Image Reconstruction by Convolution with Symmetrical Piecewise nth-Order Polynomial Kernels

Erik H. W. Meijering, Karel J. Zuiderveld, and Max A. Viergever

***Abstract*— The reconstruction of images is an important operation in many applications. From sampling theory, it is well known that the sinc-function is the ideal interpolation kernel which, however, cannot be used in practice. In order to be able to obtain an acceptable reconstruction, both in terms of computational speed and mathematical precision, it is required to design a kernel that is of finite extent and resembles the sinc-function as much as possible. In this paper, the applicability of the sinc-approximating symmetrical piecewise nth-order polynomial kernels is investigated in satisfying these requirements. After the presentation of the general concept, kernels of first, third, fifth and seventh order are derived. An objective, quantitative evaluation of the reconstruction capabilities of these kernels is obtained by analyzing the spatial and spectral behavior using different measures, and by using them to translate, rotate, and magnify a number of real-life test images. From the experiments, it is concluded that while the improvement of cubic convolution over linear interpolation is significant, the use of higher order polynomials only yields marginal improvement.**

***Index Terms*—Cubic convolution, image reconstruction, image resampling, interpolation, piecewise polynomial kernels, quintic convolution, septic convolution.**

I. Introduction

THE reconstruction[1] of images, which are in general N-dimensional signals, is an important operation in many applications. Operations such as magnification, subpixel translation, rotation, deformation, or warping of images cannot be carried out without reconstructing the image under consideration. Many interpolation schemes have been devised for that purpose. These include the very simple nearest neighbor and linear interpolators as well as the more computational expensive cubic convolution [1] and windowed sinc [2], [3] interpolation schemes.

In a particular application, the accuracy with which the original signal needs to be reconstructed from the samples determines which alternative to use. Linear interpolation is computationally very cheap and is satisfactory in many situations. Nowadays, linear interpolation is available as a standard operation in special graphics hardware, making it an attractive option. It has to be expected that as the hardware becomes faster, higher order interpolation schemes will become feasible as well. It has been demonstrated by several authors [2], [4], [5] that when further mathematical processing of the data is required, cubic convolution, although more computationally demanding, is a relatively very accurate scheme. The improvement with respect to linear interpolation stimulated us to investigate interpolation kernels consisting of higher order polynomials than just first- or third-order.

The purpose of this paper is twofold. First, we will prove that the concept of the two most popular interpolation schemes, linear interpolation and cubic convolution, can be generalized to a class of what we call *symmetrical piecewise nth-order polynomial interpolation kernels*. By utilizing this concept, higher order interpolation kernels will be derived. Second, it will be shown that while the improvement of cubic convolution over linear interpolation is substantial, the improvements of higher order schemes with respect to cubic convolution are only marginal.

The paper is organized as follows. First the concept of N-dimensional image reconstruction is presented (Section II). It will be shown that the interpolation can always be carried out by N successive *one*-dimensional (1-D) interpolations (cascaded convolution). Next, the concept of 1-D signal reconstruction by convolution with symmetrical piecewise nth-order polynomial kernels is presented (Section III). In Section IV, examples of interpolating kernels consisting of polynomials up to seventh order are given. A quantitative comparison between the reconstruction capabilities of these kernels is presented in Section V. Finally, concluding remarks are made in Section VI.

II. Reconstruction of N-Dimensional Images

A real-valued continuous image $\hat{I}(\boldsymbol{x})$ can be constructed from a real-valued N-dimensional discrete image $I_s(\boldsymbol{p})$ by

Manuscript received April 17, 1997; revised March 19, 1998. This work was supported in part by The Netherlands Ministry of Economic Affairs under IOP-Project IBV96004. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Ping Wah Wong.

E. H. W. Meijering and M. A. Viergever are with the Image Sciences Institute, Utrecht University, 3584 CX Utrecht, The Netherlands (e-mail: erik@cv.ruu.nl).

K. J. Zuiderveld is with Vital Images, Inc., Minneapolis, MN 55416-4510 USA.

Publisher Item Identifier S 1057-7149(99)00925-2.

[1] In this paper, the term "reconstruction" is used to indicate the process of retrieving the original continuous image $I(\boldsymbol{x})$ from its samples $I_s(\boldsymbol{p})$. The term "interpolation" is used to indicate the more general process of "filling in" between the samples. That is, any type of interpolation can be used to calculate values at intermediate points. However, in order to *reconstruct* the values at the intermediate points, specific constraints need to be imposed on the interpolation scheme.

means of interpolation according to

$$\hat{I}(\boldsymbol{x}) = I_s(\boldsymbol{p}) * h(\boldsymbol{x}) \tag{1}$$

where $*$ denotes (discrete) convolution, $\hat{I}$: $\mathbb{R}^N \rightarrow \mathbb{R}$ is the interpolant, $\boldsymbol{x} = [x_1, \cdots, x_N]^T \in \mathbb{R}^N$ are points in N-dimensional continuous space, $\boldsymbol{p} = [p_1, \ldots, p_N]^T \in \mathbb{Z}^N$ are points in N-dimensional discrete space (i.e., they are assumed to lie on a regular grid with unit distance between grid-points), $I_s(\boldsymbol{p})$ are the image values at those points (the samples) and h: $\mathbb{R}^N \rightarrow \mathbb{R}$ is the convolution kernel, interpolating the samples, i.e.,

$$\hat{I}(x_1, \cdots, x_N) = \sum_{p_N} \cdots \sum_{p_1} I_s(p_1, \cdots, p_N) \cdot h(x_1 - p_1, \cdots, x_N - p_N) \tag{2}$$

In the Fourier domain this becomes:

$$\tilde{I}(\boldsymbol{f}) = \tilde{I}_s(\boldsymbol{f})\tilde{H}(\boldsymbol{f}) \tag{3}$$

where $\tilde{I}$: $\mathbb{R}^N \rightarrow \mathbb{C}$ is the Fourier transform of $\hat{I}(\boldsymbol{x}), \boldsymbol{f} = [f_1, \cdots, f_N]^T \in \mathbb{R}^N$ denotes N-dimensional frequency, $\tilde{I}_s$: $\mathbb{R}^N \rightarrow \mathbb{C}$ the Fourier transform of the discrete image $I_s(\boldsymbol{x})$ (the sample data) and $\tilde{H}$: $\mathbb{R}^N \rightarrow \mathbb{C}$ the Fourier transform of the convolution kernel.

The question arises whether it is possible to let the interpolated image $\hat{I}(\boldsymbol{x})$ become exactly equal to the original continuous version $I(\boldsymbol{x})$ of the sampled image $I_s(\boldsymbol{p})$, prior to sampling. In other words: can $I(\boldsymbol{x})$ be *reconstructed* from the samples $I_s(\boldsymbol{p})$? For 1-D signals, the answer is provided by the following well-known theorem:

Theorem 1: Let I: $\mathbb{R} \rightarrow \mathbb{R}$ be a real-valued function that does not contain any frequencies higher than $F_m > 0$, i.e., the Fourier spectrum $\tilde{I}$: $\mathbb{R} \rightarrow \mathbb{C}$ of that function satisfies $\tilde{I}(f) = 0, |f| > F_m$. Then $I(x), x \in \mathbb{R}$, is completely determined by the samples $I_s(p), p \in \mathbb{Z}$ and $I_s(p) = I(p)$, if and only if the sampling frequency F_s satisfies $F_s > 2F_m$.

The essence of this theorem has been known by mathematicians in the field of interpolatory function theory since at least 1915 [6] and has later been applied to the fields of telecommunication and information theory by Nyquist [7] and Shannon [8]. An excellent tutorial review on sampling theory has been presented by Jerri [9].

The theorem can easily be extended to N-dimensions, in which case $I(\boldsymbol{x}), \boldsymbol{x} \in \mathbb{R}^N$ can be completely reconstructed from the samples $I_s(\boldsymbol{p}), \boldsymbol{p} \in \mathbb{Z}^N$ and $I_s(\boldsymbol{p}) = I(\boldsymbol{p})$, if and only if the sampling frequencies F_{s_i} satisfy $F_{s_i} > 2F_{m_i}, \forall i = 1, 2, \cdots, N$, where F_{m_i} is the highest frequency in the ith dimension, i.e., $\tilde{I}(\boldsymbol{f}) = 0, |f_i| > F_{m_i}, \forall i = 1, 2, \cdots, N$. To this end, the filter $\tilde{H}(\boldsymbol{f})$ in (3) must be an N-dimensional box-filter $\tilde{H}_B(\boldsymbol{f})$, defined by:

$$\tilde{H}_B(\boldsymbol{f}) = \begin{cases} 1, & |f_i| \leq \frac{1}{2} F_{s_i}, \forall i = 1, 2, \cdots, N \\ 0, & \text{otherwise.} \end{cases} \tag{4}$$

For a regular grid with unit distance between grid-points, the sampling frequencies F_{s_i} are equal to 1. Equation (4) can also be written as

$$\tilde{H}_B(\boldsymbol{f}) = \prod_{i=1}^{N} \tilde{H}_B(f_i) \tag{5}$$

where $\tilde{H}_B$: $\mathbb{R} \rightarrow \mathbb{R}$ is now a *one*-dimensional box-filter. Thus, $\tilde{I}(\boldsymbol{f})$ becomes:

$$\tilde{I}(\boldsymbol{f}) = \tilde{I}_s(\boldsymbol{f})\tilde{H}_B(f_1)\tilde{H}_B(f_2)\cdots\tilde{H}_B(f_N). \tag{6}$$

The inverse Fourier transform yields

$$I(\boldsymbol{x}) = [\cdots[[I_s(\boldsymbol{p}) * h_B(x_1)] * h_B(x_2)] * \cdots] * h_B(x_N) \tag{7}$$

from which it is concluded that N-dimensional reconstruction can always be carried out by N successive 1-D interpolations. Therefore, in the sequel only the merits of 1-D kernels $h(x)$ are investigated.

III. Symmetrical Piecewise nth-Order Polynomial Kernels

As shown in the previous section, exact reconstruction of an N-dimensional image is accomplished by multiplying the Fourier spectrum $\tilde{I}_s(\boldsymbol{f})$ of the sampled signal $I_s(\boldsymbol{p})$ with an N-dimensional box-filter or, equivalently, with N successive 1-D box-filters. By defining the inverse Fourier transform of a spectrum $\tilde{H}(f)$ as $\int_{-\infty}^{\infty} \tilde{H}(f)e^{\mathrm{i}2\pi fx}\, df$, it can easily be shown that the kernel of the box-filter is the sinc-function defined by

$$h_B(x) = \mathrm{sinc}(x) \triangleq \frac{\sin(\pi x)}{\pi x}. \tag{8}$$

This kernel has infinite extent and cannot be used in practice, thereby ruling out the possibility of exact reconstruction. In order to be able to obtain an acceptable reconstruction (in terms of both computational cost and mathematical precision), it is required to design a kernel of finite extent that resembles the sinc-function as much as possible. The purpose of this paper is to investigate to what extent piecewise nth-order polynomial kernels are capable of satisfying these requirements. First, the precise definition of these kernels is given below.

Definition 1: Given a regular sampling grid with unit distance between the sample points, the symmetrical piecewise nth-order polynomial interpolation kernel $h(x)$ is defined as

$$h(x) = \begin{cases} a_{ni}|x|^n + \cdots + a_{1i}|x| + a_{0i}, & i \leq |x| < i+1 \\ 0, & m \leq |x| \end{cases} \tag{9}$$

where $i = 0, 1, \cdots, m-1$, the parameter $m \in \mathbb{N}_{\setminus\{0\}}$ is the extent of the kernel, and n and m are related by[2]: $n = 2m-1$.

The $(n+1)m$ coefficients a_{ji} are to be determined by imposing constraints on $h(x)$. Since the samples $I_s(\boldsymbol{p})$ are the exact values of the original image $I(\boldsymbol{x})$ at the positions $\boldsymbol{p}$ on the sampling grid, the value of the interpolant at those positions must also be equal to the sample values $I_s(\boldsymbol{p})$. Additional constraints are derived by requiring $h(x)$, which consists of *piecewise* polynomials, to be continuous and, if possible, have continuous derivatives at the transition points $\boldsymbol{p}$. These requirements can be translated into the following constraints.

1) $h(0) = 1$ and $h(x) = 0$ for $|x| = 1, \cdots, m-1$.
2) $h^{(l)}(x)$ must be continuous at $|x| = 0, 1, \cdots, m$.

[2]This choice for n will be motivated at the end of this section.

where the superscript (l) denotes the lth derivative. The second constraint holds for $l = 0, 1, \cdots, k$, where k must be sufficiently large so as to yield a sufficient number of equations in order to be able to solve for the unknown coefficients a_{ji}.

For the determination of an upper limit for k, the following lemma is necessary.

Lemma 1: In generating a set of equations according to the aforementioned constraints, in order to solve for the unknown coefficients of the polynomials constituting an interpolation kernel $h(x)$ as given in Definition 1, the requirement that $h^{(l)}(x)$ must be continuous at $x = 0$ does not yield a nontrivial equation in the case l is even.

Proof: From the general description of symmetrical nth order polynomial kernels, given in Definition 1, it can easily be derived that $h^{(l)}(0) = (l!)a_{l0}$. The continuity constraint of $h^{(l)}(x)$ at $x = 0$, which is expressed as the requirement $\lim_{x\uparrow 0} h^{(l)}(x) = \lim_{x\downarrow 0}\ h^{(l)}(x)$, leads to $(l!)a_{l0} = (l!)a_{l0}$, which does not contain any information. Note that in the case l is odd, the constraint leads to $(l!)a_{l0} = -(l!)a_{l0}$, which results in the requirement $a_{l0} = 0$. ∎

By taking into account the result of this lemma, the maximum value for k is given by the following theorem.

Theorem 2: In generating a set of equations according to the aforementioned constraints, in order to solve for the unknown coefficients a_{ji} of the polynomials constituting an interpolation kernel $h(x)$ as given in Definition 1, the maximum allowable value of k is zero for $n = 1$ and $n - 2$ for $n > 1$.

Proof: According to Definition 1, there are $(n + 1)m$ unknown coefficients a_{ji}, which requires the same amount of independent equations in order to be able to obtain a unique solution. The first constraint will result in m equations. The second constraint will yield $(k + 1)(m + 1)$ equations and according to Lemma 1, $\lceil (k+1)/2 \rceil$ of them will not contain information. Hence, the constraints lead to a total of $m+(k+1)(m+1)-\lceil (k+1)/2 \rceil$ equations. Substituting $k = n-1$ and using the fact that, because n is odd, $n - \lceil n/2 \rceil = \lceil n/2 \rceil - 1$, we have $n(m + 1) = (n + 1)m + (m - 1)$ equations. For $n = 1, m$ and n are both equal to one and therefore the number of equations equals the number of unknowns. For $n > 1$ we have $n > m$, which results in an over-constrained problem that cannot be solved uniquely, but only, e.g., in a least-squares sense. The next smaller integer value is $k = n - 2$ which can be derived to yield $(m+1)n-m = (n+1)m-1$ equations. ∎

Corollary 1: If, in generating a set of equations according to the aforementioned constraints in order to solve for the unknown coefficients a_{ji} of the polynomials constituting an interpolation kernel $h(x)$ as given in Definition 1, the maximum value for k (according to Theorem 2) is applied, the coefficients a_{ji} can be solved uniquely for $n = 1$ and are a function of exactly *one* free parameter, say α, for $n > 1$.

Corollary 2: If, in generating a set of equations according to the aforementioned constraints in order to solve for the unknown coefficients a_{ji} of the polynomials constituting an interpolation kernel $h(x)$ as given in Definition 1, the maximum value of k (according to Theorem 2) is applied, the resulting interpolant will be an element of C^0 for $n = 1$ and of C^{n-2} for $n > 1$.

Now that it has been proven that for $n > 1$ the coefficients of the polynomials are always a function of a free parameter α, i.e. there is a *family* of possible kernels for every value of $n > 1$, it remains to derive a reasonable value for, or at least bounds on, this parameter. An obvious choice would be to require $h^{(k+1)}(x)$ to be continuous at exactly *one* of the transition points $\boldsymbol{p}$, which results in exactly one additional equation, thereby allowing the system of equations to be solved uniquely. However, this is not guaranteed to be the best choice when it comes to accurate (mathematical precise) reconstruction. It will be shown that there always exists exactly *one* optimal value for the parameter α. In order to prove this, the following lemma is necessary.

Lemma 2: The Taylor series expansion of the Fourier transform $\tilde{H}(f)$ of a symmetrical piecewise nth-order polynomial interpolation kernel $h(x)$ as given in Definition 1, in which the coefficients a_{ji} are functions of exactly one free parameter α, has the form

$$\tilde{H}(f) = \beta_0(\alpha) + \beta_2(\alpha)f^2 + \beta_4(\alpha)f^4 + \beta_6(\alpha)f^6 + \cdots, \quad (10)$$

i.e., the series only consists of even terms, in which the factors $\beta_c(\alpha)$ are linear functions of the parameter α, i.e., they have the form

$$\beta_c(\alpha) = \zeta_{c_1}\alpha + \zeta_{c_2}. \quad (11)$$

Proof: Because $h(x)$ is a real-valued function, the Fourier transform $\tilde{H}(f)$ is conjugate symmetric. Furthermore, because $h(x)$ is an even function, the Fourier transform is also even [10]. These two properties lead to the conclusion that $\tilde{H}(f)$ is a real-valued symmetrical function, from which it is known that the Taylor series only contains even terms

$$\tilde{H}(f) = \tilde{H}(0) + \tfrac{1}{2}\,\tilde{H}^{(2)}(0)f^2 + \tfrac{1}{24}\,\tilde{H}^{(4)}(0)f^4 + \cdots. \quad (12)$$

By defining the Fourier transform of $h(x)$ as $\int_{-\infty}^{\infty} h(x) \cdot e^{-\mathrm{i}2\pi fx}\,dx$, it can easily be derived that $\tilde{H}(f)$ can be written as the sum

$$\tilde{H}(f) = \sum_{j=0}^{n} \sum_{i=0}^{m-1} a_{ji}\tilde{H}_{ji}(f) \quad (13)$$

with

$$\tilde{H}_{ji}(f) = 2\int_{i}^{i+1} x^j \cos(2\pi fx)\,dx \quad (14)$$

where use is made of the symmetry property of $h(x)$. According to (13), the derivatives in (12) are calculated as

$$\tilde{H}^{(l)}(0) = \sum_{j=0}^{n} \sum_{i=0}^{m-1} a_{ji}\tilde{H}_{ji}^{(l)}(0) \quad (15)$$

from which it is concluded that the factors $\beta_c(\alpha)$ in the series are linear combinations of the coefficients a_{ji}, which are in turn linear combinations of the parameter α. ∎

Theorem 3: The value of the free parameter α that yields the most mathematical precise interpolant $h(x)$ can be determined by stating

$$\beta_2(\alpha) = 0. \quad (16)$$

Proof: In order to prevent high frequency emphasis in the Fourier spectrum $\tilde{H}(f)$ of the interpolation kernel $h(x)$, the value of the parameter α should be chosen so as to force $\tilde{H}(f)$ to be concave downward at $f = 0$. This is expressed in the requirement

$$\left.\frac{d^2\tilde{H}(f)}{df^2}\right|_{f=0} \leq 0. \tag{17}$$

In order to prevent low frequency suppression in the Fourier spectrum $\tilde{H}(f)$ of the interpolation kernel $h(x)$, the value of the parameter α should be chosen so as to force $\tilde{H}(f)$ to be concave upward at $f = 0$. This is expressed in the requirement

$$\left.\frac{d^2\tilde{H}(f)}{df^2}\right|_{f=0} \geq 0. \tag{18}$$

This leads to the conclusion that the optimal value for the free parameter α is the one that satisfies both requirements (17) and (18) and since

$$\left.\frac{d^2\tilde{H}(f)}{df^2}\right|_{f=0} = 2\beta_2(\alpha) \tag{19}$$

this implies that the optimal value can be found by stating $\beta_2(\alpha) = 0$. ∎

In Definition 1, the order n of the polynomials was related to the extent m by $n = 2m - 1$. Since $m \in \mathbb{N}_{\setminus\{0\}}$, this implies that n will always be odd. It has to be pointed out that $n = 2m-1$ is not the only possible value for n. In fact, it can be shown that $n = 2m$ or $n = 2m+1$ will also lead to a unique kernel for all $m \in \mathbb{N}_{\setminus\{0\}}$. In many cases, even higher values for n are possible. However, given the extent m, the value $n = 2m - 1$ is the *lowest possible* order of the polynomials for which a unique solution exists. In order to solve for the unknown coefficients in the case that $n > 2m - 1$, higher order derivatives are required to be continuous (i.e., k must be larger). However, a larger value for k implies that the resulting kernel will be smoother than its lower order version on the same extent, causing both its spatial and spectral behavior to be worse.

IV. Kernel Examples

The theory developed in the previous section will now be used to derive the four lowest order piecewise polynomial interpolation kernels, which provide for linear, cubic, quintic and septic interpolation.

A. Linear Interpolation

The kernel of the linear interpolator is built up of linear polynomials and approximates the ideal sinc-function in the interval $[-1, 1]$. The kernel is given by

$$h_L(x) = \begin{cases} a_{10}|x| + a_{00}, & 0 \leq |x| < 1 \\ 0, & 1 \leq |x|. \end{cases} \tag{20}$$

The two coefficients a_{ji} can be derived by imposing the following constraints:

1) $h_L(0) = 1$;
2) $h_L(x)$ must be continuous at $|x| = 0, 1$.

TABLE I
The Eight Coefficients of the Polynomials of the Cubic Convolution Kernel as a Function of the Free Parameter α

i	a_{3i}	a_{2i}	a_{1i}	a_{0i}
0	$\alpha + 2$	$-(\alpha + 3)$	0	1
1	α	-5α	8α	-4α

These constraints yield two equations in the two unknown coefficients a_{ji}. This system can be solved uniquely, yielding $a_{10} = -1$ and $a_{00} = 1$. The resulting interpolant will be an element of C^0 (i.e. is continuous).

B. Cubic Convolution

The cubic convolution kernel is built up of third-order polynomials and approximates the ideal sinc-function in the interval $[-2, 2]$. The kernel is given by

$$h_C(x) = \begin{cases} a_{30}|x|^3 + a_{20}|x|^2 + a_{10}|x| + a_{00}, & 0 \leq |x| < 1 \\ a_{31}|x|^3 + a_{21}|x|^2 + a_{11}|x| + a_{01}, & 1 \leq |x| < 2 \\ 0, & 2 \leq |x|. \end{cases} \tag{21}$$

The eight coefficients a_{ji} can be derived by imposing the following constraints:

1) $h_C(0) = 1$ and $h_C(x) = 0$ for $|x| = 1$;
2) $h_C^{(l)}(x)$ must be continuous at $|x| = 0, 1, 2$ for $l = 0, 1$.

These constraints yield seven equations in eight unknown coefficients a_{ji}. By allowing $a_{31} = \alpha$ to be a tunable parameter, the system can be solved, yielding the values for the eight coefficients as shown in Table I.

The Taylor series expansion of the Fourier spectrum of this kernel around $f = 0$ is given by

$$\begin{aligned}\tilde{H}_C(f) = 1 &- \tfrac{4}{15}(2\alpha + 1)(\pi f)^2 \\ &+ \tfrac{1}{35}(16\alpha + 1)(\pi f)^4 + \mathcal{O}(f^6).\end{aligned} \tag{22}$$

By imposing the constraint that the lower frequencies are not allowed to be suppressed and also that the higher frequencies are not allowed to be amplified, we find that the best cubic convolution kernel is the one with the value for α that forces $\tilde{H}_C(f)$ to be neither convex nor concave at $f = 0$, in other words, that causes the second-order term in the Taylor series to vanish:

$$\left.\frac{d^2\tilde{H}_C(f)}{df^2}\right|_{f=0} = -\frac{8}{15}(2\alpha + 1) = 0 \tag{23}$$

which yields $\alpha = -\frac{1}{2}$ [11], [12]. The resulting interpolant will be an element of C^1 (i.e., is continuous and has a continuous first derivative).

C. Quintic Convolution

The quintic convolution kernel is built up of fifth-order polynomials and approximates the ideal sinc-function in the

TABLE II
THE 18 COEFFICIENTS OF THE POLYNOMIALS OF THE QUINTIC CONVOLUTION KERNEL AS A FUNCTION OF THE FREE PARAMETER α

i	a_{5i}	a_{4i}	a_{3i}	a_{2i}	a_{1i}	a_{0i}
0	$10\alpha-\frac{21}{16}$	$-18\alpha+\frac{45}{16}$	0	$8\alpha-\frac{5}{2}$	0	1
1	$11\alpha-\frac{5}{16}$	$-88\alpha+\frac{45}{16}$	$270\alpha-10$	$-392\alpha+\frac{35}{2}$	$265\alpha-15$	$-66\alpha+5$
2	α	-14α	78α	-216α	297α	-162α

interval $[-3,3]$. The kernel is given by

$$h_Q(x)=\begin{cases}a_{50}|x|^5+\cdots+a_{10}|x|+a_{00}, & 0\le|x|<1\\ a_{51}|x|^5+\cdots+a_{11}|x|+a_{01}, & 1\le|x|<2\\ a_{52}|x|^5+\cdots+a_{12}|x|+a_{02}, & 2\le|x|<3\\ 0, & 3\le|x|.\end{cases} \tag{24}$$

The eighteen coefficients a_{ji} can be derived by imposing the following constraints:

1) $h_Q(0)=1$ and $h_Q(x)=0$ for $|x|=1,2$;
2) $h_Q^{(l)}(x)$ must be continuous at $|x|=0,1,2,3$ for $l=0,1,2,3$.

These constraints yield 17 equations in the 18 unknown coefficients a_{ji}. By allowing $a_{52}=\alpha$ to be a tunable parameter, the system can be solved, yielding the values for the 18 coefficients, as shown in Table II.

The Taylor series expansion of the Fourier spectrum of this kernel around $f=0$ is given by

$$\tilde{H}_Q(f)=1+\tfrac{1}{14}(64\alpha-3)(\pi f)^2-\tfrac{1}{105}(352\alpha+1)(\pi f)^4+\mathcal{O}(f^6). \tag{25}$$

By imposing the constraint that the Fourier spectrum $\tilde{H}_Q(f)$ is neither convex nor concave at $f=0$, we find that the best quintic convolution kernel is the one for which the second-order term in the Taylor series vanishes:

$$\left.\frac{d^2\tilde{H}_Q(f)}{df^2}\right|_{f=0}=\frac{2}{14}(64\alpha-3)=0 \tag{26}$$

which yields $\alpha=\frac{3}{64}$. The resulting interpolant will be an element of C^3 (i.e., is continuous and has continuous first-, second-, and third-order derivatives).

D. Septic Convolution

The septic convolution kernel is built up of seventh-order polynomials and approximates the ideal sinc-function in the interval $[-4,4]$. The kernel is given by

$$h_S(x)=\begin{cases}a_{70}|x|^7+\cdots+a_{10}|x|+a_{00}, & 0\le|x|<1\\ a_{71}|x|^7+\cdots+a_{11}|x|+a_{01}, & 1\le|x|<2\\ a_{72}|x|^7+\cdots+a_{12}|x|+a_{02}, & 2\le|x|<3\\ a_{73}|x|^7+\cdots+a_{13}|x|+a_{03}, & 3\le|x|<4\\ 0, & 4\le|x|.\end{cases} \tag{27}$$

The 32 coefficients a_{ji} can be derived by imposing the following constraints:

1) $h_S(0)=1$ and $h_S(x)=0$ for $|x|=1,2,3$;
2) $h_S^{(l)}(x)$ must be continuous at $|x|=0,1,2,3,4$ for $l=0,1,2,3,4,5$.

TABLE III
THE 32 COEFFICIENTS OF THE POLYNOMIALS OF THE SEPTIC CONVOLUTION KERNEL AS A FUNCTION OF THE FREE PARAMETER α

i	a_{7i}	a_{6i}	a_{5i}	a_{4i}
0	$245\alpha+\frac{821}{1734}$	$-621\alpha-\frac{1148}{867}$	0	$760\alpha+\frac{1960}{867}$
1	$301\alpha+\frac{1687}{6936}$	$-3309\alpha-\frac{2492}{867}$	$14952\alpha+\frac{32683}{2312}$	$-35640-\frac{128695}{3468}$
2	$57\alpha+\frac{35}{6936}$	$-1083\alpha-\frac{175}{1734}$	$8736\alpha+\frac{1995}{2312}$	$-38720\alpha-\frac{4725}{1156}$
3	α	-27α	312α	-2000α
i	a_{3i}	a_{2i}	a_{1i}	a_{0i}
0	0	$-384\alpha-\frac{1393}{578}$	0	1
1	$47880\alpha+\frac{127575}{2312}$	$-36000\alpha-\frac{13006}{289}$	$14168\alpha+\frac{120407}{6936}$	$-2352\alpha-\frac{2233}{1156}$
2	$101640\alpha+\frac{1575}{136}$	$-157632\alpha-\frac{5670}{289}$	$133336\alpha+\frac{42525}{2312}$	$-47280\alpha-\frac{8505}{1156}$
3	7680α	-17664α	22528α	-12288α

These constraints yield 31 equations in the 32 unknown coefficients a_{ji}. By allowing $a_{73}=\alpha$ to be a tunable parameter, the system can be solved, yielding the values for the 32 coefficients, as shown in Table III.

The Taylor series expansion of the Fourier spectrum of this kernel around $f=0$ is given by

$$\tilde{H}_S(f)=1-\tfrac{2}{867}(83232\alpha+71)(\pi f)^2+\tfrac{8}{8415}(97920\alpha-53)(\pi f)^4+\mathcal{O}(f^6). \tag{28}$$

Again, by imposing the constraint that the magnitude of the Fourier spectrum $\tilde{H}_S(f)$ is neither convex nor concave at $f=0$, we find that the best septic convolution kernel is the one for which the second-order term in the Taylor series vanishes:

$$\left.\frac{d^2\tilde{H}_S(f)}{df^2}\right|_{f=0}=-\frac{4}{867}(83232\alpha+71)=0 \tag{29}$$

which yields $\alpha=-(71/83232)$. The resulting interpolant will be an element of C^5 (i.e., is continuous and has continuous first-, second-, third-, fourth-, and fifth-order derivatives).

V. EXPERIMENTAL RESULTS

In this section, the reconstruction capabilities of the kernels derived in the previous section will be quantitatively evaluated and compared. First, both the spatial and spectral behavior of the kernels will be analyzed. Second, the kernels will be applied in a number of experiments on real-life images. Finally, the results will be briefly discussed.

A. Spatial and Spectral Analyses

Plots of the four interpolation kernels $h_L(x), h_C(x), h_Q(x)$, and $h_S(x)$, as presented in the previous section, are shown in Fig. 1, together with the plots of their spectra $\tilde{H}_L(f), \tilde{H}_C(f), \tilde{H}_Q(f)$, and $\tilde{H}_S(f)$. From these plots it can be observed that the cubic convolution kernel resembles the ideal box-filter $\tilde{H}_B(f)$ substantially better than the linear interpolation kernel. The spectra of the quintic and septic convolution kernels appear identical to that of the cubic convolution scheme. However, their corresponding log-plots in Fig. 1 reveal that the high frequency suppression capabilities

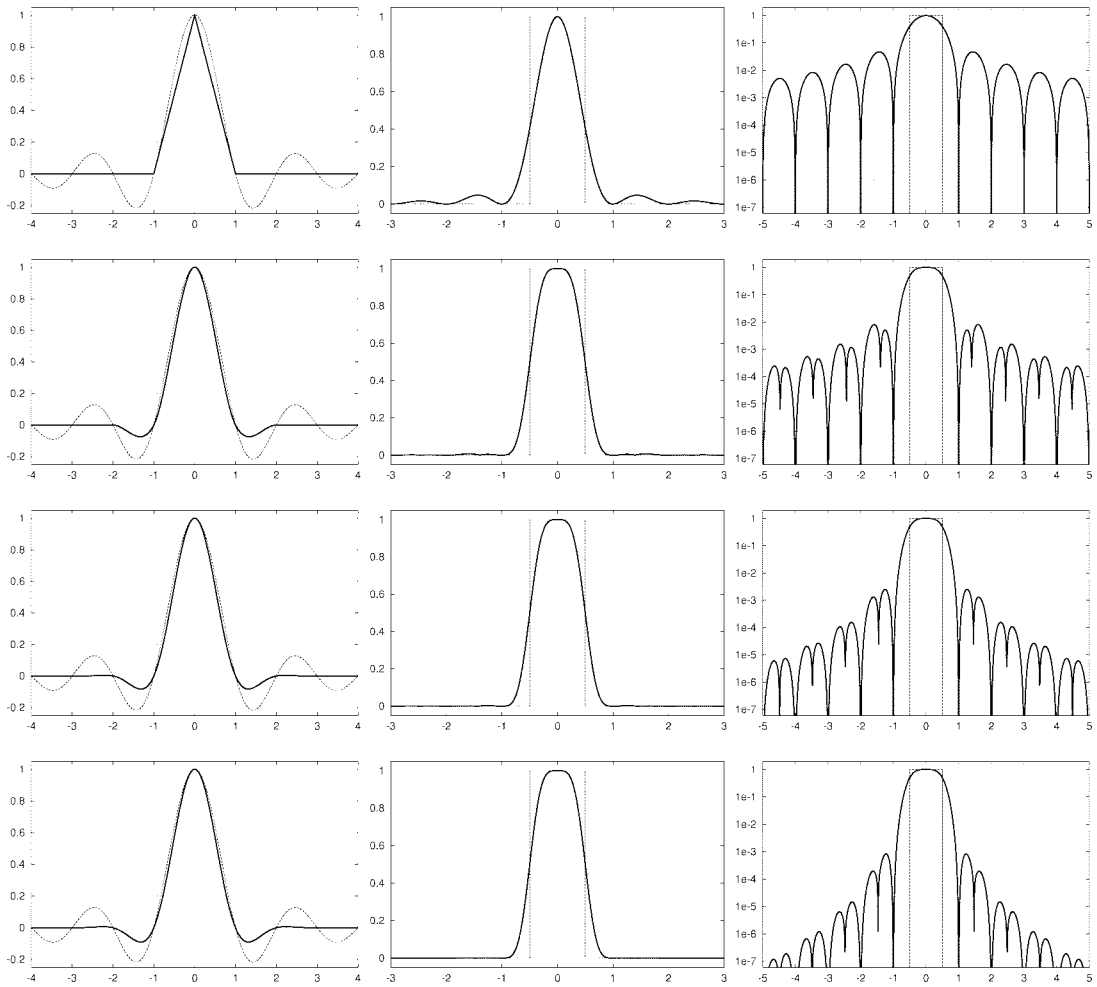

Fig. 1. Kernels of the four lowest order interpolators, compared to the ideal interpolation kernel, the sinc-function (printed as a dashed line). First column: The linear interpolation kernel $h_L(x)$, the cubic convolution kernel $h_C(x)$, the quintic convolution kernel $h_Q(x)$ and the septic convolution kernel $h_S(x)$. Second column: The corresponding Fourier spectra $\tilde{H}_L(f)$, $\tilde{H}_C(f)$, $\tilde{H}_Q(f)$, and $\tilde{H}_S(f)$, compared to the spectrum $\tilde{H}_B(f)$ of the sinc-function (printed as a dashed line). Third column: The corresponding log-plots of the spectra.

of the quintic and septic convolution kernels are, respectively, one and two orders of magnitude better than that of the cubic convolution kernel.

Since the spectra of the kernels are all flat at $f = 0$ (because of the constraint proposed in Theorem 3), they can be compared by evaluating the transition rate from stopband ($|f| > 1/2$) to passband ($-1/2 \leq f \leq 1/2$). This rate corresponds to the first derivative of the spectrum $\tilde{H}(f)$ at $f = -1/2$, and should be as large as possible. The values of this derivative for the presented kernels are shown in Table IV. The ideal kernel, the sinc-function, has an infinite transition rate, causing high frequencies to be completely suppressed and low frequencies to be completely conserved. The improvement of cubic convolution over linear interpolation is 43.2% according to this feature. Quintic and septic convolution are, respectively, 49.2% and 56.6% better than linear interpolation, but only 4.2% and 9.3% better than cubic convolution.

In the literature on sampling and reconstruction, an important error measure to study the spectral behavior of a kernel as a function of frequency is given by

$$E(f) = |1 - \tilde{H}(f)|^2 + \sum_{\eta \in \mathbb{Z} \setminus \{0\}} |\tilde{H}(f - \eta F_s)|^2 \qquad (30)$$

TABLE IV
THE FIRST DERIVATIVE OF THE FOURIER TRANSFORM $\tilde{H}(f)$ OF THE FOUR INTERPOLATION KERNELS $h_L(x), h_C(x), h_Q(x)$, AND $h_S(x)$ PRESENTED IN SECTION IV, AT THE TRANSITION FROM STOPBAND TO PASSBAND ($f = -\frac{1}{2}$), COMPARED TO THAT OF THE SINC-FUNCTION $h_B(x)$

Kernel	$\tilde{H}^{(1)}(-\frac{1}{2})$
$h_L(x)$	1.621
$h_C(x)$	2.321
$h_Q(x)$	2.419
$h_S(x)$	2.538
$h_B(x)$	∞

where $\tilde{H}(f), f \in \mathbb{R}$ denotes the spectrum of the kernel to be analyzed and $F_s \in \mathbb{R}$ is the sampling frequency. This measure was originally proposed by Park and Schowengerdt [13] in order to study the blur caused by sampling and reconstruction (SR-blur) and has since been applied by several others [14], [15] to compare their methods to cubic convolution. Note that for the sinc-function, $E(f) = 0$ for $|f| \leq (1/2)F_s$ and $E(f) = 2$ for $|f| > (1/2)F_s$. The error functions $E(f)$ for the spectra of the four interpolation kernels are shown in Fig. 2, together with that of the sinc-function. From this figure it is clear that the cubic convolution kernel is substantially better

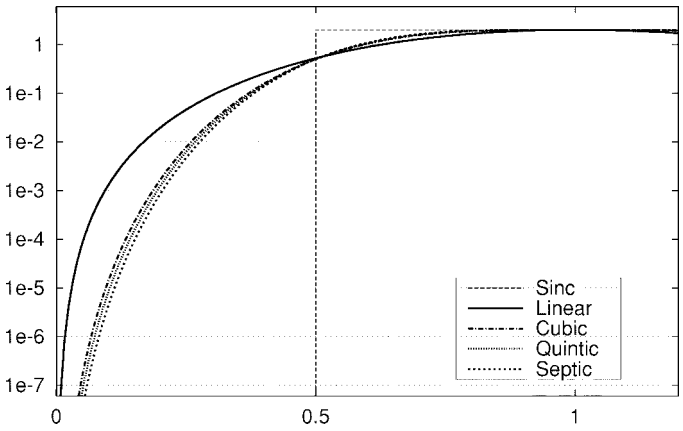

Fig. 2. Error function $E(f)$ [see (29)] for the spectra of the four interpolation kernels described in Section IV, compared to that of the sinc-function.

TABLE V
TOTAL SQUARE ERROR E_T [SEE (31)] OF THE SPECTRA OF THE FOUR INTERPOLATION KERNELS $h_L(x), h_C(x), h_Q(x)$, AND $h_S(x)$ PRESENTED IN SECTION IV, COMPARED TO THAT OF THE SPECTRUM OF THE SINC-FUNCTION $h_B(x)$

Kernel	E_T
$h_L(x)$	0.119277
$h_C(x)$	0.078894
$h_Q(x)$	0.075913
$h_S(x)$	0.072559
$h_B(x)$	0

than the linear interpolation kernel, but the improvements by the higher order kernels are only marginal.

We also computed the total square error (distance) of the spectra $\tilde{H}(f)$ with respect to the spectrum $\tilde{H}_B(f)$ of the sinc-function

$$E_T = \int_{-\infty}^{\infty} |\tilde{H}_B(f) - \tilde{H}(f)|^2 \, df = \int_{-(1/2)}^{1/2} E(f) \, df. \quad (31)$$

The results for the four interpolation kernels are listed in Table V. From these figures, it can be concluded that, according to the E_T error measure, the improvement of cubic convolution over linear interpolation is 33.9%. Quintic and septic convolution are, respectively, 36.4% and 39.2% better than linear interpolation, but only 3.8% and 8.0% better than cubic convolution.

B. Evaluation Using Real-Life Images

In order to obtain a quantitative comparison of the performance of the four interpolation kernels when applied to real-life images, the kernels were used in three different types of operations, each of which requires accurate reconstruction: i) subpixel translation; ii) rotation; iii) magnification. For these experiments, the 16 two-dimensional (2-D) test-images[3] shown in Fig. 3 were used.

[3]In all the experiments described in this paper, the original images were mirrored around the borders in order to reduce border artifacts.

To carry out a subpixel translation, the following operations have to be performed: i) reconstruction of the image; ii) sampling of the reconstructed image at a new, translated grid. The first operation requires interpolation for which, in this experiment, the four schemes presented in Section IV were used.

The translated images should be compared to the corresponding ideally translated one (in which the sinc-function is used as interpolation kernel). Since such an image cannot be obtained, the images were translated back to their initial position (using the same interpolation kernel as for the forward translation) and the mean square error with respect to the original (nontranslated) version was computed. The results of this subpixel translation experiment for the 16 test images and for a displacement vector of $[0.4, 0.7]^T$ are presented in Table VI. According to these figures, the improvement of cubic convolution over linear interpolation is (on average) 65.1%. Quintic and septic convolution are, respectively, 67.6% and 69.9% better than linear interpolation, but only 7.6% and 14.3% better than cubic convolution.

As a second experiment, the test images were rotated. In order to rotate an image, the following operations need to be carried out: i) reconstruction of the image; ii) sampling of the reconstructed image at a new, rotated grid. The first operation requires interpolation, for which the four schemes presented in Section IV were used.

The rotated images should be compared to the corresponding ideally rotated one. Since such an image cannot be obtained, the images were rotated back to their initial orientation (using the same interpolation kernel as for the forward rotation) and the mean square error with respect to the original version was computed. The results of this experiment for the 16 test images and for a rotation angle of 15° are presented in Table VI. According to these figures, the improvement of cubic convolution over linear interpolation is (on average) 68.8%. Quintic and septic convolution are, respectively, 71.2% and 73.1% better than linear interpolation, but only 7.8% and 14.3% better than cubic convolution.

Finally, the test images were magnified. In order to magnify an image, the following operations have to be carried out: i) reconstruction of the image; ii) sampling of the reconstructed image at a new, more dense, grid; iii) scaling of the resampled image. The first operation requires interpolation for which, again, the four schemes presented in Section IV were used.

The magnified images should be compared to the corresponding ideally magnified one. Since such an image cannot be obtained and since the inverse operation (i.e., subsampling) would exactly yield the original image (by definition of the kernels; recall the constraints in Section III), we had to resort to a different evaluation strategy: the magnification operation was applied to subsampled versions of the original images which, in order to reduce the influence of aliasing, were lowpass filtered prior to subsampling. That is, the lowpass filtered versions of the 16 test images (Fig. 3) were taken as the new originals and were then successively subsampled with a factor four and magnified with the same factor. The magnified images (now having the same size as the original images) were compared to the lowpass filtered versions by computing

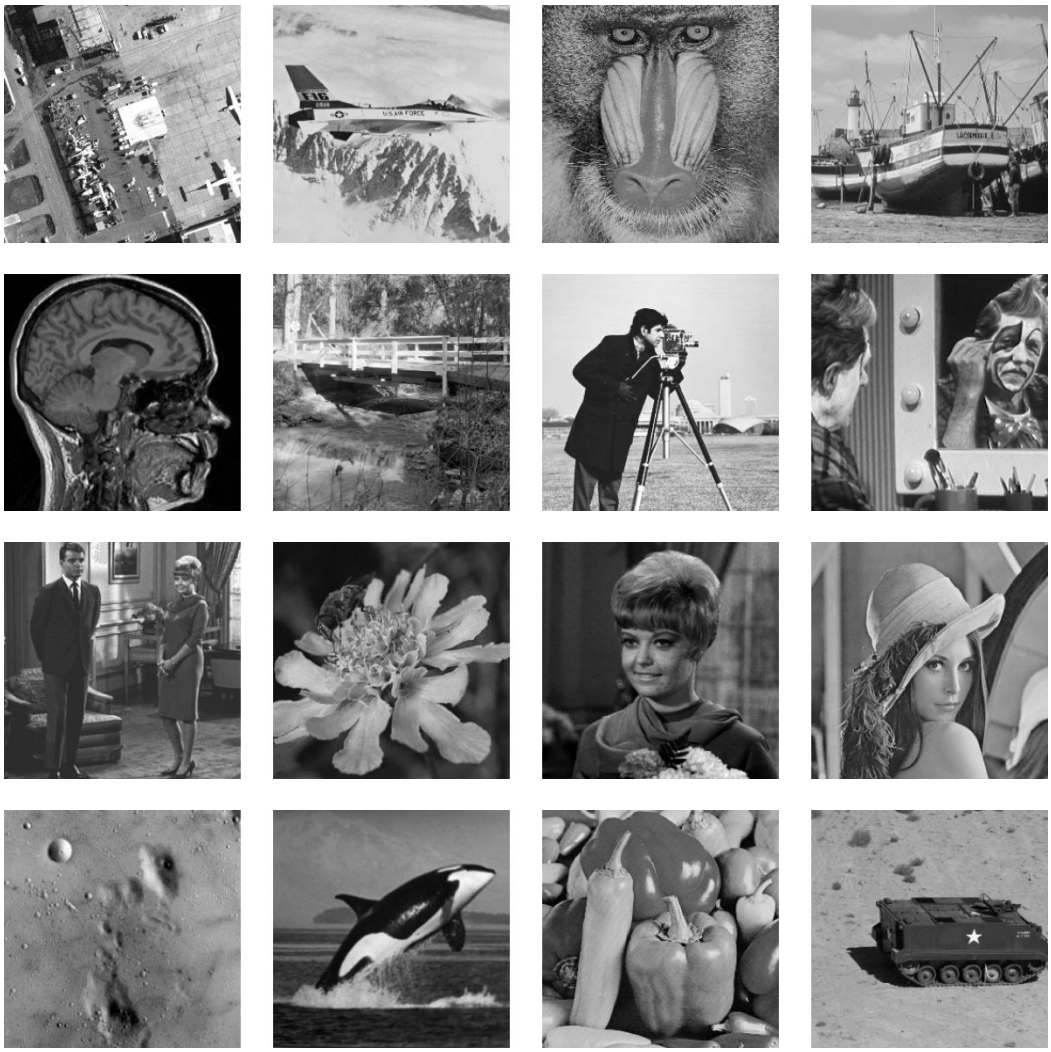

Fig. 3. The 16 test images used in the experiments for comparison of the four interpolation schemes described in Section IV. From top-left to bottom-right: Airfield, airplane, baboon, boat, brain, bridge, camera, clown, couple, flower, girl, Lena, moon, orca, peppers, and tank.

TABLE VI

MEAN SQUARE ERRORS (MSE's) MADE BY THE FOUR INTERPOLATION KERNELS, WHEN USING THEM IN SUBPIXEL TRANSLATION, ROTATION, AND MAGNIFICATION OPERATIONS, APPLIED TO THE 16 REAL-LIFE IMAGES SHOWN IN FIG. 3. THE TRANSLATION WAS CARRIED OUT OVER THE VECTOR $[0.4, 0.7]^T$. THE ROTATION WAS OVER 15° AND THE MAGNIFICATION WAS DONE WITH A FACTOR OF FOUR

Image	Experiment											
	Subpixel Translation				Rotation				Magnification			
	$h_L(x)$	$h_C(x)$	$h_Q(x)$	$h_S(x)$	$h_L(x)$	$h_C(x)$	$h_Q(x)$	$h_S(x)$	$h_L(x)$	$h_C(x)$	$h_Q(x)$	$h_S(x)$
AIRFIELD	78.85	29.88	27.78	25.97	52.21	18.13	16.83	15.79	90.79	54.74	52.74	50.53
AIRPLANE	16.07	4.09	3.73	3.39	10.29	2.34	2.12	1.93	43.09	22.81	21.69	20.45
BABOON	154.12	63.85	58.98	55.29	107.05	42.32	39.42	37.12	46.22	28.69	27.62	26.43
BOAT	28.10	9.95	9.21	8.42	17.03	4.77	4.35	3.97	42.71	25.56	24.56	23.45
BRAIN	32.26	6.20	5.54	4.90	20.09	3.47	3.09	2.75	95.57	55.10	52.84	50.36
BRIDGE	98.23	42.84	39.98	37.69	65.59	25.40	23.67	22.32	73.48	44.61	42.91	41.01
CAMERA	119.99	51.27	47.98	45.01	76.01	27.30	25.33	23.74	87.97	51.36	49.35	47.11
CLOWN	34.22	8.54	7.79	7.03	21.45	4.75	4.31	3.91	91.22	57.82	55.95	53.88
COUPLE	41.67	14.83	13.72	12.74	26.62	8.29	7.63	7.09	56.59	36.51	35.31	33.97
FLOWER	7.10	2.20	2.04	1.91	4.69	1.38	1.29	1.22	25.91	13.89	13.29	12.63
GIRL	18.09	6.33	5.87	5.48	12.14	3.92	3.63	3.41	44.14	31.57	30.89	30.13
LENA	14.98	5.02	4.67	4.36	9.49	2.78	2.57	2.39	31.75	18.78	18.10	17.36
MOON	13.52	5.92	5.53	5.19	8.46	3.06	2.83	2.64	11.14	6.78	6.53	6.27
ORCA	5.83	1.14	1.03	0.93	3.90	0.71	0.64	0.58	29.59	14.84	14.11	13.32
PEPPERS	17.43	8.19	7.74	7.38	12.00	5.25	4.99	4.80	30.23	18.17	17.59	16.94
TANK	20.76	8.81	8.19	7.72	14.33	5.62	5.25	4.96	16.50	10.35	10.02	9.65

the mean square error. The results of this experiment are listed in Table VI. According to these figures, the improvement of cubic convolution over linear interpolation is (on average) 40.1%. Quintic and septic convolution are, respectively, 42.3% and 44.8% better than linear interpolation, but only 3.7% and 7.8% better than cubic convolution.

C. Discussion of the Results

In the subpixel translation, rotation, and magnification experiments, the actual magnitude of the mean square error for a specific image depends on the information contained in that image. If large parts of the image are highly structured, the mean square error will be large, since the largest errors are made at sharp transitions (edges). This is the case in e.g. the baboon and airfield images. In, e.g., the flower and peppers images, there are large more-or-less homogeneous regions and only a relatively small part of these images contains sharp transitions, which entails that the mean square error is relatively small.

We have presented the results for one specific translation vector, rotation angle, and magnification factor. In order to make sure that the effects found were not dependent on the specific choice of these quantities, we have carried out experiments with several other vectors, angles, and factors. The results of these experiments were all very consistent with the results presented here.

As can be gathered from the results of the different experiments, the precise improvement of higher order schemes with respect to lower order ones depends on the choice for the error measure. In general, it can be stated that the specific situation in which an interpolation scheme is to be chosen also determines the error measure that must be selected in order to describe the merits of that scheme for the application. In fact, it is the type of application that determines whether a higher order scheme is to be preferred. The experiments presented in this paper show that in all cases, the improvement of quintic or septic convolution over cubic convolution is far less significant than that of cubic convolution over linear interpolation.

Finally, it should be noted that in a particular situation, the type of operations and the type of images to which they are to be applied may also determine the optimality criterion for deriving the free parameter α. In this paper, we have chosen a criterion which has been shown to yield the most mathematical precise interpolation [11], [12], i.e., the Taylor series expansion of the interpolant will be equal to that of the original signal in as many terms as possible. In a general description of these polynomial kernels, this is the most appropriate choice.

VI. Conclusions

In this paper, the sinc-approximating symmetrical piecewise nth-order polynomial kernels which can be used for the reconstruction of N-dimensional images were presented. After a derivation proving that reconstruction (or interpolation) of N-dimensional signals can be carried out by N successive one-dimensional interpolations, and the presentation of the concept of symmetrical piecewise nth-order polynomial kernels, the linear, cubic, quintic, and septic convolution kernels were derived.

An objective, quantitative comparison of the performance of the four interpolation kernels was obtained by analyzing the spatial and spectral behavior of the kernels according to several measures, as well as by using them in subpixel translation, rotation, and magnification experiments, applied to a number of real-life test images. The results of these experiments show, very consistently, that the errors made by the cubic convolution scheme are substantially smaller than those made by linear interpolation. However, higher order schemes only yield marginal improvement, at an increased computational cost.

References

[1] S. S. Rifman, "Digital rectification of ETRS multispectral imagery," in *Proc. Symp. Significant Results Obtained from the Earth Resources Technology Satellite-1*, 1973, vol. 1, Section B, NASA SP-327, pp. 1131–1142.
[2] G. Wolberg, *Digital Image Warping*. Piscataway, NJ: IEEE Computer Soc. Press, 1990.
[3] A. Antoniou, *Digital Filters: Analysis, Design, and Applications*, 2nd ed. New York: McGraw-Hill, 1993.
[4] J. A. Parker, R. V. Kenyon, and D. E. Troxel, "Comparison of interpolating methods for image resampling," *IEEE Trans. Med. Imag.*, vol. 2, pp. 31–39, 1983.
[5] G. Timmens, "Resampling of multidimensional image data," Master's thesis, Faculty Tech. Math. Informatics, Delft Univ. Technol., The Netherlands, 1991.
[6] E. T. Whittaker, "On the functions which are represented by the expansion of interpolating theory," *Proc. R. Soc. Edinburgh*, vol. 35, pp. 181–194, 1915.
[7] H. Nyquist, "Certain topics in telegraph transmission theory," *Trans. Amer. Inst. Elect. Eng.*, vol. 47, pp. 617–644, 1928.
[8] C. E. Shannon, "Communication in the presence of noise," *Proc. Inst. Radio Eng.*, vol. 37, pp. 10–21, 1949.
[9] A. J. Jerri, "The Shannon sampling theorem—Its various extensions and applications: A tutorial review," *Proc. IEEE*, vol. 65, pp. 1565–1596, 1977.
[10] H. Kwakernaak and R. Sivan, *Modern Signals and Systems*. Englewood Cliffs, NJ: Prentice-Hall, 1991.
[11] R. G. Keys, "Cubic convolution interpolation for digital image processing," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 29, pp. 1153–1160, 1981.
[12] S. K. Park and R. A. Schowengerdt, "Image reconstruction by parametric cubic convolution," *Comput. Vis., Graph. Image Process.*, vol. 23, pp. 258–272, 1983.
[13] ———, "Image sampling, reconstruction, and the effect of sample-scene phasing," *Appl. Opt.*, vol. 21, pp. 3142–3151, 1982.
[14] S. E. Reichenbach and S. K. Park, "Two-parameter cubic convolution for image reconstruction," *Proc. SPIE: Vis. Commun. Image Process.*, vol. 1199, pp. 833–840, 1989.
[15] T. E. Boult and G. Wolberg, "Local image reconstruction and subpixel restoration algorithms," *CVGIP: Graph. Models Image Processing*, vol. 55, pp. 63–77, 1993.

Erik H. W. Meijering was born in Heemskerk, The Netherlands, in 1971. He received the M.Sc. degree (cum laude) from Delft University of Technology, The Netherlands, in 1996. The subject of his graduation project was segmentation of 3-D medical images by means of volume-growing techniques. Since July 1996, he has been a Ph.D. student at the Image Sciences Institute (ISI), Utrecht University, Utrecht, The Netherlands. His research interests include image processing and image analysis.

Karel J. Zuiderveld received the M.Sc. degree electrical engineering from the Technical University Twente, Enschede, The Netherlands, in 1986, and the Ph.D. degree from the Faculty of Medicine, Utrecht University, Utrecht, The Netherlands, in 1995. His dissertation subject was visualization of multimodality medical volume data using object-oriented methods.

From 1986 to 1989, he was a Computer Scientist at the Utrecht University Hospital as a member of the Dutch Picture Archiving and Communication Systems project. In 1989, he became Research Associate at the Image Sciences Institute, Utrecht University Hospital. From 1996 to 1997, he coordinated the development of a software/hardware infrastructure for MR guided vascular interventions at Utrecht University Hospital. Since 1998, he has been a Senior Scientist at Vital Images Inc., Minneapolis, MN.

Dr Zuiderveld is a member of ACM and EuroGraphics. His professional interests include biomedical image processing, 3-D computer graphics, image guided intervention, and software engineering.

Max A. Viergever received the M.Sc. degree in applied mathematics in 1972 and the D.Sc. degree, with a thesis on cochlear mechanics, in 1980, both from Delft University of Technology, The Netherlands.

From 1972 to 1988, he was Assistant/Associate Professor of applied mathematics at Delft University of Technology. Since 1988, he has been Professor and Head of the Department of Medical Imaging at Utrecht University, and as of 1996, Scientific Director of the newly established Image Sciences Institute, Utrecht University. His research interests comprise all aspects of computer vision and medical imaging. He is co-author of over 200 refereed scientific papers on biophysics and medical image processing, and co-author or co-editor of 11 books. He is editor of the book series *Computational Imaging and Vision* (Boston, MA: Kluwer).

Dr. Viergever is a board member of IPMI and IAPR, Associate Editor-in-Chief of IEEE TRANSACTIONS ON MEDICAL IMAGING, editor of the *Journal of Mathematical Imaging and Vision*, and participates on the editorial boards of several journals.