Optimized Quasi-Interpolators for Image Reconstruction

Leonardo Sacht Diego Nehab

***Abstract*—We propose new quasi-interpolators for the continuous reconstruction of sampled images, combining a narrowly-supported piecewise-polynomial kernel and an efficient digital filter. In other words, our quasi-interpolators fit within the generalized sampling framework and are straightforward to use. We go against standard practice and optimize for approximation quality over the entire Nyquist range, rather than focusing exclusively on the asymptotic behavior as the sample spacing goes to zero. In contrast to previous work, we jointly optimize with respect to all degrees of freedom available in both the kernel and the digital filter. We consider linear, quadratic, and cubic schemes, offering different trade-offs between quality and computational cost. Experiments with compounded rotations and translations over a range of input images confirm that, due to the additional degrees of freedom and the more realistic objective function, our new quasi-interpolators perform better than the state-of-the-art, at a similar computational cost.**

***Index Terms*—image reconstruction, quasi-interpolation.**

I. Introduction

THE problem of obtaining an estimate for the value of a function at an arbitrary point, given only a discrete set of sampled values, has a long history in applied mathematics [23]. A variety of operations commonly performed on images, such as rotations, translations, warps, and resolution change, require this kind of resampling. Efficient, high-quality reconstruction is therefore of fundamental importance in computer graphics and image processing applications.

In this paper, we leverage results from the intersection of image processing and approximation theory to optimize for a new family of reconstruction schemes. As an example, figure 1 shows a standard benchmark used to evaluate reconstruction quality. An input image is repeatedly translated by subpixel offsets, thereby accumulating the errors due to each reconstruction step. The figure compares the two best-performing quadratic reconstruction schemes with the result of our method. Visual inspection suggests our method is superior at preserving high-frequency content. This impression is confirmed quantitatively by the perceptual SSIM [41] metric as well as the PSNR metric.

Such benchmarks are popular for two reasons: they magnify small differences between competing methods and the input images themselves serve as ground-truth. Nevertheless, the ability to preserve image quality under so many repeated operations can be useful, for example, in real-time reprojection strategies, where attempts to explore spatio-temporal coherence and amortize computations across multiple frames suffer with precisely this type of degradation [33].

Modern sampling and reconstruction strategies evolved from the ideal sampling theory of Shannon-Whittaker [34]. Given a uniform sample spacing T, ideal sampling *prefilters* an input function f to obtain an approximation g that is bandlimited to the *Nyquist interval* $(\text{-}\frac{0.5}{T}, \frac{0.5}{T})$:

$$g(x) = \int_{\text{-}\infty}^{\infty} f(t)\,\operatorname{sinc}\big((t-x)/T\big)\,\mathrm{d}t. \tag{1}$$

Here, $\operatorname{sinc}(x) = \sin(\pi x)/(\pi x)$ is the ideal low-pass filter. The values of $g(x)$ can then be reconstructed exactly from uniformly spaced samples $g(kT)$ of g, $k \in \mathbf{Z}$:

$$g(x) = \sum_{k\in\mathbf{Z}} g(kT)\operatorname{sinc}(x/T-k). \tag{2}$$

If f is itself already bandlimited to $(\text{-}\frac{0.5}{T}, \frac{0.5}{T})$, then prefiltering becomes both unnecessary and harmless.

Despite its importance in countless domains, ideal sampling suffers from a variety of shortcomings when applied to image processing. Typical input functions are not bandlimited, and their bandlimited approximations are often perceptually unpleasant (i.e., they *ring* too much). Finally, the infinite support and slow decay of the ideal low-pass filter make evaluating (1) and (2) impractical or inefficient.

The vast majority of sampling and reconstruction strategies, which were designed to address these shortcomings, fit within the *generalized sampling* framework [38]. A recent survey by Nehab and Hoppe [26] presents all relevant background material. Generalized sampling replaces the ideal low-pass filter in (1) by an arbitrary *analysis filter* ψ (a.k.a. *antialiasing filter*, or simply *prefilter*) and in (2) by an arbitrary *generating function* φ (a.k.a. *reconstruction kernel*). Notably, a *digital filter* $\boldsymbol{q}$ (i.e., a discrete convolution) is inserted between sampling and reconstruction to add crucial flexibility:

$$g(x) = \int_{\text{-}\infty}^{\infty} f(t)\,\psi\big((t-x)/T\big)\,\mathrm{d}t, \tag{3}$$

$$g_k = g(kT), \qquad c_i = \sum_{k\in\mathbf{Z}} g_k\, q_{i\text{-}k}, \tag{4}$$

$$\tilde{f}(x) = \sum_{i\in\mathbf{Z}} c_i\,\varphi(x/T-i). \tag{5}$$

We are interested in the situation where the input consists of samples g_k resulting from an image creation process that gives us no control over ψ. Without any attempt at inferring ψ, our goal is to obtain combinations of φ and $\boldsymbol{q}$ that enable efficient, high-quality reconstructions.

L. Sacht is with the Department of Mathematics, UFSC, Brazil.
D. Nehab is with IMPA, Brazil.
Manuscript received MMM DD, YYYY; revised MMM DD, YYYY.

(a) Input image and detail (b) Cardinal Quadratic O-MOMS [4] (c) Quadratic Condat [9] (d) Our Quadratic

Fig. 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 compute the error due to the repeated resampling operations. 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.

Our contributions are as follows:

- We provide intuition and examples to show that, contrary to widespread belief, the asymptotic behavior of approximation does *not* reliably predict image reconstruction quality;
- We propose a weaker image prior and alternative objective function that form a better model for reconstruction quality;
- To the best of our knowledge, our work is the first to *jointly* optimize for the degrees of freedom available in both the generator function φ and digital filter $\boldsymbol{q}$;
- We demonstrate improved reconstruction quality with novel linear, quadratic, and cubic generators, each paired with its own optimized digital filter.

II. Notation

When dealing with generalized sampling algorithms, it is convenient to adopt a new notation that allows us to merge equations (3)–(5) into index-free, variable-free expressions that are shorter and easier to manipulate.

Let italicized lower-case symbols (e.g., $f : \mathbf{R} \to \mathbf{C}$) denote one-variable scalar functions, and let bold lower-case symbols (e.g., $\boldsymbol{c} : \mathbf{Z} \to \mathbf{C}$) denote discrete sequences.

Let function scaling and reflection be respectively denoted by

$$f_T(x) = f(x/T) \quad \text{and} \quad f^\vee(x) = f(-x). \tag{6}$$

Define continuous, discrete, and *mixed* convolutions as

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

$$(\boldsymbol{a} * \boldsymbol{b})_n = \sum_{k\in\mathbf{Z}} a_k\, b_{n-k}, \quad \text{and} \tag{8}$$

$$(\boldsymbol{a} *_T f)(x) = \sum_{i\in\mathbf{Z}} a_i\, f(x - iT). \tag{9}$$

Note that the discrete convolution inverse of $\boldsymbol{q}$, when it exists, is another sequence (perhaps with infinite support), denoted by $\boldsymbol{q}^{-1}$ and satisfying

$$\boldsymbol{q} * \boldsymbol{q}^{-1} = \boldsymbol{\delta} = [\ldots, 0, 0, 1, 0, 0, \ldots]. \tag{10}$$

Finally, define the operation that uniformly samples a function f into a discrete sequence, with sample spacing T (assume $T = 1$ when omitted)

$$[\![f]\!]_T = \big[\ldots, f(-T),\ f(0),\ f(T), \ldots\big] \tag{11}$$

We can now write equations (3) through (5) as concisely as

$$\tilde{f} = [\![f * \psi_T^\vee]\!]_T * \boldsymbol{q} *_T \varphi_T. \tag{12}$$

The advantages of this notation are discussed in detail by Nehab and Hoppe [26]. The remaining notation used throughout the paper is standard, or otherwise defined where needed.

III. Related Work

The design of reconstruction strategies can be seen as a process in which the degrees of freedom in φ and $\boldsymbol{q}$ are eliminated according to the requirements of a target application. Different formulations for desirable properties have led to a multitude of different schemes.

For efficient random-access to each reconstructed value, the number of summation terms in (5) is limited by restricting φ to functions that vanish outside a compact support with narrow width W. Furthermore, the definition of φ is typically divided into W polynomial pieces[1], each with maximum degree N. This leads to $W \times (N+1)$ degrees of freedom in φ.

Certain applications require continuous or differentiable reconstructions. This *regularity* constraint ($\tilde{f} \in C^R$) affects only the generator: it can be traded off for degrees of freedom in φ. A symmetry restriction $\varphi = \varphi^\vee$ is also typically imposed on the generator.

It is often desired that the reconstruction *interpolates* all input samples provided:

$$\tilde{f}(kT) = g_k = g(kT) \tag{13}$$

When reconstructing directly from samples g_k, the interpolation requirement leads to equations:

$$g_k = \sum_{i\in\mathbf{Z}} g_i\, \varphi(k-i), \quad \forall k \in \mathbf{Z}. \tag{14}$$

The solution therefore imposes $\varphi(0) = 1$ and $\varphi(i) = 0$, for all $i \in \mathbf{Z} \setminus \{0\}$.

[1]Windowed-sinc approximations, such as Hamming [14] and Lanczos [13], offer an inferior trade-off between reconstruction quality and speed compared to piecewise polynomial kernels [24].

A. The digital filtering stage

The digital filtering stage first appeared when Hou and Andrews [15] proposed interpolation with the cubic B-spline kernel β^3. Such kernels do not respect the interpolation conditions above. Rather than following (14), they proposed reconstructing from coefficients c_k to be found during a preprocessing stage to satisfy the equations

$$g_k = \sum_{i \in \mathbf{Z}} c_i \, \beta^3(k-i) = \tfrac{1}{6} c_{k\text{-}1} + \tfrac{4}{6} c_k + \tfrac{1}{6} c_{k+1}, \ \forall k \in \mathbf{Z}. \quad (15)$$

This Toeplitz, symmetric, tridiagonal linear system can also be understood from the signal-processing perspective. The sample sequence $\boldsymbol{g}$ results from the discrete convolution between the unknown coefficient sequence $\boldsymbol{c}$ and the sampled cubic B-spline kernel:

$$\boldsymbol{g} = \boldsymbol{c} * \boldsymbol{[}\beta^3\boldsymbol{]}, \quad \text{where} \quad \boldsymbol{[}\beta^3\boldsymbol{]} = \begin{bmatrix} \frac{1}{6} & \frac{4}{6} & \frac{1}{6} \end{bmatrix} \quad (16)$$

The inverse operation is again a convolution, much like the inverse of a matrix product is again a matrix product. It can therefore be written in the form of the digital filtering stage of equation (4):

$$\boldsymbol{c} = \boldsymbol{g} * \boldsymbol{q}, \quad \text{where} \quad \boldsymbol{q} = \boldsymbol{[}\beta^3\boldsymbol{]}^{-1}. \quad (17)$$

In practice, the convolution inverse is instead implemented by factoring its Z-transform into two or more recursive filtering passes [40], or by solving the linear system using its precomputed LU-decomposition [21]. The computations are essentially the same: about 5 floating-point operations per sample in this example.

Conceptually, it is also possible to perform interpolation directly, by using the *cardinal* generating function φ_{int}

$$g_k = \sum_{i \in \mathbf{Z}} g_k \, \varphi_{\text{int}}(k-i), \quad \text{where} \quad \varphi_{\text{int}} = \varphi * \boldsymbol{[}\varphi\boldsymbol{]}^{-1}. \quad (18)$$

Note however that φ_{int} has infinite support, so that the indirect formulation using c_k and φ is more practical and efficient.

The digital filtering stage plays a key role in another fundamental problem. The set of all functions in the form of equation (5) is a linear subspace of L_2 (parametrized by φ and T):

$$V_{\varphi,T} = \{ \boldsymbol{c} *_T \varphi_T \mid \forall \boldsymbol{c} \in \ell_2 \}. \quad (19)$$

It is natural to look for the function in $V_{\varphi,T}$ that is closest to the input f, i.e., its orthogonal projection $P_{\varphi,T} f$. It is:

$$P_{\varphi,T} f = \boldsymbol{c}^* *_T \varphi_T, \quad \text{where} \quad (20)$$

$$\boldsymbol{c}^* = \arg\min_{\boldsymbol{c}} \| \boldsymbol{c} *_T \varphi_T - f \|_{L_2}. \quad (21)$$

This linear optimization problem was studied within computer graphics by Kajiya and Ullner [18]. It involves the Gramian matrix associated to shifts of the generator φ, which is again a Toeplitz banded matrix. The linear system can be written as the convolution

$$\boldsymbol{[}f * \varphi_T^\vee\boldsymbol{]}_T = \boldsymbol{c}^* * \boldsymbol{[}a_\varphi\boldsymbol{]}, \quad (22)$$

where $a_\varphi = \varphi * \varphi^\vee$ is the auto-correlation of the generator φ. The solution uniquely determines both the prefilter and the digital filter

$$\psi = \varphi \quad \text{and} \quad \boldsymbol{q} = \boldsymbol{[}a_\varphi\boldsymbol{]}^{-1}. \quad (23)$$

In summary, the digital filtering stage is used in preprocessing to reconstruction or post-processing to prefiltering. It is most often a combination between narrowly supported convolutions and their inverses, and brings great flexibility to the sampling pipeline at little performance cost. Implementation is trivial.

The remaining degrees of freedom in φ and $\boldsymbol{q}$ are chosen to maximize some notion of "reconstruction quality". The most prevalent notion is, by far, the *approximation order*.

B. Approximation order

Intuitively, the approximation order L measures the rate T^L at which the residual between the input f and its approximation $\tilde{f}$ vanishes as the sample spacing T tends to 0 (i.e., as the sampling rate is progressively increased).

We say that the generator φ has approximation order L when L is the largest positive integer for which there exist constants $C > 0$ such that for all $f \in \mathbf{W}_2^L$ (the Sobolev space):

$$\| f - P_{\varphi,T} f \|_{L_2} \leqslant C \cdot T^L \cdot \left\| f^{(L)} \right\|_{L_2}. \quad (24)$$

The existence and uniqueness of $P_{\varphi,T} f$ must be guaranteed. This is a mild condition on φ, easily verified in the frequency domain [1] by any of the following: $V_{\varphi,T}$ should be a closed subspace of L_2, or the shifts of φ should form Riesz sequence, or the sampled auto-correlation $\boldsymbol{[}a_\varphi\boldsymbol{]}$ should be invertible when seen as a discrete convolution operator.

In that case, we can define the *dual* $\mathring{\varphi}$ of the generator φ:

$$\mathring{\varphi} = \varphi * \boldsymbol{[}a_\varphi\boldsymbol{]}^{-1}, \quad (25)$$

so that the orthogonal projection can be expressed directly as

$$P_{\varphi,T} f = \boldsymbol{[}f * \mathring{\varphi}_T^\vee\boldsymbol{]}_T *_T * \varphi_T. \quad (26)$$

A famous result by Strang and Fix [35] shows that φ has approximation order L if and only if $V_{\varphi,T}$ contains all polynomials up to degree $L-1$. The same work shows that interpolating the sampled f (i.e., *without* prefiltering) has the same approximation order:

$$\left\| f - \boldsymbol{[}f\boldsymbol{]}_T *_T (\varphi_{\text{int}})_T \right\|_{L_2} \leqslant C_{\text{int}} \cdot T^L \cdot \left\| f^{(L)} \right\|_{L_2}, \quad (27)$$

though naturally $C_{\text{int}} \geqslant C$. Unser [37] generalized this further by showing that it is *sufficient* for ψ and $\boldsymbol{q} * \varphi$ to be biorthonormal:

$$\boldsymbol{[}\varphi * \psi^\vee\boldsymbol{]} * \boldsymbol{q} = \boldsymbol{\delta}. \quad (28)$$

It is easy to demonstrate the biorthonormality property for orthogonal projection, interpolation, and the *consistent sampling* of Unser and Aldroubi [39] (a.k.a. *oblique projection*). This approximation order guarantee is reassuring because, except for synthetic images [17, 18, 22], we are rarely able to perform orthogonal projection.

The *necessary* condition for approximation order L is that polynomials of degree up to $L-1$ must be preserved throughout the entire sampling pipeline [8, 11]:

$$\tilde{f}_T = \big[f * \psi_T^\vee\big]_T * \boldsymbol{q} *_T \varphi_T = f, \quad \forall f \in P_{L-1}. \tag{29}$$

More often than not, it is more convenient to express the residual between the input and its reconstruction using the frequency domain:

$$\|f - \tilde{f}\|_{L_2}^2 \approx \int_{-\infty}^{\infty} \left|\hat{f}(\omega)\right|^2 E(T\omega)\,\mathrm{d}\omega. \tag{30}$$

An expression for the *error kernel* $E(T\omega)$ was first obtained by Park and Schowengerdt [28], assuming no prefilter or digital filter, and later refined by Schaum [32]. The expression for arbitrary ψ, $\boldsymbol{q}$, and φ, was obtained by Blu and Unser [3], where it was shown that approximation order L is equivalent to all derivatives of E up to order $2L-1$ vanishing at 0.

C. Previous reconstruction strategies

The increasing influence of approximation theory in the development of reconstruction strategies is best illustrated by the progress in cubic reconstruction ($N = 3$). Historically, interpolation and C^1-continuity were considered vital, and the digital filtering stage was absent. The focus was on a family of cubics defined by support $W = 4$ and regularity $R = 1$, and by the interpolation requirement $[\varphi] = \boldsymbol{\delta}$. The single remaining degree of freedom, measuring the slope $\alpha = \varphi'(1)$, was set somewhat arbitrarily to match the slope of the ideal low-pass filter $\mathrm{sinc}'(1) = -1$ [30]. This choice leads to approximation order $L = 1$.

Keys [19] and Park and Schowengerdt [28] independently argued that, in order to achieve $L = 3$ under the same constraints, the value for α should be -0.5 instead. Indeed, this choice led to better results in practice [29]. Curiously, the same cubic interpolator had been identified earlier by Catmull and Rom [7] as appropriate for computer graphics applications, with no recourse to approximation theoretical arguments.

In designing their cubic, Mitchell and Netravali [25] abandoned interpolation and required approximation order $L = 2$. This left only one degree of freedom remaining. This which was selected for the best perceptual upsampling quality, as judged by a small user study.

Keys [19] proposed an increased support $W = 6$ to achieve the maximal approximation order $L = 4$ for cubics, but this never became popular, perhaps due to the additional computational cost. Instead of increasing support, Schaum [32] abandoned the regularity requirement, reaching a family of maximal-order local Lagrangian interpolators that includes a cubic ($W = 4$, $L = 4$, but $R = 0$).

Blu et al. [5] found a complete parametrization of piecewise-polynomial generating functions given N, W, R, and L. Blu et al. [4] focused on the desirable subset with maximal order and minimum support (MOMS). Their O-MOMS family results from the minimization of the asymptotic constant C in (24) under the assumption of orthogonal projection. Requiring maximal regularity consumes all degrees of freedom and leads to the B-spline family (proposed earlier by Hou and Andrews [15] with no approximation-theoretical motivation). Requiring interpolation without a digital filter also consumes all degrees of freedom, and results in the same local Lagrangian interpolators of Schaum [32].

Recent *quasi-interpolation* schemes use the B-spline family as generators, and employ the digital filtering stage to mimic the effect of orthogonal projection [2, 9, 10]. Although such schemes do not interpolate arbitrary input functions, they guarantee optimal approximation order (i.e., they interpolate *polynomials*), and differ only on the generality of the digital filters considered.

In summary, the development of reconstructions strategies has focused on approximation order and its leading constant, moved the interpolation condition to the digital filtering stage, abandoned regularity, and finally abandoned interpolation itself. We take the final step, by abandoning approximation order and focusing directly on reconstruction quality.

IV. Theory and Motivation

We base our optimization problem on the following theorem due to Blu and Unser [3]. It quantifies the L_2-error between the input function f and the output function given by (29):

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

$$\|f - \tilde{f}_T\|_{L_2} = \left(\int_{-\infty}^{\infty} |\hat{f}(\omega)|^2 E(T\omega)\,\mathrm{d}\omega\right)^{\frac{1}{2}} + e(f,T), \tag{31}$$

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

$$E(\omega) = 1 - \frac{\left|\hat{\varphi}(\omega)\right|^2}{\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)} + \widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)\left|\hat{\boldsymbol{q}}(\omega)\,\hat{\psi}(\omega) - \frac{\hat{\varphi}(\omega)}{\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)}\right|^2. \tag{32}$$

Proof: See appendix C in [3].

The residual term $e(f,T)$ vanishes in many situations such as in the case where f is band-limited in the Nyquist interval [2]. Setting this term aside, formula (31) 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 E and all its derivatives up to degree $2L-1$ vanishing at zero [2]. In turn, this causes the error kernel to behave as a polynomial with dominant power ω^{2L} near $\omega = 0$. This is the intuition that motivated the push towards schemes with high approximation order.

When two different schemes have the same approximation order, the same intuition favors the one with the smallest asymptotic constant C appearing in (24) and (27) [4, 6, 36]. This can be related to the error kernel (32), since this constant is proportional to the coefficient of the $(2L)^{\text{th}}$ power of the polynomial approximation of the error kernel around $\omega = 0$ [2].

[2] $f \in \mathbf{W}_2^r$ whenever $\int_{-\infty}^{\infty}(1+\omega^2)^r\left|\hat{f}(\omega)\right|^2\mathrm{d}\omega < \infty$.

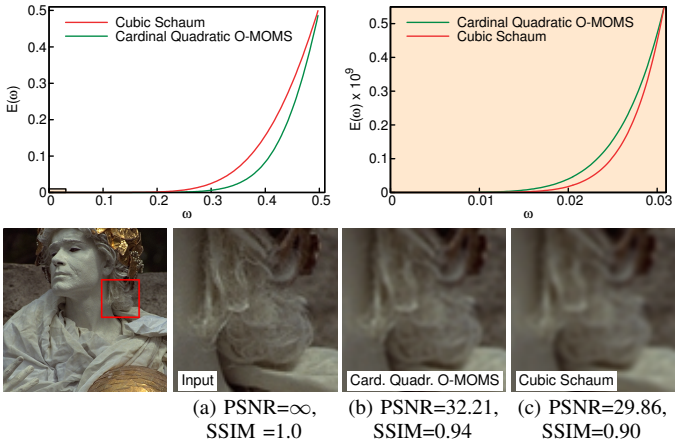

(a) PSNR=∞, SSIM =1.0 (b) PSNR=32.21, SSIM=0.94 (c) PSNR=29.86, SSIM=0.90

Fig. 2. Comparison between the quadratic O-MOMS, a 3rd-order interpolator proposed by Blu et al. [4], and a 4th-order cubic by Schaum [32]. 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.

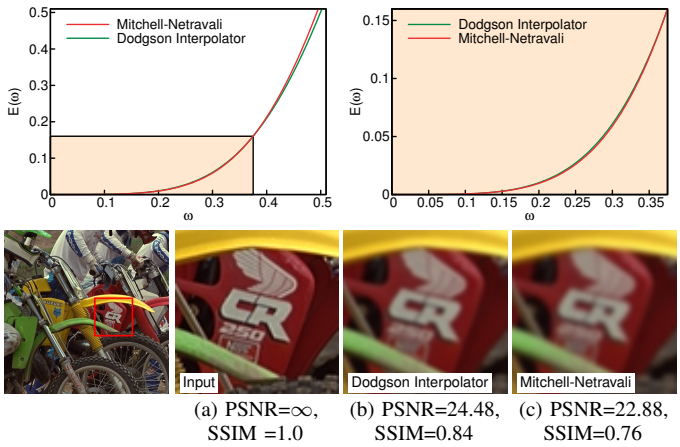

(a) PSNR=∞, SSIM =1.0 (b) PSNR=24.48, SSIM=0.84 (c) PSNR=22.88, SSIM=0.76

Fig. 3. Comparison between a quadratic interpolator proposed by Dodgson [12] and the cubic by Mitchell and Netravali [25] (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.

We agree that a higher approximation order and small asymptotic constant can be important in many applications. However, as we show in figures 2 and 3, 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 these recommendations.

Figure 2 compares the 3rd-order (cardinal) quadratic interpolator O-MOMS [4] with the 4th-order cubic local Lagrangian interpolator [32] in an experiment that consists of 30 compounded rotations. At each rotation step, the input image is interpolated, and sampled at a $\frac{360^o}{30} = 12^o$ rotation. The result is used as input for the next rotation step, and so on, until the image is back to its initial position. At this point it is compared to the original input. The PSNR and SSIM measures are higher (meaning higher quality) for the result with O-MOMS (Figure 2b), even though it has a lower approximation order. The plot of the error kernels for both approximation schemes (Figure 2 top left) show that quadratic O-MOMS has a smaller value overall in the full Nyquist interval although it is worse for low frequencies (Figure 2 top right), the latter behavior being expected since it has lower approximation order.

In figure 3, we compare the performance of the quadratic interpolator designed by Dodgson [12] with the (non-interpolating) cubic proposed by Mitchell and Netravali [25]. Both these kernels have approximation order 2, so we would expect the one with smaller asymptotic constant to perform better (The formula for obtaining the constant is available in [3].) 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 3 show that Dodgson's interpolator generates a better result (Figure 3b). This is again due to a better behaviour in the full Nyquist interval (Figure 3 top left), despite slightly worse low-frequency behavior (Figure 3 top right).

These counter-examples exist because the benchmarks violate the underlying assumption that the input frequency content is concentrated around $T\omega \to 0$. Indeed, the input power spectrum for natural images tends to behave as

$$\left|\hat{f}(\omega)\right|^2 \approx \frac{1}{\omega^p}, \tag{33}$$

where p varies from 1.6 to 3.0 [16, 31]. A photograph taken underwater tends to be blurrier, so p will be larger. A photograph of foliage contains more high-frequency content, so p will be smaller. 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 only have access to the samples of f. If f was subjected to a good prefilter prior to sampling, the frequency content outside the Nyquist interval is close to zero. If not, whatever frequencies survived 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 (31), we focus on the Nyquist interval:

$$\|f - \tilde{f}_T\|_{L_2}^2 \approx \int_{-\frac{0.5}{T}}^{\frac{0.5}{T}} \left|\hat{f}(\omega)\right|^2 E(T\omega)\,\mathrm{d}\omega. \tag{34}$$

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

We can now define our objective function for minimization:

$$\min \int_{-0.5}^{0.5} \left|\hat{f}(\omega)\right|^2 E(\omega)\,\mathrm{d}\omega. \tag{35}$$

Assuming $\hat{f}$ known (to be detailed in section V) and previously prefiltered ($\psi = \delta$ 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,

$$\text{FIR}: \boldsymbol{q} = [\ldots, 0, d_{-j}, \ldots, d_0, \ldots, d_j, 0, \ldots], \tag{36}$$

$$\text{IFIR}: \boldsymbol{q} = [\ldots, 0, e_{-k}, \ldots, e_0, \ldots, e_k, 0, \ldots]^{-1}, \text{ and} \tag{37}$$

$$\text{FIR-IFIR}: \boldsymbol{q} = \boldsymbol{d} * \boldsymbol{e}. \tag{38}$$

These formulations contain $2j+1$, $2k+1$, 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. [5] 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 $W \geqslant N$, $\varphi \in \{N, W, R, L\}$ if and only if there exists a unique set of coefficients $a_{k,\ell}$, $b_{k,\ell}$, and $c_{k,\ell}$ such that

$$\begin{aligned}\varphi\big(x-\tfrac{W}{2}\big) = &\sum_{\ell=1}^{M}\sum_{k=0}^{N\text{-}L\text{-}\ell} a_{k,\ell}\,\big(\beta_{nc}^{L+k\text{-}1} * \gamma_{\ell}^{N\text{-}L\text{-}k}\big)(x)\\ &+\sum_{\ell=0}^{M}\sum_{k=0}^{W\text{-}N+\ell\text{-}1} b_{k,\ell}\,\beta_{nc}^{N\text{-}\ell}(x-k)\\ &+\sum_{k=0}^{W\text{-}L}\sum_{\ell=0}^{L\text{-}R\text{-}2} c_{k,\ell}\,\Delta^{*\ell}\,\beta_{nc}^{L\text{-}\ell\text{-}1}(x-k),\end{aligned} \tag{39}$$

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

Proof: See [5].

In the formulas above,

$$\beta_{nc}^{n}(x) = \beta^{n}\big(x-\tfrac{n+1}{2}\big), \tag{40}$$

is the non-centered B-spline, $\Delta^{*\ell}$ is the ℓ^{th}-order finite difference, and γ_{ℓ}^{n} are linear combinations of finite difference operators applied to the polynomial simple elements defined in [5]. For example, setting $N=1$, $W=2$, $R=-1$ (meaning φ is bounded), and $L=1$ in the theorem produces

$$\varphi(x) = b_{0,0}\beta^{1}(x) + c_{0,0}\beta^{0}\big(x+\tfrac{1}{2}\big) + c_{1,0}\beta^{0}\big(x-\tfrac{1}{2}\big) \tag{41}$$

This gives us 3 additional degrees of freedom, relative to the common choice of $\varphi = \beta^{1}(x)$ [9, 10], with which we minimize our objective function.

V. Optimization

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

a) 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.

b) 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 [4, 32]. Our results confirm this. Applications requiring more regularity (e.g., for derivatives) can change this parameter in the optimization.

c) Approximation order of φ: In stark contrast to previous work, we only require first-order approximation ($L=1$). This means that frequency $\omega = 0$ will be preserved. As with the regularity constraint, these additional degrees of freedom are better left to the discretion of the optimizer.

These constraints determine the coefficients in (39) that are available for minimization. We encapsulate them into lists of coefficients $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{C}$:

$$\mathbf{A} = \{a_{k,\ell}\}, \quad \mathbf{B} = \{b_{k,\ell}\}, \quad \mathbf{C} = \{c_{k,\ell}\}. \tag{42}$$

d) Symmetry of φ and $\boldsymbol{q}$: To ensure that our quasi-interpolators have linear phase response, we require $\boldsymbol{q} * \varphi$ to be symmetric. This imposes linear relationships between the coefficients $a_{k,\ell}$, $b_{k,\ell}$, and $c_{k,\ell}$, and sets $\boldsymbol{d}_i = \boldsymbol{d}_{\text{-}i}$, and $\boldsymbol{e}_i = \boldsymbol{e}_{\text{-}i}$, for all i.

e) 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

$$\int_{\text{-}\infty}^{\infty} \varphi(x)\,\mathrm{d}x = 1, \quad \text{and} \quad \sum_{i\in\mathbf{Z}} \boldsymbol{d}_i = \sum_{i\in\mathbf{Z}} \boldsymbol{e}_i = 1. \tag{43}$$

f) 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

$$E(0) = 0. \tag{44}$$

See [2] for the proof.

g) Objective function: Recall the spectrum of natural images tends to follow (33). Since we seek input-independent quasi-interpolators, we set p to the intermediate value of $p=2$:

$$\big|\hat{f}(\omega)\big|^2 \approx \frac{1}{\omega^2}. \tag{45}$$

This choice has an additional 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 (35) to converge to a finite value at the origin.

h) The optimization problem: Given a degree N:

$$\begin{aligned}\underset{\boldsymbol{q},\mathbf{A},\mathbf{B},\mathbf{C}}{\arg\min} \quad & F(d) := \int_0^d \frac{1}{\omega^2} E(\omega)\,\mathrm{d}\omega && (46)\\ \text{subject to} \quad & \varphi \in \{N, N+1, -1, 1\}, && (47)\\ & \varphi^{\vee} = \varphi, \quad \boldsymbol{q}^{\vee} = \boldsymbol{q}, && (48)\\ & \int_{\text{-}\infty}^{\infty} \varphi(x)\,\mathrm{d}x = 1, \quad \sum_{k\in\mathbf{Z}} \boldsymbol{q}_k = 1, && (49)\\ & E(0) = 0. && (50)\end{aligned}$$

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

i) Controlling overshoot and aliasing: The natural choice for the integration limit d in (46) 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 4c.

By minimizing (46) 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

$$\begin{aligned}\hat{\varphi}_{\text{qi}}(0.5-\varepsilon) \approx 1, \quad \hat{\varphi}_{\text{qi}}(0.5+\varepsilon) \approx 0,\\ \hat{\varphi}_{\text{qi}}(\text{-}0.5+\varepsilon) \approx 1, \quad \hat{\varphi}_{\text{qi}}(\text{-}0.5-\varepsilon) \approx 0.\end{aligned} \tag{51}$$

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 [5] for the expression), $\hat{\boldsymbol{q}}(\omega)$ must approximate the discontinuities near $\omega = \pm 0.5$.

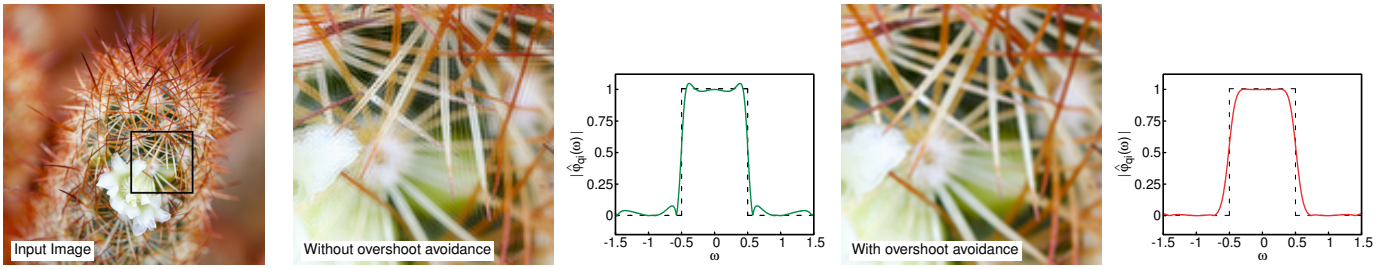

(a) PSNR = ∞, SSIM = 1.0 (b) PSNR = 34.662, SSIM = 0.995 (c) Frequency response (d) PSNR = 43.812, SSIM = 0.999 (e) Frequency response

Fig. 4. 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 the value $d = 0.34$ was automatically obtained by a binary search. Plots in figures (c) and (e) show the frequency responses of the associated quasi-interpolators.

Given the filter has a finite support, its DTFT is a finite sum of sinusoids (or the reciprocal of it), and this leads to the Gibbs phenomenon in $\hat{\boldsymbol{q}}(\omega)$. It is modulated by $\hat{\varphi}(\omega)$ and manifests itself as ringing in the reconstructed images (figure 4b).

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

$$\hat{\varphi}_{\mathrm{qi}}(\omega) \leqslant 1.0025, \quad \forall \omega \in [-0.5, 0.5] \quad \text{and} \quad (52)$$

$$|\hat{\varphi}_{\mathrm{qi}}(\omega)| \leqslant 0.025, \quad \forall \omega \in [-\infty, -0.75] \cup [0.75, \infty] \,. \quad (53)$$

Intuitively, condition (52) prevents overshoot and condition (53) 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 (46) that leads to an admissible quasi-interpolator.

Note that condition (53) excludes the interval $(0.5, 0.75)$. In fact, for $N = 1$, even this relaxed condition is too restrictive, so we test only for condition (52). Degrees $N = 2$ and 3 have larger parameter spaces, and we find values for d that satisfy both constraints.

The practical effect of the admissibility conditions can be seen in the example of figure 4. 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 overshooting is mostly gone.

j) Length and type of $\boldsymbol{q}$: We optimized (46)–(50) under FIR, IFIR and FIR-IFIR digital filter formulations. IFIR filters achieved higher quality at similar computational costs to FIR and FIR-IFIR filters. The wider $\boldsymbol{q}$ is (i.e., the more degrees of freedom it offers), the lower objective function values are obtained. However, little is gained for widths greater than 5. We therefore always use width 5 and an IFIR formulation for the digital filter.

VI. Results and Discussion

We implemented our numerical optimization framework within Mathematica. The optimization method by Nelder and Mead [27], which is suitable for constrained non-linear problems, was the most reliable for our objective function. The objective function is somewhat brittle, due to the integrand in (46) being unstable near the origin. We were careful to keep the error kernel in a simple algebraic 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 (46). To reduce the risk of finding poor local minima, we solve each optimization problem 40 times, with randomization, and select the best result. The optimization completes in a few minutes. The values for all resulting parameters for degrees 1, 2, and 3 are given in appendix A. For convenience, we also provide the source-code for the generators and the digital filter coefficients in the supplementary materials.

Figure 5 shows plots of our quadratic and cubic generators φ and their associated quasi-interpolators $\varphi_{\mathrm{qi}} = \boldsymbol{q} * \varphi$ (b). Like the local Lagrangian interpolators of Schaum [32] and the O-MOMS of Blu et al. [4], our quadratic is not continuous, and the cubic is not differentiable. These point discontinuities and irregularities are not detrimental to approximation quality. The figure also shows a comparison between the absolute value of the frequency response $\hat{\varphi}_{\mathrm{qi}}$ of our quasi-interpolators with the state-of-the art. Our interpolators are clearly the sharpest. Furthermore, the error kernel plots show that our quasi-interpolators have lower error overall in the Nyquist interval. Please note that the plots of the previous state-of-the-art are closer to each other than to our new quasi-interpolators.

We compare the quality of our reconstruction schemes against previous quasi-interpolators by performing a variety of practical experiments.

Figure 6 shows the results of our quantitative experiments on two image benchmarks. The first benchmark consists of the 24 photographs in the Kodak [20] dataset. The second benchmark is composed of a selection of medical images of different sources (some of them are shown in the supplementary materials). The experiment consists of applying 75 compound randomized translations to the images. The randomization is added to rule out the possibility of errors being cancelled by negative correlations. The sequence of translations is such that, after every 3 translations, the net translation is zero. At these points, we can measure PSNR against the original input, and this is what the plots show. To obtain a single number, we average the PSNR results over all images in each dataset. Results show that our new cubic quasi-interpolator performs best, even when compared to quintic quasi-interpolators. Similarly, our new quadratic quasi-interpolator performs better than all other quadratics and cubics.

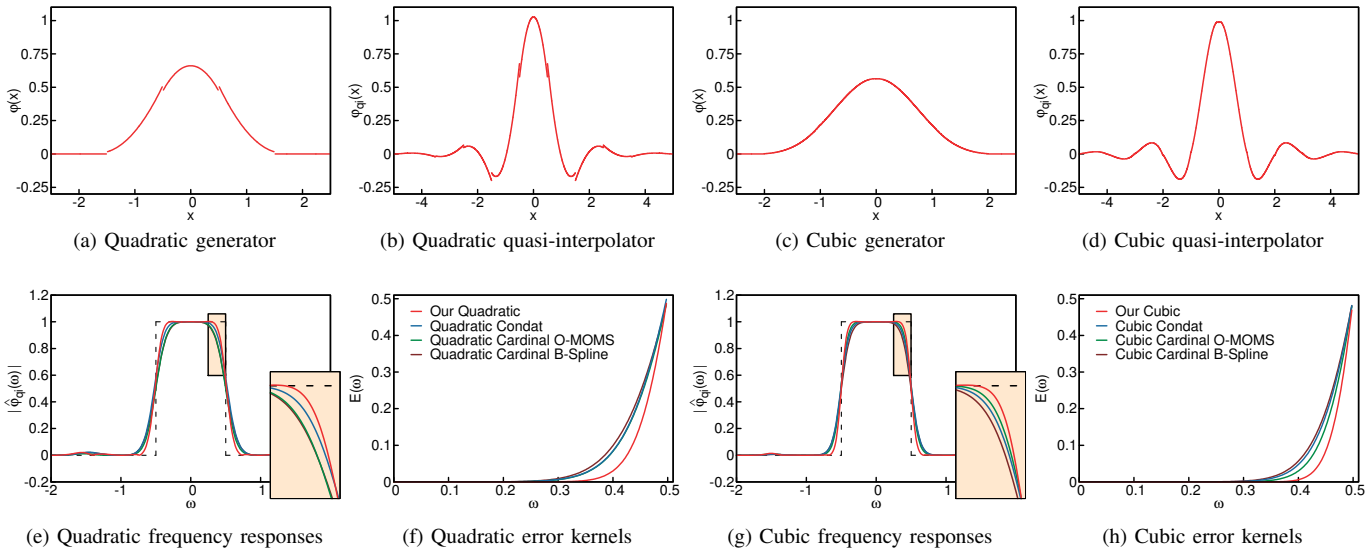

(a) Quadratic generator (b) Quadratic quasi-interpolator (c) Cubic generator (d) Cubic quasi-interpolator

(e) Quadratic frequency responses (f) Quadratic error kernels (g) Cubic frequency responses (h) Cubic error kernels

Fig. 5. Impulse responses of our quadratic and cubic generators (a,c) and the corresponding quasi-interpolators (b,d). Comparison of the frequency responses the quasi-interpolators against the best quadratics and cubics (e,g) show ours to be the sharpest. The associated error kernels (f,h) are the lowest over most of the Nyquist interval.

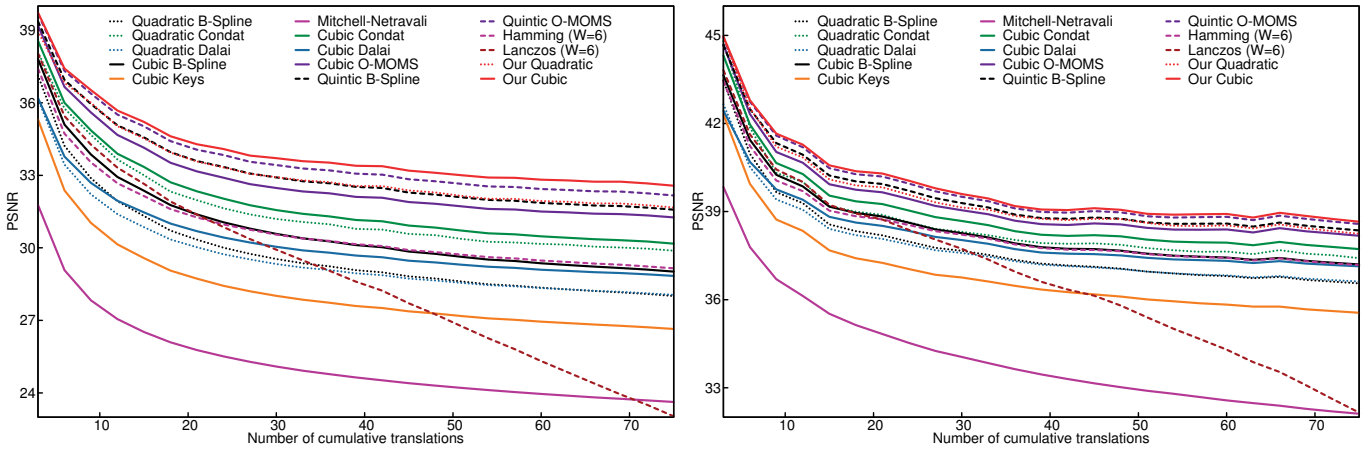

Fig. 6. Average PSNR of applying 75 randomized translations to (left) the Kodak dataset and (right) a set of medical images. The best and most popular quasi-interpolators in the literature are compared. Our new cubic quasi-interpolator (solid red) reaches the best quality, better even than the *quintic* O-MOMS (dashed purple). Our new quadratic (dotted red) reaches quality superior to any other quadratic or cubic.

Figure 7 shows visual results for our linear quasi-interpolator. Our results are significantly sharper than those obtained by the linear state-of-the art [9]. In fact, our results compare favourably even against the cardinal *quadratic* B-spline. Comparisons with other quasi-interpolators for this setting are provided in the supplementary materials.

The example in figure 8 shows the improvement of our quadratic quasi-interpolator relative to the quadratic by Condat et al. [9]. Our quadratic compares favourably even against the state-of-the-art *cubic* O-MOMS interpolator [36]. Additional comparisons are shown in the supplementary materials.

Figure 9 tests the quality 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 [2] (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. Please see more comparisons in the supplementary materials. 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.

The slight disagreement between PSRN and SSIM in figure 9 (c,d) (even clearer in figure 6 (d,e) in the supplementary materials) gives further evidence that the L_2 metric is not the best proxy for the perceptual differences detected by the human visual system. This is the reason why we must tweak our objective function. More work is required in this area.

We recommend the use of our solutions for degrees 2 and 3, given their superior performance and moderate computational cost (please see a performance comparison with previous methods in the supplementary materials). Tasks requiring even more speed can use the degree 1 solution, which we take as a proof-of-concept. Please see the supplementary materials for full resolution images of all these experiments. We also provide additional videos containing other interpolation sequences.

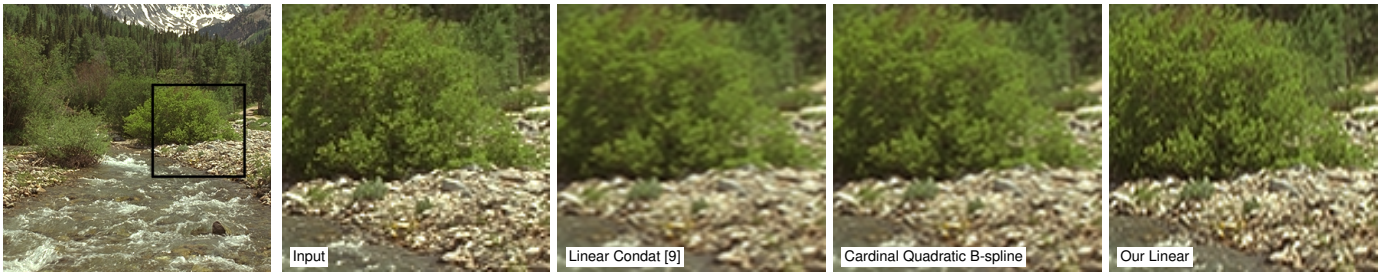

(a) PSNR = ∞, SSIM = 1.0 (b) PSNR = 25.873, SSIM = 0.862 (c) PSNR = 27.000, SSIM = 0.898 (d) PSNR = 28.584, SSIM = 0.942

Fig. 7. 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 [9] (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

Fig. 8. 30 (randomized) rotations for different quasi-interpolators. The quadratic proposed by Condat et al. [9] 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).

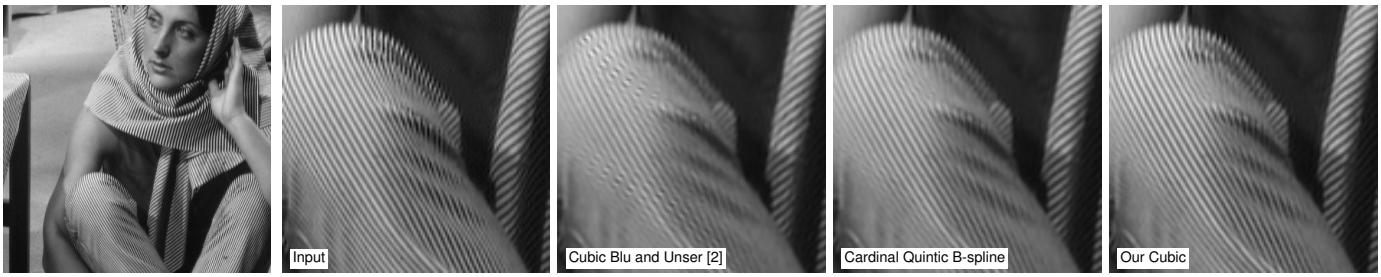

(a) PSNR = ∞, SSIM = 1.0 (b) PSNR = 13.013, SSIM = 0.749 (c) PSNR = 13.047, SSIM = 0.775 (d) PSNR = 13.044, SSIM = 0.784

Fig. 9. 40 compounded rotations. The result produced by one of the best cubic quasi-interpolators [2] has aliasing of high frequencies. Using the cardinal *quintic* B-spline leads to the same problem at a smaller magnitude. Our cubic almost completely removes the artefacts, while keeping the result sharp.

A. Limitations

One limitation of our method can be seen in figure 10, 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 our approach (c). The plots in figure 10 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.

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.

VII. 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.

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.

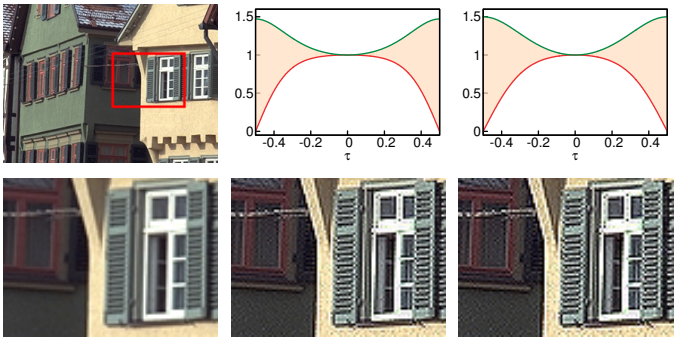

(a) Input (b) Our linear (c) Linear Condat [9]

Fig. 10. 4 repeated translations by exactly half pixel. Our linear quasi-interpolator (b) and the one proposed by Condat et al. [9] magnify frequencies too much. Shaded regions in the plots show the range of frequency amplitude scaling for each translation τ. The worst case is $\tau = 0.5$.

APPENDIX A
QUASI-INTERPOLATORS

a) Linear: $d = 0.5$.

$$b_{0,0} = 0.79076352 \qquad c_{0,0} = c_{1,0} = 0.10461824,$$
$$e_0 = 0.77412669, \qquad e_1 = e_{-1} = 0.11566267,$$
$$e_2 = e_{-2} = -0.00272602$$

b) Quadratic: $d \approx 0.34$.

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

c) Cubic: $d \approx 0.35$.

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

APPENDIX B
PROOFS

We can take $T = 1$ in (34) because, for any fixed $T > 0$, and assuming input spectra as (33):

$$\begin{aligned} \|f - \tilde{f}_T\|_{L_2}^2 &\approx \int_{-0.5/T}^{0.5/T} \left|\hat{f}(\omega)\right|^2 E(T\omega)\,\mathrm{d}\omega \\ &= T \int_{-0.5}^{0.5} \left|f(\omega/T)\right|^2 E(\omega)\,\mathrm{d}\omega \\ &= T^{p+1} \int_{-0.5}^{0.5} 1/\omega^p\, E(\omega)\,\mathrm{d}\omega \\ &= T^{p+1} \int_{-0.5}^{0.5} \left|\hat{f}(\omega)\right|^2 E(\omega)\,\mathrm{d}\omega. \end{aligned} \tag{54}$$

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

Simplifying:

$$\begin{aligned} E(\omega) &= 1 - \frac{\varphi(\omega)^2}{\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)} + \widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)\left(\hat{\boldsymbol{q}}(\omega)^2 - 2\frac{\hat{\boldsymbol{q}}(\omega)\hat{\varphi}(\omega)}{\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)} + \frac{\hat{\varphi}(\omega)^2}{\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega)^2}\right) \\ &= 1 - 2\hat{\boldsymbol{q}}(\omega)\hat{\varphi}(\omega) + \hat{\boldsymbol{q}}(\omega)^2\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega) \\ &= \big(1 - \hat{\boldsymbol{q}}(\omega)\hat{\varphi}(\omega)\big)^2 - \hat{\boldsymbol{q}}(\omega)^2\hat{\varphi}(\omega)^2 + \hat{\boldsymbol{q}}(\omega)^2 \sum_n \hat{\varphi}(\omega+n)^2 \\ &= \big(1 - \hat{\boldsymbol{q}}(\omega)\hat{\varphi}(\omega)\big)^2 + \sum_{n\neq 0} \hat{\boldsymbol{q}}(\omega+n)^2\hat{\varphi}(\omega+n)^2 \\ &= \big(1 - \hat{\varphi}_{\mathrm{qi}}(\omega)\big)^2 + \sum_{n\neq 0} \hat{\varphi}_{\mathrm{qi}}(\omega+n)^2. \end{aligned} \tag{55}$$

Above, we used the following equalities

$$\hat{\varphi}_{\mathrm{qi}}(\omega) = \hat{\boldsymbol{q}}(\omega)\hat{\varphi}(\omega) \tag{56}$$

$$\hat{\boldsymbol{q}}(\omega) = \hat{\boldsymbol{q}}(\omega + n), \forall n \in \mathbf{N}, \quad \text{and} \tag{57}$$

$$\widehat{\boldsymbol{a}_{\boldsymbol{\varphi}}}(\omega) = \sum_n \hat{\varphi}(\omega + n)^2. \tag{58}$$

The sum of non-negative terms in (55) shows us that

$$E(0.5-\varepsilon) \approx 0 \Rightarrow \begin{cases} \hat{\varphi}_{\mathrm{qi}}(0.5-\varepsilon) \approx 1, \\ \hat{\varphi}_{\mathrm{qi}}(\text{-}0.5-\varepsilon) \approx 0. \end{cases} \tag{59}$$

Symmetry of E implies $E(\text{-}0.5+\varepsilon) \approx 0$. From (55), we have

$$E(\text{-}0.5+\varepsilon) \approx 0 \Rightarrow \begin{cases} \hat{\varphi}_{\mathrm{qi}}(\text{-}0.5+\varepsilon) \approx 1, \\ \hat{\varphi}_{\mathrm{qi}}(0.5+\varepsilon) \approx 0. \end{cases} \tag{60}$$

ACKNOWLEDGMENT

This work has been funded in part by a PhD fellowship and a PQ 305479/2014-8 grant from CNPq.

REFERENCES

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

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

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

[4] T. Blu, P. Thévenaz, and M. Unser. MOMS: Maximal-order interpolation of minimal support. *IEEE Transactions on Image Processing*, 10(7):1069–1080, 2001.

[5] T. Blu, P. Thévenaz, and M. Unser. Complete parametrization of piecewise-polynomial interpolation kernels. *IEEE Transactions on Image Processing*, 12(11):1297–1309, 2003.

[6] T. Blu, P. Thévenaz, and M. Unser. Linear interpolation revitalized. *IEEE Transactions on Image Processing*, 13(5): 710–719, 2004.

[7] E. Catmull and R. Rom. A class of local interpolating splines. In *Computer Aided Geometric Design*, pages 317–326, 1974.

[8] C. Chui and H. Diamond. A characterization of multivariate quasi-interpolation formulas and its applications. *Numerische Mathematik*, 57(2):105–121, 1990.

[9] L. Condat, T. Blu, and M. Unser. Beyond interpolation: optimal reconstruction by quasi-interpolation. In *Proceedings of the*

IEEE International Conference on Image Processing, volume 1, pages 33–36, 2005.

[10] M. Dalai, R. Leonardi, and P. Migliorati. Efficient digital prefiltering for least-squares linear approximation. In *Visual Content Processing and Representation*, volume 3893 of *Lecture Notes in Computer Science*, pages 161–169. Springer Berlin Heidelberg, 2006.

[11] C. de Boor. Quasiinterpolants and the approximation power of multivariate splines. In M. Gasca and C. A. Micchelli, editors, *Computation of Curves and Surfaces*. Kluwer Academic, 1990.

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

[13] C. E. Duchon. Lanczos filtering in one and two dimensions. *Journal of Applied Meteorology*, 18(8):1016–1022, 1979.

[14] R. W. Hamming. *Digital Filters*. Prentice Hall, 1977.

[15] H. S. Hou and H. C. Andrews. Cubic splines for image interpolation and digital filtering. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 25(6):508–517, 1978.

[16] W. H. Hsiao and R. P. Millane. 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, 2005.

[17] R. Hummel. Sampling for spline reconstruction. *SIAM Journal on Applied Mathematics*, 43(2):278–288, 1983.

[18] J. Kajiya and M. Ullner. Filtering high quality text for display on raster scan devices. *Computer Graphics (Proceedings of ACM SIGGRAPH 1981)*, 15(3):7–15, 1981.

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

[20] Kodak. True colour kodak images, 2010. URL http://r0k.us/graphics/kodak/.

[21] M. A. Malcolm and J. Palmer. A fast method for solving a class of tridiagonal linear systems. *Communications of the ACM*, 17 (1):14–17, 1974.

[22] M. D. McCool. Analytic antialiasing with prism splines. In *Proceedings of ACM SIGGRAPH 1995*, pages 429–436, 1995.

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

[24] E. H. W. Meijering, W. J. Niessen, and M. A. Viergever. Quantitative evaluation of convolution-based methods for medical image interpolation. *Medical Image Analysis*, 5(2):111–126, 2001.

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

[26] D. Nehab and H. Hoppe. A fresh look at generalized sampling. *Foundations and Trends in Computer Graphics and Vision*, 8(1): 1–84, 2014.

[27] J. A. Nelder and R. Mead. A simplex method for function minimization. *The Computer Journal*, 7(4):308–313, 1965.

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

[29] J. A. Parker, R. V. Kenyon, and D. E. Troxel. Comparison of interpolating methods for image resampling. *IEEE Transactions on Medical Imaging*, MI-2(1):31–39, 1983.

[30] S. S. Rifman. Evaluation of digital correction techniques for ERTS images. Technical Report TRW 20634-6003-TU-00, NASA Goddard Space Flight Center, 1973.

[31] Daniel L. Ruderman. Origins of scaling in natural images. *Vision Research*, 37(23):3385–3398, 1997.

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

[33] D. Scherzer, L. Yang, O. Mattausch, D. Nehab, P. V. Sander, M. Wimmer, and E. Eisemann. Temporal coherence methods in real-time rendering. *Computer Graphics Forum*, 31(8):2378–2408, 2012.

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

[35] G. Strang and G. Fix. 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, 1971.

[36] P. Thévenaz, T. Blu, and M. Unser. Interpolation revisited. *IEEE Transactions on Medical Imaging*, 19(17):739–758, 2000.

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

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

[39] M. Unser and A. Aldroubi. A general sampling theory for nonideal acquisition devices. *IEEE Transactions on Signal Processing*, 42(11):2915–2925, 1994.

[40] M. Unser, A. Aldroubi, and M. Eden. Fast B-spline transforms for continuous image representation and interpolation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13 (3):277–285, 1991.

[41] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.

Leonardo Sacht is an adjunct professor at Federal University of Santa Catarina (UFSC) in Florianopolis, Brazil. He received a bachelor degree in Mathematics and Scientific Computing from UFSC in 2008 and MSc and DSc degrees in Mathematics from the Institute for Pure and Applied Mathematics (IMPA) in 2010 and 2014, respectively. He also spent one year between 2012 and 2013 as a visiting student at ETH Zurich, in Switzerland. His main areas of research are image processing, geometry processing and numerical analysis.

Diego Nehab is an associate professor at the National Institute for Pure and Applied Mathematics (IMPA) in Rio de Janeiro, Brazil. He received BEng and MSc degrees in Computer Science from PUC-Rio in 2000 and 2002, respectively, and a PhD degree also in Computer Science from Princeton University in 2007. Before joining IMPA in 2010, he worked as a post-doctoral researcher at Microsoft Research in Redmond. He is interested in most topics related to Computer Graphics, but focuses on parallelism, real-time rendering, and image processing.