Optimized Quasi-Interpolators for Image Reconstruction

Leonardo Sacht Diego Nehab
IMPA, Instituto Nacional de Matemática Pura e Aplicada

(a) Input image and detail (b) Blu et al. [2001] (O-MOMS) (c) Condat et al. [2005] (d) Our result

Figure 1: *Comparison between state-of-the-art quadratic quasi-interpolators with similar computational cost. The test consists of applying 40 cumulative translations to the input. The net effect brings the image back in alignment with the input, so we can compare them. Our new quadratic is better at preserving detail. SSIM: (a) 1.0, (b) 0.977, (c) 0.987, (d) 0.995. PSNR: (a) ∞, (b) 32.938, (c) 34.149, (d) 36.443.*

Abstract

We obtain new quasi-interpolators for continuous reconstruction of sampled images by minimizing a new objective function that takes into account the approximation error over the full Nyquist interval. To achieve this goal, we optimize with respect to all possible degrees of freedom in the approximation scheme. We consider three study cases offering different trade-offs between quality and computational cost: a linear, a quadratic, and a cubic scheme. Experiments with compounded rotations and translations confirm that our new quasi-interpolators perform better than the state-of-the-art for a similar computational cost.

1 Introduction

The problem of obtaining an estimate for the value of a function at an arbitrary point, when given only a discrete set of sampled values, has a long history in applied mathematics [Meijering 2002]. A variety of operations commonly performed on images, such as rotations, translations, warps, and resolution change, require resampling. Efficient, high-quality reconstruction is therefore of fundamental importance in computer graphics and image processing applications.

In this paper, we leverage recent results from the intersection of image processing and approximation theory to optimize for a new family of reconstruction schemes. Figure 1 shows a typical benchmark used to evaluate reconstruction quality. An input image is repeatedly translated so as to accumulate the errors due to multiple compound reconstruction steps. The figure compares the two best performing quadratic reconstruction schemes with the result of our method. Visual inspection suggests our method is better at preserving high-frequency content, and this is confirmed quantitatively by the perceptual SSIM [Wang et al. 2004] metric as well as the PSNR metric. This success is the result of the greater number of degrees of freedom and the more realistic objective function we use in our optimization framework.

Figure 2 shows the modern approach to sampling and reconstruction [Blu et al. 1999]. The precise definition of each stage in the pipeline is given in section 2. Intuitively, an approximation $\tilde{f}_T$ to f is obtained as follows. In the first two stages, f is subjected to a continuous convolution with an *analysis filter* ψ (a.k.a. *prefilter*), and then sampled with constant sample spacing T. The traditional

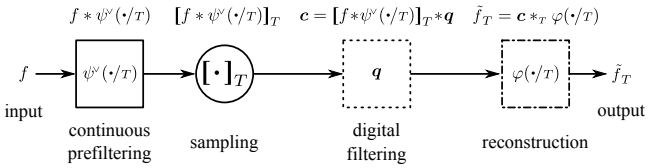

Figure 2: *The sampling pipeline. The input f is first convolved with a scaled prefilter ψ and then sampled on a grid with fixed spacing T. We assume the image is available to us after sampling and we optimize for the quality of the reconstruction $\tilde{f}_T$ using all the degrees of freedom of the digital filter $\boldsymbol{q}$ and of the generator φ.*

role of the prefiltering stage is to eliminate from f frequencies above the Nyquist rate $\frac{0.5}{T}$ so as to avoid aliasing in the sampled sequence. In the applications we discuss in this paper, we assume no knowledge or control over the prefilter ψ. In other words, either there was no prefilter (equivalently, $\psi = \delta$, the Dirac delta) or our goal is to approximate a previously prefiltered signal $f * \psi^\vee$ instead of f itself.

The remaining stages apply a *digital filter* $\boldsymbol{q}$ to the samples (by discrete convolution), and then build $\tilde{f}_T$ by combining shifted copies of a *generating function* φ, each scaled by a filtered sample. The digital filtering stage $\boldsymbol{q}$ is a recent addition to the sampling pipeline [Unser 2000]. It brings several advantages: it increases the range of approximation techniques that can be expressed, and gives more freedom to the design of generators φ with desirable approximation properties. Furthermore, it incurs no significant performance penalty.

The ideal sampling of Shannon [1949] is represented in the sampling pipeline by setting both the prefilter and generating function to the ideal low-pass filter (i.e., $\psi = \varphi = \mathrm{sinc}$, the *sinus cardinalis*), and omitting the digital filtering stage (or equivalently, setting $\boldsymbol{q} = \boldsymbol{\delta}$, the Kronecker delta). For reasons that include its wide support even when windowed, its high computational cost, and results with an excessive amount of ringing, sinc has progressively lost favor to narrowly supported piecewise polynomial kernels, which bring performance and quality advantages [Meijering et al. 2001].

A typical use for the digital filtering stage is in interpolation. Absent $\boldsymbol{q}$, the interpolation property eliminates degrees of freedom from φ that could be used for better purposes. These constraints can be moved to $\boldsymbol{q}$ instead [Thévenaz et al. 2000], as in the case of interpolation by B-splines [Unser et al. 1991]. Another use for the digital filtering stage is in obtaining the best approximation of f

for a given choice of generator φ. Since the functions of interest in many applications have considerable bandwidth outside the Nyquist interval $(-\frac{0.5}{T}, \frac{0.5}{T})$, there is no hope of reaching an exact reconstruction. The goal is instead to minimize the L_2 norm of the residual $\|f - \tilde{f}_T\|_{L_2}$, and this goal uniquely determines both ψ and $\boldsymbol{q}$. Setting $\boldsymbol{q}^\vee * \psi = \mathring{\varphi}$, the dual of the generator φ, we obtain the orthogonal projection of f into the space of functions spanned by shifted copies of the generator φ.

The order with which the residual of the orthogonal projection vanishes as we progressively reduce the sample spacing T is a property of the generating function φ [Strang and Fix 1971]. Piecewise polynomial generators with minimal support and optimal approximation order were completely characterized by Blu et al. [2001]. To achieve the same approximation order as the orthogonal projection without access to ψ, as is our assumption, we must wisely select $\boldsymbol{q}$ [de Boor 1990; Blu and Unser 1999a]. The results are *quasi-interpolators*. Optimal order digital filters $\boldsymbol{q}$ for a given φ have also been obtained [Blu and Unser 1999b; Condat et al. 2005; Dalai et al. 2006].

Previous work has either assumed orthogonal projection and optimized for the generator, or assumed a given generator and optimized for the digital filter. As far as we know, our work is the first to jointly optimize all degrees of freedom in $\boldsymbol{q}$ and φ for the best quasi-interpolating scheme. Furthermore, inspired by Schaum [1993], our optimization framework takes into account the entire Nyquist interval. This is in contrast to the dominant strategy of focusing on the asymptotic behavior of the residual in the limit as $T \to 0$.

Our work results in a new family of piecewise polynomial quasi-interpolating schemes. Each scheme is given by a combination of digital filter $\boldsymbol{q}$ and generating function φ, and is optimal with regard to our metric. We run a variety of empirical tests to demonstrate that the resulting schemes yield superior reconstruction quality in typical tasks, when compared to the state-of-the-art, while maintaining competitive performance.

The remainder of the paper is structured in layers to help readers unfamiliar with the topic. Section 2 details our notation. In addition to a historical review of related work, section 3 presents important concepts from approximation theory and the motivation for our work. Section 4 delves deeper into theory and substantiates our motivation with concrete examples. Section 5 presents our optimization framework. The resulting interpolators and comparisons against the state-of-the-art appear in section 6. We conclude in section 7 with directions for future research on this topic.

2 Notation

Let $f : \mathbb{R} \to \mathbb{C}$ be a function. We say that $f \in L_2$ if

$$\int_{-\infty}^{\infty} \left|f(x)\right|^2 \mathrm{d}x < \infty. \tag{1}$$

In this case, we denote the quantity in (1) by $\|f\|_{L_2}$. The Fourier transform of f and its inverse are given by

$$\hat{f}(\omega) = \int_{-\infty}^{\infty} f(x) e^{-2\pi i x \omega} \,\mathrm{d}x, \quad \text{and} \tag{2}$$

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{f}(\omega) e^{2\pi i x \omega} \,\mathrm{d}\omega. \tag{3}$$

A sequence $\boldsymbol{q} : \mathbb{Z} \to \mathbb{C}$ belongs to ℓ_2 if $\sum_{k\in\mathbb{Z}} |q_k|^2 < \infty$. The discrete-time Fouriter transform (DTFT) of $\boldsymbol{q}$ is defined by

$$\hat{\boldsymbol{q}}(\omega) = \sum_{k\in\mathbb{Z}} q_k e^{-2\pi i \omega k}. \tag{4}$$

For a given real number $r > 0$ we define the Sobolev space $\mathbf{W}_2^r$ as the set of functions f that satisfy $\int_{-\infty}^{\infty} (1+\omega^2)^r |\hat{f}(\omega)|^2 \,\mathrm{d}\omega < \infty$.

Sampling f on a grid with spacing T amounts to obtaining

$$\boldsymbol{[}f\boldsymbol{]}_T := \big[\ldots, f(-T),\, f(0),\, f(T), \ldots\big]. \tag{5}$$

The flip of a function f and of a digital filter $\boldsymbol{q} = [\ldots, q_{-1},\, q_0,\, q_1, \ldots]$ are defined as

$$f^\vee(x) = f(-x) \quad \text{and} \quad (\boldsymbol{q}^\vee)_i = \boldsymbol{q}_{-i} \tag{6}$$

Given functions f and g, and sequences $\boldsymbol{c}$ and $\boldsymbol{q}$, the continuous, discrete, and mixed convolutions are given respectively by

$$(f * g)(x) = \int_{-\infty}^{\infty} f(t) g(x-t) \,\mathrm{d}t, \tag{7}$$

$$(\boldsymbol{c} * \boldsymbol{q})_n = \sum_{k\in\mathbb{Z}} c_k \, q_{n-k}, \quad \text{and} \tag{8}$$

$$(\boldsymbol{q} *_T f)(x) = \sum_{k\in\mathbb{Z}} q_k \, f(x - kT). \tag{9}$$

With these notations, we can conveniently express the output of the pipeline described in figure 2:

$$\tilde{f}_T = \boldsymbol{[}f * \psi^\vee(\,\cdot/T)\boldsymbol{]}_T * \boldsymbol{q} *_T \varphi(\,\cdot/T) \tag{10}$$

The auto-correlation of a function φ can be written equivalently as

$$a_\varphi(x) = (\varphi * \varphi^\vee)(x) = \int_{-\infty}^{\infty} \varphi(t)\, \varphi(t-x) \,\mathrm{d}t. \tag{11}$$

It is frequently sampled into a discrete sequence

$$\boldsymbol{a_\varphi} = \boldsymbol{[}a_\varphi\boldsymbol{]}. \tag{12}$$

The convolution inverse of $\boldsymbol{q}$, when it exists, is another sequence denoted by $\boldsymbol{q}^{-1}$ such that

$$\boldsymbol{q} * \boldsymbol{q}^{-1} = \boldsymbol{\delta} = [\ldots, 0, 0, 1, 0, 0, \ldots]. \tag{13}$$

The digital filter $\boldsymbol{q}$ is said do be FIR (for Finite Impulse Response) if it has a finite number of non-zero entries. It is said to be IFIR (for Inverse of Finite Impulse Response) if it is the inverse of a FIR digital filter. Finally, $\boldsymbol{q}$ is said to be FIR-IFIR if it can be written as $\boldsymbol{q} = \boldsymbol{q}_1 * \boldsymbol{q}_2$, where $\boldsymbol{q}_1$ is FIR and $\boldsymbol{q}_2$ is IFIR.

The subspace of functions generated by shifted copies of the generator φ widened to match the grid with spacing T is defined by

$$V_{\varphi,T} = \big\{\tilde{f} = \boldsymbol{c} *_T \varphi(\,\cdot/T) \;\big|\; \forall \boldsymbol{c} \in \ell_2\big\}. \tag{14}$$

To have the orthogonal projection of any function $f \in L_2$ on $V_{\varphi,T}$ well-defined, we require the latter to be a closed subspace of L_2. This is equivalent to imposing the existence of constants $A, B > 0$ such that $A \leqslant \widehat{\boldsymbol{a_\varphi}}(\omega) \leqslant B$, where $\boldsymbol{a_\varphi} = \boldsymbol{[}a_\varphi\boldsymbol{]}$ is the sampled auto-correlation of φ [Aldroubi 1996]. In this case, the orthogonal projection of f on $V_{\varphi,T}$ has the following expression:

$$P_{\varphi,T}(f) = \boldsymbol{[}f * \varphi^\vee(\,\cdot/T)\boldsymbol{]} * \boldsymbol{[}a_\varphi\boldsymbol{]}^{-1} *_T \varphi(\,\cdot/T). \tag{15}$$

The generator φ has approximation order L if L is the greatest positive integer for which there exists a constants $C > 0$ such that

$$\|f - P_{\varphi,T}(f)\|_{L_2} \leqslant C \cdot T^L \cdot \|f^{(L)}\|_{L_2}, \forall f \in \mathbf{W}_2^L. \tag{16}$$

Intuitively, this means that as T goes to zero, the "distance" between f and its closest function in $V_{\varphi,T}$ decays to zero as fast as T^L. The

same concept extends to any operator in the form of (10). We say an operator has approximation order $\tilde{L}$ (necessarily $\leqslant L$) if

$$\|f - \tilde{f}_T\|_{L_2} \leqslant \tilde{C} \cdot T^{\tilde{L}} \cdot \|f^{(\tilde{L})}\|_{L_2}, \forall f \in \mathbf{W}_2^L. \qquad (17)$$

The B-spline family of generators can be recursively defined as

$$\beta^0(x) = \begin{cases} 1, & |x| < 0.5 \\ 0.5, & |x| = 0.5 \\ 0, & |x| > 0.5 \end{cases}, \quad \text{and}$$
$$\beta^n = \beta^{n-1} * \beta^0, \quad n \geqslant 1. \qquad (18)$$

We use several concepts when describing a generator $\rho : \mathbb{R} \to \mathbb{R}$. It is *symmetric* if $\rho^\vee = \rho$; it has *support* $W > 0$ if W is the length of the smallest interval I for which $\rho(x) = 0, \forall x \notin I$; it is *interpolating* if $[\![\rho]\!] = \boldsymbol{\delta}$; it has regularity R if R is the greatest integer for which $f \in C^R$. Finally, if ρ is a piecewise polynomial function, its degree N is the greatest degree of its polynomial pieces.

3 Related work

The classic result by Shannon [1949] states that the scheme described in figure 2 can result in $\tilde{f}_T = f$ when f's bandwidth is restricted to the Nyquist interval $(-\frac{0.5}{T}, \frac{0.5}{T})$ when we select $\psi = \delta$, $\boldsymbol{q} = \delta$ and $\varphi = \mathrm{sinc}$. In the practical setting, when input functions are not necessarily band-limited and φ is required to be piecewise polynomial and compactly supported, Shannon's result does not apply and $f \neq \tilde{f}_T$.

For such cases, Strang and Fix [1971] established conditions under which the error $\|f - \tilde{f}_T\|_{L_2}$ goes to zero as a power of the sampling spacing T, when $\tilde{f}_T$ is the orthogonal projection of f into V_φ. They prove that $\tilde{f}_T$ has approximation order L if and only if V_φ contains all polynomials up to degree $L-1$. They also prove that interpolating f in V_φ has the same approximation order. Unser [1996] proved that if ψ is $\boldsymbol{q} * \varphi$ are bi-orthogonal then the scheme in figure 2 has the same approximation order of φ. Blu and Unser [1999a] completed the characterization of approximation order L by proving it to be equivalent to ψ and $\boldsymbol{q} * \varphi$ being quasi-biorthonormal of order L.

Over the years, many different approaches have been proposed for designing good generators φ. When only samples of f are available ($\psi = \delta$), the typical choice is to design interpolating generators φ, with the claim that this would lead to better approximations [Keys 1981; Schaum 1993; Dodgson 1997; German 1997]. These early works did not include the digital filter $\boldsymbol{q}$ in the approximation scheme (equivalent to taking $\boldsymbol{q} = \delta$).

Interpolation has the same approximation order L as φ [Strang and Fix 1971]. Unser et al. [1991] and Blu et al. [1999] advocate that the interpolation condition $\tilde{f}_T(k) = f(k), \forall k \in \mathbb{Z}$ is best enforced by the introduction of a digital filter $\boldsymbol{q} = [\![\varphi]\!]^{-1}$ to the sampling pipeline. This digital filtering stage can be performed very efficiently [Unser et al. 1991; Nehab et al. 2011]. The addition of the digital filtering stage leaves more freedom to design φ for increased approximation quality [Thévenaz et al. 2000]. The popular example is interpolation using B-splines ($\varphi = \beta^n$). B-splines have high approximation order, short support, and high regularity [Unser 1999]. The *cardinal* B-splines $\beta^n_{int} = [\![\beta^n]\!]^{-1} * \beta^n$ converge to sinc as n goes to infinity [Aldroubi and Unser 1994]. Additionally, they are very efficient to use as pre-filters [Heckbert 1986], digital filters [Unser et al. 1993], and generators [Sigg and Hadwiger 2005].

Explicit formulas for the *approximation constant* $\tilde{C}$ in (17), associated with an approximation scheme were developed [Möller et al. 1997; Blu and Unser 1999a]. Blu et al. [2001] parametrized all generators with minimum support and optimal approximation order in terms of B-splines and their derivatives. Using this parametrization, they obtained excellent generators that minimize the asymptotic constant of the orthogonal projection scheme (the O-MOMS). Although designed to be optimal for orthogonal projection, the O-MOMS were shown to be good cardinal interpolators as well [Thévenaz et al. 2000]. Blu et al. [2003] later provided a complete parametrization for generators φ in terms of their degree R, support W, regularity R, and order L. The general expression is a linear combination of B-splines and their convolution with certain distributions. Along this line of research, Blu et al. [2004] determined an optimal shift in the linear interpolation scheme such that it reaches the asymptotic constant of the orthogonal projection.

Interpolation is too strong a constraint. It completely defines the digital filter $\boldsymbol{q} = [\![\varphi]\!]^{-1}$. Giving up interpolation allows us to use the digital filter to improve reconstruction quality even further. These so called *quasi-interpolation* schemes were shown to have approximation order L whenever they exactly reproduce polynomials up to degree $L-1$ (i.e., $\tilde{f}_T(x) = f(x), \forall x \in \mathbb{R}$ if f is a polynomial of degree less than L) [de Boor 1990; Chui and Diamond 1990]. This in turn will be true whenever the combination of prefilter and digital filter $\psi^\vee * \boldsymbol{q}$ have the same moments as the dual $\tilde{\varphi}^\vee$ up to order $L-1$ [Blu and Unser 1999a]. This equivalence was explored in the design of digital filters for quasi-interpolators based on B-spline generators $\varphi = \beta^n$: Condat et al. [2005] proposes an IFIR design, Dalai et al. [2006] an FIR design, and Blu and Unser [1999b] propose a combination of FIR and IFIR filters.

The approximation order L and constant C describe the asymptotic behavior of the residual as $T \to 0$. In practice, this will be the dominant effect only when we are able to reduce T arbitrarily, or when the input signal has a narrow band around zero. The main motivation for our work is our belief that neither of these conditions apply in typical image processing and computer graphics applications. A better goal is to minimize the residual under some appropriate metric. Although a perceptual metric would be ideal in some applications [Zhang and Wandell 1996; Wang et al. 2004], more powerful tools are available to work with the L_2 metric.

Error kernels $E(\omega)$ allow us to separate, in the computation of the value of the residual $\|f - \tilde{f}_T\|_{L_2}$, the influence of the input f and the influence of the approximation scheme. The general result states

$$\|f - \tilde{f}_T\|^2_{L_2} \approx \int_{-\infty}^{\infty} \left|\hat{f}(\omega)\right|^2 E(\omega)\, \mathrm{d}\omega. \qquad (19)$$

Park and Schowengerdt [1983] obtained an expression for the error kernel when $\psi = \delta$, $\boldsymbol{q} = \delta$, and used it to determine optimal generators φ in the family of interpolating cubics. Schaum [1993] obtained a similar expression, but searched for more general generators φ and considered different classes of input spectra $\hat{f}$. A complete result in the form (19) for arbitrary ψ, $\boldsymbol{q}$, φ and $\hat{f}$ was obtained by Blu and Unser [1999a] using multiple generators, proving the equivalence between approximation order and quasi-biorthonormality. A version of (19) for a single generator φ (the case of interest in our work) was further detailed and analyzed by Blu and Unser [1999b].

Different works have had some success in optimizing for high approximation order L and low approximation constant C (which control the behavior of the residual in the limit $T \to 0$) as a proxy for lowering the magnitude of the error kernel in (19) (e.g., [Thévenaz et al. 2000; Blu et al. 2004]). In our work (see section 4) we provide concrete examples that show there is no direct connection between these goals. This is why we define our objective functions to minimize the expression in (19).

Unlike previous work, we obtain optimal quasi-interpolators by optimizing for, in addition to the degrees of freedom in the generator parametrization by [Blu et al. 2003], all degrees of freedom in the digital filter. In other words, we jointly optimize for both

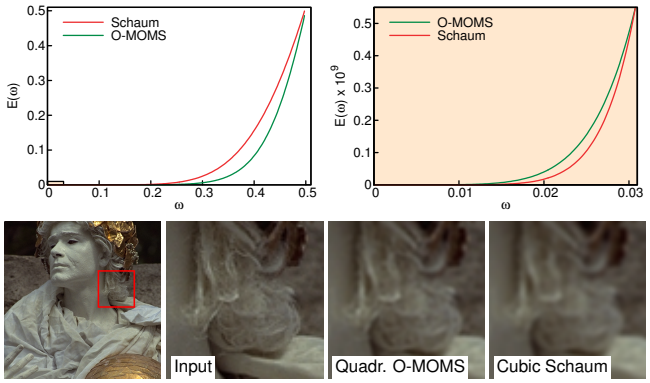

(a) PSNR = ∞, SSIM =1.0 (b) PSNR = 32.21, SSIM = 0.94 (c) PSNR = 29.86, SSIM = 0.90

Figure 3: *Comparison between the quadratic O-MOMS, a 3[rd]-order interpolator proposed by [Blu et al. 2001], and a 4[th]-order cubic by Schaum [1993]. Even with its lower order, O-MOMS's error kernel shows a better behavior overall in most of the Nyquist interval (top left). Detail (top right) shows that Schaum's is only better for a tiny portion of the spectrum near the origin. Comparison of 30 consecutive rotations confirm the better approximation qualities of the O-MOMS interpolator.*

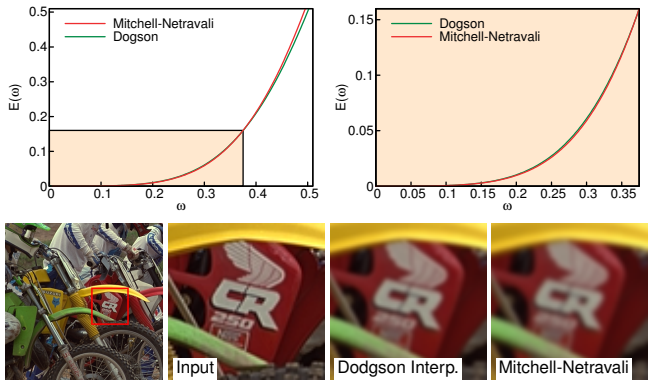

(a) PSNR = ∞, SSIM =1.0 (b) PSNR = 24.48, SSIM = 0.84 (c) PSNR = 22.88, SSIM = 0.76

Figure 4: *Comparison between a quadratic interpolator proposed by Dodgson [1997] and the cubic by Mitchell and Netravali [1988] (not interpolating), both with approximation order 2. Error kernels show the overall better behaviour of Dodgson's interpolator in the full Nyquist interval. This is despite its poorer behaviour near the origin (top right), as predicted by its higher asymptotic constant. Comparison of 15 consecutive translations show the higher quality achieved by Dodgson's interpolator.*

φ and $\boldsymbol{q}$. Our generators and digital filters do not change depending on the input f, nor on the image-processing operation being performed. Interpolation schemes that are input-dependent include [El-Khamy et al. 2005; Kopf et al. 2013].

4 Theory and motivation

We base our optimization problem on the following theorem due to Blu and Unser [1999a]. It quantifies the L_2-error between the input and output functions in the approximating scheme of figure 2:

Theorem 1: For all $f \in \mathbf{W}_2^r$ with $r > \frac{1}{2}$, the approximation error is given by

$$\|f - \tilde{f}_T\|_{L_2} = \left(\int_{-\infty}^{\infty} |\hat{f}(\omega)|^2 E(T\omega)\, \mathrm{d}\omega \right)^{\frac{1}{2}} + e(f,T), \quad (20)$$

where $e(f,T) = o(T^r)$ and

$$E(\omega) = 1 - \frac{\left|\hat{\varphi}(\omega)\right|^2}{\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)} + \widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega) \left| \hat{\boldsymbol{q}}(\omega)\hat{\psi}(\omega) - \frac{\hat{\varphi}(\omega)}{\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)} \right|^2. \quad (21)$$

Proof: See appendix C in [Blu and Unser 1999a].

The residual term $e(f,T)$ vanishes in many situations such as in the case where f is band-limited in the Nyquist interval [Blu and Unser 1999b]. Setting this term aside, formula (20) tells us that when most of the energy of the input is concentrated at low frequencies relative to the sampling spacing (i.e., for frequencies such that $T\omega \to 0$), we can obtain a small residual by simply requiring the error kernel E to vanish near $T\omega \to 0$.

This condition is satisfied by schemes with $L > 0$. Indeed, approximation order L is equivalent to all derivatives of E up to degree $2L-1$ vanishing at zero [Blu and Unser 1999b]. In turn, this causes the error kernel to behave as a polynomial of degree $2L$ near $\omega = 0$. This intuition led to significant effort being devoted towards the development of high approximation order schemes [Keys 1981; Unser 1996; German 1997; Blu et al. 2001].

When two different schemes have the same approximation order, the same intuition suggests using the asymptotic constant $\tilde{C}$ that appears in (17) to select the best ones [Blu et al. 2001; Thévenaz et al. 2000; Blu et al. 2004]. This can again be related to the error kernel (21), since this constant is proportional to the coefficient of the leading $(2L)^{\text{th}}$ power of the polynomial approximation of the error kernel around $\omega = 0$ [Blu and Unser 1999b].

We agree that a higher approximation order and small asymptotic constant can be important in many applications. However, as we show in figures 3 and 4, it is possible to find counter-examples for *both* criteria for exactly the applications that are typically used to showcase the approximation quality achieved by following them.

Figure 3 compares the 3[rd]-order (cardinal) quadratic interpolator OMOMS-2 [Blu et al. 2001] with the 4[th]-order cubic local Lagrangian interpolator [Schaum 1993] in an experiment that consists of 30 compounded rotations. At each rotation step, the input image is interpolated, and sampled at a $\frac{360^{\text{o}}}{30} = 12^{\text{o}}$ angle. The result is used as input for next rotation step and so on until the image is back to its initial position, at which point it is compared to the original input. The PSNR and SSIM [Wang et al. 2004] measures are higher (meaning higher quality) for the result with OMOMS-2 (Figure 3b), that has a lower order. The plot of the error kernels for both approximation schemes (Figure 3 top left) show that OMOMS-2 has a smaller value overall in the full Nyquist interval although it is worse for low frequencies (Figure 3 top right), the latter behavior being expected since it has lower approximation order.

In figure 4, we compare the performance of the quadratic interpolator designed by Dodgson [1997] with the (non-interpolating) cubic proposed by Mitchell and Netravali [1988]. Both these kernels have approximation order 2, so we would expect the one with smaller asymptotic constant to be better (the formula the constant is provided in [Blu and Unser 1999a]). In this case the constant for Dodgson's interpolator is slightly larger than Mitchell-Netravali's cubic's (by about 0.0004). Nevertheless, the compounded 15-translations in figure 4 show that Dodgson's interpolator generates a better result (Figure 4b). This is again due to a better behaviour in the full Nyquist interval (Figure 4 top left), although it is a bit worse for low frequencies (Figure 4 top right).

These counter-examples exist because the benchmarks violate the underlying assumption that the input frequency content is concentrated around $T\omega \to 0$. As is obvious from the images, they have significant frequency content away from $\omega \to 0$. Indeed, the input

power spectrum for natural images tend to behave as

$$\left|\hat{f}(\omega)\right|^2 \approx \frac{1}{\omega^p}, \tag{22}$$

where p varies from 1.6 to 3.0 [Field and Brady 1997; Ruderman 1997; Hsiao and Millane 2005]. A photograph taken underwater tends to be blurrier, so p will be higher. A photograph taken in the woods, where foliage produces high-frequency content, will have a smaller p. While the idea of taking $T \to 0$ *is* valid for numerical analysis applications that control the sampling spacing, we are not afforded the same freedom in most image-processing applications: we must therefore analyze the error kernel E over the entire frequency domain.

Recall we assume we only have access to the samples of f. If f was filtered by a good prefilter prior to sampling, the frequency content outside the Nyquist interval is close to zero. If not, whatever frequencies were outside the Nyquist interval have already been aliased back into it when the image was sampled. Therefore, rather than integrating on the real line as in (20), we focus on the Nyquist interval:

$$\|f - \tilde{f}_T\|_{L_2}^2 \approx \int_{-\frac{0.5}{T}}^{\frac{0.5}{T}} \left|\hat{f}(\omega)\right|^2 E(T\omega)\,\mathrm{d}\omega. \tag{23}$$

Since T is fixed, we may assume $T = 1$ with no loss of generality (see appendix B for proof).

We can now define our minimization problem:

$$\arg\min \int_{-0.5}^{0.5} \left|\hat{f}(\omega)\right|^2 E(\omega)\,\mathrm{d}\omega. \tag{24}$$

Assuming $\hat{f}$ known (to be detailed in section 5) and no prefiltering in the quasi-interpolation scheme ($\psi = \delta$ or $\hat{\psi} \equiv 1$), the degrees of freedom lie in the definitions of the digital filter $\boldsymbol{q}$ and the generator φ.

We explore three options for the form of digital filter $\boldsymbol{q}$, FIR, IFIR, and FIR-IFIR. Formally,

$$\text{FIR}: \ \boldsymbol{q} = [\ldots, 0, d_{-j}, \ldots, d_0, \ldots, d_j, 0, \ldots], \tag{25}$$

$$\text{IFIR}: \ \boldsymbol{q} = [\ldots, 0, e_{-k}, \ldots, e_0, \ldots, e_k, 0, \ldots]^{-1}, \quad \text{and} \tag{26}$$

$$\text{FIR-IFIR}: \ \boldsymbol{q} = \boldsymbol{d} * \boldsymbol{e}. \tag{27}$$

These formulations provide us with $2j+1$, $2k+1$, and $2(j+k+1)$ degrees of freedom, respectively.

To isolate the degrees of freedom in the generator in a meaningful way, we use the parametrization by Blu et al. [2003] in terms of its degree N, support W, regularity R, and approximation order L (for simplicity, we write $\varphi \in \{N, W, L, R\}$).

Theorem 2: Given $W \geqslant \mathrm{N}$, $\varphi \in \{N, W, R, L\}$ if and only if there exists a unique set of coefficients $a_{k,\ell}$, $b_{k,\ell}$, and $c_{k,\ell}$ such that

$$\begin{aligned}\varphi\big(x - \tfrac{W}{2}\big) = &\sum_{\ell=1}^{M} \sum_{k=0}^{N\text{-}L\text{-}\ell} a_{k,\ell} \left(\beta_{nc}^{L+k\text{-}1} * \gamma_\ell^{N\text{-}L\text{-}k}\right)(x) \\ &+ \sum_{\ell=0}^{M} \sum_{k=0}^{W\text{-}N+\ell\text{-}1} b_{k,\ell}\, \beta_{nc}^{N\text{-}\ell}(x-k) \\ &+ \sum_{k=0}^{W\text{-}L} \sum_{\ell=0}^{L\text{-}R\text{-}2} c_{k,\ell}\, \Delta^{*\ell}\, \beta_{nc}^{L\text{-}\ell\text{-}1}(x-k),\end{aligned} \tag{28}$$

where $M = N - \max(R+1, L)$.

Proof: See [Blu et al. 2003].

In the formulas above,

$$\beta_{nc}^n(x) = \beta^n\big(x - \tfrac{n+1}{2}\big), \tag{29}$$

is the non-centered B-spline, $\Delta^{*\ell}$ is the ℓ^{th}-order finite difference, and γ_ℓ^n are distributions (e.g., derivatives and shifts).

For example, setting $N = 1$, $W = 2$, $R = -1$ (meaning φ is bounded), and $L = 1$ in the decomposition theorem produces

$$\varphi(x) = b_{0,0}\beta^1(x) + c_{0,0}\beta^0\big(x + \tfrac{1}{2}\big) + c_{1,0}\beta^0\big(x - \tfrac{1}{2}\big) \tag{30}$$

This gives us 3 additional degrees of freedom, relative to the common choice of $\varphi = \beta^1(x)$ [Condat et al. 2005; Dalai et al. 2006], with which we minimize our objective function.

5 Optimization

We now state the minimization problem that will result in optimal quasi-interpolators. Before the objective function itself, we detail the constraints.

Degree and width of φ The degree N is the guiding parameter in our method. We set the width to $W = N + 1$ to match the run-time efficiency of generators such as B-splines and O-MOMS.

Regularity of φ The only restriction we impose is boundedness ($R = -1$). Several authors have observed that regularity is not fundamental for good approximation quality Schaum [1993]; Blu et al. [2001]. Our results confirm this. Applications requiring more regularity (e.g., for derivatives) can change this parameter in the optimization.

Approximation order of φ In stark contrast to previous work, we only first-order approximation ($L = 1$). This means that frequency $\omega = 0$ (i.e., DC or the average input value) will be preserved, but nothing else. In analogy to the regularity constraint, our results show that these additional degrees of freedom are better left to the discretion of the optimizer.

These constraints determine the coefficients in (28) that are available for minimization. We encapsulate them into lists of coefficients $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{C}$:

$$\mathbf{A} = \{a_{k,\ell}\}, \quad \mathbf{B} = \{b_{k,\ell}\}, \quad \mathbf{C} = \{c_{k,\ell}\}. \tag{31}$$

Symmetry of φ and $\boldsymbol{q}$ To guarantee linear phase, we require our quasi-interpolators $\boldsymbol{q} * \varphi$ to be symmetric. The condition imposes simple linear relationships between the coefficients $a_{k,\ell}$, $b_{k,\ell}$, and $c_{k,\ell}$, and sets $\boldsymbol{d}_i = \boldsymbol{d}_{-i}$, and $\boldsymbol{e}_i = \boldsymbol{e}_{-i}$, for all i.

Unit scale for φ and $\boldsymbol{q}$ There is a scale ambiguity within the remaining degrees of freedom. Scaling φ by s and $\boldsymbol{q}$ by $\frac{1}{s}$ leaves the quasi-interpolator $\boldsymbol{q} * \varphi$ unchanged. We therefore impose

$$\int_{-\infty}^{\infty} \varphi(x)\,\mathrm{d}x = 1, \quad \text{and} \quad \sum_{i\in\mathbb{Z}} \boldsymbol{d}_i = \sum_{i\in\mathbb{Z}} \boldsymbol{e}_i = 1. \tag{32}$$

Approximation order of the scheme We also require the *scheme as a whole* to have first order of approximation. The generator φ satisfies the restriction by construction, but a misguided choice of $\boldsymbol{q}$ could ruin it. The equivalent condition on the error kernel is

$$E(0) = 0. \tag{33}$$

See [Blu and Unser 1999b] for the proof.

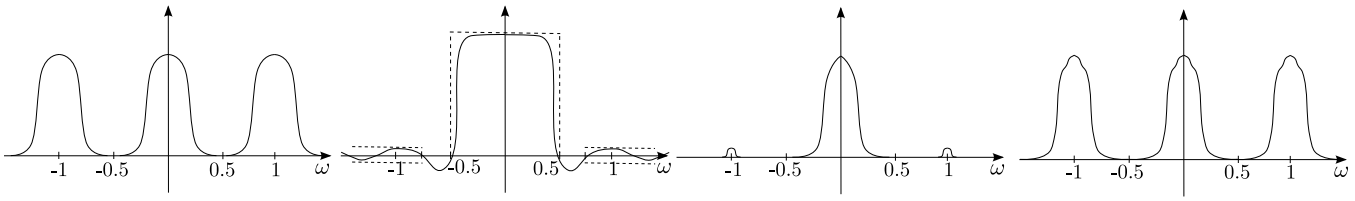

(a) $DTFT([\![f * \psi^\vee]\!])(\omega)$ (b) $\hat{\varphi}_{\text{qi}}(\omega)$ (c) $DTFT([\![f * \psi^\vee]\!])(\omega) \cdot \hat{\varphi}_{\text{qi}}(\omega)$ (d) $DTFT([\![[\![f * \psi^\vee]\!] * \varphi_{\text{qi}}]\!])(\omega)$

Figure 5: *By imposing $\hat{\varphi}_{qi}(\omega) \leqslant 1.0025,\ \forall\omega \in [-0.5, 0.5]$ and $|\hat{\varphi}_{qi}(\omega)| \leqslant 0.025,\ \forall\omega \in [-\infty, -0.75] \cup [0.75, \infty]$ we control overshoot and aliasing in the re-sampled image. The DTFT of the input $[\![f * \psi^\vee]\!]$ (a) is multiplied by $\hat{\varphi}_{qi}$ (b) and results in a spectrum with small out-band (aliasing) spectrum and non-amplified in-band spectrum (c). Resampling it leads to $[\![[\![f * \psi^\vee]\!] * \varphi_{qi}]\!]$ (whose DTFT is shown in figure d) with controlled frequency amplification.*

Objective function Recall the spectrum of natural images tend to follow (22). Since we seek input-independent quasi-interpolators, we set p to the intermediate value of $p = 2$:

$$\left|\hat{f}(\omega)\right|^2 \approx \frac{1}{\omega^2}. \tag{34}$$

This choice has an extra advantage in our formulation. Since we are imposing $E(0) = 0$ and since $E'(0) = 0$ is automatically satisfied due to the symmetry of the error kernel, we have $E(\omega)$ proportional to ω^2 near the origin. This causes the integrand in (24) to converge to a finite value at the origin.

The optimization problem Given a degree N:

$$\underset{\boldsymbol{q},\mathbf{A},\mathbf{B},\mathbf{C}}{\arg\min} \quad F(d) := \int_0^d \frac{1}{\omega^2} E(\omega)\,\mathrm{d}\omega \tag{35}$$

$$\text{subject to} \quad \varphi \in \{N, N+1, -1, 1\}, \tag{36}$$

$$\varphi^\vee = \varphi, \quad \boldsymbol{q}^\vee = \boldsymbol{q}, \tag{37}$$

$$\int \varphi(x)dx = 1, \quad \sum \boldsymbol{q}_k = 1, \tag{38}$$

$$E(0) = 0. \tag{39}$$

(We can restrict the integral to positive ω because of symmetry.)

Controlling overshoot and aliasing The natural choice for the integration limit d in (35) is 0.5, since we only have access to samples of f. Unfortunately, this often results in quasi-interpolators with highly oscillating spectra, such as the one presented in figure 6c.

By minimizing (35) with $d = 0.5$ we are requiring the error kernel to be small near $\omega = 0.5$, say $E(0.5 - \varepsilon) \approx 0$. As shown in appendix B, this implies

$$\begin{aligned} \hat{\varphi}_{\text{qi}}(0.5 - \varepsilon) &\approx 1, \quad \hat{\varphi}_{\text{qi}}(0.5 + \varepsilon) \approx 0, \\ \hat{\varphi}_{\text{qi}}(\text{-}0.5 + \varepsilon) &\approx 1, \quad \hat{\varphi}_{\text{qi}}(\text{-}0.5 - \varepsilon) \approx 0. \end{aligned} \tag{40}$$

Thus $E(0.5 - \varepsilon) \approx 0$ leads to $\hat{\varphi}_{\text{qi}}(\omega) = \hat{\boldsymbol{q}}(\omega)\hat{\varphi}(\omega)$ that approximates a function with discontinuities near $\omega = \pm 0.5$. Since $\hat{\varphi}(\omega)$ cannot oscillate much (see [Blu et al. 2003] for the expression), $\hat{\boldsymbol{q}}(\omega)$ is responsible for approximating the discontinuities near $\omega = \pm 0.5$. Since the filter has a finite support in the form of (4) or its reciprocal, this leads to the Gibbs phenomenon in $\hat{\boldsymbol{q}}(\omega)$, which is modulated by $\hat{\varphi}(\omega)$ and manifests itself as ringing in the reconstructed images (figure 6b).

To prevent this issue, we only consider quasi-interpolators that satisfy the following admissibility conditions:

$$\hat{\varphi}_{\text{qi}}(\omega) \leqslant 1.0025, \quad \forall\omega \in [-0.5, 0.5] \quad \text{and} \tag{41}$$

$$|\hat{\varphi}_{\text{qi}}(\omega)| \leqslant 0.025, \quad \forall\omega \in [-\infty, -0.75] \cup [0.75, \infty]\,. \tag{42}$$

Intuitively, condition (41) prevents overshoot and condition (42) prevents aliasing. The values 1.0025, 0.025 and 0.75 were empirically determined. To solve the optimization problem, we relax the objective function by performing a binary search for the largest value of $d \in [0, 0.5]$ in (35) that leads to an admissible quasi-interpolator.

Figure 5 illustrates the importance of conditions (41) and (42) by following the effects of each stage of the sampling pipeline, in the frequency domain. The DTFT of the input $[\![f * \psi^\vee]\!]$ is modulated by the spectrum of the quasi-interpolator $\hat{\varphi}_{\text{qi}}$. The resampling step then replicates this spectrum. Conditions (41) and (42) control magnification of both in-band and out-band spectra.

Note that condition (42) skips interval $(0.5, 0.75)$. In fact, for $N = 1$, even this relaxed condition is too restrictive. We therefore test only condition (41). Degrees $N = 2$ and 3 have larger parameter spaces, and we can find a value for d that satisfies both constraints.

The practical effect of the admissibility conditions can be seen in the example of figure 6. There, the quasi-interpolator that results from the optimization with $d = 0.5$ leads to overshoot in high frequencies (note ringing surrounding thorns). The binary search finds the value $d \approx 0.34$. The resulting quasi-interpolator is softer, but is still sharp enough. The overshooting is mostly gone.

Length and type of $\boldsymbol{q}$ We solve (35)–(39) using FIR, IFIR and FIR-IFIR digital filters. FIR filters led to the worst results, both w.r.t. the objective function (24) and our interpolation experiments. Since IFIR and FIR-IFIR formulations lead to similar results, we prefer the lower cost IFIR. The wider $\boldsymbol{q}$ is (i.e., the more degrees of freedom it offers), the lower objective function values are obtained. However, very little is gained for widths greater than 5. All our results assume a width 5 and an IFIR formulation for the digital filter.

6 Results and discussion

We have implemented this optimization framework in Mathematica, selecting the optimization method by Nelder and Mead [1965]. This method is suitable for constrained non-linear problems, and worked best in practice with our objective function. To reduce the risk of finding poor local minima, we solve each optimization problem 40 times and select the best result.

The objective function is somewhat brittle, due to the integrand in (35) being unstable near the origin. We were careful to keep the error kernel in its simplest possible form to avoid numerical round-off errors. All calculations were performed with 20-digit precision. The Quasi-Monte Carlo method gave the most robust results for the numerical integration of (35).

The values for all arguments in the solution to the optimization problem of degrees 1, 2, and 3 are given in appendix A. For convenience, we also provide the source-code for the generators in the supplemental materials and the digital filter entries. We compare the practical performance of our reconstruction schemes against previous quasi-interpolators by performing a variety of experiments.

Figure 7 shows a plot of our quadratic generator φ (a) and the quasi-interpolator $\varphi_{\text{qi}} = \boldsymbol{q} * \varphi$ (b). As has been observed in previous work,

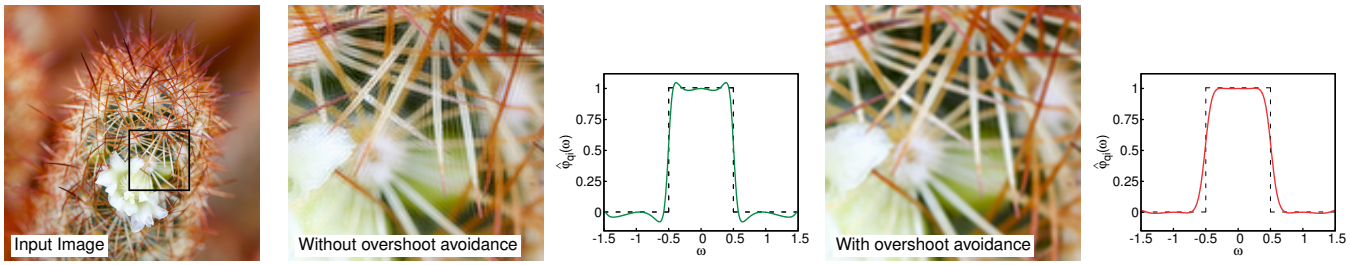

(a) PSNR = ∞, SSIM =1.0 (b) PSNR = 34.662, SSIM = 0.995 (c) Frequency response (d) PSNR = 43.812, SSIM = 0.999 (e) Frequency response

Figure 6: *Quadratic interpolation result for 20 compounded translations. Minimizing $F(0.5)$ leads to a quasi-interpolator φ_{qi} that overshoots high frequencies (b). This problem is avoided by minimizing $F(0.34)$ (d), where $d = 0.34$ is automatically obtained by a binary search. Plots in figures (c) and (e) show the frequency response of the associated quasi-interpolators compared with the frequency response of the ideal interpolator.*

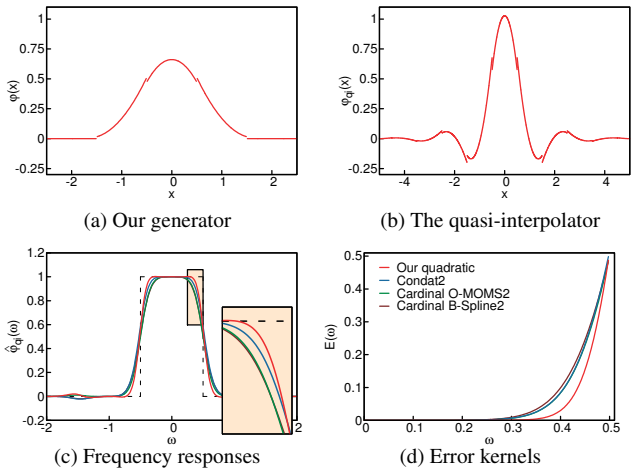

(a) Our generator (b) The quasi-interpolator (c) Frequency responses (d) Error kernels

Figure 7: *Impulse responses of our quadratic generator (a) and quasi-interpolator (b). We compare the frequency response of our quasi-interpolator with the best quadratics (c), showing ours to be closer to the ideal interpolator. Plot in (d) shows that the error kernel associated to our method is smaller in most of the Nyquist interval.*

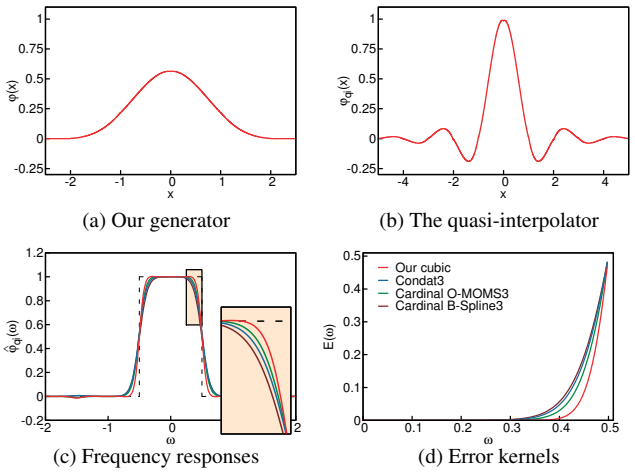

(a) Our generator (b) The quasi-interpolator (c) Frequency responses (d) Error kernels

Figure 8: *Impulse responses of our cubic generator (a) and quasi-interpolator (b). Its frequency response compared with best cubics (c) and the associated error kernels are shown in (d).*

regularity is not fundamental for achieving good approximation quality. Like the local Lagrangian interpolators of Schaum [1993] and the OMOMS-2 of Blu et al. [2001], our quadratic quasi-interpolator is not even continuous. The figure also shows a comparison between the frequency response $\hat{\varphi}_{qi}$ of our quasi-interpolator with the state-of-the art in quadratics (c). It is clear our interpolator is sharper. Furthermore, the error kernel plots (d) show that our quasi-interpolator has a lower error overall in the Nyquist interval. Figure 8 shows the same analysis, this time for our cubic quasi-interpolator. Similar conclusions can be drawn. Please note that the improvements due to our new quasi-interpolators is more marked than the quality differences between the previous state-of-the-art.

Figure 9 shows results for our linear quasi-interpolator. The tests add a random perturbation to each compounded translation offset in order to rule out the possibility of errors being cancelled by negative correlations. Our results are significantly sharper than those obtained by state-of-the art linear quasi-interpolator proposed by Condat et al. [2005]. In fact, our results compare favourably even against the cardinal *quadratic* B-spline.

The example in figure 10 shows that our quadratic quasi-interpolator performs better than the one proposed by Condat et al. [2005]. In fact, our quadratic compares favourably against the cardinal *cubic* O-MOMS, which is the state-of-the-art in cubic interpolation [Thévenaz et al. 2000].

Figure 13 tests the performance of our cubic quasi-interpolator with a challenging task of rotating a high-frequency pattern consisting of parallel lines. Our result shows almost perfect reconstruction. The cubic quasi-interpolator proposed by Blu and Unser [1999b] (which uses a wider FIR-IFIR formulation) and, to a lesser extent, the quintic cardinal B-spline, show aliasing in the form of spurious slanted lines. This final example helps emphasize one of the key points in our paper: the quintic cardinal B-spline has approximation order 6, and our cubic has only approximation order 1. Nevertheless, our cubic performs better.

To put all these results in context, we ran an additional experiment. We applied 90 randomized translations to the images in [Kodak 2010] in such a way that after every 3 translations it goes back to the initial position. At these points, we can measure PSNR against the input. To obtain a single number, we average the PSNR results over the 24 input images. Results can be seen in figure 11. Our cubic quasi-interpolator performs best, even when compared to quintic quasi-interpolators. Our quadratic quasi-interpolator performed better than any other quadratic and cubic.

We recommend the use of our solutions for degrees 2 and 3, given their superior performance and moderate computational cost. Tasks requiring even more speed can use the degree 1 solution, which we take as a proof-of-concept. We refer the reader to the supplemental material for full resolution images of all these experiments. We also provide additional videos containing other interpolation sequences.

(a) PSNR = ∞, SSIM = 1.0 (b) PSNR = 25.873, SSIM = 0.862 (c) PSNR = 27.000, SSIM = 0.898 (d) PSNR = 27.642, SSIM = 0.913

Figure 9: *Result of 9 repeated (randomized) translations. Our linear quasi-interpolator (d) produces a result sharper than the one produced by one of the best linear quasi-interpolators [Condat et al. 2005] (b). The output by our method is also slightly better than the one produced with the cardinal quadratic B-spline (c).*

(a) PSNR = ∞, SSIM =1.0 (b) PSNR = 30.429, SSIM = 0.932 (c) PSNR = 32.221, SSIM = 0.951 (d) PSNR = 33.291, SSIM = 0.960

Figure 10: *Comparison of 30 (randomized) rotations for different quasi-interpolators. The quadratic proposed by Condat et al. [2005] distorts the vertical aspect of the fence, while ours better preserves the geometry of the scene. Our result is competitive even if compared with the one produces by the cardinal cubic O-MOMS (considered the best cubic in the literature).*

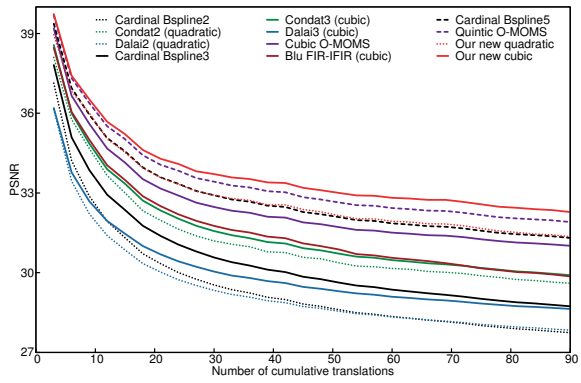

Figure 11: *Average PSNR of applying 90 randomized translations to 24 input images. Translations were applied in a way that after every 3 translations, the image was back to its initial position and we could measure PSNR. The best quasi-interpolators in the literature were compared. Our new cubic quasi-interpolator (solid red) reaches the best quality, better than the* quintic *O-MOMS (dashed purple). Our new quadratic (dotted red) reaches higher quality than any other quadratic and cubic.*

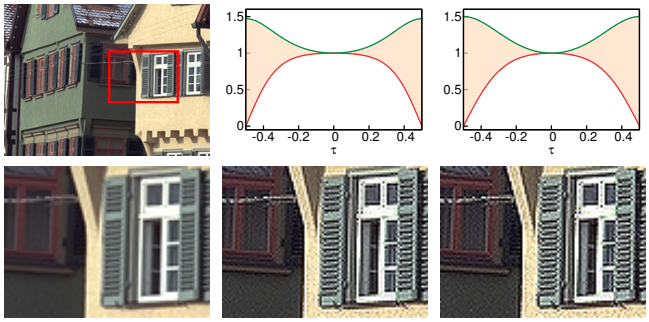

(a) Input (b) Our linear (c) Condat et al. [2005]

Figure 12: *Result of 4 repeated translations by exactly half pixel. Our linear quasi-interpolator (b) and the one proposed by Condat et al. [2005] magnify frequencies too much. Shaded regions in the plots show the range of frequency amplitude scaling for each translation τ. The worst case is $\tau = 0.5$.*

6.1 Limitations

One limitation of our method can be seen in figure 12, which uses our linear quasi-interpolator (b). The figure shows the result of 4 compound translations by exactly half a pixel. It is clear that high-frequencies have been excessively magnified. This limitation is not specific to to our approach (c). The plots in figure 12 explain the problem: for each translation τ, the shaded region illustrates the minimum and maximum possible frequency amplitude scaling. The worst behavior happens in the unfortunate case $\tau = 0.5$. This problem practically disappears when random translations are applied.

As future work, we will incorporate new criteria in our optimization framework to reduce this effect.

We have also noticed that the sharpness of our results comes at the cost additional mild ringing (for instance, see figure 1d). For hundreds of repeated translations, our linear and quadratic quasi-interpolators showed excessive ringing. In this extreme case, other methods presented either a similar behavior or excessive blurring.

7 Conclusion and future work

We have presented a new class of quasi-interpolators for image processing that are optimal with respect to a non-asymptotic criterion. In contrast, previous strategies focused on making them optimal only around $\omega = 0$. Additionally, we used all available degrees of freedom in the approximation problem to reach higher quality.

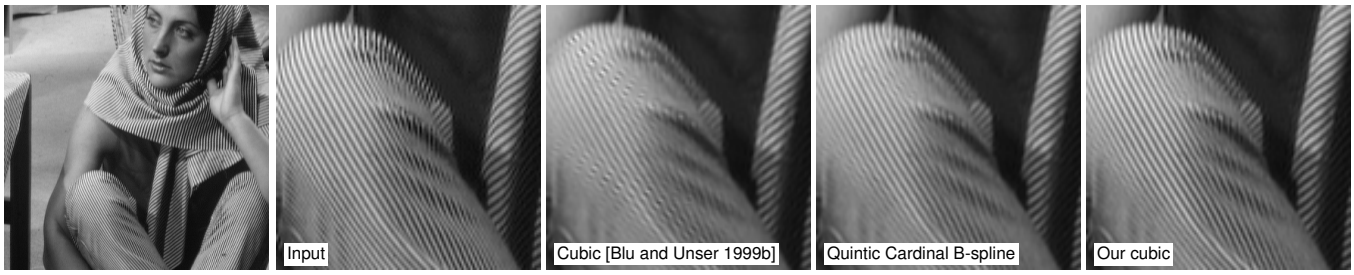

(a) PSNR = ∞, SSIM = 1.0 (b) PSNR = 13.013, SSIM = 0.749 (c) PSNR = 13.047, SSIM = 0.775 (d) PSNR = 13.044, SSIM = 0.784

Figure 13: *40 compounded rotations. The result produced by one of the best cubic quasi-interpolators [Blu and Unser 1999b] has aliasing of high frequencies. Using the cardinal quintic B-spline presents the same problem at a smaller magnitude. Our cubic almost completely removes the artefacts, while keeping the result sharp.*

An improvement to our optimization would be to consider a metric other than L_2. On one hand, it could lead to a more natural treatment for the overshoot problem, but it could also add additional difficulties to the optimization.

In this work we have considered a 1D formulation of the approximation problem, but applied it to images in a separable fashion. We believe that considering non-separable 2D quasi-interpolators will increase approximation quality, and we also consider this direction for future work.

A Quasi-Interpolators

Linear $d = 0.5$.

$$\begin{aligned} b_{0,0} &= 0.79076352 & c_{0,0} = c_{1,0} &= 0.10461824,\\ e_0 &= 0.77412669, & e_1 = e_{\text{-}1} &= 0.11566267,\\ e_2 = e_{\text{-}2} &= -0.00272602 \end{aligned}$$

Quadratic $d \approx 0.34$.

$$\begin{aligned} a_{0,1} &= 0, & b_{0,0} &= 0.75627421,\\ b_{0,1} = b_{1,1} &= 0.11798097, & c_{0,0} = c_{2,0} &= 0.01588197,\\ c_{1,0} &= -0.02400002, & e_0 &= 0.65314970,\\ e_1 = e_{\text{-}1} &= 0.17889730, & e_2 = e_{\text{-}2} &= -0.00547216 \end{aligned}$$

Cubic $d \approx 0.35$.

$$\begin{aligned} a_{0,1} &= 0.07922533, & a_{0,2} &= 0,\\ a_{1,1} = -2.25\, a_{0,1} &= -0.17825701, & b_{0,0} &= 0.53954836,\\ b_{0,1} &= 0.32092636, & b_{0,2} = b_{2,2} &= 0.02593862,\\ b_{1,1} = -1.5\, a_{0,1} + b_{0,1} &= 0.20208835, & b_{1,2} &= -0.01871558,\\ c_{0,0} = c_{3,0} &= 0.001940114, & c_{1,0} = c_{2,0} &= -0.00028665,\\ e_0 &= 0.56528428, & e_1 = e_{\text{-}1} &= 0.21523558,\\ e_2 = e_{\text{-}2} &= 0.00212228. \end{aligned}$$

B Proofs

We can take $T = 1$ in (23) because, for any fixed $T > 0$, and assuming input spectra as (22):

$$\begin{aligned} \|f - \tilde{f}_T\|_{L_2}^2 &\approx \textstyle\int_{-0.5/T}^{0.5/T} \left|\hat{f}(\omega)\right|^2 E(T\omega)\,\mathrm{d}\omega\\ &= T \textstyle\int_{-0.5}^{0.5} \left|f(\omega/T)\right|^2 E(\omega)\,\mathrm{d}\omega\\ &= T^{p+1} \textstyle\int_{-0.5}^{0.5} 1/\omega^p\, E(\omega)\,\mathrm{d}\omega\\ &= T^{p+1} \textstyle\int_{-0.5}^{0.5} \left|\hat{f}(\omega)\right|^2 E(\omega)\,\mathrm{d}\omega. \end{aligned} \tag{43}$$

To see why $E(0.5-\varepsilon) \approx 0$ implies $\hat{\varphi}_{\text{qi}}(0.5-\varepsilon) \approx 1$, $\hat{\varphi}_{\text{qi}}(0.5+\varepsilon) \approx 0$, $\hat{\varphi}_{\text{qi}}(\text{-}0.5+\varepsilon) \approx 1$, $\hat{\varphi}_{\text{qi}}(\text{-}0.5-\varepsilon) \approx 0$, recall $\hat{\psi} = 1$, and both $\hat{\boldsymbol{q}}$ and $\hat{\varphi}$ are real due to symmetry. The error kernel simplifies to:

$$\begin{aligned} E(\omega) &= 1 - \frac{\varphi(\omega)^2}{\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)} + \widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)\left(\hat{\boldsymbol{q}}(\omega)^2 - 2\frac{\hat{\boldsymbol{q}}(\omega)\hat{\varphi}(\omega)}{\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)} + \frac{\hat{\varphi}(\omega)^2}{\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)^2}\right)\\ &= 1 - 2\hat{\boldsymbol{q}}(\omega)\hat{\varphi}(\omega) + \hat{\boldsymbol{q}}(\omega)^2\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)\\ &= \big(1-\hat{\boldsymbol{q}}(\omega)\hat{\varphi}(\omega)\big)^2 - \hat{\boldsymbol{q}}(\omega)^2\hat{\varphi}(\omega)^2 + \hat{\boldsymbol{q}}(\omega)^2\sum_n \hat{\varphi}(\omega+n)^2\\ &= \big(1-\hat{\boldsymbol{q}}(\omega)\hat{\varphi}(\omega)\big)^2 + \sum_{n\neq 0} \hat{\boldsymbol{q}}(\omega+n)^2\hat{\varphi}(\omega+n)^2\\ &= \big(1-\hat{\varphi}_{\text{qi}}(\omega)\big)^2 + \sum_{n\neq 0} \hat{\varphi}_{\text{qi}}(\omega+n)^2. \end{aligned} \tag{44}$$

Above, we used the following equalities

$$\hat{\varphi}_{\text{qi}}(\omega) = \hat{\boldsymbol{q}}(\omega)\hat{\varphi}(\omega) \tag{45}$$

$$\hat{\boldsymbol{q}}(\omega) = \hat{\boldsymbol{q}}(\omega+n), \forall n \in \mathbb{N}, \quad \text{and} \tag{46}$$

$$\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega) = \sum_n \hat{\varphi}(\omega+n)^2. \tag{47}$$

The sum of non-negative terms in (44) shows us that

$$E(0.5-\varepsilon) \approx 0 \Rightarrow \hat{\varphi}_{\text{qi}}(0.5-\varepsilon) \approx 1, \hat{\varphi}_{\text{qi}}(\text{-}0.5-\varepsilon) \approx 0. \tag{48}$$

The symmetry of E implies $E(\text{-}0.5+\varepsilon) \approx 0$. From (44), we have

$$E(\text{-}0.5+\varepsilon) \approx 0 \Rightarrow \hat{\varphi}_{\text{qi}}(\text{-}0.5+\varepsilon) \approx 1, \hat{\varphi}_{\text{qi}}(0.5+\varepsilon) \approx 0. \tag{49}$$

References

ALDROUBI, A. 1996. Oblique projections in atomic spaces. *Proceedings of the American Mathematical Society*, 1244(7):2051–2060.

ALDROUBI, A. and UNSER, M. 1994. Sampling procedures in function spaces and asymptotic equivalence with Shannon's sampling theory. *Numerical Functional Analysis and Optimization*, 15(1–2):1–21.

BLU, T., THÉNAVAZ, P., and UNSER, M. 1999. Generalized interpolation: Higher quality at no additional cost. In *Proceedings of the IEEE International Conference on Image Processing*, volume 3, pages 667–671.

BLU, T., THÉVENAZ, P., and UNSER, M. 2001. MOMS: Maximal-order interpolation of minimal support. *IEEE Transactions on Image Processing*, 10(7):1069–1080.

BLU, T., THÉVENAZ, P., and UNSER, M. 2003. Complete parametrization of piecewise-polynomial interpolation kernels. *IEEE Transactions on Image Processing*, 12(11):1297–1309.

BLU, T., THÉVENAZ, P., and UNSER, M. 2004. Linear interpolation revitalized. *IEEE Transactions on Image Processing*, 13(5):710–719.

BLU, T. and UNSER, M. 1999. Approximation error for quasi-interpolators and (multi-)wavelet expansions. *Applied and Computational Harmonic Analysis*, 6(2):219–251.

BLU, T. and UNSER, M. 1999. Quantitative Fourier analysis of approximation techniques: Part I—Interpolators and projectors. *IEEE Transactions on Signal Processing*, 47(10):2783–2795.

CHUI, C. and DIAMOND, H. 1990. A characterization of multivariate quasi-interpolation formulas and its applications. *Numer. Math.*, 57(2):105–121.

CONDAT, L., BLU, T., and UNSER, M. 2005. Beyond interpolation: optimal reconstruction by quasi-interpolation. In *Proceedings of the IEEE International Conference on Image Processing*, volume 1, pages 33–36.

DALAI, M., LEONARDI, R., and MIGLIORATI, P. 2006. Efficient digital pre-filtering for least-squares linear approximation. In *Visual Content Processing and Representation*, volume 3893 of *Lecture Notes in Computer Science*, pages 161–169.

DE BOOR, C. 1990. In M. Gasca and C. A. Micchelli, editors, *Computation of Curves and Surfaces*, chapter Quasiinterpolants and the approximation power of multivariate splines. Kluwer Academic.

DODGSON, N. A. 1997. Quadratic interpolation for image resampling. *IEEE Transactions on Image Processing*, 6(9):1322–1326.

EL-KHAMY, S. E., HADHOUD, M. M., DESSOUKY, M. I., SALAM, B. M., and ABD EL-SAMIE, F. E. 2005. An adaptive cubic convolution image interpolation approach. *International Journal of Machine Graphics & Vision*, 14(3):235–258.

FIELD, D. J. and BRADY, N. 1997. Visual sensitivity, blur and the sources of variability in the amplitude spectra of natural scenes. *Vision Research*, 37(23):3367–3383.

GERMAN, I. 1997. Short kernel fifth-order interpolation. *IEEE Transactions on Signal Processing*, 45(5):1355–1359.

HECKBERT, P. S. 1986. Filtering by repeated integration. *Computer Graphics (Proceedings of ACM SIGGRAPH 1986)*, 20(4):315–321.

HSIAO, W. H. and MILLANE, R. 2005. Effects of occlusion, edges, and scaling on the power spectra of natural images. *Journal of the Optical Society of America A*, 22(9):1789–1797.

KEYS, R. G. 1981. Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 29(6):1153–1160.

KODAK. 2010. True colour kodak images. URL `http://r0k.us/graphics/kodak/`.

KOPF, J., SHAMIR, A., and PEERS, P. 2013. Content-adaptive image downscaling. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH Asia 2013)*, 32(6):173.

MEIJERING, E. H. W. 2002. A chronology of interpolation: From ancient astronomy to modern signal processing. *Proceedings of the IEEE*, 90(3):319–342.

MEIJERING, E. H. W., NIESSEN, W. J., and VIERGEVER, M. A. 2001. Quantitative evaluation of convolution-based methods for medical image interpolation. *Medical Image Analysis*, 5(2):111–126.

MITCHELL, D. P. and NETRAVALI, A. N. 1988. Reconstruction filters in computer graphics. *Computer Graphics (Proceedings of ACM SIGGRAPH 1988)*, 22(4):221–228.

MÖLLER, T., MACHIRAJU, R., MUELLER, K., and YAGEL, R. 1997. Evaluation and design of filters using a Taylor series expansion. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):184–199.

NEHAB, D., MAXIMO, A., LIMA, R. S., and HOPPE, H. 2011. GPU-efficient recursive filtering and summed-area tables. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH Asia 2011)*, 30(6):176.

NELDER, J. A. and MEAD, R. 1965. A simplex method for function minimization. *The Computer Journal*, 7(4):308–313.

PARK, S. K. and SCHOWENGERDT, R. A. 1983. Image reconstruction by parametric cubic convolution. *Computer Vision, Graphics & Image Processing*, 23(3):258–272.

RUDERMAN, D. L. 1997. Origins of scaling in natural images. *Vision Research*, 37(23):3385–3398.

SCHAUM, A. 1993. Theory and design of local interpolators. *Computer Vision, Graphics & Image Processing*, 55(6):464–481.

SHANNON, C. E. 1949. Communication in the presence of noise. *Proceedings of the Institute of Radio Engineers*, 37(1):10–21.

SIGG, C. and HADWIGER, M. 2005. Fast third-order texture filtering. In M. Pharr, editor, *GPU Gems 2*, chapter 20, pages 313–329. Addison Wesley Professional.

STRANG, G. and FIX, G. 1971. A Fourier analysis of the finite element variational method. In G. Geymonat, editor, *Constructive Aspects of Functional Analysis*, volume 57 of *C.I.M.E. Summer Schools, 2011*, pages 793–840. Springer Berlin Heidelberg.

THÉVENAZ, P., BLU, T., and UNSER, M. 2000. Interpolation revisited. *IEEE Transactions on Medical Imaging*, 19(17):739–758.

UNSER, M. 1996. Approximation power of biorthogonal wavelet expansions. *IEEE Transactions on Signal Processing*, 44(3): 519–527.

UNSER, M. 1999. Splines: A perfect fit for signal and image processing. *IEEE Signal Processing Magazine*, 16(6):22–38.

UNSER, M. 2000. Sampling—50 years after Shannon. *Proceedings of the IEEE*, 88(4):569–587.

UNSER, M., ALDROUBI, A., and EDEN, M. 1991. Fast B-spline transforms for continuous image representation and interpolation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(3):277–285.

UNSER, M., ALDROUBI, A., and EDEN, M. 1993. B-spline signal processing: Part II—efficient design and applications. *IEEE Transactions on Signal Processing*, 41(2):834–848.

WANG, Z., BOVIK, A., SHEIKH, H., and SIMONCELLI, E. 2004. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612.

ZHANG, X. and WANDELL, B. A. 1996. A spatial extension to CIELAB for digital color image reproduction. In *Society for Information Display Symposium Technical Digest*, volume 27, pages 731–734.