigenvalues of matrix $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ in the following proposition.

Proposition 7. *Let $\boldsymbol{x} \in \mathbb{R}^n_+$ be a local minimum of an objective function J satisfying Assumption 1, and suppose that Assumption 2 holds. Let $\boldsymbol{\phi}$ be the mapping introduced in Definition 1, which is continuously differentiable in a neighborhood of $\boldsymbol{x}$. Moreover, let us define the positive semi-definite matrix*

$$\boldsymbol{P}(\boldsymbol{x}) = \boldsymbol{D}(\boldsymbol{x})\boldsymbol{\nabla}^2 J(\boldsymbol{x})\boldsymbol{D}(\boldsymbol{x}) \tag{12}$$

with $\boldsymbol{D}(\boldsymbol{x}) = \operatorname{diag}(\boldsymbol{x}/\boldsymbol{p}(\boldsymbol{x}))^{\frac{1}{2}}$ (where $\boldsymbol{p}(\boldsymbol{x})$, defined in Assumption 1, has no zero entry), and the positive scalar

$$\eta^\star = \frac{2}{\|\boldsymbol{P}(\boldsymbol{x})\|_2}, \tag{13}$$

where $\|.\|_2$ denotes the matrix 2-norm or spectral norm[7] (if $\|\boldsymbol{P}(\boldsymbol{x})\|_2 = 0$, $\eta^\star$ is infinite).

Then if $\eta = 0$, all the eigenvalues of the Jacobian matrix $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ are equal to 1. Otherwise, the multiplicity of $\lambda = 1$ as an eigenvalue of $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ is equal to the dimension of the kernel of matrix $[\boldsymbol{P}(\boldsymbol{x})]_+$ (with the use of Notation 1). Moreover, the other eigenvalues of $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ are as follows:

- *If $\eta \in]0, \eta^\star[$, all the other eigenvalues have a magnitude lower than 1;*
- *If $\eta < 0$, all the other eigenvalues are greater than 1;*
- *If $\eta > \eta^\star$, at least one eigenvalue is lower than -1;*
- *If $\eta = \eta^\star$, at least one eigenvalue is equal to -1.*

Proposition 7 is proved in Appendix B. Note that in the case $\eta = 0$, $\boldsymbol{\phi}$ is the identity transform. In other respects, if $[\boldsymbol{P}(\boldsymbol{x})]_+$ is non-singular, 1 is *not* an eigenvalue of $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$. Now we can state the stability properties of mapping (10):

Proposition 8. *Let $\boldsymbol{x} \in \mathbb{R}^n_+$ be a local minimum of an objective function J satisfying Assumption 1, and suppose that Assumption 2 holds. Let $\boldsymbol{\phi}$ be the mapping introduced in Definition 1, which is continuously differentiable in a neighborhood of $\boldsymbol{x}$.*

Then, following the notation introduced in Proposition 7,

- *$\boldsymbol{x}$ is an exponentially stable fixed point if and only if $\eta \in]0, \eta^\star[$ and matrix $[\boldsymbol{P}(\boldsymbol{x})]_+$ is non-singular;*
- *if $\eta \notin [0, \eta^\star]$, $\boldsymbol{x}$ is an unstable fixed point;*
- *if $\eta = 0$, $\boldsymbol{x}$ is a marginally stable fixed point.*

Proof: The first and second assertions are corollaries of Theorem 6 and Proposition 7. The third assertion is trivial, since if $\eta = 0$, mapping $\boldsymbol{\phi}$ is the identity transform. ■

Note that the non-singularity of matrix $[\boldsymbol{P}(\boldsymbol{x})]_+$ is equivalent to the combination of the two following properties:

$$\forall i \text{ such that } x_i = 0,\ \nabla_i J(\boldsymbol{x}) > 0, \tag{14}$$

$$\text{matrix } [\boldsymbol{\nabla}^2 J(\boldsymbol{x})]_+ \text{ is positive definite.} \tag{15}$$

If $\eta = \eta^\star$, or if $\eta \in]0, \eta^\star[$ and matrix $[\boldsymbol{P}(\boldsymbol{x})]_+$ is singular, Lyapunov's first method does not permit to conclude, since there is at least one eigenvalue of magnitude 1.

Finally, the following proposition completes Proposition 8, and proves the equivalence between the exponentially stable fixed points of mapping $\boldsymbol{\phi}$, and the local minima of function J satisfying both properties (14) and (15).

Proposition 9. *Let J be an objective function such that Assumption 1 holds. Let $\boldsymbol{x} \in \mathbb{R}^n_+$ be a vector satisfying Assumption 2. Let $\boldsymbol{\phi}$ be the mapping introduced in Definition 1, which is continuously differentiable in a neighborhood of $\boldsymbol{x}$. Assume that $\eta > 0$ and that $\boldsymbol{x}$ is an exponentially stable fixed point of mapping $\boldsymbol{\phi}$. Then $\boldsymbol{x}$ is a local minimum of function J, which additionally satisfies properties* (14) *and* (15)*.*

Proposition 9 is proved in Appendix B.

[5] For $1 \leq i, j \leq n$, the $(i,j)^{\text{th}}$ coefficient of matrix $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ is $\frac{\partial \phi_j}{\partial x_i}$.

[6] The spectral radius of a matrix is the maximum among the magnitudes of its eigenvalues.

[7] The spectral norm of a matrix is equal to its greatest singular value. In the particular case of Hermitian matrices, the spectral norm is equal to the spectral radius.

C. Lyapunov's second (direct) method

For a fixed point which is not exponentially stable, Lyapunov's second method permits to further investigate its stability properties, and possibly prove its Lyapunov or asymptotic stability. In this section, we will prove the following result, which completes that of Proposition 8: if $\eta \in]0, \eta^\star[$, even if assumption (14) does not stand, assumption (15) alone is sufficient for guaranteeing the asymptotic stability of the dynamical system. Let us first recall the principle of Lyapunov's second method, that we apply in the domain $\mathbb{R}_+^n$.

Definition 6 (Lyapunov function)**.** For any $\boldsymbol{x} \in \mathbb{R}_+^n$, a Lyapunov function $\boldsymbol{y} \mapsto V(\boldsymbol{x}, \boldsymbol{y})$ is a continuous scalar function defined on a neighborhood of $\boldsymbol{x}$ included in $\mathbb{R}_+^n$, which is positive-definite (in the sense that $V(\boldsymbol{x}, \boldsymbol{x}) = 0$, and $V(\boldsymbol{x}, \boldsymbol{y}) > 0$ for all $\boldsymbol{y} \neq \boldsymbol{x}$).

Theorem 10 (Lyapunov's second stability theorem)**.** *Let $\boldsymbol{x} \in \mathbb{R}_+^n$ be a fixed point of a continuous mapping $\boldsymbol{\phi} : \mathbb{R}_+^n \to \mathbb{R}_+^n$.*

- *If there is a Lyapunov function V such that $V(\boldsymbol{x}, \boldsymbol{\phi}(\boldsymbol{y})) \leq V(\boldsymbol{x}, \boldsymbol{y})$ for all $\boldsymbol{y}$ in a neighborhood of $\boldsymbol{x}$, then $\boldsymbol{x}$ is Lyapunov stable.*
- *If there is a Lyapunov function V such that $V(\boldsymbol{x}, \boldsymbol{\phi}(\boldsymbol{y})) < V(\boldsymbol{x}, \boldsymbol{y})$ for all $\boldsymbol{y} \neq \boldsymbol{x}$ in a neighborhood of $\boldsymbol{x}$, then $\boldsymbol{x}$ is asymptotically stable.*

If $\boldsymbol{x}$ is a local minimum of a continuous function J, a natural candidate Lyapunov function for the mapping $\boldsymbol{\phi}$ defined in Definition 1 would be $V(\boldsymbol{x}, \boldsymbol{y}) = J(\boldsymbol{y}) - J(\boldsymbol{x})$. However, this choice raises two problems:

- In some cases, a fixed-point $\boldsymbol{x}$ is Lyapunov-stable, whereas the objective function J is not globally monotonically decreasing[8].
- The condition that V is positive-definite may not be satisfied in a neighborhood of $\boldsymbol{x}$ [9].

For these reasons, we propose an alternative Lyapunov function in the following lemma.

Lemma 11. *Consider an objective function J satisfying Assumption 1, and $\boldsymbol{x} \in \mathbb{R}_+^n$ such that Assumption 2 holds. Then the function*

$$V(\boldsymbol{x}, \boldsymbol{y}) = \frac{1}{2}(\boldsymbol{y} - \boldsymbol{x})^T \operatorname{diag}\left(\frac{\boldsymbol{p}(\boldsymbol{x}) + \boldsymbol{p}(\boldsymbol{y})}{\boldsymbol{x} + \boldsymbol{y}}\right)(\boldsymbol{y} - \boldsymbol{x}) \quad (16)$$

defines a symmetric Lyapunov function on $\mathbb{R}_+^n \times \mathbb{R}_+^n$.

Proof: The definition and continuity of function V on the borders of the domain $\mathbb{R}_+^n$ follow from the inequality $\left|\frac{y_i - x_i}{x_i + y_i}\right| \leq 1 \ \forall x_i, y_i \in \mathbb{R}_+^\star$. ∎

We can now state our main result:

Proposition 12. *Let $\boldsymbol{x} \in \mathbb{R}_+^n$ be a local minimum of a function J satisfying Assumption 1, and suppose that Assumption 2 and property* (15) *hold. If $\eta \in]0, \eta^\star[$ (where $\eta^\star$ was defined in equation* (13)*), then the mapping $\boldsymbol{\phi}$ introduced in Definition 1 makes the Lyapunov function $\boldsymbol{y} \mapsto V(\boldsymbol{x}, \boldsymbol{y})$ defined in equation* (16) *strictly decrease in a neighborhood of $\boldsymbol{x}$. As a consequence, $\boldsymbol{x}$ is an asymptotically stable fixed point of mapping $\boldsymbol{\phi}$.*

Proposition 12 is proved in Appendix C. Considering Proposition 8, it can be noticed that if all hypotheses in Proposition 12 stand, but property (14) is not satisfied, then $\boldsymbol{x}$ is an asymptotically stable, but not exponentially stable fixed point of mapping $\boldsymbol{\phi}$. This means that, depending on the initialization, the dynamical system generally has a sub-linear speed of convergence.

IV. Application to NMF

In this section, we show how Lyapunov's stability theory can be applied to the particular problem of NMF with an objective function based on the β-divergence, which was introduced in section II-A. We first focus on the simple case of supervised NMF, which is a direct application of the theory presented in section III, and then study the more complex case of unsupervised NMF. Due to the lack of space, the complete proofs have been moved to a separate document [1]. Nevertheless, the key ideas are provided whenever possible in the following discussion.

Notation 2. In the following, the entries of the $F \times K$ matrix $\boldsymbol{W}$ and the $K \times T$ matrix $\boldsymbol{H}$ are remapped into vectors $\boldsymbol{w}$ and $\boldsymbol{h}$ of dimensions KF and KT, respectively. Vector $\boldsymbol{x}$ is formed by concatenating $\boldsymbol{w}$ and $\boldsymbol{h}$. Then let us define the objective function

$$J(\boldsymbol{x}) = D(\boldsymbol{V} | \boldsymbol{W}\boldsymbol{H}) \quad (17)$$

where function D was defined in equation (1), and involves the β-divergence (2). The gradient of J w.r.t. $\boldsymbol{w}$ is decomposed as the difference of two non-negative functions $\boldsymbol{p}^w(\boldsymbol{x})$ and $\boldsymbol{m}^w(\boldsymbol{x})$ whose coefficients have been defined in equation (5), and similar notation is used for $\boldsymbol{h}$ (see equation (32)). Vector $\boldsymbol{p}(\boldsymbol{x})$ is formed by concatenating $\boldsymbol{p}^w(\boldsymbol{x})$ and $\boldsymbol{p}^h(\boldsymbol{x})$. The Hessian matrices of function J w.r.t. $\boldsymbol{x}$, $\boldsymbol{w}$ and $\boldsymbol{h}$ are denoted $\boldsymbol{\nabla}^2_{xx} J(\boldsymbol{x})$, $\boldsymbol{\nabla}^2_{ww} J(\boldsymbol{x})$ and $\boldsymbol{\nabla}^2_{hh} J(\boldsymbol{x})$, respectively. Let

$$\boldsymbol{\phi}^w(\boldsymbol{w}, \boldsymbol{h}) = \boldsymbol{\Lambda}_w(\boldsymbol{x})^\eta \, \boldsymbol{w} \quad (18)$$

$$\boldsymbol{\phi}^h(\boldsymbol{w}, \boldsymbol{h}) = \boldsymbol{\Lambda}_h(\boldsymbol{x})^\eta \, \boldsymbol{h} \quad (19)$$

where $\boldsymbol{\Lambda}_w(\boldsymbol{x}) = \operatorname{diag}\left(\frac{\boldsymbol{m}^w(\boldsymbol{x})}{\boldsymbol{p}^w(\boldsymbol{x})}\right)$ and $\boldsymbol{\Lambda}_h(\boldsymbol{x}) = \operatorname{diag}\left(\frac{\boldsymbol{m}^h(\boldsymbol{x})}{\boldsymbol{p}^h(\boldsymbol{x})}\right)$.

A. Application to supervised NMF

In the context of supervised NMF, vector $\boldsymbol{w}$ is kept unchanged, and vector $\boldsymbol{h}$ only is updated according to $\boldsymbol{h}^{(p+1)} = \boldsymbol{\phi}^h(\boldsymbol{w}, \boldsymbol{h}^{(p)})$, where mapping $\boldsymbol{\phi}^h$ was defined in equation (19), which is equivalent to the multiplicative update (8). The parameter $\eta_h^\star$ introduced in the following lemma will play the same role as $\eta^\star$ in sections III-B and III-C.

Lemma 13. *Given a constant vector $\boldsymbol{w}$, let $\boldsymbol{h}$ be a local minimum of the NMF objective function $\boldsymbol{h} \mapsto J(\boldsymbol{w}, \boldsymbol{h})$ defined in equation* (17)*. Function $\boldsymbol{h} \mapsto J(\boldsymbol{w}, \boldsymbol{h})$ satisfies Assumption 1, and we assume that $\boldsymbol{h}$ satisfies Assumption 2. Following Notation 2, let us define*

$$\eta_h^\star = \frac{2}{\|\boldsymbol{P}^h(\boldsymbol{x})\|_2}, \quad (20)$$

[8] This is the case of NMF with $1 < \eta < 2$.

[9] In the case of unsupervised NMF, because of the invariances of the factorization, there is a continuous set of fixed points $\boldsymbol{y}$ satisfying $J(\boldsymbol{y}) = J(\boldsymbol{x})$ (cf. section IV).

where $\boldsymbol{P}^h(\boldsymbol{x})$ *is the positive semi-definite matrix*

$$\boldsymbol{P}^h(\boldsymbol{x}) = \boldsymbol{D}^h(\boldsymbol{x})\boldsymbol{\nabla}^2_{hh}J(\boldsymbol{x})\boldsymbol{D}^h(\boldsymbol{x}). \tag{21}$$

with

$$\boldsymbol{D}^h(\boldsymbol{x}) = \mathrm{diag}(\boldsymbol{h}/\boldsymbol{p}^h(\boldsymbol{x}))^{\frac{1}{2}}. \tag{22}$$

Then $\forall \beta \in \mathbb{R}$*, we have* $0 < \eta^\star_h \leq 2$*, and if* $\beta \in [1,2]$*,* $\eta^\star_h = 2$*.*

Proof: This lemma is proved by exhibiting an eigenvector of the positive semi-definite matrix $\boldsymbol{P}^h(\boldsymbol{x})$, whose eigenvalue is equal to 1. If, additionally, $\beta \in [1,2]$, the convexity of function J w.r.t. $\boldsymbol{h}$ permits to prove that all the eigenvalues of $\boldsymbol{P}^h(\boldsymbol{x})$ are lower than or equal to 1. ■

Following Lemma 13, Propositions 8 and 12 directly prove the exponential or the asymptotic stability of the local minima of function $\boldsymbol{h} \mapsto J(\boldsymbol{w},\boldsymbol{h})$ for all $\eta \in]0, \eta^\star_h[$, under mild conditions (numerical examples are presented in section V-A). Of course, the same result stands for the reciprocal algorithm, where $\boldsymbol{h}$ is kept unchanged, and vector $\boldsymbol{w}$ only is updated according to $\boldsymbol{w}^{(p+1)} = \boldsymbol{\phi}^w(\boldsymbol{w}^{(p)},\boldsymbol{h})$, where mapping $\boldsymbol{\phi}^w$ was defined in equation (18), which is equivalent to the multiplicative update (7).

Lemma 14. *Given a constant vector* $\boldsymbol{h}$*, let* $\boldsymbol{w}$ *be a local minimum of the NMF objective function* $\boldsymbol{w} \mapsto J(\boldsymbol{w},\boldsymbol{h})$ *defined in equation* (17)*. Function* $\boldsymbol{w} \mapsto J(\boldsymbol{w},\boldsymbol{h})$ *satisfies Assumption 1, and we assume that* $\boldsymbol{w}$ *satisfies Assumption 2. Following Notation 2, let us define*

$$\eta^\star_w = \frac{2}{\|\boldsymbol{P}^w(\boldsymbol{x})\|_2}, \tag{23}$$

where $\boldsymbol{P}^w(\boldsymbol{x})$ *is the positive semi-definite matrix*

$$\boldsymbol{P}^w(\boldsymbol{x}) = \boldsymbol{D}^w(\boldsymbol{x})\boldsymbol{\nabla}^2_{ww}J(\boldsymbol{x})\boldsymbol{D}^w(\boldsymbol{x}); \tag{24}$$

with

$$\boldsymbol{D}^w(\boldsymbol{x}) = \mathrm{diag}(\boldsymbol{w}/\boldsymbol{p}^w(\boldsymbol{x}))^{\frac{1}{2}}; \tag{25}$$

Then $\forall \beta \in \mathbb{R}$*, we have* $0 < \eta^\star_w \leq 2$*, and if* $\beta \in [1,2]$*,* $\eta^\star_w = 2$*.*

This lemma is proved in the same way as lemma 13.

B. Application to unsupervised NMF

Actually, analyzing the stability of the algorithm which alternates multiplicative updates (7) and (8) is particularly difficult for the following reasons:

- It is well known that unsupervised NMF admits several invariances (the problem of the uniqueness of unsupervised NMF has been addressed in [23] for instance). Indeed, the product $\boldsymbol{WH}$ is unchanged by replacing matrices $\boldsymbol{W}$ and $\boldsymbol{H}$ by the non-negative matrices $\boldsymbol{W}' = \boldsymbol{WD}$ and $\boldsymbol{H}' = \boldsymbol{D}^{-1}\boldsymbol{H}$, where $\boldsymbol{D}$ is any diagonal matrix with positive diagonal coefficients. For this simple reason, the *local minima of the objective function are never isolated* (any local minimum is reached on a *continuum* of matrices $\boldsymbol{W}'$ and $\boldsymbol{H}'$ whose product is equal to $\boldsymbol{WH}$). The consequence is that assumption (15) never stands.
- Anyway, recursion (7)-(8) cannot be implemented with a mapping of the form (10), since it switches between updates for $\boldsymbol{W}$ and $\boldsymbol{H}$.

For these reasons, the results presented in section III cannot be straightforwardly applied to recursion (7)-(8), and the local minima of the objective function can never be exponentially stable (the Jacobian matrix always admits $\lambda = 1$ as an eigenvalue). Thus Lyapunov's first method will not permit to conclude on the stability of multiplicative updates. Nevertheless, this approach still provides an interesting insight into the stability properties of the algorithm. We summarize below the main results that we obtained by applying this approach.

Recursion (7)-(8) is rewritten in the form

$$\left\{ \begin{array}{rcl} \boldsymbol{w}^{(p+1)} & = & \boldsymbol{\phi}^w(\boldsymbol{w}^{(p)},\boldsymbol{h}^{(p)}) \\ \boldsymbol{h}^{(p+1)} & = & \boldsymbol{\phi}^h(\boldsymbol{w}^{(p+1)},\boldsymbol{h}^{(p)}) \end{array} \right. \tag{26}$$

where mappings $\boldsymbol{\phi}^w$ and $\boldsymbol{\phi}^h$ were defined in equations (18) and (19). Equivalently, we can write $\boldsymbol{x}^{(p+1)} = \boldsymbol{\phi}(\boldsymbol{x}^{(p)})$, with

$$\boldsymbol{\phi}(\boldsymbol{x}) = \left[\boldsymbol{\phi}^w(\boldsymbol{w},\boldsymbol{h}); \boldsymbol{\phi}^h(\boldsymbol{\phi}^w(\boldsymbol{w},\boldsymbol{h}),\boldsymbol{h})\right]. \tag{27}$$

The parameter $\eta^\star_x$ introduced in the following lemma will play the same role as $\eta^\star$ in sections III-B and III-C.

Lemma 15. *Let* $\boldsymbol{x}$ *be a local minimum of the NMF objective function* J *defined in equation* (17)*. Function* J *satisfies Assumption 1, and we assume that* $\boldsymbol{x}$ *satisfies Assumption 2. Let us define*

$$\eta^\star_x = \min\left(\eta^\star_h, \eta^\star_w\right), \tag{28}$$

where $\eta^\star_h$ *and* $\eta^\star_w$ *were defined in equations* (20) *and* (23)*. Then* $\forall \beta \in \mathbb{R}$*, we have* $0 < \eta^\star_x \leq 2$*, and if* $\beta \in [1,2]$*,* $\eta^\star_x = 2$*.*

This lemma is a corollary of lemmas 13 and 14. Note that the definition of $\eta^\star_x$ does not follow the same scheme as the definitions of $\eta^\star_h$ and $\eta^\star_w$ in equations (20) and (23). Indeed in Proposition 16, which characterizes the eigenvalues of the Jacobian matrix $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$, the upper bound $\eta^\star_x$ is not equal to $\frac{2}{\|\boldsymbol{P}(\boldsymbol{x})\|_2}$ (with $\boldsymbol{P}(\boldsymbol{x})$ defined in equation (29)).

Proposition 16. *Let* $\boldsymbol{x}$ *be a local minimum of function* J *defined in equation* (17)*. Function* J *satisfies Assumption 1, and we assume that* $\boldsymbol{x}$ *satisfies Assumption 2. Consider the mapping* $\boldsymbol{\phi}$ *defined in equation* (27)*, which is continuously differentiable in a neighborhood of* $\boldsymbol{x}$*.*

Then if $\eta = 0$*, all the eigenvalues of the Jacobian matrix* $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ *are equal to* 1*. Otherwise,* $\lambda = 1$ *is always an eigenvalue of* $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$*. Following Notations 1 and 2, its multiplicity is greater than or equal to the dimension of the kernel of matrix* $[\boldsymbol{P}(\boldsymbol{x})]_+$*, where*

$$\boldsymbol{P}(\boldsymbol{x}) = \boldsymbol{D}(\boldsymbol{x})\boldsymbol{\nabla}^2_{xx}J(\boldsymbol{x})\boldsymbol{D}(\boldsymbol{x}) \tag{29}$$

and $\boldsymbol{D}(\boldsymbol{x}) = \mathrm{diag}(\boldsymbol{x}/\boldsymbol{p}(\boldsymbol{x}))^{\frac{1}{2}}$*. Moreover, the other eigenvalues of* $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ *are as follows:*

- *If* $\eta \notin [0,2]$*, there is at least one eigenvalue greater than* 1*;*
- *If* $\eta \in]0, \eta^\star_x[$ *(where* $\eta^\star_x$ *was defined in equation* (28)*), the multiplicity of the eigenvalue 1 is equal to the dimension*

of the kernel of matrix $[\boldsymbol{P}(\boldsymbol{x})]_+$*, and all other eigenvalues have a magnitude lower than* 1.

Proof: The proof of Proposition 16 follows the same outline as the proof of Proposition 7. It additionally relies on the observation that $\lambda = 1$ and $\lambda = (1-\eta)^2$ are always eigenvalues of the Jacobian matrix $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ (which is proved by exhibiting the corresponding eigenvectors). ■

As mentioned above, Proposition 16 does not permit to conclude on the stability of multiplicative updates. Nevertheless, it enables a formal proof of the first and second assertions in the following proposition (the third one is trivial):

Proposition 17. *Let* $\boldsymbol{x}$ *be a local minimum of function* J *defined in equation* (17)*. Function* J *satisfies Assumption 1, and we assume that* $\boldsymbol{x}$ *satisfies Assumption 2. Consider the mapping* $\boldsymbol{\phi}$ *defined in equation* (27)*, which is continuously differentiable in a neighborhood of* $\boldsymbol{x}$*.*

Then

- $\forall\eta \in \mathbb{R}$*,* $\boldsymbol{x}$ *is not an exponentially stable fixed point;*
- *if* $\eta \notin [0,2]$*,* $\boldsymbol{x}$ *is an unstable fixed point;*
- *if* $\eta = 0$*,* $\boldsymbol{x}$ *is a marginally stable fixed point.*

Proposition 17 does not tell us what happens if $\eta \in]0,2]$. This is because 1 is always an eigenvalue of the Jacobian matrix $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$, as shown in Proposition 16. If $\eta \in]0, \eta_x^\star[$, its multiplicity is equal to the dimension of the kernel of matrix $[\boldsymbol{P}(\boldsymbol{x})]_+$, which we suppose accounts for the invariances of the factorization[10].

V. Simulation results

In this section we propose some numerical simulations which illustrate the theoretical results presented in sections III and IV. We first focus on supervised NMF, and then investigate the case of unsupervised NMF.

A. Supervised NMF

First, we study the stability of the multiplicative update (8) applied to matrix $\boldsymbol{H}$, while keeping matrix $\boldsymbol{W}$ unchanged.

1) Example of sub-linear convergence speed: In this first experiment, the dimensions are $F = 3$, $T = 3$ and $K = 2$. The multiplicative update (8) is applied to the Kullback-Leibler divergence ($\beta = 1$) with a step size $\eta = 1$ (which corresponds to the standard multiplicative update). The matrix $\boldsymbol{V}$ to be decomposed is defined as a square non-negative Hankel matrix: $\boldsymbol{V} = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 4 \\ 3 & 4 & 5 \end{bmatrix}$, and the matrix $\boldsymbol{W}$ is defined as

$$\boldsymbol{W} = \begin{bmatrix} 1 & 1 \\ 2 & 1 \\ 3 & 1 \end{bmatrix}. \tag{30}$$

[10]Actually the invariances of the factorization imply some conditions on the first and second order derivatives of the objective function J. We managed to prove that if all the hypotheses in Proposition 17 and property (14) are satisfied, then the dimensionality of such conditions is equal to the dimension of the kernel of matrix $[\boldsymbol{P}(\boldsymbol{x})]_+$.

It can be noticed that $\boldsymbol{V}$ is singular, and that it can be exactly factorized as the product $\boldsymbol{V} = \boldsymbol{W}\boldsymbol{H}$, where

$$\boldsymbol{H} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \end{bmatrix}. \tag{31}$$

Thus we readily know that the lowest value of $D(\boldsymbol{V}|\boldsymbol{W}\boldsymbol{H})$ w.r.t. $\boldsymbol{H}$ is 0, and that this global minimum is reached for $\boldsymbol{H}$ defined in equation (31). This particular example was chosen so that $\widehat{\boldsymbol{V}} = \boldsymbol{V}$, thus $\boldsymbol{p}^h = \boldsymbol{m}^h$, as can be noticed in equation (32). Consequently, property (14) does not stand, since $h_{21} = 0$ and $\frac{\partial J}{\partial h_{21}} = 0$. Therefore Proposition 8, which would have proven the exponential stability of the global minimum, cannot be applied. Nevertheless all the hypotheses in Proposition 12 and Lemma 13 are satisfied, which proves the asymptotic stability. Thus the speed of convergence of the multiplicative update (8) may be sub-linear, which will be confirmed by the following simulation results.

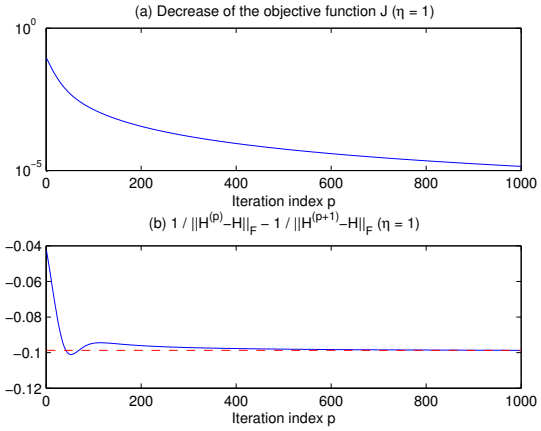

Fig. 1. Example of sub-linear convergence speed

Fig. 1 shows the results obtained by initializing (8) with a matrix $\boldsymbol{H}$ having all coefficients equal to 2. As can be noticed in Fig. 1-(a), the objective function J monotonically converges to 0 (its global decrease was proven in Proposition 1). Besides, Fig. 1-(b) represents the sequence $\frac{1}{\|\boldsymbol{H}^{(p)}-\boldsymbol{H}\|_F} - \frac{1}{\|\boldsymbol{H}^{(p+1)}-\boldsymbol{H}\|_F}$ (where $\|.\|_F$ denotes the Frobenius norm, $\boldsymbol{H}^{(p)}$ is the matrix computed at iteration p and $\boldsymbol{H}$ is the matrix defined in equation (31)) as a solid blue line. It can be noticed that this sequence converges to a finite negative value (represented by the dashed red line), which shows that $\|\boldsymbol{H}^{(p)} - \boldsymbol{H}\|_F = O(1/p)$. As predicted by the theoretical analysis, the convergence speed happens to be sub-linear (at least for the proposed initialization).

2) Example of linear convergence speed: In this second experiment, all variables are defined as in section V-A1, except that the top left coefficient of $\boldsymbol{V}$ is replaced by 0.9. Consequently, this matrix is no longer singular, thus the global minimum of the objective function w.r.t. $\boldsymbol{H}$ cannot be zero. Instead, a local (possibly global) minimum w.r.t. $\boldsymbol{H}$ can be computed by means of multiplicative update (8), initialized as

in section V-A1 [11]. Numerically, we observed that the local minimum $\boldsymbol{x}$ is still such that $h_{21} = 0$, but $\frac{\partial J}{\partial h_{21}} > 0$, thus property (14) now stands, and Proposition 8 and Lemma 13 prove the exponential stability of this local minimum, with a convergence rate equal to the spectral radius $\rho\left(\boldsymbol{\nabla}_h \boldsymbol{\phi}^{h^T}(\boldsymbol{x})\right)$, which will be confirmed by the following simulation results.

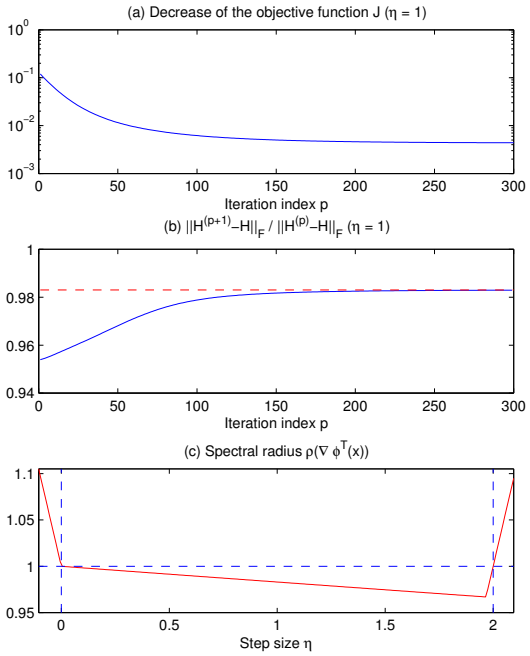

Fig. 2. Example of linear convergence speed

Fig. 2-(a) shows that the objective function J is monotonically decreasing. Besides, Fig. 2-(b) represents the sequence $\frac{\|\boldsymbol{H}^{(p+1)}-\boldsymbol{H}\|_F}{\|\boldsymbol{H}^{(p)}-\boldsymbol{H}\|_F}$ as a solid blue line, and the value $\rho\left(\boldsymbol{\nabla}_h \boldsymbol{\phi}^{h^T}(\boldsymbol{x})\right)$ as a dashed red line[12]. It can be noticed that this sequence converges to $\rho\left(\boldsymbol{\nabla}_h \boldsymbol{\phi}^{h^T}(\boldsymbol{x})\right)$, which shows that $\|\boldsymbol{H}^{(p)} - \boldsymbol{H}\|_F = O\left(\rho\left(\boldsymbol{\nabla}_h \boldsymbol{\phi}^{h^T}(\boldsymbol{x})\right)^p\right)$. As predicted by the theoretical analysis, the convergence speed is linear, with a convergence rate equal to $\rho\left(\boldsymbol{\nabla}_h \boldsymbol{\phi}^{h^T}(\boldsymbol{x})\right)$.

3) Optimal step size: In this third experiment, all variables are defined as in section V-A2, and we are looking for an optimal step size η. Since $\beta = 1$, Proposition 8 and Lemma 13 prove that the local minimum is exponentially stable if and only if $0 < \eta < 2$. In Fig. 2-(c), the solid red line presents the spectral radius $\rho\left(\boldsymbol{\nabla}_h \boldsymbol{\phi}^{h^T}(\boldsymbol{x})\right)$ as a function of η, for all $\eta \in]-0.1, 2.1[$. This simulation result confirms that $\rho\left(\boldsymbol{\nabla}_h \boldsymbol{\phi}^{h^T}(\boldsymbol{x})\right) < 1$ if and only if $0 < \eta < 2$, and it shows that there is an optimal value of parameter η, for which the rate of convergence is optimal. In particular, we note that the standard step size $\eta = 1$ is not optimal. Besides, we observed that a value of η outside the range $[0, 2]$ results in a divergence of the objective function J (for $\eta = 0$ the objective function is constant, and for $\eta = 2$ it oscillates between two values).

[11]Note that in this second experiment, the value of matrix $\boldsymbol{W}$ defined in equation (30), which remains unchanged throughout the iterations, does no longer correspond to a local minimum of the objective function w.r.t. $\boldsymbol{W}$.

[12]The spectral radius $\rho\left(\boldsymbol{\nabla}_h \boldsymbol{\phi}^{h^T}(\boldsymbol{x})\right)$ was computed from the closed form expression of matrix $\boldsymbol{\nabla}_h \boldsymbol{\phi}^{h^T}(\boldsymbol{x})$ presented in equation (52) in the supporting document [1].

B. Unsupervised NMF

We now study the case of unsupervised NMF, which alternates multiplicative updates (7) and (8) for $\boldsymbol{W}$ and $\boldsymbol{H}$. In this fourth experiment, all variables are defined as in section V-A2, and the algorithm is initialized in the same way.

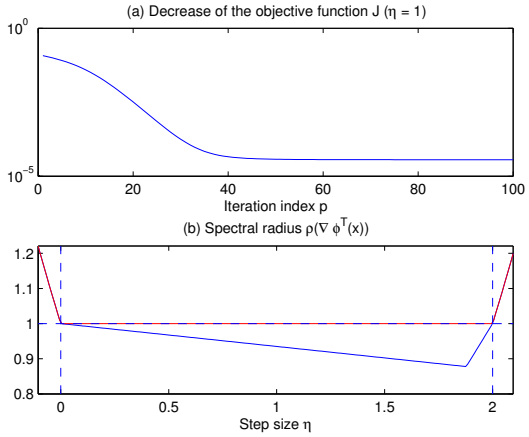

Fig. 3. Unsupervised NMF

Fig. 3-(a) shows that the objective function J is monotonically decreasing (to a non-zero value). As in section V-A3, the solid red line in Fig. 3-(b) represents the spectral radius $\rho\left(\boldsymbol{\nabla}\boldsymbol{\phi}^T(\boldsymbol{x})\right)$ as a function of the step size η, for all $\eta \in]-0.1, 2.1[$ [13]. As proved in Lemma 15 and Proposition 16, $\rho\left(\boldsymbol{\nabla}\boldsymbol{\phi}^T(\boldsymbol{x})\right) > 1$ if $\eta \notin [0, 2]$, and $\rho\left(\boldsymbol{\nabla}\boldsymbol{\phi}^T(\boldsymbol{x})\right) = 1$ in the range $\eta =]0, 2[$, which confirms that the local minimum is not exponentially stable, as stated in Proposition 17. Finally, the solid blue line represents the maximum among the magnitudes of the eigenvalues of matrix $\boldsymbol{\nabla}\boldsymbol{\phi}^T(\boldsymbol{x})$ *which are different from 1* [14]. This suggests an optimal value $\eta \approx 1.875$, which is again different from the standard step size $\eta = 1$. Indeed it can be verified that the lowest value of the objective function J (after 100 iterations) is reached when the algorithm is run with this optimal value of η.

VI. Conclusions

In this paper, we analyzed the convergence properties of general multiplicative update algorithms, designed to solve optimization problems with non-negativity constraints, where

[13]The spectral radius $\rho\left(\boldsymbol{\nabla}\boldsymbol{\phi}^T(\boldsymbol{x})\right)$ was computed from the closed form expression of matrix $\boldsymbol{\nabla}\boldsymbol{\phi}^T(\boldsymbol{x})$ presented in equation (55) in the supporting document [1].

[14]This maximum value discards the eigenvalues equal to 1, which are due to the invariances of the factorization.

we introduced an exponent step size η. We have applied Lyapunov's first and second methods to find some criteria which guarantee the exponential or asymptotic stability of the local minima of the objective function, either by analyzing the eigenvalues of the Jacobian matrix of the mapping, or by introducing a suitable Lyapunov function. We noticed that exponential stability ensures a linear convergence speed, whereas asymptotic stability generally leads to a sub-linear convergence speed (depending on the initialization). In both cases, we provided the closed-form expression of an upper bound $\eta^\star$ such that the exponential or asymptotic stability is guaranteed for $\eta \in]0, \eta^\star[$. This study was straightforwardly applied to supervised NMF algorithms based on β-divergences, which update all variables at once, instead of switching between two multiplicative updates. We have thus presented some criteria which guarantee the exponential or asymptotic stability of NMF multiplicative updates, and proved that the upper bound is such that $\forall \beta \in \mathbb{R}$, $\eta^\star \in]0, 2]$, and if $\beta \in [1, 2]$, $\eta^\star = 2$ (note that in this last case, Lee and Sung's multiplicative updates correspond to the particular case $\eta = 1$). We then applied the same methodology to study the more complex case of unsupervised NMF. Analyzing the stability properties of the multiplicative updates happened to be particularly difficult because of the invariances of the factorization, which make the local minima of the objective function non-isolated, thus non-asymptotically stable. Nevertheless Lyapunov's first method has provided some interesting insights into the convergence properties of those updates. In particular, we proved their unstability if $\eta \notin [0, 2]$. Finally, the theoretical results presented in the paper were confirmed by numerical simulations involving both supervised and unsupervised NMF. Those simulations showed that the convergence rate depends on the value of η, and that there exists an optimal value of η (generally different from 1), which provides the fastest convergence rate (thus faster than that of Lee and Sung's multiplicative updates). Finally, this study can also be applied to non-negative tensor factorization *via* unfolding [11].

A possible extension of this stability analysis would focus on the hybrid case of constrained unsupervised NMF. If the constraint is expressed *via* a modification of the underlying model like in [13], [31], [33], the convergence analysis may be similar to that of unconstrained NMF. If the constraint is expressed *via* additional penaly terms in the objective function like in [7], [26], [34], [36], the convergence analysis may paradoxically be simpler than in unconstrained NMF: the algorithm still switches between two multiplicative updates, but the modified objective function may have isolated, thus asymptotically stable, local minima. In other respects, an algorithmic outlook of this work would be the design of multiplicative update algorithms with an optimal or an adaptive exponent step size. Similarly, since functions $\boldsymbol{p}$ and $\boldsymbol{m}$ are not unique, it would be interesting to investigate the impact of the choice of $\boldsymbol{p}$ and $\boldsymbol{m}$ onto the convergence rate.

APPENDIX

A. Multiplicative update algorithms for NMF

Proof of Proposition 1: We remark that $d_\beta(x|y)$ defined in (2) can be written in the form $d_\beta(x|y) = x^\beta \delta_\beta(\frac{y}{x})$, where

$$\begin{aligned}\delta_\beta(u) &= \tfrac{1}{\beta(\beta-1)}(1 + (\beta-1)u^\beta - \beta u^{\beta-1}),\\ \tfrac{d\delta_\beta}{du} &= u^{\beta-1} - u^{\beta-2},\\ \tfrac{d^2\delta_\beta}{du^2} &= (\beta-1)u^{\beta-2} + (2-\beta)u^{\beta-3}.\end{aligned}$$

Function δ_β is strictly convex in $\mathbb{R}_+$ if and only if $\forall u > 0$, $\frac{d^2\delta_\beta}{du^2} > 0$, which is equivalent to $1 \leq \beta \leq 2$. Then we write

$$D(\boldsymbol{V}|\boldsymbol{W}\boldsymbol{H}) = \sum_{f=1}^{F}\sum_{t=1}^{T} v_{ft}^\beta \delta_\beta\left(\frac{\hat{v}_{ft}}{v_{ft}}\right)$$

where $\hat{v}_{ft}$ was defined in equation (6). Because of the strict convexity of δ_β, we have $\forall \boldsymbol{h}'_t \in \mathbb{R}_+^K$

$$\begin{aligned}\delta_\beta\left(\frac{\sum_{k=1}^{K} w_{fk}h'_{kt}}{v_{ft}}\right) &= \delta_\beta\left(\sum_{k=1}^{K} \tfrac{w_{fk}h_{kt}}{\hat{v}_{ft}} \times \tfrac{\hat{v}_{ft}h'_{kt}}{v_{ft}h_{kt}}\right)\\ &\leq \sum_{k=1}^{K} \tfrac{w_{fk}h_{kt}}{\hat{v}_{ft}}\delta_\beta\left(\tfrac{\hat{v}_{ft}h'_{kt}}{v_{ft}h_{kt}}\right)\end{aligned}$$

with equality if and only if $\forall k \in \{1 \dots K\}$, $h'_{kt} = h_{kt}$.

Thus $\forall \boldsymbol{H}' \in \mathbb{R}_+^{K\times T}$, $D(\boldsymbol{V}|\boldsymbol{W}\boldsymbol{H}') \leq G^{W,H}(\boldsymbol{H}')$, where $G^{W,H}(\boldsymbol{H}') = \sum_{k=1}^{K}\sum_{t=1}^{T} G_{kt}^{W,H}(h'_{kt})$ and $G_{kt}^{W,H}(h'_{kt}) = \sum_{f=1}^{F} \frac{w_{fk}h_{kt}}{\hat{v}_{ft}} v_{ft}^\beta \delta_\beta\left(\frac{\hat{v}_{ft}}{v_{ft}}\frac{h'_{kt}}{h_{kt}}\right)$ with equality if and only if $\boldsymbol{H}' = \boldsymbol{H}$. Then let $h'_{kt}(\eta) = h_{kt}\left(\frac{m^h_{kt}}{p^h_{kt}}\right)^\eta$ where p^h_{kt} and m^h_{kt} are defined similarly to p^w_{fk} and m^w_{fk} in equation (5):

$$\begin{aligned}p^h_{kt} &= \sum_{f=1}^{F} w_{fk}\hat{v}_{ft}^{\beta-1}\\ m^h_{kt} &= \sum_{f=1}^{F} w_{fk}v_{ft}\hat{v}_{ft}^{\beta-2}\end{aligned} \tag{32}$$

(here we assume $p^h_{kt} \neq 0$). Define the function $F_{kt}(\eta) = G_{kt}^{W,H}(h'_{kt}(\eta))$. Then $\frac{dF_{kt}}{d\eta} = \frac{dG_{kt}^{W,H}}{dh'_{kt}}\frac{dh'_{kt}}{d\eta} = h_{kt}p^h_{kt}\ln\left(\frac{m^h_{kt}}{p^h_{kt}}\right)\left(\frac{m^h_{kt}}{p^h_{kt}}\right)^{\eta\beta}\left[1 - \left(\frac{m^h_{kt}}{p^h_{kt}}\right)^{1-\eta}\right]$ (where we assume $m^h_{kt} \neq 0$). If $h_{kt} \neq 0$ and $m^h_{kt} \neq p^h_{kt}$, then

- $\frac{dF_{kt}}{d\eta} < 0$ for all $\eta \in]-\infty, 1[$,
- $\frac{dF_{kt}}{d\eta} = 0$ for $\eta = 1$,
- $\frac{dF_{kt}}{d\eta} > 0$ for all $\eta \in]1, +\infty[$.

In particular, $\forall \eta \in]0, 1]$, $F_{kt}(\eta) < F_{kt}(0)$. Finally, let $F(\eta) = \sum_{k=1}^{K}\sum_{t=1}^{T} F_{kt}(\eta) = G^{W,H}(\boldsymbol{H}'(\eta))$. If $\boldsymbol{H}'(\eta) \neq \boldsymbol{H}$, then $\forall \eta \in]0, 1]$, $F(\eta) < F(0)$. Consequently, $\forall \eta \in]0, 1]$, $D(\boldsymbol{V}|\boldsymbol{W}\boldsymbol{H}'(\eta)) \leq F(\eta) < F(0) = D(\boldsymbol{V}|\boldsymbol{W}\boldsymbol{H})$. Thus $\boldsymbol{H}'(\eta) \neq \boldsymbol{H} \Rightarrow D(\boldsymbol{V}|\boldsymbol{W}\boldsymbol{H}'(\eta)) < D(\boldsymbol{V}|\boldsymbol{W}\boldsymbol{H})$.

The same proof can be applied to the update of $\boldsymbol{W}$. ■

B. Lyapunov's first method

Proof of Proposition 7: By differentiating equation (10), we obtain the expression of the Jacobian matrix $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$:

$$\begin{aligned}&\nabla\boldsymbol{\phi}^T(\boldsymbol{x}) = \boldsymbol{\Lambda}(\boldsymbol{x})^\eta +\\ &\eta\left(\nabla\boldsymbol{m}^T(\boldsymbol{x})\mathrm{diag}(\tfrac{1}{\boldsymbol{m}(\boldsymbol{x})}) - \nabla\boldsymbol{p}^T(\boldsymbol{x})\mathrm{diag}(\tfrac{1}{\boldsymbol{p}(\boldsymbol{x})})\right)\mathrm{diag}(\boldsymbol{\phi}(\boldsymbol{x})).\end{aligned} \tag{33}$$

If $\boldsymbol{x}$ is a local minimum of the objective function J, Lemma 5 proves that it is a fixed point of $\boldsymbol{\phi}$, thus equation (33) yields

$$\nabla\boldsymbol{\phi}^T(\boldsymbol{x}) = \boldsymbol{\Lambda}(\boldsymbol{x})^\eta - \eta\boldsymbol{\nabla}^2 J(\boldsymbol{x})\,\mathrm{diag}(\boldsymbol{x}/\boldsymbol{p}(\boldsymbol{x})) \tag{34}$$

Now let us have a look at the eigenvalues of matrix $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$.

- For all i such that $x_i = 0$, it is easy to see that the i^{th} column of the identity matrix, that we denote by $\boldsymbol{u}_i$, is a right eigenvector of $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$, associated to the eigenvalue $\lambda_i = \left(\frac{m_i(\boldsymbol{x})}{p_i(\boldsymbol{x})}\right)^\eta$ (since the product of the last term in equation (34) and vector $\boldsymbol{u}_i$ is zero). We can conclude that if $\nabla_i J(\boldsymbol{x}) = 0$ or $\eta = 0$, then $\lambda_i = 1$; otherwise $\lambda_i \in [0,1[$ if $\eta > 0$, and $\lambda_i > 1$ if $\eta < 0$.
- Let $\boldsymbol{u}$ be a right eigenvector of $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ which does not belong to the subspace spanned by the previous ones, associated to an eigenvalue λ. Then

$$\boldsymbol{\Lambda}(\boldsymbol{x})^\eta\boldsymbol{u} - \eta\boldsymbol{\nabla}^2 J(\boldsymbol{x})\mathrm{diag}(\boldsymbol{x}/\boldsymbol{p}(\boldsymbol{x}))\boldsymbol{u} = \lambda\boldsymbol{u}. \tag{35}$$

  Let $\boldsymbol{v} = \boldsymbol{D}(\boldsymbol{x})\,\boldsymbol{u}$; this vector is non-zero, otherwise $\boldsymbol{u}$ would belong the subspace spanned by the previous set of eigenvectors $\boldsymbol{u}_i$. Left multiplying equation (35) by $\boldsymbol{D}(\boldsymbol{x})$ yields $\boldsymbol{\Lambda}(\boldsymbol{x})^\eta\boldsymbol{v} - \eta\boldsymbol{P}(\boldsymbol{x})\boldsymbol{v} = \lambda\boldsymbol{v}$, where the positive semi-definite matrix $\boldsymbol{P}(\boldsymbol{x})$ was defined in equation (12). Then noting that $\boldsymbol{\Lambda}(\boldsymbol{x})^\eta\boldsymbol{v} = \boldsymbol{v}$ (since for any index i, either $\Lambda_i(\boldsymbol{x}) = 1$ or $v_i = 0$), we obtain $(\boldsymbol{I}_n - \eta\,\boldsymbol{P}(\boldsymbol{x}))\,\boldsymbol{v} = \lambda\boldsymbol{v}$, where $\boldsymbol{I}_n$ denotes the $n \times n$ identity matrix. This proves that λ is an eigenvalue of $\boldsymbol{I}_n - \eta\,\boldsymbol{P}(\boldsymbol{x})$. It is easy to see that the previous set of vectors $\boldsymbol{u}_i$ are also eigenvectors of $\boldsymbol{P}(\boldsymbol{x})$, associated to the eigenvalue 0, but they cannot be colinear to $\boldsymbol{v}$, since $v_i = 0$ for all i such that $x_i = 0$. Thus $\lambda = 1 - \eta\mu$, where μ is an eigenvalue of $[\boldsymbol{P}(\boldsymbol{x})]_+^\star$ (with the use of Notation 1). We can conclude that if $\mu = 0$, then $\lambda = 1$. Otherwise we note that $\eta^\star$ defined in equation (13) is equal to $\frac{2}{\|[\boldsymbol{P}(\boldsymbol{x})]_+^\star\|_2}$, and
  - if $\eta = 0$, all the other eigenvalues are equal to 1;
  - if $0 < \eta < \eta^\star$, all the other eigenvalues belong to $]-1,1[$;
  - if $\eta < 0$, all the other eigenvalues are greater than 1;
  - if $\eta > \eta^\star$, there is at least one eigenvalue $\lambda < -1$;
  - if $\eta = \eta^\star$, there is at least one eigenvalue $\lambda = -1$.

Finally, the total number of eigenvalues equal to 1 (if $\eta \neq 0$) is the number of coefficients i such that $x_i = 0$ and $\nabla_i J(\boldsymbol{x}) = 0$, plus the dimension of the kernel of matrix $[\boldsymbol{P}(\boldsymbol{x})]_+^\star$. In other words, it is equal to the dimension of the kernel of $[\boldsymbol{P}(\boldsymbol{x})]_+$. ■

Proof of Proposition 9: Since $\boldsymbol{x}$ is an exponentially stable fixed point of mapping $\boldsymbol{\phi}$, all the eigenvalues of $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$ have magnitude lower than 1. Moreover, $\boldsymbol{\phi}(\boldsymbol{x}) = \boldsymbol{x}$ and $\forall i$, either $x_i = 0$ or $m_i(\boldsymbol{x}) = p_i(\boldsymbol{x})$. Thus equation (33) still yields equation (34). Again, let us have a look at the eigenvalues of matrix $\nabla\boldsymbol{\phi}^T(\boldsymbol{x})$:

- For all i such that $x_i = 0$, the eigenvalue $\lambda_i = \left(\frac{m_i(\boldsymbol{x})}{p_i(\boldsymbol{x})}\right)^\eta$ associated to the eigenvector $\boldsymbol{u}_i$ is lower than 1. Thus $m_i(\boldsymbol{x}) < p_i(\boldsymbol{x})$ and $\nabla_i J(\boldsymbol{x}) > 0$.
- Previous developments in the Proof of Proposition 7 show that the others eigenvalues λ can be written in the form $\lambda = 1 - \eta\,\mu$, where μ is an eigenvalue of $[\boldsymbol{P}(\boldsymbol{x})]_+^\star$. Since $\lambda < 1$, we conclude that $\mu > 0$, thus $[\boldsymbol{P}(\boldsymbol{x})]_+^\star$ is a positive definite matrix, and so is $[\boldsymbol{\nabla}^2 J(\boldsymbol{x})]_+^\star = [\boldsymbol{\nabla}^2 J(\boldsymbol{x})]_+$.

We have thus proved that properties (14) and (15) stand. This proves that $\boldsymbol{x}$ is a local minimum of function J. ■

C. Lyapunov's second method

Proof of Proposition 12: Function $V(\boldsymbol{x},\boldsymbol{y})$ defined in equation (16) can be decomposed as follows:

$$\begin{aligned} V(\boldsymbol{x},\boldsymbol{y}) = & \sum_{i/x_i=0} y_i\,\tfrac{p_i(\boldsymbol{x})+p_i(\boldsymbol{y})}{2} \\ & + \sum_{i/x_i>0} \tfrac{1}{2}(y_i - x_i)^2\,\tfrac{p_i(\boldsymbol{x})+p_i(\boldsymbol{y})}{x_i+y_i} \\ = & \sum_{i/x_i=0} y_i\,(p_i(\boldsymbol{x}) + O(\|\boldsymbol{y}-\boldsymbol{x}\|)) \\ & + \sum_{i/x_i>0} \tfrac{1}{2}(y_i - x_i)^2\left(\tfrac{p_i(\boldsymbol{x})}{x_i} + O(\|\boldsymbol{y}-\boldsymbol{x}\|)\right). \end{aligned} \tag{36}$$

Note that, since $\boldsymbol{\phi}(\boldsymbol{x}) = \boldsymbol{x}$ and $\boldsymbol{\phi}$ is continuously differentiable at $\boldsymbol{x}$, $\|\boldsymbol{\phi}(\boldsymbol{y}) - \boldsymbol{x}\| = O(\|\boldsymbol{y}-\boldsymbol{x}\|)$. Therefore replacing $\boldsymbol{y}$ by $\boldsymbol{\phi}(\boldsymbol{y})$ in equation (36) yields

$$\begin{aligned} V(\boldsymbol{x},\boldsymbol{\phi}(\boldsymbol{y})) = & \sum_{i/x_i=0} \phi_i(\boldsymbol{y})\,(p_i(\boldsymbol{x}) + O(\|\boldsymbol{y}-\boldsymbol{x}\|)) \\ & + \sum_{i/x_i>0} \tfrac{1}{2}(\phi_i(\boldsymbol{y}) - x_i)^2\left(\tfrac{p_i(\boldsymbol{x})}{x_i} + O(\|\boldsymbol{y}-\boldsymbol{x}\|)\right) \end{aligned} \tag{37}$$

Then subtracting equation (36) to equation (37) yields

$$\begin{aligned} & V(\boldsymbol{x},\boldsymbol{\phi}(\boldsymbol{y})) - V(\boldsymbol{x},\boldsymbol{y}) \\ & = \sum_{i/x_i=0} (\phi_i(\boldsymbol{y}) - y_i)\,(p_i(\boldsymbol{x}) + O(\|\boldsymbol{y}-\boldsymbol{x}\|)) + \\ & \sum_{i/x_i>0} (\phi_i(\boldsymbol{y}) - y_i)(y_i - x_i + \tfrac{\phi_i(\boldsymbol{y})-y_i}{2})(\tfrac{p_i(\boldsymbol{x})}{x_i} + O(\|\boldsymbol{y}-\boldsymbol{x}\|)) \end{aligned} \tag{38}$$

However, equation (10) proves that $\phi_i(\boldsymbol{y}) - y_i = y_i\left(\left(\frac{m_i(\boldsymbol{y})}{p_i(\boldsymbol{y})}\right)^\eta - 1\right)$; in particular:

- if $x_i = 0$ and $\nabla_i J(\boldsymbol{x}) > 0$,

$$\phi_i(\boldsymbol{y}) - y_i = -y_i\left(1 - \left(\frac{m_i(\boldsymbol{x})}{p_i(\boldsymbol{x})}\right)^\eta + O(\|\boldsymbol{y}-\boldsymbol{x}\|)\right)$$

- if $x_i = 0$ and $\nabla_i J(\boldsymbol{x}) = 0$,

$$\phi_i(\boldsymbol{y}) - y_i = -\frac{\eta y_i}{p_i(\boldsymbol{x})}\left([\boldsymbol{\nabla}^2 J(\boldsymbol{x})(\boldsymbol{y}-\boldsymbol{x})]_i + O(\|\boldsymbol{y}-\boldsymbol{x}\|^2)\right)$$

- if $x_i > 0$ (and $\nabla_i J(\boldsymbol{x}) = 0$),

$$\phi_i(\boldsymbol{y}) - y_i = -\frac{\eta\,x_i}{p_i(\boldsymbol{x})}[\boldsymbol{\nabla}^2 J(\boldsymbol{x})(\boldsymbol{y}-\boldsymbol{x})]_i + O(\|\boldsymbol{y}-\boldsymbol{x}\|^2).$$

Substituting these three equalities into equation (38), plus a few manipulations show that (with the notation $[.]_0$ and $[.]_+$ introduced in Notation 1)

$$\begin{aligned} & V(\boldsymbol{x},\boldsymbol{\phi}(\boldsymbol{y})) - V(\boldsymbol{x},\boldsymbol{y}) = -[\boldsymbol{y}]_0^T\,([\boldsymbol{v}(\boldsymbol{x})]_0 + O(\|\boldsymbol{y}-\boldsymbol{x}\|)) \\ & -\eta\,[\boldsymbol{y}-\boldsymbol{x}]_+^T[M(\boldsymbol{x})]_+[\boldsymbol{y}-\boldsymbol{x}]_+ + O(\|\boldsymbol{y}-\boldsymbol{x}\|^3) \end{aligned} \tag{39}$$

where $[\boldsymbol{v}(\boldsymbol{x})]_0 = [\boldsymbol{I}_n - \boldsymbol{\Lambda}(\boldsymbol{x})^\eta]_0\,[\boldsymbol{p}(\boldsymbol{x})]_0$ is a vector with (strictly) positive coefficients since $\eta > 0$, and

$$\begin{aligned} [M(\boldsymbol{x})]_+ = & [\boldsymbol{\nabla}^2 J(\boldsymbol{x})]_+ \\ & -\tfrac{\eta}{2}[\boldsymbol{\nabla}^2 J(\boldsymbol{x})]_+\left[\mathrm{diag}\left(\tfrac{\boldsymbol{x}}{\boldsymbol{p}(\boldsymbol{x})}\right)\right]_+[\boldsymbol{\nabla}^2 J(\boldsymbol{x})]_+ \end{aligned} \tag{40}$$

is a positive definite matrix since $\eta < \eta^\star$ (cf. lemma 18 below). Equation (39) finally proves that there is a neighborhood of $\boldsymbol{x}$ such that $\forall \boldsymbol{y} \neq \boldsymbol{x}$, $V(\boldsymbol{x}, \boldsymbol{\phi}(\boldsymbol{y})) - V(\boldsymbol{x}, \boldsymbol{y}) < 0$.

■

Lemma 18. *Let $\boldsymbol{x} \in \mathbb{R}_+^n$ be a local minimum of a function J satisfying Assumption 1, and suppose that Assumption 2 and property* (15) *hold. Then the matrix $[\boldsymbol{M}(\boldsymbol{x})]_+$ defined in equation* (40) *is positive definite if and only if $\eta < \eta^\star$ (where $\eta^\star$ was defined in equation* (13)*).*

Proof of Lemma 18: Matrix $[\boldsymbol{M}(\boldsymbol{x})]_+$ is positive definite if and only if $[\boldsymbol{I}_n]_+ - \frac{\eta}{2}[\boldsymbol{P}'(\boldsymbol{x})]_+$ is positive definite, where

$$[\boldsymbol{P}'(\boldsymbol{x})]_+ = \left([\boldsymbol{\nabla}^2 J(\boldsymbol{x})]_+\right)^{\frac{1}{2}} \left[\operatorname{diag}\left(\frac{\boldsymbol{x}}{\boldsymbol{p}(\boldsymbol{x})}\right)\right]_+ \left([\boldsymbol{\nabla}^2 J(\boldsymbol{x})]_+\right)^{\frac{1}{2}}$$

which is equivalent to $\eta < \frac{2}{\|[\boldsymbol{P}'(\boldsymbol{x})]_+\|_2}$. However, it is easy to prove that the eigenvalues of $[\boldsymbol{P}'(\boldsymbol{x})]_+$ are equal to those of

$$[\boldsymbol{P}(\boldsymbol{x})]_+ = [\boldsymbol{D}(\boldsymbol{x})]_+ [\boldsymbol{\nabla}^2 J(\boldsymbol{x})]_+ [\boldsymbol{D}(\boldsymbol{x})]_+ .$$

Consequently, $\|[\boldsymbol{P}'(\boldsymbol{x})]_+\|_2 = \|[\boldsymbol{P}(\boldsymbol{x})]_+\|_2 = \|\boldsymbol{P}(\boldsymbol{x})\|_2$. ■

REFERENCES

[1] R. Badeau, N. Bertin, and E. Vincent, "Supporting document for the paper "Stability analysis of multiplicative update algorithms and application to non-negative matrix factorization"," Télécom ParisTech, Paris, France, Tech. Rep. 2009D023, Nov. 2009.

[2] M. Berry, M. Browne, A. Langville, V. Pauca, and R. Plemmons, "Algorithms and applications for approximate nonnegative matrix factorization," *Computational Statistics & Data Analysis*, vol. 52, no. 1, pp. 155–173, 2007.

[3] N. Bertin and R. Badeau, "Initialization, distances and local minima in audio applications of the non-negative matrix factorization," in *Proc. of Acoustics'08*. Paris, France: JASA, Jul. 2008.

[4] N. Bertin, R. Badeau, and E. Vincent, "Fast Bayesian NMF algorithms enforcing harmonicity and temporal continuity in polyphonic music transcription," in *Proc. of IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, New Paltz, NY, USA, Oct. 2009, pp. 29–32.

[5] ——, "Enforcing harmonicity and smoothness in Bayesian non-negative matrix factorization applied to polyphonic music transcription," *IEEE Trans. Audio, Speech, Language Process.*, vol. 18, no. 3, pp. 538–549, Mar. 2010.

[6] N. Bertin, C. Févotte, and R. Badeau, "A tempering approach for Itakura-Saito non-negative matrix factorization. With application to music transcription," in *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing*, Taipei, Taiwan, Apr. 2009, pp. 1545–1548.

[7] S. Choi, "Algorithms for orthogonal nonnegative matrix factorization," in *Proc. of IEEE International Joint Conference on Neural Networks*, Orlando, Florida, USA, Jun. 2008, pp. 1828–1832.

[8] M. Chu, F. Diele, R. Plemmons, and S. Ragni, "Optimality, computation, and interpretation of nonnegative matrix factorizations," Wake Forest University, North Carolina, USA, Tech. Rep., Oct. 2004.

[9] A. Cichocki, S. Amari, R. Zdunek, R. Kompass, G. Hori, and Z. He, "Extended SMART algorithms for non-negative matrix factorization," *Springer LNAI*, vol. 4029, pp. 548–562, 2006.

[10] A. Cichocki, R. Zdunek, and S. Amari, "Csiszar's divergences for non-negative matrix factorization: Family of new algorithms," in *Proc. of the 6th International Conference on Independent Component Analysis and Blind Signal Separation*, Charleston, SC, USA, Mar. 2006, pp. 32–39.

[11] ——, "Nonnegative matrix and tensor factorization," *IEEE Signal Processing Magazine*, vol. 25, no. 1, pp. 142–145, Jan. 2008.

[12] A. Cont, "Realtime multiple pitch observation using sparse non-negative constraints," in *Proc. of International Symposium on Music Information Retrieval*, Victoria, Canada, Oct. 2006, pp. 206–212.

[13] J. Eggert, H. Wersing, and E. Korner, "Transformation-invariant representation and NMF," in *Proc. of IEEE International Joint Conference on Neural Networks*, Budapest, Hungary, Jul. 2004, pp. 2535–2539.

[14] S. Eguchi and Y. Kano, "Robustifying maximum likelihood estimation," Tokyo Institute of Statistical Mathematics, Tokyo, Japan, Tech. Rep., 2001. [Online]. Available: http://www.ism.ac.jp/~eguchi/pdf/Robustify_MLE.pdf

[15] C. Févotte, N. Bertin, and J.-L. Durrieu, "Nonnegative matrix factorization with the Itakura-Saito divergence. With application to music analysis," *Neural Computation*, vol. 21, no. 3, pp. 793–830, Mar. 2009.

[16] L. Finesso and P. Spreij, "Approximate nonnegative matrix factorization via alternating minimization," in *Proc. of the 16th International Symposium on Mathematical Theory of Networks and Systems*, Leuven, Belgium, Jul. 2004.

[17] E. Gonzalez and Y. Zhang, "Accelerating the Lee-Seung algorithm for nonnegative matrix factorization," TR-05-02, Rice University, Houston, Texas, USA, Tech. Rep., Mar. 2005.

[18] D. Guillamet, J. Vitrià, and B. Schiele, "Introducing a weighted non-negative matrix factorization for image classification," *Pattern Recognition Letters*, vol. 24, no. 14, pp. 2447–2454, 2003.

[19] F. Itakura and S. Saito, "Analysis synthesis telephony based on the maximum likelihood method," in *Proc. of the 6th International Congress on Acoustics*, Tokyo, Japan, Aug. 1968, pp. C17–C20.

[20] R. Kompass, "A generalized divergence measure for nonnegative matrix factorization," *Neural Computation*, vol. 19, no. 3, pp. 780–791, Mar. 2007.

[21] X. Kong, C. Hu, and C. Han, "On the discrete-time dynamics of a class of self-stabilizing MCA extraction algorithms," *IEEE Trans. Neural Netw.*, vol. 21, no. 1, pp. 175–181, Jan. 2010.

[22] J. Lasalle, *The stability and control of discrete processes*. New York, NY, USA: Springer-Verlag, 1986.

[23] H. Laurberg, "Uniqueness of non-negative matrix factorization," in *Proc. of IEEE Workshop on Statistical Signal Processing*, Madison, WI, USA, Aug. 2007, pp. 44–48.

[24] D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, pp. 788–791, Oct. 1999.

[25] ——, "Algorithms for non-negative matrix factorization," in *Proc. of Conference on Advances in Neural Information Processing Systems*, vol. 13. Vancouver, British Columbia, Canada: MIT Press, Dec. 2001, pp. 556–562.

[26] S. Li, X. Hou, H. Zhang, and Q. Cheng, "Learning spatially localized, parts-based representation," in *Proc. of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Hawaii, USA, Dec. 2001, pp. 207–212.

[27] C.-J. Lin, "On the convergence of multiplicative update algorithms for nonnegative matrix factorization," *IEEE Trans. Neural Netw.*, vol. 18, no. 6, pp. 1589–1596, Nov. 2007.

[28] ——, "Projected gradient methods for non-negative matrix factorization," *Neural Computation*, vol. 19, no. 10, pp. 2756–2779, 2007.

[29] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed., ser. Springer Series in Operations Research and Financial Engineering. New York Inc.: Springer-Verlag, Aug. 2006.

[30] G. Peyré, "Non-negative sparse modeling of textures," in *Proc. of Scale Space and Variational Methods in Computer Vision*, Ischia, Italy, May 2007, pp. 628–639.

[31] P. Smaragdis, "Non-negative matrix factor deconvolution; extraction of multiple sound sources from monophonic inputs," in *Proc. of International Conference on Independent Component Analysis and Blind Signal Separation*, Granada, Spain, Sep. 2004, pp. 494–499.

[32] E. Vincent, "Musical source separation using time-frequency source priors," *IEEE Trans. Audio, Speech, Language Process.*, vol. 14, no. 1, pp. 91–98, Jan. 2006.

[33] E. Vincent, N. Bertin, and R. Badeau, "Adaptive harmonic spectral decomposition for multiple pitch estimation," *IEEE Trans. Audio, Speech, Language Process.*, vol. 18, no. 3, pp. 528–537, Mar. 2010.

[34] T. Virtanen, "Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria," *IEEE Trans. Audio, Speech, Language Process.*, vol. 15, no. 3, pp. 1066–1074, Mar. 2007.

[35] Z. Wu, H. Su, J. Chu, and W. Zhou, "Improved delay-dependent stability condition of discrete recurrent neural networks with time-varying delays," *IEEE Trans. Neural Netw.*, vol. 21, no. 4, pp. 692–697, Apr. 2010.

[36] Y. Zhang and Y. Fang, "A NMF algorithm for blind separation of uncorrelated signals," in *Proc. of IEEE International Conference on Wavelet Analysis and Pattern Recognition*, Beijing, China, Nov. 2007, pp. 999–1003.

[37] C.-D. Zheng, H. Zhang, and Z. Wang, "An augmented LKF approach involving derivative information of both state and delay," *IEEE Trans. Neural Netw.*, vol. 21, no. 7, pp. 1100–1109, Jul. 2010.

Enforcing Harmonicity and Smoothness in Bayesian Non-negative Matrix Factorization Applied to Polyphonic Music Transcription.

Nancy Bertin*, Roland Badeau, *Member, IEEE* and Emmanuel Vincent, *Member, IEEE*

***Abstract*—This article presents theoretical and experimental results about constrained non-negative matrix factorization (NMF) in a Bayesian framework. A model of superimposed Gaussian components including harmonicity is proposed, while temporal continuity is enforced through an inverse-Gamma Markov chain prior. We then exhibit a space-alternating generalized expectation-maximization (SAGE) algorithm to estimate the parameters. Computational time is reduced by initializing the system with an original variant of multiplicative harmonic NMF, which is described as well. The algorithm is then applied to perform polyphonic piano music transcription. It is compared to other state-of-the-art algorithms, especially NMF-based. Convergence issues are also discussed on a theoretical and experimental point of view.**

Bayesian NMF with harmonicity and temporal continuity constraints is shown to outperform other standard NMF-based transcription systems, providing a meaningful mid-level representation of the data. However, temporal smoothness has its drawbacks, as far as transients are concerned in particular, and can be detrimental to transcription performance when it is the only constraint used. Possible improvements of the temporal prior are discussed.

***Index Terms*—Non-negative matrix factorization (NMF), music transcription, audio source separation, unsupervised machine learning, Bayesian regression.**

I. Introduction

NON-NEGATIVE matrix factorization (NMF) is a powerful, unsupervised decomposition technique allowing the representation of two-dimensional non-negative data as a linear combination of meaningful elements in a basis.

NMF has been widely and successfully used to process audio signals, including various tasks such as monaural sound source separation [1], audio stream separation [2], audio-to-score alignment [3], drum transcription [4]. In particular, it has been efficiently used to separate notes in polyphonic music [5], [6] and transcribe it in a symbolic format such as MIDI. In this case, a time-frequency representation of the signal is factored as the product between a basis (or dictionary) of pseudo-spectra and a matrix (decomposition) of time-varying gains. When obtained from harmonic instruments sounds, the basis is shown to partially retain harmonic components, with a pitched structure, that can be interpreted as musical notes, while the decomposition gives information about the onset and offset times of the associated notes.

Meaningful is here a key word: we expect the basis to be formed of interpretable elements, exhibiting certain semantics. The non-negativity constraint is a first step towards this interpretability, compared to other well-known techniques such as Singular Value Decomposition (SVD). For instance, the basis learnt by NMF from an image database is expected to contain meaningful images (the so-called "part-based representation" [7]). This interpretability is often observed in practice, which is certainly one of the reasons for NMF's popularity; but it is not always as satisfying as expected (see, for instance, facial images in [8], that are expected to retain facial parts like eyes, nose, mouth, but do not exactly). As some other desirable characteristics of the decomposition, it is more observed as a welcome side-effect, than enforced and controlled.

To alleviate this lack of control on the decomposition properties, most authors have proposed constrained variants of NMF, ensuring and enhancing those side-effects of baseline NMF: sparsity, spatial localization, temporal continuity for instance. The typical approach for such constrained variants is to add a penalty term to the usual cost function (reconstruction error) and minimize their sum, see *e.g.* [1], [8], [9].

On the other hand, several authors have imported the idea of a non-negative constraint in other frameworks than NMF, in particular statistical framework. We can cite non-negative variants of Independent Component Analysis (ICA) [10] and non-negative sparse coding [11]. The Bayesian framework offers both a strong theoretical framework, and the possibility to manage constraints through models and priors.

In this paper, we focus on a Bayesian approach of NMF that allows to enforce harmonicity of the dictionary components (a desired property for music transcription task) and temporal smoothness of the decomposition, preserving however the adaptiveness of NMF, which is purely data-driven, and the interest of the provided mid-level representation for other potential applications. The paper is organised as follows. Section II recalls the baseline NMF model and state-of-the-art constrained NMF algorithms. In particular, constraints of harmonicity and temporal continuity are discussed and Bayesian approaches for NMF are presented. Our model, and an EM-like algorithm for NMF with harmonicity and temporal smoothness are proposed in section III, including a multiplicative initialization phase that updates our previous work on harmonic NMF. Section IV is devoted to experimental results in the transcription task context. Conclusion and perspectives

N. Bertin and R. Badeau are with the Département Traitement du Signal et des Images, Institut TELECOM, TELECOM ParisTech, LTCI CNRS, Paris, France (e-mail: {nancy.bertin,roland.badeau}@telecom-paristech.fr).

E. Vincent is with the METISS Project, IRISA-INRIA, 35 042 Rennes Cedex, France (e-mail: emmanuel.vincent@irisa.fr)

are drawn in section V.

II. Constrained non-negative matrix factorization

A. Notations

Matrices are denoted by straight bold letters, for instance, $\mathbf{V}$, $\mathbf{W}$, $\mathbf{H}$. Lowercase bold letters denote column vectors, such as $\mathbf{w}_k = (w_{1k} \dots w_{Fk})^T$, while lowercase plain letters with a single index denote rows, such that $\mathbf{H} = (h_1^T \dots h_K^T)^T$. We also define the matrix $\hat{\mathbf{V}} = \mathbf{WH}$.

We use the binary operators $\triangleq$ to denote definitions and $\overset{c}{=}$ to denote equality up to an additive constant.

Calligraphic uppercase letters are used to denote probability distributions: $\mathcal{N}$, $\mathcal{P}$, $\mathcal{IG}$ denote Gaussian, Poisson and inverse-Gamma distributions. Their expressions are recalled in appendix A.

B. Baseline model and algorithms

Out of any applicative context, the NMF problem is expressed as follows: given a matrix $\mathbf{V}$ of dimensions $F \times N$ with non-negative entries, NMF is the problem of finding a factorization

$$\mathbf{V} \approx \mathbf{WH} = \hat{\mathbf{V}} \tag{1}$$

where $\mathbf{W}$ and $\mathbf{H}$ are non-negative matrices of dimensions $F \times K$ and $K \times N$, respectively. K is usually chosen such that $FK + KN \ll FN$, hence reducing the data dimension. In typical audio applications, the matrix $\mathbf{V}$ is chosen as a time-frequency representation (*e.g.* magnitude or power spectrogram), f denoting the frequency bin and n the time frame.

The factorization (1) is generally obtained by minimizing a cost function defined by

$$D(\mathbf{V}|\hat{\mathbf{V}}) = \sum_{f=1}^{F} \sum_{n=1}^{N} d(v_{fn}|\hat{v}_{fn}) \tag{2}$$

where $d(a|b)$ is a function of two scalar variables. d is typically non-negative and takes value zero if and only if (iff) $a = b$. The most popular cost functions for NMF are the Euclidean (EUC) distance and the generalized Kullback-Leibler (KL) divergence, which were particularly popularized (as NMF itself) by Lee and Seung, see, *e.g.*, [7]. They described multiplicative update rules under which $D(\mathbf{V}|\mathbf{WH})$ is shown to be non-increasing, while ensuring non-negativity of $\mathbf{W}$ and $\mathbf{H}$. The update rules are obtained by using a simple heuristics, which can be seen as a gradient descent algorithm with an appropriate choice of the descent step. By expressing the gradient of the cost function ∇D as the difference of two positive terms $\nabla^+ D$ and $\nabla^- D$, the cost function is shown (in particular cases) or observed to be nonincreasing under the rules:

$$\begin{cases} \mathbf{W} \leftarrow \mathbf{W} \otimes \dfrac{\nabla^-_{\mathbf{W}} D(\mathbf{V}|\mathbf{WH})}{\nabla^+_{\mathbf{W}} D(\mathbf{V}|\mathbf{WH})} \\ \mathbf{H} \leftarrow \mathbf{H} \otimes \dfrac{\nabla^-_{\mathbf{H}} D(\mathbf{V}|\mathbf{WH})}{\nabla^+_{\mathbf{H}} D(\mathbf{V}|\mathbf{WH})} \end{cases} \tag{3}$$

Sparsity	$\sum_{k=1}^{K} \frac{1}{\sqrt{N}-1} \left(\sqrt{N} - \sum_{n=1}^{N} \lvert h_{kn}\rvert / \sqrt{\sum_{n=1}^{N} h_{kn}^2}\right)$	[12]
Spatial localization	$\lambda_1 \sum_{k=1}^{K} \sum_{k'=1}^{K} \left[\mathbf{W}^T\mathbf{W}\right]_{kk'} - \lambda_2 \sum_{k=1}^{K} \left[\mathbf{HH}^T\right]_{kk}$	[8]
Least correlation	$\sum_{k=1}^{K} \log \left[\mathbf{HH}^T\right]_{kk} - \log \lvert\mathbf{HH}^T\rvert$	[9]
Temporal continuity	$\sum_{k=1}^{K} \sum_{n=1}^{N} \lvert h_{kn} - h_{k(n-1)}\rvert^2$	[1]

TABLE I
Some state-of-the-art constraints D_c in NMF problem.

For some choices of d, like EUC or KL, monotonicity of the criterion under these rules can be proven [7], but in the general case, these updates do not guarantee any convergence.

C. Constrained approaches

1) Constraints imposed via penalty terms: In standard NMF, the only constraint is the elementwise non-negativity of all matrices. All other properties of the decomposition, as satisfying as it is, come as uncontrolled side-effects and in a way, the fact that the decomposition retains certain semantics of the original signal, performs separation or provides meaningful and interpretable components is just "good news". It sounds thus natural to try to improve this potential by adding explicit constraints to the factorization problem, in order to enhance and control desired properties.

Then, several constraints have been introduced to get NMF solutions that better fit certain expectancies. Among other proposed constraints, we can cite sparsity [12], spatial localization [8], least correlation between sources [9] or temporal continuity [1], [13].

The common point between those algorithms, whichever constraint is considered, is the "penalty term approach". Rather than minimizing only a reconstruction error term D_r (EUC or KL, typically), the minimized cost function includes a term D_c that quantifies the desired property. The constrained NMF problem is then expressed as:

$$\min_{\mathbf{W},\mathbf{H}} D_r(\mathbf{V}|\mathbf{WH}) + \lambda D_c(\mathbf{V}|\mathbf{WH})$$

where λ is a weight parameter. Table I gives a few examples of literature penalty terms. Temporal smoothness is one of these examples. In standard NMF and most of its variants, time frames are considered as independent, non-related observations, which is obviously not true for real-world sounds and in particular for music. In the case of musical notes, the main part of the note (the sustain and decay parts, after the attack) possesses a slowly time-varying spectrum. When expressed as the product between a template spectrum $\mathbf{w}_k$ and a time-varying gain h_k, according to NMF formulation, it is equivalent to saying that the row h_k is smooth, or, in other words, that the coefficient h_{kn} is not that different from $h_{k(n-1)}$. [1] and [13] thus introduce penalty terms in the NMF cost function to take into account this temporal continuity. In

[1] the term is directly linked to the differences $h_{kn} - h_{k(n-1)}$, while [13] variant relies on a ratio between short-time and long-time variance of h_k. Those terms are shown to favor smoothness in lines of $\mathbf{H}$. Another possible approach is the statistical approach from [14]. Temporal continuity is favored through putting an appropriate prior on $\mathbf{H}$. This solution will be exposed with more details and adapted to our case in section III-C.

It is interesting to notice that non-smoothness may also be an objective (see for instance [15]), depending on the data and the application. [15] points out that smoothness of one of the NMF factors (*i.e.* $\mathbf{W}$ or $\mathbf{H}$) may enhance sparsity of the other one, thus establishing a link between those two popular constraints. On the other hand, [1] combines sparsity and temporal continuity constraints on $\mathbf{H}$, but concludes to the non-efficiency of the sparsity constraint in his particular case.

The penalty approach has several drawbacks. First, a criterion quantifying the desired property must be found. Second, no general proof of convergence is available for the update scheme (3). Moreover, the parameter λ has to be chosen empirically. These reasons motivated our approach for harmonicity constraint in previous and current work; this approach is exposed in section II-C2.

2) Deterministic constraints: Musical notes, excluding transients, are pseudo-periodic. Their spectra are then comb-alike, with regularly spaced frequency peaks. As we wish to use NMF to separate musical notes in a polyphonic recording, we expect that elements in the basis $\mathbf{W}$ are as near as possible from a harmonic distribution. This property is yet not easily quantified by a penalty term.

In [16], we rather proposed an alternative model to baseline NMF problem, enforcing the basis harmonicity. We impose the basis components to be expressed as the linear combination of narrow-band harmonic spectra (patterns), which are arbitrarily fixed:

$$w_{fk} = \sum_{m=1}^{M} e_{mk} \mathbf{P}_{km}(f) \tag{4}$$

For a given component number k, all the patterns $\mathbf{P}_{km}$ share the same pitch (fundamental frequency f_0); they are defined by summation of the spectra of a few adjacent individual partials at harmonic frequencies of f_0, scaled by the spectral shape of subband k. This spectral envelope is chosen according to perceptual modelling [16]. Figure 1 illustrates the patterns for one note and the corresponding atom $\mathbf{w}_k$. Coefficients e_{mk} are learned by NMF as well as the decomposition $\mathbf{H}$. Update rules are obtained by minimizing the same cost function as in baseline NMF, except that it is minimized with respect to (wrt) $\mathbf{E}$ and $\mathbf{H}$ rather than $\mathbf{W}$ and $\mathbf{H}$.

3) Statistical constraints: Another way to induce properties in the NMF is to switch to a statistical framework and introduce adequate prior distributions. Let us consider the following model, proposed in [17], [18]: $\forall n = 1, ..., N$,

$$\mathbf{x}_n = \sum_{k=1}^{K} \mathbf{c}_{kn} \in \mathbb{C}^F \tag{5}$$

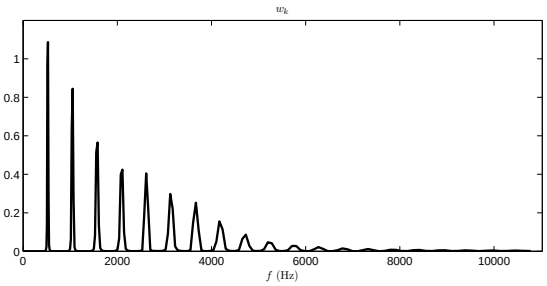

(a) Dictionary atom $\mathbf{w}_k$

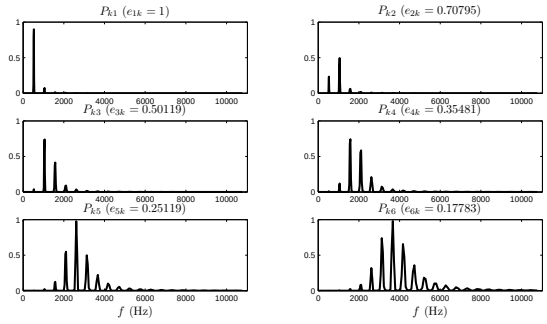

(b) Corresponding patterns P_{km}

Fig. 1. Example harmonic basis spectrum $\mathbf{w}_k$ corresponding to the note $C4$ (MIDI pitch 72), with underlying narrowband spectra $\mathbf{P}_{km}$ and spectral envelope coefficients e_{mk} (with M= 6).

where latent variables $\mathbf{c}_{kn}$ are independent and follow a multivariate Gaussian distribution

$$\mathbf{c}_{kn} \sim \mathcal{N}(0, h_{kn}\text{diag}(\mathbf{w}_k)) \tag{6}$$

In [14], the estimation of the parameter $\boldsymbol{\theta} = \{\mathbf{W}, \mathbf{H}\}$, in a maximum likelihood (ML) sense is shown to be equivalent to solving the NMF problem $\mathbf{V} \approx \mathbf{WH}$, when observing $\mathbf{V} = (|x_{fn}|^2)_{fn}$ and choosing the underlying cost function d as the Itakura-Saito divergence:

$$d_{IS}(a|b) = \frac{a}{b} - \log\frac{a}{b} - 1 \tag{7}$$

Other authors, like [19], have proven similar equivalences between NMF with KL cost and ML estimation in the model:

$$|\mathbf{x}_n| = \sum_{k=1}^{K} |\mathbf{c}_{kn}| \tag{8}$$

under the assumption $|\mathbf{c}_{kn}(f)| \sim \mathcal{P}(w_{fk}h_{kn})$, where $\mathcal{P}(\lambda)$ is the Poisson distribution.

In [20], the authors propose a model where the factors $\mathbf{W}$ and $\mathbf{H}$ are expressed as two functions f_h and f_w (called "link-functions") of Gaussian latent variables. It can be seen as a generalization of the previous model for appropriate choices of f_h and f_w (relatively soft assumptions are put on them). It is another example of the power of the statistical approach to incorporate constraints or knowledge in the NMF problem.

One main advantage of this statistical approach is the possibility to switch from ML estimation to maximum a

posteriori (MAP) estimation, thanks to Bayes rule:

$$p(\mathbf{W},\mathbf{H}|\mathbf{V}) = \frac{p(\mathbf{V}|\mathbf{W},\mathbf{H})p(\mathbf{W})p(\mathbf{H})}{p(\mathbf{V})} \quad (9)$$

Thus, choosing adequate prior distributions $p(\mathbf{W})$ and $p(\mathbf{H})$ is a way to induce desired properties in the decomposition. Furthermore, the statistical framework provides a strong theoretical basis and efficient algorithms with proven convergence, like the expectation-maximization (EM) algorithm and its variants, to estimate NMF factors.

In next section, we propose to combine this framework and the previous model (4) to enforce both harmonicity in columns of $\mathbf{W}$ and smoothness in rows of $\mathbf{H}$, which are desired properties of the NMF of musical signals.

III. Proposed algorithm

A. Probabilistic harmonic model

The direct usage of formulation (4) in the model (5) is possible, but leads to computational issues. An equivalent model is obtained by assuming:

$$\mathbf{x}_n = \sum_{k=1}^{K}\sum_{m=1}^{M} \mathbf{d}_{kmn} \quad (10)$$

with

$$\begin{aligned} \mathbf{x}_n &\in \mathbb{C}^F \\ \mathbf{d}_{kmn} &\sim \mathcal{N}(0, h_{kn}e_{mk}\text{diag}(\mathbf{P}_{km})) \\ \mathbf{P}_{km} &= [P_{km}(1)...P_{km}(F)]^T \end{aligned}$$

Assuming the equality $\mathbf{c}_{kn} = \sum_m \mathbf{d}_{kmn}$ and the independence of $\mathbf{d}_{kmn}$, we can verify that $\mathbf{c}_{kn} \sim \mathcal{N}(0, h_{kn}\sum_m e_{mk}\text{diag}(\mathbf{P}_{km}))$.F

From [14], we can establish the equivalence between ML estimation in this generative model (10) and minimization of d_{IS}, which will offer a good coherence and comparability between algorithms for our test. [14] also shows that Itakura-Saito divergence, whose expression is recalled in equation (7), is well-suited to NMF decomposition of audio signals. Advantages of d_{IS} also include a good fit between the representation and the observation on a log scale **(due to the shape of the function $d_{\mathbf{V}}(\hat{v}_{fn}) = d_{IS}(v_{fn}|\hat{v}_{fn})$ at fixed energy scale v_{fn}, and the strong cost of representing a bin v_{fn} by $\hat{v}_{fn} = 0$)** and then, **better chances to represent low-energy residual noise** (if components are devoted to it, see future work suggestions in section V). This motivates our model and the choice of IS cost (and not, for instance, the weighted Euclidean cost from [16]) in this work.

B. Maximum likelihood estimation

We now describe an EM-based algorithm for the estimation of the parameters $\boldsymbol{\theta} = \{\mathbf{E},\mathbf{H}\}$. This algorithm is adapted from ML estimation proposed in [14] for the model (10). In ML estimation, the criterion to be maximized is the log-likelihood of the observations:

$$C_{ML}(\boldsymbol{\theta}) \triangleq \log p(\mathbf{V}|\boldsymbol{\theta}) \quad (11)$$

We partition the set of all parameters into disjoint subsets $\boldsymbol{\theta}_k = \{\{e_{mk}\}_m, h_k\}$ so that $\boldsymbol{\theta} = \bigcup_{k=1}^{K}\boldsymbol{\theta}_k$. This partition, and the additive form of the model (10) where the latent variables are supposed independent, allow the usage of the Space Alternating Generalized EM algorithm (SAGE), introduced in [21], to estimate the parameters. The hidden data-space associated with each subset $\boldsymbol{\theta}_k$ is $\mathbf{D}_k = [\mathbf{D}_{k1}\dots\mathbf{D}_{kN}]$, where $\mathbf{D}_{km} = [\mathbf{d}_{km1}\dots\mathbf{d}_{kmN}] \in \mathbb{C}^{F\times N}$. The use of SAGE implies maximizing the functional $Q_k^{ML}(\boldsymbol{\theta}_k|\boldsymbol{\theta'})$ which is the conditional expectation of the log likelihood of $\mathbf{D}_k$:

$$Q_k^{ML}(\boldsymbol{\theta}_k|\boldsymbol{\theta'}) \triangleq \int_{\mathbf{D}_k} \log p(\mathbf{D}_k|\boldsymbol{\theta}_k)p(\mathbf{D}_k|\mathbf{V},\boldsymbol{\theta'})d\mathbf{D}_k \quad (12)$$

where $\boldsymbol{\theta'}$ contains the most up-to-date estimated values of all parameters.

We can however notice that Q_k^{ML} can be expressed as the sum (over m) of auxiliary functionals Q_{km}^{ML} expressed as:

$$Q_{km}^{ML}(\boldsymbol{\theta}_{km}|\boldsymbol{\theta'}) \triangleq \int_{\mathbf{D}_{km}} \log p(\mathbf{D}_{km}|\boldsymbol{\theta}_{km})p(\mathbf{D}_{km}|\mathbf{V},\boldsymbol{\theta'})d\mathbf{D}_{km} \quad (13)$$

where we define subsets $\boldsymbol{\theta}_{km} = \{e_{mk}, h_k\}$. The problem reduces to maximizing each $Q_{km}^{ML}(\boldsymbol{\theta}_{km}|\boldsymbol{\theta'})$ wrt e_{mk}, and the sum $Q_k^{ML}(\boldsymbol{\theta}_k|\boldsymbol{\theta'})$ wrt h_{kn} iteratively. Maximizing these functionals makes the criterion $C_{ML}(\boldsymbol{\theta})$ increase, according to [21].

At each iteration and for each k, the functionals Q_{km}^{ML} are computed. The sum of the functionals over m is formed and maximized by computing and zeroing its derivative wrt h_{kn}. The derivative wrt e_{mk} of each functional is computed and zeroed, resulting in an update rule for each e_{mk}. Details of the computations are available in appendix B. Updates rules can be then expressed as follows:

$$h_{kn}^{(\ell+1)} = h_{kn}^{(\ell)} \times \left(1 + \frac{1}{FM}\sum_f\sum_m \frac{h_{kn}^{(\ell)}e_{mk}^{(\ell)}P_{km}(f)}{\hat{v}_{fn}}\left(\frac{v_{fn}}{\hat{v}_{fn}} - 1\right)\right) \quad (14)$$

$$e_{mk}^{(\ell+1)} = e_{mk}^{(\ell)} \times \left(1 + \frac{1}{FN}\sum_n\sum_f \frac{h_{kn}^{(\ell+1)}e_{mk}^{(\ell)}P_{km}(f)}{\hat{v}_{fn}}\left(\frac{v_{fn}}{\hat{v}_{fn}} - 1\right)\right) \quad (15)$$

where the superscript ℓ denotes the value at iteration ℓ and where $\hat{v}_{fn}$ is the current reconstruction of v_{fn}, *i.e.* $\hat{v}_{fn} = \sum_{k=1}^{K}\sum_{m=1}^{M} h_{kn}e_{mk}P_{km}(f)$ with the most up-to-date values of the parameter, either (ℓ) or $(\ell+1)$ depending on the most recent available values. In SAGE formalism, we update separately each row h_k, one after the other, *but*, during this step, all h_{kn} for n from 1 to N are updated simultaneously.[1].

Using SAGE framework guarantees the monotonicity of the criterion $C^{ML}(\boldsymbol{\theta})$. Moreover, [21] proves the existence of a region of monotone convergence in norm, i.e., $\boldsymbol{\theta}$ converges in norm to a local minimum, provided the algorithm was initialized in an appropriate neighborhood of that minimum.

[1]With a more explicit notation, at iteration $(\ell+1)$, the coefficient $h_{kn}^{(\ell+1)}$ is determined using, for all p, coefficients $h_{jp}^{(\ell+1)}$ for $j < k$ and $h_{jp}^{(\ell)}$ for $j \geq k$.

C. Enforcing temporal smoothness

In terms of computational cost, this maximum likelihood estimation of $\mathbf{E}$ and $\mathbf{H}$ has no practical interest, compared to multiplicative gradient descent update rules: as observed in [14] for a similar case (multiplicative vs. SAGE algorithm for standard NMF with Itakura-Saito divergence), it is computationnally slower and even more sensitive to local minima than usual multiplicative algorithms. However, it has two main advantages: first, the theoretical framework guarantees convergence to a local minimum; second, it opens the possibility of including priors on the parameters, possibly in a hierarchical fashion, and then constraining NMF solutions in an elegant way.

In [14], this framework is exploited to enforce temporal smoothness over the rows of $\mathbf{H}$. We provide a priori information on $\boldsymbol{\theta}$, expressed as a prior distribution $p(\boldsymbol{\theta})$. Thanks to the Bayes rule, recalled in equation (9), we get a maximum a posteriori (MAP) estimator by maximizing the following criterion:

$$C_{MAP}(\boldsymbol{\theta}) \triangleq \log p(\boldsymbol{\theta}|\mathbf{V}) \quad (16)$$
$$\stackrel{\text{c}}{=} C_{ML}(\boldsymbol{\theta}) + \log p(\boldsymbol{\theta}) \quad (17)$$

We choose here to use the Markov chain prior structure proposed in [14]:

$$p(h_k) = p(h_{k1}) \prod_{n=2}^{N} p(h_{kn}|h_{k(n-1)}) \quad (18)$$

where $p(h_{kn}|h_{k(n-1)})$ reaches its maximum at $h_{k(n-1)}$, thus favoring a slow variation of h_k in time. We proposed for instance the following choice:

$$p(h_{kn}|h_{k(n-1)}) = \mathcal{IG}(h_{kn}|\alpha_k, (\alpha_k+1)h_{k(n-1)}) \quad (19)$$

where $\mathcal{IG}(u|a,b)$ denotes the inverse-Gamma distribution with shape parameter a and scale parameter b, whose mode is $b/(a+1)$; the initial distribution $p(h_{k1})$ is Jeffrey's non-informative prior (see appendix C).

Several reasons motivated the choice of this prior. First, non-negativity arises naturally from this modelling. Secondly, this prior is conjugate with respect to the Gaussian observation model, which brings computational simplicity. Moreover, it seems appropriate to the modelling of temporal envelopes in music signal. First, it favours the smoothness by the appropriate choice of the mode. Secondly, the assymmetry of the inverse-Gamma distribution around the mode (if $\delta h \geq 0$, $p(h_{kn}+\delta h|h_{kn}) \leq p(h_{kn}-\delta h|h_{kn})$) constraints more smoothness on decrease parts ($h_{k(n+1)} \leq h_{kn}$) than on increase parts ($h_{k(n+1} \geq h_{kn}$). Thus, it favours smoothness in silence and sustain/decay parts of the notes, but still allows correct reprensentation of the attacks.

Parameters α_k are here arbitrarily fixed, depending on the desired degree of smoothness (the higher α_k, the smoother h_k), but we could consider in future work the possibility to learn it as well.

As the prior respects the scheme $p(\mathbf{H}) = \prod_{k=1}^{K} p(h_k)$, we can still use the SAGE formalism. The functional (12) to minimize is now written:

$$Q_k^{MAP}(\boldsymbol{\theta}_k|\boldsymbol{\theta}') \stackrel{\text{c}}{=} \sum_{m=1}^{M} Q_{km}^{ML}(e_{mk}, h_k|\boldsymbol{\theta}') + \log p(h_k) \quad (20)$$

Q_{km}^{ML} being unchanged, we just have to incorporate the contribution of the prior in the computation and zeroing of the gradients. In Appendix C, this is shown to be proportional to a second-order polynomial:

$$\nabla_{h_{kn}} Q_k^{MAP}(e_{mk}, h_k|\boldsymbol{\theta}') = \frac{-FM}{h_{kn}^2}(p_2\, h_{kn}^2 + p_1\, h_{kn} - p_0) \quad (21)$$

The values of p_0, p_1, p_2 are common for each $n \in [2 \dots N-1]$ and take different values at the borders of the Markov chain ($n = 1$ and $n = N$). They obviously depend on k, n and ℓ (though the notation doesn't mention it, for readability purpose). Their expressions are given in Table II and the detailed computations are available in Appendix C. The resulting update rule is given by the only non-negative root of the polynomial:

$$h_{kn}^{(\ell+1)} = \frac{2p_0}{\sqrt{p_1^2 + 4p_2p_0} + p_1} \quad (22)$$

(written here in a form avoiding possible division by zero) and the ML update of $\mathbf{E}$ (15) is unchanged.

	$n=1$	$n=2\dots N-1$	$n=N$
p_0	$\tilde{h}_{k1}$	$\tilde{h}_{kn} + \frac{\alpha_k+1}{FM} h_{k(n-1)}$	$\tilde{h}_{kN} + \frac{(\alpha_k+1)}{FM} h_{k(N-1)}$
p_1	$1 + \frac{1-\alpha_k}{FM}$	$1 + \frac{1}{FM}$	$1 + \frac{1+\alpha_k}{FM}$
p_2	$\frac{1}{FM}\frac{\alpha_k+1}{h_{k2}}$	$\frac{1}{FM}\frac{\alpha_k+1}{h_{k(n+1)}}$	0

TABLE II
COEFFICIENTS OF THE ORDER 2 POLYNOMIAL TO BE SOLVED IN ORDER TO UPDATE h_{kn} IN BAYESIAN HARMONIC NMF WITH AN INVERSE-GAMMA MARKOV CHAIN PRIOR. $\tilde{h}_{kn}$ DENOTES THE ML UPDATE, GIVEN BY THE RIGHT MEMBER OF EQUATION (14).

In the following, we refer to this algorithm as "Harmonic Smooth NMF" (or, in short form, "HS-NMF").

We can also consider the current model of temporal smoothness, but without harmonicity constraint, leading to the regularized NMF algorithm proposed in [14]. In the following, this algorithm will be denoted as "S-NMF".

D. Multiplicative initialization with harmonicity

Due to the slow convergence of EM-like algorithms, HS-NMF needs to be efficiently initialized. Theoretical results from [21] also suggest the interest of smart initialization in terms of convergence of the algorithm. Harmonic multiplicative NMF could then be used to "bootstrap" SAGE algorithm. However, the multiplicative algorithm of [22] was originally designed for a perceptually weighted Euclidean distance, which would not be coherent with HS-NMF criterion (based on IS divergence (7)). For this reason, we wish to adapt

harmonic NMF with multiplicative update rules from [22] to this distance. The criterion to be minimized is written:

$$D_{IS}(\mathbf{V}|\mathbf{WH}) = \sum_{f=1}^{F}\sum_{n=1}^{N} d_{IS}(v_{fn}|\sum_{k=1}^{K}\sum_{m=1}^{M} h_{kn}e_{mk}P_{km}(f)) \tag{23}$$

We compute its derivative wrt h_{kn}, which is expressed as the difference of two positive terms:

$$\nabla_{h_{kn}} D_{IS}(\mathbf{V}|\mathbf{WH}) = \sum_{f=1}^{F}\frac{w_{fk}}{\hat{v}_{fn}} - \sum_{f=1}^{F}\frac{v_{fn}w_{fk}}{\hat{v}_{fn}^2} \tag{24}$$

where $\hat{v}_{fn} = \sum_{k'=1}^{K} w_{fk'}h_{k'n} = \sum_{k'=1}^{K}\sum_{m'=1}^{M} e_{m'k'}P_{k'm'}(f)h_{k'n}$. The derivative wrt e_{mk} fits in the same scheme:

$$\nabla_{e_{mk}} D_{IS}(\mathbf{V}|\mathbf{WH}) = \sum_{f=1}^{F}\sum_{n=1}^{N}\frac{h_{kn}P_{km}(f)}{\hat{v}_{fn}} - \sum_{f=1}^{F}\sum_{n=1}^{N}\frac{v_{fn}h_{kn}P_{km}(f)}{\hat{v}_{fn}^2} \tag{25}$$

The update rules are derived from the heuristics (3) and write:

$$h_{kn} \leftarrow h_{kn} \times \frac{\sum_{f=1}^{F} v_{fn}w_{fk}/\hat{v}_{fn}^2}{\sum_{f=1}^{F} w_{fk}/\hat{v}_{fn}} \tag{26}$$

$$e_{mk} \leftarrow e_{mk} \times \frac{\sum_{f=1}^{F}\sum_{n=1}^{N} v_{fn}h_{kn}P_{km}(f)/\hat{v}_{fn}^2}{\sum_{f=1}^{F}\sum_{n=1}^{N} h_{kn}P_{km}(f)/\hat{v}_{fn}} \tag{27}$$

In the following, this algorithm will be referred to as "H-NMF/MU".

IV. Application to music transcription

Music transcription consists in converting a raw music signal into a symbolic representation of the music within: for instance a score, or a MIDI file. Here, we focus on information strictly related to musical notes, *i.e.* musical pitch, onset and offset time, discarding high level information usually available in a full music sheet, such as bar lines or key signature. Automatic transcription is a very active field of research, known to be difficult, in particular because of note overlapping in the time-frequency plane. Various methods have been proposed to address the transcription issue, including neural network modelling [23], parametric signal modelling and HMM tracking [24] or Bayesian approaches [25]. We propose here to assess the efficiency of Bayesian harmonic and smooth NMF for this task.

A. Experimental setup

1) Database: To evaluate and quantify transcription performance, we need a set of polyphonic music pieces with accurate MIDI references. The two most simple ways to get such data are either to record a MIDI instrument (the acquisition of audio and MIDI being simultaneous), or to synthetize sound from given MIDI files. For the sake of timbre realism and ease of acquisition, the piano is an instrument of choice: very high quality software synthetizers are available on sale, and an acoustic piano can be equipped to play mechanically, and produce a MIDI output, while retaining the timbre of a real instrument. In his thesis [26], Valentin Emiya collected such a database. *MAPS* (MIDI-Aligned Piano Sounds) includes isolated notes, random and tonal chords, pieces from the piano repertoire, recordings on an upright DisKlavier and high quality software synthesis. From this very complete database, we excerpted two subsets to evaluate our algorithms: a synthetic subset, produced by Native Instruments' Akoustik Piano ("Bechstein Bach" preset, from samples recorded on a Bechstein D280 piano), and a real audio subset, recorded at TELECOM ParisTech on a Yamaha Mark III (upright DisKlavier). Each subset is composed of 30 pieces of 30 seconds each (original pieces from *MAPS* were truncated). The piano was chosen for practical reasons, but it can be stressed that nothing in the method constraints it to be applied only to piano signals.

2) Structure of NMF-based transcription: All NMF-based transcription systems used here follow the same workflow:

1) Computation of an adapted time-frequency representation of the signal, $\mathbf{V}$;
2) Factorization $\mathbf{V} \approx \mathbf{WH}$;
3) Attribution of a MIDI pitch to each basis spectrum $\mathbf{w}_k$ (either from original labelling of columns, when the algorithm includes the harmonicity constraint, or by performing a pitch estimation);
4) Onset/offset detection applied to each time envelope h_k.

In [22], it is observed that using a nonlinear frequency scale resulted in a representation of smaller size, with better temporal resolution in the higher frequency range, than usual Short-Time Fourier Transform (STFT), while preserving the subsequent transcription performance. We then pass the signal through a filterbank of 257 sinusoidally modulated Hanning windows with frequencies linearly spaced between 5 Hz and 10.8 kHz on the Equivalent Rectangular Bandwidth (ERB) scale. We then split each subband into disjoint 23 ms time frames and compute the power within each frame.

Abbr.	Description	Reference
NMF/MU	Baseline NMF minimizing IS divergence Multiplicative update rules	[14]
S-NMF	SAGE algorithm for NMF With smoothness constraint on $\mathbf{H}$	[14]
Virtanen'07	Multiplicative NMF With temporal continuity constraint Minimizing KL div. plus a penalty term	[1]
Vincent'08	Multiplicative NMF With weighted Euclidean distance and harmonicity constraint	[16]
Marolt'04	Neural network based transcription	[23]

TABLE III
Reference algorithms.

Pitch estimation of basis spectra is superfluous in NMF with harmonicity constraint, since each basis component can be labelled from the beginning with the pitch of the patterns $\mathbf{P}_{km}$ used to initialize it. For NMF without this constraint, pitch identification is performed on each column of $\mathbf{W}$ by the harmonic comb-based technique used in [16].

Note onsets and offsets are determined by a simple threshold-based detection, followed by a minimum-duration pruning, see [16]. The detection threshold is denoted by A_{dB} and expressed in dB under $\mathbf{H}$ maximum.

3) Evaluation: Transcription performance is quantitatively evaluated according to usual information retrieval scores [27]. **Precision rate** ($\mathcal{P}$) is the proportion of correct notes among all transcribed notes (quantifying the number of notes that are transcribed, but should not). **Recall rate** ($\mathcal{R}$) is the proportion of notes from the MIDI reference which are correctly transcribed (thus quantifying the number of notes that should be transcribed, but are not). **F-measure** ($\mathcal{F}$) aggregates the two former criteria in one unique score and is defined as $\mathcal{F} = 2\mathcal{P}\mathcal{R}/(\mathcal{P}+\mathcal{R})$. A transcribed note is considered as correct if its pitch is identical to the ground truth, and its onset time is within 50ms of the ground truth, according to community standards (see, for instance, the MIREX competition). Note offset detection is also evaluated through the mean overlap ratio ($\mathcal{MOR}$) defined in [28]. For a correctly transcribed note, the overlap ratio o_{note} between the original note and its transcription is the ratio between the length of the intersection and union of their temporal widths:

$$o_{note} = \frac{\min(t_{off}) - \max(t_{on})}{\max(t_{off}) - \min(t_{on})} \tag{28}$$

where t_{on} and t_{off} are the vectors of onset times (respectively offset times) of the original and corresponding transcribed note. **Mean Overlap Ratio** ($\mathcal{MOR}$) is the mean of overlap ratios for all correct notes.

The original algorithms (H-NMF/MU and HS-NMF) previously proposed are compared to several state-of-the-art algorithms listed in Table III.

H-NMF/MU, HS-NMF and S-NMF were implemented by the authors for this work. Virtanen'07 and NMF/MU are run from their author's implementation, which they nicely shared, and Marolt'04 is run from the SONIC software, distributed by its author. Vincent'08 is tuned with the optimal parameters determined in [16]. The order K is set to 88 (the number of components, i.e. of columns in $\mathbf{W}$, is naturally taken as the number of keys on a piano) for all NMF-based algorithms. For algorithms with harmonicity constraint, we take take one component (fundamental frequency) per pitch. The maximum number of patterns per note is $M = 10$. When a multiplicative initialization is needed (HS-NMF and S-NMF), 10 iterations of the associated multiplicative algorithm (H-NMF/MU and NMF/MU respectively) are performed before switching to the tested algorithm. Note detection thresholds A_{dB} are manually tuned algorithm per algorithm (and reported in Tables IV and V), by maximizing the average F-measure on each dataset. The minimum duration for a transcribed note is fixed to 50ms.

B. Results

1) Convergence: We monitor the values of C^{MAP} and D_{IS} at each iteration of HS-NMF, in order to evaluate its speed and efficiency of convergence, and to assess the impact of initializing HS-NMF by H-NMF/MU. Then, we compare the evolution of the criteria between "pure" HS-NMF, and HS-NMF preceded by 10 iterations of H-NMF/MU, on the same example piece from the dataset and with the same random initialization. Figure 2 presents this evolution in these two cases. Though C^{MAP} decreases sharply during the initialization (10 first iterations), the multiplicative initialization phase allows the algorithm to reach a higher value of the criterion for the same number of iterations, as well as a lower value of the reconstruction error term D_{IS} (which is equal to the minus log-likelihood up to a constant). After a few hundreds of iterations, the reconstruction error changes very little, while the contribution from the prior still increases slowly, resulting in very few changes in the transcription performance. More decisive, on the presented excerpt (one 30s piece from the real audio subset), HS-NMF with multiplicative initialization reaches a good transcription performance ($\mathcal{F}$=54.5%), while its counterpart without HS-NMF/MU initialization is totally inefficient in separating notes in the same time ($\mathcal{F}$=0% after 500 iterations). An explanation for this is the relative weights between the two terms in C^{MAP}: the first goal is to reach a good reconstruction, smoothness is a bonus; but if the contribution from the prior takes the most part of the criterion, reconstruction will be poor. Multiplicative initialization allows to optimize first the reconstruction error term, then to focus on the refinement that is the smoothness constraint.

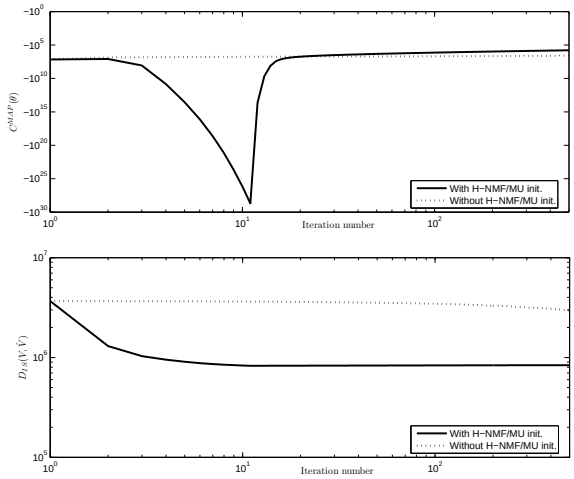

Fig. 2. Evolution of the criteria C^{MAP} and D_{IS} wrt the iteration number.

Algorithm	$\mathcal{P}$	$\mathcal{R}$	$\mathcal{F}$	$\mathcal{MOR}$	A_{dB}
NMF/MU	63.4	56.1	54.9	51.2	-62
Vincent'08	60.7	60.0	58.4	54.8	-32
H-NMF/MU	58.7	59.1	52.4	46.0	-33
S-NMF	62.4	43.3	49.5	50.7	-51
Virtanen'07	55.9	56.4	53.6	52.1	-22
HS-NMF	65.8	64.5	60.7	44.3	-38
Marolt'04	83.5	70.1	75.8	53.5	-

TABLE IV
TRANSCRIPTION SCORES ON SYNTHETIC DATA.

2) Overall transcription performance: Tables IV and V report the transcription performance of tested algorithms on the synthetic and recorded datasets respectively. HS-NMF

outperforms other NMF-based algorithms in both cases, but remains less performant than SONIC software. Smoothness constraint used alone seems detrimental to transcription performance, may it be implemented by a multiplicative algorithm (Virtanen'07) or by a Bayesian algorithm (S-NMF), but improves the performance of harmonically constrained NMF (H-NMF vs. HS-NMF).

Results are comparable to scores from [24] obtained on a database including ours, and place our algorithm performance at the state-of-the-art level.

Algorithm	$\mathcal{P}$	$\mathcal{R}$	$\mathcal{F}$	$\mathcal{MOR}$	A_{dB}
NMF/MU	43.3	43.4	40.8	47.7	-60
Vincent'08	38.7	37.4	36.1	50.0	-30
H-NMF/MU	43.0	42.7	41.3	44.6	-30
S-NMF	46.2	32.0	36.6	45.6	-49
Virtanen'07	34.2	34.8	33.6	47.1	-21
HS-NMF	46.6	45.3	45.0	43.2	-32
Marolt'04	63.7	53.6	58.0	50.0	-

TABLE V
TRANSCRIPTION SCORES ON REAL AUDIO DATA.

3) Harmonicity of the basis: On figure 3, we display bases $\mathbf{W}$ after convergence, with columns sorted by increasing pitch. We can see that non-harmonically constrained NMF exhibits a dictionary that has a pitched structure but a rather noisy look, whereas bases from harmonically constrained algorithms are much cleaner. S-NMF produces a much less sparser dictionary than unconstrained NMF, which is coherent with observations from [15] and could explain its lower performance. Another noticeable result is the pitch repartition in the basis. In NMF without harmonicity constraint, as the basis is completely free, pitch repartition in the basis follows the same trend as pitch repartition in the original piece; NMF tends to use more components to represent faithfully the most frequent notes, while possibly neglecting rare passing tones. Moreover, some components do not exhibit a pitched structure (5, in average). On the contrary, NMF with harmonicity constraint have a fixed number of components per pitch (one, in our case). This guarantees representation of all notes, including notes played only a few times in the piece, but implies also useless computation on components corresponding to absent notes in the piece, and does not allow representation of non-harmonic parts of the signal. This could be alleviated by adding unconstrained components to the harmonic dictionary, updated separately under usual multiplicative rules, for instance.

4) Smoothness of components: Temporal envelopes h_k, for k corresponding to the note $C4$, obtained by NMF/MU (without constraint), H-NMF, S-NMF and HS-NMF are displayed on Figure 4. The ground truth pianoroll (time-pitch representation) is displayed as well. S-NMF and HS-NMF produce indeed smoother envelopes, which can be noticed in particular when the note is supposed to be off. We can notice several spurious peaks in NMF/MU and H-NMF/MU, for instance during the first 750 milliseconds (region **(a)**) or around $t = 10s$ (region **(b)**), whose amplitude is reduced or zeroed by the associate smooth version (S-NMF and HS-NMF respectively). Another noticeable result is that harmonicity constraint seems to disfavour smoothness of the envelopes.

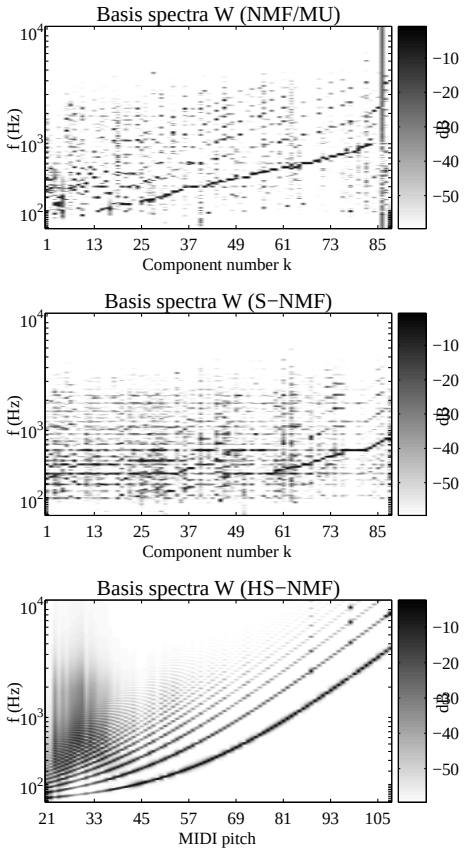

Fig. 3. Example basis matrices $\mathbf{W}$ for algorithms without and with harmonicity contraint. Columns are sorted by increasing pitch.

We also briefly investigated on the impact of the choice of α_k on smoothness and performance; values of α_k between 5 and 15 resulted in a loss of less than 2 points in the F-measure and a barely noticeable difference in the smoothness of rows.

We could have expected a positive influence of the smoothness constraint on the $\mathcal{MOR}$ values, which would mean a better offset detection. Though we observe slightly better $\mathcal{MOR}$ for Virtanen'07 and S-NMF compared to, for instance, HS-NMF, it seems here difficult to draw a straightforward conclusion, partially because of the previously observed interactions between harmonicity and smoothness contraints.

5) Detection threshold: In Tables IV and V, the optimal detection threshold is manually determined to get the best mean F-measure over the test database. Varying this threshold allows to display Precision-Recall curves and have a deeper insight on algorithms performance. Figure 5 presents these curves for NMF/MU, H-NMF/MU, S-NMF and HS-NMF. The curve confirms the good performance of HS-NMF. It reaches a better trade-off between precision and recall and is more robust to the choice of the threshold. Both multiplicative algorithms (H-NMF/MU and NMF/MU) are comparable around the optimal F-measure. S-NMF gives the poorest results at every threshold. We can also notice that a 100% recall is never reached, even at very low threshold, which points a limit of NMF-based transcription algorithms.

These curves, as well as Tables IV and V, are obtained by averaging the scores over the dataset, but it is important

to note an important variability between pieces, in terms of performance and optimal threshold. At fixed threshold A_{dB}, $\mathcal{F}$ standard deviation is worth about 12% for all NMF-based algorithms (from 9% for Virtanen'07, to 16% for HS-NMF).

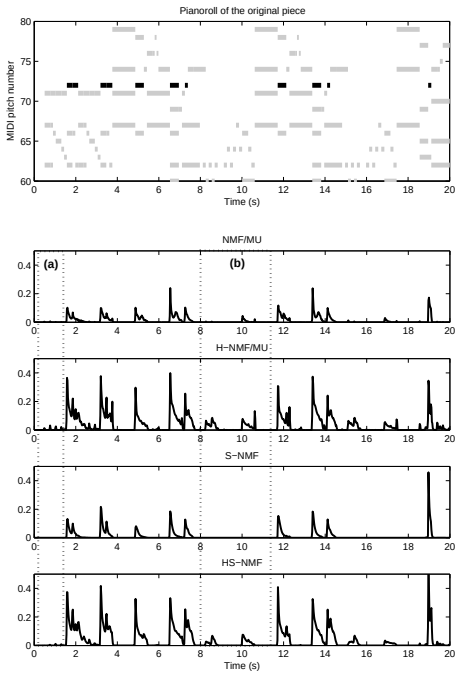

Fig. 4. Temporal activation of note $C4$ for four different algorithms (NMF/MU, H-NMF/MU, S-NMF and HS-NMF from top to bottom) on the same excerpt. The pianoroll of the corresponding excerpt is on top, with $C4$ in black and neighbour notes in gray. Regions of interest are framed with dotted lines.

V. Conclusion and perspectives

In this paper, we proposed an original model for including harmonicity and temporal smoothness constraints in non-negative matrix factorization of time-frequency representations, in a unified framework. The resulting algorithm we propose, HS-NMF, is derived from a Bayesian framework and outperforms other benchmarked NMF approaches in a task of polyphonic music transcription, evaluated on a realistic music database. The Bayesian framework also offers theoretical results about convergence, that are generally not available in usual multiplicative approaches of NMF. We also proposed a novel multiplicative NMF with harmonicity constraint, minimizing Itakura-Saito divergence, which has links with the exposed statistical approach and was shown to suit well for the representation of audio signals in this context [14]. Thus, the contributions of this paper are theoretical, algorithmic and experimental at a time, in the very active domains of music transcription and NMF study.

NMF-based methods remain here less performant than other finely tuned state-of-the-art methods, especially methods implying a training phase, the use of learning data and musicologically inspired post-processing. However, NMF is totally data-driven, it requires no training and then adapts itself to the data while avoiding the risk of a mismatch between training and test data. It also provides a semantically meaningful mid-level representation of the data. Its potential here assessed is clear, letting the hope of very good performance with better tuning and improvements. The temporal smoothness constraint

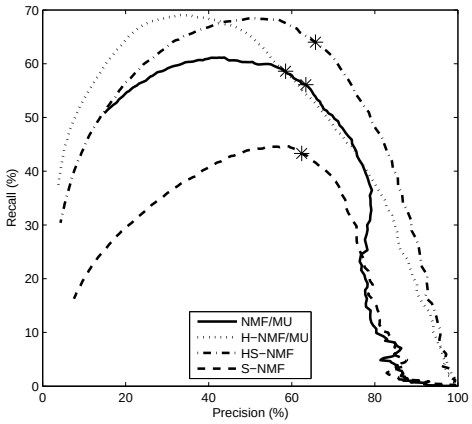

Fig. 5. Precision-Recall curves for four different algorithms. The detection threshold varies from 0 to -100 dB under $\mathbf{H}$ maximum. The couple $(\mathcal{P},\mathcal{R})$ realizing the $\mathcal{F}$ maximum is represented with a star.

does not bring all improvements we could expect, in particular in terms of robustness to the detection threshold and efficiency of the note duration estimation. However, it seems useful to compensate the tendency of NMF with harmonicity constraint to produce non-smooth decomposition, and lead therefore to a better transcription performance when both constraints are used. A limitation of our common NMF framework (NMF core algorithm plus detection threshold based post-processing) appears here, as a 100% recall rate is never reached, for any value of the threshold or any tested algorithm.

Using a statistical model relies of course on the fact that the ground truth actually follows this model. Performance obtained here let hope it is more or less the case, but adequation between the data and the model should be further investigated on. In particular, the choice of the shape parameter α of the inverse-Gamma prior put on temporal envelopes should be discussed, and its learning, as well as NMF factors are learnt, should be considered.

Possible improvements include a refinement of the temporal prior, which suits for modelling the sustain and decay parts of the note, but disfavour attacks and silences. An option to alleviate this mismatch between the model and the data could be the use of switching state models for the rows of $\mathbf{H}$, that would explicitly model the possibility for h_{kn} to vary quickly (attack) or to be strictly zero (absence of the note). As far as $\mathbf{W}$ is concerned, transients are badly represented in an entirely harmonic dictionary, but this could be solved by adding a few unconstrained (non harmonic) components in the representation, which would hopefully be well captured thanks to d_{IS} scale-invariance. At last, as many EM-based algorithms, HS-NMF remains very slow compared to multiplicative gradient descent approaches; an alternative to it could be the direct minimization of the criterion it optimizes by the usual multiplicative heuristics (3), possibly losing the

proof of convergence but reducing computational time.

APPENDIX A
STANDARD DISTRIBUTIONS

Complex valued Gaussian $\mathcal{N}(\mathbf{u}|\boldsymbol{\mu},\boldsymbol{\Sigma}) = |\pi\,\boldsymbol{\Sigma}|^{-1}\,\exp -(\mathbf{u}-\boldsymbol{\mu})^H\,\boldsymbol{\Sigma}^{-1}\,(\mathbf{u}-\boldsymbol{\mu})$
Poisson $\mathcal{P}(u|\lambda) = \exp(-\lambda)\,\frac{\lambda^u}{u!}$
Inverse-Gamma $\mathcal{IG}(u|\alpha,\beta) = \frac{\beta^\alpha}{\Gamma(\alpha)}\,u^{-(\alpha+1)}\,\exp(-\frac{\beta}{u}),\ u \geq 0$
Jeffrey's

APPENDIX B
SAGE UPDATE RULES WITH HARMONICITY

In this appendix we detail the derivations leading to update rules of equations (14) and (15). The functional $Q_{km}^{ML}(\boldsymbol{\theta}_{km}|\boldsymbol{\theta'})$ defined in equation (13) may be processed in two steps. First, we write the hidden data log-likelihood:

$$\log p(\mathbf{D}_{km}|\boldsymbol{\theta}_{km}) = \log \prod_{n=1}^{N}\prod_{f=1}^{F} p(d_{kmn}(f)|\boldsymbol{\theta}_{km}) \qquad (29)$$

As $d_{kmn}(f) \sim \mathcal{N}(0, h_{kn}e_{mk}\mathbf{P}_{km}(f))$, we have:

$$\log p(\mathbf{D}_{km}|\boldsymbol{\theta}_{km}) \stackrel{c}{=} -\sum_{n=1}^{N}\sum_{f=1}^{F}\log(h_{kn}e_{mk}\mathbf{P}_{km}(f)) + \qquad (30)$$

The second term to be computed is the hidden data posterior $p(\mathbf{D}_{km}|\mathbf{V},\boldsymbol{\theta'})$. It may be obtained by writing $\mathbf{x}_n = \mathbf{d}_{kmn} + \sum\sum_{(k',m')\neq(k,m)}\mathbf{d}_{k'm'n}$ and using the Wiener filtering method proposed in [17] for the separation of two sources. According to it, the posterior mean and variance of $d_{kmn}(f)$ write respectively:

$$\mu_{kmn}^{post}(f) = \frac{h_{kn}e_{mk}\mathbf{P}_{km}(f)}{\hat{v}_{fn}}x_n(f) \qquad (31)$$

$$\lambda_{kmn}^{post}(f) = \frac{h_{kn}e_{mk}\mathbf{P}_{km}(f)}{\hat{v}_{fn}}\sum_{(k',m')\neq}\sum_{(k,m)} h_{k'n}e_{k'm'}\mathbf{P}_{k'm'}(f) \qquad (32)$$

Then, by taking the expectation of the log-likelihood with regard to the posterior, we get the functional expression:

$$Q_{km}^{ML}(\boldsymbol{\theta}_{km}|\boldsymbol{\theta'}) = -\sum_{n=1}^{N}\sum_{f=1}^{F}\log(h_{kn}e_{mk}\mathbf{P}_{km}(f)) + \frac{\left|\mu_{kmn}^{post}(f)\right|^2 + \lambda_{kmn}^{post}(f)}{h_{kn}e_{mk}\mathbf{P}_{km}(f)} \qquad (33)$$

Zeroing the gradients of Q_{km}^{ML} wrt e_{mk} and the gradient of their sum over m wrt h_{kn} leads to the update rules:

$$h_{kn}^{(\ell+1)} = \frac{1}{FM}\sum_f\sum_m \frac{\left|\mu_{kmn}^{post'}(f)\right|^2 + \lambda_{kmn}^{post'}(f)}{e_{mk}^{(\ell)}\mathbf{P}_{km}(f)} \qquad (34)$$

$$e_{mk}^{(\ell+1)} = \frac{1}{FN}\sum_n\sum_f \frac{\left|\mu_{kmn}^{post'}(f)\right|^2 + \lambda_{kmn}^{post'}(f)}{h_{kn}^{(\ell+1)}\mathbf{P}_{km}(f)} \qquad (35)$$

where the superscript $'$ indicates that $\lambda_{kmn}^{post'}$ and $\mu_{kmn}^{post'}$ are computed with most up-to-date values of $\mathbf{E}$ and $\mathbf{H}$. This form lets appear possible numeric errors if $h_{kn} = 0$ or $e_{mk} = 0$. This can be avoided by replacing λ_{kmn}^{post} and μ_{kmn}^{post} by their expressions (31) and (32). This leads to update rules proposed in equations (14) and (15).

APPENDIX C
SAGE UPDATE RULES WITH HARMONICITY AND TEMPORAL SMOOTHNESS

We write the functional $Q_k^{MAP} = \sum_{m=1}^{M} Q_{km}^{ML} + \log p(h_k)$ as the sum of the ML functional and contributions from the prior. For $n = 2 \ldots N-1$:

$$\nabla_{h_{kn}} Q_k^{MAP}(\boldsymbol{\theta}_k|\boldsymbol{\theta'}) = \nabla_{h_{kn}}\left(\sum_{m=1}^{M} Q_{km}^{ML}(\boldsymbol{\theta}_{km}|\boldsymbol{\theta'})\right) + \nabla_{h_{kn}}\left(\log p(h_{k(n+1)}|h_{kn}) + \log p(h_{kn}|h_{k(n-1)})\right) \qquad (36)$$

As $\log \mathcal{IG}(u|\alpha,\beta) \stackrel{c}{=} \alpha\log\beta - (\alpha+1)\log u - \beta/u$, we have:

$$\nabla_{h_{kn}} Q_k^{MAP}(\boldsymbol{\theta}_k|\boldsymbol{\theta'}) = -\frac{\alpha_k+1}{h_{k(n+1)}} - \frac{FM+1}{h_{kn}} + \frac{1}{h_{kn}^2}\left(\sum_{f=1}^{F}\sum_{m=1}^{M}\frac{\left|\mu_{kmn}^{post}(f)\right|^2 + \lambda_{kmn}^{post}(f)}{e_{mk}P_{km}(f)} + (\alpha_k+1)h_{k(n-1)}\right) \qquad (37)$$

Then, this gradient is proportional to a second-order polynomial:

$$\begin{aligned}\nabla_{h_{kn}} Q_k^{MAP}(\boldsymbol{\theta}_k|\boldsymbol{\theta'}) &= \frac{-FM}{h_{kn}^2}\left(p_2 h_{kn}^2 + p_1 h_{kn} - p_0\right)\\ \text{with } p_2 &= \frac{1}{FM}\frac{\alpha_k+1}{h_{k(n+1)}}\\ p_1 &= 1 + \frac{1}{FM}\\ p_0 &= \tilde{h}_{kn} + \frac{\alpha_k+1}{FM}h_{k(n-1)}\end{aligned}$$

where $\tilde{h}_{kn}$ is the ML estimator (see equation (34)). For $n = N$ the term $p(h_{k(n+1)}|h_{kn})$ is simply removed from equation (36). For $n = 1$, the Markov chain structure imposes to choose a prior $p(h_{k1})$. We take Jeffreys' non-informative prior: $p(h_{k1}) \propto 1/h_{k1}$. The corresponding gradients are written:

$$\nabla_{h_{k1}} Q_k^{MAP}(\boldsymbol{\theta}_k|\boldsymbol{\theta'}) = -\frac{FM}{h_{k1}} + \frac{1}{h_{k1}^2}\left(\sum_{f=1}^{F}\sum_{m=1}^{M}\frac{\left|\mu_{kmn}^{post}(f)\right|^2 + \lambda_{kmn}^{post}(f)}{e_{mk}\mathbf{P}_{km}(f)}\right) - \frac{\alpha_k-1}{h_{k1}} - \frac{\alpha_k+1}{h_{k2}}$$

This leads to p_0, p_1 and p_2 values reported in Table II.

ACKNOWLEDGMENT

The research leading to this paper was supported by the European Commission under contract FP6-027026, Knowledge Space of semantic inference for automatic annotation and retrieval of multimedia content (K-SPACE), and by

the French GIP ANR under contract ANR-06-JCJC-0027-01, Décomposition en Éléments Sonores et Applications Musicales (DESAM).

The authors would like to thank Cédric Févotte for his decisive influence on the Bayesian orientation of this work, wise advice on literature review, and support. We also wish to credit Valentin Emiya for its incredible work on collecting and sharing MAPS database, and Tuomas Virtanen for gently sharing code and usage advice.

REFERENCES

[1] T. Virtanen, "Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 3, pp. 1066–1074, mar 2007.

[2] B. Wang and M. Plumbley, "Musical audio stream separation by non-negative matrix factorization," in *Proceedings of the DMRN Summer Conference*, Glasgow, UK, July 23-24, 2005.

[3] A. Cont, "Realtime audio to score alignment for polyphonic music instruments using sparse non-negative constraints and hierarchical HMMs," in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP'06)*, Toulouse, France, May 14-17, 2006.

[4] J. Paulus and T. Virtanen, "Drum transcription with non-negative spectrogram factorisation," in *Proc. of the 13th European Signal Processing Conference (EUSIPCO)*, Antalya, Turkey, Sep 4-8, 2005.

[5] P. Smaragdis and J. Brown, "Non-negative matrix factorization for polyphonic music transcription," in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA'03)*, New Paltz, New York, USA, Oct. 19-22 2003, pp. 177–180.

[6] N. Bertin, R. Badeau, and G. Richard, "Blind signal decompositions for automatic transcription of polyphonic music: NMF and K-SVD on the benchmark," in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP'07)*, vol. 1, Honolulu, Hawaii, USA, Apr. 15-20, 2007, pp. 65–68.

[7] D. Lee and H. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, pp. 788–791, Oct. 1999.

[8] S. Li, X. Hou, H. Zhang, and Q. Cheng, "Learning spatially localized, parts-based representation," in *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 1, Hawaii, USA, Dec. 11-13, 2001, pp. 207–212.

[9] Y. Zhang and Y. Fang, "A NMF algorithm for blind separation of uncorrelated signals," in *Proc. of the 2007 International Conference on Wavelet Analysis and Pattern Recognition*, Beijing, China, Nov. 2-4, 2007, pp. 999–1003.

[10] M. Plumbley, "Algorithms for nonnegative independent component analysis," *IEEE Transactions on Neural Networks*, vol. 14, no. 3, pp. 534–543, Mar 2003.

[11] S. Abdallah and M. Plumbley, "Polyphonic music transcription by non-negative sparse coding of power spectra," in *Proceedings of the 5th International Conference on Music Information Retrieval (ISMIR'04)*, Barcelona, Spain, Oct. 10-14, 2004, pp. 318–325.

[12] P. Hoyer, "Non-negative matrix factorization with sparseness constraints," *Journal of Machine Learning Research*, vol. 5, pp. 1457–1469, Nov. 2004.

[13] Z. Chen, A. Cichocki, and T. M. Rutkowski, "Constrained non-negative matrix factorization method for EEG analysis in early detection of Alzheimers disease," in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP'06)*, vol. 5, Toulouse, France, May 14-19, 2006, pp. 893–896.

[14] C. Févotte, N. Bertin, and J.-L. Durrieu, "Nonnegative matrix factorization with the Itakura-Saito divergence. With application to music analysis," *Neural Computation*, 2008, in press. http://www.tsi.enst.fr/~fevotte/TechRep/techrep08_is-nmf.pdf.

[15] A. Pascual-Montano, J. Carazo, K. Kochi, D. Lehmann, and R. Pascual-Marqui, "Nonsmooth nonnegative matrix factorization (nsNMF)," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 28, no. 3, pp. 403–415, Mar 2006.

[16] E. Vincent, N. Bertin, and R. Badeau, "Harmonic and inharmonic nonnegative matrix factorization for polyphonic pitch transcription," in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP'08)*, Las Vegas, Nevada, USA, March 30 - April 4, 2008, pp. 109–112.

[17] L. Benaroya, L. McDonagh, R. Gribonval, and F. Bimbot, "Non negative sparse representation for Wiener based source separation with a single sensor," in *Proc. of International Conference on Acoustics, Speech and Signal Processing (ICASSP'03)*, Hong Kong, China, Apr. 6-10, 2003, pp. 613–616.

[18] L. Benaroya, R. Blouet, C. Févotte, and I. Cohen, "Single sensor source separation using multiple-window STFT representation," in *Proc. of the International Workshop on Acoustic Echo and Noise Control (IWAENC06)*, Paris, France, Sep. 12-14, 2006.

[19] T. Virtanen, A. T. Cemgil, and S. Godsill, "Bayesian extensions to non-negative matrix factorisation for audio signal modelling," in *Proc. of the International Conference on Acoustics, Speech and Signal Processing (ICASSP'08)*, Las Vegas, Nevada, USA, March 30 - April 4, 2008, pp. 1825–1828.

[20] M. Schmidt and H. Laurberg, "Nonnegative matrix factorization with gaussian process priors," *Computational Intelligence and Neuroscience*, In press. 2008.

[21] J. A. Fessler and A. O. Hero, "Space-alternating generalized expectation maximization algorithm," *IEEE Transactions on Signal Processing*, vol. 42, no. 10, pp. 2664–2677, Oct 1994.

[22] E. Vincent, N. Bertin, and R. Badeau, "Two nonnegative matrix factorization methods for polyphonic pitch transcription," in *Proc. Music Information Retrieval Evaluation eXchange (MIREX)*, University of Vienna, Austria, September 23-30, 2007.

[23] M. Marolt, "A connectionist approach to automatic transcription of polyphonic piano music," *IEEE Trans. on Multimedia*, vol. 6, no. 3, pp. 439–449, June 2004.

[24] V. Emiya, R. Badeau, and B. David, "Automatic transcription of piano music based on HMM tracking of jointly-estimated pitches," in *Proc. Eur. Conf. Sig. Proces. (EUSIPCO)*, Lausanne, Switzerland, Aug 25-29, 2008.

[25] A. Cemgil, H. Kappen, and D. Barber, "A generative model for music transcription," *IEEE Trans. on Audio, Speech and Language Processing*, vol. 14, no. 2, pp. 679–694, March 2006.

[26] V. Emiya, "Transcription automatique de la musique de piano," Ph.D. dissertation, Institut TELECOM; TELECOM ParisTech, 2008.

[27] C. van Rijsbergen, *Information retrieval*, 2nd ed. London, UK: Butterworths, 1979.

[28] M. Ryynänen and A. Klapuri, "Polyphonic music transcription using note event modeling," in *Proc. 2005 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, New Paltz, New York, USA, Oct. 2005, pp. 319–322.

NMF with time-frequency activations to model non stationary audio events

Romain Hennequin, *Student Member, IEEE,* Roland Badeau, *Member, IEEE,* and Bertrand David, *Member, IEEE*

***Abstract*—Real world sounds often exhibit time-varying spectral shapes, as observed in the spectrogram of a harpsichord tone or that of a transition between two pronounced vowels. Whereas the standard Non-negative Matrix Factorization (NMF) assumes fixed spectral atoms, an extension is proposed where the temporal activations (coefficients of the decomposition on the spectral atom basis) become frequency dependent and follow a time-varying ARMA modeling. This extension can thus be interpreted with the help of a source/filter paradigm and is referred to as source/filter factorization. This factorization leads to an efficient single-atom decomposition for a single audio event with strong spectral variation (but with constant pitch). The new algorithm is tested on real audio data and shows promising results.**

***Index Terms*—music information retrieval, non-negative matrix factorization, unsupervised machine learning.**

I. Introduction

THE decomposition of audio signals in terms of elementary atoms has been a large field of research for years. As we usually encounter very dissimilar audio events (both in their spectral and temporal characteristics), the decomposition on a single basis (such as the Fourier basis) is generally not sufficient to accurately explain the content of a large class of signals. Sparse decomposition techniques [1] use a redundant dictionary of vectors (called atoms) and try to decompose a signal using few of them (much less than the dimension of the space): thus the signal can be accurately decomposed with few elements. When atoms are designed to fit the signal (for instance harmonic atoms for musical signals [2]), these elements become more meaningful, and then a supervised classification can be performed to cluster atoms corresponding to a real event in the signal [3]. These methods are quite powerful and give good results. However, since the dictionary is fixed, it must be designed to fit all possible signals, which is not achievable in practice. Recently methods of data *factorization* were proposed to simultaneously extract atoms from the signal and provide a decomposition on these atoms. These techniques that we call *factorization* make use of the natural redundancy of the signal, mimicking human cognition which utilizes this redundancy to understand visual and audio signals: Principal Component Analysis, Independent Component Analysis [4], [5], sparse coding [6] or NMF [7] have been introduced both to reduce the dimensionality and to explain the whole data set by a few meaningful elementary objects. Thanks to the non-negativity constraint, NMF is able to provide a meaningful representation of the data: applied to musical spectrograms it will hopefully decompose them into elementary notes or impulses. The technique is widely used in audio signal processing, with a number of applications such as automatic music transcription [8], [9], [10] and sound source separation [11], [12], [13].

However, the standard NMF is shown to be efficient when the elementary components (notes) of the analyzed sound are nearly stationary, i.e. when the envelope of the spectra of these components does not change over time. Nevertheless, in several situations, elementary components can be strongly non stationary. In this article, we will focus on timbral variability, i.e. variability of the spectral shape that we can find in plucked strings sounds or singing voice (sounds of different vowels present greatly dissimilar spectral shapes). However, we will not address pitch variability that is encountered in vibrato or prosody. In case of a noticeable spectral variability, the standard NMF will likely need several non-meaningful atoms to decompose a single event, which often leads to a necessary post-processing (to cluster the different parts of a single source [14]). To overcome this drawback, Smaragdis [15] proposes a shift-invariant extension of NMF in which time/frequency templates are factorized from the original data: each atom then corresponds to a time-frequency musical event able to include spectral variations over time. This method gives good results, but does not permit any variation between different occurrences of the same event (atom), its duration and spectral content evolution being fixed.

In this paper, an extension of NMF is proposed in which temporal activation becomes frequency dependent: it thus can be interpreted with the help of the classical source/filter paradigm as a source/filter factorization. Our method includes AutoRegressive Moving Average (ARMA) filters estimated from the data, associates a time-varying filter with each source and learns the sources (atoms) in a totally unsupervised way. this method presents some similarity with Durrieu's work [16], [17] in which a source/filter model is used in a NMF framework to extract the main melody of musical pieces. This model permits to efficiently take the strong spectral variations of the human voice into account.

However, our approach is quite different since, in opposition to Durrieu's work, sources are learnt, a time-varying filter is associated with each source and the class of filter we use is more standard.

In section II, we introduce the source/filter decomposition as an extension of NMF. In section III, we derive an iterative

R. Hennequin, R. Badeau and B. David are with the Institut Telecom, Telecom ParisTech, CNRS LTCI, 46, rue Barrault - 75634 Paris Cedex 13 - France (email: <forename>.<surname>@telecom-paristech.fr)

The research leading to this paper was supported by the French GIP ANR under contract ANR-06-JCJC-0027-01, DESAM, and by the Quaero Program, funded by OSEO, French State agency for innovation.

Manuscript received December 8, 2009

algorithm similar to those used for NMF to compute this decomposition. In section IV, we present experiments of source/filter decomposition of the spectrogram of three different sounds, and compare this decomposition to the standard NMF. Conclusions are drawn in section V.

II. Model

A. NMF and extension

Given an $F \times T$ non-negative matrix $\mathbf{V}$ and an integer R such that $FR + RT \ll FT$, NMF approximates $\mathbf{V}$ by the product of an $F \times R$ non-negative matrix $\mathbf{W}$ and an $R \times T$ non-negative matrix $\mathbf{H}$:

$$\mathbf{V} \approx \hat{\mathbf{V}} = \mathbf{WH} \quad \left(\text{i.e. } V_{ft} \approx \hat{V}_{ft} = \sum_{r=1}^{R} w_{fr} h_{rt}\right) \tag{1}$$

This approximation is generally quantified by a cost function $\mathcal{C}(\mathbf{V}, \mathbf{W}, \mathbf{H})$ to be minimized with respect to (wrt) $\mathbf{W}$ and $\mathbf{H}$. A common class of cost functions is designed element-wise:

$$\mathcal{C}(\mathbf{V}, \mathbf{W}, \mathbf{H}) = \sum_{f=1}^{F} \sum_{t=1}^{T} d(V_{ft} | \hat{V}_{ft})$$

where d is a scalar divergence (i.e. a function such that $\forall (a,b) \in \mathbb{R}^2$ $d(a|b) \geq 0$ and $d(a|b) = 0$ if and only if $a = b$). Several classes of such divergences have been proposed, for instance Bregman divergences [18] and β-divergences [19], [20]. In this paper, we will focus on the β-divergence which includes usual measures (Euclidean distance, Kullback-Lieber divergence and Itakura-Saito divergence). The β-divergence is defined for $\beta \in \mathbb{R} \backslash \{0, 1\}$ by:

$$d_\beta(x|y) = \frac{1}{\beta(\beta-1)} (x^\beta + (\beta-1) y^\beta - \beta x y^{\beta-1}) \tag{2}$$

For $\beta = 0$ and $\beta = 1$, the β-divergence is defined by continuity:

$$d_0(x|y) = \frac{x}{y} - \log \frac{x}{y} - 1$$
$$d_1(x|y) = x(\log x - \log y) + (y - x)$$

For these values, the β-divergence respectively corresponds to Itakura-Saito divergence and Kullback-Leibler divergence, and for $\beta = 2$, it corresponds to the Euclidean distance.

One could notice that the singularities in the definition of the β-divergence for $\beta = 0$ and $\beta = 1$ no longer appear in the partial derivative wrt y: this partial derivative is useful for designing descent methods in order to minimize the cost function $\mathcal{C}$. The first order partial derivative of the β-divergence wrt y is for all $\beta \in \mathbb{R}$:

$$\frac{\partial d_\beta(x|y)}{\partial y} = y^{\beta-2}(y - x)$$

When applied to power (squared magnitude) spectrograms, NMF factorizes data into a matrix (or basis) of frequency templates which are the R columns of $\mathbf{W}$ and a matrix $\mathbf{H}$ whose R rows are the temporal vectors of activations corresponding to each template. For a musical signal made of several notes played by the same instrument, it is hoped that the decomposition leads to spectral templates corresponding to single notes or percussive sounds. $\mathbf{H}$ will then display a representation similar to a "piano-roll" (cf. [8]).

This factorization however does not yield an effective representation of a sound presenting a noticeable spectral evolution. For instance a single note of a plucked string instrument most of the time shows high frequency components which decrease faster than low frequency components. This characteristic is not well modeled with a single frequency template. Several templates are needed which results in a less meaningful decomposition: roughly one for low frequency partials and one for high frequency partials. The meaning of each template is lost (a template no longer corresponds to a musical event such as a note).

To address this issue, we propose an extension of NMF where temporal activations become time/frequency activations. The factorization (1) becomes:

$$V_{ft} \approx \hat{V}_{ft} = \sum_{r=1}^{R} w_{fr} h_{rt}(f) \tag{3}$$

where the activation coefficients are now frequency dependent. To avoid an increase of the problem dimensionality the $h_{rt}(f)$ coefficients are further parameterized by means of ARMA models (section II-B).

Equation (3) can be interpreted with the help of the source/filter paradigm: the spectrum of each frame of the signal results from the combination of filtered templates (sources). $h_{rt}(f)$ corresponds to the time-varying filter associated to the source r. The decomposition thus benefits from the versatility of the source/filter model proved well suited for numerous sound objects.

B. AutoRegressive Moving Average (ARMA) Modeling

$h_{rt}(f)$ is parameterized following the general ARMA model:

$$h_{rt}^{ARMA}(f) = \sigma_{rt}^2 \frac{\left| \sum_{q=0}^{Q} b_{rt}^q e^{-i2\pi\nu_f q} \right|^2}{\left| \sum_{p=0}^{P} a_{rt}^p e^{-i2\pi\nu_f p} \right|^2}$$

where $\nu_f = \frac{f-1}{2(F-1)}$ is the normalized frequency associated to frequency index $f \in \{1, ..., F\}$ (as audio signal are real valued, we only consider frequencies between 0 and the Nyquist frequency). b_{rt}^q are the coefficients of the MA part of the filter and a_{rt}^p those of the AR part. σ_{rt}^2 is the global gain of the filter: in order to avoid identifiability problems, the first coefficient of all filters is imposed to be equal to 1. For $P = Q = 0$, $h_{rt}^{ARMA}(f)$ no longer depends on f and the decomposition corresponds to a standard NMF with temporal activations σ_{rt}^2.

Defining $\mathbf{a}_{rt} = (a_{rt}^0, \ldots, a_{rt}^P)^T$ and $\mathbf{b}_{rt} = (b_{rt}^0, \ldots, b_{rt}^Q)^T$, time/frequency activations can be rewritten as:

$$h_{rt}^{ARMA}(f) = \sigma_{rt}^2 \frac{\mathbf{b}_{rt}^T \mathbf{T}(\nu_f) \mathbf{b}_{rt}}{\mathbf{a}_{rt}^T \mathbf{U}(\nu_f) \mathbf{a}_{rt}}$$

where $\mathbf{T}(\nu)$ is the $(Q+1)\times(Q+1)$ Toeplitz matrix with $[\mathbf{T}(\nu)]_{pq}=\cos(2\pi\nu(p-q))$ and $\mathbf{U}(\nu)$ is similar to $\mathbf{T}(\nu)$ but of dimension $(P+1)\times(P+1)$. MA only and AR only models are included by respectively taking $P=0$ and $Q=0$. It is worth noting that $h_{rt}^{ARMA}(f)$ is always non-negative while there exists no non-negativity constraint on $\mathbf{b}_{rt}^q$ or on $\mathbf{a}_{rt}^p$.

The parameterized power spectrogram given in equation (3) then becomes:

$$\hat{V}_{ft}=\sum_{r=1}^{R} w_{fr}\sigma_{rt}^2\frac{\mathbf{b}_{rt}^T\mathbf{T}(\nu_f)\mathbf{b}_{rt}}{\mathbf{a}_{rt}^T\mathbf{U}(\nu_f)\mathbf{a}_{rt}} \tag{4}$$

III. Algorithm

We choose a general β-divergence cost function:

$$\mathcal{C}(\mathbf{W},\mathbf{A},\mathbf{B},\mathbf{\Sigma})=\sum_{f=1}^{F}\sum_{t=1}^{T} d_\beta(V_{ft},\hat{V}_{ft})$$

with $[\mathbf{W}]_{fr}=w_{fr}$, $[\mathbf{\Sigma}]_{rt}=\sigma_{rt}^2$, $[\mathbf{A}]_{rtp}=a_{rt}^p$ and $[\mathbf{B}]_{rtq}=b_{rt}^q$ and the expression of d_β is given in equation (2).

The partial derivative of the cost function wrt any variable θ (θ being any coefficient of $\mathbf{W}$, $\mathbf{\Sigma}$, $\mathbf{A}$ or $\mathbf{B}$) is:

$$\frac{\partial\mathcal{C}(\mathbf{W},\mathbf{A},\mathbf{B},\mathbf{\Sigma})}{\partial\theta}=\sum_{f=1}^{F}\sum_{t=1}^{T}\hat{V}_{ft}^{\beta-2}(\hat{V}_{ft}-V_{ft})\frac{\partial\hat{V}_{ft}}{\partial\theta} \tag{5}$$

The expression of the gradient of $\mathcal{C}$ wrt a vector θ of several coefficients of $\mathbf{A}$ or $\mathbf{B}$ is the same, replacing the partial derivative by a gradient ∇_θ in (5).

This leads to update rules for a multiplicative gradient descent algorithm similar to those used in [7], [21], [15]. In such an iterative algorithm, the update rule associated to one of the parameters is obtained from the partial derivative of the cost function wrt this parameter, written as a difference of two positive terms: $\frac{\partial\mathcal{C}}{\partial\theta}=G_\theta-F_\theta$

The update rule for θ is then:

$$\theta\leftarrow\theta\times\frac{F_\theta}{G_\theta} \tag{6}$$

This rule ensures that θ remains non-negative, becomes constant if the partial derivative is zero and evolves in the opposite direction of the partial derivative (thus in the descent direction).

A. Update of frequency templates:

We derive multiplicative update rules for w_{fr} from the expression of the partial derivative of the cost function with respect to w_{fr}.

The partial derivative of the parameterized spectrogram defined in equation (4) with respect to $w_{f_0r_0}$ is:

$$\frac{\partial\hat{V}_{ft}}{\partial w_{f_0r_0}}=h_{r_0t}^{ARMA}(f_0)\delta_{f_0f}$$

where δ_{f_0f} is a Kronecker delta.

Then, by replacing this expression in equation (5) with $\theta=w_{fr_0}$, we obtain the partial derivative of the cost function with respect to $w_{f_0r_0}$:

$$\frac{\partial\mathcal{C}(\mathbf{W},\mathbf{A},\mathbf{B},\mathbf{\Sigma})}{\partial w_{f_0r_0}}=\sum_{t=1}^{N}h_{r_0t}^{ARMA}(f_0)\hat{V}_{f_0t}^{\beta-2}(\hat{V}_{f_0t}-V_{f_0t})$$

This derivative is written as a difference of two positive terms:

$$G_{w_{f_0r_0}}=\sum_{t=1}^{N}h_{r_0t}^{ARMA}(f_0)\hat{V}_{f_0t}^{\beta-1}$$

and

$$F_{w_{f_0r_0}}=\sum_{t=1}^{N}h_{r_0t}^{ARMA}(f_0)\hat{V}_{f_0t}^{\beta-2}V_{f_0t}$$

Then the update rule of $w_{f_0r_0}$ is:

$$w_{f_0r_0}\leftarrow w_{f_0r_0}\frac{F_{w_{f_0r_0}}}{G_{w_{f_0r_0}}} \tag{7}$$

B. Update of temporal activation gain:

In the same way as for w_{fr}, we derive multiplicative update rules for σ_{rt}^2 from the expression of the partial derivative of the cost function with respect to σ_{rt}^2.

The partial derivative of the parameterized spectrogram defined in equation (4) with respect to $\sigma_{r_0t_0}^2$ is:

$$\frac{\partial\hat{V}_{ft}}{\partial\sigma_{r_0t_0}^2}=w_{fr_0}\frac{\mathbf{b}_{r_0t_0}^T\mathbf{T}(\nu_f)\mathbf{b}_{r_0t_0}}{\mathbf{a}_{r_0t_0}^T\mathbf{U}(\nu_f)\mathbf{a}_{r_0t_0}}\delta_{t_0t}$$

where δ_{t_0t} is a Kronecker delta.

Then, by substituting this expression into equation (5) with $\theta=\sigma_{r_0t_0}^2$, we obtain the partial derivative of the cost function with respect to $\sigma_{r_0t_0}^2$:

$$\frac{\partial\mathcal{C}(\mathbf{W},\mathbf{A},\mathbf{B},\mathbf{\Sigma})}{\partial\sigma_{r_0t_0}^2}=\sum_{f=1}^{M}\frac{w_{fr_0}}{\sigma_{r_0t_0}^2}h_{r_0t_0}^{ARMA}(f)\hat{V}_{ft_0}^{\beta-2}(\hat{V}_{ft_0}-V_{ft_0})$$

This derivative is written as a difference of two positive terms:

$$G_{\sigma_{r_0t_0}^2}=\sum_{f=1}^{M}\frac{w_{fr_0}}{\sigma_{r_0t_0}^2}h_{r_0t_0}^{ARMA}(f)\hat{V}_{ft_0}^{\beta-1}$$

and

$$F_{\sigma_{r_0t_0}^2}=\sum_{f=1}^{M}\frac{w_{fr_0}}{\sigma_{r_0t_0}^2}h_{r_0t_0}^{ARMA}(f)\hat{V}_{ft_0}^{\beta-2}V_{ft_0}$$

Then the update rule of $\sigma_{r_0t_0}^2$ is:

$$\sigma_{r_0t_0}^2\leftarrow\sigma_{r_0t_0}^2\frac{F_{\sigma_{r_0t_0}^2}}{G_{\sigma_{r_0t_0}^2}} \tag{8}$$

We can notice that when $Q=0$ and $P=0$ (i.e. there is no filter), the update rules of w_{fr} and σ_{rt}^2 are the same as the ones given in [21] which are the standard NMF rules for a β-divergence cost function where w_{fr} corresponds to frequency templates and σ_{rt}^2 to temporal activations.

C. Update of filters

The update rules of the coefficients of the filters are derived in a similar way, but the updates are not element-wise, but rather "vector-wise": we derive an update rule for each $\mathbf{b}_{rt}$ and for each $\mathbf{a}_{rt}$.

Update of $\mathbf{b}_{rt}$: The gradient of the parameterized spectrogram $\hat{V}_{ft}$ wrt $\mathbf{b}_{r_0t_0}$ is:

$$\nabla_{\mathbf{b}_{r_0t_0}} \hat{V}_{ft} = \delta_{t_0t} \frac{2w_{fr_0}\sigma^2_{r_0t_0}}{\mathbf{a}^T_{r_0t_0}\mathbf{U}(\nu_f)\mathbf{a}_{r_0t_0}} \mathbf{T}(\nu_f)\mathbf{b}_{r_0t_0}$$

Then, by substituting this expression into equation (5) with $\theta = \mathbf{b}_{r_0t_0}$, we obtain the gradient of the cost function wrt $\mathbf{b}_{r_0t_0}$:

$$\begin{aligned}
\nabla_{\mathbf{b}_{r_0t_0}} \mathcal{C} &= 2\sum_{f=1}^{F} \frac{w_{fr_0}\sigma^2_{r_0t_0}\hat{V}^{\beta-2}_{ft_0}(\hat{V}_{ft_0} - V_{ft_0})}{\mathbf{a}^T_{r_0t_0}\mathbf{U}(\nu_f)\mathbf{a}_{r_0t_0}} \mathbf{T}(\nu_f)\mathbf{b}_{r_0t_0} \\
&= 2\sigma^2_{r_0t_0}\left(\mathbf{R}_{r_0t_0} - \mathbf{R}'_{r_0t_0}\right)\mathbf{b}_{r_0t_0}
\end{aligned}$$

where:

$$\mathbf{R}_{r_0t_0} = \sum_{f=1}^{F} \frac{w_{fr_0}\hat{V}^{\beta-1}_{ft_0}}{\mathbf{a}^T_{r_0t_0}\mathbf{U}(\nu_f)\mathbf{a}_{r_0t_0}} \mathbf{T}(\nu_f)$$

$$\mathbf{R}'_{r_0t_0} = \sum_{f=1}^{F} \frac{w_{fr_0}\hat{V}^{\beta-2}_{ft_0}V_{ft_0}}{\mathbf{a}^T_{r_0t_0}\mathbf{U}(\nu_f)\mathbf{a}_{r_0t_0}} \mathbf{T}(\nu_f)$$

Both matrices $\mathbf{R}_{r_0t_0}$ and $\mathbf{R}'_{r_0t_0}$ are positive definite under mild assumptions: these matrices are clearly positive semi-definite and it can easily be shown that if there are at least $Q+1$ different indexes f such that $w_{fr_0}\hat{V}^{\beta}_{ft_0} \neq 0$ then $R_{r_0t_0}$ is non-singular (for $R'_{r_0t_0}$, the assumption is very similar). This assumption is always true in practice as long as the frame with index t_0 is not a null vector (*i.e.* with all samples equal to 0): in this particular case, the decomposition is trivial, and the global gains $\sigma_{r_0t_0}$ should be equal to 0.

Then, we follow the approach given in [22] and derive the following update rule for the MA part of the filter:

$$\mathbf{b}_{r_0t_0} \leftarrow \mathbf{R}^{-1}_{r_0t_0}\mathbf{R}'_{r_0t_0}\mathbf{b}_{r_0t_0} \tag{9}$$

As $\mathbf{R}_{r_0t_0}$ and $\mathbf{R}'_{r_0t_0}$ are both non singular, $\mathbf{R}^{-1}_{r_0t_0}$ is well defined and $\mathbf{b}_{r_0t_0}$ is ensured to never be zero.

Update of $\mathbf{a}_{rt}$:

The update rules of $\mathbf{a}_{rt}$ are derived in the same way as for $\mathbf{b}_{rt}$. The partial gradient of the parameterized spectrogram $\hat{V}_{ft}$ with respect to $\mathbf{a}_{r_0t_0}$ is:

$$\nabla_{\mathbf{a}_{r_0t_0}} \hat{V}_{ft} = -2\delta_{tt_0} w_{fr_0} \frac{h^{ARMA}_{r_0t_0}(f)}{\mathbf{a}^T_{r_0t_0}\mathbf{U}(\nu_f)\mathbf{a}_{r_0t_0}} \mathbf{U}(\nu_f)\mathbf{a}_{r_0t_0}$$

Then, by substituting this expression into equation (5) with $\theta = \mathbf{a}_{r_0t_0}$, we obtain the partial gradient of the cost function with respect to $\mathbf{a}_{r_0t_0}$:

$$\nabla_{\mathbf{a}_{r_0t_0}} \mathcal{C}(\mathbf{W}, \mathbf{A}, \mathbf{B}, \mathbf{\Sigma}) = 2\sigma^2_{r_0t_0}\left(\mathbf{S}'_{r_0t_0} - \mathbf{S}_{r_0t_0}\right)\mathbf{a}_{r_0t_0}$$

where:

$$\mathbf{S}_{r_0t_0} = \sum_{f=1}^{M} w_{fr_0}\hat{V}^{\beta-1}_{ft_0} \frac{\mathbf{b}^T_{r_0t_0}\mathbf{T}(\nu_f)\mathbf{b}_{r_0t_0}}{(\mathbf{a}^T_{r_0t_0}\mathbf{U}(\nu_f)\mathbf{a}_{r_0t_0})^2} \mathbf{U}(\nu_f)$$

and

$$\mathbf{S}'_{r_0t_0} = \sum_{f=1}^{M} w_{fr_0}\hat{V}^{\beta-2}_{ft_0}V_{ft_0} \frac{\mathbf{b}^T_{r_0t_0}\mathbf{T}(\nu_f)\mathbf{b}_{r_0t_0}}{(\mathbf{a}^T_{r_0t_0}\mathbf{U}(\nu_f)\mathbf{a}_{r_0t_0})^2} \mathbf{U}(\nu_f)$$

Both matrices $\mathbf{S}_{r_0t_0}$ and $\mathbf{S}'_{r_0t_0}$ are positive definite under mild assumptions.

Thus we derive the following update rule for the AR part of the filter:

$$\mathbf{a}_{r_0t_0} \leftarrow \mathbf{S}'^{-1}_{r_0t_0}\mathbf{S}_{r_0t_0}\mathbf{a}_{r_0t_0} \tag{10}$$

D. Description of the algorithm

The update rules (7), (8), (9) and (10) are applied successively to all the coefficients of $\mathbf{W}$, all the coefficients of $\mathbf{\Sigma}$, all the coefficients of $\mathbf{B}$ and all the coefficients of $\mathbf{A}$. Between the updates of each of these matrices (and tensors), the parameterized spectrogram $\hat{\mathbf{V}}$ is recomputed: as for the standard NMF algorithm, this recomputation between each update is necessary to ensure the convergence.

Identifiability: As for the standard NMF, the decomposition (4) which minimizes the cost function is not unique. To cope with identifiability issues, we impose constraints on $\mathbf{W}$, $\mathbf{\Sigma}$, $\mathbf{B}$ and $\mathbf{A}$:

- for all r and t, we impose that $\mathbf{b}_{rt}$ and $\mathbf{a}_{rt}$ (considered as polynomials) have all their roots inside the unit circle.
- for all r, we impose $\|\mathbf{w}_r\| = 1$ for some norm $\|.\|$.
- for all r and t, we impose $b^0_{rt} = 1$ and $a^0_{rt} = 1$.

Thus, at the end of each iteration of our algorithm, we transform $\mathbf{b}_{r,t}$ and $\mathbf{a}_{r,t}$ by replacing roots outside the unit circle by the conjugate of their inverse and accordingly adapting the gain, normalize each column of $\mathbf{W}$, divide $\mathbf{b}_{r,t}$ and $\mathbf{a}_{r,t}$ by their first coefficient and update $\mathbf{\Sigma}$ in order not to change $\hat{V}_{ft}$ by these modifications. All these transformations have no influence on the values of the parameterized spectrogram.

Another choice of filters normalization has been tested: rather than imposing $b^0_{rt} = 1$ and $a^0_{rt} = 1$, we can impose for all (r,t), $\sum_f \mathbf{b}^T_{rt}\mathbf{T}(\nu_f)\mathbf{b}_{rt} = 1$ and $\sum_f \mathbf{a}^T_{rt}\mathbf{U}(\nu_f)\mathbf{a}_{rt} = 1$. It corresponds to a power normalization and then it is more meaningful.

Our algorithm is detailed in Algorithm 1. In the remainder of the article, we will refer to this algorithm by the expression "source/filter factorization".

E. Dimensionality

Since our algorithm is a dimensionality reduction technique like NMF, one should take care of the dimension of the decomposition provided. The dimension of the original data is FT. In the standard NMF with R atoms, the dimension of the parameterized spectrogram is $\dim \mathbf{W} + \dim \mathbf{H} = R(F+T)$ (where $\dim \mathbf{\Phi}$ stands for the number of coefficients of matrix or tensor $\mathbf{\Phi}$). With our algorithm, the dimension of the parameters is: $\dim \mathbf{W} + \dim \mathbf{\Sigma} + \dim \mathbf{A} + \dim \mathbf{B} = RF + RT(P+Q+1)$. Thus, one should have $RF + RT(P+Q+1) \ll FT$, i.e.

Algorithm 1 Source/filter spectrogram factorization

Input: **V**, R, Q, P, n_{iter}, β
Output: **W**, $\mathbf{\Sigma}$, **B**, **A**
  Initialize **W**, $\mathbf{\Sigma}$ with non-negative values
  Initialize **B**, **A** with flat filters
  for j = 1 to n_{iter} **do**
    compute $\hat{\mathbf{V}}$
    for all f and r **do**
      compute $F_{w_{fr}}$ and $G_{w_{fr}}$
      $w_{fr} \leftarrow w_{fr} \frac{F_{w_{fr}}}{G_{w_{fr}}}$
    end for
    compute $\hat{\mathbf{V}}$
    for all r and t **do**
      compute $F_{\sigma^2_{rt}}$ and $G_{\sigma^2_{rt}}$
      $\sigma^2_{rt} \leftarrow \sigma^2_{rt} \frac{F_{\sigma^2_{rt}}}{G_{\sigma^2_{rt}}}$
    end for
    compute $\hat{\mathbf{V}}$
    for all r and t **do**
      compute $\mathbf{R}_{rt}$, $\mathbf{R}'_{rt}$, $\mathbf{S}_{rt}$ and $\mathbf{S}'_{rt}$
      $\mathbf{b}_{rt} \leftarrow \mathbf{R}_{rt}^{-1}\mathbf{R}'_{rt}\mathbf{b}_{rt}$
      $\mathbf{a}_{rt} \leftarrow \mathbf{S}_{rt}'^{-1}\mathbf{S}_{rt}\mathbf{a}_{rt}$
    end for
    bring back roots of all filters inside the unit circle
    divide the coefficients of all filters by the first one
    normalize **W**
    update appropriately $\mathbf{\Sigma}$
  end for

$R \ll T$ and $R(P+Q+1) \ll F$, so P and Q must be small. As in practice $F \leq T$, the condition to be respected is $R(P+Q+1) \ll F$.

One should notice that our decomposition allows a significant reduction of the number of atoms R needed to accurately fit the data when the parts of the sounds present strong spectral variations. Then the total dimension of the parameters obtained with our decomposition remains comparable to the one obtained with standard NMF as will be shown in section IV.

Besides, one should notice that a large number of the coefficients of the filters are useless and therefore do not need to be retained: when the global gain of one of the filter σ_{rt} becomes close to zero, these coefficients ($\mathbf{b}_{rt}$ and $\mathbf{a}_{rt}$) become meaningless and then are useless in the decomposition and can be removed without affecting the values of the parameterized spectrogram.

Finally, in the decomposition (4), all atoms are associated to filters of the same order, but it is also possible to implement a larger model where filters do not have the same characteristics for all atoms. This larger model is not presented in this paper for readability reasons.

F. Computational complexity

The computational complexity of one iteration of source/filter factorization depends on P, Q, R, F and T. The computational complexity of each step of the algorithm is given here:

- Computation of $\hat{\mathbf{V}}$: $\mathcal{O}\left((P+Q)\,RFT\right)$ operations.
- Update of $\mathbf{W}$ and $\mathbf{\Sigma}$: $\mathcal{O}\left(RFT\right)$ operations each.
- Update of $\mathbf{B}$: $\mathcal{O}\left(RT(FP+P^3)\right)$ operations.
- Update of $\mathbf{A}$: $\mathcal{O}\left(RT(FQ+Q^3)\right)$ operations.
- Normalization/stabilization: $\mathcal{O}\left(RT(F+P^3+Q^3)\right)$ operations.

The total complexity of a single iteration of the algorithm is then $\mathcal{O}\left((P+Q)\,RFT\right)$ operations. With our current implementation in Matlab, 100 iterations of our algorithm applied to a 1025×550 spectrogram (corresponding to a 6.5s signal sampled at $f_s = 22050$Hz with 2048-sample-long windows and 75% overlap) with $R = 10$ atoms, $P = 2$ and $Q = 2$ last about 300s (on an Intel®Core™2 Duo E8400 @3.00GHz, with Matlab's multithreaded math libraries). In comparison, 100 iterations of standard NMF with the same spectrogram, the same parameters ($R = 10$, but $P = 0$ and $Q = 0$) on the same computer last about 9s: our algorithm appears to be slower than standard NMF. However, this comparison puts at a disadvantage our algorithm, which is designed to work with fewer atoms than standard NMF: in this case our algorithm deals with many more parameters than NMF. If we compare execution times with the same number of parameters, the difference is smaller: for the same spectrogram, 100 iterations of our algorithm with $R = 2$ atoms, $P = 2$ and $Q = 2$ (i.e. with the dimensionality of a standard NMF with $R = 10$) last about 60s.

About a third of the computation time is due to the computation of the roots of all filters (considered as polynomials) during the stabilization process. Some improvements could be made by considering another regularization method. The inversion of the matrices R_{rt} and S'_{rt} (a bit more than 10% of the total computation time) and the computation the frequency response of all filters (slighly less than 10% of the total computation time) are also very costly.

G. Practical implementation and choice of β

We empirically observed the monotonic decrease of the cost function and the convergence of the algorithm for $0 \leq \beta \leq 2$ over a large set of tests: this decrease and this convergence are illustrated in particular cases in section IV-D.

However, the algorithm is unstable for $1 \leq \beta \leq 2$: some numerical instabilities appear while poles of the filters come close to the unit circle. These instabilities are very common when β becomes close to 2 (Euclidean distance); however, the strong dynamics of audio data is better fitted when β becomes close to 0 as stated in [21]. In order to avoid these instabilities, we limit the modulus of poles: the monotonic decrease of the cost function is no longer observed but it permits to avoid non desirable behavior of the decomposition (very resonant order 2 filters only permit to fit one partial).

For the examples in the next section, we chose $\beta = 0.5$ since the numerical instabilities were almost negligible and the results were more accurate than with $\beta = 0$ (Itakura-Saito divergence).

IV. EXAMPLES

In this section several experiments are presented to show that our algorithm is well adapted to decompose sounds having

strong spectral variations. All the spectrograms used in these experiments are power spectrograms obtained from recorded signals by means of a short time Fourier transform (STFT).

Algorithms (standard NMF and source/filter factorization) were initialized with random values (except for the filters which were initially flat) and were run until apparent convergence. The algorithms have been rerun with 100 different initializations in order to maximize the chances to come close to a "good" minimum. Despite these different starting points, the reached solutions were similar in terms of qualitative aspect of the reconstructed spectrograms.

A. Didgeridoo

1) Description of the excerpt: In this section our algorithm is applied to a short didgeridoo excerpt. The didgeridoo is an ethnic wind instrument from northern Australia. It makes a continuous modulated sound produced by the vibrations of the lips. The modulations result from the mouth and throat configuration with the help of which the player is able to control several resonances. Figure 1(a) represents the spectrogram of the excerpt: the sound produced is almost harmonic (with some noise) and a strong moving resonance appears in the spectrogram. We can thus consider that this signal is composed of a single event encompassing spectral variations, and try to decompose it with a single atom ($R = 1$). The sampling rate of the excerpt is $f_s = 11025Hz$. We chose a 1024-sample-long Hann window with 75% overlap for the STFT.

2) Experiment and results: The spectrogram of the excerpt is decomposed with a standard NMF algorithm for $R = 1$ atom and $R = 5$ atoms, and with source/filter factorization for $R = 1$ atom, with an order 3 AR modeling ($Q = 0$ and $P = 3$). Reconstructed spectrograms are respectively represented in figures 1(b), 1(c) and 1(d).

Although the didgeridoo is played alone in the analyzed spectrogram, the standard NMF needs many atoms to accurately decompose the power spectrogram. With 1 atom, NMF does not accurately represent the moving resonance (figure 1(b)). With 5 atoms, some spectral variations appear (figure 1(c)), but the resonance trajectory remains a bit unclear. Besides, the signal is not decomposed in a meaningful way (each atom is a part of the sound which has no perceptual meaning) and the dimensionality of the parameters is large ($FR + RT = 3290$).

In opposition to the standard NMF, source/filter factorization permits to accurately represent the spectral variability of the sound (figure 1(d)) with a single atom, keeping the dimensionality low ($FR + TR(Q + 1) = 1093$): the moving resonance of the original sound is well tracked, and the total error $\mathcal{C}$ is smaller than that of the standard NMF with $R = 5$. In this case, the decomposition is more efficient and relevant than the standard NMF.

B. Harpsichord

1) Description of the excerpt: In this section our algorithm is applied to a short harpsichord excerpt, composed of two different notes ($C2$ and $E\flat2$): first, the $C2$ is played alone, then the $E\flat2$, and at last, both notes are played simultaneously.

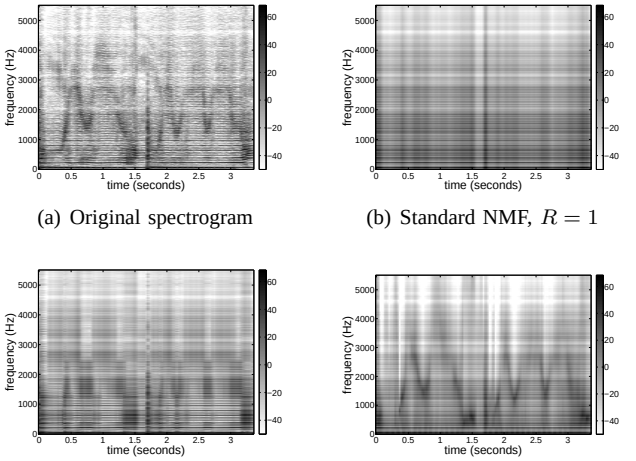

(a) Original spectrogram (b) Standard NMF, $R = 1$

(c) Standard NMF, $R = 5$ (d) Srce/filter factorization, $R = 1$

Fig. 1. Original power spectrogram of the extract of didgeridoo 1(a) and reconstructed spectrograms 1(b), 1(c) and 1(d)

The spectrogram of the extract is represented in figure 2(a). As for most of plucked string instruments, high frequency partials of a harpsichord tone decay faster than low frequency partials. This phenomenon clearly occurs in the L-shaped spectrograms of figure 2(a). The sampling rate of the excerpt is $f_s = 44100Hz$. We chose a 2048-sample-long Hann window with 75% overlap for the STFT.

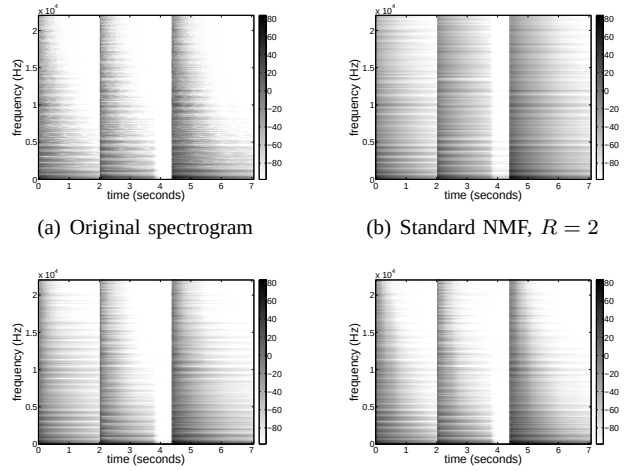

(a) Original spectrogram (b) Standard NMF, $R = 2$

(c) Standard NMF, $R = 6$ (d) Srce/filter factorization, $R = 2$

Fig. 2. Original power spectrogram of the extract of harpsichord 2(a) and reconstructed spectrograms 2(b), 2(c) and 2(d)

2) Experiment and results: The spectrogram of the excerpt was decomposed with a standard NMF algorithm for $R = 2$ atoms (1 atom per note) and $R = 6$ atoms, and with source/filter factorization for $R = 2$ atoms, with an ARMA modeling ($Q = 1$ and $P = 1$). Reconstructed spectrograms are respectively represented in figures 2(b), 2(c) and 2(d).

The standard NMF needs several atoms per note to accurately decompose the L-shaped power spectrograms: with only 2 atoms (1 per note played), the faster decay of high frequency content does not appear at all (figure 2(b)). With 6 atoms, the attenuation of high frequency partials appears (figure 2(c)),

but each atom is a part of a note spectrum and has no real perceptual meaning.

The ARMA modeling included in our algorithm leads to a good description of the overall spectrogram shape. 2 atoms (1 per note) are enough to accurately fit the original short time spectrum: each atom is harmonic (figure 3(a)) and corresponds to one note while the decay of high frequency partials is clearly well described by the ARMA modeling (see time/frequency activations $h_{rt}^{ARMA}(f)$ in figure 3(b)). The dimensionality of the data provided by our algorithm ($FR+TR(Q+P+1) = 5704$) remains lower than with a standard NMF with 6 atoms ($FR+RT = 9804$) and the global error $\mathcal{C}$ between the original and the reconstructed spectrogram is approximately the same as the one obtained with the standard NMF with $R = 6$.

Thus the decomposition provided by source/filter factorization seems to give a more meaningful representation of the given spectrogram than the one obtained with the standard NMF.

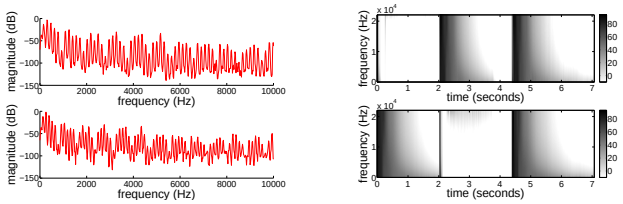

(a) Frequency templates (b) Time/frequency activations

Fig. 3. Source/filter decomposition ($R = 2$, $Q = 1$ and $P = 1$) of the power spectrogram of the harpsichord excerpt

C. Guitar with wah pedal

1) Description of the excerpt: In this section our algorithm is used to decompose a short extract of electric guitar processed by a wah pedal. The wah pedal (or wah-wah pedal) is a popular guitar effect which consists of a resonant filter, the resonant frequency of which is controlled by means of a foot pedal. This effect is named by emphasizing the resemblance with the human onomatopoeia "Wah". A single note of electric guitar processed by a moving wah pedal presents strong spectral variations and therefore cannot be well represented by a single atom in a standard NMF.

As a wah pedal is well modeled by an AR filter with two complex conjugates poles, we chose to decompose the extract with $Q = 0$ and $P = 2$. The chosen extract represented in figure 4(a) is composed of three different notes played successively (the first note is played a second time at the end of the extract). Each note can be viewed as a harmonic pattern which is filtered by a resonant filter, the resonant frequency of which varies between $400Hz$ and $1200Hz$: this resonance clearly appears in the power spectrogram. The sampling rate of the excerpt is $f_s = 11025Hz$. We chose a 1024-sample-long Hann window with 75% overlap for the STFT.

2) Experiment and results: As the analyzed sound presents strong spectral variations, the standard NMF needs many atoms to accurately decompose the power spectrogram. Thus one atom no longer corresponds to one note, and the decomposition does not correspond to the analysis that could

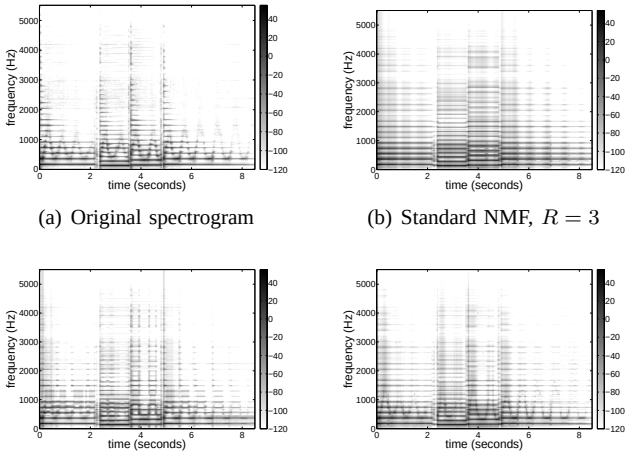

(a) Original spectrogram (b) Standard NMF, $R = 3$

(c) Standard NMF, $R = 10$ (d) Srce/filter factorization, $R = 3$

Fig. 4. Original power spectrogram of the extract of electric guitar processed by a wah pedal 4(a) and reconstructed spectrograms 4(b), 4(c) and 4(d)

be performed by a human listener. Figure 4(b) represents the power spectrogram reconstructed from the NMF of the original spectrogram with 3 atoms and figure 4(c) with 10 atoms. With 3 atoms, NMF is not able to track the resonance of the wah pedal. With 10 atoms, the resonance appears, but the signal is not well explained (each atom is a part of a note and then has no perceptual meaning) and the dimensionality is higher ($FR + RT = 8790$).

With source/filter factorization, the strong spectral variations of each note can be accurately represented in the filter activations taking an order 2 AR modeling ($Q = 0$ and $P = 2$) as shown in figure 4(d). Then 3 atoms (one for each note) are enough to correctly fit the original spectrogram. Indeed, the global β-divergence error between the original and the reconstructed spectrogram obtained with source/filter factorization is approximately the same as the one obtained with standard NMF with 10 atoms; this β-divergence (obtained with source/filter factorization) is also approximately half that obtained with standard NMF with 3 atoms. Each atom is harmonic and corresponds to one note, and the resonance of the wah pedal clearly appears. The dimensionality of the representation obtained with source/filter factorization remains about half that of NMF with 10 atoms: $MR+R(Q+1)N = 4833$. The decomposition provided by our algorithm distinguishes a stationary spectrum representing "average" guitar sounds (contained in $\mathbf{W}$) from the non-stationary effect due to the wah pedal (described by time/frequency activations). The 3 frequency templates (columns of $\mathbf{W}$) obtained are represented in figure 5(a): each template is harmonic with its own fundamental frequency, thus it corresponds to a note (standard NMF with 3 atoms provides similar templates). The time/frequency activations ($h_{rt}^{ARMA}(f)$) are represented in figure 5(b): the resonance of the Wah pedal appears clearly where the notes are played. Thus the decomposition provided by our algorithm seems to give a more meaningful representation of the given spectrogram.

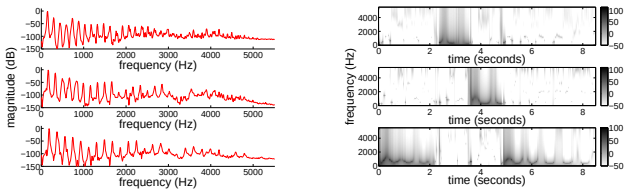

(a) Frequency templates (b) Time/frequency activations

Fig. 5. Source/filter decomposition ($R = 3$ and $P = 2$) of the power spectrogram of the Wah processed guitar excerpt

D. Convergence of the algorithm

The evolution of the cost function over iterations for source filter/factorization is represented in figure 6 with 8 different random initializations, for the decomposition of excerpts presented in sections IV-B (harpsichord excerpt) and IV-C (Wah guitar excerpt). The value of the β-divergence is represented after each iteration. Figures show a monotonic decrease of the cost function and an apparent convergence. In figure 6(a), all initializations lead to the same final value of the cost function and the shape of the evolution is very similar for all initializations. On the other hand, in figure 6(b), all initializations do not lead to the same value of the cost function, showing that multi-initialization is useful.

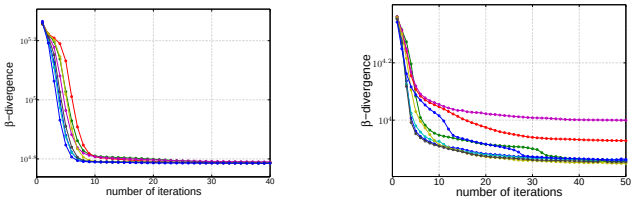

(a) Harpsichord excerpt (b) Wah guitar excerpt

Fig. 6. Evolution of the cost function over iterations (decomposition of excerpts in sections IV-B and IV-C)

V. Conclusion and Future Work

In this paper, we proposed a new iterative algorithm which is an extended version of the Non-negative Matrix Factorization based on a source/filter representation. We showed that this representation is particularly suitable to efficiently and meaningfully decompose non stationary sound objects including noticeable spectral variations.

In the future, this extended decomposition should be further developed to deal with small pitch variations (like vibrato), for instance using a constant-Q spectrogram like in [23], [24]. Besides, we plan to introduce harmonic constraints in the basis spectra following [25], [26], [27]. Finally, we plan to investigate the introduction of continuity constraints between filters from one frame to another following the approach given in [12], [26], [21].

References

[1] Stéphane Mallat and Zhifeng Zhang, "Matching pursuit with time-frequency dictionnaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, December 1993.

[2] Rémi Gribonval and Emmanuel Bacry, "Harmonic decomposition of audio signals with matching pursuit," *IEEE Transactions on Signal Processing*, vol. 51, no. 1, pp. 101–111, January 2003.

[3] Pierre Leveau, Emmanuel Vincent, Gaël Richard, and Laurent Daudet, "Instrument-specific harmonic atoms for mid-level music representation," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 16, no. 1, pp. 116–128, January 2008.

[4] P. Comon, "Independent component analysis, a new concept," *Signal Processing special issue Higher-Order Statistics*, vol. 36, pp. 287–314, April 1994.

[5] S. Makino, S. Araki, R. Mukai, and H. Sawada, "Audio source separation based on independent component analysis," in *International Symposium on Circuits and Systems*, Vancouver, Canada, May 2004, pp. V–668–V–671.

[6] Samer A. Abdallah and Mark D. Plumbley, "Unsupervised analysis of polyphonic music by sparse coding," *IEEE Transactions on neural Networks*, vol. 17, no. 1, pp. 179 – 196, January 2006.

[7] D.D. Lee and H.S. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, no. 6755, pp. 788–791, October 1999.

[8] P. Smaragdis and J.C. Brown, "Non-negative matrix factorization for polyphonic music transcription," in *Workshop on Applications of Signal Processing to Audio and Acoustics*, New Paltz, NY, USA, October 2003, pp. 177 – 180.

[9] J. Paulus and T. Virtanen, "Drum transcription with non-negative spectrogram factorization," in *European Signal Processing Conference (EUSIPCO)*, Antalya, Turkey, September 2005.

[10] N. Bertin, R. Badeau, and G. Richard, "Blind signal decompositions for automatic transcription of polyphonic music: NMF and K-SVD on the benchmark," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Honolulu, Hawaii, USA, April 2007, vol. 1, pp. I–65 – I–68.

[11] A. Cichocki, R. Zdunek, and S. ichi Amari, "New algorithms for non-negative matrix factorization in applications to blind source separation," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Toulouse, France, May 2006, vol. 5, pp. 621 – 625.

[12] T. Virtanen, "Monaural sound source separation by nonnegative matrix factorization with temporal continuity," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 15, no. 3, pp. 1066–1074, March 2007.

[13] Alexey Ozerov and Cédric Févotte, "Multichannel nonnegative matrix factorization in convolutive mixtures for audio source separation," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 18, no. 3, pp. 550–563, 2010.

[14] Derry FitzGerald, Matt Cranitch, and Eugene Coyle, "Non-negative tensor factorisation for sound source separation," in *Irish Signals and Systems Conference*, 2005.

[15] P. Smaragdis, "Non-negative matrix factor deconvolution; extraction of multiple sound sources from monophonic inputs," in *Conference on Independent Component Analysis and Blind Source Separation (ICA)*, Grenada, Spain, September 2004, pp. 494–499.

[16] J.-L. Durrieu, G. Richard, and B. David, "An iterative approach to monaural musical mixture de-soloing," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Taipei, Taiwan, April 2009, pp. 105 – 108.

[17] J.-L. Durrieu, G. Richard, and B. David, "Singer melody extraction in polyphonic signals using source separation methods," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Las Vegas, Nevada, USA, August 2008, pp. 169 – 172.

[18] I. Dhillon and S. Sra, "Generalized nonnegative matrix approximations with Bregman divergences," in *Neural Information Processing Systems conference (NIPS)*, Y. Weiss, B. Schölkopf, and J. Platt, Eds., pp. 283–290. MIT Press, Cambridge, MA, 2006.

[19] R. Kompass, "A generalized divergence measure for nonnegative matrix factorization," *Neural Computation*, vol. 19, no. 3, pp. 780–791, March 2007.

[20] A. Cichocki, R. Zdunek, and S. Amari, "Csiszars divergences for non-negative matrix factorization: Family of new algorithms," in *Conference on Independent Component Analysis and Blind Source Separation (ICA)*, Charleston, SC, USA, March 2006, pp. 32 – 39.

[21] C. Févotte, N. Bertin, and J.-L. Durrieu, "Nonnegative matrix factorization with the Itakura-Saito divergence. With application to music analysis," *Neural Computation*, vol. 11, no. 3, pp. 793–830, March 2009.

[22] R. Badeau and B. David, "Weighted maximum likelihood autoregressive and moving average spectrum modeling," in *International Conference on*

Acoustics, Speech and Signal Processing (ICASSP), Las Vegas, Nevada, USA, March 2008, pp. 3761 – 3764.

[23] P. Smaragdis, B. Raj, and M. Shashanka, "Sparse and shift-invariant feature extraction from non-negative data," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Las Vegas, Nevada, USA, March 2008, pp. 2069 – 2072.

[24] M. Schmidt and M. Mørup, "Nonnegative matrix factor 2-D deconvolution for blind single channel source separation," in *Conference on Independent Component Analysis and Blind Source Separation (ICA)*, Paris, France, April 2006, vol. 3889 of *Lecture Notes in Computer Science (LNCS)*, pp. 700–707, Springer.

[25] Emmanuel Vincent, Nancy Bertin, and Roland Badeau, "Harmonic and inharmonic nonnegative matrix factorization for polyphonic pitch transcription," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Las Vegas, Nevada, USA, March 2008, pp. 109 – 112.

[26] Nancy Bertin, Roland Badeau, and Emmanuel Vincent, "Enforcing harmonicity and smoothness in bayesian non-negative matrix factorization applied to polyphonic music transcription," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 18, no. 3, pp. 538–549, 2010.

[27] Emmanuel Vincent, Nancy Bertin, and Roland Badeau, "Adaptive harmonic spectral decomposition for multiple pitch estimation," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 18, no. 3, pp. 528–537, 2010.

Dépôt légal : 2011 – 1er trimestre
Imprimé à TELECOM ParisTech (Ecole Nationale Supérieure des Télécommunications) – Paris
ISSN 0751-1345 ENST D (Paris) (France 1983-9999)

TELECOM ParisTech

Institut TELECOM (Groupe des écoles des télécommunications) - membre de ParisTech

46, rue Barrault - 75634 Paris Cedex 13 - Tél. + 33 (0)1 45 81 77 77 - www.telecom-paristech.fr

Département TSI

© Institut TELECOM - Télécom ParisTech 2011