

Sampling Curves with Finite Rate of Innovation

Hanjie Pan*, *Student Member, IEEE* Thierry Blu, *Fellow, IEEE* and
Pier Luigi Dragotti, *Senior Member, IEEE*

EDICS: DSP-SAMP, DSP-RECO, DSP-TFSR

Abstract

In this paper, we extend the theory of sampling signals with finite rate of innovation (FRI) to a specific class of two-dimensional curves, which are defined implicitly as the zeros of a mask function. Here the mask function has a parametric representation as a weighted summation of a finite number of complex exponentials, and therefore, has *finite rate of innovation* [1]. An associated edge image, which is discontinuous on the predefined parametric curve, is proved to satisfy a set of linear *annihilation equations*. We show that it is possible to reconstruct the parameters of the curve (i.e. to detect the exact edge positions in the continuous domain) based on the annihilation equations. Robust reconstruction algorithms are also developed to cope with scenarios with model mismatch. Moreover, the annihilation equations that characterize the curve are *linear* constraints that can be easily exploited in optimization problems for further image processing (e.g., image up-sampling). We demonstrate one potential application of the annihilation algorithm with examples in edge-preserving interpolation. Experimental results with both synthetic curves as well as edges of natural images clearly show the effectiveness of the annihilation constraint in preserving sharp edges, and improving SNRs.

Index Terms

Sampling, signals with finite rate of innovation (FRI), annihilating filter, image up-sampling

Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Hanjie Pan is with the School of Computer and Communication Sciences, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland (e-mail: hanjie.pan@ieee.org). Thierry Blu is with the Department of Electronic Engineering, the Chinese University of Hong Kong, Shatin, N.T., Hong Kong (e-mail: thierry.blu@m4x.org). Pier Luigi Dragotti is with the Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2AZ, United Kingdom (e-mail: p.dragotti@imperial.ac.uk).

I. INTRODUCTION

Sampling plays an essential role in signal processing and communications, which aims at representing a continuous domain signal with a few discrete samples. A major concern with sampling is whether this set of discrete samples is a faithful representation of the original signal. For bandlimited signals, the result is given by Shannon's fundamental sampling theorem [2], which requires that the sampling rate be at least twice the signal bandwidth. The classical sampling theorem can be extended to non-bandlimited signals, which live in shift-invariant spaces, such as for example uniform splines and wavelets [3]–[6]. We refer the readers to [6] for a comprehensive review of modern developments of sampling theory.

These theories cannot deal with more general cases, when the signal is neither bandlimited nor belongs to a shift-invariant space. A recent development by Vetterli *et al.* [1] (see also [7]) shows that it is possible to develop sampling schemes for specific classes of parametric non-bandlimited signals. A common feature of this type of signals is that they are completely characterized by a finite number of parameters per unit time/space, and hence have *finite rate of innovation* (FRI). Typical examples include streams of Diracs [1], [7], piecewise polynomials [1], and piecewise sinusoids [8]. Previous efforts to extend the FRI framework to multi-dimensional cases give rise to various sampling schemes for 2-D Diracs [9], lines of finite length [9], and polygons [9]–[12]. These multi-dimensional signals, however, are usually of very simple geometry and may not have enough descriptive power to cope with various shapes encountered in higher dimensions. Note also an inspiring approach to reconstruct polygons, and then quadrature domains from the complex moments of their indicator image [13]–[15].

In this paper, we propose a sampling scheme and a perfect reconstruction algorithm for a certain class of 2-D curves, which have a parametric representation with finite number of degrees of freedom. Depending on the number of degrees of freedom, this model can describe a wide set of curves that are much more complex than former multi-dimensional FRI-type signals. Robust algorithms to reconstruct the FRI curve model in the presence of noise will be developed. We will demonstrate that our curve model is able to describe edges in natural images, which is of particular interest for practical image processing problems.

This paper is organized as follows. In Section II, we first define an edge image associated with curves that specify the discontinuity positions (i.e. edges) in the 2D plane. Next we prove that under certain assumptions on the parametric form of the curve model, the edge image should satisfy a set of linear equations known as *annihilation equations*. A new more intuitive spatial domain interpretation will also be presented. Based on the annihilation equation, we develop, in Section III, a perfect reconstruction




Fig. 1. A continuous domain edge image, which is discontinuous on some predefined curve. The discontinuity across the curve can be (a) uniform or (b) varying in general.

algorithm of the curve model from a finite number of (ideal) samples of the edge image. Depending on the level of noise, two different algorithms, namely the least-square method and Cadzow's iterative denoising method, are proposed to reconstruct the curve model robustly. Further in Section IV, we show how to use the curve model in natural images. Specifically, we exemplify the potential of the curve annihilation, which can be imposed as an edge-preserving constraint in a practical image processing problem, e.g. image up-sampling. Experimental results with both synthetic curves, which satisfy the exact model, and natural images are shown in Section V. We then conclude in Section VI.

II. EDGE IMAGES WITH FINITE RATE OF INNOVATION

A. Edge Images

For each curve C defined on a 2D plane, we introduce a continuously-defined *edge image* according to the following contour integral (counterclockwise integration):

$$I_C(x, y) = \frac{1}{2j\pi} \int_C \frac{f_0(x_0, y_0) dz_0}{z_0 - x - jy}, \quad (1)$$

where $f_0(x, y)$ is an arbitrary function and $z_0 = x_0 + jy_0$. It can be shown that, for most general $f_0(x, y)$ and contours C (directed, but not necessarily closed), the complex-valued image is analytic almost everywhere except on the curve C where it is discontinuous with jump value given by $f_0(x, y)$.

If we further restrict $f_0(x, y)$ to be analytical within the interior of a closed contour C , then by virtue of Cauchy's integral formula, the edge image $I_C(x, y)$ becomes equal to this analytic continuation in the interior of the contour, and vanishes outside: see Fig. 1 (a) for $f_0 = \text{Constant}$ and Fig. 1 (b) for a more general analytic function.

From the contour integral representation of the edge image (1), we can calculate¹ the exact Fourier transform of the edge image (1) as:

$$\hat{I}_C(\omega_x, \omega_y) = \frac{1}{\omega_x + j\omega_y} \int_C f_0(x_0, y_0) e^{-jx_0\omega_x - jy_0\omega_y} dz_0. \quad (2)$$

Note that here (2) is valid without any restrictions on the curve C , which defines the edge locations in the continuous image domain. We will see in Section II-C1 that, by imposing that the curve C satisfies some specific parametric forms, we can derive a set of linear equations that the edge image should satisfy.

B. Curves with Finite Rate of Innovation

We choose to represent curves implicitly through the zeros of some mask function μ : $\mu(x, y) = 0$. In order to have a limited number of degrees of freedom, we approximate this mask function using a periodic Fourier expansion with few terms:

$$C : \overbrace{\sum_{k=-K_0}^{K_0} \sum_{l=-L_0}^{L_0} c_{k,l} e^{j\frac{2\pi k}{\tau_x} x + j\frac{2\pi l}{\tau_y} y}}^{\mu(x,y)} = 0, \quad \text{where} \quad \begin{cases} 0 \leq x \leq \tau_x, \\ 0 \leq y \leq \tau_y. \end{cases} \quad (3)$$

Here τ_x and τ_y are some positive real numbers that specify the x - y dimensions of the rectangle that contains the curve. We can uniquely define C within this rectangle with $(2K_0 + 1)(2L_0 + 1)$ coefficients $c_{k,l}$, which are also known as “signal innovations” in the framework of *Finite Rate of Innovation (FRI)* [1].

Note however that, if no restrictions on $c_{k,l}$ are enforced, (3) is equivalent to *two real equations* (real + imaginary part) for two real unknowns (x, y) , which is unlikely to result in more than a few isolated points. Hence, in order to allow curves as solutions, we choose to require that $\mu(x, y)$ is a *real-valued* expression (i.e., one real equation for two real unknowns), which is equivalent to enforcing the Hermitian symmetry of the coefficients $c_{k,l}$:

$$c_{k,l} = c_{-k,-l}^*.$$

This formulation leads to curves with very diverse topologies: we may have multiply connected curves, open curves, crossings, and non-smooth connections. Moreover, this representation is potentially very rich: if we increase the number of degrees of freedom, we can approximate arbitrary mask functions as accurately as we want. A closer look at (3) suggests that the coefficients $c_{k,l}$ *annihilate* the complex exponentials, which are uniformly sampled at discrete frequencies $(\omega_x, \omega_y) = (\frac{2\pi k}{\tau_x}, \frac{2\pi l}{\tau_y})$ for $k = -K_0, \dots, K_0$ and $l = -L_0, \dots, L_0$. This is why we term the curves defined in (3) as “annihilable curves”.

¹It can be proved that the Fourier transform of $1/(x + jy)$ is $-2j\pi/(\omega_x + j\omega_y)$.

We will then denote by “edge images with Finite Rate of Innovation” those edge images that are based on FRI curves.

C. Annihilation

1) *Annihilation of the Derivative of the Edge Image*: The annihilating filter method, which is related to Prony’s method in spectral estimation [16]–[18], was shown to be an effective technique to reconstruct FRI-type signals [1], [7]–[10]. Here we will apply the same strategy to annihilate the derivative of an edge image associated with the parametric curve defined in (3). The complex derivative involved here is known to be of *Wirtinger* type in complex analysis: $I'_C = \frac{\partial I_C}{\partial x} + j \frac{\partial I_C}{\partial y}$.

Theorem 1: Consider an annihilable curve C as defined in (3) and its associated edge image $I_C(x, y)$, then for *any* frequencies ω_x and ω_y we have:

$$\sum_{k=-K_0}^{K_0} \sum_{l=-L_0}^{L_0} c_{k,l} \widehat{I}'_C \left(\omega_x - \frac{2\pi k}{\tau_x}, \omega_y - \frac{2\pi l}{\tau_y} \right) = 0, \quad (4)$$

where $\widehat{I}'_C(\omega_x, \omega_y) = j(\omega_x + j\omega_y) \widehat{I}_C(\omega_x, \omega_y)$.

Proof: From (2), $\widehat{I}'_C(\omega_x, \omega_y) = j \int_C f_0(x_0, y_0) e^{-jx_0\omega_x - jy_0\omega_y} dz_0$. Therefore,

$$\begin{aligned} \text{lhs} &= j \sum_{k,l} c_{k,l} \int_C f_0(x_0, y_0) e^{-jx_0\omega_x - jy_0\omega_y} e^{j\frac{2\pi k}{\tau_x}x_0 + j\frac{2\pi l}{\tau_y}y_0} dz_0 \\ &= j \int_C f_0(x_0, y_0) e^{-jx_0\omega_x - jy_0\omega_y} \underbrace{\sum_{k,l} c_{k,l} e^{j\frac{2\pi k}{\tau_x}x_0 + j\frac{2\pi l}{\tau_y}y_0}}_{=0 \text{ from (3)}} dz_0 = 0. \quad \blacksquare \end{aligned}$$

Hence, any $c_{k,l}$ that annihilates the curve C is automatically an annihilator of the derivative of the associated edge image. Observe that the annihilation equation (4) is a linear equation with respect to both the annihilation coefficients and the image data, \widehat{I}'_C , which is completely specified by the Fourier transform of the edge image. Therefore, in principle we should be able to reconstruct the annihilation coefficients by solving a *linear* system of equations, as long as we have access to the Fourier transform of the edge image at sufficiently many frequencies. This fact will be discussed in details in Section III-B.

2) *Spatial Domain Interpretation of the Annihilation Equation*: The annihilation of the derivative of the edge image (4) is formulated as a Fourier domain convolution, which may appear obscure at first glance. However, the underlying principle is actually very simple. We can transform every element in the annihilation equation to the spatial domain (see Table I).

Note that the inverse Fourier transform of the annihilation coefficients $c_{k,l}$ is (up to a constant $4\pi^2$) the same function $\mu(x, y)$ whose roots define the annihilable curve in (3) in the first place. The image

TABLE I
FOURIER DOMAIN ANNIHILATION EQUATION AND ITS SPATIAL DOMAIN INTERPRETATIONS.

Fourier Domain	Spatial Domain
$\hat{\mu}(\omega_x, \omega_y) = 4\pi^2 \sum_{k,l} c_{k,l} \delta\left(\omega_x - \frac{2\pi k}{\tau_x}, \omega_y - \frac{2\pi l}{\tau_y}\right)$ $\widehat{I}'_C(\omega_x, \omega_y) = j(\omega_x + j\omega_y) \hat{I}_C(\omega_x, \omega_y)$ <p style="text-align: center;">convolution *</p> $\sum_{k,l} c_{k,l} \widehat{I}'_C\left(\omega_x - \frac{2\pi k}{\tau_x}, \omega_y - \frac{2\pi l}{\tau_y}\right)$	$\mu(x, y) = \sum_{k,l} c_{k,l} e^{j\frac{2\pi k}{\tau_x}x + j\frac{2\pi l}{\tau_y}y}$ $I'_C(x, y) = \left(\frac{\partial}{\partial x} + j\frac{\partial}{\partial y}\right) I_C(x, y)$ <p style="text-align: center;">multiplication \times</p> $\mu(x, y) \times \underbrace{I'_C(x, y)}_{\text{derivative of edge image } I_C}$




Fig. 2. Spatial domain interpretation of the annihilation equation (4). The function $\mu(x, y)$, whose roots defines the curve, serves as a mask to annihilate whatever is different from zero in the derivative image I'_C (see text in Section II-C2).

data $\widehat{I}'_C(\omega_x, \omega_y)$ are obtained by taking derivatives in the spatial domain. Since the edge image is analytic almost everywhere except on the defined curve C , the derivative image I'_C is different from zero only on the annihilable curve. Therefore, the annihilation equation, which is a Fourier domain convolution, corresponds to a simple spatial domain multiplication. Here the function $\mu(x, y)$ serves as a “mask” that automatically *annihilates* whatever is different from zero in the derivative image I'_C (see, Fig. 2).

III. RECONSTRUCTION OF AN EDGE IMAGE

A. Ideal Samples

1) *Acquisition of the Edge Image Samples:* Consider the general two-dimensional signal sampling setup (as shown in Fig. 3), where a (τ_x, τ_y) -periodic version of the continuous domain edge image, $I_{C_{\text{per}}}(x, y)$, is low-pass filtered with an anti-aliasing filter $\varphi(x, y)$ and uniformly sampled. The question at hand is: can we reconstruct the two dimensional curve (or equivalently the mask coefficients $c_{k,l}$) from the given finite number of samples of the edge image?




Fig. 3. Sampling of edge image $I_{C_{per}}$ associated with an annihilable curve C . The periodized continuous domain edge image is low-pass filtered and then uniformly sampled.

For simplicity, assume that the edge image is convolved with a 2D sinc window (ideal low-pass) of bandwidth² B_x and B_y and sampling step $T_x = \tau_x/N_x$ and $T_y = \tau_y/N_y$ along x, y -directions. Then the samples are (see the derivations in Appendix A-A):

$$\begin{aligned} g_{m,n} &= \langle I_{C_{per}}(x, y), \text{sinc}(B_x(mT_x - x))\text{sinc}(B_y(nT_y - y)) \rangle \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I_{C_{per}}(x, y) \text{sinc}(B_x(mT_x - x)) \text{sinc}(B_y(nT_y - y)) dx dy \end{aligned} \quad (5)$$

$$= \int_0^{\tau_x} \int_0^{\tau_y} I_C(x, y) \varphi(mT_x - x, nT_y - y) dx dy, \quad (6)$$

where $m = 0, 1, \dots, N_x - 1$, $n = 0, 1, \dots, N_y - 1$ and $\varphi(x, y) = \frac{\sin(\pi B_x x) \sin(\pi B_y y)}{B_x B_y \tau_x \tau_y \sin(\pi x/\tau_x) \sin(\pi y/\tau_y)}$ is the 2D Dirichlet kernel. We will stick to the sinc kernel in our analysis in the subsequent sections. Note that, in order to get the same ideal samples, we can also use other non-ideal (even compactly supported) anti-aliasing filters as long as they can reproduce some complex exponentials exactly—see, e.g. the exponential reproducing kernels in [19], which avoid periodizing the edge image. Yet, in practice, we can simply approximate these ideal samples by applying an additional post-filtering when a non-ideal anti-aliasing filter, e.g. a Gaussian kernel $\varphi(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$, is used (see details in Section III-C).

2) *Exact Fourier Samples*: Notice that in order to reconstruct the annihilation coefficients, we need to know the exact Fourier transform of the edge image at certain frequencies (so that we can build a linear system of equations based on the annihilation equation (4)). However, we do not have such direct access—we only have the spatial domain samples $g_{m,n}$ from the sampling process of Fig. 3. Therefore, the difficulty now lies on how we can obtain uniform samples of the Fourier transform of the edge image from the given spatial domain samples.

The periodized edge image $I_{C_{per}}$ can be expressed as a Fourier series:

$$I_{C_{per}}(x, y) = \sum_{k,l \in \mathbb{Z}} \hat{I}_{k,l} e^{j2\pi kx/\tau_x + j2\pi ly/\tau_y}, \quad (7)$$

²Here $B_x\tau_x, B_y\tau_y$ are some *odd* integers for the consideration of the convergence of sinc function summation—see [7] for details.

Algorithm 1: Annihilation of the derivative of the edge image

- 1 Compute the 2-D DFT $\hat{g}_{k,l}$ of the ideally low-pass filtered samples $g_{m,n}$;
 - 2 Evaluate the Fourier transform of the edge image I_C at uniformly sampled discrete frequencies $(\frac{2\pi k}{\tau_x}, \frac{2\pi l}{\tau_y})$ using (9);
 - 3 Build a discrete image data samples $\widehat{I}'_{k,l}$ according to (11);
 - 4 Solve for the annihilation coefficients based on (10) under the constraint of Hermitian symmetry of the coefficients $c_{k,l}$.
-

where $\hat{I}_{k,l}$ are the Fourier series coefficients of $I_{C_{\text{per}}}$ and τ_x, τ_y are the periods of the edge image along x, y -axis respectively. Moreover, thanks to Poisson's summation formula, we also know that

$$\hat{I}_{k,l} = \frac{1}{\tau_x \tau_y} \hat{I}_C \left(\frac{2\pi k}{\tau_x}, \frac{2\pi l}{\tau_y} \right).$$

Assume that the total number of samples within one period is such that $N_x \geq B_x \tau_x$ and $N_y \geq B_y \tau_y$, then from (5) and (7), we have (see the derivations in Appendix A-B):

$$\begin{aligned} g_{m,n} &= \frac{1}{B_x B_y} \sum_{\substack{|k| \leq \lfloor B_x \tau_x / 2 \rfloor \\ |l| \leq \lfloor B_y \tau_y / 2 \rfloor}} \hat{I}_{k,l} e^{j2\pi mk/N_x + j2\pi nl/N_y} \\ &= \frac{1}{B_x B_y \tau_x \tau_y} \sum_{\substack{|k| \leq \lfloor B_x \tau_x / 2 \rfloor \\ |l| \leq \lfloor B_y \tau_y / 2 \rfloor}} \hat{I}_C \left(\frac{2\pi k}{\tau_x}, \frac{2\pi l}{\tau_y} \right) e^{j2\pi mk/N_x + j2\pi nl/N_y} \end{aligned} \quad (8)$$

A careful look at (8) reveals that the spatial domain samples $g_{m,n}$ are simply the *band-limited* inverse discrete Fourier transform (IDFT) of the Fourier series $\hat{I}_{k,l}$. Therefore, the DFT of $g_{m,n}$ coincides (up to a scaling factor) with the Fourier transform of $I_C(x, y)$ for frequencies (ω_x, ω_y) bounded inside the box $\frac{2\pi}{N_x} [-\lfloor B_x \tau_x / 2 \rfloor, \lfloor B_x \tau_x / 2 \rfloor] \times \frac{2\pi}{N_y} [-\lfloor B_y \tau_y / 2 \rfloor, \lfloor B_y \tau_y / 2 \rfloor]$:

$$\hat{g}_{k,l} = \begin{cases} \frac{1}{B_x B_y \tau_x \tau_y} \hat{I}_C \left(\frac{2\pi k}{\tau_x}, \frac{2\pi l}{\tau_y} \right) & \text{for } |k| \leq \lfloor B_x \tau_x / 2 \rfloor \text{ and } |l| \leq \lfloor B_y \tau_y / 2 \rfloor, \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

B. Curve Coefficients from Ideal Samples

We have proved in (4) that the convolution between the curve coefficients and the image data, which is completely determined by the Fourier transform of the associated edge image, is zero. Specifically, if we take $\omega_x = \frac{2\pi k_0}{\tau_x}$, $\omega_y = \frac{2\pi l_0}{\tau_y}$ for $k_0, l_0 \in \mathbb{Z}$, then (4) reduces to a discrete convolution between the annihilation coefficients $c_{k,l}$ and a *discrete* sequence of image data $\widehat{I}'_{k,l}$:

$$c_{k,l} * \widehat{I'_{k,l}} = \sum_{k=-K_0}^{K_0} \sum_{l=-L_0}^{L_0} c_{k,l} \widehat{I'_{k_0-k, l_0-l}} = 0, \quad (10)$$

where

$$\widehat{I'_{k,l}} = \widehat{I'_C}(\omega_x, \omega_y) \Big|_{\substack{\omega_x=2\pi k/\tau_x \\ \omega_y=2\pi l/\tau_y}} = \left(\frac{2\pi k}{\tau_x} + j \frac{2\pi l}{\tau_y} \right) \widehat{I}_C \left(\frac{2\pi k}{\tau_x}, \frac{2\pi l}{\tau_y} \right). \quad (11)$$

Therefore, we can reconstruct the coefficients $c_{k,l}$ (i.e. the annihilable curve) *exactly* from the Fourier transform samples $\widehat{I}_C \left(\frac{2\pi k}{\tau_x}, \frac{2\pi l}{\tau_y} \right)$, which are shown to have one to one correspondence with the spatial domain samples (see, Section III-A2). The annihilating filter method is summarized in Algorithm 1.

Note that we can rewrite the annihilation equation (10) as a linear system of equations:

$$\mathbf{I}' \mathbf{c} = \mathbf{0}, \quad (12)$$

where \mathbf{I}' is a block-circulant convolution matrix associated with the image data $\widehat{I'_{k,l}}$ in (11) and \mathbf{c} is the annihilation coefficients $c_{k,l}$ rearranged lexicographically as a column vector. To avoid any trivial solution (i.e. $\mathbf{c} = \mathbf{0}$), we may want to solve (12) under a non-zero constraint, e.g. $c_{0,0} = 1$. The solution to this linear system of equations gives an exact reconstruction of the coefficients $c_{k,l}$ (up to a multiplicative factor). With the annihilating filter method, we actually separate the reconstruction of *positions* (i.e. where the curve C locates) from *amplitudes* (i.e. the amount of discontinuity across C in the edge image)—see [1], [7] for a few examples in 1D. Any scaling to the reconstructed annihilation coefficients $c_{k,l}$ does not alter the curve it defines, which is uniquely determined by the roots of the mask function $\mu(x, y)$ in (3). Consequently, we can re-synthesize the curve C from $c_{k,l}$, which gives the *exact* position of the discontinuity (in the continuous domain) of the edge image. From this perspective, we can consider the restoration algorithm here as an *edge detector* with infinite accuracy³. Moreover, we will see that the reconstructed mask function finds more extensive applications than the binary edge map provided by conventional edge detection algorithms (see, e.g. Section III-E and Section IV-A).

The next issue is to understand what is the minimum number of samples required in order to have perfect reconstruction. In the 1D cases, e.g. streams of K -periodic Diracs, the minimum sampling rate is $2K + 1$ samples per period [1], which is the next best thing for a signal with the rate of innovation $2K$ (K positions and K amplitudes). For the annihilable curve (3), the innovation parameters are completely determined by the $(2K_0 + 1)(2L_0 + 1)$ complex-valued annihilation coefficients $c_{k,l}$ or, taking Hermitian symmetry into account, $(2K_0 + 1)(2L_0 + 1)$ real-valued degrees of freedom. However, in order to retrieve these degrees of freedom, it is necessary to use more samples of the edge image, as shown in the following theorem.

³Because what is reconstructed is a *continuous domain* edge model.




Fig. 4. The total number of shifts we can have when sliding a $K_1 \times L_1$ window over $K_2 \times L_2$ data block is $(K_2 - K_1 + 1)(L_2 - L_1 + 1)$.




Fig. 5. Approximate ideal low-pass filtered samples obtained by applying additional post-filtering to non-ideal samples (see text in Section III-C).

Theorem 2: The minimum number of samples required for the annihilating filter method to perfectly reconstruct an annihilable curve C , which is generated from $(2K_0 + 1) \times (2L_0 + 1)$ coefficients $c_{k,l}$, is at least $(4K_0 + 1)(4L_0 + 1)$.

Proof: The annihilation equation (10) is a convolution between the curve coefficients $c_{k,l}$ and the image data $\widehat{I'_{k,l}}$, which amounts to sliding $c_{k,l}$ over the image data block. With $(4K_0 + 1)(4L_0 + 1)$ samples, we have $(2K_0 + 1)(2L_0 + 1)$ different shifts (see Fig. 4) of the image data (11), each of which satisfies the annihilation equation. Thus we can uniquely determine the $(2K_0 + 1)(2L_0 + 1)$ unknowns (i.e. the annihilation coefficients $c_{k,l}$) from the $(2K_0 + 1)(2L_0 + 1)$ independent equations—except in the unlikely event that these equations are linearly dependent. ■

C. Non-ideal Low-pass Samples

Our analysis in the previous sections assumes that the continuous domain edge image is ideally low-pass filtered prior to sampling. However, the sinc kernel has infinite spatial support and hence cannot be used in practice. Instead, we must consider more practical kernels for the anti-aliasing filter in the sampling setup (Fig. 3), e.g. the Gaussian kernel $\varphi(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$.

It is important to note that in the annihilation algorithm, we need to have access to the Fourier coefficients of the edge image at certain frequencies, which are shown to have one to one correspondence




Fig. 6. (a) Samples with ideal sinc kernel. (b) Samples with Gaussian kernel. (c) Samples obtained by applying additional post-filtering. The SNR of (b) and (c) compared with the ground truth (a) are 13.47 dB and 38.39 dB respectively.

with the ideal low-pass filtered samples (via the truncated DFT) in Section III-A2. But such a one-to-one correspondence between the spatial and Fourier domain samples is no longer valid in general when kernels that do not reproduce the required sinusoids are used in the sampling setup. Therefore, in order to apply the same annihilation algorithm, we need to find a way to process these non-ideal samples in such a way that we can have a good estimate of the ideal low-pass filtered samples.

In [4], Blu and Unser showed that one can design an optimal (in the least-square sense) digital filter such that the post-filtered samples $\tilde{g}_{m,n}$ are the closest match with the ideal samples (Fig. 5). This optimal filter $\tilde{\varphi}_{m,n}$ is specified with its discrete space Fourier transform (DSFT) as (see details in Appendix A-C):

$$\hat{\tilde{\varphi}}(\omega_x, \omega_y) = \frac{\hat{\varphi}(\omega_x, \omega_y)^*}{\sum_{k,l \in \mathbb{Z}} \left| \hat{\varphi} \left(\omega_x + \frac{2k\pi}{T_x}, \omega_y + \frac{2l\pi}{T_y} \right) \right|^2}$$

for $(\omega_x, \omega_y) \in \left[-\frac{\pi}{T_x}, \frac{\pi}{T_x} \right] \times \left[-\frac{\pi}{T_y}, \frac{\pi}{T_y} \right]$. (13)

In many cases (typ., if φ has compact support), the infinite summation in (13) can be computed exactly. But in practice, we may simply use a few k, l -s to approximate the infinite summation accurately (typ., when $\hat{\varphi}$ is fastly decreasing).

We illustrate the effectiveness of the additional post-filtering with Gaussian low-pass filtered samples. Here σ in the Gaussian kernel is chosen in such a way that one sampling step corresponds to a full width at half maximum (FWHM). The processed samples get much closer to the ideal low-pass filtered ones because of the additional post-filtering (Fig. 6).

D. Curve Coefficients from Inexact Samples

In Section III-B, we have shown that the annihilable curve position can be perfectly reconstructed from a finite number of ideally low-pass filtered samples of the associated edge image. However, the restoration algorithm presented in Section III-B is very sensitive to sample inaccuracies, which are ubiquitous in signal processing.

In the case of annihilable curves, “noise” may arise from two different sources: either from the imperfection of the sampling process i.e., noise corruption (see examples in Section V-A2), quantization, non-ideal lowpass filtering; or more generally, as a result of model mismatch, where the curves that we try to reconstruct are not built from the exact annihilation model (3) (see examples in Section V-A3 and Section V-C1). In either scenario, we need to make a projection to an annihilable curve model that is the best fit for the given “noisy” samples.

Depending on the severity of deviation from an exact annihilable curve model (3), two different denoising algorithms are at our disposal: 1) least square denoising method for mild noise perturbation (typically for SNR above 20 dB); and 2) Cadzow’s method [20], [21] for much more severe distortions.

1) *Least Square Denoising*: In the presence of noise, the image data matrix \mathbf{I}' is only approximately singular. However, it is reasonable to minimize the annihilation results $\mathbf{I}'\mathbf{c}$ in the least-square sense [1], [7] under a non-zero constraint, e.g. $c_{0,0} = 1$:

$$\begin{aligned} \min_{\mathbf{c}} \quad & \|\mathbf{I}'\mathbf{c}\|_2^2 \\ \text{subject to} \quad & c_{0,0} = 1. \end{aligned}$$

The reconstructed coefficients are $\mathbf{c} = (\mathbf{e}_0^H (\mathbf{I}'^H \mathbf{I}')^{-1} \mathbf{e}_0)^{-1} ((\mathbf{I}'^H \mathbf{I}')^{-1} \mathbf{e}_0)$, where \mathbf{e}_j is an indicator vector, whose component is 1 for the j -th position and 0 otherwise.

2) *Cadzow’s Iterative Denoising*: The least square denoising method works quite well for moderate levels of noise. However, as the distortion gets more severe, the reconstructed annihilation coefficients become less reliable. From our previous experience in the 1D FRI signals [7], it is desirable to exploit the *consistencies* among samples (via the annihilation equation) over a wider range to further denoise the samples before applying the least square denoising method. This usually amounts to looking for a larger set of annihilation coefficients as if the signal description (i.e. the annihilable curve model (3)) required more innovation parameters than the ground-truth.

Corollary 1: Suppose the image data matrix \mathbf{I}' is built as if the annihilation coefficients $\tilde{c}_{k,l}$ were $(2\tilde{K}_0 + 1) \times (2\tilde{L}_0 + 1)$, then the last $(2(\tilde{K}_0 - K_0) + 1)(2(\tilde{L}_0 - L_0) + 1)$ smallest singular values of \mathbf{I}' are zero.

Proof: Note that we always have one set of annihilation coefficients $\tilde{c}_{k,l}$ (which is of larger size compared with $c_{k,l}$)

$$\tilde{c}_{k,l} = \begin{cases} c_{k,l} & \text{for } |k| \leq K_0 \text{ and } |l| \leq L_0, \\ 0 & \text{otherwise,} \end{cases}$$

such that the annihilation equation is satisfied

$$\sum_{k=-\tilde{K}_0}^{\tilde{K}_0} \sum_{l=-\tilde{L}_0}^{\tilde{L}_0} \tilde{c}_{k,l} \widehat{I}'_{k_0-k, l_0-l} = \sum_{k=-K_0}^{K_0} \sum_{l=-L_0}^{L_0} c_{k,l} \widehat{I}'_{k_0-k, l_0-l} = 0. \quad (14)$$

On the other hand, the annihilation equation (4) is satisfied for any frequencies ω_x and ω_y . Arbitrary frequency shifts of the image data or integer shifts of the annihilation coefficients do not change the outcome of the annihilation (10) and (14). Specifically, consider another set of coefficients

$$\tilde{c}_{k,l}^{\text{shift}} = \begin{cases} c_{k-k', l-l'} & \text{for } |k-k'| \leq K_0 \text{ and } |l-l'| \leq L_0, \\ 0 & \text{otherwise,} \end{cases}$$

which is obtained by shifting the non-zero portion in $\tilde{c}_{k,l}$ (specified by $c_{k,l}$) by k' and l' in the horizontal and vertical directions respectively. The shifted coefficients $c_{k,l}^{\text{shift}}$ still satisfy the same annihilation equation (14). The total number of shifts we can have are $(2(\tilde{K}_0 - K_0) + 1)(2(\tilde{L}_0 - L_0) + 1)$ (see Fig. 4). All the shifted sets of annihilation coefficients form the basis vectors of the null space of \mathbf{I}' , i.e. the last $(2(\tilde{K}_0 - K_0) + 1)(2(\tilde{L}_0 - L_0) + 1)$ smallest singular values of \mathbf{I}' should be zero. ■

In general, the convolution matrix built from the noisy samples does not satisfy the rank constraint in Corollary 1. Like in the 1D case, we would like to find a (block)-circulant convolution matrix with the rank that is expected from Corollary 1. Cadzow's solution consists in iterating back and forth the following two steps (see Algorithm 2):

- (i) Thresholding step — Set the smallest singular values of the convolution matrix to zero (based on Corollary 1);
- (ii) Projection step — Find the closest block-circulant convolution matrix to the thresholded matrix obtained in step (i).

Once the samples are “denoised” (i.e., satisfy Corollary 1), the annihilation coefficients are reconstructed by solving the standard least square minimization problem of Section III-D1.

E. Edge Image Amplitudes

The annihilation algorithm developed in the previous subsections addresses only half of the restoration problem: what we reconstruct is the (curve) location where the discontinuity occurs in the edge image.

Algorithm 2: Cadzow's Iterative Denoising

- 1 Build the image data matrix \mathbf{I}' as if the annihilating coefficients $\tilde{c}_{k,l}$ were $(2\tilde{K}_0 + 1) \times (2\tilde{L}_0 + 1)$ according to (11) and perform the SVD of $\mathbf{I}' = \mathbf{U}\mathbf{S}\mathbf{V}^H$;
 - 2 Build the diagonal matrix \mathbf{S}_0 by setting the last $(2(\tilde{K}_0 - K_0) + 1)(2(\tilde{L}_0 - L_0) + 1)$ smallest eigenvalues in \mathbf{S} to be zero and re-synthesize $\mathbf{I}'_0 = \mathbf{U}\mathbf{S}_0\mathbf{V}^H$;
 - 3 Obtain the denoised image data matrix $\hat{I}'_{0k,l}$ by averaging diagonals in \mathbf{I}'_0 based on the block-circulant structure of a 2D convolution matrix;
 - 4 Repeat step 1 to 3 until the $[(2(\tilde{K}_0 - K_0) + 1)(2(\tilde{L}_0 - L_0) + 1)]^{\text{th}}$ eigenvalue is smaller than a predefined threshold;
 - 5 Take the denoised image data matrix $\hat{I}'_{0k,l}$ and use the least square minimization in Section III-D1 to reconstruct the annihilation coefficients.
-

With the annihilating filter method, the reconstruction of amplitudes is separated from that of locations. In the one dimensional FRI case, e.g. streams of K Diracs, the amplitudes can be recovered via data-fitting by solving an unconstrained least square minimization problem once we obtain the position of the Diracs [1]. This approach was feasible because only a *finite* number of amplitudes have to be retrieved (K , if there are K Diracs). However, the same strategy fails for the retrieval of the jump values of the edge image, because they require the retrieval of a complete function; i.e., an infinite number of degrees of freedom.

In order to reconstruct the amplitudes, we may either 1) make the assumption that the amplitude function can be represented explicitly in a parametric form (with finite number of parameters); or 2) approximate the edge amplitudes by a smoothly varying function which can be retrieved, e.g., through an implicit regularization framework. Since the amplitude functions may not satisfy any explicit parametric form exactly, we adopt the second implicit approach in the rest of this paper.

For simplicity, let us assume that what we have are the ideally low-pass filtered samples I_{LP} (for instance, the $g_{m,n}$ samples of Subsection III-A) of the edge image. Instead of considering a functional approximation problem, we may lower our ambition and consider an up-sampling problem: given the low-resolution samples I_{LP} , find higher resolution samples I (for instance, samples that would be obtained for a smaller value of T_x and T_y in (6)). Specifically, the up-sampling problem is formulated as a Fourier

domain extrapolation:

$$\min_{\hat{I}_{k,l}} \sum_{k,l} \left| \left(\left(\frac{2\pi k}{\tau_x} \right)^2 + \left(\frac{2\pi l}{\tau_y} \right)^2 \right) \hat{I}_{k,l} \right|^2$$

$$\text{subject to } \hat{I}_{k,l} = [\hat{I}_{\text{LP}}]_{k,l} \text{ for } k, l \in \text{low frequency regions}, \quad (15)$$

$$c_{k,l} * \left(\frac{2\pi k}{\tau_x} + j \frac{2\pi l}{\tau_y} \right) \hat{I}_{k,l} = 0, \quad (16)$$

where \hat{I} and \hat{I}_{LP} are the DFT coefficients of the up-sampled image I and the known low-resolution samples I_{LP} respectively. Here we use the Laplacian of the up-sampled image (in the Fourier domain) as a smoothness-inducing regularizer. Observe that in addition to the usual data-fidelity constraint (15), annihilation of the up-sampled image (16) is also enforced. The reason is that the DFTs of both the up-sampled image and the low-resolution one are samples of the exact Fourier transform of the *same* continuously defined edge image. Therefore they should all satisfy the annihilation equation (4) and be annihilated by the *same* annihilation coefficients $c_{k,l}$.

When the edge curve cannot be annihilated exactly (model mismatch, noise, etc.), it is reasonable to integrate the annihilation constraint as a part of the objective functional to be minimized:

$$\min_{\hat{I}_{k,l}} \sum_{k,l} \left| \left(\left(\frac{2\pi k}{\tau_x} \right)^2 + \left(\frac{2\pi l}{\tau_y} \right)^2 \right) \hat{I}_{k,l} \right|^2 + \lambda \left\| c_{k,l} * \left(\frac{2\pi k}{\tau_x} + j \frac{2\pi l}{\tau_y} \right) \hat{I}_{k,l} \right\| \quad (17)$$

$$\text{subject to } \hat{I}_{k,l} = [\hat{I}_{\text{LP}}]_{k,l} \text{ for } k, l \in \text{low frequency regions},$$

where λ is a tunable weight and $\| \cdot \|$ is an arbitrary (unspecified) norm (e.g., ℓ_2). In this sense, we can claim that the annihilation constraint behaves as an “edge-preserving” interpolation technique in the up-sampling problem.

IV. GENERALIZATION FROM EDGE IMAGES TO NATURAL IMAGES

A. Annihilation-driven Interpolation

Until this section, our focus has been on synthetic curves that satisfy the annihilable curve model (3) *exactly*, which can be perfectly reconstructed from a finite number of samples of the associated edge image. In that context, we have proposed to use the annihilation equation in a regularization framework to retrieve a good approximation of the FRI edge image (see Section III-E). We would now like to apply this two-stage algorithm (annihilation, followed by annihilation-driven upsampling) to natural images. A direct implementation of the Fourier algorithm that solves (17) yields readily good results as can be seen in Fig. 8(a)). Even though the images that we consider do not satisfy the analyticity




Fig. 7. Natural images satisfies *approximately* the analyticity assumption requirement for the annihilation. (a) The original image and (b) the modulus of the complex derivative/gradient image.

condition, their (Wirtinger) derivatives can still outline the distinction between the edges and smooth areas—images usually have significantly larger gradients for edges (see, e.g. Fig. 7). That is why the annihilable curve model still holds approximately and can be obtained with the denoising algorithms presented in Section III-D.

B. Actual Implementation with a Spatial Domain Formulation

Even though we can directly apply the up-sampling algorithm (17) to natural images, it usually demands enormous amount of memory with our current implementation⁴. Instead, by adopting the new interpretation of annihilation in Section II-C2, we can *approximate* the exact annihilation constraint, which is a Fourier domain convolution, with a spatial domain multiplication and reformulate (17) directly in the spatial domain:

$$\begin{aligned} \min_{\mathbf{I}} \quad & \|\Delta \mathbf{I}\|_2^2 + \lambda \overbrace{\|\mathbf{M} \mathbf{D} \mathbf{I}\|_k^k}^{\text{annihilation}} \\ \text{subject to} \quad & \Phi \mathbf{I} = \mathbf{I}_{\text{LP}}, \end{aligned} \quad (18-\ell_k)$$

where $k = 1, 2$ depending on how we would like to enforce the annihilation constraint (see examples in Section V-C2). Here \mathbf{I} is the up-sampled image and \mathbf{I}_{LP} is the given low-resolution one (both are rearranged as column vectors); Δ and \mathbf{D} are the convolution matrices associated with a discrete Laplacian filter and the first-order derivatives respectively; Φ models the sampling process (i.e. low-pass filtering

⁴It is because natural edges are more complicated and hence requires an annihilation model with many more degrees of freedom compared with the simple synthetic curves that we have considered so far.




Fig. 8. Upsampling by a factor 3 of the small inset image in (a) and (b) using two methods: in (a) we use the exact Fourier domain formulation, in (b) we use the approximate spatial domain formulation. Either method yields a similar result: SNR between (a) and (b) is 34.44 dB (order of the FRI curve model: 31×31).

and down-sampling); and \mathbf{M} is a diagonal matrix whose diagonal entries are specified by the annihilation mask $\mathbf{M}_{i,i} = \mu_i$, which is obtained directly from the given low-resolution image (i.e. samples).

Experimentally, the approximated spatial domain approach gives very similar up-sampling results compared to the exact Fourier domain formulation (Fig. 8). However, the spatial domain formulation (18) allows us to cope with more general cases without introducing too many difficulties in the implementation: all we need to do is adapt the matrix Φ to account for different convolution kernels. On the contrary, the Fourier-domain formulation is somewhat “restricted” to the ideal low-pass kernel in (17), which ensures that the up-sampled and the known low-resolution images share the same low-frequency discrete Fourier coefficients (Fourier-domain extrapolation). Note that our current implementation is not optimized: we either solve directly a linear system of equations (ℓ_2 case), or use an iterated reweighted least-squares algorithm (ℓ_1 case), with the help of Matlab `sparse` toolbox. In this context, the computation time is of the order of 10 seconds (ℓ_2 case) and 40 seconds (ℓ_1 case), for images of size 255×255 on a MacBook Pro, 2.3GHz Intel Core i7, 16Gb RAM.




Fig. 9. Exact reconstruction of the annihilable curve C with minimum number of samples of the associated edge image I_C (see, Theorem 2). The discontinuity across the curve is $f_0 = x + jy$. (a) Modulus of the continuous domain edge image. (b) Modulus of the noiseless samples of the edge image (size: 9×9). (c) Perfectly reconstructed annihilable curve C .

V. EXPERIMENTAL RESULTS

A. Synthetic FRI Curves

1) *Exact Reconstruction from Noiseless Samples:* We exemplify the perfect reconstruction of the annihilable curve by sampling the associated edge image at the minimum sampling rate specified in Theorem 2. Consider an annihilable curve defined in (3) with 5×5 annihilation coefficients $c_{k,l}$. The discontinuity across the curve is $f_0(x, y) = x + jy$. The edge image is then ideally low-pass filtered and uniformly sampled. In our case here, the minimum number of samples required to perfectly reconstruct the annihilable curve is 81 (9×9). The reconstructed annihilable curve is exact (Fig. 9).

2) *Reconstruction from Noisy Samples:* Consider the same annihilable curve and the associated edge image as in Section V-A1, where we demonstrated perfect reconstruction from noiseless samples. Different levels of additive Gaussian white noise are added to the ideally low-pass filtered samples such that the signal-to-noise ratio (SNR) is 35 dB (mild distortion) and 15 dB (severe distortion) respectively. The reconstructed annihilable curves with the least square method (Fig. 10 (d) and (g)) and with the additional Cadzow's denoising (Fig. 10 (e) and (h)) are compared. The extra Cadzow's denoising improves the robustness of the annihilating filter method particularly in the presence of high noise level.

Such an improvement does not come without cost: the additional Cadzow's iteration is more computationally intensive as it computes the SVD of a larger convolution matrix several times (because of the iterations) compared with the simple one-step least square minimization. Also, we observe that Cadzow's algorithm usually requires many more iterations to meet the same stopping criterion as the one in the 1D cases [1], [7]. A direction for future research is to devise a more efficient algorithm that would be




Fig. 10. Reconstruction of an annihilable curve from noisy samples. (a) Noisy samples (SNR = 35 dB). (d) Noisy samples (SNR = 15 dB). (b) and (e) Reconstruction with least square minimization. (c) and (f) Reconstruction with additional Cadzow's denoising. The solid black line is the reconstructed annihilable curve while the dotted red line is the ground truth.

able to cope with severe model mismatch cases.

3) *Reconstruction of Non-annihilable Curves:* One potential source of uncertainties arises from the fact that curves that we want to reconstruct may not always satisfy the exact annihilable curve model (3). Yet, we want to show that this curve model is flexible enough to cope with more general curves. Specifically, we create a “non-annihilable” curve by interpolating (with cubic spline) a few points on an annihilable curve (Fig. 11). Then we obtain an annihilable curve model approximation⁵ of the new curve by applying the annihilation algorithm to the samples of the edge image. The reconstructed curve model is a rather accurate descriptor of the spline interpolated curve (which is *not* annihilable).

4) *Reconstruction of Annihilable Curves for Non-analytic Amplitudes:* It is essential for the exact annihilation that the amplitudes on the image be specified by a function that is analytic except on the image edges as discussed in Section II-C2. However, natural images are real-valued; hence, if this function is not piecewise constant, analyticity will be violated and the annihilation equation is *not* expected to be

⁵In the experiment, we have assumed that the annihilable curve approximation of the spline curve has *the same* degrees of freedom as the original annihilable curve which it derives from.




Fig. 11. The annihilable curve model is able to represent accurately non-annihilable curve (see, text in Section V-A3). (a) Consider a non-annihilable curve (dashed blue line) built using a cubic spline interpolation of a few samples (red circles) of an annihilable curve (solid black line). (b) Edge image associated with the non-annihilable curve. (c) Samples of the edge image (9×9 pixels). (d) The reconstructed annihilable curve (solid black line) based on the samples in (c) is barely distinguishable from the ground truth non-annihilable curve (dashed blue line).




Fig. 12. Reconstruction of an annihilable curve from samples of a real-valued edge image. (a) Edge image. (b) and (d) Samples of the edge image with different sizes ((b): 13×13 and (d): 21×21). (c) and (e) Reconstructions with Cadzow's iterative denoising from the given samples.

satisfied exactly even with synthetic annihilable curves. We want to evaluate the influence of this (strong) model mismatch on the accuracy of curve reconstruction.

To this end, we consider the same annihilable edge image as in the previous section except that




Fig. 13. Implicit reconstruction of the amplitudes on an annihilable curve via image up-sampling. (a) The given low-resolution samples (size: 9×9). (b) Amplitude reconstructions via image up-sampling (size: 315×315). (c) Actual (ideal low-pass filtered) samples of the edge image at the same resolution as (b). The SNR between (b) and (c) is 34.90 dB.

the discontinuity across the curve is now real-valued: $f_0(x, y) = \Re\{x + jy\} = x$. The reconstructed annihilation models from two different numbers of samples with Cadzow's iterative algorithm are shown in Fig. 12. We observe that more samples are required in order to have a more accurate reconstruction when a real-valued edge image is involved.

B. Synthetic FRI Edge Image Amplitudes

We demonstrate the amplitude reconstruction of an edge image in Fig. 13 where we have used the same edge image as the one in Section V-A1. The annihilation coefficients $c_{k,l}$ are first retrieved from the given 9×9 samples before we enforce the FRI curve model in the up-sampling problem. The high signal-to-noise ratio (SNR = 34.90 dB) between the up-sampled image (Fig. 13 (b)) and the actual ground truth (Fig. 13 (c)) samples at the same resolution clearly shows that we have a good estimate of the amplitudes on the annihilable curve.

C. Natural Images

1) *FRI Curve Approximation of Natural Edges*: In the experiments, we apply the annihilation algorithm to a low-resolution image, which are samples of the ground truth high-resolution image (low-pass filtered and down-sampled)⁶. In the case when non-ideal samples are involved, we first obtain an estimate of the ideal low-pass filtered samples by applying additional post-filtering (see details in Section III-C). The

⁶We consider both the sinc and bicubic kernels with a 3 : 1 down-sampling ratio, i.e. we take every pixel out of three from the low-pass filtered image.




Fig. 14. Reconstruction of FRI curve model for natural images from (a) ideal and (d) Bicubic low-pass filtered samples of the high-resolution ground truth image. (b) and (e) Reconstructed mask function. (c) and (f) Roots of the mask function, i.e. the annihilable curve. The order of the FRI model is 29×29 .

annihilating filter is then built from the Fourier coefficients of the image, which are directly related to the ideal low-pass filtered samples.

The reconstructed annihilation models, including the mask function and its roots (i.e. the reconstructed curves), are shown in Fig. 14. If we look at the reconstructed curve alone, then the results may not appear very promising. We need to emphasize, however, that the reconstructed edge model means more than a simple binary edge map (contrary to most cases with conventional edge detection algorithms): the reconstructed mask function is used to assign position-dependent weights to preserve edges in further processing (see, e.g. Section IV-A).

2) *Annihilation-driven Image Up-sampling*: We exemplify the edge-preserving interpolation algorithm (18) with several images shown in Fig. 15. The ground truth high-resolution image is first low-pass




Fig. 15. Test images considered in the edge-preserving interpolation experiments. The given low-resolution image, which is a low-pass filtered and down-sampled version of the ground truth high-resolution image, is also shown on the lower right corner of each image.

TABLE II

COMPARISONS OF DIFFERENT UP-SAMPLED IMAGES WITH THE GROUND TRUTH HIGH-RESOLUTION IMAGE.

Image	bicubic interpolation (<code>imresize</code>)	minimizing total variation	w/o annihilation constraint (19)	w/ annihilation constraint ($18-\ell_2$)	w/ annihilation constraint ($18-\ell_1$)	learning-based algorithm [25], [27]	learning-based algorithm [29]
<i>chip</i>	26.43 dB	27.98 dB	27.61 dB	28.90 dB	29.14 dB	28.24 dB	28.29 dB
<i>peppers</i>	28.41 dB	30.31 dB	29.62 dB	30.27 dB	30.89 dB	30.00 dB	30.43 dB
<i>bank</i>	21.45 dB	23.06 dB	22.33 dB	22.80 dB	23.02 dB	22.26 dB	22.91 dB
<i>MRI</i>	23.95 dB	24.63 dB	25.25 dB	25.57 dB	25.70 dB	25.15 dB	25.40 dB

filtered with bicubic kernel and then down-sampled with a 3 : 1 down-sampling ratio⁷. The given low-resolution image (i.e. samples) is then up-sampled 3 times and compared with the ground truth image (measured in terms of PSNR). Both the ℓ_2 and ℓ_1 formulations (18) are considered. The regularization weight λ for the annihilation constraint is set to $\lambda_{\ell_2} = 5 \times 10^5$ and $\lambda_{\ell_1} = 2 \times 10^8$ for the ℓ_2 and ℓ_1 formulations respectively, and for all the images. The up-sampled images are compared with images obtained by minimizing the smoothness regularization subject to the data-fidelity constraint only, *without* enforcing the annihilation constraint:

$$\begin{aligned} \min_{\mathbf{I}} \quad & \|\Delta \mathbf{I}\|_2^2 \\ \text{subject to} \quad & \Phi \mathbf{I} = \mathbf{I}_{LP}. \end{aligned} \tag{19}$$

Another standard nonlinear image up-sampling formulation is also included for comparisons, where the up-sampled image is obtained by minimizing the total-variation (TV), which is known to encourage

⁷We have observed an offset (variable, depending on the sampling ratio) between the samples and the `imresize` output. We have adjusted the sampling grid accordingly such that the samples match the `imresize` result.




Fig. 16. Comparisons of different image up-sampling results. (a) Ground truth image. (b) Standard total variation minimizing result (PSNR = 30.31 dB). (c) Up-sampled image without annihilation constraint (PSNR = 29.62 dB). (d) Up-sampled image with annihilation constraint ($18-\ell_1$) (PSNR = 30.89 dB).

piecewise smoothness without jeopardizing sharp edges [22], [23]. We have implemented the majorization minimization algorithm described in [24] to solve the TV minimization problem.

Table II summarizes the up-sampling results obtained with different algorithms. The gain due to the annihilation constraint ranges from 0.5 dB to 1.5 dB compared with the results from (19) (when the ℓ_1 formulation ($18-\ell_1$) is used). It not only makes the edges of the up-sampled image sharper but also reduces the ringing artifacts significantly (Fig. 16). Compared with total-variation minimization, the annihilation constraint has a better balance⁸ between preserving sharp edges *without* introducing artificial edges. Note that recent developments in computer vision [25]–[29] have given rise to some promising training-based approaches for image up-sampling, e.g. single-image super-resolution [26], [29]. Our preliminary experiments show that the annihilation-driven approach is still quite competitive with these algorithms. However, our intention here is not to have comprehensive comparisons with state-of-the-art algorithms but rather to exemplify that the curve annihilation is a method that has the potential to bring substantial improvements in, e.g., super-resolution applications.

Observe that enforcing the annihilation constraint in the ℓ_1 -norm sense generally leads to 0.2 ~ 0.7 dB further gain compared with the least-square ℓ_2 formulation ($18-\ell_2$). One explanation might be that the ℓ_2 -norm penalizes the annihilation errors *equally* for all pixels in the whole image. With other alternatives, e.g. the ℓ_1 -norm, the annihilation mask gets more efficient in differentiating between the edge regions and the smooth areas in the image, hence yielding a higher SNR for the up-sampled image.

⁸Arguably, this is related to the regularization weight λ that we choose for the annihilation constraint.

VI. CONCLUSION

In this paper, we have extended one-dimensional sparse annihilation-based algorithms to two-dimensional annihilation algorithms in order to retrieve a specific class of curves, which are defined implicitly as the roots of a mask function. We have shown that it is possible to have a perfect reconstruction of the implicit curve representation from a finite number of samples of the associated edge image. Further, the exact curve annihilation model is generalized to describe edges of natural images, which can be used as a *linear* constraint in a practical image processing problem, e.g. image up-sampling.

The mask function we considered in this paper is expressed as a weighted summation of sinusoids, which is not a local representation. Hence, it may not be well adapted to characterize images with edges that are non-uniformly distributed. We may consider replacing the complex exponentials with other locally-varying basis, e.g. spline kernels [10], [19], in the future. Another possible extension of the curve annihilation is to obtain local annihilation models from image patches.

APPENDIX A

PROOFS AND DERIVATIONS

A. Derivations of the Sampling Formula (6)

The ideal low-pass filtered samples $g_{m,n}$ of the (τ_x, τ_y) -periodized edge image $I_{C_{\text{per}}}$ are:

$$g_{m,n} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I_{C_{\text{per}}}(x, y) \text{sinc}(B_x(mT_x - x)) \text{sinc}(B_y(nT_y - y)) dx dy.$$

Using the relation between $I_{C_{\text{per}}}$ and I_C

$$I_{C_{\text{per}}}(x, y) = \sum_{k \in \mathbb{Z}} \sum_{l \in \mathbb{Z}} I_C(x + k\tau_x, y + l\tau_y),$$

we find that

$$\begin{aligned} g_{m,n} &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I_C(x, y) \sum_{k \in \mathbb{Z}} \sum_{l \in \mathbb{Z}} \text{sinc}(B_x(mT_x - x - k\tau_x)) \text{sinc}(B_y(nT_y - y - l\tau_y)) dx dy \\ &= \int_0^{\tau_x} \int_0^{\tau_y} I_C(x, y) \varphi(mT_x - x, nT_y - y) dx dy, \end{aligned}$$

where

$$\begin{aligned} \varphi(x, y) &= \sum_{k \in \mathbb{Z}} \sum_{l \in \mathbb{Z}} \text{sinc}(B_x(x - k\tau_x)) \text{sinc}(B_y(y - l\tau_y)) \\ &= \frac{\sin(\pi B_x x) \sin(\pi B_y y)}{B_x B_y \tau_x \tau_y \sin(\pi x / \tau_x) \sin(\pi y / \tau_y)}. \end{aligned}$$

B. Correspondence between the Spatial and Fourier Domain Samples

By replacing the Fourier series representation (7) of the edge image in (5), we have

$$\begin{aligned}
g_{m,n} &= \sum_{k \in \mathbb{Z}} \sum_{l \in \mathbb{Z}} \hat{I}_{k,l} e^{j \frac{2k\pi}{\tau_x} x + j \frac{2l\pi}{\tau_y} y} * (\text{sinc}(B_x x) \text{sinc}(B_y y)) \Big|_{\substack{x=mT_x \\ y=nT_y}} \\
&= \sum_{k \in \mathbb{Z}} \sum_{l \in \mathbb{Z}} \hat{I}_{k,l} e^{j \frac{2k\pi}{\tau_x} mT_x + j \frac{2l\pi}{\tau_y} nT_y} \frac{1}{B_x B_y} \text{rect}\left(\frac{k}{\tau_x B_x}\right) \text{rect}\left(\frac{l}{\tau_y B_y}\right) \\
&= \frac{1}{B_x B_y} \sum_{\substack{|k| \leq \lfloor B_x \tau_x / 2 \rfloor \\ |l| \leq \lfloor B_y \tau_y / 2 \rfloor}} \hat{I}_{k,l} e^{j 2\pi m k / N_x + j 2\pi n l / N_y},
\end{aligned}$$

where we have used the fact that $\tau_x = N_x T_x$, $\tau_y = N_y T_y$, and that the Fourier transform of $\text{sinc}(x)$ is $\text{rect}(\omega/(2\pi))$. On the other hand, we have the discrete Fourier transform (DFT) representation of $g_{m,n}$ as

$$g_{m,n} = \frac{1}{N_x N_y} \sum_{k=0}^{N_x-1} \sum_{l=0}^{N_y-1} \hat{g}_{k,l} e^{j 2\pi m k / N_x + j 2\pi n l / N_y}.$$

Hence, we have

$$\hat{g}_{k,l} = \begin{cases} \frac{N_x N_y}{B_x B_y} \hat{I}_{k,l} & \text{for } |k| \leq \lfloor B_x \tau_x / 2 \rfloor \text{ and } |l| \leq \lfloor B_y \tau_y / 2 \rfloor, \\ 0 & \text{otherwise,} \end{cases}$$

C. Optimal Post-filter Applied to Non-ideal Samples




Fig. 17. Optimal post-filter applied to non-ideal low-pass samples (1-D case).

For the sake of simplicity, let us consider the one-dimensional case in the derivations. The 2-D result is a direct extension of the derived results here. We use notations that are as consistent as possible with the 2-D case.

Our strategy to estimate the ideal low-pass filtered samples is first to obtain a continuous domain description of the underlying signal from the given non-ideal samples via interpolation and then to apply the standard ideal low-pass filtering and sampling to this continuous signal. Consider the sampling setup in Fig. 17, where φ is an anti-aliasing filter, ψ is an interpolation kernel, and w is the ideal low-pass

sampling kernel. From the sampling process, we have the samples are:

$$g_n = \int_{-\infty}^{\infty} I(t)\varphi(nT - t)dt.$$

Based on the given discrete samples g_n we can estimate a continuous domain description of the original signal $I(t)$ via interpolation:

$$J(t) = \sum_{k \in \mathbb{Z}} g_k \psi(t - kT),$$

for some interpolation kernel $\psi(t)$.

The primary consideration in the design of the optimal post-filter is to ensure the *consistency* constraint— if we resample the interpolated signal, we should get the same samples which we started from:

$$\int_{-\infty}^{\infty} J(t)\varphi(nT - t)dt = g_n.$$

Hence,

$$\begin{aligned} g_n &= \int_{-\infty}^{\infty} \sum_{k \in \mathbb{Z}} g_k \psi(t - kT) \varphi(nT - t) dt \\ &= \sum_{k \in \mathbb{Z}} g_k \int_{-\infty}^{\infty} \psi(t) \varphi(nT - kT - t) dt = \sum_{k \in \mathbb{Z}} (\psi * \varphi)((n - k)T) g_k. \end{aligned} \quad (20)$$

This means that $(\psi * \varphi)(kT) = \delta_k$. The optimal choice for ψ (in the least-square sense, see [4]) is the *dual* of the sampling kernel φ

$$\hat{\psi}(\omega) = \frac{T \hat{\varphi}(\omega)^*}{\sum_{n \in \mathbb{Z}} |\hat{\varphi}(\omega + \frac{2n\pi}{T})|^2}.$$

For the second stage of sampling with the ideal low-pass sampling kernel, we have a similar equation as in (20)

$$\tilde{g}_n = \sum_{k \in \mathbb{Z}} \underbrace{(\psi * w)((n - k)T)}_{\tilde{\varphi}_{n-k}} g_k.$$

Using Poisson's summation formula, the DTFT of $\tilde{\varphi}_n$ is

$$\begin{aligned} \sum_{n \in \mathbb{Z}} \tilde{\varphi}_n e^{-jnT\omega} &= \sum_{n \in \mathbb{Z}} (\psi * w)((n - k)T) e^{-jnT\omega} \\ &= \frac{1}{T} \sum_{n \in \mathbb{Z}} \hat{\psi}\left(\omega + \frac{2n\pi}{T}\right) \hat{w}\left(\omega + \frac{2n\pi}{T}\right) \\ &= \frac{\sum_{n \in \mathbb{Z}} \hat{\varphi}\left(\omega + \frac{2n\pi}{T}\right)^* \hat{w}\left(\omega + \frac{2n\pi}{T}\right)}{\sum_{n \in \mathbb{Z}} |\hat{\varphi}\left(\omega + \frac{2n\pi}{T}\right)|^2}. \end{aligned} \quad (21)$$

This is the frequency response of the post-filter needed to process the non-ideal samples g_n to obtain close to ideal samples. For the 2-D case, the result is a direct extension of (21) as shown in (13).

ACKNOWLEDGMENT

T. Blu and H. Pan are supported in part by an RGC grant #CUHK410110 of the Hong Kong University Grant Council and in part by the Prof. Charles K. Kao Research Exchange scheme. P. L. Dragotti is supported by the European Research Council (ERC) starting investigation award #277800 (RECOAMP).

REFERENCES

- [1] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1417–1428, 2002.
- [2] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, 1948.
- [3] A. Aldroubi and K. Gröchenig, "Nonuniform sampling and reconstruction in shift-invariant spaces," *SIAM review*, vol. 43, no. 4, pp. 585–620, 2001.
- [4] T. Blu and M. Unser, "Quantitative Fourier analysis of approximation techniques: Part I—Interpolators and projectors," *IEEE Transactions on Signal Processing*, vol. 47, no. 10, pp. 2783–2795, October 1999.
- [5] —, "Quantitative Fourier analysis of approximation techniques: Part II—Wavelets," *IEEE Transactions on Signal Processing*, vol. 47, no. 10, pp. 2796–2806, October 1999.
- [6] M. Unser, "Sampling—50 years after Shannon," *Proceedings of the IEEE*, vol. 88, no. 4, pp. 569–587, 2000.
- [7] T. Blu, P. L. Dragotti, M. Vetterli, P. Marziliano, and L. Coulot, "Sparse sampling of signal innovations," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 31–40, 2008.
- [8] J. Berent, P. L. Dragotti, and T. Blu, "Sampling piecewise sinusoidal signals with finite rate of innovation methods," *IEEE Transactions on Signal Processing*, vol. 58, no. 2, pp. 613–625, 2010.
- [9] I. Maravić and M. Vetterli, "Exact sampling results for some classes of parametric nonbandlimited 2-D signals," *IEEE Transactions on Signal Processing*, vol. 52, no. 1, pp. 175–189, 2004.
- [10] P. Shukla and P. L. Dragotti, "Sampling schemes for multidimensional signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 55, no. 7, pp. 3670–3686, 2007.
- [11] I. Maravić and M. Vetterli, "A sampling theorem for the Radon transform of finite complexity objects," in *2002 IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP)*, vol. 2. IEEE, 2002, pp. II–1197.
- [12] C. Chen, P. Marziliano, and A. C. Kot, "2D finite rate of innovation reconstruction method for step edge and polygon signals in the presence of noise," *IEEE Transactions on Signal Processing*, vol. 60, no. 6, pp. 2851–2859, 2012.
- [13] P. Milanfar, G. C. Verghese, W. C. Karl, and A. S. Willsky, "Reconstructing polygons from moments with connections to array processing," *IEEE Transactions on Signal Processing*, vol. 43, no. 2, pp. 432–443, 1995.
- [14] B. Gustafsson, C. He, P. Milanfar, and M. Putinar, "Reconstructing planar domains from their moments," *Inverse Problems*, vol. 16, no. 4, p. 1053, 2000.
- [15] M. Elad, P. Milanfar, and G. H. Golub, "Shape from moments—An estimation theory perspective," *IEEE Transactions on Signal Processing*, vol. 52, no. 7, pp. 1814–1829, 2004.
- [16] S. M. Kay, *Modern spectral estimation: Theory and application*. Prentice Hall, NJ, 1988.
- [17] P. Stoica and R. L. Moses, *Introduction to spectral analysis*. Prentice Hall Upper Saddle River, NJ, 1997, vol. 89.
- [18] S. M. Kay and S. L. Marple Jr, "Spectrum analysis—A modern perspective," *Proceedings of the IEEE*, vol. 69, no. 11, pp. 1380–1419, 1981.

- [19] P. L. Dragotti, M. Vetterli, and T. Blu, "Sampling moments and reconstructing signals of finite rate of innovation: Shannon meets Strang-Fix," *IEEE Transactions on Signal Processing*, vol. 55, no. 5, pp. 1741–1757, 2007.
- [20] D. W. Tufts and R. Kumaresan, "Estimation of frequencies of multiple sinusoids: Making linear prediction perform like maximum likelihood," *Proceedings of the IEEE*, vol. 70, no. 9, pp. 975–989, 1982.
- [21] J. A. Cadzow, "Signal enhancement—A composite property mapping algorithm," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 36, no. 1, pp. 49–62, 1988.
- [22] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1, pp. 259–268, 1992.
- [23] A. Chambolle, "An algorithm for total variation minimization and applications," *Journal of Mathematical imaging and vision*, vol. 20, no. 1-2, pp. 89–97, 2004.
- [24] J. M. Bioucas-Dias, M. A. Figueiredo, and J. P. Oliveira, "Total variation-based image deconvolution: a majorization-minimization approach," in *2006 IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP)*, vol. 2. IEEE, 2006, pp. 861–864.
- [25] J. Yang, J. Wright, T. Huang, and Y. Ma, "Image super-resolution as sparse representation of raw image patches," in *IEEE Conference on Computer Vision and Pattern Recognition, 2008. CVPR 2008*. IEEE, 2008, pp. 1–8.
- [26] M. Elad and D. Datsenko, "Example-based regularization deployed to super-resolution reconstruction of a single image," *The Computer Journal*, vol. 52, no. 1, pp. 15–30, 2009.
- [27] J. Yang, J. Wright, T. S. Huang, and Y. Ma, "Image super-resolution via sparse representation," *IEEE Transactions on Image Processing*, vol. 19, no. 11, pp. 2861–2873, 2010.
- [28] J. Yang, Z. Wang, Z. Lin, S. Cohen, and T. Huang, "Coupled dictionary training for image super-resolution," *IEEE Transactions on Image Processing*, vol. 21, no. 8, pp. 3467–3478, 2012.
- [29] R. Zeyde, M. Elad, and M. Protter, "On single image scale-up using sparse-representations," in *Curves and Surfaces*. Springer, 2012, pp. 711–730.



Hanjie Pan (S'11) was born in Jiangsu, China, in 1988. In 2010, he received the B.Eng. degree with first honor in Electronic Engineering from The Chinese University of Hong Kong (CUHK), Shatin, Hong Kong. He received the M.Phil degree from the same institution in 2013. He is currently a PhD student with the School of Computer and Communication Sciences (IC), École Polytechnique Fédérale de Lausanne (EPFL). In the summer 2010, he was a visiting student at Imperial College, London, United Kingdom. His research interests include image restorations with sparsity constraints and sampling signals with finite rate of innovation.



Thierry Blu (F'12) was born in Orléans, France, in 1964. He received the “Diplôme d’ingénieur” from École Polytechnique, France, in 1986 and from Télécom Paris (ENST), France, in 1988. In 1996, he obtained a Ph.D in electrical engineering from ENST for a study on iterated rational filterbanks, applied to wideband audio coding.

Between 1998 and 2007, he was with the Biomedical Imaging Group at the Swiss Federal Institute of Technology (EPFL) in Lausanne, Switzerland. He is now a Professor in the Department of Electronic Engineering, The Chinese University of Hong Kong.

Dr. Blu was the recipient of two best paper awards from the IEEE Signal Processing Society (2003 and 2006). He is also coauthor of a paper that received a Young Author best paper award (2009) from the same society. He has been an Associate Editor for the IEEE Transactions on Image Processing (2002–2006), the IEEE Transactions on Signal Processing (2006–2010), Elsevier Signal Processing (2008–2011). He is currently on the board of EURASIP J. on Image and Video Processing (since 2010) and is a member of the IEEE Signal Processing Theory and Methods Technical Committee (since 2008).

Research interests: (multi)wavelets, multiresolution analysis, multirate filterbanks, interpolation, approximation and sampling theory, sparse sampling, image denoising, psychoacoustics, biomedical imaging, optics, wave propagation...



Pier Luigi Dragotti is currently a Reader (Associate Professor) with the Electrical and Electronic Engineering Department at Imperial College London, U.K. He received the Laurea degree (summa cum laude) in electrical and electronic engineering from the University of Naples Federico II, Naples, Italy, in 1997; the master’s degree in communications systems from the Swiss Federal Institute of Technology of Lausanne (EPFL), Lausanne, Switzerland, in 1998, and the Ph.D. degree from EPFL in April 2002. He has held several visiting positions at different universities and research centers including: Visiting Student

with Stanford University, CA, USA, in 1996, Visiting Researcher with the Mathematics of Communications Department, Bell Labs, Lucent Technologies, Murray Hill, NJ, USA, in 2000, Visiting Professor with the Politecnico of Milan, Italy, in 2010, and Visiting Scientist with the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2011. He was Technical Co-Chair for the European Signal Processing Conference in 2012, Associate Editor of the IEEE TRANSACTIONS ON IMAGE PROCESSING from 2006 to 2009 and an Elected Member of the IEEE Image, Video and Multidimensional Signal Processing Technical Committee and the IEEE Signal Processing Theory and Methods Technical Committee. He is a recipient of the ERC Starting Investigator Award for the project RecoSamp. His research interests include sampling theory, wavelet theory and its applications, image and video compression, image-based rendering, and image super-resolution.