The Shannon Total Variation

Rémy Abergel · Lionel Moisan

(revised preprint, to appear in 2017 in *Journal of Mathematical Imaging and Vision*)

Abstract Discretization schemes commonly used for total variation regularization lead to images that are difficult to interpolate, which is a real issue for applications requiring subpixel accuracy and aliasing control. In the present work, we reconciliate total variation with Shannon interpolation and study a Fourier-based estimate that behaves much better in terms of grid invariance, isotropy, artifact removal, and sub-pixel accuracy. We show that this new variant (called Shannon total variation) can be easily handled with classical primal-dual formulations, and illustrate its efficiency on several image processing tasks, including deblurring, spectrum extrapolation, and a new aliasing reduction algorithm.

Keywords total variation, image interpolation, Shannon theory, Legendre-Fenchel duality, aliasing, image restoration.

1 Introduction

Since total variation (TV) regularization was proposed by Rudin, Osher and Fatemi for image denoising [62], it has proven extremely useful for many applications (and beyond image data, for that matter) like image deblurring [70,19], inpainting [20], interpolation [36], spectral extrapolation [60], image decomposition [69], super-resolution [5], stereovision [48], and much more (see [15] and references therein for more examples). In the last decade, the development of dual and primal-dual formulations [13,6,72,31,17] and graph-cuts methods [24] has provided efficient algorithms for TV-based minimization problems, thus increasing even further the popularity of TV regularization.

Université Paris Descartes, MAP5 (CNRS UMR 8145), Sorbonne Paris Cité, France.

A modern way to explain the efficiency of TV is to see it as a sparsity-promoting model: being defined by a L^1 norm (of the gradient), TV minimization tends to favor solutions whose gradient is sparse (that is, often takes the value 0), which corresponds to the so-called cartoon images. Of course, real-life photographs are not cartoons, but outside textured regions (which can be ignored in many image analysis tasks) they are close to that. Another explanation of the usefulness of TV is its ability to penalize oscillations (which is typically the kind of structures one wants to avoid when solving an ill-posed inverse problem) while allowing discontinuities at the same time.

When it comes to implementing an optimization problem involving a TV regularization term, like, e.g., TV denoising of an image u_0 by

$$\operatorname*{argmin}_u \|u-u_0\|^2 + \lambda \mathrm{TV}(u), \tag{1}$$

(where $\lambda > 0$ is a positive parameter selecting the desired amount of regularization), the issue of TV discretization arises. Most algorithms choose to approximate the continuous TV by a sum (over all pixels) of the ℓ^2 norm of a discrete finite-difference estimate of the image gradient, that is,

$$\mathrm{TV}^{\mathrm{d}}(u) = \sum_{(k,l)\in\Omega} \sqrt{(\partial_1 u(k,l))^2 + (\partial_2 u(k,l))^2} \tag{2}$$

$$\text{where} \quad \begin{cases} \partial_1 u(k,l) = u(k+1,l) - u(k,l), \\ \partial_2 u(k,l) = u(k,l+1) - u(k,l), \end{cases} \tag{3}$$

and $u : \Omega \to \mathbb{R}$ is a discrete gray-level image defined on the finite domain $\Omega \subset \mathbb{Z}^2$ (we purposely ignore boundary issues here, as they are not related to our discussion). In the following, we shall refer to (2) as the discrete TV. In some situations, an anisotropic scheme (ℓ^1

norm) may be used [14,43,1], leading to the anisotropic discrete TV

$$\mathrm{TV}^{\mathrm{d}}_{\mathrm{ani}}(u) = \sum_{(k,l)\in\Omega} |\partial_1 u(k,l)| + |\partial_2 u(k,l)|.$$

Curiously enough, as popular as they are, these numerical schemes present strong drawbacks in terms of image quality at pixel and subpixel scales. Indeed, an image obtained by minimizing TV^{d}-based energies is very difficult to interpolate, or, said differently, badly sampled according to Shannon theory. In practice, this means that trying to interpolate such an image will result in the appearance of undesired artifacts (see Fig. 1), generally a mix between blockiness and ringing depending on the interpolation method. This strongly limits the possibility of exploiting an image delivered by a TV^{d}-based scheme, as usual operations like geometric transformations, registration, sub-pixel shape matching, derivative estimates (not to mention others) require well-interpolable images. New discrete schemes have been recently proposed [16,23] to improve the isotropy of the discrete TV, but they do not solve (nor address) the interpolability issue we consider here.

In the present paper, we study a new formulation of the discrete TV, which reconciliates TV minimization and Shannon theory. This variant, which we shall name *Shannon Total Variation* (STV), first appeared in [44], and was later explicitly considered in [49] and then used in [30,56] under the name *Spectral Total Variation* (but we shall not keep this name since it would introduce a confusion with [34]). The STV variant consists in estimating the true total variation of the exact (continuous) total variation of the Shannon interpolate of u by using a Riemann sum approximation of the associated integral. We show that STV successfully addresses the above-mentioned issues and delivers images on which the discrete sinc and spline interpolations behave nicely, while preserving the desired properties of TV regularization. The lack of isotropy observed with classical finite difference schemes is also naturally avoided with STV. This comes at the expense of a few Fourier Transforms at each iteration of the optimization process, which is, in most applications, an affordable cost considering the strong benefits in terms of image quality.

The paper is organized as follows. In Section 2, we present the discrete sinc interpolation as a consequence of Shannon sampling Theorem, and discuss in particular the (generally overlooked) difficulties encountered with Nyquist frequencies in the case of even image dimensions. We also give an independent justification of discrete sinc interpolation as the unique linear interpolation that defines invertible subpixellic translations,

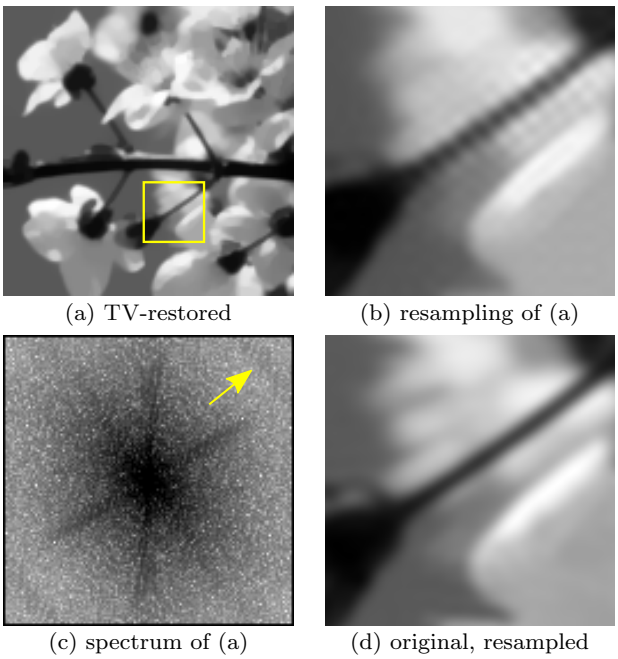

(a) TV-restored (b) resampling of (a) (c) spectrum of (a) (d) original, resampled

Fig. 1 Discrete TV produces aliasing. An image denoised with a classical discrete implementation of TV denoising (a) is improperly sampled, as attested by the aliasing artifact appearing in its Fourier spectrum ((c), red arrow), which is responsible for the undesired oscillating patterns that appear when magnifying the image using Shannon interpolation (b). Note that this artifact is not present on the original image (d). This experiment illustrates the difficulty of manipulating images at a subpixel scale after a processing involving the discrete TV.

and discuss the link with B-spline interpolation. In Section 3, we define STV and discuss the choice of the upsampling factor used to discretize the continuous TV integral into a Riemann sum. We then show in Section 4 that STV-based algorithms can be efficiently implemented by deriving a dual formulation which can be used in the powerful Chambolle-Pock optimization procedure. In Section 5, we illustrate the use of STV regularization in the case of several classical applications (denoising and more general inverse problems like deblurring, image magnification with spectrum extrapolation, tomography). We then present a new STV-based image restoration model involving a weight function in Fourier domain, which leads to interesting applications in terms of de-aliasing and can be viewed as an "image Shannonizer" as it provides a way to approximate a given image by a well-sampled one according to Shannon interpolation (Section 6). We finally conclude in Section 7 and present some perspectives.

2 Shannon interpolation

We here propose a self-contained and detailed exposition of what we shall refer as "discrete Shannon interpolation" (that is, the Shannon interpolation of a discrete signal or image in a periodic setting), which is the starting point of the Shannon Total Variation presented in Section 3. The general form has been known and used for a long time, and most results reported here are not new (see, e.g., [45,63,12]), but several often neglected details (like, in particular, the treatment of Nyquist frequencies in the case of even dimensions) are here clarified and discussed to yield non-ambiguous definitions. We also give a new justification of Shannon interpolation that does not rely on Shannon Sampling Theorem, discuss the link with B-spline interpolation, and show how Shannon interpolation can be combined with B-spline interpolation to implement geometrical transformations of images with efficiency and strong accuracy.

2.1 Shannon Sampling Theorem

A classical way to understand the relation between a (d-dimensional) continuous signal and its sampled version is Shannon Sampling Theorem, which can be considered in some way as the foundation of the digital era. It has been formulated in different contexts by Whittaker [74] and Shannon [64], and also independently by Kotelnikov [40]. In the following, we write $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^{d} x_i y_i$ the canonical Euclidean inner product between two vectors $\mathbf{x} = (x_i)$ and $\mathbf{y} = (y_i)$ of $\mathbb{R}^d$.

Theorem 1 (Shannon Sampling Theorem [64]) *Consider a positive real number δ and an absolutely integrable function $f : \mathbb{R}^d \to \mathbb{R}$ whose Fourier Transform*

$$\widehat{f}(\xi) = \int_{\mathbb{R}^d} f(\mathbf{x})\, e^{-i\langle \mathbf{x}, \xi \rangle}\, dxdy \tag{4}$$

satisfies $$\forall \xi \notin \left[-\frac{\pi}{\delta}, \frac{\pi}{\delta}\right]^d, \quad \widehat{f}(\xi) = 0. \tag{5}$$

Then, f is continuous and uniquely determined by its values on $\delta\mathbb{Z}^d$, as for any $\mathbf{x} \in \mathbb{R}^d$,

$$f(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} f(\delta\, \mathbf{k}) \mathrm{sinc}\left(\frac{\mathbf{x}}{\delta} - \mathbf{k}\right) \tag{6}$$

where the cardinal sine function is defined on $\mathbb{R}^d$ by

$$\mathrm{sinc}(\mathbf{x}) = \prod_{i=1}^{d} \frac{\sin(\pi x_i)}{\pi x_i} \tag{7}$$

with the continuity-preserving convention $\frac{\sin(0)}{0} = 1$.

In the present paper, we will focus on one-dimensional signals ($d = 1$) and two-dimensional images ($d = 2$), but the extension to higher dimensions is straightforward. Apart from establishing a clear correspondence between the support of the Fourier spectrum of the bandlimited function f and the critical sampling step δ permitting its exact reconstruction from discrete samples, Shannon Sampling Theorem provides with Equation 6 (for $\delta = 1$) an interpolation formula that extends to $\mathbb{R}^d$ a discrete signal initially defined on $\mathbb{Z}^d$. However, this formula cannot be used as such in practice since it involves an infinite number of samples. We first discuss that issue in the simpler case $d = 1$.

2.2 Discrete Shannon interpolation of 1-D signals

Let us consider a discrete signal $s : I_M \to \mathbb{R}$ where $M \in \mathbb{N}^*$ and $I_M = \{0, 1, \dots M-1\}$. In order to define the Shannon interpolate $S : \mathbb{R} \to \mathbb{R}$ of s using (6), we first need to extend s into an infinite signal $\tilde{s} : \mathbb{Z} \to \mathbb{R}$, so that

$$S(x) = \sum_{k \in \mathbb{Z}} \tilde{s}(k)\, \mathrm{sinc}(x - k). \tag{8}$$

Extending s with 0 in $\mathbb{Z} \setminus I_M$ would be a poor solution, as it would interpolate a constant discrete signal s by an oscillating function. Instead, the classical solution consists in extending s as a M-periodic function $\tilde{s}(k) = s(k \mod M)$. Using such a periodic extension is not completely straightforward as it does not fit the hypotheses of Shannon Sampling Theorem (a M-periodic $\tilde{s} : \mathbb{Z} \to \mathbb{R}$ cannot be the sampled version of an absolutely integrable bandlimited function), but we can formally write

$$\begin{aligned} S(x) &= \sum_{k \in \mathbb{Z}} \tilde{s}(k)\, \mathrm{sinc}(x - k) \\ &= \sum_{p \in \mathbb{Z}} \sum_{k \in I_M} s(k)\, \mathrm{sinc}(x - k - pM) \\ &= \sum_{k \in I_M} s(k) \left(\sum_{p \in \mathbb{Z}} \mathrm{sinc}(x - k - pM) \right), \end{aligned}$$

and the factor of $s(k)$ can be explicitly computed with

Proposition 1 (discrete cardinal sine) *Define the discrete cardinal sine of order M as the M-periodization of the cardinal sine function, that is,*

$$\mathrm{sincd}_M(x) := \lim_{n \to +\infty} \sum_{p=-n}^{n} \mathrm{sinc}(x - pM). \tag{9}$$

Then, one has

$$\mathrm{sincd}_M(x) = \begin{cases} \dfrac{\sin(\pi x)}{M \sin\left(\frac{\pi x}{M}\right)} & \text{if } M \text{ is odd,} \\ \dfrac{\sin(\pi x)}{M \tan\left(\frac{\pi x}{M}\right)} & \text{if } M \text{ is even,} \end{cases} \tag{10}$$

where the indeterminate forms $0/0$ *are solved by continuity, that is,* $\mathrm{sincd}_M(x) = 1$ *for any* $x \in M\mathbb{Z}$.

The proof is given in Appendix A. In view of Proposition 1, we can rewrite the interpolation of s as

$$S(x) = \sum_{k \in I_M} s(k)\, \mathrm{sincd}_M(x-k). \tag{11}$$

Note that the kernel defined in (10) has been known for a long time (see for example [77,41]). For small values of $|x|$ (more precisely, when $|x| \ll M$), we have $M \sin \frac{\pi x}{M} \simeq M \tan \frac{\pi x}{M} \simeq \pi x$, so that $\mathrm{sincd}_M(x) \simeq \mathrm{sinc}(x)$, which formally shows the asymptotic equivalence between sinc and sincd_M interpolation as $M \to +\infty$.

In practice, (11) is barely used, since there is an equivalent (but numerically more efficient) formulation due to the fact that sincd_M is a trigonometric polynomial.

Proposition 2 *The function* sincd_M *is a trigonometric polynomial, which can be written*

$$\mathrm{sincd}_M(x) = \mathrm{Re}\left(\frac{1}{M} \sum_{\alpha \in \widehat{I}_M} e^{2i\pi \frac{\alpha x}{M}} \right) \tag{12}$$

where $\widehat{I}_M = \left[-\frac{M}{2}, \frac{M}{2}\right) \cap \mathbb{Z}$ *and the real part in (12) is required only if* M *is even.*

Proof The set $\widehat{I}_M$ is made of M consecutive integer values, and can thus be written

$$\widehat{I}_M = \{a, a+1, \dots a+M-1\},$$

where $a = -\lfloor \frac{M}{2} \rfloor$ and $\lfloor \frac{M}{2} \rfloor$ denotes the (lower) integer part of $\frac{M}{2}$. Thus, if $x \notin M\mathbb{Z}$ we have

$$\begin{aligned} \sum_{\alpha \in \widehat{I}_M} e^{2i\pi \frac{\alpha x}{M}} &= \sum_{\alpha=a}^{a+M-1} \left(e^{2i\pi \frac{x}{M}}\right)^{\alpha} \\ &= e^{2i\pi \frac{ax}{M}} \cdot \frac{1 - e^{2i\pi x}}{1 - e^{2i\pi \frac{x}{M}}} \\ &= e^{i\pi x \frac{2a+M-1}{M}} \cdot \frac{\sin(\pi x)}{\sin \pi \frac{x}{M}}. \end{aligned}$$

If M is odd, $2a + M - 1 = 0$ and we get

$$\frac{1}{M} \sum_{\alpha \in \widehat{I}_M} e^{2i\pi \frac{\alpha x}{M}} = \frac{\sin(\pi x)}{M \sin \pi \frac{x}{M}} = \mathrm{sincd}_M(x)$$

as expected. If M is even, $2a + M - 1 = -1$ and we now obtain

$$\begin{aligned} \mathrm{Re}\left(\frac{1}{M} \sum_{\alpha \in \widehat{I}_M} e^{2i\pi \frac{\alpha x}{M}} \right) &= \frac{\sin(\pi x)}{M \sin \pi \frac{x}{M}} \cdot \mathrm{Re}(e^{-i \frac{\pi x}{M}}) \\ &= \frac{\sin(\pi x)}{M \tan \pi \frac{x}{M}} \\ &= \mathrm{sincd}_M(x) \end{aligned}$$

as well. □

A consequence of Proposition 2 is that the Shannon interpolation formula (11) can be rewritten using the Discrete Fourier Transform recalled below.

Definition 1 The discrete Fourier Transform (DFT) of a signal $s : I_M \to \mathbb{R}$ is the M-periodic complex-valued signal $\widehat{s}$ defined by

$$\forall \alpha \in \mathbb{Z}, \quad \widehat{s}(\alpha) = \sum_{k \in I_M} s(k) e^{-2i\pi \frac{\alpha k}{M}}.$$

Proposition 3 *The discrete Shannon interpolation of a signal* $s : I_M \to \mathbb{R}$ *can be written*

$$S(x) = \mathrm{Re}\left(\frac{1}{M} \sum_{\alpha \in \widehat{I}_M} \widehat{s}(\alpha)\, e^{2i\pi \frac{\alpha x}{M}} \right), \tag{13}$$

and the real part is required only if M *is even.*

Proof Thanks to Proposition 2, the Shannon interpolate of s defined by (11) can be rewritten

$$\begin{aligned} S(x) &= \sum_{k \in I_M} s(k)\, \mathrm{Re}\left(\frac{1}{M} \sum_{\alpha \in \widehat{I}_M} e^{2i\pi \frac{\alpha(x-k)}{M}} \right) \\ &= \mathrm{Re}\left(\frac{1}{M} \sum_{\alpha \in \widehat{I}_M} \left(\sum_{k \in I_M} s(k) e^{-2i\pi \frac{\alpha k}{M}} \right) e^{2i\pi \frac{\alpha x}{M}} \right) \end{aligned}$$

from which (13) directly follows. □

Note that if $x \in I_M$, the function $\alpha \mapsto \widehat{s}(\alpha)\, e^{2i\pi \frac{\alpha x}{M}}$ is M-periodic, and since $\widehat{I}_M$ is an interval of M consecutive values, we have

$$\frac{1}{M} \sum_{\alpha \in \widehat{I}_M} \widehat{s}(\alpha)\, e^{2i\pi \frac{\alpha x}{M}} = \frac{1}{M} \sum_{\alpha \in I_M} \widehat{s}(\alpha)\, e^{2i\pi \frac{\alpha x}{M}} = s(x)$$

as we recognize the inverse DFT of $\widehat{s}$. As expected, the Shannon interpolation defined by (13) is exact (that is, the restriction of S to I_M is exactly s).

Also remark that when M is even, we need a real part to cancel the imaginary part of the term $\alpha = -\frac{M}{2}$ in the sum (13) since the conjugate term (which would correspond to $\alpha = \frac{M}{2}$) is not present in the sum. The real part can be avoided when $\widehat{s}(-\frac{M}{2}) = 0$, or by considering instead a sum with $M+1$ terms, as stated by

Proposition 4 *Define, for integer M,*

$$\varepsilon_M(\alpha) = \begin{cases} 1/2 & \text{if } |\alpha| = \frac{M}{2}, \\ 1 & \text{otherwise.} \end{cases} \tag{14}$$

The discrete Shannon interpolate of a signal $s : I_M \to \mathbb{R}$ can be written

$$S(x) = \frac{1}{M} \sum_{\substack{\alpha\in\mathbb{Z} \\ -\frac{M}{2}\le\alpha\le\frac{M}{2}}} \varepsilon_M(\alpha) \cdot \widehat{s}(\alpha)\, e^{2i\pi\frac{\alpha x}{M}}. \tag{15}$$

Note that if M is odd, ε_M is identically equal to 1. This asymmetry between the case M odd and M even can be simply explained. Let us define as T_M the real vector space of real-valued trigonometric polynomials that can be written as complex linear combinations of $(x \mapsto e^{2i\pi\frac{\alpha x}{M}})_{-\frac{M}{2}\le\alpha\le\frac{M}{2}}$. If M is odd, $\dim T_M = M$ and there is a unique element S of T_M that exactly interpolates s, and it is given by (13). If M is even, $\dim T_M = M+1$ and any element of T_M that exactly interpolates s can be written under the form $S(x) + \lambda \sin(\pi x)$ with $\lambda \in \mathbb{R}$, and the interpolation formula (13) corresponds to the implicit (minimal norm) choice $\lambda = 0$.

2.3 Shannon interpolation of 2-D images

Let $u : I_M \times I_N \to \mathbb{R}$ be a discrete $M \times N$ image. Its 2-dimensional DFT $\widehat{u} : \mathbb{Z}^2 \to \mathbb{C}$ is defined by

$$\widehat{u}(\alpha,\beta) = \sum_{\substack{k\in I_M \\ l\in I_N}} u(k,l) e^{-2i\pi\left(\frac{\alpha k}{M}+\frac{\beta l}{N}\right)}, \tag{16}$$

and the natural extension of (11) is

Definition 2 The discrete Shannon interpolate of an image $u : I_M \times I_N \to \mathbb{R}$ is $U : \mathbb{R}^2 \to \mathbb{R}$ defined by

$$U(x,y) = \sum_{\substack{k\in I_M \\ l\in I_N}} u(k,l)\, \mathrm{sincd}_M(x-k)\, \mathrm{sincd}_N(y-l). \tag{17}$$

As in the 1-D case, Definition 2 can be reformulated in the Fourier domain.

Proposition 5 *The discrete Shannon interpolate of an image $u : I_M \times I_N \to \mathbb{R}$ can be written*

$$U(x,y) = \frac{1}{MN} \times \sum_{\substack{\alpha,\beta\in\mathbb{Z} \\ -\frac{M}{2}\le\alpha\le\frac{M}{2} \\ -\frac{N}{2}\le\beta\le\frac{N}{2}}} \varepsilon_M(\alpha)\varepsilon_N(\beta) \cdot \widehat{u}(\alpha,\beta) e^{2i\pi\left(\frac{\alpha x}{M}+\frac{\beta y}{N}\right)}, \tag{18}$$

where ε_M and ε_N are defined in (14).

Proof Simply remark that (12) can be rewritten

$$\mathrm{sincd}_M(x) = \frac{1}{M} \sum_{-\frac{M}{2}\le\alpha\le\frac{M}{2}} \varepsilon_M(\alpha) e^{2i\pi\frac{\alpha x}{N}} \tag{19}$$

and (18) follows quite directly from (16) and (17). □

Note that if both M and N are odd, (18) boils down to

$$U(x,y) = \frac{1}{MN} \sum_{\substack{\alpha\in\widehat{I}_M \\ \beta\in\widehat{I}_N}} \widehat{u}(\alpha,\beta) e^{2i\pi\left(\frac{\alpha x}{M}+\frac{\beta y}{N}\right)}, \tag{20}$$

which is exactly the definition of the inverse DFT of $\widehat{u}$ for integer values of x and y. Thus, one could wonder whether in the general case (M, N even or odd) the generalization of (13), that is,

$$U'(x,y) = \mathrm{Re}\left(\frac{1}{MN} \sum_{\substack{\alpha\in\widehat{I}_M \\ \beta\in\widehat{I}_N}} \widehat{u}(\alpha,\beta)\, e^{2i\pi\left(\frac{\alpha x}{M}+\frac{\beta y}{N}\right)} \right), \tag{21}$$

would be an equivalent definition of U as in the 1-D case. In fact, (17) and (21) both define bivariate trigonometric polynomials of $T_M \otimes T_N$ that exactly interpolate u in $I_M \times I_N$, but they differ when both M and N are even. In that case, $U'(x,y)$ can still be rewritten in a form similar to (18), but we have to change the coefficient $\varepsilon_M(\alpha)\varepsilon_N(\beta)$ into

$$\varepsilon'_{M,N}(\alpha,\beta) = \begin{cases} \frac{1}{2} & \text{if } (\alpha,\beta) = \pm(\frac{M}{2},\frac{N}{2}), \\ 0 & \text{if } (\alpha,\beta) = \pm(-\frac{M}{2},\frac{N}{2}), \\ \varepsilon_{M,N}(\alpha,\beta) & \text{otherwise.} \end{cases} \tag{22}$$

Thus, one easily shows that

$$U'(x,y) = U(x,y) - \widehat{u}\left(\frac{M}{2},\frac{N}{2}\right) \sin(\pi x)\sin(\pi y). \tag{23}$$

Even if this difference is expected to be small for natural images (the Fourier coefficients of a natural image decrease rather quickly as the frequency increases), the true interpolate U is to be preferred to U' as it is separable and more invariant; in particular, the transform $u \mapsto U'$ does not commute with the plane transforms $(x,y) \mapsto (-x,y)$ and $(x,y) \mapsto (x,-y)$.

In the literature, most papers involving 2-D discrete Shannon interpolation either do not mention this issue [33,44], or restrict their study to odd numbers of samples [45,65]. In some other papers (see, e.g., the first version of [10]), the slightly incorrect variant U' is used in place of U, probably because taking the real part is the most simple way to get rid of the imaginary part that naturally appears when Nyquist frequencies are not carefully handled.

2.4 Dealing with periodization artifacts

Using discrete Shannon interpolation requires a careful handling of edge effects, as the implicit periodization of the image may produce interpolation artifacts (that is, undesired oscillations) near the boundary of the image domain if the intensity values on the opposite edges of the image domain do not match well. This issue is discussed in detail in [50], and an efficient solution is proposed that consists in decomposing the original image into the sum of a periodic image and a smooth image. Other solutions exist like symmetrization or apodization using an appropriate weight function (e.g., a Hamming window), but they appear to be less efficient in general. In all the experiments presented throughout this paper (and in particular in Section 5 and 6), the periodic plus smooth decomposition of [50] will systematically be used.

2.5 Shannon interpolation and reversible transforms

As we saw earlier, Shannon Sampling Theorem provides a nice theoretical framework that establishes a one-to-one correspondence between continuous bandlimited and discrete images, which naturally leads to the discrete Shannon interpolation we just presented. Interestingly, there is another justification for Shannon interpolation, that does not explicitly refer to Shannon Sampling Theorem: basically, it is the only linear interpolation that defines invertible subpixellic translations (in a periodic setting). In the following, we assume for simplicity that M is an odd integer, and write $\mathcal{S}$ the space of M-periodic signals $s : \mathbb{Z} \to \mathbb{R}$.

Theorem 2 *There exists a unique family of linear operators $(T_z)_{z\in\mathbb{R}}$ on $\mathcal{S}$ such that :*

(i) $z \mapsto T_z$ *is continuous,*

(ii) $\forall k, z \in \mathbb{Z},\ T_z s(k) = s(k - z)$,

(iii) $\forall w, z \in \mathbb{R},\ T_{w+z} = T_w \circ T_z$,

(iv) $\lim\limits_{z\to 0} |z|^{-1}\|T_z - id\|_2$ *is minimal.*

It is defined by

$$T_z s(k) = S(k - z), \tag{24}$$

where S is the discrete Shannon interpolate of s defined in (11) or equivalently in (13).

The Proof is given in Appendix B. Theorem 2 remains true for M even, provided that we define $\mathcal{S}$ in this case by

$$\mathcal{S} = \left\{ s : I_M \to \mathbb{R},\ \sum_{k\in I_M} (-1)^k s(k) = 0 \right\}. \tag{25}$$

(Note that it is equivalent to assume $\widehat{s}(M/2) = 0$). This restriction is needed to exclude from $\mathcal{S}$ the alternated signal $k \mapsto (-1)^k$, which clearly cannot be translated in a way compatible with Hypotheses (ii) and (iii).

Theorem 2 shows that the only minimal continuous semi-group extending the integer (periodic) translations is given by Shannon interpolation. This result is interesting in the sense that it brings another justification to Shannon interpolation without referring to Shannon Sampling Theorem (or to the Fourier Transform, for that matter): among linear interpolation methods, only Shannon interpolation is able to translate images without information loss.

From Equation (74), we can see that a subpixellic translation with Shannon interpolation can be implemented with two DFTs, as

$$\widehat{T_z s}(\alpha) = e^{-2i\pi\alpha z/M}\, \widehat{s}(\alpha). \tag{26}$$

Moreover, T_z is a linear isometry ($\|T_z s\|_2 = \|s\|_2$), which is another way to explain that no information loss occurs.

Signal and image magnification is also very easy to perform with discrete Shannon interpolation, as it essentially boils down to a *zero-padding* in the Fourier domain (for even dimensions, it is also necessary to split the coefficients corresponding to Nyquist frequencies $\alpha = \pm\frac{M}{2}$ or $\beta = \pm\frac{N}{2}$). More surprisingly, image rotation can also be implemented efficiently with the DFT (see [75]), thanks to the following factorization of a rotation matrix into a product of shear matrices:

$$\begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} = \begin{pmatrix} 1 & -t \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ \sin\theta & 1 \end{pmatrix} \begin{pmatrix} 1 & -t \\ 0 & 1 \end{pmatrix} \tag{27}$$

with $t = \tan\frac{\theta}{2}$. As a shear transform like

$$v(x, y) = u(x - ty, y) \tag{28}$$

consists in applying 1-D translations to each line of u, a 2-D rotation can be decomposed as a combination of 1-D translations, which can be implemented in the Fourier domain. For that reason, image rotation with discrete Shannon interpolation is a linear isometry, and can thus be considered as a lossless transform.

2.6 Link with spline interpolation

A popular alternative to Shannon interpolation is spline interpolation. Without going too much into details (see [67,66] and the references therein), it is worth mentioning the relation between spline and Shannon interpolation, and to understand how they can be combined to

yield what is probably the most accurate and efficient linear interpolation of bandlimited signals.

The spline interpolation of order n ($n \in \mathbb{N}$) of a signal $s \in \ell^2(\mathbb{Z})$ can be written

$$S^n(x) = \sum_{k\in\mathbb{Z}} c(k)\,\beta^n(x-k), \tag{29}$$

where $\beta^n : \mathbb{R} \to \mathbb{R}$ is the spline of order n defined by induction by $\beta^0 = \mathbb{1}_{[-\frac{1}{2},\frac{1}{2})}$ and $\beta^{k+1} = \beta^k * \beta^0$ for all $k \in \mathbb{N}$. It can be shown that the signal $c : \mathbb{Z} \to \mathbb{R}$ is uniquely defined by the interpolation constraint $S^n(k) = s(k), k \in \mathbb{Z}$. When $n \in \{0,1\}$, one has $c = s$ and spline interpolation corresponds to piecewise constant ($n = 0$) or piecewise affine ($n = 1$) interpolation. When $n > 1$, c depends linearly on s and can be efficiently computed using recursive filtering [67]. As remarked in [68], spline interpolation achieves an optimal trade-off between complexity (the support of β^n is an interval with length $n+1$) and asymptotic accuracy (rate of convergence towards the unsampled signal as the sampling step tends to 0). How does spline interpolation compare with Shannon interpolation? Indeed, (29) can be rewritten as

$$S^n(x) = \sum_{k\in\mathbb{Z}} s(k)\,\beta^n_{\text{card}}(x-k), \tag{30}$$

where $\beta^n_{\text{card}} : \mathbb{R} \to \mathbb{R}$ is the cardinal spline of order n defined in the Fourier domain by

$$\widehat{\beta^n_{\text{card}}}(\xi) = \frac{\left(\operatorname{sinc}\frac{\xi}{2\pi}\right)^{n+1}}{\sum_{k\in\mathbb{Z}} \beta^n(k) e^{-ik\xi}}. \tag{31}$$

This provides a nice interpretation of spline interpolation in the Fourier domain, as the Fourier transform of (30) yields

$$\widehat{S}^n(\alpha) = \widehat{s}(\alpha)\widehat{\beta^n_{\text{card}}}(\alpha), \tag{32}$$

where $\widehat{s}(\alpha) = \sum_{k\in\mathbb{Z}} s(k)e^{-ik\alpha}$ is the Fourier Transform of the discrete signal s. Thus, if S is a bandlimited signal ($\operatorname{supp}\widehat{S} \subset [-\pi,\pi]$) and $s(k) = S(k)$ for all $k \in \mathbb{Z}$, the Fourier transform of S_n is deduced from $\widehat{S}$ by periodization and multiplication by $\widehat{\beta^n_{\text{card}}}$. This is to be compared to Shannon interpolation, that recovers the exact signal S since

$$\widehat{S}(\alpha) = \widehat{s}(\alpha)\mathbb{1}_{[-\pi,\pi]}. \tag{33}$$

In fact, $\widehat{\beta^n_{\text{card}}} \to \mathbb{1}_{[-\pi,\pi]}$ as $n \to +\infty$ [2] (or, equivalently, $\beta^n_{\text{card}} \to \operatorname{sinc}$), which means that spline interpolation can be viewed as an approximation of Shannon interpolation (the equivalence being asymptotically obtained for $n = +\infty$). For finite n however, the effect of spline interpolation in the Fourier domain is questionable: it creates high frequencies aliases (by spectrum periodization), and then attenuates the whole spectrum (the known part $[-\pi,\pi]$ included) by an apodization function that is a smooth approximation of $\mathbb{1}_{[-\pi,\pi]}$. This apodization function (that is, $\widehat{\beta^n_{\text{card}}}$) is represented in Fig. 2 for various values of n.

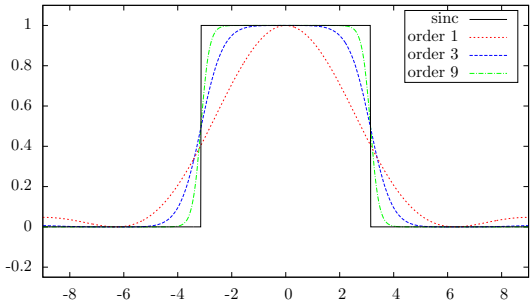

Fig. 2 Cardinal splines in the Fourier domain. The Fourier transform of the interpolation kernels β^n_{card} is represented for $n = 1, 3, 9$. As n increases, they get closer to the ideal low-pass filter obtained with the sinc kernel. The approximation is responsible for blur (attenuation of known frequencies) and aliasing (creation of high frequencies duplicated from existing low frequencies) on spline-interpolated images.

On the one hand, spline interpolation is computationally efficient, and also versatile: it can be used to magnify an image by an arbitrary factor, or to apply an homography or a non-rigid transform to an image. On the other hand, Shannon interpolation is very accurate, as it does not attenuate known Fourier coefficients or create high-frequency aliases. Getting the best of the two worlds (that is, the accuracy of exact Shannon interpolation and the efficiency of spline interpolation) is easy: magnify the original image by a small factor (e.g. 2), and then use spline interpolation on the magnified image. Fig. 3 illustrates the interest of such a combination in the case of a homographic transform.

In this section, we gave a precise definition of Shannon interpolation (with a careful treatment of Nyquist frequencies in the case of even dimensions), and saw how it provides a nice framework for interpolating bandlimited images with a high degree of accuracy. It is particularly useful for imaging sciences that require an accurate treatment of subpixel scales and a strict control of artifacts (in particular, satellite imaging). As we shall see in the next sections, Shannon interpolation can be made compatible with total variation regularization, provided that we use what we shall call the *Shannon total variation*.

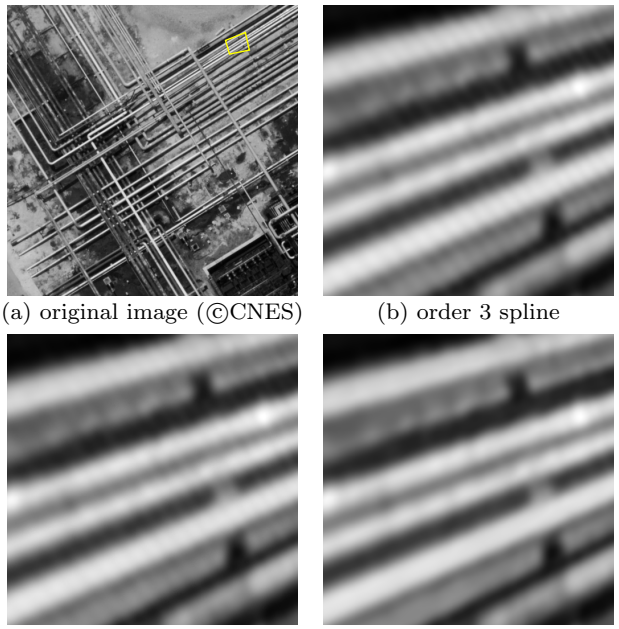

(a) original image (©CNES) (b) order 3 spline

(c) order 9 spline (d) Shannon + order 3 spline

Fig. 3 High quality homographic transforms using a combination of Shannon and spline interpolations. Applying an homographic transform to an image (a) requires the use of an interpolation scheme. Spline kernels are interesting but may produce undesired artifacts (the slight superimposed line hatch patterns in b,c) due to the creation of spurious high frequencies. Applying the same transform with Shannon interpolation alone would be computationally very expensive, but a simple $\times 2$ magnification with Shannon interpolation followed by an homographic transform implemented by a spline of order 3 produces an artifact-free image for a computational cost equivalent to spline interpolation.

3 The Shannon total variation

3.1 Definition

Let $|\cdot|$ denotes the ℓ^2 norm over $\mathbb{R}^2$, let $\Omega = I_M \times I_N$ denote a 2-D discrete domain of size $M \times N$ and $u \in \mathbb{R}^\Omega$ a discrete gray-level image with domain Ω. We define the Shannon total variation of u by

$$\mathrm{STV}_\infty(u) = \int_{[0,M]\times[0,N]} |\nabla U(x,y)|\, dxdy\,, \tag{34}$$

where U is the Shannon interpolation of u defined in (17), and $\nabla U : \mathbb{R}^2 \to \mathbb{R}^2$ denotes the gradient of the trigonometric polynomial U. No closed-form formula exist for (34), but we can approximate this continuous integral with the Riemann sum

$$\mathrm{STV}_n(u) = \frac{1}{n^2} \sum_{(k,l)\in\Omega_n} |\nabla_n u(k,l)|\,, \tag{35}$$

where $n \in \mathbb{N}^*$, $\Omega_n = I_{nM} \times I_{nN}$ and

$$\forall (k,l) \in \Omega_n, \quad \nabla_n u(k,l) = \nabla U\left(\tfrac{k}{n}, \tfrac{l}{n}\right)\,.$$

In order to compute $\mathrm{STV}_n(u)$, we need to focus on the practical computation of $\nabla_n u$. By differentiating (18), we get the gradient of U, that is, $\forall (x,y) \in \mathbb{R}^2$,

$$\nabla U(x,y) = \frac{1}{MN} \sum_{\substack{-\frac{M}{2}\le\alpha\le\frac{M}{2}\\ -\frac{N}{2}\le\beta\le\frac{N}{2}}} e^{2i\pi\left(\frac{\alpha x}{M}+\frac{\beta y}{N}\right)}\, g_{\widehat{u}}(\alpha,\beta)\,, \tag{36}$$

where

$$g_{\widehat{u}}(\alpha,\beta) = 2i\pi\, \varepsilon_M(\alpha)\varepsilon_N(\beta)\, \widehat{u}(\alpha,\beta) \begin{pmatrix} \alpha/M \\ \beta/N \end{pmatrix}\,. \tag{37}$$

Therefore, $\nabla_n u$ can be efficiently computed in the Fourier domain for $n \ge 2$ with the following

Proposition 6 *For any $n \ge 2$ and any $(\alpha,\beta) \in \widehat{\Omega_n} := \widehat{I}_{nN} \times \widehat{I}_{nM}$, we have*

$$\widehat{\nabla_n u}(\alpha,\beta) = \begin{cases} n^2 g_{\widehat{u}}(\alpha,\beta) \text{ if } |\alpha| \le \frac{M}{2},\ |\beta| \le \frac{N}{2}\,, \\ 0 \qquad \text{otherwise,} \end{cases} \tag{38}$$

where $g_{\widehat{u}}$ is given by (37).

Proof The result comes directly when writing (36) with $(x,y) = \left(\frac{k}{n}, \frac{l}{n}\right)$, and extending the sum to the frequency domain $\widehat{\Omega_n}$ by adding zero terms. Note that $\widehat{\Omega_n}$ contains all the frequencies (α,β) such that $-\frac{M}{2} \le \alpha \le \frac{M}{2}$ and $-\frac{N}{2} \le \beta \le \frac{N}{2}$ involved in (36) since $n > 1$. □

The next Proposition establishes an upper-bound for the induced ℓ^2 norm (noted $|||\cdot|||$) of the ∇_n operator, which will be useful later.

Proposition 7 *For any $n \ge 2$, we have*

$$|||\nabla_n||| \le n\pi\sqrt{2}\,. \tag{39}$$

Proof Let $u \in \mathbb{R}^\Omega$, from (38) we deduce

$$\|\widehat{\nabla_n u}\|^2 = \left\|n^2 g_{\widehat{u}}\right\|^2 \le 4\pi^2 n^4 \|\widehat{u}\|^2 \left(\frac{1}{4}+\frac{1}{4}\right)\,, \tag{40}$$

since for any (α,β) such as $|\alpha| \le \frac{M}{2}$ and $|\beta| \le \frac{N}{2}$, we have $|\varepsilon_M(\alpha)\varepsilon_N(\beta)\frac{\alpha}{M}|^2 \le \frac{1}{4}$ and $|\varepsilon_M(\alpha)\varepsilon_N(\beta)\frac{\beta}{N}|^2 \le \frac{1}{4}$. Then, using the Parseval identity in (40), that is,

$$\|\nabla_n u\|^2 = \frac{1}{n^2 MN}\|\widehat{\nabla_n u}\|^2 \quad \text{and} \quad \frac{1}{MN}\|\widehat{u}\|^2 = \|u\|^2\,,$$

yields $\|\nabla_n u\|^2 \le 2\pi^2 n^2 \|u\|^2$ and consequently (39). □

Similarly to Proposition 6, we can compute the adjoint of ∇_n in the Fourier domain (the proof is detailed in Appendix C).

Proposition 8 *Let* $\mathrm{div}_n = -\nabla_n^*$, *then for any* $n \geq 2$, $p = (p_x, p_y) \in \mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}$, *and* $(\alpha, \beta) \in \widehat{\Omega} := \widehat{I_M} \times \widehat{I_N}$, *we have*

$$\widehat{\mathrm{div}_n(p)}(\alpha,\beta) = 2i\pi \left(\frac{\alpha}{M} h_{\widehat{p_x}}(\alpha,\beta) + \frac{\beta}{N} h_{\widehat{p_y}}(\alpha,\beta) \right),$$

with $h_{\widehat{p_x}}(\alpha,\beta) =$

$$\begin{cases} \widehat{p_x}(\alpha,\beta) & \text{if } |\alpha| < \frac{M}{2},\ |\beta| < \frac{N}{2} \\ \frac{1}{2}\left(\widehat{p_x}(\alpha,\beta) - \widehat{p_x}(-\alpha,\beta)\right) & \text{if } \alpha = -\frac{M}{2},\ |\beta| < \frac{N}{2} \\ \frac{1}{2}\left(\widehat{p_x}(\alpha,\beta) + \widehat{p_x}(\alpha,-\beta)\right) & \text{if } |\alpha| < \frac{M}{2},\ \beta = -\frac{N}{2} \\ \frac{1}{4}\sum\limits_{\substack{s_1=\pm 1 \\ s_2=\pm 1}} s_1\, \widehat{p_x}(s_1\alpha, s_2\beta) & \text{if } (\alpha,\beta) = (-\frac{M}{2}, -\frac{N}{2}), \end{cases}$$

and $h_{\widehat{p_y}}(\alpha,\beta) =$

$$\begin{cases} \widehat{p_y}(\alpha,\beta) & \text{if } |\alpha| < \frac{M}{2},\ |\beta| < \frac{N}{2} \\ \frac{1}{2}\left(\widehat{p_y}(\alpha,\beta) + \widehat{p_y}(-\alpha,\beta)\right) & \text{if } \alpha = -\frac{M}{2},\ |\beta| < \frac{N}{2} \\ \frac{1}{2}\left(\widehat{p_y}(\alpha,\beta) - \widehat{p_y}(\alpha,-\beta)\right) & \text{if } |\alpha| < \frac{M}{2},\ \beta = -\frac{N}{2} \\ \frac{1}{4}\sum\limits_{\substack{s_1=\pm 1 \\ s_2=\pm 1}} s_2\, \widehat{p_y}(s_1\alpha, s_2\beta) & \text{if } (\alpha,\beta) = (-\frac{M}{2}, -\frac{N}{2}). \end{cases}$$

Notice that Propositions 6 to 8 can be easily adapted to the case $n = 1$. However, we shall not need to consider this case as STV_1 happens to be a poor approximation of STV_∞ (see next section). Note also that similar definitions and propositions could be established for the U' variant of Shannon interpolation mentioned in (21). This variant yields somewhat simpler formulas (no weights are required to handle Nyquist frequencies in the case of even dimensions) since all operators can be obtained by taking the real part of complex-valued images. However, in addition to being less invariant (as discussed in the end of Section 2.3), U' is also computationally less efficient as it requires the computation of DFTs of complex-valued images.

3.2 Choice of the oversampling factor n

When estimating $\mathrm{STV}_\infty(u)$ with $\mathrm{STV}_n(u)$, which value of the oversampling factor n should we choose? We experimentally observed on many images that the convergence with respect to n is extremely fast, so that in practice choosing $n = 2$ or $n = 3$ is enough. Note that an estimate of $\mathrm{STV}_\infty(u)$ could also be obtained by using a finite difference scheme on the image magnified with Shannon interpolation, that is, $n^{-1}\mathrm{TV}^{\mathrm{d}}(Z_n u)$ with

$$\forall (k,l) \in \Omega_n, \quad Z_n u(k,l) = U\left(\frac{k}{n}, \frac{l}{n}\right).$$

Both estimates are consistent in the sense that

$$\lim_{n\to+\infty} \mathrm{STV}_n(u) = \lim_{n\to+\infty} n^{-1}\mathrm{TV}^{\mathrm{d}}(Z_n u) = \mathrm{STV}_\infty(u).$$

However, the convergence speed is much worse for the latter, which comforts us in the choice of STV_n (see Table 1).

n	$n^{-1}\mathrm{TV}^{\mathrm{d}}(Z_n u)$	$\mathrm{STV}_n(u)$
1	$1.6 \cdot 10^{-1}$	$\mathbf{1.8 \cdot 10^{-2}}$
2	$4.2 \cdot 10^{-2}$	$\mathbf{1.3 \cdot 10^{-3}}$
3	$2.1 \cdot 10^{-2}$	$\mathbf{1.7 \cdot 10^{-4}}$
5	$8.6 \cdot 10^{-3}$	$\mathbf{7.3 \cdot 10^{-5}}$
10	$2.8 \cdot 10^{-3}$	$\mathbf{3.4 \cdot 10^{-6}}$

Table 1 Relative errors of two STV_∞ estimates. We compare two estimates of $\mathrm{STV}_\infty(u)$ when u is the classical "*Lena*" image. As we can observe, the relative errors are much smaller with $\mathrm{STV}_n(u)$ (third column) than with $n^{-1}\mathrm{TV}^{\mathrm{d}}(Z_n u)$ (second column), and the convergence with respect to n is faster. Even for $n = 2$, the STV_2 estimate is very accurate with a relative error of 0.1% or so. This experiment has been repeated on many other images, including pure noise images, and yielded similar conclusions for all of them.

As concerns the idea of estimating $\mathrm{STV}_\infty(u)$ with $\mathrm{STV}_1(u)$, the following result shows that it could lead to incorrect results, as controlling $\mathrm{STV}_1(u)$ is not sufficient to control $\mathrm{STV}_\infty(u)$. We believe that, on the contrary, such a control is ensured as soon as $n \geq 2$, even though we have no proof of this affirmation yet.

Theorem 3 *There exists no constant* C *such that*

$$\mathrm{STV}_\infty(u) \leq C \cdot \mathrm{STV}_1(u)$$

for any positive integer M *and any discrete image* u *of size* $M \times M$.

The proof is given in Appendix D. It consists in building a sequence of discrete images u_M with size $M \times M$ such that $\mathrm{STV}_1(u_M)$ is fixed but $\mathrm{STV}_\infty(u_M)$ increases to $+\infty$ with M.

In all the experiments reported in this paper, we used STV_n with $n = 3$, but we observed only very slight improvements (and sometimes none) compared to the case $n = 2$, which should probably be preferred when computational issues are important. Note also that one could choose non-integer values of n (only nM and nN have to be integers), which could also be interesting for computational issues.

4 Duality tools for handling the STV regularizer in a variational framework

4.1 Recall of convex analysis

We here briefly recall some classical convex analysis results needed for non-smooth convex optimization. We refer to [28] for a more detailed presentation.

Consider a finite-dimensional real vector space E and let $E^\star$ denotes its dual space, that is, the set of linear mappings from E to $\mathbb{R}$. Let $\overline{\mathbb{R}}$ denotes the set $\mathbb{R} \cup \{-\infty, +\infty\}$ and $\langle \cdot, \cdot \rangle : E^\star \times E \to \mathbb{R}$ the bilinear mapping defined by

$$\forall \varphi \in E^\star,\ \forall u \in E, \quad \langle \varphi, u \rangle = \varphi(u)\,.$$

An affine function on E is a function $\mathcal{A} : u \mapsto \langle \varphi, u \rangle + \alpha$, where $\varphi \in E^\star$ is called the slope of $\mathcal{A}$ and $\alpha \in \mathbb{R}$ the constant term. We denote by $\Gamma(E)$ the set of functions $F : E \to \overline{\mathbb{R}}$ which are the pointwise supremum of a family of affine functions over E. One can show that F is an element of $\Gamma(E)$ if and only if it is convex and lower semi-continuous (l.s.c.) and does not take the value $-\infty$ unless it is constant. In order to dismiss singular cases, we say that F is proper if it never assumes the value $-\infty$ and is different from the constant $+\infty$. We denote by $\Gamma_0(E)$ the set of proper elements of $\Gamma(E)$.

Given a function $F : E \to \overline{\mathbb{R}}$, the Γ-regularization of F is the largest element of $\Gamma(E)$ which lower bounds F, or, equivalently, the pointwise supremum of all affine functions that lower bound F. Note that an affine function $\mathcal{A}$ with slope $\varphi \in E^\star$ and constant term $\alpha \in \mathbb{R}$ satisfies $\mathcal{A} \leq F$ if and only if $\alpha \leq -F^\star(\varphi)$, where

$$F^\star(\varphi) = \sup_{u \in \mathrm{dom} F} \langle \varphi, u \rangle - F(u)\,, \tag{41}$$

and $\mathrm{dom} F = \{u \in E,\ F(u) < +\infty\}$. The function $F^\star : E^\star \to \overline{\mathbb{R}}$ is called the Legendre-Fenchel transform of F (or the polar, or the conjugate of F). It is an element of $\Gamma(E^\star)$, as it can be seen as the pointwise supremum over the dual space $E^\star$ of all affine functions $\{\mathcal{A}_u\}_{u \in \mathrm{dom} F}$ defined by

$$\forall u \in \mathrm{dom} F, \quad \mathcal{A}_u : \varphi \mapsto \langle \varphi, u \rangle - F(u)\,.$$

Since here E has finite dimension, it is a reflexive space and the Legendre-Fenchel transform of $F^\star$ (noted $F^{\star\star}$) is an element of $\Gamma(E^{\star\star})$ (and thus an element of $\Gamma(E)$), which happens to be exactly the Γ-regularization of F. In particular $F^{\star\star} \leq F$ and we have the characterization

$$F \in \Gamma(E) \Leftrightarrow F^{\star\star} = F\,, \tag{42}$$

which is very useful to derive a primal-dual reformulation of an optimization problem when the cost function decomposes as a sum with at least one term in $\Gamma(E)$. Besides, since E (endowed with the Euclidean inner product) is a Hilbert space, it is self-dual in the sense that any element of $E^\star$ can be represented as the inner product with an element of E, which is very useful in practical computations.

4.2 Chambolle-Pock Algorithm

The recent use in imaging of those powerful convex analysis tools based on duality allowed to properly handle total variation-based variational problems (see e.g. [13,78]). This initiated some flourishing theoretical research (see e.g. [4,31]) as well as the development of efficient numerical schemes [17,22,7,72,54,27,57] dedicated to nonsmooth optimization. We will here briefly recall the formulation of the celebrated first order primal-dual algorithm of Chambolle and Pock [17], which can be used to address various total variation based image processing tasks and comes with nice convergence theorems.

Consider X and Y two finite-dimensional real vector spaces, an inner product $\langle \cdot, \cdot \rangle$ over Y and the generic saddle-point problem

$$\min_{x \in X} \max_{y \in Y}\ G(x) + \langle Kx, y \rangle - F^\star(y)\,, \tag{43}$$

where $F \in \Gamma_0(Y)$, $G \in \Gamma_0(X)$ and $K : X \to Y$ denotes a linear operator. We set $H : (x, y) \mapsto G(x) + \langle Kx, y \rangle - F^\star(y)$ and we assume that problem (43) has at least one solution (i.e. a saddle-point of H). Recall that thanks to (42), for any $x \in X$ we have

$$F(Kx) = F^{\star\star}(Kx) = \sup_{y \in Y} \langle Kx, y \rangle - F^\star(y)\,, \tag{44}$$

therefore one can interpret Equation (43) as a primal-dual formulation of the primal problem

$$\min_{x \in X}\ G(x) + F(Kx) \tag{45}$$

as soon as the $\sup_{y \in Y}$ is indeed a maximum (which will be the case in practice). The proximal splitting algorithm proposed by Chambolle and Pock in [17] (see also [51,59], or more recently [55,21] for more details about proximity operators and proximal algorithms) for solving problem (43) is described in Algorithm 1 below.

In the case $\theta = 0$, one iteration k of Algorithm 1 consists in a proximal ascent of $y \mapsto H(x^k, y)$ followed by a proximal descent of $x \mapsto H(x, y^{k+1})$, yielding a semi-implicit variant of the classical Arrow-Hurwicz algorithm [3]. In the case $\theta > 0$, the iterate $\overline{x}^{k+1} = x^{k+1} + \theta\left(x^{k+1} - x^k\right)$ represents a linear approximation (or extrapolation) of the next iterate x^{k+2} based on the current and the previous iterates x^{k+1} and x^k; it is used to make the scheme more implicit and prove the convergence (in the case $\theta = 1$ and $\tau\sigma|||K|||^2 < 1$) of the sequence $(x^k, y^k)_{k \geq 0}$ toward a saddle-point of H, with an estimate of the convergence rate in $\mathcal{O}(1/N)$ after N iterations (see Theorem 1 in [17]). Notice that some accelerated variants of this algorithm were also proposed

Algorithm 1: Chambolle-Pock resolvent algorithm for problem (43)

Initialization: *Choose* $\tau, \sigma > 0$, $\theta \in [0,1]$, $x^0 \in X$, $y^0 \in Y$, *and set* $\overline{x}^0 = x^0$
(note: for $\theta = 1$, *convergence toward a solution of (43) was proven in [17] when* $\tau\sigma|||K|||^2 < 1$*).*

Iterations: *For* $k \geq 0$, *update* x^k, y^k *and* $\overline{x}^k$ *as follows,*

$$y^{k+1} = \operatorname*{argmin}_{y \in Y} \frac{1}{2\sigma} \left\| y - (y^k + \sigma K \overline{x}^k) \right\|_2^2 + F^\star(y)$$
$$x^{k+1} = \operatorname*{argmin}_{x \in X} \frac{1}{2\tau} \left\| x - (x^k - \tau K^* y^{k+1}) \right\|_2^2 + G(x)$$
$$\overline{x}^{k+1} = x^{k+1} + \theta \left(x^{k+1} - x^k \right)$$

by the same authors, which under regularity assumptions on $F^\star$ and G achieve better convergence rates, thanks to Nesterov-like acceleration strategies [52] (see Algorithms 2 and 3 in [17]).

4.3 Dual formulation of the Shannon total variation

The STV_n operator defined in (35) can be rewritten under the form $\mathrm{STV}_n(u) = \frac{1}{n^2} \|\nabla_n u\|_{1,2}$, noting $\|\cdot\|_{1,2}$ the norm over the space $\mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}$ defined by

$$\forall g \in \mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}, \quad \|g\|_{1,2} = \sum_{(x,y) \in \Omega_n} |g(x,y)| \,.$$

One easily checks that the dual norm of $\|\cdot\|_{1,2}$ is the norm $\|\cdot\|_{\infty,2}$ defined by

$$\forall p \in \mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}, \ \|p\|_{\infty,2} = \max_{(x,y) \in \Omega_n} |p(x,y)| \,.$$

Consequently (see e.g. [8]), the Legendre-Fenchel transform of $\|\cdot\|_{1,2}$, noted $\|\cdot\|_{1,2}^\star$, is the indicator function of the closed unit ball for the norm $\|\cdot\|_{\infty,2}$, defined by

$$\forall p \in \mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}, \ \delta_{\|\cdot\|_{\infty,2} \leq 1}(p) = \begin{cases} 0 & \text{if } \|p\|_{\infty,2} \leq 1 \,, \\ +\infty & \text{otherwise.} \end{cases}$$

We will now use the duality tools described in Section 4.1 to derive a dual formulation of our STV_n operator.

Proposition 9 (dual formulation of STV_n) *For any integer* $n \geq 1$ *and for any image* $u \in \mathbb{R}^\Omega$,

$$\mathrm{STV}_n(u) = \max_{p \in \mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}} \langle \tfrac{1}{n^2} \nabla_n u, p \rangle - \delta_{\|\cdot\|_{\infty,2} \leq 1}(p) \,.$$

Proof Since $\|\cdot\|_{1,2}$ is convex and l.s.c. over $\mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}$, it is an element of $\Gamma(\mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n})$, thereby $\|\cdot\|_{1,2} = \|\cdot\|_{1,2}^{\star\star}$ thanks to (42). Besides, given any image $u \in \mathbb{R}^\Omega$, one has $\mathrm{STV}_n(u) = \frac{1}{n^2} \|\nabla_n u\|_{1,2} = \|\frac{1}{n^2} \nabla_n u\|_{1,2}$. Therefore, $\mathrm{STV}_n(u) = \left\| \frac{1}{n^2} \nabla_n u \right\|_{1,2}^{\star\star}$, i.e.

$$\mathrm{STV}_n(u) = \sup_{p \in \mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}} \langle \tfrac{1}{n^2} \nabla_n u, p \rangle_{\mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}} - \|p\|_{1,2}^\star \,,$$

and $\|p\|_{1,2}^\star$ is exactly $\delta_{\|\cdot\|_{\infty,2} \leq 1}(p)$. Last, one can see that the supremum is attained, since it is nothing but the maximum of the inner product term over the closed unit ball for the dual norm $\|\cdot\|_{\infty,2}$. □

4.4 The Huber STV

The use of TV^{d} as a regularizer for image processing applications has a well-known drawback, the so-called *staircasing effect*, which is the creation of piecewise constant regions with artificial boundaries where one would have expected smooth intensity variations (see for instance [53,18,58] for theoretical results about the staircasing). Several variants of TV^{d} have been proposed in order to avoid this undesirable effect (see for instance [9,42,43]). In the numerical experiments that will be presented in Section 5, we observed that although this staircasing effect is significantly attenuated when using the STV_n variant of TV^{d}, it remains present (at least visually) in the processed images.

In the case of TV^{d}, a classical way to get rid of the staircasing effect consists in replacing the ℓ^2 norm $|\cdot|$ of the gradient in the definition of the TV operator by its smooth Huber approximation with parameter $\alpha > 0$ (coming from the statistical literature [37,38], and used for instance in [71,73,17]). It is defined by

$$\forall y \in \mathbb{R}^2, \quad \mathcal{H}_\alpha(y) = \begin{cases} \frac{|y|^2}{2\alpha} & \text{if } |y| \leq \alpha \,, \\ |y| - \frac{\alpha}{2} & \text{otherwise.} \end{cases} \tag{47}$$

The same adaptation can be easily done in the case of STV by replacing the ℓ^2 norm by the Huber-function $\mathcal{H}_\alpha$ in Equations (34) and (35), which in the case of the Riemann approximation leads to

$$\mathrm{HSTV}_{\alpha,n}(u) = \frac{1}{n^2} \sum_{(x,y) \in \Omega_n} \mathcal{H}_\alpha \left(\nabla_n u(x,y) \right) \,, \tag{48}$$

for any image $u \in \mathbb{R}^\Omega$. Next Proposition establishes a dual reformulation of (48).

Proposition 10 (dual formulation of $\mathrm{HSTV}_{\alpha,n}$) *Let* $\alpha > 0$ *and* $n \geq 1$. *For any image* $u \in \mathbb{R}^\Omega$, *one has*

$$\mathrm{HSTV}_{\alpha,n}(u) = \max_{p \in \mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}} \langle \tfrac{1}{n^2} \nabla_n u, p \rangle - \delta_{\|\cdot\|_{\infty,2} \leq 1}(p) - \tfrac{\alpha}{2n^2} \|p\|_2^2 \,.$$

The Proof is given in Appendix E. In the following, we shall use the dual formulations of STV_n and $\mathrm{HSTV}_{\alpha,n}$ provided by Propositions 9 and 10 in order to reformulate many optimization problems frequently considered in image restoration in their primal-dual form (43).

5 Image processing applications

In this section, we illustrate the interest of STV in the case of several TV-based image processing applications. As we shall see, replacing the classical discrete TV by STV does not raise any theoretical nor numerical difficulty, and brings clear improvements regarding subpixellic scales.

5.1 Image denoising

The STV variant of the denoising model (1) proposed by Rudin, Osher and Fatemi (ROF) in [62] writes

$$\operatorname*{argmin}_{u\in\mathbb{R}^\Omega} \|u-u_0\|_2^2 + \lambda\,\mathrm{STV}_n(u)\,, \tag{49}$$

where $u_0\in\mathbb{R}^\Omega$ denotes the observed image with (discrete) domain Ω, and $\lambda\geq 0$ is the so-called regularity parameter that controls the trade-off between the data-fidelity term (the square ℓ^2 distance to u_0) and the regularity term $\mathrm{STV}_n(u)$ in the minimization process. Using Proposition 9, we immediately get a primal-dual reformulation of (49),

$$\operatorname*{argmin}_{u\in\mathbb{R}^\Omega}\ \max_{p\in\mathbb{R}^{\Omega_n}\times\mathbb{R}^{\Omega_n}} \|u-u_0\|_2^2 + \langle\tfrac{\lambda}{n^2}\nabla_n u, p\rangle - \delta_{\|\cdot\|_{\infty,2}\leq 1}(p)\,, \tag{50}$$

which has exactly the form of (43) with $(x,y)=(u,p)$, $G(u)=\|u-u_0\|_2^2$, $K=\frac{\lambda}{n^2}\nabla_n$ (with adjoint $K^*=-\frac{\lambda}{n^2}\mathrm{div}_n$), and $F^\star(p)=\delta_{\|\cdot\|_{\infty,2}\leq 1}(p)$.

Notice that replacing $\mathrm{STV}_n(u)$ by $\mathrm{HSTV}_{\alpha,n}(u)$ into (49) leads to the Huber STV_n variant of ROF. In view of Proposition 10, it amounts to replace the term $F^\star(p)=\delta_{\|\cdot\|_{\infty,2}\leq 1}(p)$ by $F^\star(p)=\delta_{\|\cdot\|_{\infty,2}\leq 1}(p)+\frac{\lambda\alpha}{2n^2}\|p\|_2^2$ into the primal-dual problem (50).

For both STV_n and $\mathrm{HSTV}_{\alpha,n}$ regularizers, the corresponding primal-dual problem can be numerically solved by specializing Algorithm 1, which yields Algorithm 2 below. Notice that (39) yields the upper bound $|||K|||\leq\frac{\lambda\pi\sqrt{2}}{n}$, which is useful to set the parameters τ and σ of the algorithm. The images resulting from the different (discrete or Shannon, Huber or usual) TV-based image denoising models are compared in Fig. 4 and 5: we illustrate in Fig. 4 the improved behavior of STV over the classical discrete TV regarding posterior interpolation, and do the same in Fig. 5 for the Huber variant.

Algorithm 2: Chambolle-Pock resolvent Algorithm for Problem (49)

Initialization: *Choose* $\tau,\sigma>0$, $\theta\in[0,1]$, $u^0\in\mathbb{R}^\Omega$, $p^0\in\mathbb{R}^{\Omega_n}\times\mathbb{R}^{\Omega_n}$, *set* $\overline{u}^0=u^0$ *and set* $\nu=1$ *when using the* STV_n *regularizer and* $\nu=1+\sigma\frac{\alpha\lambda}{n^2}$ *when using the* $\mathrm{HSTV}_{\alpha,n}$ *regularizer. Define the projection operator* $\pi_{\infty,2}$ *by*

$$\forall(x,y)\in\Omega_n,\quad \pi_{\infty,2}(p)(x,y)=\frac{p(x,y)}{\max\left(1,|p(x,y)|\right)}\,,$$

for any $p\in\mathbb{R}^{\Omega_n}\times\mathbb{R}^{\Omega_n}$.

Iterations: *For* $k\geq 0$, *update* p^k, u^k *and* $\overline{u}^k$ *with*

$$p^{k+1}=\pi_{\infty,2}\left(\left(p^k+\tfrac{\sigma\lambda}{n^2}\nabla_n\overline{u}^k\right)/\nu\right) \tag{51a}$$

$$u^{k+1}=\frac{u^k+\frac{\tau\lambda}{n^2}\mathrm{div}_n p^{k+1}+2\tau u_0}{1+2\tau} \tag{51b}$$

$$\overline{u}^{k+1}=u^{k+1}+\theta\left(u^{k+1}-u^k\right) \tag{51c}$$

5.2 Inverse problems

Let us now consider the more general case of a linear inverse problem addressed with quadratic data fidelity and STV regularization. It writes

$$\tilde{u}\in\operatorname*{argmin}_{u\in\mathbb{R}^\Omega} \|Au-u_0\|_2^2+\lambda\,\mathrm{STV}_n(u)\,, \tag{52}$$

where $u_0\in\mathbb{R}^\omega$ denotes the observed image (ω being a finite subset of $\mathbb{Z}^2$, possibly $\omega=\Omega$) and $A:\mathbb{R}^\Omega\to\mathbb{R}^\omega$ is a linear operator which may model the convolution with the impulse response of an acquisition device (defocus or motion blur for instance) or other linear observation mechanisms such as tomography, downsampling, loss of image regions, etc.

Proposition 11 (primal-dual formulation of (52)) *Any solution* $\tilde{u}$ *of Problem* (52) *satisfies*

$$\tilde{u}\in\operatorname*{argmin}_{u\in\mathbb{R}^\Omega}\ \max_{\substack{p\in\mathbb{R}^{\Omega_n}\times\mathbb{R}^{\Omega_n}\\ q\in\mathbb{R}^\omega}} G(u)+\langle Ku,(p,q)\rangle-F^\star(p,q)\,,$$

where $G(u)=0$, $F^\star(p,q)=\delta_{\|\cdot\|_{\infty,2}\leq 1}(p)+\|\frac{q}{2}+u_0\|_2^2$ *and* $K:\mathbb{R}^\Omega\to\left(\mathbb{R}^{\Omega_n}\times\mathbb{R}^{\Omega_n}\right)\times\mathbb{R}^\omega$ *is the linear operator defined by* $Ku=\left(\frac{\lambda}{n^2}\nabla_n u, Au\right)$ *for any* $u\in\mathbb{R}^\Omega$.

Proof Writing $f(v)=\|v-u_0\|_2^2$, one easily gets the expression of the Legendre-Fenchel transform of f, that is $f^\star(q)=\|\frac{q}{2}+u_0\|_2^2-\|u_0\|_2^2$. Besides, since $f\in\Gamma_0(\mathbb{R}^\omega)$, we have

$$\begin{aligned}\|Au-u_0\|_2^2 &= f(Au)=f^{\star\star}(Au)\\ &=\sup_{q\in\mathbb{R}^\omega}\langle Au,q\rangle-\|\tfrac{q}{2}+u_0\|_2^2+\|u_0\|_2^2\,,\end{aligned} \tag{53}$$

and the supremum is attained since the cost functional is concave, differentiable, and its gradient vanishes at the point $q=2(Au-u_0)$. Replacing the quadratic term accordingly into (52), removing the constant $\|u_0\|_2^2$ (which does not change the set of minimizers), and replacing as well the STV_n term by its dual formulation using Proposition 9, we obtain the desired result. □

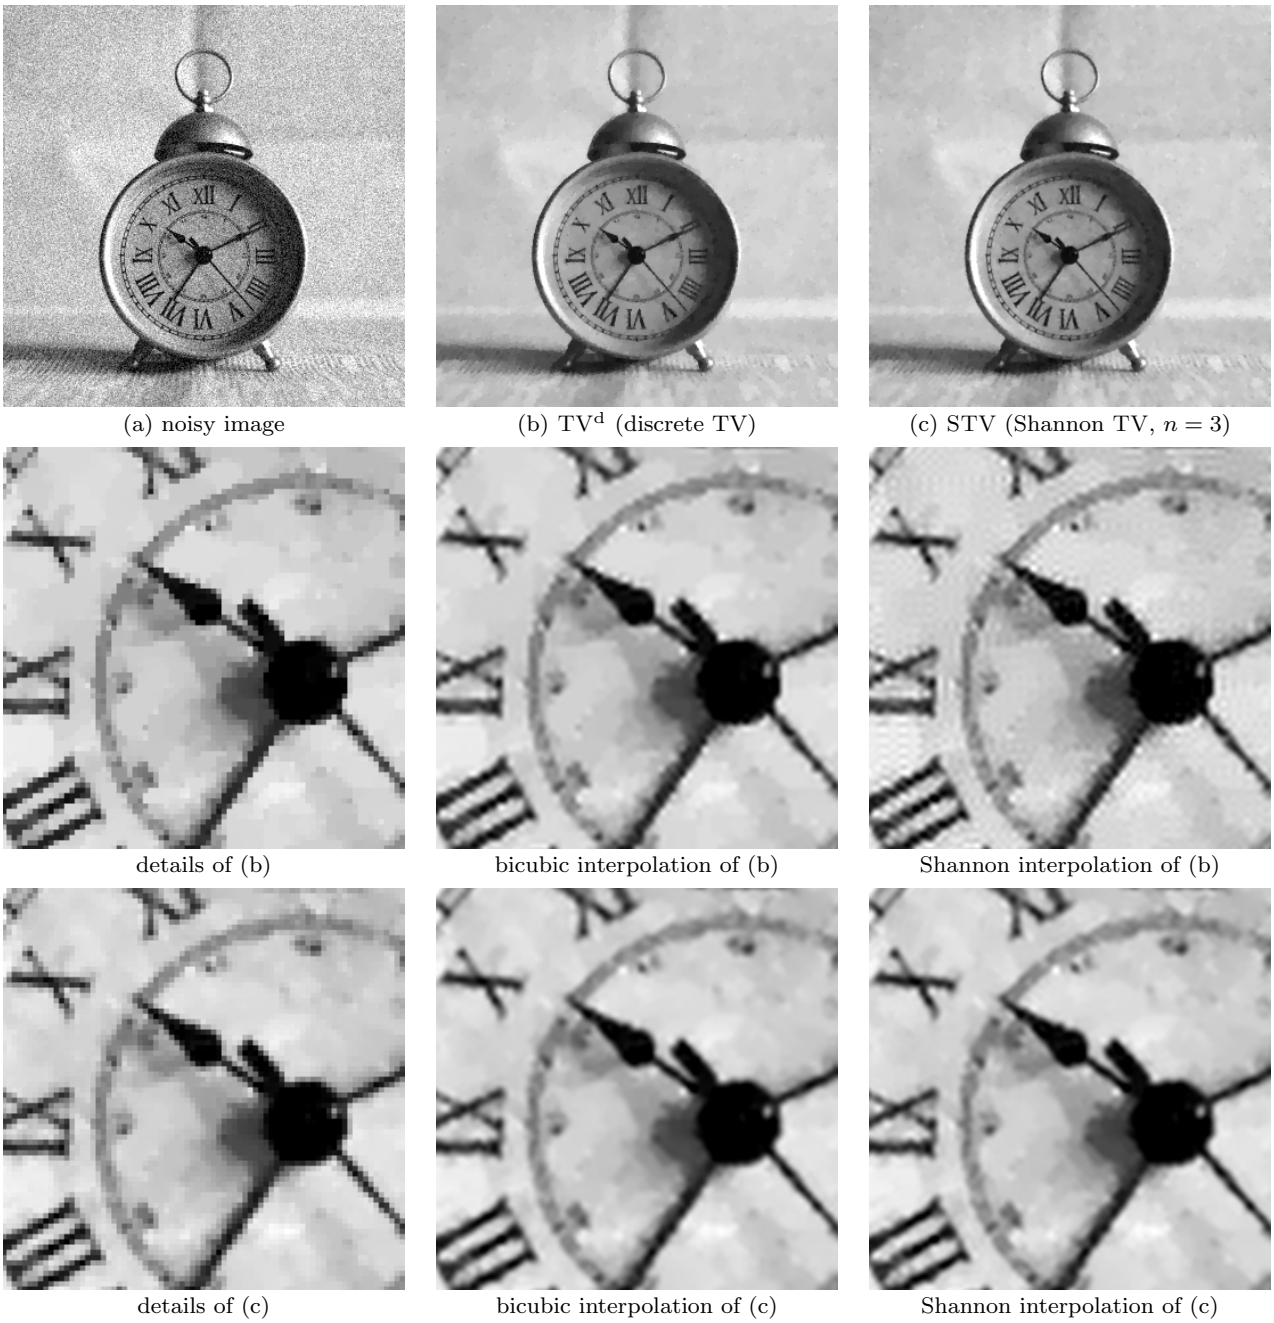

(a) noisy image (b) TV^{d} (discrete TV) (c) STV (Shannon TV, $n = 3$)

details of (b) bicubic interpolation of (b) Shannon interpolation of (b)

details of (c) bicubic interpolation of (c) Shannon interpolation of (c)

Fig. 4 Comparison of discrete TV and Shannon TV for image denoising. A noisy image (top, left) undergoing additive white Gaussian noise with zero mean and standard deviation $\sigma = 20$ (see also the reference image in Fig. 5) was processed with the ROF model using the TV^{d} (top, center) and STV_3 (top, right) discretizations. The regularity parameter λ was set in order to get the same norm of the estimated noise (the difference between the noisy and the restored image) in each simulation. In the second row we display a cropping of the TV^{d}-restored image oversampled with factor 3 using different interpolation methods (from left to right: nearest neighbor, bicubic spline and Shannon interpolation). In the third row, the same operation is realized on the STV-restored image. We can see that images TV^{d} and STV images look globally similar. The details on the left of rows 2 and 3 reveal the presence of staircasing in both cases, but this artifact is significantly attenuated in the case of STV. Looking at the second row, we see that the TV^{d} image cannot be interpolated in a satisfying way, since both bicubic and Shannon interpolation methods yield images with undesirables oscillations (*ringing*) localized near objects contours. This is not the case with the STV image, that can be interpolated without creating artifacts with both bicubic and Shannon interpolations (row 3).

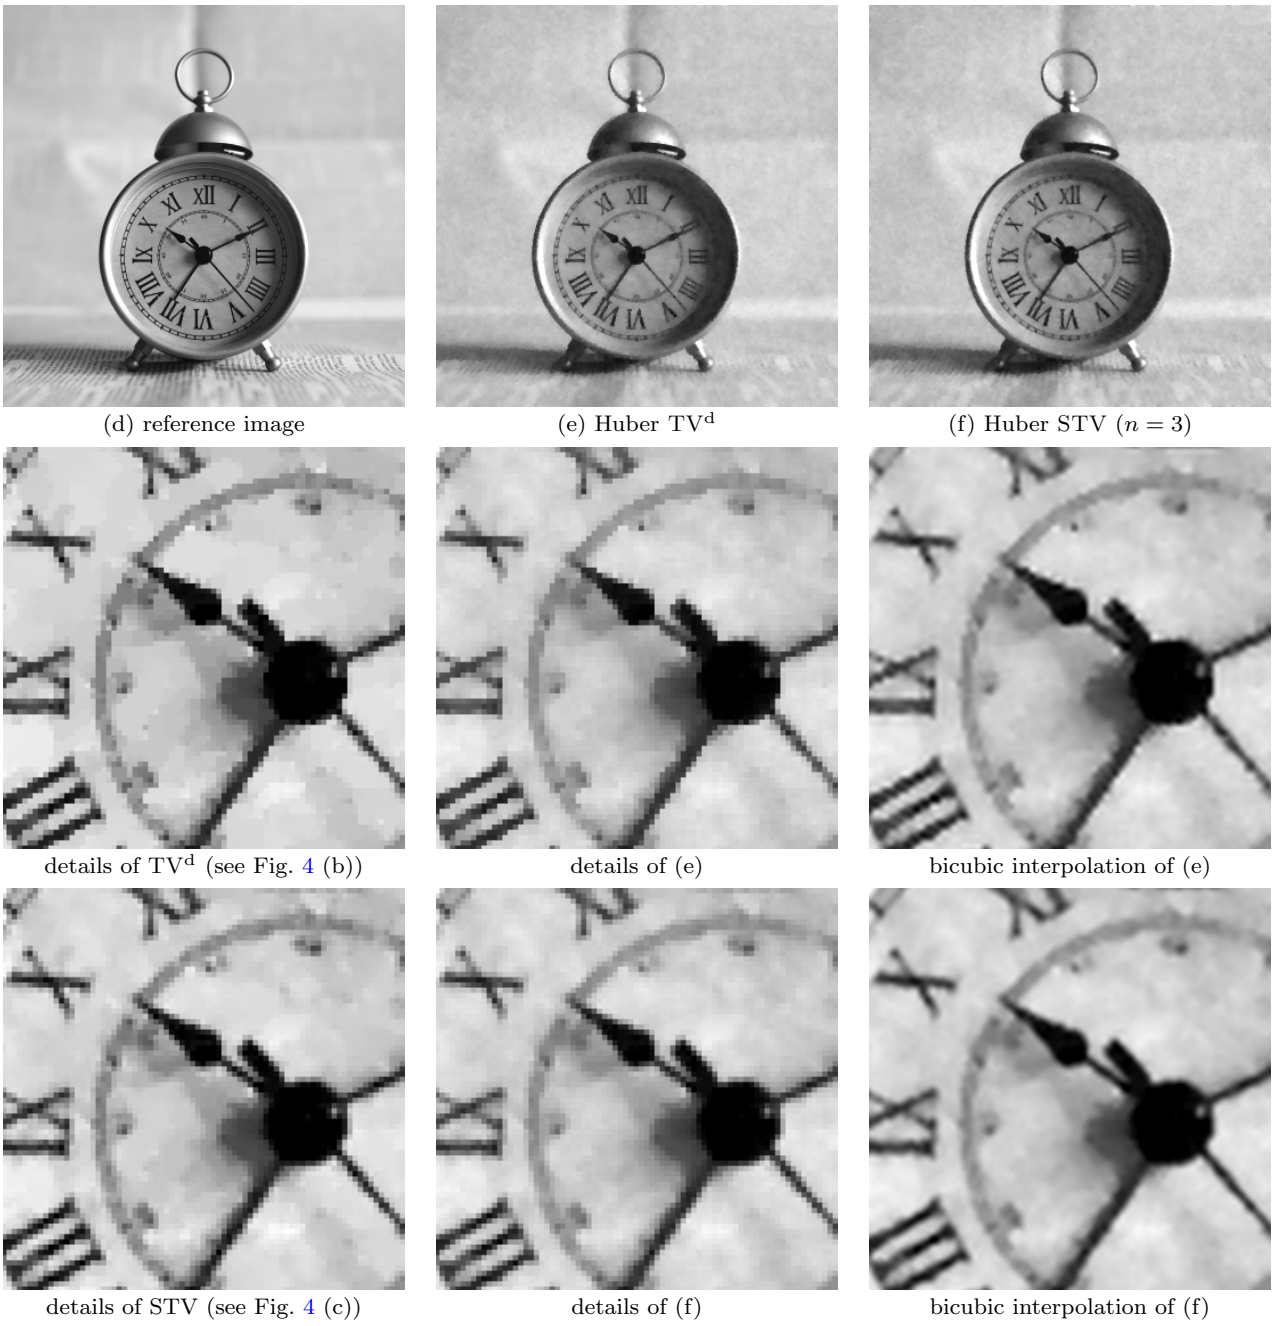

(d) reference image (e) Huber TV^d (f) Huber STV ($n = 3$)

details of TV^d (see Fig. 4 (b)) details of (e) bicubic interpolation of (e)

details of STV (see Fig. 4 (c)) details of (f) bicubic interpolation of (f)

Fig. 5 Image denoising with Huber-TV and Huber-STV. This experiment is similar to Fig. 4, except that we here consider the Huber variant (with $\alpha = 5$) of ROF denoising, both for the TV^d and STV discretizations. As expected, the Huber variant avoids the staircasing effect for both discretizations (TV^d and STV). However, it does not solve the interpolability issue for TV^d: the bicubic interpolation of Huber TV^d presents several ringing artifacts (like the non-Huber TV^d displayed in Fig. 4), and these artifacts are again completely avoided by considering the STV discretization. The same phenomenon is observed on the Shannon interpolates of (e) and (f) (not displayed here).

Again, the Huber version of (52) is obtained by replacing the $\mathrm{STV}_n(u)$ term by $\mathrm{HSTV}_{\alpha,n}(u)$, which simply changes $F^\star(p,q) = \delta_{\|\cdot\|_{\infty,2}\leq 1}(p) + \|\frac{q}{2} + u_0\|_2^2$ into $F^\star(p,q) = \delta_{\|\cdot\|_{\infty,2}\leq 1}(p) + \frac{\alpha\lambda}{2n^2}\|p\|_2^2 + \|\frac{q}{2} + u_0\|_2^2$.

Note that the adjoint of K (defined in Proposition 11) is $K^*(p,q) = -\frac{\lambda}{n^2}\mathrm{div}_n p + A^*q$, and its induced ℓ^2 norm satisfies

$$|||K|||^2 \leq |||\tfrac{\lambda}{n^2}\nabla_n|||^2 + |||A|||^2 \leq 2\left(\tfrac{\pi\lambda}{n}\right)^2 + |||A|||^2 .$$

Thus, Chambolle-Pock Algorithm can be rewritten in the present case as Algorithm 3 below. The update of the dual variable (here the tuple (p,q)) in the generic Algorithm 1 was split into two independent updates thanks to the additive separability with respect to p and q of the function $(p,q) \mapsto \langle Ku, (p,q)\rangle - F^\star(p,q)$.

Algorithm 3: Chambolle-Pock resolvent Algorithm for Problem (52)

Initialization: *Choose* $\tau, \sigma > 0$, $\theta \in [0,1]$, $u^0 \in \mathbb{R}^\Omega$, $p^0 \in \mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}$, $q^0 \in \mathbb{R}^\omega$, *set* $\overline{u}^0 = u^0$ *and set* ν *as in Algorithm 2.*

Iterations: *For* $k \geq 0$, *update* p^k, u^k *and* $\overline{u}^k$ *with*

$$p^{k+1} = \pi_{\infty,2}\left(\left(p^k + \tfrac{\sigma\lambda}{n^2}\nabla_n \overline{u}^k\right)/\nu\right)$$
$$q^{k+1} = \frac{2\,q^k + 2\sigma\left(A\overline{u}^k - u_0\right)}{2+\sigma}$$
$$u^{k+1} = u^k + \tfrac{\tau\lambda}{n^2}\mathrm{div}_n p^{k+1} - \tau A^* q^{k+1}$$
$$\overline{u}^{k+1} = u^{k+1} + \theta\left(u^{k+1} - u^k\right)$$

5.2.1 Application to image deconvolution

In the case of image deconvolution, the linear operator A in (52) is the convolution with a point spread function k_A (modeling for instance some blurring phenomenon such as diffraction, defocus, motion blur, ...). Let us consider such a discrete convolution kernel $k_A \in \mathbb{R}^{\omega_A}$ with finite domain $\omega_A \subset \mathbb{Z}^2$, and define the associated operator $A : \mathbb{R}^\Omega \to \mathbb{R}^\omega$ by

$$Au(x,y) = \sum_{(a,b)\in\omega_A} k_A(a,b)\, u(x-a, y-b)\ , \tag{55}$$

where ω denotes the subset of Ω made of all the pixels $(x,y) \in \Omega$ such as $(x,y) - \omega_A \subset \Omega$. In order to use Algorithm 3, we need the explicit expression of $A^* : \mathbb{R}^\omega \to \mathbb{R}^\Omega$, which writes

$$A^*v(x,y) = \sum_{(a,b)\in\omega_A} k_A(a,b)\, v(x+a, y+b)\ , \tag{56}$$

for $v \in \mathbb{R}^\omega$ and $(x,y) \in \Omega$, with the convention that $v(x+a, y+b) = 0$ when $(x+a, y+b) \notin \omega$. One easily checks that $|||A||| \leq ||k_A||_1$ as well.

Most authors define the convolution with kernel k_A as an operator $A : \mathbb{R}^\Omega \to \mathbb{R}^\Omega$ at the cost of an extension of u outside of Ω, usually a periodic or a mirroring condition, or a zero-extension. Such a convention simplifies the analysis (and the computations, especially in the periodic case where the convolution can be implemented with the DFT), but we shall not use it here as it is unrealistic and thus of little help to process real data. Experiments illustrating STV deblurring are displayed in Fig. 6 (motion blur) and 7 (out of focus).

5.2.2 Application to image zooming and inpainting

The variational formulation (52) can be used to perform many other image processing tasks: as soon as we can derive a closed-form expression for A, its adjoint A^*, and estimate an upper bound for $|||A|||$, Algorithm (3) can be implemented without difficulty. We here mention two more examples of applications (zoom and inpainting), each corresponding to a particular choice of A. We experimentally checked that, in both cases, the use of STV_n instead of TV yields nicely interpolable images.

In the case of image zooming, the operator A is often assumed to be a blurring kernel followed by a subsampling procedure (see [44,17]). A simple particular case is the discrete captor integration model $A : \mathbb{R}^\Omega \to \mathbb{R}^\omega$ defined by

$$Au(x,y) = \frac{1}{\delta^2}\sum_{(a,b)\in I_\delta^2} u(\delta x + a, \delta y + b)\ , \tag{57}$$

where $\omega = I_M \times I_N$ denotes a small discrete domain and $\Omega = I_{\delta M} \times I_{\delta N}$ a bigger one, δ (the magnification factor) being an integer at least equal to 2. In that case, we easily obtain the relation $|||A||| = \frac{1}{\delta}$ and the expression of the adjoint operator $A^* : \mathbb{R}^\omega \to \mathbb{R}^\Omega$ as

$$A^*v(x,y) = \frac{1}{\delta^2}\, v\left(\lfloor\tfrac{x}{\delta}\rfloor, \lfloor\tfrac{y}{\delta}\rfloor\right)\ . \tag{58}$$

Another example is image inpainting, which aims at estimating plausible image intensities in a (nonempty) subpart ω_0 of the image domain Ω where the information is missing. In that case, $\omega = \Omega$, the operator $A : \mathbb{R}^\Omega \to \mathbb{R}^\Omega$ is defined by

$$Au(x,y) = \mathbb{1}_{\omega_0}(x,y)\cdot u(x,y)\,,$$

and one easily checks that $A^* = A$ (A is a diagonal operator) and $|||A||| = 1$.

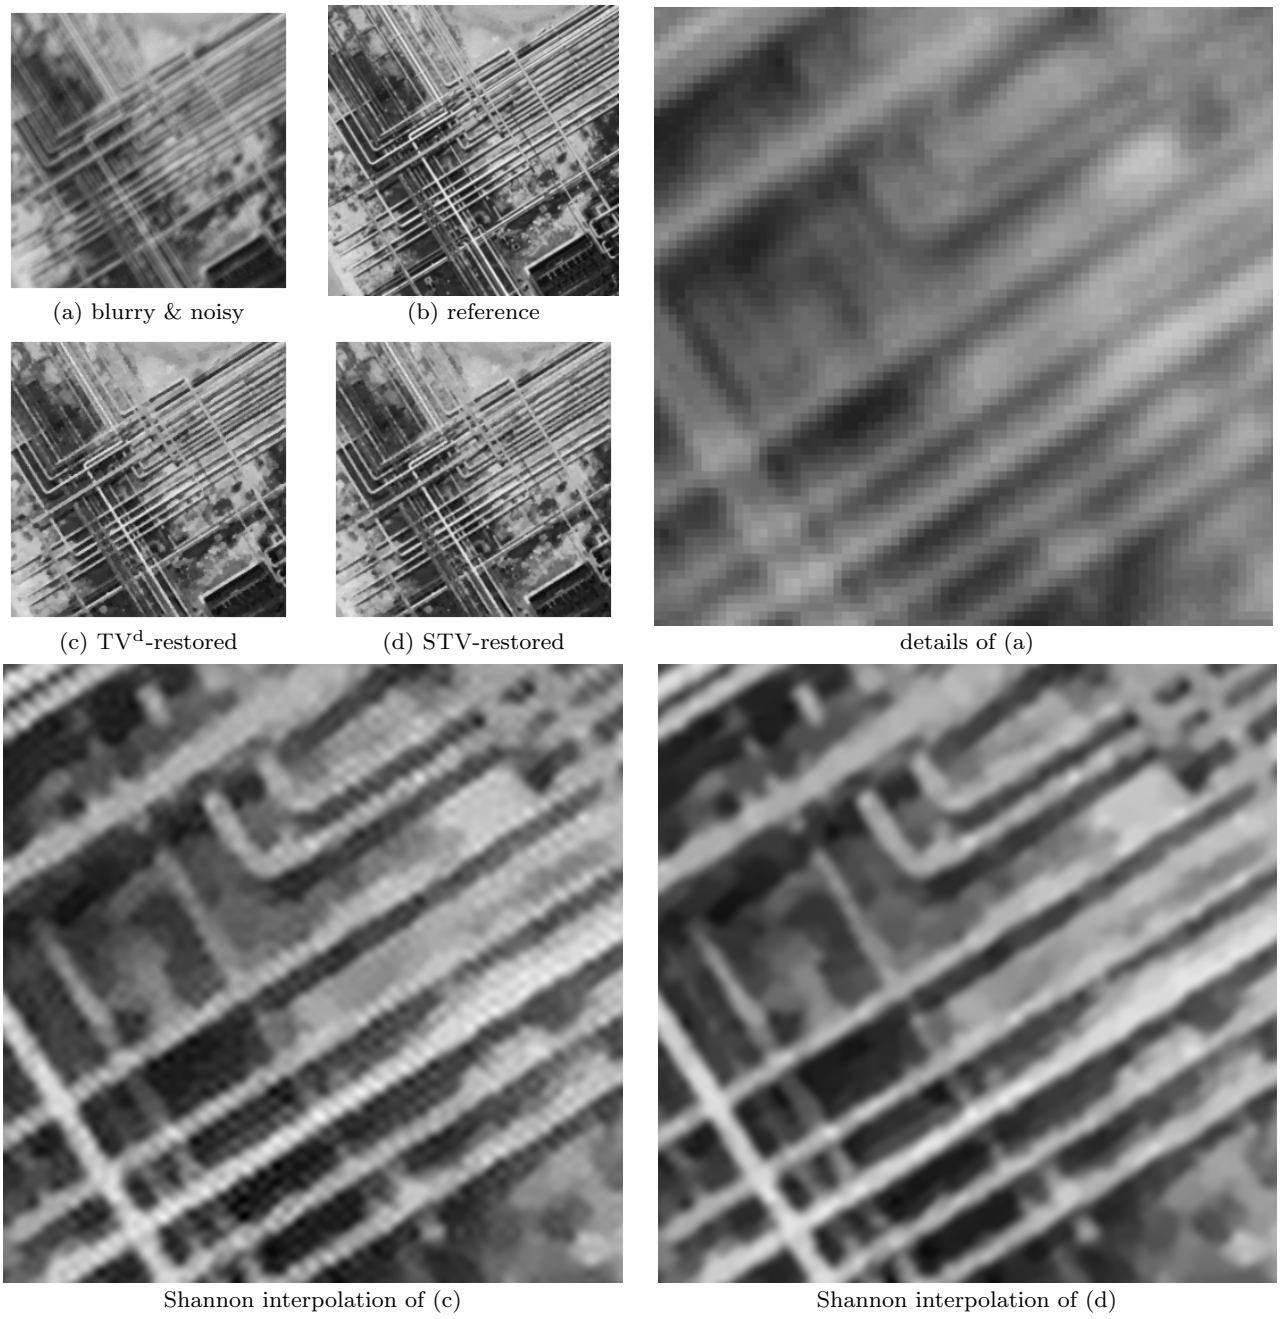

(a) blurry & noisy (b) reference details of (a)

(c) TV^d-restored (d) STV-restored

Shannon interpolation of (c) Shannon interpolation of (d)

Fig. 6 Motion deblurring with discrete TV and Shannon TV. A degraded (blurry and noisy) image (a) is synthesized by convolving the reference image (b) with a real-data motion blur kernel and then adding a white Gaussian noise with zero-mean and standard deviation $\sigma = 2$. The degraded image (a) is then processed by solving the corresponding TV^d and STV_3 regularized inverse problems (Equation (52)). As in Fig. 4, the regularization parameter λ is set in such a way that the amount of estimated noise (here the quantity $\|A\tilde{u} - u_0\|_2$, where $\tilde{u}$ is the restored image) is the same for both methods. The resulting images (c) and (d) are quite similar, but the magnified views on the bottom row (magnification of factor 4 with Shannon interpolation) clearly show that they strongly differ in terms of interpolability: as in the denoising case, the interpolated TV^d image exhibits strong ringing artifacts, whereas the interpolated STV image does not.

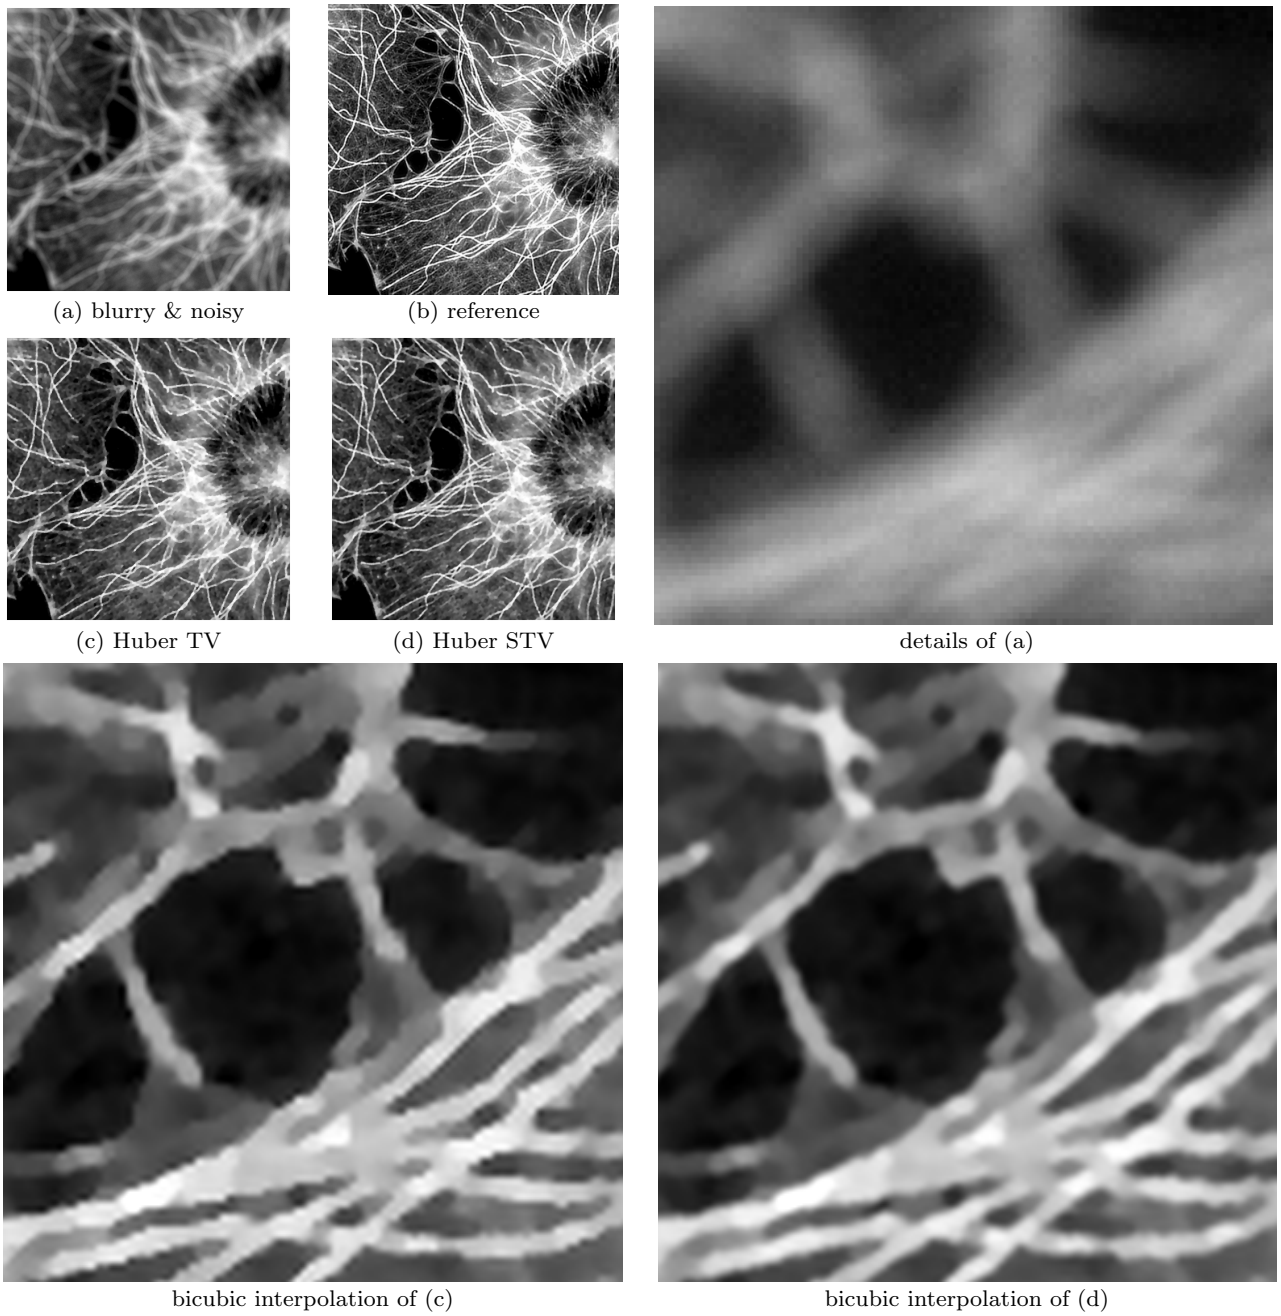

(a) blurry & noisy (b) reference details of (a)

(c) Huber TV (d) Huber STV

bicubic interpolation of (c) bicubic interpolation of (d)

Fig. 7 Out-of-focus deblurring using Huber TV and Huber STV. This experiment is similar to Fig. 6, except that we here used a fluorescence microscopy image of actin filaments and microtubules in interphase cells (source cellimagelibrary.org, CIL number 240, first channel), a synthetic out-of-focus blur kernel defined by the indicator of a disk with radius 7 pixels, and we replaced the TV^{d} and STV_3 regularizers by their Huber versions ($\alpha = 5$). The conclusions are identical.

5.3 Constrained minimization

In some situations, it is desirable to consider constrained minimization problems of the type

$$\tilde{u} \in \operatorname*{argmin}_{u\in\mathbb{R}^\Omega} \mathrm{STV}_n(u) \quad \text{subject to} \quad Au = u_0\,, \tag{59}$$

where u_0 denotes the observed image with discrete domain ω, $\tilde{u}$ denotes the reconstructed image with discrete domain Ω, and A denotes again a linear operator from $\mathbb{R}^\Omega$ to $\mathbb{R}^\omega$. In other words, we are interested in the computation of an image $\tilde{u}$ having the smallest Shannon TV among those satisfying the constraint $Au = u_0$. Remark that the inverse problem (52) is none other than a relaxed version of (59). In the presence of noise, it is better to use the relaxed formulation, but the constrained model (59) may be interesting when the level of noise in u_0 is low, especially because it does not require the setting of any regularization parameter λ.

Using Proposition 9, we obtain a primal-dual reformulation of (59),

$$\tilde{u} \in \operatorname*{argmin}_{u\in\mathbb{R}^\Omega} \max_{p\in\mathbb{R}^{\Omega_n}\times\mathbb{R}^{\Omega_n}} \delta_{A^{-1}(u_0)}(u) + \langle \tfrac{1}{n^2}\nabla_n u, p\rangle - \delta_{\|\cdot\|_{\infty,2}\leq 1}(p)\,, \tag{60}$$

where the (closed and convex) set

$$A^{-1}(u_0) := \left\{u \in \mathbb{R}^\Omega,\ Au = u_0\right\}$$

is assumed to be nonempty, and $\delta_{\mathcal{P}}$ denotes the indicator function of a set $\mathcal{P}$ (that is, $\delta_{\mathcal{P}}(p) = 0$ if $p \in \mathcal{P}$, $+\infty$ otherwise). A solution of Problem (60) can be numerically computed using Algorithm 4, taking $G = \delta_{A^{-1}(u_0)}$, $F^\star = \delta_{\|\cdot\|_{\infty,2}\leq 1}$ and $K = \frac{1}{n^2}\nabla_n$ in Chambolle-Pock Algorithm.

Algorithm 4: Chambolle-Pock resolvent Algorithm for Problem (60)

Initialization: *Choose* $\tau, \sigma > 0$, $\theta \in [0,1]$, $u^0 \in \mathbb{R}^\Omega$, $p^0 \in \mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}$, *set* $\overline{u}^0 = u^0$ *and define* ν *and* $\pi_{\infty,2}$ *as in Algorithm 2. Denote by* π_0 *the* ℓ^2 *projection from* $\mathbb{R}^\Omega$ *onto the (closed and convex) set* $A^{-1}(u_0) = \left\{u \in \mathbb{R}^\Omega,\ Au = u_0\right\}$.

Iterations: *For* $k \geq 0$, *update* p^k, u^k *and* $\overline{u}^k$ *with*

$$p^{k+1} = \pi_{\infty,2}\left(\left(p^k + \tfrac{\sigma}{n^2}\nabla_n \overline{u}^k\right)/\nu\right)$$
$$u^{k+1} = \pi_0\left(u^k + \tfrac{\tau}{n^2}\mathrm{div}_n p^{k+1}\right)$$
$$\overline{u}^{k+1} = u^{k+1} + \theta\left(u^{k+1} - u^k\right)$$

To illustrate the general framework above, we will consider in the next section the problem of reconstructing an image from partial measurements in the Fourier domain. A particular case is image magnification (assuming that the original low-resolution image does not suffer from aliasing), which corresponds to the recovery of high-frequency components only, but other situations (like tomography) require spectrum interpolation in a more complicated domain. Note also that many other applications, such as image inpainting or image zooming presented in Section 5.2.2, can be easily handled as well with the constrained formulation (59).

5.3.1 Application to spectrum extrapolation

Given an image $u_0 \in \mathbb{R}^\Omega$ whose spectrum $\widehat{u}_0$ is known on a certain (symmetric) subdomain $\widehat{\omega}_0$ of $\widehat{\Omega}$, how to extend this spectrum to the whole spectral domain $\widehat{\Omega}$? The trivial *zero-padding* approach, which amounts to extending the spectrum with the constant zero, yields a very oscillatory image in general, in reason of the irregularity (missing Fourier coefficients) of the extrapolated spectrum. A more satisfying reconstruction can be obtained with a variational approach: among all possible spectrum extensions, choose the one that minimizes a given energy. This kind of approach was used by Rougé and Seghier [60], who considered the Burg entropy, and by Guichard and Malgouyres [36,44], who used the discrete TV (but in a slightly different framework, since they take as input a subsampled image which suffers from aliasing). We here consider the energy STV_n; in a constrained formulation, this is a particular case of (59), since the frequency constraint ($\widehat{u}$ and $\widehat{u}_0$ are equal on $\widehat{\omega}_0$) can be enforced under the form $Au = u_0$ where $A = \mathcal{F}^{-1} \circ M_{\widehat{\omega}_0} \circ \mathcal{F}$ ($\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the direct and inverse discrete Fourier transforms respectively, the operator $M_{\widehat{\omega}_0}$ denotes the pointwise multiplication of a element of $\mathbb{C}^{\widehat{\Omega}}$ with $\mathbb{1}_{\widehat{\omega}_0}$, and $\widehat{u}_0$ is implicitly set to zero outside $\widehat{\omega}_0$). Note that the ℓ^2 projection π_0 onto the set $A^{-1}(u_0)$ is simply obtained in the Fourier domain with

$$\forall u \in \mathbb{R}^\Omega,\ \forall (\alpha,\beta) \in \widehat{\Omega},\quad \widehat{\pi_0(u)}(\alpha,\beta) = \begin{cases} \widehat{u}_0(\alpha,\beta) \text{ if } (\alpha,\beta) \in \widehat{\omega}_0 \\ \widehat{u}(\alpha,\beta) \text{ otherwise}\,. \end{cases}$$

Some examples of spectrum extrapolations are proposed in Fig. 8 and 9.

5.3.2 Link with Compressed Sensing

The theory of compressed sensing [11,26] brings an interesting perspective on the role and interest of TV in image processing: under certain conditions, cartoon images (that is, discrete images whose discrete gradient is sparse) can be reconstructed *exactly* from a limited number of linear measurements using TV mini-

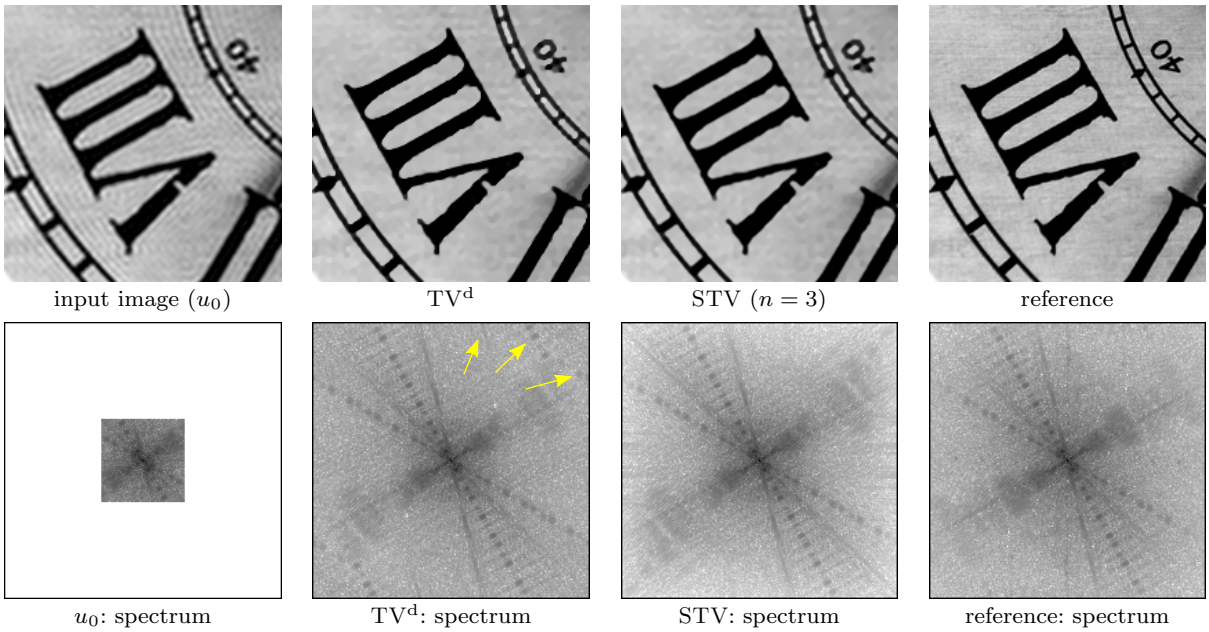

Fig. 8 Image zooming with spectrum extrapolation. An input image (1st column) is synthesized by setting to 0 the high frequency components (that is, outside a square $\widehat{\omega}_0$) of a reference image (4th column). Spectrum extrapolation is then realized using either the discrete TV (2nd column) or the STV (3rd column). For each image of the first row, the spectrum (Fourier modulus, in log scale) is displayed below on the second row. As we can observe, the constrained TV minimization framework (59) is efficient for spectrum extrapolation: both discretizations manage to reconstruct part of the missing high frequencies and remove the ringing patterns observed in the input image. However, STV is to be preferred to discrete TV as it manages to avoid the aliasing artifacts of the latter (red arrows), and delivers nicely interpolable images.

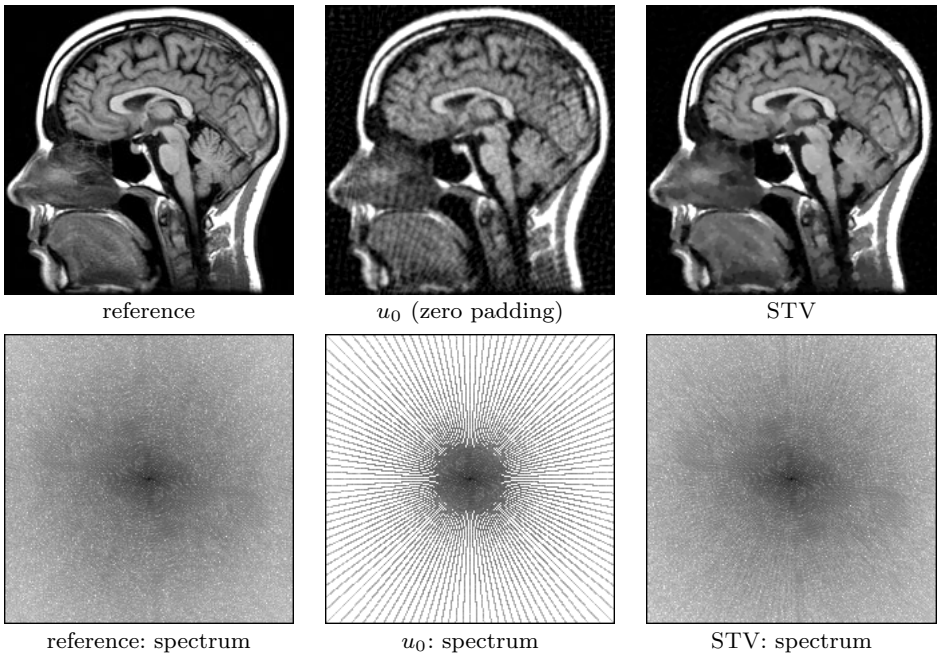

Fig. 9 Image reconstruction from partial measurements in the Fourier domain. We here reproduce a simplified tomography inversion experiment: a reference image (1st column) is sampled in the Fourier domain along several discrete rays (covering around 35% of the whole frequency domain), and two image reconstruction methods are compared. The first one consists in setting the missing Fourier coefficients to 0 (2nd column), which produces severe ringing artifacts. Extrapolating the missing Fourier coefficients with the constrained STV minimization framework (59) yields a much nicer image (3rd column) which can be easily interpolated. As in Fig. 8, the spectrum of each image of the first row is displayed on the second row.

mization. Indeed, the problem of spectrum extrapolation we just described is a fundamental experiment of the Compressed Sensing theory. Candès, Romberg and Tao remarked in [11] that a discrete sampled image of the Shepp-Logan phantom (a cartoon synthetic analytical image made by the addition of 10 ellipses) could be exactly reconstructed from a limited number of its DFT coefficients by minimizing the discrete TV of the latent image, constrained with the known DFT coefficients (see Fig. 10 (b)). The theory was simultaneously developed by Donoho [26], and continuously refined later, in particular to extend the initial probabilistic proofs to the deterministic case [25,39]. Also note the works on simple model sets/quasicrystals by Matei and Meyer [47,46], which establish links between Compressed Sensing and irregular sampling. Please consider [29] and the references therein for more details about the Compressed Sensing theory.

In Fig. 10, we reproduce a classical experiment of the Compressed Sensing theory by trying to reconstruct the whole *Shepp-Logan* phantom image from a limited number of its discrete Fourier coefficients. This spectrum extrapolation problem is solved using TV minimization (either with TV^{d} or the STV variant), as described in Section 5.3.1. Contrary to TV^{d}, which leads to a perfect reconstruction (Fig. 10 (b)), the STV variant yields a distorted approximation (Fig. 10 (c)). The reason for this seeming failure is simple: since the reference image is strongly aliased (as a sampled version of a piecewise constant image), the use of the Shannon interpolation (and hence STV minimization) is particularly inadequate. Notice incidentally that this experiment (conceived to mimic the reconstruction framework in Magnetic Resonance Imaging) is not realistic, because in the physical world one would not be able to measure Fourier coefficients of the discrete (aliased) phantom image, but only Fourier coefficients of the continuous phantom. The difference is clearly analyzed in [35]: *"Counterintuitively, the reconstructions out of rasterized simulations lead to aliasing effects that have a positive impact on visual quality. This situation, which occurs when the same model is used for both simulation and reconstruction, is sometimes referred to as "inverse crime." It arises because of the artificially imposed consistency between the computational forward models used for simulation and reconstruction. In such an inverse-crime situation, the continuous nature of the underlying physical model is not taken into account."*

Is it nonetheless possible to obtain results similar to compressed sensing on interpolable images? To that aim, we have to give up strict sparsity and consider smooth versions of cartoon images (that is, the kind of bandlimited image that could be observed in a realistic optical system undergoing diffraction). In Fig. 11, we used a smooth and non-aliased version of Shepp-Logan phantom (obtained by keeping the low frequency components of a 10-fold higher resolution image smoothed by a Gaussian kernel). With this reference image, which is well-suited to Shannon (and spline) interpolation, and with acquisition conditions similar to those of Fig. 10 (same measured Fourier coefficients), STV regularization leads to a more faithful reconstruction than TV^{d}.

Thus, we can see that even in a compressed sensing framework (for which STV is a priori less suited because the Shannon interpolate of a non-zero discrete image cannot have a sparse gradient), the use of STV remains interesting, in particular if the output image is to be interpolated. Another example of this is the experiment of Fig. 9, for which the TV^{d} reconstruction (not shown) and the STV reconstruction are visually similar, but only the STV image can be well interpolated.

6 Regularization with weighted frequencies

Using STV as a regularizer leads to iterative algorithms that operate in the Fourier domain. This has a non-negligible computational cost, even though this kind of algorithms is common nowadays and there exist very efficient implementations of the Fourier Transform, like FFTW [32]. We now consider an image restoration model that benefits from the availability of the Fourier transform of the current image at each iteration.

6.1 Model

Let $u_0 : \Omega \to \mathbb{R}$ be an input image (with $\Omega = I_M \times I_N$) and $\gamma : \widehat{\Omega} \to \mathbb{R}_+$ a non-negative map that is symmetric (in the sense that $\widehat{\gamma}$ is real-valued). We consider the minimization problem

$$\operatorname*{argmin}_{u \in \mathbb{R}^{\Omega}} \|\widehat{u} - \widehat{u_0}\|_{\gamma}^2 + \lambda\, \mathrm{STV}_n(u)\,, \tag{62}$$

where $\lambda > 0$ is a regularization parameter and

$$\|\widehat{u} - \widehat{u_0}\|_{\gamma}^2 = \frac{1}{|\Omega|} \sum_{(\alpha,\beta)\in\widehat{\Omega}} \gamma(\alpha,\beta) \cdot |\widehat{u}(\alpha,\beta) - \widehat{u_0}(\alpha,\beta)|^2$$

is a weighted squared distance between u and u_0 (strictly speaking, it defines a distance only if γ does not vanish). Model (62) generalizes two other models considered above. Indeed, STV image denoising (49) is obtained with $\gamma \equiv 1$, while the choice $\gamma = \mathbb{1}_{\widehat{\omega_0}}$ leads

(a) frequency mask (b) reconstruction with TV^{d} (c) reconstruction with STV

Fig. 10 Compressed sensing under "inverse crime" conditions. We here simulate a compressed-sensing acquisition of the synthetic cartoon *Shepp-Logan* image (sampled on a 187×187 grid) by observing its DFT coefficients along 20 discrete lines (a) that cover only 12% of the full frequency domain. The image is then reconstructed using constrained TV^{d} (b) or STV (c) minimization, as described in Section 5.3.1. The reconstruction obtained using TV^{d} is indistinguishable from the reference image (not displayed here), whereas the reconstruction obtained using STV is a poor approximation. This illustrates the inability of STV to reconstruct cartoon (and thus severely aliased) images in an "inverse crime" (unrealistic) compressed sensing framework (see discussion in Section 5.3.2).

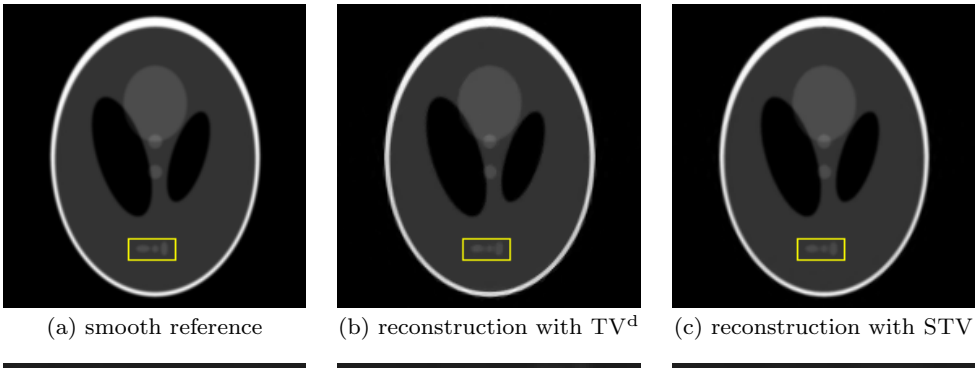

(a) smooth reference (b) reconstruction with TV^{d} (c) reconstruction with STV

resampling of (a) resampling of (b) resampling of (c)

Fig. 11 Compressed sensing under realistic conditions. We here consider a smooth reference image (a), obtained by keeping the 187×187 low-frequency content of a 1870×1870 image of the analytical Shepp-Logan phantom smoothed with a Gaussian filter (with standard deviation $\sigma = 7$). This setup is more realistic than the "inverse crime" situation of Fig. 10, since the DFT coefficients of the reference image are now much more faithful to the Fourier coefficients of the underlying continuous image. We reproduced the same experiments as in Fig. 10 with this smooth reference image, and obtained the corresponding TV^{d} (b) and STV (c) reconstructed image. We can see that STV regularization now clearly outperforms the discrete TV: the reconstructed discrete image (c) is more faithful to the reference than (b) (visually, and in terms of PSNR: 54.4 dB with STV, versus 37.5 dB with TV^{d}), and the comparison of the interpolated STV and TV^{d} images (second row) shows that the STV result is geometrically more precise at a sub-pixel scale.

to a relaxed version of spectrum extrapolation considered in Section 5.3.1. Choosing a more general (non-binary) weight map γ provides a way to selectively regularize the Fourier coefficients of the input image u_0: when $\gamma(\alpha, \beta)$ is large, one expects to obtain $\widehat{u}(\alpha, \beta) \approx \widehat{u_0}(\alpha, \beta)$; on the contrary, the coefficients $\widehat{u}(\alpha, \beta)$ corresponding to small (or zero) values of $\gamma(\alpha, \beta)$ are essentially driven by STV regularization.

6.2 Algorithm

Replacing the STV_n term by its dual formulation (Proposition 9) into (62) yields the primal-dual problem

$$\underset{u\in\mathbb{R}^{\Omega}}{\operatorname{argmin}}\ \max_{p\in\mathbb{R}^{\Omega_n}\times\mathbb{R}^{\Omega_n}} \|\widehat{u}-\widehat{u_0}\|_\gamma^2 + \langle \tfrac{\lambda}{n^2}\nabla_n u, p\rangle - \delta_{\|\cdot\|_{\infty,2}\le 1}(p)\,. \quad (63)$$

In order to apply Algorithm 1 to (63), one needs to perform at each iteration k the primal update

$$u^{k+1} = \underset{u\in\mathbb{R}^{\Omega}}{\operatorname{argmin}} \tfrac{1}{2\tau}\left\|u-u^{k+1/2}\right\|_2^2 + \|\widehat{u}-\widehat{u_0}\|_\gamma^2\,, \quad (64)$$

where $u^{k+1/2} = u^k + \frac{\tau\lambda}{n^2}\mathrm{div}_n p^{k+1}$. Thanks to Parseval Identity, this can be rewritten

$$\widehat{u^{k+1}} = \underset{u\in\mathbb{R}^{\Omega}}{\operatorname{argmin}} \tfrac{1}{2\tau|\Omega|}\left\|\widehat{u}-\widehat{u^{k+1/2}}\right\|_2^2 + \|\widehat{u}-\widehat{u_0}\|_\gamma^2\,, \quad (65)$$

from which we easily obtain the explicit formula for the update given in Algorithm 5.

6.3 Image Shannonization

One interesting application of Model (62) is its ability to (partly or fully) remove aliasing from a given image, thus providing what we could call an "Image Shannonizer". We did not thoroughly investigate this phenomenon yet but the first results we obtained using the simple Gaussian weight function

$$\gamma(\alpha,\beta) = e^{-2\pi^2\sigma^2\left(\frac{\alpha^2}{M^2}+\frac{\beta^2}{N^2}\right)} \quad (66)$$

seem interesting enough to be mentioned here.

Aliasing arises when a continuous image is not sampled in accordance with Shannon Theorem, that is, when the sampling step is too large compared to the highest frequency component that the image contains. In that case, the sampled image will be *aliased*, which means that its discrete Fourier coefficients will be the sum of one correct value and several incorrect values arising from higher frequencies that cannot be represented in the available discrete Fourier domain. In practice, since the power spectrum of natural images tends to exhibit a power-law decrease (see [61]), aliasing mostly impacts the highest frequencies of the discrete image in general; it is thus logical to choose for γ a decreasing function of the distance to the origin. The isotropic map (66) is a possibility, but it would certainly be worth exploring other choices.

The Shannon interpolate of an aliased image is very oscillatory in general, because the aliased component define a trigonometric polynomial with improper aliased frequencies. Therefore, we can expect Model (62) to show interesting aliasing removal performances, as STV is strongly affected by oscillations. Indeed, we can observe in Fig. 12 and 13 that the aliasing of the input image u_0 (which is clearly visible on its spectrum) is completely removed after processing through the Image Shannonizer, without introducing noticeable blur on the image.

Algorithm 5: Chambolle-Pock resolvent algorithm for problem (62)

Initialization: *Choose* $\tau,\sigma>0$, $\theta\in[0,1]$, $u^0\in\mathbb{R}^{\Omega}$, $p^0\in\mathbb{R}^{\Omega_n}\times\mathbb{R}^{\Omega_n}$, *set* $\overline{u}^0=u^0$ *and define* ν *and* $\pi_{\infty,2}$ *as in Algorithm 2.*

Iterations: *For* $k\ge 0$, *update* p^k, u^k *and* $\overline{u}^k$ *with*

$$\begin{aligned} p^{k+1} &= \pi_{\infty,2}\left(\left(p^k+\tfrac{\sigma\lambda}{n^2}\nabla_n\overline{u}^k\right)/\nu\right)\\ u^{k+1/2} &= u^k + \tfrac{\tau\lambda}{n^2}\mathrm{div}_n p^{k+1}\\ u^{k+1} &= \mathcal{F}^{-1}\left(\frac{\widehat{u^{k+1/2}} + 2\tau\gamma\cdot\widehat{u_0}}{1+2\tau\gamma}\right)\\ \overline{u}^{k+1} &= u^{k+1}+\theta\left(u^{k+1}-u^k\right) \end{aligned}$$

7 Conclusion

In this paper we showed that images delivered by variational TV-based models could not be easily interpolated when the TV is discretized with a classical finite difference scheme. However, we demonstrated on several examples that a variant called STV (for Shannon TV) successfully addresses this issue, and can be efficiently handled using Legendre-Fenchel duality and Chambolle-Pock Algorithm. We easily adapted the STV variant to Huber-TV regularization, which let us believe that STV could be easily applied to other variants of the discrete TV as well; for example, the Total Generalized Variation (TGV) proposed in [9] involves higher order derivatives that could be computed exactly as in the STV approach.

The choice of the upsampling factor n used to estimate STV with a Riemann sum was discussed and it was shown that $n=1$ was inadequate. However, it would be interesting to further investigate this issue and prove that $n=2$ (or intermediate values between 1 and 2) guarantees a close correspondence between the true STV and its estimate STV_n.

We also presented a new STV-based restoration model relying on a weight map in the Fourier domain, and showed that in certain cases it could be used as an "Image Shannonizer", which transforms an image

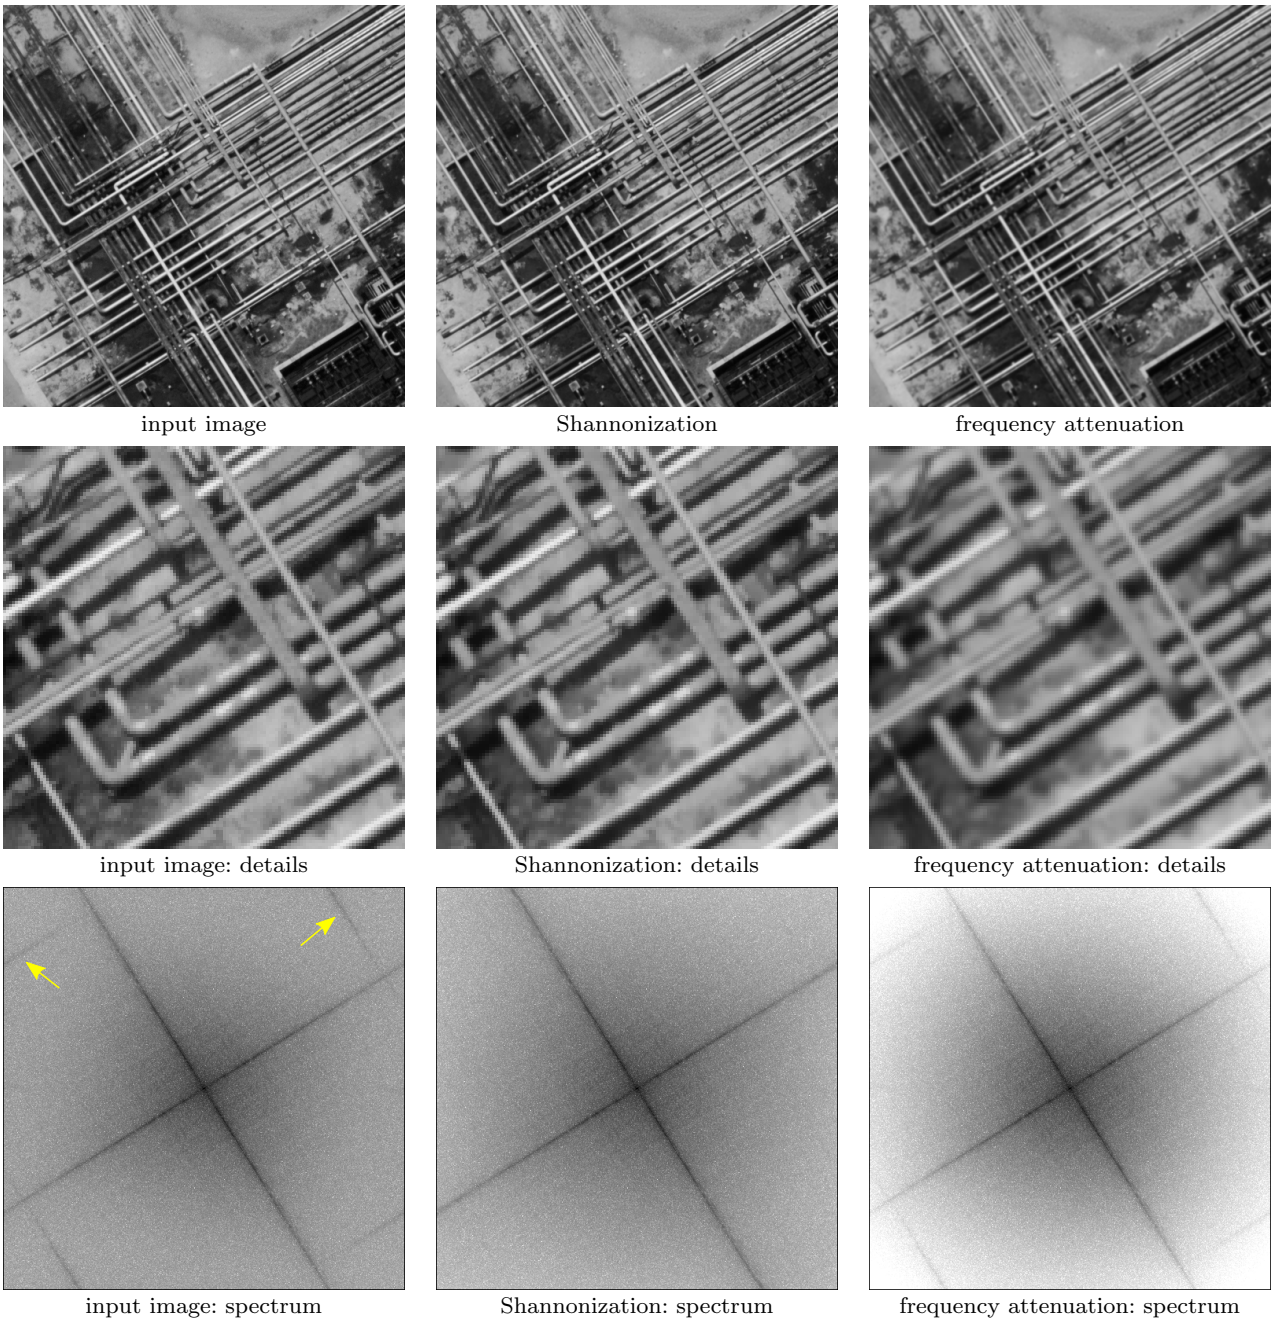

Fig. 12 Image "Shannonization". The input image (left column) is slightly aliased, as indicated by the periodic continuation patterns (see red arrows) that appear in its Fourier spectrum (3rd row). Processing this image with the "Image Shannonizer" (62) results in a visually similar image (middle column) that seems aliasing-free (the patterns are not visible any more on the 3rd row). In comparison, a generic frequency attenuation process (on the right column, with a Gaussian attenuation map) produces a large amount of blur while being less efficient in terms of aliasing removal.

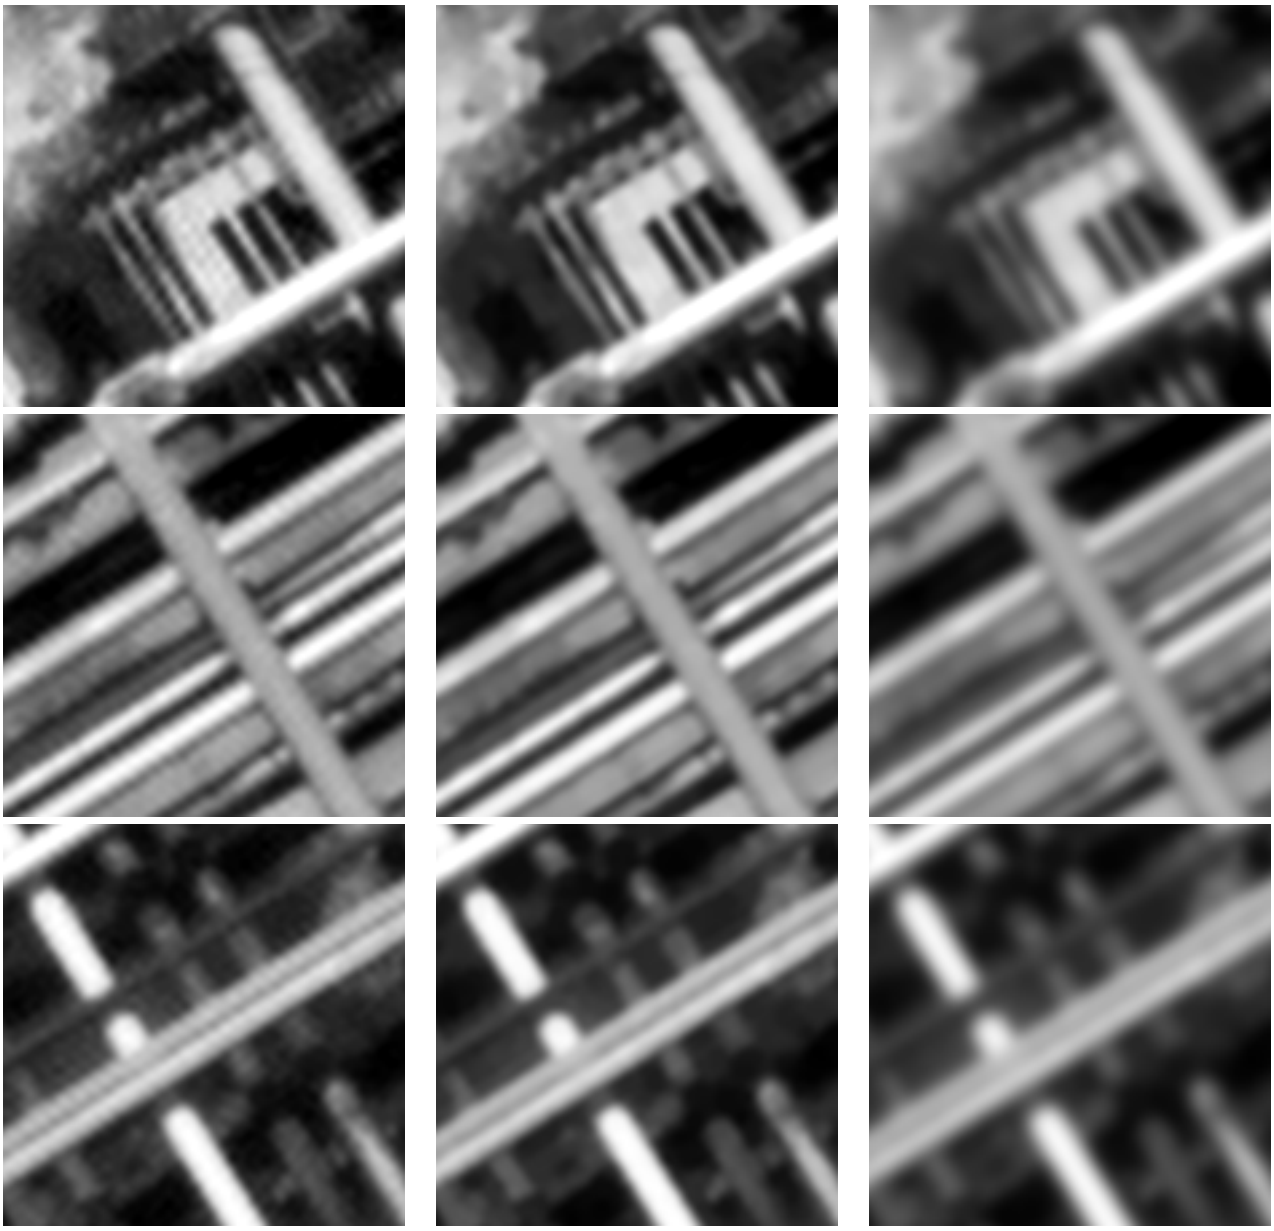

Fig. 13 Details of Fig. 12 with Shannon resampling. Different parts of the three images of the first row of Fig. 12 are shown after Shannon interpolation. As expected, the output of the "Image Shannonizer" (middle) is well interpolable, contrary to the input image (left) on which oscillations appear. A simple frequency attenuation (right) is not efficient, since it introduces a large amount of undesired blur.

into a very similar one that can be easily interpolated (with Shannon interpolation or spline interpolation for example). This seems particularly interesting, as most images are not perfectly sampled (and hence difficult to interpolate) and would hence benefit a lot from this process. This opens new perspectives on aliasing removal (and thus super-resolution from a single image), but several questions are still to be answered, in particular concerning the choice of the weight map.

Source codes

Source codes corresponding to the algorithms described in this paper are freely available on the web pages of the authors.

Acknowledgments

We would like to thank the anonymous reviewers for their valuable suggestions.

References

1. R. Abergel, C. Louchet, L. Moisan, and T. Zeng. Total variation restoration of images corrupted by poisson noise with iterated conditional expectations. In *Proceedings of the International Conference on Scale Space and Variational Methods in Computer Vision, Lecture Notes in Computer Science*, volume 9087, pages 178–190. Springer, 2015.
2. A. Aldroubi, M. Unser, and M. Eden. Cardinal spline filters: Stability and convergence to the ideal sinc interpolator. *Signal Processing*, 28(2):127–138, 1992.
3. K. J. Arrow, L. Hurwicz, H. Uzawa, and H. B. Chenery. *Studies in linear and non-linear programming*. Stanford University Press, 1958.
4. J.-F. Aujol and A. Chambolle. Dual norms and image decomposition models. *International Journal of Computer Vision*, 63(1):85–104, 2005.
5. S. Babacan, R. Molina, and A. Katsaggelos. Total variation super resolution using a variational approach. In *Proceedings of the International Conference on Image Processing*, pages 641–644, 2008.
6. A. Beck and M. Teboulle. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *IEEE Transactions on Image Processing*, 18(11):2419–2434, 2009.
7. A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
8. S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
9. K. Bredies, K. Kunisch, and T. Pock. Total generalized variation. *SIAM Journal on Imaging Sciences*, 3(3):492–526, 2010.
10. T. Briand and J. Vacher. How to Apply a Filter Defined in the Frequency Domain by a Continuous Function. *Image Processing On Line*, 6:183–211, 2016.
11. E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on information theory*, 52(2):489–509, 2006.
12. F. Candocia and J. C. Principe. Comments on "sinc interpolation of discrete periodic signals". *IEEE Transactions on signal processing*, 46(7):2044–2047, 1998.
13. A. Chambolle. An algorithm for total variation minimization and applications. *Journal of Mathematical Imaging and Vision*, 20(1-2):89–97, 2004.
14. A. Chambolle. Total variation minimization and a class of binary MRF models. In *Energy minimization methods in computer vision and pattern recognition*, pages 136–152. Springer, 2005.
15. A. Chambolle, V. Caselles, D. Cremers, M. Novaga, and T. Pock. An introduction to total variation for image analysis. *Theoretical foundations and numerical methods for sparse recovery*, 9:263–340, 2010.
16. A. Chambolle, S. E. Levine, and B. J. Lucier. An upwind finite-difference method for total variation-based image smoothing. *SIAM Journal on Imaging Sciences*, 4(1):277–299, 2011.
17. A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
18. T. Chan, A. Marquina, and P. Mulet. High-order total variation-based image restoration. *SIAM Journal on Scientific Computing*, 22(2):503–516, 2000.
19. T. F. Chan and C.-K. Wong. Total variation blind deconvolution. *IEEE transactions on Image Processing*, 7(3):370–375, 1998.
20. T. F. Chan, A. M. Yip, and F. E. Park. Simultaneous total variation image inpainting and blind deconvolution. *International Journal of Imaging Systems and Technology*, 15(1):92–102, 2005.
21. P. L. Combettes and J.-C. Pesquet. Proximal splitting methods in signal processing. In *Fixed-point algorithms for inverse problems in science and engineering*, pages 185–212. Springer, 2011.
22. P. L. Combettes and V. R. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling & Simulation*, 4(4):1168–1200, 2005.
23. L. Condat. Discrete Total Variation: New Definition and Minimization. Preprint GIPSA-lab, 2016.
24. J. Darbon and M. Sigelle. Image restoration with discrete constrained total variation part i: Fast and exact optimization. *Journal of Mathematical Imaging and Vision*, 26(3):261–276, 2006.
25. R. A. DeVore. Deterministic constructions of compressed sensing matrices. *Journal of Complexity*, 23(4):918–925, 2007.
26. D. L. Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.
27. Y. Drori, S. Sabach, and M. Teboulle. A simple algorithm for a class of nonsmooth convex-concave saddle-point problems. *Operations Research Letters*, 43(2):209–214, 2015.
28. I. Ekeland and R. Témam. *Convex Analysis and Variational Problems*, volume 28. SIAM, 1999.
29. Y. C. Eldar and G. Kutyniok. *Compressed sensing: theory and applications*. Cambridge University Press, 2012.
30. G. Facciolo Furlan, A. Almansa, J.-F. Aujol, and V. Caselles. Irregular to regular sampling, denoising and deconvolution. *Multiscale Modeling & Simulation*, 7(4):1574–1608, 2009.

31. J. M. Fadili and G. Peyré. Total variation projection with first order schemes. *IEEE Transactions on Image Processing*, 20(3):657–669, 2011.
32. M. Frigo and S. G. Johnson. The design and implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231, 2005. Special issue on "Program Generation, Optimization, and Platform Adaptation".
33. P. Getreuer. Linear methods for image interpolation. *Image Processing On Line*, 1, 2011.
34. G. Gilboa. A spectral approach to total variation. In *Proceedings of the International Conference on Scale Space and Variational Methods in Computer Vision, Lecture Notes in Computer Science*, volume 7893, pages 36–47, 2013.
35. M. Guerquin-Kern, L. Lejeune, K. P. Pruessmann, and M. Unser. Realistic analytical phantoms for parallel magnetic resonance imaging. *IEEE Transactions on Medical Imaging*, 31(3):626–636, 2012.
36. F. Guichard and F. Malgouyres. Total variation based interpolation. In *Proceedings of the European signal processing conference*, volume 3, pages 1741–1744, 1998.
37. P. J. Huber. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.
38. P. J. Huber. Robust regression: Asymptotics, conjectures and monte carlo. *The Annals of Statistics*, 1(5):799–821, 1973.
39. P. Indyk. Explicit constructions for compressed sensing of sparse signals. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 30–33. Society for Industrial and Applied Mathematics, 2008.
40. V. A. Kotelnikov. On the capacity of the 'ether' and of cables in electrical communication. In *Proceedings of the 1st All-Union Conference on Technological Reconstruction of the Community Sector and Low-Current Engineering*, 1933.
41. S. Lanzavecchia and P. L. Bellon. A moving window shannon reconstruction algorithm for image interpolation. *Journal of Visual Communication and Image Representation*, 5(3):255–264, 1994.
42. C. Louchet and L. Moisan. Posterior expectation of the total variation model: properties and experiments. *SIAM Journal on Imaging Sciences*, 6(4):2640–2684, 2013.
43. C. Louchet and L. Moisan. Total variation denoising using iterated conditional expectation. In *Proceedings of the European signal processing conference*, pages 1592–1596. IEEE, 2014.
44. F. Malgouyres and F. Guichard. Edge direction preserving image zooming: a mathematical and numerical analysis. *SIAM Journal on Numerical Analysis*, 39(1):1–37, 2001.
45. R. Marks. *Introduction to Shannon sampling and interpolation theory*. Springer Science & Business Media, 2012.
46. B. Matei. Model sets and new versions of shannon sampling theorem. In *New Trends in Applied Harmonic Analysis*, pages 215–279. Springer International Publishing, 2016.
47. B. Matei and Y. Meyer. A variant of compressed sensing. *Revista Matemática Iberoamericana*, 25(2):669–692, 2009.
48. W. Miled, J. Pesquet, and M. Parent. A convex optimization approach for depth estimation under illumination variation. *IEEE Transactions on Image Processing*, 18(4):813–830, 2009.
49. L. Moisan. How to discretize the total variation of an image? In *the 6th International Congress on Industrial Applied Mathematics, Proceedings in Applied Mathematics and Mechanics*, volume 7(1), pages 1041907–1041908, 2007.
50. L. Moisan. Periodic plus smooth image decomposition. *Journal of Mathematical Imaging and Vision*, 39(2):161–179, 2011.
51. J.-J. Moreau. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société mathématique de France*, 93:273–299, 1965.
52. Y. Nesterov. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. In *Soviet Mathematics Doklady*, volume 27, pages 372–376, 1983.
53. M. Nikolova. Local strong homogeneity of a regularized estimator. *SIAM Journal on Applied Mathematics*, 61(2):633–658, 2000.
54. P. Ochs, Y. Chen, T. Brox, and T. Pock. iPiano: inertial proximal algorithm for nonconvex optimization. *SIAM Journal on Imaging Sciences*, 7(2):1388–1419, 2014.
55. N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in optimization*, 1(3):123–231, 2013.
56. J. Preciozzi, P. Musé, A. Almansa, S. Durand, A. Khazaal, and B. Rougé. SMOS images restoration from L1A data: A sparsity-based variational approach. In *IEEE International Geoscience and Remote Sensing Symposium*, pages 2487–2490. IEEE, 2014.
57. H. Raguet, J. M. Fadili, and G. Peyré. A generalized forward-backward splitting. *SIAM Journal on Imaging Sciences*, 6(3):1199–1226, 2013.
58. W. Ring. Structural properties of solutions to total variation regularization problems. *ESAIM: Modélisation Mathématique et Analyse Numérique*, 34(4):799–810, 2000.
59. R. T. Rockafellar. *Convex analysis*. Princeton university press, 1997. Reprint of the 1970 original, Princeton Paperbacks.
60. B. Rougé and A. Seghier. Nonlinear spectral extrapolation: new results and their application to spatial and medical imaging. In *Proceedings of the SPIE's 1995 International Symposium on Optical Science, Engineering, and Instrumentation*, pages 279–289. International Society for Optics and Photonics, 1995.
61. D. L. Ruderman. The statistics of natural images. *Network: computation in neural systems*, 5(4):517–548, 1994.
62. L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.
63. T. Schanze. Sinc interoplation of discrete periodic signals. *IEEE Transactions on Signal Processing*, 43(6):1502–1503, 1995.
64. C. E. Shannon. Communication in the presence of noise. *Proceedings of the Institute of Radio Engineers*, 37(1):10–21, 1949.
65. L. Simon and J.-M. Morel. Influence of unknown exterior samples on interpolated values for band-limited images. *SIAM Journal on Imaging Sciences*, 9(1):152–184, 2016.
66. M. Unser. Sampling-50 years after shannon. *Proceedings of the IEEE*, 88(4):569–587, 2000.
67. 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.
68. M. A. Unser. Ten good reasons for using spline wavelets. In *Optical Science, Engineering and Instrumentation'97*, pages 422–431. International Society for Optics and Photonics, 1997.

69. L. A. Vese and S. J. Osher. Modeling textures with total variation minimization and oscillating patterns in image processing. *Journal of scientific computing*, 19(1-3):553–572, 2003.
70. C. Vogel and M. Oman. Fast, robust total variation-based reconstruction of noisy, blurred images. *IEEE Transactions on Image Processing*, 7(6):813–824, 1998.
71. P. Weiss and L. Blanc-Féraud. A proximal method for inverse problems in image processing. In *Proceedings of the European signal processing conference*, pages 1374–1378. IEEE, 2009.
72. P. Weiss, L. Blanc-Féraud, and G. Aubert. Efficient schemes for total variation minimization under constraints in image processing. *SIAM journal on Scientific Computing*, 31(3):2047–2080, 2009.
73. M. Werlberger, W. Trobin, T. Pock, A. Wedel, D. Cremers, and H. Bischof. Anisotropic Huber-L1 Optical Flow. In *Proceedings of the British Machine Vision Conference*, volume 1, page 3, 2009.
74. E. T. Whittaker. On the functions which are represented by the expansion of interpolating theory. In *Proceedings of the Royal Society of Edinburgh*, volume 35, pages 181–194, 1915.
75. L. Yaroslavsky. Signal sinc-interpolation: a fast computer algorithm. *Bioimaging*, 4(4):225–231, 1996.
76. K. Yosida. *Functional Analysis*. Springer Berlin Heidelberg, 1980. Originally published as volume 123 in the series: Grundlehren der mathematischen Wissenschaften, 1968.
77. C. Yuen and D. Fraser. Digital spectral analysis. *Digital spectral analysis., by Yuen, CK; Fraser, D.. London (UK): Pitman Publishing, 156 p.*, 1, 1979.
78. M. Zhu and T. Chan. An efficient primal-dual hybrid gradient algorithm for total variation image restoration. UCLA CAM Report, 2008.

A Proof of Proposition 1

Let us consider, for $x \in \mathbb{R} \setminus \mathbb{Z}$,

$$\begin{aligned} S_n(x) &= \sum_{p=-n}^{n} \operatorname{sinc}(x - pM) \\ &= \operatorname{sinc}(x) + \sum_{p=1}^{n} \operatorname{sinc}(x - pM) + \operatorname{sinc}(x + pM) \\ &= \frac{\sin \pi x}{\pi x} + \sum_{p=1}^{n} (-1)^{pM} \left(\frac{\sin \pi x}{\pi(x - pM)} + \frac{\sin \pi x}{\pi(x + pM)} \right). \end{aligned}$$

Writing $x = \frac{Mt}{\pi}$, we obtain

$$S_n(x) = \frac{\sin Mt}{M} \left(\frac{1}{t} + \sum_{p=1}^{n} (-1)^{pM} \left(\frac{1}{t - p\pi} + \frac{1}{t + p\pi} \right) \right)$$

and the limit $\operatorname{sincd}_M(x) = \lim_{n \to \infty} S_n(x)$ can be computed explicitly using classical series expansions (due to Euler):

$$\begin{aligned} \forall t \in \mathbb{R} \setminus \pi\mathbb{Z}, \quad \frac{1}{\tan t} &= \frac{1}{t} + \sum_{p=1}^{\infty} \frac{1}{t - p\pi} + \frac{1}{t + p\pi}, \\ \frac{1}{\sin t} &= \frac{1}{t} + \sum_{p=1}^{\infty} (-1)^p \left(\frac{1}{t - p\pi} + \frac{1}{t + p\pi} \right). \end{aligned}$$

If M is odd, $(-1)^{pM} = (-1)^p$ and we obtain

$$\operatorname{sincd}_M(x) = \frac{\sin Mt}{M \sin t} = \frac{\sin \pi x}{M \sin \frac{\pi x}{M}},$$

and if M is even, $(-1)^{pM} = 1$ and the other series yields

$$\operatorname{sincd}_M(x) = \frac{\sin Mt}{M \tan t} = \frac{\sin \pi x}{M \tan \frac{\pi x}{M}}$$

as announced. □

B Proof of Theorem 2

Since each operator T_z is linear and translation-invariant (Hypothesis (ii)), it can be written as a convolution, that is,

$$T_z s(k) = (\psi_z \star s)(k) := \sum_{l \in I_M} \psi_z(k - l) s(l), \tag{67}$$

where ψ_z is an element of $\mathcal{S}$. Taking the DFT of (67), we obtain

$$\forall \alpha \in \mathbb{Z}, \qquad \widehat{T_z s}(\alpha) = \widehat{\psi}_z(\alpha) \widehat{s}(\alpha). \tag{68}$$

Now, from Hypothesis (iii) we immediately get

$$\forall z, w \in \mathbb{R}, \ \forall \alpha \in \mathbb{Z}, \qquad \widehat{\psi}_{z+w}(\alpha) = \widehat{\psi}_z(\alpha) \widehat{\psi}_w(\alpha), \tag{69}$$

and by continuity of $z \mapsto \widehat{\psi}_z(\alpha)$ (deduced from Hypothesis (i)) we obtain

$$\forall \alpha \in \mathbb{Z}, \qquad \widehat{\psi}_z(\alpha) = e^{\gamma(\alpha) z} \tag{70}$$

for some $\gamma(\alpha) \in \mathbb{C}$. Since $\widehat{\psi}_1(\alpha) = e^{\frac{-2i\pi\alpha}{M}}$, we have

$$\gamma(\alpha) = -2i\pi \left(\frac{\alpha}{M} + p(\alpha) \right), \tag{71}$$

where $p(\alpha) \in \mathbb{Z}$ and $p(-\alpha) = -p(\alpha)$ (the fact that $T_z u$ is real-valued implies that $\widehat{\psi}_z(-\alpha) = \widehat{\psi}_z(\alpha)^*$).

Last, we compute

$$\begin{aligned} &\|T_z - id\|_2^2 \\ &= \sup_{\|s\|_2 = 1} \|T_z s - s\|_2^2 \\ &= \frac{1}{M} \sup_{\|s\|_2 = 1} \|\widehat{T_z s} - \hat{s}\|_2^2 \\ &= \frac{1}{M} \sup_{\|\hat{s}\|_2^2 = M} \sum_{\alpha \in \widehat{I}_M} |e^{-2i\pi\left(\frac{\alpha}{M} + p(\alpha)\right) z} - 1|^2 \cdot |\hat{s}(\alpha)|^2 \\ &= 4 \max_{\alpha \in \widehat{I}_M} \sin^2 \left(\pi \left(\frac{\alpha}{M} + p(\alpha) \right) z \right) \\ &= 4\pi^2 z^2 \max_{\alpha \in \widehat{I}_M} \left(\frac{\alpha}{M} + p(\alpha) \right)^2 + \underset{z \to 0}{o}(z^2). \end{aligned}$$

Hence,

$$\lim_{z \to 0} |z|^{-1} \|T_z - id\|_2 = 2\pi \max_{\alpha \in \widehat{I}_M} \left| \frac{\alpha}{M} + p(\alpha) \right| \tag{72}$$

and since $\frac{\alpha}{M} \in (-\frac{1}{2}, \frac{1}{2})$ and $p(\alpha) \in \mathbb{Z}$ for any $\alpha \in \widehat{I}_M$, the right-hand term of (72) is minimal if and only if $p(\alpha) = 0$ for all $\alpha \in \widehat{I}_M$. We conclude from (71) and (70) that

$$\forall \alpha \in \widehat{I}_M, \qquad \widehat{\psi}_z(\alpha) = e^{-2i\pi\alpha z / M}, \tag{73}$$

and thus (68) can be rewritten as

$$T_z s(k) = \frac{1}{M} \sum_{\alpha \in \widehat{I}_M} \widehat{s}(\alpha) e^{-2i\pi\alpha z/M} e^{-2i\pi\alpha k/M}, \tag{74}$$

which is exactly $S(k-z)$ thanks to (13) (recall that the real part is not needed because M is odd). Therefore, (24) is a necessary form for a set of operators (T_z) satisfying Hypotheses (i) to (iv).

Conversely, one easily checks that the operators (T_z) defined by (24) satisfy the Hypotheses (i) to (iv). □

C Proof of Proposition 8

Let us denote by $\nabla_{n,x} u$ and $\nabla_{n,y} u$ the two elements of $\mathbb{R}^{\Omega_n}$ such that $\nabla_n u = (\nabla_{n,x} u, \nabla_{n,y} u)$. In the following, the notation $\langle \cdot, \cdot \rangle_X$ stands for the usual Euclidean (respectively Hermitian) inner product over the real (respectively complex) Hilbert space X. We have

$$\langle \nabla_n u, p \rangle_{\mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}} = \langle \nabla_{n,x} u, p_x \rangle_{\mathbb{R}^{\Omega_n}} + \langle \nabla_{n,y} u, p_y \rangle_{\mathbb{R}^{\Omega_n}} .$$

Recall that we defined $\mathrm{div}_n = -\nabla_n^*$, the opposite of the adjoint of ∇_n. Noting $\mathrm{div}_{n,x} = -\nabla_{n,x}^*$ and $\mathrm{div}_{n,y} = -\nabla_{n,y}^*$, we have

$$\langle \nabla_n u, p \rangle_{\mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}} = \langle u, -\mathrm{div}_{n,x}(p_x) - \mathrm{div}_{n,y}(p_y) \rangle_{\mathbb{R}^{\Omega}} .$$

so that we identify $\mathrm{div}_n(p) = \mathrm{div}_{n,x}(p_x) + \mathrm{div}_{n,y}(p_y)$. Let us focus on the computation of $\mathrm{div}_{n,x}(p_x)$. Let $\widehat{\Omega_1}$, $\widehat{\Omega_2}$, $\widehat{\Omega_3}$, $\widehat{\Omega_4}$ be the sets defined by

$$\begin{aligned}
\widehat{\Omega_1} &= \left\{(\alpha,\beta) \in \mathbb{R}^2,\ |\alpha| < \tfrac{M}{2},\ |\beta| < \tfrac{N}{2}\right\} \cap \mathbb{Z}^2 \\
\widehat{\Omega_2} &= \left\{(\pm\tfrac{M}{2},\beta) \in \mathbb{R}^2,\ |\beta| < \tfrac{N}{2}\right\} \cap \mathbb{Z}^2 \\
\widehat{\Omega_3} &= \left\{(\alpha,\pm\tfrac{N}{2}) \in \mathbb{R}^2,\ |\alpha| < \tfrac{M}{2}\right\} \cap \mathbb{Z}^2 \\
\widehat{\Omega_4} &= \left\{(\pm\tfrac{M}{2},\pm\tfrac{N}{2})\right\} \cap \mathbb{Z}^2 .
\end{aligned}$$

Notice that some sets among $\widehat{\Omega_2}$, $\widehat{\Omega_3}$ and $\widehat{\Omega_4}$ may be empty according to the parity of M and N. Now, let $h_{\widehat{p_x}}$ be the function defined in Proposition 8 and let us show that

$$\forall (\alpha,\beta) \in \widehat{\Omega}, \quad \widehat{\mathrm{div}_{n,x}(p_x)}(\alpha,\beta) = 2i\pi \frac{\alpha}{M} h_{\widehat{p_x}}(\alpha,\beta). \tag{75}$$

Given $z \in \mathbb{C}$, we denote as usual by z^* the conjugate of z. Thanks to Parseval identity, and using Proposition 6 (because we assumed $n \geq 2$), we have

$$\begin{aligned}
\langle \nabla_{n,x} u, p_x \rangle_{\mathbb{R}^{\Omega_n}} &= \frac{1}{n^2 MN} \langle \widehat{\nabla_{n,x} u}, \widehat{p_x} \rangle_{\mathbb{C}^{\Omega_n}} \\
&= \frac{1}{n^2 MN} \sum_{(\alpha,\beta) \in \widehat{\Omega}_n} \widehat{\nabla_{n,x} u}(\alpha,\beta)\ (\widehat{p_x}(\alpha,\beta))^* \\
&= \frac{1}{MN} \sum_{\substack{-\frac{M}{2} \leq \alpha \leq \frac{M}{2} \\ -\frac{N}{2} \leq \beta \leq \frac{N}{2}}} -\widehat{u}(\alpha,\beta) \left(2i\pi\varepsilon_M(\alpha)\varepsilon_N(\beta)\frac{\alpha}{M}\widehat{p_x}(\alpha,\beta)\right)^* .
\end{aligned}$$

It follows that

$$\langle \nabla_{n,x} u, p_x \rangle_{\mathbb{R}^{\Omega_n}} = S_1 + S_2 + S_3 + S_4 ,$$

where for all $k \in \{1,2,3,4\}$, we have set

$$S_k = \frac{1}{MN} \sum_{(\alpha,\beta) \in \widehat{\Omega}_k} -\widehat{u}(\alpha,\beta) \left(2i\pi\varepsilon_M(\alpha)\varepsilon_N(\beta)\frac{\alpha}{M}\widehat{p_x}(\alpha,\beta)\right)^* .$$

Consider S_1 first. Since we have $\varepsilon_M(\alpha) = \varepsilon_N(\beta) = 1$ and $h_{\widehat{p_x}}(\alpha,\beta) = \widehat{p_x}(\alpha,\beta)$ for all $(\alpha,\beta) \in \widehat{\Omega_1}$, we recognize

$$S_1 = \frac{1}{MN} \sum_{|\alpha| < \frac{M}{2},\ |\beta| < \frac{N}{2}} -\widehat{u}(\alpha,\beta) \left(2i\pi\frac{\alpha}{M} h_{\widehat{p_x}}(\alpha,\beta)\right)^* .$$

Now consider S_2. If M is odd, $\widehat{\Omega_2}$ is empty and $S_2 = 0$. Otherwise, since $\varepsilon_M(\alpha)\varepsilon_N(\beta) = 1/2$ for all $(\alpha,\beta) \in \widehat{\Omega_2}$, by grouping together the terms $(-\frac{M}{2},\beta)$ and $(\frac{M}{2},\beta)$, we get

$$\begin{aligned}
S_2 &= \frac{1}{MN} \sum_{\alpha = -\frac{M}{2}, |\beta| < \frac{N}{2}} -\widehat{u}(\alpha,\beta) \\
&\qquad \times \left(2i\pi\frac{1}{2}\frac{\alpha}{M}\widehat{p_x}(\alpha,\beta) - 2i\pi\frac{1}{2}\frac{\alpha}{M}\widehat{p_x}(-\alpha,\beta)\right)^* \\
&= \frac{1}{MN} \sum_{\alpha = -\frac{M}{2}, |\beta| < \frac{N}{2}} -\widehat{u}(\alpha,\beta) \left(2i\pi\frac{\alpha}{M} h_{\widehat{p_x}}(\alpha,\beta)\right)^* ,
\end{aligned}$$

since we have set $h_{\widehat{p_x}}(-\frac{M}{2},\beta) = \frac{1}{2}\left(\widehat{p_x}(-\frac{M}{2},\beta) - \widehat{p_x}(\frac{M}{2},\beta)\right)$ for $|\beta| < N/2$.

Similarly for the term S_3. When N is odd, $\widehat{\Omega_3} = \emptyset$ and $S_3 = 0$. Otherwise, when N is even, we have $\varepsilon_M(\alpha)\varepsilon_N(\beta) = 1/2$ for all $(\alpha,\beta) \in \widehat{\Omega_3}$, thus, by grouping together the terms $(\alpha,-\frac{N}{2})$ and $(\alpha,\frac{N}{2})$, we get

$$\begin{aligned}
S_3 &= \frac{1}{MN} \sum_{|\alpha| < \frac{M}{2}, \beta = -\frac{N}{2}} -\widehat{u}(\alpha,\beta) \\
&\qquad \times \left(2i\pi\frac{1}{2}\frac{\alpha}{M}\widehat{p_x}(\alpha,\beta) + 2i\pi\frac{1}{2}\frac{\alpha}{M}\widehat{p_x}(\alpha,-\beta)\right)^* \\
&= \frac{1}{MN} \sum_{|\alpha| < \frac{M}{2}, \beta = -\frac{N}{2}} -\widehat{u}(\alpha,\beta) \left(2i\pi\frac{\alpha}{M} h_{\widehat{p_x}}(\alpha,\beta)\right)^* ,
\end{aligned}$$

since we have set $h_{\widehat{p_x}}(\alpha,-\frac{N}{2}) = \frac{1}{2}\left(\widehat{p_x}(\alpha,-\frac{N}{2}) + \widehat{p_x}(\alpha,\frac{N}{2})\right)$ for $|\alpha| < M/2$.

Lastly, let us consider S_4. When M and N are both even (otherwise $\widehat{\Omega_4} = \emptyset$ and $S_4 = 0$), we immediately get, for $\alpha = -\frac{M}{2}$ and $\beta = -\frac{N}{2}$,

$$\begin{aligned}
S_4 &= -\widehat{u}(\alpha,\beta) \left(\sum_{s_1=\pm1, s2=\pm1} 2i\pi\frac{1}{4} s_1 \frac{\alpha}{M}\widehat{p_x}(s_1\alpha, s_2\beta)\right)^* \\
&= -\widehat{u}(\alpha,\beta) \left(2i\pi\frac{\alpha}{M} h_{\widehat{p_x}}(\alpha,\beta)\right)^* ,
\end{aligned}$$

since for all $(\alpha,\beta) \in \widehat{\Omega_4}$, we have $\varepsilon_M(\alpha)\varepsilon_N(\beta) = 1/4$ and we have set $h_{\widehat{p_x}}(\alpha,\beta) = \frac{1}{4}\sum_{s_1=\pm1, s2=\pm1} s_1 \widehat{p_x}(s_1\alpha, s_2\beta)$.

Finally, we can write $S_1 + S_2 + S_3 + S_4$ as a sum over $\widehat{\Omega}$, indeed,

$$\begin{aligned}
\langle \nabla_{n,x} u, p_x \rangle_{\mathbb{R}^{\Omega}} &= S_1 + S_2 + S_3 + S_4 \\
&= \frac{1}{MN} \sum_{(\alpha,\beta) \in \widehat{\Omega}} -\widehat{u}(\alpha,\beta) \left(2i\pi\frac{\alpha}{M} h_{\widehat{p_x}}(\alpha,\beta)\right)^* ,
\end{aligned}$$

and using again the Parseval identity, we get (75). With a similar approach, one can check that

$$\forall (\alpha,\beta) \in \widehat{\Omega}, \quad \widehat{\mathrm{div}_{n,y}(p_y)}(\alpha,\beta) = 2i\pi\frac{\beta}{N} h_{\widehat{p_y}}(\alpha,\beta) ,$$

where $h_{\widehat{p_y}}$ is defined in Proposition 8. Consequently, for any $(\alpha,\beta) \in \widehat{\Omega}$, we have

$$\widehat{\mathrm{div}_n(p)}(\alpha,\beta) = 2i\pi \left(\frac{\alpha}{M} h_{\widehat{p_x}}(\alpha,\beta) + \frac{\beta}{N} h_{\widehat{p_y}}(\alpha,\beta)\right) ,$$

which ends the proof of Proposition 8. □

D Proof of Theorem 3

Recall that for any integer M, we denote by T_M the real vector space of real-valued trigonometric polynomials that can be written as complex linear combination of the family $(x \mapsto e^{2i\pi\frac{\alpha x}{M}})_{-\frac{M}{2}\leq\alpha\leq\frac{M}{2}}$. In order to prove Theorem 3 we need the following Lemma.

Lemma 1 *Let $M = 2m+1$ be an odd positive integer. The functions F and G defined by,*

$$\forall x \in \mathbb{R}, \quad F(x) = \frac{1}{M}\sum_{\alpha=-m}^{m} e^{\frac{2i\pi\alpha x}{M}}, \; G(x) = F(x) - F(x-1)\,,$$

are both in T_M and G satisfies

$$\sum_{k=0}^{M-1}|G(k)| = 2\,, \quad \int_1^M |G(x)|\,dx \geq \frac{8}{\pi^2}\log\left(\frac{2M}{\pi}\right) - 2\,.$$

Proof F is in T_M by construction, and so is G as the difference of two elements of T_M. Writing $\omega = \frac{\pi}{M}$, we can notice that $F(0) = 1$ and

$$\forall x \in (0, M), \quad F(x) = \frac{e^{2i\omega(-m)x}}{M}\cdot\frac{1-e^{2i\pi x}}{1-e^{2i\omega x}} = \frac{\sin(\pi x)}{M\sin(\omega x)}\,,$$

so that $F(k) = 0$ for all integers $k \in [1, M-1]$. Consequently, $G(0) = 1$, $G(1) = -1$ and $G(k) = 0$ for all integers $k \in [2, M-1]$, thus

$$\sum_{k=0}^{M-1}|G(k)| = |G(0)| + |G(1)| = 2\,,$$

yielding the first announced result of the Lemma. Now, remark that the sign changes of G in $(0, 2m+1)$ occur at integer points $2, 3, \dots 2m$ and in $\frac{1}{2}$ (by symmetry). Thus, we have

$$\begin{aligned} J := \int_1^M |G(x)|\,dx &= \sum_{k=1}^{2m}(-1)^k\int_k^{k+1} G(x)\,dx \\ &= 2\sum_{k=0}^{2m-1}(-1)^k\int_k^{k+1} F(x)\,dx\,, \end{aligned}$$

since for all $x \in [0, M]$, we have $G(x) = F(x) - F(x-1)$ and (because M is odd) $F(x) = F(M-x)$. It follows that

$$J \geq 2\left(\sum_{k=0}^{2m}(-1)^k\int_k^{k+1} F(x)\,dx\right) - 2\,,$$

since $|F| \leq 1$ everywhere.

Consequently, by isolating the index $\alpha = 0$ in the definition of F, we get $J \geq 2\left(J' + \frac{1}{M}\right) - 2$, with

$$J' = \sum_{k=0}^{2m}\frac{(-1)^k}{M}\sum_{\substack{-m\leq\alpha\leq m\\ \alpha\neq 0}}\int_k^{k+1} e^{2i\omega\alpha x}\,dx\,.$$

By exchanging the sums and grouping identical terms, we obtain

$$\begin{aligned} J' &= \frac{1}{M}\sum_{\substack{-m\leq\alpha\leq m\\ \alpha\neq 0}}\sum_{k=0}^{2m}(-1)^k\cdot\frac{e^{2i\omega\alpha(k+1)} - e^{2i\omega\alpha k}}{2i\omega\alpha} \\ &= \sum_{\substack{-m\leq\alpha\leq m\\ \alpha\neq 0}}\frac{-1}{i\pi\alpha}\sum_{k=1}^{2m}\left(-e^{2i\omega\alpha}\right)^k. \end{aligned} \tag{76}$$

After summation of the geometric progression

$$\begin{aligned} \sum_{k=1}^{2m}\left(-e^{2i\omega\alpha}\right)^k &= -e^{2i\omega\alpha}\cdot\frac{1-e^{2i\omega\alpha(2m)}}{1+e^{2i\omega\alpha}} \\ &= e^{i\pi\alpha}\frac{i\sin(2\omega m\alpha)}{\cos(\omega\alpha)} = \frac{i\sin(2\omega m\alpha - \pi\alpha)}{\cos(\omega\alpha)} = -i\tan(\omega\alpha)\,, \end{aligned}$$

Equation (76) finally leads to

$$J' = \sum_{\substack{-m\leq\alpha\leq m\\ \alpha\neq 0}}\frac{1}{\pi\alpha}\cdot\tan(\omega\alpha) = \frac{2}{M}\sum_{\alpha=1}^{m} g(\omega\alpha)$$

where $g = t \mapsto \frac{\tan t}{t}$. Now since g is positive and increasing on $(0, \frac{\pi}{2})$, we have

$$\sum_{\alpha=1}^{m} g(\omega\alpha) \geq \int_0^m g(\omega x)\,dx = \frac{1}{\omega}\int_0^{\omega m} g(t)\,dt\,.$$

Using the lower bound $g(t) \geq \frac{2}{\pi}\tan t$ for $t \in (0, \frac{\pi}{2})$, we finally get

$$J' \geq \frac{4}{\pi^2}\int_0^{\omega m}\tan t\,dt = -\frac{4}{\pi^2}\log\cos(\omega m) = -\frac{4}{\pi^2}\log\sin\left(\frac{\omega}{2}\right)$$

and thus $J' \geq \frac{4}{\pi^2}\log\left(\frac{2}{\omega}\right)$, from which the inequality announced in Lemma 1 follows. □

Now, let us prove the Theorem 3 by building a discrete image u such that $\mathrm{STV}_1(u)$ is fixed but $\mathrm{STV}_\infty(u)$ increases with the image size. We consider the function H defined by

$$\forall x \in \mathbb{R}, \quad H(x) = \int_0^x G(t)\,dt\,,$$

where $G \in T_M$ is the real-valued M-periodic trigonometric polynomial defined in Lemma 1 ($M = 2m+1$). Since the integral of G over one period is zero ($\int_0^M G(t)\,dt = 0$), H is also an element of T_M. Consequently, the bivariate trigonometric polynomial defined by

$$\forall (x, y) \in \mathbb{R}^2, \quad U(x, y) = \frac{1}{M}H(x)\,,$$

belongs to $T_M \otimes T_M$, and since M is odd it is exactly the Shannon interpolate of the discrete image defined by

$$\forall (k, l) \in I_M \times I_M, \quad u(k, l) = U(k, l). \tag{77}$$

In particular, by definition of STV_1 and STV_∞, we have

$$\mathrm{STV}_1(u) = \sum_{(k,l)\in\Omega}|\nabla U(k, l)|\,,$$
$$\text{and} \quad \mathrm{STV}_\infty(u) = \int_{[0,M]^2}|\nabla U(x, y)|\,dxdy\,.$$

From Lemma 1, we have on the one hand,

$$\begin{aligned} \mathrm{STV}_1(u) &= \sum_{(k,l)\in\Omega}|\nabla U(k, l)| \\ &= \sum_{k=0}^{2m}|H'(k)| = \sum_{k=0}^{2m}|G(k)| = 2\,, \end{aligned}$$

and on the other hand,

$$\begin{aligned} \mathrm{STV}_\infty(u) = \int_{[0,M]^2}|\nabla U(x, y)|\,dxdy &= \int_0^M |H'(x)|\,dx \\ &= \int_0^M |G(x)|\,dx \geq \frac{8}{\pi^2}\log\left(\frac{2M}{\pi}\right) - 2\,. \end{aligned}$$

which cannot be bounded from above by a constant independent of M. □

E Proof of Proposition 10

Let $u \in \mathbb{R}^{\Omega}$, $n \in \mathbb{N}$ and $\alpha \in \mathbb{R}$ such that $n \geq 1$ and $\alpha > 0$. One can rewrite $\mathrm{HSTV}_{\alpha,n}(u) = \frac{1}{n^2} H_\alpha(\nabla_n u)$, where

$$\forall g \in \mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}, \quad H_\alpha(g) = \sum_{(x,y)\in\Omega_n} \mathcal{H}_\alpha(g(x,y)) \,.$$

Let us show that the Legendre-Fenchel transform of H_α is

$$H_\alpha^\star(p) = \delta_{\|\cdot\|_{\infty,2}\leq 1}(p) + \tfrac{\alpha}{2}\|p\|_2^2 \,.$$

One easily checks that $\mathcal{H}_\alpha \in \Gamma(\mathbb{R}^2)$, and it follows that $H_\alpha \in \Gamma(\mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n})$. Thus, for any image $u \in \mathbb{R}^{\Omega}$, we have $H_\alpha(\nabla_n u) = H_\alpha^{\star\star}(\nabla_n u)$ and

$$H_\alpha^{\star\star}(\nabla_n u) = \sup_{p\in\mathbb{R}^{\Omega_n}\times\mathbb{R}^{\Omega_n}} \langle \nabla_n u, p\rangle - H_\alpha^\star(p) \,. \tag{78}$$

Besides, we have $H_\alpha^\star(p) = \sum_{(x,y)\in\Omega_n} \mathcal{H}_\alpha^\star(p(x,y))$, and the Legendre-Fenchel transform of $\mathcal{H}_\alpha$ is the function $\mathcal{H}_\alpha^\star(z) = \delta_{|\cdot|\leq 1}(z) + \frac{\alpha}{2}|z|^2$, where $\delta_{|\cdot|\leq 1}$ denotes the indicator function of the unit ball for the ℓ^2 norm in $\mathbb{R}^2$. Indeed, it is proven in [55] that $\mathcal{H}_\alpha$ is the Moreau envelope (or Moreau-Yosida regularization) [51,76] with parameter α of the ℓ^2 norm $|\cdot|$, or equivalently the infimal convolution (see [59]) between the two proper, convex and l.s.c functions $f_1(x) = |x|$ and $f_2(x) = \frac{1}{2\alpha}|x|^2$, that is

$$\forall y \in \mathbb{R}^2, \quad \mathcal{H}_\alpha(y) = (f_1 \square f_2)(y) := \inf_{x\in\mathbb{R}^2} f_1(x) + f_2(y-x) \,.$$

Thus, we have $\mathcal{H}_\alpha^\star = (f_1\square f_2)^\star = f_1^\star + f_2^\star$ (see [59,55]), leading exactly to $\mathcal{H}_\alpha^\star(z) = \delta_{|\cdot|\leq 1}(z) + \frac{\alpha}{2}|z|^2$ for any $z \in \mathbb{R}^2$, since we have $f_1^\star = z \mapsto \delta_{|\cdot|\leq 1}(z)$ and $f_2^\star = z \mapsto \frac{\alpha}{2}|z|^2$. It follows that for any $p \in \mathbb{R}^{\Omega_n} \times \mathbb{R}^{\Omega_n}$, we have

$$H_\alpha^\star(p) = \sum_{(x,y)\in\Omega_n} \mathcal{H}_\alpha^\star(p(x,y)) = \delta_{\|\cdot\|_{\infty,2}\leq 1}(p) + \tfrac{\alpha}{2}\|p\|_2^2 \,, \tag{79}$$

and the supremum (78) is a maximum for the same reason as in the proof of Proposition 9. Finally, writing $\mathrm{HSTV}_{\alpha,n}(u) = \frac{1}{n^2} H_\alpha(\nabla_n u) = \frac{1}{n^2} H_\alpha^{\star\star}(\nabla_n u)$ using (78) and (79) leads to the announced result. □