



ADVANCES IN LEARNING-BASED IMAGE RESTORATION



Journée thématique soutenue par RT-MAIAGES et GdR-IASIS

Institut Henri Poincaré, December, 9th, 2024

- 9h15 - OPENING
- 9h20 - **Jean-Christophe Pesquet** (CentraleSupélec)
A focus on Lipschitz properties of neural networks
- 10h - **Pierre Weiss** (CNRS, Université de Toulouse)
Statistical comparison of plug-and-play and unrolled networks
- 10h20 - **Marien Renaud** (Institut de Mathématiques de Bordeaux)
Equivariant Denoisers for Image Restoration
- 10h40 - COFFEE BREAK
- 11h10 - **Nelly Pustelnik** (CNRS, ENS Lyon)
On the stability and performance of plug-and-play methods in image reconstruction
- 11h50 - **Antoine Guennec** (Institut de Mathématiques de Bordeaux)
Joint structure-texture image decomposition using a plug-and-play framework
- 12h10 - **Sophie Carneiro Esteves** (CREATIS Lyon)
A plug-and-play framework for curvilinear structure segmentation based on a learned reconnecting regularization
- 12h30 - LUNCH BREAK
- 14h - **Pierre Chainais** (Centrale Lille Institut, CRISAL)
Plug-and-Play Split Gibbs Sampler: embedding deep generative priors in Bayesian inference
- 14h40 - **Maud Biquard** (ISAE-Supaero, CNES)
Variational Bayes image restoration with compressive autoencoders applied to satellite image restoration
- 15h - **Hubert Leterme** (Ensicaen, CEA Paris-Saclay)
A plug-and-play approach with conformal predictions for weak lensing mass mapping
- 15h20 - COFFEE BREAK
- 15h50 - **Andrés Almansa** (CNRS, Université de Paris)
Posterior sampling in imaging with learnt priors: from Langevin to diffusion models
- 16h40 - **Corentin Vazia** (Université de Bretagne Sud)
Guidance of a diffusion model for material decomposition in photon-counting computed tomography
- 17h - **Liam Moroy** (ONERA)
Evaluating the Posterior Sampling Ability of Plug-and-Play Diffusion Methods in Sparse-View CT
- 17h20 - **Marcelo Pereyra** (Heriot-Watt University)
Tackling fundamental challenges in hypothesis testing in imaging inverse problems



ADVANCES IN LEARNING-BASED IMAGE RESTORATION



Journée thématique
soutenue par RT-MAIAGES et GdR-IASIS

Institut Henri Poincaré, 9 Décembre 2024

- 9h15 - OPENING
- 9h20 - **Jean-Christophe Pesquet** (CentraleSupélec)
A focus on Lipschitz properties of neural networks
- 10h - **Pierre Weiss** (CNRS, Université de Toulouse)
Statistical comparison of plug-and-play and unrolled networks

Plug-and-play, diffusion models and unrolled networks have emerged in the last ten years and are gaining a lot of popularity for the resolution of inverse problems. However, they are often studied by different partly hermetic communities. In this talk, I will summarize a few insights on their similarities and differences, strengths and limitations. In particular, we can relate plug and play methods to Maximum A Posteriori estimators, while unrolled methods rather coincide with Minimum Mean Square Estimators. This has some important consequences and should be taken into account when targeting a specific application.

- 10h20 - **Marien Renaud** (Institut de Mathématiques de Bordeaux)
Equivariant Denoisers for Image Restoration

One key ingredient of image restoration is to define a realistic prior on clean images to complete the missing information in the observation. State-of-the-art restoration methods rely on a neural network to encode this prior. Moreover, typical image distributions are invariant to some set of transformations, such as rotations or flips. However, most deep architectures are not designed to represent an invariant image distribution. Recent works have proposed to overcome this difficulty by including equivariance properties within a Plug-and-Play paradigm. In this presentation, we will present an unified framework named Equivariant Regularization by Denoising (ERED) based on equivariant denoisers and stochastic optimization. We will analyze the convergence of this algorithm and discuss its practical benefit.

- 10h40 - COFFEE BREAK
- 11h10 - **Nelly Pustelnik** (CNRS, ENS Lyon)
On the stability and performance of plug-and-play methods in image reconstruction

Plug-and-play algorithms constitute a popular framework for solving inverse imaging problems, utilizing a denoiser to implicitly define an image prior. By employing powerful pre-trained denoisers, these methods can address a broad spectrum of imaging tasks without the need for task-specific model training. However, plug-and-play approaches often suffer from instability,

which limits their versatility and results in suboptimal reconstruction quality. In this presentation, we will first assess the robustness of proposed proximal unfolded neural networks when plugged in a forward-backward algorithm for an image deblurring problem. Second, we show that enforcing equivariance to certain groups of transformations on the denoiser strongly improves the stability of the algorithm as well as its reconstruction quality. Third, we provide a multilevel PnP framework to improve the stability and performance in the context of inpainting.

- 11h50 - **Antoine Guennec** (Institut de Mathématiques de Bordeaux)

Joint structure-texture image decomposition using a plug-and-play framework

Afin de résoudre le problème de la séparation des images en un composant structure et un composant texture, nous introduisons un modèle de structure-texture joint. Plutôt que de considérer des fonctions de régularisation distinctes pour chacune des composantes, nous introduisons une unique fonction de régularisation qui les considère de manière couplée. Ce choix permet à notre modèle de mieux capturer les interactions entre la structure et la texture, tout en supprimant les paramètres à ajuster.

Afin de construire cette fonction de régularisation, nous nous appuyons sur le cadre plug-and-play. Nous entraînons le prior via des données synthétiques générées aléatoirement d'exemples de ce modèle. Nos résultats expérimentaux démontrent que ce modèle joint surpasse les approches traditionnelles basées sur des régularisations séparées. De plus, bien qu'entraîné sur des données synthétiques, notre modèle se révèle efficace pour décomposer des images naturelles et peut être appliqué pour résoudre des problèmes inverses, tels que l'inpainting.

- 12h10 - **Sophie Carneiro Esteves** (CREATIS Lyon)

A plug-and-play framework for curvilinear structure segmentation based on a learned reconnecting regularization

Les structures curvilinéaires sont présentes dans divers domaines du traitement d'images, tels qu'en imagerie médicale (vaisseaux sanguins, neurones, ...) ou encore en télédétection (routes, rivières, ...). Leur détection est cruciale pour de nombreuses applications. Nous proposons une méthode plug-and-play pour la segmentation des structures curvilinéaires, mettant l'accent sur la préservation de leur connectivité. Nous avons développé un algorithme pour générer des paires réalistes de structures connectées/déconnectées, afin d'entraîner un terme de régularisation reconnecteur à partir de données synthétiques. Une fois appris, ce modèle peut être intégré dans un schéma de segmentation variationnelle et utilisé pour segmenter des images de structures curvilinéaires sans nécessiter d'annotations. Nous démontrons l'intérêt de notre approche sur deux applications distinctes, dans des images 2D et 3D, et comparons ses résultats avec ceux d'approches classiques non supervisées et basées sur l'apprentissage profond. Les évaluations comparatives mettent en évidence les performances supérieures de notre méthode, montrant des améliorations significatives dans la préservation de la connectivité des structures curvilinéaires (environ 90 % en 2D et 70% en 3D). Enfin, nous illustrons la capacité de généralisation de notre méthode sur deux applications différentes : la segmentation de fissures présentes sur des routes et la segmentation de cellules cornéennes porcines.

- 12h30 - LUNCH BREAK [newpage](#)

- 14h - **Pierre Chainais** (Centrale Lille Institut, CRISAL)

Plug-and-Play Split Gibbs Sampler: embedding deep generative priors in Bayesian inference

This work introduces a stochastic plug-and-play (PnP) sampling algorithm that leverages variable splitting to efficiently sample from a posterior distribution. The algorithm based on split

Gibbs sampling (SGS) draws inspiration from the half quadratic splitting method (HQS) and the alternating direction method of multipliers (ADMM). It divides the challenging task of posterior sampling into two simpler sampling problems. The first problem depends on the likelihood function, while the second is interpreted as a Bayesian denoising problem that can be readily carried out by a deep generative model. Specifically, for an illustrative purpose, the proposed method is implemented in this paper using state-of-the-art diffusion-based generative models. Akin to its deterministic PnP-based counterparts, the proposed method exhibits the great advantage of not requiring an explicit choice of the prior distribution, which is rather encoded into a pre-trained generative model. However, unlike optimization methods (e.g., PnP-ADMM and PnP-HQS) which generally provide only point estimates, the proposed approach allows conventional Bayesian estimators to be accompanied by confidence intervals at a reasonable additional computational cost. Experiments on commonly studied image processing problems illustrate the efficiency of the proposed PnP-SGS sampling strategy. Its performance is compared to recent state-of-the-art optimization and sampling methods.

- 14h40 - **Maud Biquard** (ISAE-Supaero, CNES)

Variational Bayes image restoration with compressive autoencoders applied to satellite image restoration

The ability of neural networks to learn efficient image representations has been recently exploited to design powerful data-driven regularizers for image restoration. While state-of-the-art plug-and-play methods rely on an implicit regularization provided by neural denoisers, alternative Bayesian approaches consider Maximum A Posteriori (MAP) estimation in the latent space of a generative model, thus with an explicit regularization. We introduce a variational Bayes framework enabling to approximate the posterior distribution of the inverse problem within variational autoencoders (VAEs). First, the Variational Bayes Latent Estimation (VBLE) algorithm models the posterior only in the latent space of a VAE. This leads to fast and straightforward posterior sampling abilities. Then, the joint latent and image posterior approximation (VBLE-xz) improves the approximated posterior by taking into account the reconstruction error of the VAE. At last, we choose to use compressive VAEs, as their light structure enables to keep a scalable restoration process, while regularizing efficiently the inverse problem with their hyperprior. Experiment results on natural images and satellite images demonstrate that VBLE and VBLE-xz reach similar performance than state-of-the-art plug-and-play methods, while being able to quantify uncertainties faster than other existing posterior sampling techniques.

- 15h - **Hubert Leterme** (Ensicaen, CEA Paris-Saclay)

A plug-and-play approach with conformal predictions for weak lensing mass mapping

In this talk, I will present a plug-and-play (PnP) approach for estimating the distribution of dark matter from noisy shear measurements, in the context of weak gravitational lensing. The method aims to provide accurate estimates efficiently while eliminating the need to train a deep learning model for each observed region of the sky. Instead, the approach requires training a model just once on simulated convergence maps corrupted with a Gaussian white noise. Additionally, we propose to apply a distribution-free uncertainty quantification (UQ) method, namely, conformalized quantile regression (CQR), to this mass mapping framework. Using a calibration set also derived from simulations, CQR provides coverage guarantees independent of any specific prior data distribution. We benchmark our results against CQR applied to existing mass mapping approaches such as Kaiser-Squires, Wiener, MCALens, and DeepMass. Our results reveal that, while the miscoverage rate remains constant across methods, the choice of such method significantly impacts the size of the error bars.

- 15h20 - COFFEE BREAK

- 15h50 - **Andrés Almansa** (CNRS, Université de Paris)

Posterior sampling in imaging with learnt priors: from Langevin to diffusion models

In this talk we explore some recent techniques to perform posterior sampling for ill-posed inverse problems in imaging when the likelihood is known explicitly, and the prior is only known implicitly via a denoising neural network that has been pretrained on a large collection of images. We show how to extend the Unadjusted Langevin Algorithm (ULA) to this particular setting leading to Plug & Play ULA. We explore the convergence properties of PnP-ULA, the crucial role of the stepsize and its relationship with the smoothness of the prior and the likelihood. In order to relax stringent constraints on the stepsize, annealed Langevin algorithms have been proposed, which are tightly related to generative denoising diffusion probabilistic models (DDPM). The image prior that is implicit in these generative models can be adapted to perform posterior sampling, by a clever use of Gaussian approximations, with varying degrees of accuracy, like in Diffusion Posterior Sampling (DPS) and Pseudo-Inverse Guided Diffusion Models (PiGDM). We conclude with an application to blind deblurring, where DPS and PiGDM are used in combination with an Expectation Maximization algorithm to jointly estimate the unknown blur kernel, and sample sharp images from the posterior.

Collaborators (in alphabetical order) Guillermo Carbajal, Eva Coupeté, Valentin de Bortoli, Julie Delon, Alain Durmus, Ulugbek Kamilov, Charles Laroche, Rémy Laumont, Jiaming Liu, Pablo Musé, Marcelo Pereyra, Marien Renaud, Matias Tassano,

- 16h40 - **Corentin Vazia** (Université de Bretagne Sud)

Guidance of a diffusion model for material decomposition in photon-counting computed tomography

Computed tomography image reconstruction is a well-known inverse problem in medical imaging that aims to retrieve the linear attenuation map of the scanned object or patient. In the context of spectral computed tomography, and more specifically with photon-counting detectors, we can now obtain energy-dependent linear attenuation coefficients (LAC) from transmission measurements at different energy levels. This (already ill-posed) inverse problem can be followed by material decomposition where we decompose the LAC images into a material attenuation basis. Recent advances in diffusion models can be used as priors and for regularization of inverse problems by guiding the sampling process (e.g. with diffusion posterior sampling) to a solution that fits with the measurements. In this talk, we present and compare multiple methods using diffusion posterior sampling for material decomposition.

- 17h - **Liam Moroy** (ONERA)

Evaluating the Posterior Sampling Ability of Plug-and-Play Diffusion Methods in Sparse-View CT

Plug-and-Play (PnP) diffusion models are state-of-the-art methods in computed tomography (CT) reconstruction. Such methods usually consider applications where the sinogram contains a sufficient amount of information for the posterior distribution to be peaked, and consequently are evaluated using image-to-image metrics such as PSNR/SSIM. Instead, we are interested in reconstructing compressible flow images from sinograms having a small number of projections, which results in a posterior distribution no longer peaked or even multimodal. Thus, in this paper, we aim at evaluating the approximate posterior of PnP diffusion models and introduce two posterior evaluation criteria. We quantitatively evaluate three PnP diffusion methods on three different datasets for several numbers of projections. We surprisingly find that, for each method,

the approximate posterior deviates from the true posterior when the number of projections decreases.

- 17h20 - **Marcelo Pereyra** (Heriot-Watt University)

Tackling fundamental challenges in hypothesis testing in imaging inverse problems

Despite decades of sustained progress in image estimation accuracy, most imaging methods cannot reliably support hypothesis testing and statistical significance arguments, which are essential for the rigorous interpretation of experiments and robust interfacing of imaging pipelines with decision-making processes. This critical limitation hinders the value of images as evidence for decision making and science. In this regard, imaging sciences are far behind other data disciplines that routinely report uncertainty and significance information and strive to surpass p-values, correlation, and significance as declarations of truth. The reasons for this are manifold: 1) it is extremely difficult to use reliably the observed measurement data to simultaneously reconstruct an image, formulate a hypothesis from this image, and evaluate the significance of the hypothesis from the image; 2) hypotheses in imaging problems are often semantic in nature and more conveniently expressed through language, as written propositions, rather than as quantitative statements about pixel values - as a result it is hard to identify the set of images associated with the null and alternative hypotheses; and 3) hypothesis testing requires establishing a null distribution, which is challenging without resorting to oversimplifying assumptions. This talk presents a novel approach to hypothesis testing for imaging that seeks to address these fundamental difficulties by leveraging ideas from unsupervised machine learning, vision-language models, and non-parametric statistics.