

Anisotropic Interpolation of Sparse, Generalized Image Samples

Aurélien Bourquard, *Student Member, IEEE*, and Michael Unser, *Fellow, IEEE*

Abstract—Practical image-acquisition systems are often modeled as a continuous-domain prefilter followed by an ideal sampler, where generalized samples are obtained after convolution with the impulse response of the device. In this paper, our goal is to interpolate images from a given subset of such samples. We express our solution in the continuous domain, considering consistent resampling as a data-fidelity constraint. To make the problem well-posed and ensure edge-preserving solutions, we develop an efficient anisotropic regularization approach that is based on an improved version of the edge-enhancing anisotropic diffusion equation. Following variational principles, our reconstruction algorithm minimizes successive quadratic cost functionals. To ensure fast convergence, we solve the corresponding sequence of linear problems by using multigrid iterations that are specifically tailored to their sparse structure. We conduct illustrative experiments and discuss the potential of our approach both in terms of algorithmic design and reconstruction quality. In particular, we present results that use as little as two percent of the image samples.

Index Terms—Anisotropic diffusion, diffusion tensors, edge-enhancing diffusion, generalized sampling, image interpolation, image magnification, image reconstruction, inverse problems, iteratively reweighted least squares, multigrid techniques, PDE-based methods.

I. INTRODUCTION

Shannon's sampling theorem [1] states that a bandlimited signal can be perfectly reconstructed from its samples, provided that Nyquist's criterion is satisfied. In that case, the solution can be found by sinc interpolation. However, in an imaging context, bandlimited functions do not correctly match the physical reality [2]. This classical assumption can thus lead to inaccurate results for such classes of problems. Specifically, optical acquisition systems do deviate from an ideal sampler in practice as they involve filtering associated with their impulse response prior to sampling. When taken into account, those effects impose *consistency constraints* on generalized samples [3], which makes the reconstruction more intricate. Several works have successfully dealt with this non-ideality [4], [5], [6], [7], [8], [9], yielding results that are visually sharper as compared to the standard sampling paradigm. Nevertheless, these approaches are typically focused on pure magnification cases, and have not been applied to sparse-interpolation problems.

Copyright (c) 2012 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.

The authors are with the Biomedical Imaging Group (BIG), École polytechnique fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland (email: aurelien.bourquard@epfl.ch; michael.unser@epfl.ch).

Digital Object Identifier xx.xxxx/TIP.2012.xxxxx

In this paper, we introduce a novel interpolation approach that simultaneously handles generalized and sparse image sampling. The objective of our method is to reconstruct a continuous image from a subset of its generalized samples. As a first step towards specifying our problem, we define a data-fidelity measure that is based on consistent resampling. The unknowns being under-constrained, regularization is needed to find a suitable solution.

Variational formulations are commonly employed for regularization in the literature. In particular, quadratic regularization has been previously considered to develop fast sparse-interpolation approaches [10], [11]. Despite their efficiency, these linear methods tend to blur image contours [4]. In order to produce edge-preserving reconstructions, nonquadratic functionals are required [4]. A classical solution is the total-variation (TV) norm [12], which is also associated with fast algorithms such as primal-dual approaches [13], [14]. However, the behavior of standard edge-preserving regularization techniques is not adapted to sparse interpolation for it produces singular points [15]. Thus, recent works on image interpolation involve more advanced formulations, especially when the sampling rate is low, or when the available data consist in a reduced set of samples as in our problem. Promising results have been obtained in the variational framework using a non-local approach [16]. Unfortunately, the associated computation time tends to be prohibitive.

Some of the most efficient regularization strategies in the area of inpainting and sparse interpolation involve partial differential equations (PDEs) that are based on *anisotropic diffusion*¹ [15], [17], [18]. The behavior of these methods can be fine-tuned via the specification of diffusion tensors. A high-quality technique based on an *edge-enhancing-anisotropic-diffusion* (EED) PDE has been devised by Galić *et al.* to interpolate sparse image samples [15]. This approach enjoys remarkable edge-reconstruction performance even at very high sparsity levels.

Nonlinear PDEs can be solved using explicit or semi-implicit schemes that are based on finite-difference approximations of the original formulation [19]. An alternate approach involving *lagged-diffusivity fixed-point iterations* has been also investigated for the TV flow, and subsequently for other types of isotropic diffusion equations [20], [21]. In order to obtain rapid convergence, each iteration involves a linearized version of the flow where the diffusivity terms are fixed according to some current solution estimate. Despite their

¹In this paper, *anisotropy* of a given diffusion process is understood in the sense of [15], implying tensor diffusivities. Though nonquadratic, the TV functional only acts as an isotropic regularizer following that definition.

increased complexity [22], tensor diffusivities can also be handled efficiently in large-scale problems. In that regard, a novel class of algorithms based on fast explicit diffusion (FED) has been devised in [23]. This approach follows a coarse-to-fine strategy, and allows to implement advanced PDE-based methods such as [15].

The above PDE-based regularization approaches are most efficient for interpolation and inpainting problems. Variational approaches, on the other hand, do result in efficient implementations for a larger class of inverse problems, including image restoration [24]. In particular, variational formulations are most adapted to our extended interpolation model. They allow to efficiently handle our specific data-fidelity constraint that involves an analysis kernel before sampling.

The distinction between both types of methods, however, is not clear-cut. Indeed, regularization-based methods are related to PDE formulations through the Euler-Lagrange equation [12]. Consequently, the basic steepest-descent method applied to the given cost functional is equivalent to the corresponding gradient flow. Similarly, the iteratively reweighted least-squares (IRLS) technique [25] that is used for nonquadratic regularization is associated with linearized versions of the gradient of the original functional [21], [26]. This provides an interpretation that relates IRLS to lagged-diffusivity fixed-point iterations.

In this work, we develop a hybrid regularization framework that combines advantageous aspects of both PDE and variational formulations using similar principles based on lagged diffusivities. The specificity of our design is that it stems directly from the definition of an anisotropic-diffusion equation. As a consequence, our regularizer consists in a series of quadratic functionals that are based on successive tensor-valued diffusivity estimates. While being adapted to our particular problem, this cost-functional approach exhibits similarities with the FED method of [23] where first-order approximations of the underlying diffusion process are taken. Regarding the actual specification of the flow, our central contribution is to propose our own PDE as an extension of the EED solution considered in [15]. In particular, we redefine the associated tensor diffusivities so as to further improve edge-reconstruction capability on natural images.

Although not strictly originating from a variational formulation, our regularization approach yields an anisotropic version of the IRLS technique [25]. Starting from a quadratic data-fidelity constraint, our reconstruction algorithm called *anisotropic IRLS* (AIRLS) entails the partial resolution of successive weighted linear problems. Since the diffusivity estimation of our method is constrained to weight updates, they do not affect the overall algorithmic performance significantly. We also devise a fast multigrid solver that is adapted to the sparse structure of our linear problems and that is inspired from previous works [10], [11]. Note that the obtained AIRLS framework can then be used to implement distinct regularization PDEs as well, including the one of [15].

The paper is organized as follows: In Section II, we present our continuous interpolation framework where the unknowns are expressed as coefficients in a shift-invariant basis. In Section III, we consider classical variational approaches to express

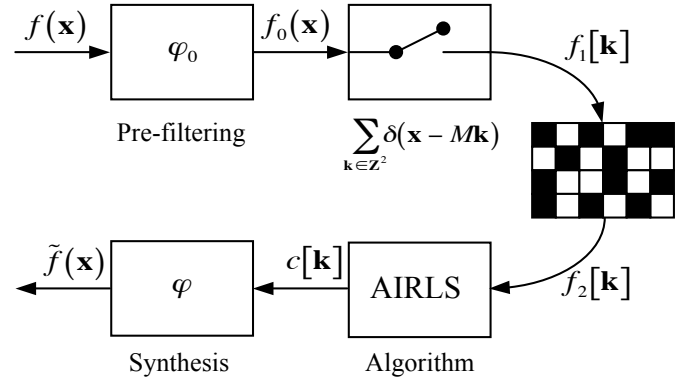


Fig. 1. The continuously defined image f is convolved with φ_0 (the impulse response of the acquisition device) before being sampled at integer intervals M along each dimension. The resulting sequence f_1 is then masked, which yields f_2 . Starting from these masked samples, our algorithm outputs the coefficients c of the reconstructed image. The continuously defined solution \tilde{f} can be then be obtained from these coefficients according to (3).

our reconstruction problem. Our actual strategy resulting in an IRLS-type procedure is introduced in Section IV. The associated linear problems are then specified in Section V, and their iterative resolution using our own multilevel approach is addressed in Section VI. In the experiments of Section VII, we consider distinct interpolation cases where our method is compared with the state of the art, both quantitatively and qualitatively. Implications of our results are finally discussed in Section VIII.

II. SAMPLING AND RECONSTRUCTION

A. Forward Model

As represented in Figure 1, the input signal of our model consists in the continuous-domain image f . Assuming generalized sampling, the latter is first convolved with a prefilter φ_0 that corresponds to the impulse response of the acquisition device [3]. The intermediate image is then sampled at integer intervals M along each dimension, which results in the sequence

$$f_1[\mathbf{k}] = (f * \varphi_0)(\mathbf{x})|_{\mathbf{x}=M\mathbf{k}}, \quad (1)$$

where $*$ denotes continuous-domain convolution. Only a subset of the sequence f_1 is retained through the binary mask w , which yields the masked samples

$$f_2[\mathbf{k}] = w[\mathbf{k}]f_1[\mathbf{k}]. \quad (2)$$

Assuming that the sampling process and the binary mask are known, we define our interpolation problem as the task of accurately reconstructing the original image f from the available samples in f_2 . We elaborate on our reconstruction approach below.

B. Reconstruction Space

Following the generalized sampling theory of [3], our reconstruction space is in the continuous domain and spanned by normalized translates of an analog generating kernel φ . Specifically, the reconstructed image \tilde{f} takes the form

$$\tilde{f}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^2} c[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k}), \quad (3)$$

where c is a discrete sequence of coefficients that describes the solution exactly. The reconstruction is defined on a grid that is M times finer than the acquisition in each dimension. In our implementation, the image data are defined over some rectangular domain Ω and are extended periodically outside.

The formulation (3) enables the solution to be computed and stored in terms of its coefficients, despite its continuous character. In this framework, we specify φ as a B-spline function [27] of order η , which makes straightforward sub-pixel post-processing (e.g., registration) of the reconstructed data possible, and allows to properly define our reconstruction approach. The interpolating B-spline is differentiable and has suitable approximation properties, such as reproduction of polynomials [27].

C. Constraints

In order to be accurate, the solution (3) has to be consistent with the available samples f_2 . While adopting the consistent-measurement principle of [3], we nevertheless want to accommodate for noise and model imperfections. Therefore, we propose a soft form of this constraint, demanding that \tilde{f} reintroduced in place of f into the generalized-sampling system of Figure 1 results in measurements \tilde{f}_2 that are close to f_2 . The image \tilde{f} and the samples \tilde{f}_2 are related in the same way as f and f_2 through (1) and (2). Accordingly, we propose to define the data discrepancy measure as the squared ℓ_2 -norm

$$\mathcal{D}(c) = \left\| f_2 - \tilde{f}_2 \right\|_{\ell_2}^2, \quad (4)$$

which is an implicit function of the expansion coefficients in (3). Note that, in the sequel, we shall use implicit functions of the solution coefficients when appropriate. As a soft constraint, we impose

$$\mathcal{D}(c) < \mathcal{K}, \quad (5)$$

where \mathcal{K} is a positive constant. The sequence \tilde{f}_2 is derived from the above relations as

$$\tilde{f}_2[\mathbf{k}] = w[\mathbf{k}] \left(\varphi_0 * \sum_{\mathbf{n} \in \mathbb{Z}^2} c[\mathbf{n}] \varphi(\cdot - \mathbf{n}) \right) (\mathbf{x}) \Big|_{\mathbf{x} = M\mathbf{k}}. \quad (6)$$

Given (6) and the binary nature of the mask, (4) can be rewritten in the explicit form

$$\mathcal{D}(c) = \sum_{\mathbf{k} \in \mathbb{Z}^2} w[\mathbf{k}] |f_2[\mathbf{k}] - \{b \star c\}_{1M}[\mathbf{k}]|^2, \quad (7)$$

where \star denotes discrete convolution, and where the sequence b is defined as

$$b[\mathbf{k}] = (\varphi_0 * \varphi)(\mathbf{x})|_{\mathbf{x}=\mathbf{k}}. \quad (8)$$

Therefore, we want any valid solution \tilde{f} to have a low discrepancy measure \mathcal{D} . In this work, we assume that the

prefilter φ_0 used in our generalized-sampling model (1) is nonnegative, which itself implies the non-negativity of b given the definition of φ as a B-spline.

Under (5), our reconstruction problem is still ill-posed. We thus have to define additional regularity constraints that make the problem well-posed. In the sequel, we discuss reconstruction approaches that satisfy this requirement while ensuring low data discrepancy.

III. EXISTING VARIATIONAL APPROACHES

The variational framework lends itself well to the specification of our reconstruction problem. It allows the solution to satisfy a constraint of the form (5) under suitable regularity criteria. In this section, we review some classical regularization functionals that are used for image reconstruction. Their properties as well as their links with IRLS and PDE formulations also serve as background for our own reconstruction method introduced in Section IV. Their expressions are readily introduced within our reconstruction framework, which allows to predefine relevant quantities and relations for the sequel. In this variational setting, the generic reconstruction problem is to minimize the functional

$$\mathcal{J}(c) = \mathcal{D}(c) + \lambda \mathcal{R}(c), \quad (9)$$

where \mathcal{D} is the quadratic data-fidelity term defined in (7), and where \mathcal{R} is a generic term that penalizes non-desired solutions. The constant $\lambda > 0$ balances the influence of both terms, which determines an implicit \mathcal{K} value for (5). While \mathcal{D} is fixed according to our forward model, the choice of \mathcal{R} strongly determines the quality of the solution. Its proper specification is therefore extensively discussed below.

A. Quadratic Regularization

When applied to our framework, an extended class of quadratic functionals can be written as the Sobolev-type norm

$$\mathcal{R}_S(c) = \left\| \mathbf{L} \tilde{f}(\mathbf{x}) \right\|_{L_2}^2, \quad (10)$$

where \mathbf{L} is a linear differential operator. These regularizers penalize high responses of \mathbf{L} at each spatial location, which promotes regular solutions. Given (3) and the quadratic nature of both functionals (7) and (10), the associated minimization problems consist in discrete sets of linear equations.

When the data-fidelity term reduces to the denoising case (i.e., $M = 1$, φ_0 is the Dirac distribution $\delta(\cdot)$, and $w = 1$), these regularization functionals are linked with the standard form of the so-called *smoothing splines*. Indeed, for appropriate \mathbf{L} , the minimizer of (9) defined in the spline space (3) coincides with the optimum among all possible functions [28]. Note also that the noiseless magnification case (i.e., $\varphi_0 = \delta(\cdot)$, $w = 1$, and $\lambda \rightarrow 0$) has been specifically addressed in [7]. In a similar framework, some authors have proposed fast linear solutions to interpolate very sparse samples [10], [11]. These approaches exploit the B-spline expansion of the solution in a multigrid fashion, yielding fast iterative algorithms.

B. Nonquadratic Regularization

Edge-preserving reconstruction is achievable with non-quadratic regularizers. In this section, we review a class of such functionals described in [26], [29] and defined in our framework as

$$\mathcal{R}_N^0(c) = \int_{\mathbb{R}^2} \Psi(\|\nabla \tilde{f}(\mathbf{x})\|) d\mathbf{x}, \quad (11)$$

where $\Psi : \mathbb{R}_+ \rightarrow \mathbb{R}$ denotes a *potential function*. In addition, we discuss the associated IRLS technique that is widely used for reconstruction. As shown in the sequel, the structure of the latter is closely related to our approach developed in Section IV. The edge-preserving potential Ψ grows less fast than a quadratic function [26] unless \mathcal{R}_N^0 degenerates to \mathcal{R}_S with $L = \nabla$. For instance, TV regularization corresponds to the L_1 -norm of the image gradient, *i.e.*, to the choice $\Psi(t) = |t|$. This case has already been considered in the context of generalized sampling for image magnification [4], [5], [13], [14].

In order to yield a tractable reconstruction problem, (11) is typically discretized before minimization. This approach is standard when dealing with sampled data. For instance, a discrete form of TV based on a graph model is formulated in [30]. Following a similar idea, the gradient values entering in the original continuous-domain definition are approximated using first-order difference filters in [31]. In our context, all discrete quantities are a natural outcome of the B-spline expansion (3) after replacing the integral (11) by a sum. Here, we select a configuration where the partial derivatives of the gradient are evaluated in-between the grid nodes. This avoids the creation of spurious oscillations or divergence of the solution². Accordingly, we define

$$\mathcal{R}_N(c) = \sum_{\mathbf{k} \in \mathbb{Z}^2} \Psi(\|\overset{\circ}{\nabla} \tilde{f}(\mathbf{x})\|_{\mathbf{x}=\mathbf{k}}). \quad (12)$$

The upper-ring notation modifies the gradient—and similar vector operators—as $\overset{\circ}{\nabla} = \mathcal{S}\nabla$, where \mathcal{S} shifts a continuous-domain vector function \mathbf{v} as

$$\mathcal{S}\mathbf{v}(\mathbf{x}) = (v_1(x_1 + 1/2, x_2), v_2(x_1, x_2 + 1/2)). \quad (13)$$

The explicit form of (12) in terms of the solution coefficients is then

$$\mathcal{R}_N(c) = \sum_{\mathbf{k} \in \mathbb{Z}^2} \Psi(\|(c \star \mathbf{g})[\mathbf{k}]\|), \quad (14)$$

where \mathbf{g} is the discrete multivariate filter

$$\mathbf{g}[\mathbf{k}] = \overset{\circ}{\nabla} \varphi(\mathbf{x})|_{\mathbf{x}=\mathbf{k}}. \quad (15)$$

In order to minimize (9) with the nonquadratic term (12), we can resort to an IRLS approach. Starting from an initial solution estimate $c^{(0)}$, this method consists in the partial minimization of augmented quadratic functionals $\mathcal{J}(\cdot|c^{(n)})$ that are based on the original $\mathcal{J}(\cdot)$ and successively updated according to the current solution estimate $c^{(n)}$. Following

²Similar schemes ensure numerical stability in computational fluid dynamics [32].

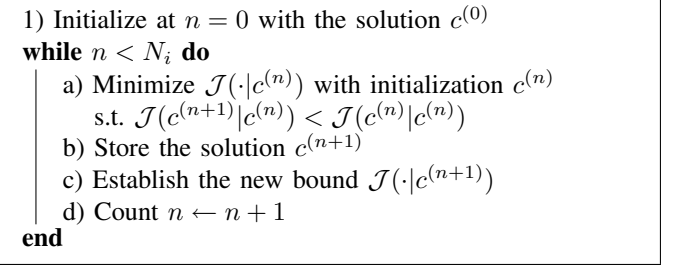


Fig. 2. Generic IRLS procedure.

the *multiplicative form* of half-quadratic minimization [26], [33], we obtain the iterative procedure given in Figure 2. The augmented functionals yielding the solution are defined as

$$\mathcal{J}_N(c|c^{(n)}) = \mathcal{D}(c) + \lambda \mathcal{R}_N(c|c^{(n)}). \quad (16)$$

This form of minimization is called multiplicative because the structure of $\mathcal{R}_N(\cdot|c^{(n)})$ involves multiplications with weights. Specifically, each augmented regularizer is defined as the quadratic functional

$$\mathcal{R}_N(c|c^{(n)}) = \frac{1}{2} \sum_{\mathbf{k} \in \mathbb{Z}^2} \theta_N(c^{(n)}, \psi)[\mathbf{k}] \|(c \star \mathbf{g})[\mathbf{k}]\|^2, \quad (17)$$

where the weight sequence θ_N at each index \mathbf{k} depends on the current solution $c^{(n)}$ as

$$\theta_N(c^{(n)}, \psi)[\mathbf{k}] = \psi(\|(c^{(n)} \star \mathbf{g})[\mathbf{k}]\|). \quad (18)$$

The scalar function $\psi : \mathbb{R}_+ \rightarrow \mathbb{R}$ is derived from the potential function Ψ of the regularizer (14) through the constraint that the successive $\mathcal{J}(\cdot|c^{(n)})$ constitute valid upper bounds of $\mathcal{J}(\cdot)$. This constraint corresponds to

$$\mathcal{J}_N(c|c^{(n)}) + \text{const.} \geq \mathcal{J}_N(c), \forall c, \quad (19)$$

with equality at $c = c^{(n)}$; the scalar constant in (19) is independent from c . In the TV case, we obtain $\psi(t) = (t^2 + \epsilon)^{-1/2}$, where the small positive parameter ϵ is added to overcome the non-differentiability of the original functional [21]. Note that the use of this constant is avoidable with nonsmooth solution initializations [31]. Minimizing an instance of (17) amounts to solving a linear problem. As specified in Figure 2, each quadratic cost need only be decreased slightly in IRLS. This implies that each of these linear problems must only be solved partially, which is doable using iterative methods.

As a prerequisite to our approach discussed in the next section, let us now draw the link between IRLS and the fixed-point interpretation discussed in [20], [21], [26]. Expanding the regularization part of (16), its minimum with respect to the coefficients c satisfies the first-order condition

$$\lambda^{-1} \mathcal{D}'(c) + \mathbf{g}^T \star \left(\theta_N(c^{(n)}, \psi)(c \star \mathbf{g}) \right) = 0. \quad (20)$$

The solution of (20) depends on the current estimate $c^{(n)}$ through the weights θ_N defined from ψ in (18). Following the terminology of [20], [26], the latter quantities are identified as *lagged diffusivities*. By extension, the IRLS algorithm of

Figure 2 can be recast as a discretized *lagged-diffusivity fixed-point iteration*, which consists in alternating between the partial resolution of (20) and the update of $c^{(n)}$. Any sequence c which satisfies

$$\lambda^{-1}\mathcal{D}'(c) + \mathbf{g}^T \star (\theta_N(c, \psi)(c \star \mathbf{g})) = 0 \quad (21)$$

is thus a fixed point of the IRLS process. Given (15), the regularization part of (20) that involves the filters \mathbf{g} and \mathbf{g}^T corresponds to the spatially discretized form of

$$\operatorname{div} \left(\psi(\|\nabla u^{(n)}\|) \nabla u \right), \quad (22)$$

where u is the continuous-domain solution with its current estimate $u^{(n)}$. Remarkably, the above expression is directly related to isotropic-diffusion flows of the form

$$\partial_t u = \operatorname{div} (\psi(\|\nabla u\|) \nabla u), \quad (23)$$

where isotropy is defined in the sense of [19]. Specifically, the divergence term (22) is similar to the right-hand-side term of (23), except that the diffusivities ψ are lagged in the former.

For general functions Ψ , it holds from [26] that the left-hand-side terms of (20) correspond to the linearized gradient of the original nonquadratic functional, up to discretization. As a consequence, (23) corresponds to the gradient flow of the original regularizer as specified by the Euler-Lagrange equation. The IRLS method used for variational image reconstruction can thus be interpreted as fixed-point iterations where the successive augmented regularization functionals are related to lagged versions of the corresponding PDE.

IV. PROPOSED APPROACH

The above discussion emphasizes the theoretical pathway that relates the regularization part of the IRLS structure to the corresponding PDE formulation. In our approach developed below, we first specify our own continuous-domain PDE, and then transpose it into an IRLS-type framework using similar concepts. Note that this section mainly deals with our regularization strategy; after derivation, our method shall involve a series of regularization functionals combined with the same data-fidelity term as in (16).

A. Edge-Enhancing Anisotropic Diffusion

The EED equation has first been applied to interpolation problems by Galić *et al.* [34], [15]. This PDE is divergence-based as in (23), and involves tensor-valued diffusivities that are determined from a smoothed gradient map $\mathcal{G}u$ of the current solution u . As compared to the regularizers presented above, the anisotropic character of EED is associated with better reconstruction properties. Equations of a similar structure have been proposed for multichannel images by Roussos and Maragos [8], [24], considering denoising and magnification applications with a generalized-sampling model. The use of tensor-driven diffusion equations has also been investigated by Tschumperlé and Deriche for general imaging problems [35], [17]. In this paper, we consider the original EED definition, which we write as

$$\partial_t u = \operatorname{div} (\mathbf{M}(\mathcal{G}u, \psi) \nabla u), \quad (24)$$

where $\mathbf{M} \in \mathbb{R}^{2 \times 2}$ denotes a symmetric and positive-definite tensor-diffusivity function. The first argument of \mathbf{M} is a smoothed gradient $\mathcal{G}u(\mathbf{x})$, while its last one is the scalar function ψ . The operator \mathcal{G} denotes a modified gradient that includes additional smoothing. As discussed in Section IV-C, our contribution is to extend the definition of \mathcal{G} so as to better preserve certain image features. The definition of \mathbf{M} distinguishes EED from the other divergence-based anisotropic PDEs considered in [35], [8], [24]. According to [15],

$$\mathbf{M}(\mathbf{v}, \psi) = \psi(\|\mathbf{v}\|) \mathcal{P}_{\mathbf{v}} + \mathcal{P}_{\mathbf{v}}^{\perp}, \quad (25)$$

where \mathcal{P} and \mathcal{P}^{\perp} are projectors onto the subscripted vector and the perpendicular directions, respectively. While anisotropic flows linked to an energy function (*e.g.*, Beltrami flow) have been studied for imaging applications³ [24], [35], [36], there is not any known energy interpretation of (24). For example, EED does not comply with the structure of [24] where the diffusion tensor involves two distinct convolutions with the same kernel. This absence of global interpretation is common in the literature [29].

In order to ensure the positive-semidefiniteness of the diffusivities and the stability of EED, we impose the function $\psi(t)$ to be nonnegative and nonincreasing in t with $\psi(0) = 1$. Given these constraints, the tensor $\mathbf{M}(\mathbf{0}, \psi)$ is well-defined and corresponds to the identity matrix. Similar to its role in (23), the purpose of ψ in (24) is to reduce smoothing across edges. In this tensor case, however, the associated flow modification is anisotropic. The action of \mathbf{M} at each position is to decompose the corresponding gradient $\nabla u(\mathbf{x})$ into the sum of two orthogonal vectors that are respectively parallel and perpendicular to $\mathcal{G}u(\mathbf{x})$. The magnitude of the parallel part of this gradient is reduced by multiplication with $\psi(\|\mathcal{G}u(\mathbf{x})\|)$, while the perpendicular one is left untouched; this permits stronger diffusion along edges, hence the EED effect. From Definition (25), the elements M_{ij} of the tensor \mathbf{M} are expressed as

$$\begin{aligned} M_{11}(\mathbf{v}, \psi) &= \|\mathbf{v}\|^{-2} (\psi(\|\mathbf{v}\|) v_1^2 + v_2^2), \\ M_{22}(\mathbf{v}, \psi) &= \|\mathbf{v}\|^{-2} (v_1^2 + \psi(\|\mathbf{v}\|) v_2^2), \\ M_{12}(\mathbf{v}, \psi) &= \|\mathbf{v}\|^{-2} (\psi(\|\mathbf{v}\|) - 1) v_1 v_2, \end{aligned} \quad (26)$$

where v_1 and v_2 are the coordinates of the vector \mathbf{v} .

B. Specification of the Penalty Function

In our approach, we consider three diffusivity functions ψ among those proposed in the literature. The first is the *Charbonnier diffusivity* used in [15] that is defined as

$$\psi_0(t) = (1 + t^2/\alpha^2)^{-1/2}, \quad (27)$$

³These flows can remain anisotropic (*e.g.*, tensor TV in [24]) or degenerate [35], [36] when applied to single-channel data.

where $\alpha \in \mathbb{R}_+^*$ is a constant. The second one is linked to the Huber potential through the multiplicative form of half-quadratic minimization [33], and is defined as

$$\psi_1(t) = \begin{cases} 1, & |t| \leq \alpha, \\ \alpha|t|^{-1}, & \text{otherwise.} \end{cases} \quad (28)$$

Since $\psi_1(t) \propto |t|^{-1}$ for large t , it tampers cross-edge smoothing the same way as TV. Meanwhile, this function can restore smoothly varying regions, because the associated potential is quadratic for $|t| < \alpha$. The functions ψ_0 and ψ_1 are closely related because their values coincide when t tends to zero or to infinity. The third alternative that we propose is the Perona-Malik diffusivity

$$\psi_2(t) = \exp(-t^2/\beta^2), \quad (29)$$

where $\beta \in \mathbb{R}_+^*$ is a constant. This function is well-known for its contrast-enhancing properties in the case of isotropic [37] as well as of anisotropic diffusion [19].

C. Specification of the Modified Gradient Operator

As discussed above, the map $\mathcal{G}u$ is used to specify anisotropic diffusivities. The operator \mathcal{G} serves in the EED equation (24) as an edge-information estimate that is more robust than the standard gradient. In [15], this operator corresponds to the Gaussian-smoothed gradient

$$\mathcal{G}_0 = \nabla * \phi_\sigma, \quad (30)$$

where ϕ_σ is an isotropic Gaussian filter of standard deviation σ . The associated results reported in [15] are very promising, even when dealing with very sparse data. Note that, although $\mathcal{G}u$ only enters in the definition of \mathbf{M} , it nonetheless determines the character of the edge-enhancing effect in terms of flow regulation. As shown in Section VII, the potential differences in terms of reconstruction behavior are important, which suggests introducing some better estimates.

Anisotropic flows are obtainable without the requirement of gradient smoothing in general [24]. This operation is nevertheless necessary in our PDE to guarantee the anisotropy of the diffusivities⁴. The smoothing process generally tends to wipe out fine-scale edge information. Our contribution is to specify an operator \mathcal{G}_1 that yields a *directionally smoothed* version of the image gradient, which better preserves fine-scale information compared to the Gaussian solution \mathcal{G}_0 . In order to determine the corresponding smoothing directions, we compute orientation estimates $\tau \in [0, \pi[$ that are matched with the local edge features of the image argument u at each position. These estimates are obtained as the solution of the optimization problem described below.

Let us consider the class of segment cross-sections of constant size Σ , centered at positions \mathbf{x} , and with orientations $\tau_0 \in [0, \pi[$. Given u , we associate this class to the local oriented-mean measure

⁴Replacing \mathcal{G} by ∇ in (24) would cause EED to degenerate to the isotropic flow (23).

$$\mu^*(u, \mathbf{x}, \tau_0) = \Sigma^{-1} \int_{-\Sigma/2}^{\Sigma/2} u(\gamma(t)) dt, \quad (31)$$

where

$$\gamma(t) = \mathbf{x} + t(\cos(\tau_0), \sin(\tau_0)). \quad (32)$$

The corresponding variance measure is given as

$$\text{Var}^*(u, \mathbf{x}, \tau_0) = \Sigma^{-1} \int_{-\Sigma/2}^{\Sigma/2} (u(\gamma(t)) - \mu^*(u, \gamma(t), \tau_0))^2 dt. \quad (33)$$

Given (33), we choose our estimates τ to match the minimum-variance orientations of the image. In that context, Σ can be interpreted as a scale parameter that is approximately determined from the characteristic oriented-feature size. At each position, the solution thus corresponds to a local edge-orientation estimate that is expressed as

$$\tau(u, \mathbf{x}) = \arg \min_{\tau_0} \text{Var}^*(u, \mathbf{x}, \tau_0), \quad (34)$$

which satisfies translation and rotation invariance with respect to u .

At given scale Σ , our variance measure quantifies the image fluctuations along each orientation τ_0 . In that respect, the minimum argument of (34) is conceptually similar to the coherence direction defined in [19] that is based on structure tensors. The advantage of variance-based criteria is their ability to estimate the orientation even in the vicinity of a contrast change [38]. As shown in Figure 3, the map τ provides accurate data on the local feature orientations of the image. Adaptive filtering of the gradient map along those directions can thus reduce the loss of information associated with cross-edge smoothing. Because no closed-form solution of (34) exists to compute the weights, we propose to optimize τ among a discrete set of N_o orientations. Note that similar discretization approaches have been considered using candidate stencils for the evaluation of image variations along oriented paths [9].

Prior to filtering, we make the gradient map of u consistent with the estimated orientation map τ pointwise. Accordingly, we only keep the component of each gradient vector $\nabla u(\mathbf{x})$ which is perpendicular to the orientation $\tau(u, \mathbf{x})$. This projection operation yields the corrected gradients

$$\nabla^c u(\mathbf{x}) = ((\nabla u)(\mathbf{x}))^T \mathbf{e}^\perp(u, \mathbf{x}) \mathbf{e}^\perp(u, \mathbf{x}), \quad (35)$$

where

$$\mathbf{e}^\perp(u, \mathbf{x}) = (-\sin(\tau(u, \mathbf{x})), \cos(\tau(u, \mathbf{x}))). \quad (36)$$

We finally smooth these corrected gradients along the corresponding τ , using directional averaging filters with the same Σ as in the mean and variance measures of (31) and (33). The operator \mathcal{G}_1 thus acts as

$$\mathcal{G}_1 u(\mathbf{x}) = \mu^*(\nabla^c u, \mathbf{x}, \tau(u, \mathbf{x})). \quad (37)$$

The invariances of (31) and (34) and the pointwise character of (35) imply that \mathcal{G}_1 is intrinsically translation and rotation




Fig. 3. Illustration of our orientation-estimation method. On the left, *Fingerprint*. On the right, minimum-variance directions τ obtained according to (34) with $\Sigma = 25$. The directions $[0, \pi[$ are mapped to the grayscale range [black, white[. This map contains two large zones that are associated with the two main feature directions of the original image. Note that abrupt black-white transitions appear because orientations are only defined modulo π .

invariant as in \mathcal{G}_0 . These important characteristics guarantee that \mathcal{G} is not biased towards particular orientations or positions in the image.

D. AIRLS Algorithm

As previously discussed, the quadratic functionals of IRLS are linked to the lagged-diffusivity forms of the gradient flow. Similarly, we propose an IRLS procedure that is based on lagged versions of the EED flow, following the concepts introduced in Section III. Despite the fact that there is no underlying maximization principle in our case, the successive functionals that we define constitute linear approximations of (24). Given one current estimate $u^{(n)}$, the lagged EED equation is

$$\partial_t u = \operatorname{div} \left(\mathbf{M}(\mathcal{G}u^{(n)}, \psi) \nabla u \right). \quad (38)$$

According to the Euler-Lagrange equation, (38) is the gradient flow which originates from the functional

$$\mathcal{R}_A^0(u|u^{(n)}) = \int_{\mathbb{R}^2} \left\langle \mathbf{M}(\mathcal{G}u^{(n)}(\mathbf{x}), \psi) \nabla u(\mathbf{x}), \nabla u(\mathbf{x}) \right\rangle d\mathbf{x}, \quad (39)$$

up to a factor of 2 that we drop for convenience. Note that the regularizer proposed in [39] is of similar form; it includes structure tensors that are based on fixed estimates, but it is nonquadratic unlike (39).

Expression (39) can be rewritten in a more intelligible form. Indeed, defining \mathbf{D} and \mathbf{D}^\perp as directional derivatives along the subscript vector argument and the direction perpendicular to it, respectively,

$$\begin{aligned} \langle \mathcal{P}_\mathbf{v} \nabla u(\mathbf{x}), \nabla u(\mathbf{x}) \rangle &= (\mathbf{D}_\mathbf{v} u)^2(\mathbf{x}), \\ \langle \mathcal{P}_\mathbf{v}^\perp \nabla u(\mathbf{x}), \nabla u(\mathbf{x}) \rangle &= (\mathbf{D}_\mathbf{v}^\perp u)^2(\mathbf{x}), \end{aligned} \quad (40)$$

which implies from (25) that $\mathcal{R}_A^0(\cdot|u^{(n)})$ expands as

$$\begin{aligned} \mathcal{R}_A^0(u|u^{(n)}) &= \int_{\mathbb{R}^2} \psi(\|\mathcal{G}u^{(n)}(\mathbf{x})\|) (\mathbf{D}_{\mathcal{G}u^{(n)}(\mathbf{x})} u)^2(\mathbf{x}) d\mathbf{x} \\ &\quad + \int_{\mathbb{R}^2} (\mathbf{D}_{\mathcal{G}u^{(n)}(\mathbf{x})}^\perp u)^2(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (41)$$

Equation (41) gives further insight on how each quadratic regularizer—which is linked to the linearized form of the EED flow—penalizes u . In the first integral term, the directional derivatives of u that are parallel to $\mathcal{G}u^{(n)}$, i.e., perpendicular to the edge features, are weakly penalized given their multiplication with ψ . Edges are therefore well-preserved as in nonquadratic regularization. Simultaneously, regularity along those same edge estimates is strongly enforced in the second integral term, which is akin to an L_2 -norm. These two simultaneous constraints favor curvature regularity of the solution.

In order to obtain an IRLS procedure compatible with the algorithmic framework of Figure 2, we discretize the cost (39) and combine it with our data term. Indeed, as for standard regularization techniques, computationally tractable approaches require a discretized version of the continuous quadratic costs $\mathcal{R}_A^0(\cdot|u^{(n)})$. Using the same discretization as in (12), the expression

$$\mathcal{R}_A(c|c^{(n)}) = \sum_{\mathbf{k} \in \mathbb{Z}^2} \left\langle \mathbf{M}(\hat{\mathcal{G}}\tilde{f}^{(n)}(\mathbf{x}), \psi) \hat{\nabla} \tilde{f}(\mathbf{x}), \hat{\nabla} \tilde{f}(\mathbf{x}) \right\rangle_{\mathbf{x}=\mathbf{k}} \quad (42)$$

is obtained for our solution coefficients. Along with (15), this definition allows to rewrite (42) as

$$\mathcal{R}_A(c|c^{(n)}) = \sum_{\mathbf{k} \in \mathbb{Z}^2} (c \star \mathbf{g})[\mathbf{k}]^T \boldsymbol{\theta}(c^{(n)}, \psi)[\mathbf{k}] (c \star \mathbf{g})[\mathbf{k}], \quad (43)$$

where the tensor weights $\boldsymbol{\theta}$ are determined as

$$\boldsymbol{\theta}(c^{(n)}, \psi)[\mathbf{k}] = \mathbf{M}(\hat{\mathcal{G}}\tilde{f}^{(n)}(\mathbf{x}), \psi)|_{\mathbf{x}=\mathbf{k}}. \quad (44)$$

Equation (44) evaluates $\hat{\mathcal{G}}\tilde{f}^{(n)}$ over a sequence of points $\mathbf{x} = \mathbf{k}$. The corresponding means and variances used for \mathcal{G}_1 in (34) and (37) can be determined with arbitrary precision, using the continuous-line integrals (31) and (33) for $\tilde{f}^{(n)}$ in the same way as for u in Section IV-C. This holds because $\tilde{f}^{(n)}$ is continuously defined from $c^{(n)}$ given our spline model (3). For computational reasons, we approximate all integrals by finite sums depending on uniformly spaced samples. This discretization also allows to compute (34) recursively: each of the N_o oriented-variance maps linked to $\hat{\mathcal{G}}\tilde{f}^{(n)}$ can be estimated using parallel sliding windows of size Σ and orientation τ_0 . The samples of $\tilde{f}^{(n)}$ and the local variance assigned to each window are then updated recursively⁵. Equations (7), (43), and (44) define our total quadratic functionals as

$$\mathcal{J}_A(c|c^{(n)}) = \mathcal{D}(c) + \lambda \mathcal{R}_A(c|c^{(n)}). \quad (45)$$

⁵The computational performance is optimal when these updates are performed along rows or columns of values. We maintain this condition for $\tau_0 \notin \{0, \pi/2\}$ by applying approximate pre- and post-shearing transformations on the sample lattices.

According to the generic optimization framework of Figure 2, our algorithm is first initialized with the solution $c^{(0)}$. We define this sequence as the masked image samples f_2 upsampled by M using zero padding and smoothed by the filter $\frac{1}{4}[1 \ 2 \ 1]$ along each dimension. Subsequently, we partially minimize N_i successive costs (45) to obtain a solution. Each of these minimization steps corresponds to a fixed-point iteration where the next estimate $c^{(n+1)}$ is found from $c^{(n)}$ using fixed tensor diffusivities θ . In that regard, our approach is similar to the anisotropic technique proposed in [39] where a series of convex problems is solved to denoise images. In our case, however, the successive functionals to minimize are quadratic and thus easier to tackle. Moreover, since (43) is an ℓ_2 -norm whose weights θ depend on $c^{(n)}$, our reconstruction method is of IRLS type. We call it AIRLS according to the anisotropic nature of the tensor weights in (44). Each cost $\mathcal{J}_A(\cdot|c^{(n)})$ has a unique minimum that satisfies the first-order condition

$$\lambda^{-1}\mathcal{D}'(c) + \mathbf{g}^T \star (\theta(c^{(n)}, \psi)(c \star \mathbf{g})) = 0. \quad (46)$$

Ultimately, any fixed point of our global iterative process satisfies

$$\lambda^{-1}\mathcal{D}'(c) + \mathbf{g}^T \star (\theta(c, \psi)(c \star \mathbf{g})) = 0. \quad (47)$$

In analogy with the case of Section III, Condition (47) is the steady state of a discretized PDE whose regularization part corresponds to the EED flow (24). This shows that, in terms of asymptotic solutions, our approach is similar to conventional PDE-based methods. It is, however, more attractive computationally because the computation of the smoothed-gradient map as well as the nonlinear operations involved in the diffusivity estimations (44) are restrained to the reweightings. As confirmed in our experiments, satisfactory results are obtained with a small amount of reweightings N_i , precisely as in standard IRLS.

The AIRLS algorithm is specified by the parameters $\{\mathcal{G}, \psi\}$. The first argument \mathcal{G} corresponds either to the Gaussian-smoothed gradient \mathcal{G}_0 defined in (30), or to our modified operator \mathcal{G}_1 defined in (37). Similarly, the function ψ can be chosen as the Charbonnier diffusivity ψ_0 defined in (27), the Huber diffusivity ψ_1 defined in (28), or the Perona-Malik diffusivity ψ_2 defined in (29). Our specific EED settings denoted by EED₁ and EED₂ correspond to $\{\mathcal{G}_1, \psi_1\}$ and $\{\mathcal{G}_1, \psi_2\}$, respectively. When using $\{\mathcal{G}_0, \psi_0\}$, our algorithmic framework reproduces the PDE-based method of [15]; the related technique is then called EED₀.

V. LINEAR PROBLEMS

The global minimum of each weighted quadratic cost (45) corresponds to the solution of a linear system. In this section, we derive the explicit form of those systems using matrix notation, where lowercase bold-symbol vectors relate to the corresponding non-bold sequences through lexicographic ordering, and where each matrix, as a vector multiplier, implements some linear operation acting on the related sequence. Accordingly, the generic form of each matrix system is

$$\mathbf{A}^{(n)}\mathbf{c}^{(\min)} = \mathbf{d}, \quad (48)$$

where $\mathbf{c}^{(\min)}$ contains the lexicographically ordered coefficients of the minimizer, where $\mathbf{A}^{(n)}$ is the system matrix that depends on our estimate $\mathbf{c}^{(n)}$, and where \mathbf{d} is a vector whose components are constant.

The first step towards specifying $\mathbf{A}^{(n)}$ and \mathbf{d} is to reformulate (45) using matrix notation. Accordingly,

$$\mathcal{D}(c) = (\mathbf{f}_2 - \mathbf{D}_M\mathbf{B}\mathbf{c})^T\mathbf{W}(\mathbf{f}_2 - \mathbf{D}_M\mathbf{B}\mathbf{c}), \quad (49)$$

where each vector is specified as discussed above. According to the value of M used in our generalized sampling model (1), the matrices \mathbf{D}_M and its transpose \mathbf{U}_M implement two-dimensional M -fold downsampling and upsampling, respectively, while \mathbf{B} is a convolution matrix associated with the filter b in (8). Finally, the diagonal matrix \mathbf{W} is associated with point-wise multiplication with the weights w used for masking in (2). Similarly, we can write the regularization term (43) in the compact form

$$\mathcal{R}_A(c|\mathbf{c}^{(n)}) = (\mathbf{R}\mathbf{c})^T\Theta(c^{(n)}, \psi)\mathbf{R}\mathbf{c}, \quad (50)$$

where \mathbf{R} and Θ concatenate convolution and diagonal matrices, respectively. The rectangular matrix $\mathbf{R} = (\mathbf{R}_1, \mathbf{R}_2)$ implements the gradient. Specifically, each convolution matrix \mathbf{R}_i relates to the derivative component g_i of the multivariate filter \mathbf{g} defined in (15). The square matrix Θ is updated according to the current estimate $\mathbf{c}^{(n)}$; it decomposes as $(\Theta_{11} \ \Theta_{12}, \Theta_{12} \ \Theta_{22})$, where each diagonal matrix Θ_{ij} is associated with point-wise multiplication with the corresponding scalar sequence θ_{ij} related to the tensor weights θ of (44). Note that Θ structurally extends its counterpart in IRLS, as it concatenates distinct and off-diagonal sub-matrices Θ_{ij} . Based on (49) and (50), we write the total cost as

$$\mathcal{J}_A(c|\mathbf{c}^{(n)}) = \mathcal{D}(c) + \lambda\mathcal{R}_A(c|\mathbf{c}^{(n)}). \quad (51)$$

Since (51) is quadratic, its gradient with respect to \mathbf{c} vanishes at $\mathbf{c}^{(\min)}$. Enforcing this condition, and using matrix differentiation, we obtain

$$\mathbf{A}^{(n)} = \mathbf{B}^T\tilde{\mathbf{W}}_1\mathbf{B} + \mathbf{R}^T\tilde{\mathbf{W}}_2\mathbf{R}, \quad (52)$$

where $\tilde{\mathbf{W}}_1 = \mathbf{U}_M\mathbf{W}\mathbf{D}_M$, and where $\tilde{\mathbf{W}}_2 = \lambda\Theta(c^{(n)}, \psi)$. Finally, the vector \mathbf{d} in (48) corresponds to

$$\mathbf{d} = \mathbf{B}^T\mathbf{U}_M\mathbf{f}_2. \quad (53)$$

The matrix system (52) is extremely large due to the considerable number of unknowns, which implies that Problem (48) cannot be solved exactly. However, as stated in Section III-B, the corresponding quadratic functional (45) needs only be partially minimized with respect to the current solution estimate to yield $\mathbf{c}^{(n+1)}$. In Section VI, we propose to partially solve (48) iteratively, initializing the next solution $\mathbf{c}^{(n+1)}$ to the current estimate $\mathbf{c}^{(n)}$.

VI. ITERATIVE SOLUTION

As mentioned above, our variational approach allows to restrain the diffusivity estimations associated with the regularization problem in the reweightings. This allows to focus on the resolution of the subproblems derived in Section V. In the sequel, we devise a fast iterative method to partially solve these linear systems, considering their sparse structure. In particular, the diagonal matrices entering the definition of (52) are well suited for multigrid solvers.

A. Multigrid Approach

The multigrid strategy consists in solving problems by iterating not only at their nominal scale, but also at coarser ones, adapting their discretization accordingly. Multigrid iterative methods are beneficial for certain types of problems where the additional lower-resolution iterates are of negligible computational cost compared to their overall contribution in terms of convergence rates [40]. In particular, iterating at successively downscaled Cartesian grids is efficient for linear image reconstruction from spatially sparse samples [10], [11]. In this context, according to the original formulation (48), we define our linear problems at H distinct grids, using the notation

$$\mathbf{A}^h \mathbf{c}^h = \mathbf{d}^h, \quad (54)$$

where the superscripts $h \in \{0, \dots, H-1\}$ relate quantities to a specific grid Ω^h . Each grid is constructed with a regular step 2^h in each dimension, which means that the number of elements of \mathbf{c}^h scales as 4^{-h} .

From (54), we use Full-Multigrid V-cycles [40] as an iterative scheme to find the solution $\mathbf{c}^{(n+1)}$. This method is standard in the literature and involves transfer operations as well as relaxations with an iterator \mathcal{I} at each grid, as described in Appendix A with the relevant definitions. In order to maximize the performance of our approach, we use the obtained \mathbf{c}^h to initialize the next linear problem at all grids. The problem at grid Ω^0 is (48), which implies that $\mathbf{A}^0 = \mathbf{A}^{(n)}$, $\mathbf{d}^0 = \mathbf{d}$, and $\mathbf{c}^0 = \mathbf{c}^{(n+1)}$. At coarser grids, \mathbf{A}^h are scaled versions of \mathbf{A}^0 , while \mathbf{d}^h are the *residuals* produced from the iterative process itself.

Since the solution is expressed as coefficients in a B-spline basis, the corresponding *prolongation* and *restriction* operators \mathbf{I}^\uparrow and \mathbf{I}^\downarrow exploit the two-scale relations [27]. Specifically, they correspond to the B-spline scaling filter h_2 of degree η following upsampling by 2, and to the B-spline scaling filter h_2^T followed by downsampling by 2, respectively. In matrix form, we write

$$\begin{aligned} \mathbf{I}^\uparrow &= \mathbf{H}_2 \mathbf{U}_2, \\ \mathbf{I}^\downarrow &= \mathbf{D}_2 \mathbf{H}_2^T. \end{aligned} \quad (55)$$

We now have to specify \mathbf{A} at each grid. In order not to complexify the problem formulation, we impose similar matrix structures at all grids, decomposing \mathbf{A}^h as

$$\mathbf{A}^h = \mathbf{B}^{hT} \tilde{\mathbf{W}}_1^h \mathbf{B}^h + \mathbf{R}^{hT} \tilde{\mathbf{W}}_2^h \mathbf{R}^h, \quad (56)$$

where the separate terms at Ω^0 correspond to the ones of (52).

In order to specify the data part of \mathbf{A}^h , we build a weight pyramid, starting from the available fine-scale matrix $\tilde{\mathbf{W}}_1^0$. Simplifying the coarser-scale convolution matrices \mathbf{B}^h as identity, we express the diagonal elements of $\tilde{\mathbf{W}}_1^1$ as

$$\tilde{w}_1^1[\mathbf{k}] = \{\tilde{w}_1^0 \star h_2^\vee \star b^\vee\}_{\downarrow 2}[\mathbf{k}], \quad (57)$$

where \vee flips a given sequence as $\cdot^\vee[\mathbf{k}] = \cdot[-\mathbf{k}]$. The convolutive effect of \mathbf{B} is thus taken into account at this first weight level. For $h > 1$, the expression of \tilde{w}_1^h takes the simpler form

$$\tilde{w}_1^{h+1}[\mathbf{k}] = \{\tilde{w}_1^h \star h_2^\vee\}_{\downarrow 2}[\mathbf{k}]. \quad (58)$$

Regarding the regularization term, the components of the coarse-scale diagonal sub-matrices $\tilde{\mathbf{W}}_{2,ij}^h$ involved in (56) are obtained as in (58) through the relations

$$\tilde{w}_{2,ij}^{h+1}[\mathbf{k}] = \{\tilde{w}_{2,ij}^h \star h_2^\vee\}_{\downarrow 2}[\mathbf{k}], \quad (59)$$

while \mathbf{R}^h is defined as

$$\mathbf{R}^h = 2^{-h} \mathbf{R}, \quad (60)$$

according to the scaling properties of the gradient operator.

B. Successive Over-Relaxation

The phase and grid parameters given to the iterator \mathcal{I} in the Full-Multigrid V-cycles allow to specify distinct numbers of iterations $\nu_\downarrow(h)$ and $\nu_\uparrow(h)$, as well as distinct relaxation constants $\omega_\downarrow(h)$ and $\omega_\uparrow(h)$. For convenience, we denote the relaxation constants by ω when referring to them in a generic sense.

Given the symmetry and positive-definiteness of the \mathbf{A}^h in (54), a certain class of iterative methods can be used to specify \mathcal{I} , including the well-known Conjugate Gradient (CG). The successive over-relaxation (SOR) technique [41] is especially efficient for our multigrid problem. It corresponds to a damped version of the Gauss-Seidel iterative method, and its convergence is guaranteed for $\omega \in [0, 2]$. Given \mathbf{A}^h and \mathbf{d}^h , the SOR iterate at grid Ω^h is defined as

$$\mathbf{c}^h \leftarrow \mathbf{c}^h + \omega(\mathbf{D}^h + \omega \mathbf{L}^h)^{-1} \mathbf{R}(\Omega^h), \quad (61)$$

where \mathbf{D}^h and \mathbf{L}^h stand for the diagonal and strictly lower triangular parts of the matrix \mathbf{A}^h , respectively.

Unlike exact resolution, this iterative approach only involves partial matrix inversions. Structurally, each iteration is performed by updating the coefficient vector \mathbf{c}^h componentwise. The sparse structure of \mathbf{A}^h makes these updates correspond to space-domain operations of complexity $\mathcal{O}(N \log N)$ at most.

VII. EXPERIMENTS

In this section, we compare our AIRLS approach with respect to the state of the art, considering interpolation experiments on grayscale images⁶. Our implementation has been coded in Java, and run on Max OS X with a Quad-Core Xeon 2×2.8 GHz and 4 GB of DDR2 memory. The computation of the optimal orientations in (34) for \mathcal{G}_1 and the SOR updates in (61) are parallelized using multithreading. The state-of-the-art methods that are considered for comparison are also based on parallel implementations⁷.

We consider B-splines of order $\eta = 2$, and compute $N_o = 16$ orientations when using the operator \mathcal{G}_1 in (24). Following the generic IRLS procedure of Figure 2, our algorithm reconstructs images by solving $N_i = 10$ successive linear problems that are defined on $H = 4$ grids in (54). For each linear problem, the iteration and relaxation constants common to all experiments are $\omega_\uparrow(h) = 1.5, \nu_1(h) = 2, \nu_\uparrow(h) = 1, \forall h$, and $\omega_1(h) = 1.5, \forall h > 0$. The constants α and β involved in the diffusivities (28) and (29) are scaled according to the dynamic range I_d of the image under consideration. For convenience, we display unit-resampled versions of the continuous reconstructions [6]. The PSNR measures are evaluated on the central portion of the images (80%) so as to discard the influence of the boundary conditions.

A. Sparse Interpolation of Ideal Samples

In these experiments, we interpolate grayscale images from 2% of their samples, according to one realization of a binary random mask. The sampling is ideal, meaning that the prefilter φ_0 used in our generalized model (1) reduces to the Dirac distribution $\delta(\cdot)$. The other parameters are $\lambda = 0.01, M = 1, \Sigma = 25, \sigma = 4, \nu_0 = 1, \omega_\downarrow(0) = 1.95, \alpha = 2 \cdot 10^{-3} I_d$, and $\beta = 8 \cdot 10^{-2} I_d$. We consider sets of 256×256 images, the corresponding binary mask being shown in Figure 4.

We compare our three EED_i methods with the fast PDE-based approach of Tschumperlé [17] implemented in version 2.9 of *GREYcStorage*⁸ (GREY). We have also implemented a PDE-based version of the EED flow using explicit time steps; this approach is referred to as EED in the sequel.

Quantitative and visual results for these algorithms are provided for the interpolation of several images in Table I and in Figure 5, respectively. Observe that our specific tensor-estimation approach in EED_1 restores directional features better than GREY, EED, and EED_0 , and yields higher SNR values. Our EED_2 method is quantitatively inferior to EED_1 but restores very sharp edges as can be seen in Figure 5. Its regularization behavior is consistent with the properties of the Perona-Malik diffusivity mentioned in Section IV-B. Our

⁶The *CT* image is part of the Dicom stack CT HEAD-NK 5.0 B30s (Keith E. Blackwell, M.D.). The standard *Bird* test image is found at <http://www2.isye.gatech.edu/~brani/images/bird.gif>, while the rest of the original data belongs to the GCF-BM3D set found at <http://www.cs.tut.fi/~foi/GCF-BM3D>. The image histograms are rescaled to $[0, 255]$ for all experiments.

⁷The implementations that are distinct from our method sometimes run on different platforms, and can differ in their level of optimization and parallelization. Caution should therefore be exerted to not overinterpret the runtime results reported in Table I.

⁸This code is run with 15 iterations under default settings. The resulting PSNR tends to degrade when iterating further.




Fig. 4. Masking process. The pseudo-random binary mask shown above is applied before interpolation and only keeps 2% of the original data; the average gap between the retained samples corresponds to 4.3 pixels.

PDE-based implementation of EED requires 5000 time steps for convergence. It is much slower than EED_0 , as shown in Table I, but yields similar results in terms of PSNR and visual appearance. This corroborates the fixed-point interpretations discussed in Sections III and IV.

These results demonstrate the suitability of our EED_i methods to restore geometrical information from few image samples. In general, our approach is less successful at restoring textures because, in several cases, the latter are composed of repetitive patches rather than well-defined oriented features.

B. Sparse Interpolation of Generalized Samples

In this second part, we wish to reconstruct the same images from 2% of their samples, considering here generalized sampling $\varphi_0 = a^{-2} \text{rect}(a^{-1} \cdot)$, with $a = 7$. The EED approach of [15] is not applicable here; it can be substituted with our EED_0 method as implemented in our more general AIRLS framework.

Using the same mask and parameters as above, the results are provided in Table I for our three EED_i methods and shown in Figure 6 for EED_0 and EED_1 . The improvement of EED_1 relative to the other methods is comparable with the ideal case in terms of PSNR and visual quality. Remarkably, all results are better in terms of PSNR than their counterparts in the ideal-sampling setting. As a matter of fact, the analysis function φ_0 acts as an anti-aliasing filter before sampling, which causes the overall reconstructed features to be more consistent with the original image. This emphasizes the interest of using generalized sampling for sparse interpolation.

C. Image Magnification

In our framework, image magnification corresponds to a particular instance of sparse interpolation where the sparsity is regular and typically low. Our approach is well suited to that problem because the presence of a prefilter before sampling is inherent in practical acquisition devices [3]. For that case, we compare our EED_i methods with quadratic regularization as well as with distinct magnification algorithms that also handle generalized sampling. The quadratic approach regularizes the L_2 -norm of the image gradient. It is implemented in AIRLS as the limit case $\psi = 1$ where one single unweighted linear problem has to be solved. We further consider iterative TV




Fig. 5. Ideal interpolation of *CT* and *Lena* (256×256 crop) from 2% of samples.




Fig. 6. Generalized interpolation of *CT* and *Lena* (256×256 crop) from 2% of samples acquired with a prefilter φ_0 .

reconstruction with an implementation⁹ of the primal-dual method [14] as well as the PDE-based method of Roussos and Maragos (R&M) provided by Getreuer as an online demo [42]. Choosing a magnification factor of $M = 4$, we model the sensor integration φ_0 as a 2D Gaussian of standard deviation $0.35M$ in each dimension as can be specified in the R&M demo. The parameters specific to AIRLS are $\lambda = 0.01$, $\Sigma = 9$, $\nu_0 = 2$, $\omega_1(0) = 1.5$, $\alpha = 2 \cdot 10^{-3}I_d$, and $\beta = 2 \cdot 10^{-2}I_d$. The other algorithms are used with their default settings.

The available data consists in images that are primarily downsampled according to the above generalized-sampling

⁹This Matlab implementation is based on a publicly available source code of Y. Chen and T. Pock, Graz University of Technology, Austria. We have modified the original version so as to handle the φ_0 under consideration. While the code structure is not optimized for multithreading, most low-level Matlab functions that are involved are intrinsically parallelized (e.g., elementary operators). The algorithm is run until 400 iterations are reached, or until the PSNR increase per iteration is lower than 10^{-3} .

settings. The corresponding reconstructions are obtained by magnification and compared with the known oracles; the results are shown in Table I and Figure 7. In Figure 8, we perform a similar magnification experiment on an image that is provided as such without any prior downscaling.

We observe that the combination of consistent-resampling constraints with edge-preserving regularization yields high-quality reconstructions in terms of feature preservation, as discussed in [5]. This emphasizes the interest of taking the sensor integration into account for image magnification. The results of Figures 7 and 8 demonstrate that all nonlinear methods restore sharper edges than quadratic regularization.

As compared to the isotropic TV solution, the anisotropic EED_0 , EED_1 , and R&M methods better preserve certain fine structures as well as the curvature of the objects, thanks to the associated flows. They also avoid staircasing artifacts but introduce some edge smearing. The EED_2 method yields the highest-quality results because it benefits from the desirable properties of anisotropic diffusion while preserving image sharpness nearly at the same level as TV. It also yields the highest PSNR values for several images as shown in Table I.

VIII. CONCLUSIONS

We have designed a method that reconstructs continuous images from a sparse set of generalized samples. Combined with consistent data-fidelity constraints, our anisotropic regularization approach was designed to preserve the edge information accurately, and to be functional at high sparsity levels. In the experiments, promising results have been obtained with nonuniform interpolation of images as well as with consistent image magnification.

From an algorithmic perspective, the low computational cost of our method demonstrates that approaches based on iteratively reweighted least squares (IRLS) can successfully handle anisotropic regularization in sparse inverse problems. This low cost legitimates the use of linear-multigrid approaches. As in IRLS, we have successfully maximized the algorithmic performance by restricting the diffusivity estimations to the

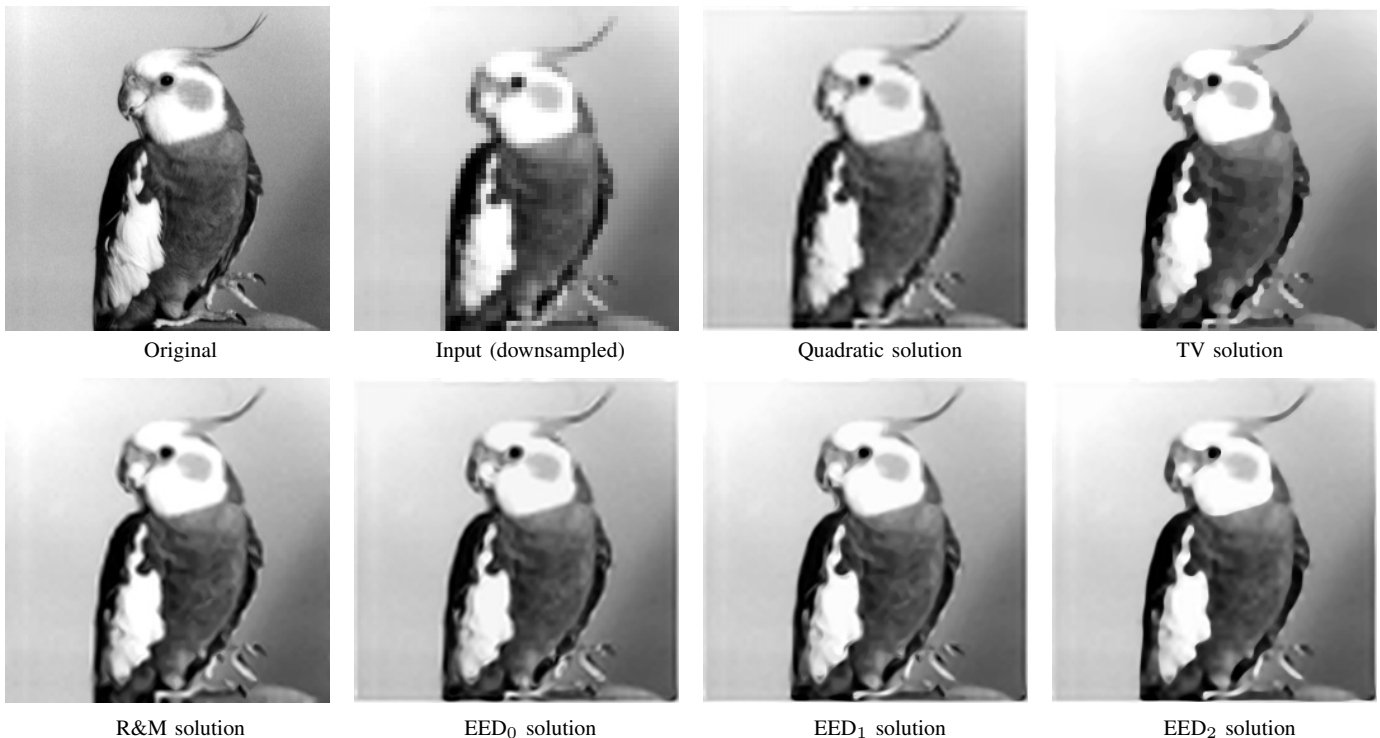


Fig. 7. Magnification of *Bird* after it was downsampled by a factor 4 along each dimension. The input downsampled image is shown at the same scale using pixel duplication.

TABLE I
NUMERICAL RESULTS (PSNR VALUES AND AVERAGE RUNTIMES) OBTAINED IN OUR EXPERIMENTS.

Experiment	Ideal Interpolation					Generalized Interpolation			Magnification					
	GREY	EED	EED ₀	EED ₁	EED ₂	EED ₀	EED ₁	EED ₂	Quad.	TV	R&M	EED ₀	EED ₁	EED ₂
Runtime [s]	6.09*	134.90	3.62	5.73	5.00	4.22	6.52	6.60	1.31	5.32*	1.88*	5.46	7.45	7.59
<i>Bird</i>	19.86	20.45	20.77	21.02	20.42	23.28	22.98	22.49	27.22	28.16	28.31	27.99	28.43	28.68
<i>Cameraman</i>	17.37	17.68	17.58	17.90	17.40	19.58	19.73	19.50	21.73	22.37	22.16	21.99	22.18	22.42
<i>CT</i>	20.66	21.21	21.14	21.88	21.28	22.76	23.68	23.59	27.11	27.67	28.75	27.53	28.80	28.67
<i>House</i>	18.47	18.76	18.71	19.73	19.60	21.07	21.62	21.31	25.74	26.46	26.60	25.98	26.55	26.54
<i>Lena (crop)</i>	18.38	19.39	19.34	20.22	19.61	21.12	21.87	21.70	24.81	25.19	25.70	25.54	25.85	25.70
<i>Montage</i>	17.66	18.57	18.52	18.67	18.15	20.11	20.38	20.17	21.85	22.20	22.04	22.03	22.28	22.30
<i>Peppers</i>	19.00	18.97	19.12	19.15	18.49	20.77	21.21	20.80	25.36	25.80	26.21	25.85	26.38	26.54

* These runtimes are obtained from implementations on distinct platforms.

reweightings. This has led to a simple and efficient design that drastically reduces redundancy in terms of operations. Our algorithm has been optimized for sparse interpolation. It is comparable to state-of-the-art implementations in terms of computational efficiency.

Because it is based on IRLS, our variational approach is modular. Its data term can potentially be redefined so as to solve several types of inverse problems, such as multi-image super-resolution or deconvolution. The structure of the proposed regularization framework is also able to handle alternate definitions of the anisotropic flow. Extensions to higher dimensions are straightforward because the edge-enhancing anisotropic diffusion that we formulated in Section IV is generic. This emphasizes the flexibility of our method in terms of applicability and leaves room for further improvements.

ACKNOWLEDGMENTS

This work was funded (in part) by the Swiss National Science Foundation under Grant 200020-121763. The authors

wish to thank Dr. Kunal N. Chaudhury (Princeton University, New Jersey) for the fruitful discussions and critiques.

APPENDIX A FULL-MULTIGRID V-CYCLES

Each Full-Multigrid V-cycle $\text{FMG}(h)$ parameterized by $\nu_0 \in \mathbb{N}^*$ corresponds to the recursive function shown on the left of of Figure 9, the residual at a given grid Ω^h being defined as

$$\mathbf{R}(\Omega^h) = \mathbf{d}^h - \mathbf{A}^h \mathbf{c}^h. \quad (62)$$

The expression \mathbf{I}^\downarrow denotes a *restriction operator*, which transfers a sequence from a grid Ω^h towards a coarser grid Ω^{h+1} , while \mathbf{I}^\uparrow denotes a *prolongation operator*, which transfers a sequence from a grid Ω^{h+1} towards a finer grid Ω^h . The operator $\mathbf{V}(h)$ performs one V-Cycle at grid Ω^h , which is itself a recursive function. The two arguments given to the iterator \mathcal{I} correspond to the descending or ascending phase

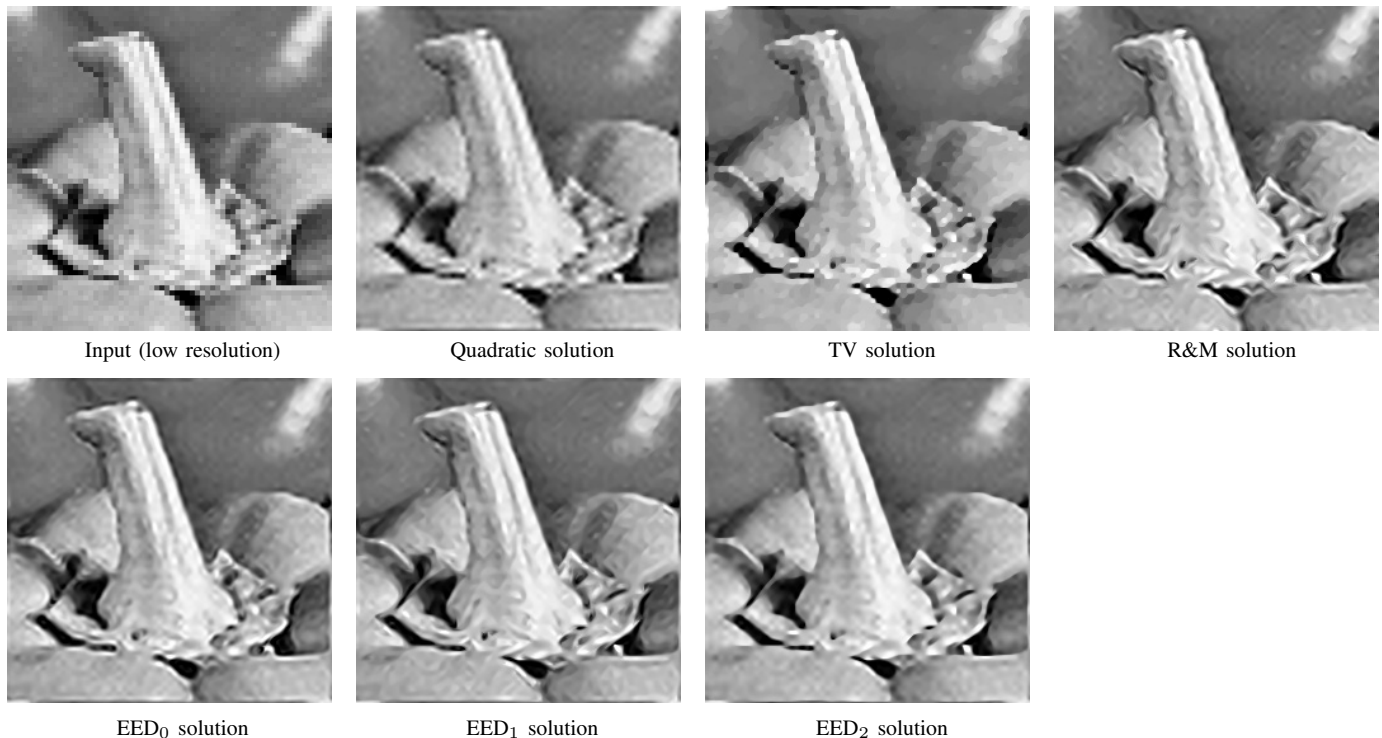


Fig. 8. Magnification of *Peppers* (64×64 crop) by a factor 4 along each dimension. The input low-resolution image is shown at the same scale using pixel duplication.

of the V-Cycle, and to the grid level, respectively. They determine the number of iterations to perform, as well as their parameterization.

REFERENCES

- [1] C.E. Shannon, "Communication in the presence of noise," *Proceedings of the IEEE*, vol. 72, no. 9, pp. 1192–1201, September 1984.
- [2] P. Thevenaz, T. Blu, and M. Unser, "Interpolation revisited," *IEEE Transactions on Medical Imaging*, vol. 19, no. 7, pp. 739–758, July 2000.
- [3] M. Unser and A. Aldroubi, "A general sampling theory for nonideal acquisition devices," *IEEE Transactions on Signal Processing*, vol. 42, no. 11, pp. 2915–2925, November 1994.
- [4] F. Malgouyres and F. Guichard, "Edge direction preserving image zooming: A mathematical and numerical analysis," *SIAM Journal on Numerical Analysis*, vol. 39, no. 1, pp. 1–37, February 2001.
- [5] H.A. Aly and E. Dubois, "Image up-sampling using total-variation regularization with a new observation model," *IEEE Transactions on Image Processing*, vol. 14, no. 10, pp. 1647–1659, October 2005.
- [6] A. Hirabayashi and M. Unser, "Consistent sampling and signal recovery," *IEEE Transactions on Signal Processing*, vol. 55, no. 8, pp. 4104–4115, August 2007.
- [7] A. Bourquard, P. Thévenaz, K. Balać, and M. Unser, "Consistent and regularized magnification of images," in *Proceedings of the 2008 IEEE International Conference on Image Processing*, San Diego, CA, USA, October 12–15, 2008, pp. 325–328.
- [8] A. Roussos and P. Maragos, "Reversible interpolation of vectorial images by an anisotropic diffusion-projection PDE," *International Journal on Computer Vision*, vol. 84, no. 2, pp. 130–145, August 2009.
- [9] P. Getreuer, "Contour stencils: Total variation along curves for adaptive image interpolation," *SIAM Journal on Imaging Sciences*, vol. 4, no. 3, pp. 954–979, September 2011.
- [10] M. Unser, "Multigrid adaptive image processing," in *Proceedings of the International Conference on Image Processing*, Washington, DC, USA, October 23–26, 1995, vol. 1, pp. 49–52.
- [11] M. Arigovindan, M. Sühling, P. Hunziker, and M. Unser, "Variational image reconstruction from arbitrarily spaced samples: A fast multiresolution spline solution," *IEEE Transactions on Image Processing*, vol. 14, no. 4, pp. 450–460, April 2005.
- [12] L. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D*, vol. 60, no. 1–4, pp. 259–268, November 1992.
- [13] T. Pock, D. Cremers, H. Bischof, and A. Chambolle, "Global solutions of variational models with convex regularization," *SIAM Journal on Imaging Sciences*, vol. 3, no. 4, pp. 1122–1145, December 2010.
- [14] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *Journal of Mathematical Imaging and Vision*, vol. 40, no. 1, pp. 120–145, May 2011.
- [15] I. Galic, J. Weickert, M. Welk, A. Bruhn, A. Belyaev, and H.-P. Seidel, "Image compression with anisotropic diffusion," *Journal of Mathematical Imaging and Vision*, vol. 31, no. 2–3, pp. 255–269, July 2008.
- [16] G. Facciolo, P. Arias, V. Caselles, and G. Sapiro, "Exemplar-based interpolation of sparsely sampled images," in *Lecture Notes in Computer Science*, Bonn, Germany, August 24–27, 2009, pp. 331–344.
- [17] D. Tschumperlé, "Fast anisotropic smoothing of multi-valued images using curvature-preserving PDE's," *International Journal of Computer Vision*, vol. 68, no. 1, pp. 65–82, June 2006.
- [18] F. Bornemann and T. März, "Fast image inpainting based on coherence transport," *Journal of Mathematical Imaging and Vision*, vol. 28, no. 3, pp. 259–278, July 2007.
- [19] J. Weickert, *Anisotropic Diffusion in Image Processing*, B.G. Teubner, Stuttgart, Germany, 1998.
- [20] C.R. Vogel and M.E. Oman, "Iterative methods for total variation denoising," in *SIAM Journal on Scientific Computing*, Breckenridge, CO, USA, April 5–9, 1994, vol. 17, pp. 227–238.
- [21] T.F. Chan and P. Mulet, "On the convergence of the lagged diffusivity fixed point method in total variation image restoration," *SIAM Journal on Numerical Analysis*, vol. 36, no. 2, pp. 354–367, February 1999.
- [22] A. Douiri, M. Schweiger, J. Riley, and S.R. Arridge, "Anisotropic diffusion regularization methods for diffuse optical tomography using edge prior information," *Measurement Science & Technology*, vol. 18, no. 1, pp. 87–95, January 2007.
- [23] S. Grevenig, J. Weickert, and A. Bruhn, "From box filtering to fast explicit diffusion," in *Pattern Recognition - 32nd DAGM Symposium, Proceedings*, Darmstadt, Germany, September 22–24, 2010, pp. 533–542.
- [24] A. Roussos and P. Maragos, "Tensor-based image diffusions derived from generalizations of the total variation and beltrami functionals,"

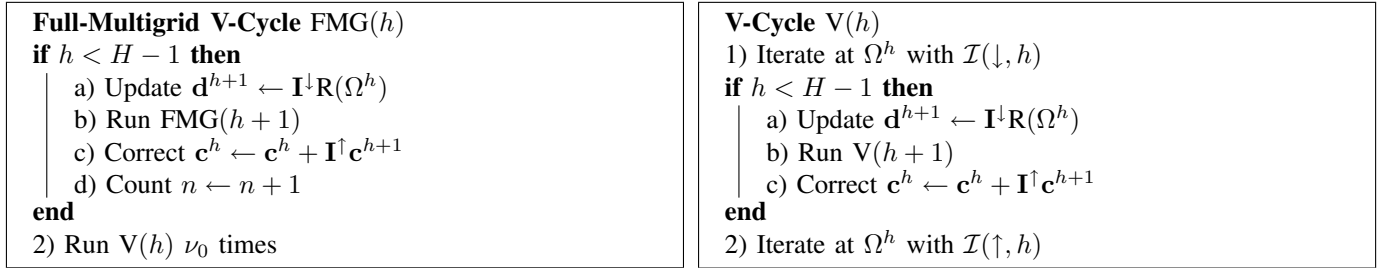



Fig. 9. Multigrid algorithms used in our implementation.

in *IEEE International Conference on Image Processing*, Hong Kong, China, September 26-29, 2010, pp. 4141–4144.

- [25] D. Geman and Y. Chengda, “Nonlinear image recovery with half-quadratic regularization,” *IEEE Transactions on Image Processing*, vol. 4, no. 7, pp. 932–946, July 1995.
- [26] M. Nikolova and R.H. Chan, “The equivalence of half-quadratic minimization and the gradient linearization iteration,” *IEEE Transactions on Image Processing*, vol. 16, no. 6, pp. 1623–1627, June 2007.
- [27] M. Unser, “Splines: A perfect fit for signal and image processing,” *IEEE Signal Processing Magazine*, vol. 16, no. 6, pp. 22–38, November 1999.
- [28] M. Unser and T. Blu, “Generalized smoothing splines and the optimal discretization of the Wiener filter,” *IEEE Transactions on Signal Processing*, vol. 53, no. 6, pp. 2146–2159, June 2005.
- [29] D. Tschumperlé and R. Deriche, “Anisotropic diffusion partial differential equations for multichannel image regularization: Framework and applications,” *Advances in Imaging and Electron Physics*, vol. 145, no. suppl., pp. 149–209, March 2007.
- [30] T.F. Chan, S. Osher, and J. Shen, “The digital TV filter and nonlinear denoising,” *IEEE Transactions on Image Processing*, vol. 10, no. 2, pp. 231–241, February 2001.
- [31] J.P. Oliveira, J.M. Bioucas-Dias, and M.A.T. Figueiredo, “Adaptive total variation image deblurring: A majorization-minimization approach,” *Signal Processing (Netherlands)*, vol. 89, no. 9, pp. 1683–93, September 2009.
- [32] F.H. Harlow and J.E. Welch, “Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface,” *Physics of Fluids*, vol. 8, no. 12, pp. 2182–2189, December 1965.
- [33] M. Nikolova and M.K. Ng, “Analysis of half-quadratic minimization methods for signal and image recovery,” *SIAM Journal of Scientific Computing*, vol. 27, no. 3, pp. 937–966, December 2005.
- [34] I. Galic, J. Weickert, M. Welk, A. Bruhn, A. Belyaev, and H.-P. Seidel, “Towards PDE-Based Image Compression,” in *Variational, Geometric, and Level Set Methods in Computer Vision. Third International Workshop, VLSM 2005. Proceedings*, Beijing, China, October 16, 2005, pp. 37–48.
- [35] D. Tschumperlé and R. Deriche, “Vector-valued image regularization with PDEs: A common framework for different applications,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 4, pp. 506–517, April 2005.
- [36] N. Sochen, R. Kimmel, and R. Malladi, “A general framework for low level vision,” *IEEE Transactions on Image Processing*, vol. 7, no. 3, pp. 310–318, March 1998.
- [37] P. Perona and J. Malik, “Scale-space and edge detection using anisotropic diffusion,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 7, pp. 629–639, July 1990.
- [38] A. Selinger, R.C. Nelson, and A.C. Nelson, “Using directional variance to extract curves in images, thus improving object recognition in clutter,” Tech. Rep. UR CSD / TR712, Department of Computer Science, University of Rochester, 1999.
- [39] M. Grasmair and F. Lenzen, “Anisotropic total variation filtering,” *Applied Mathematics and Optimization*, vol. 62, no. 3, pp. 323–339, December 2010.
- [40] W.L. Briggs, V.E. Henson, and S.F. McCormick, *A multigrid tutorial*, Society for Industrial and Applied Mathematics, University City Science Center, Philadelphia, PA, USA, 2nd edition, July 1, 2000.
- [41] D.M. Young, “Iterative methods for solving partial difference equations of elliptic type,” *Transactions of the American Mathematical Society*, vol. 76, no. 1, pp. 92–111, January 1954.
- [42] P. Getreuer, “Roussos-Maragos tensor-driven diffusion for image interpolation,” *Image Processing On Line*, September 2011, DOI: http://dx.doi.org/10.5201/ipol.2011.g_rmdi.



Aurélien Bourquard (M’08) received the M.Sc. degree in microengineering from the Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland, in 2008. He is currently working toward the Ph.D. degree with the Biomedical Imaging Group, EPFL.

His research interests include image reconstruction using convex optimization and multigrid techniques, as well as new acquisition methods in the framework of computational optics.



Michael Unser (M’89-SM’94-F’99) received the M.S. (*summa cum laude*) and Ph.D. degrees in electrical engineering from the Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland, in 1981 and 1984, respectively.

From 1985 to 1997, he worked as a Scientist with the National Institutes of Health, Bethesda, MD. He is currently a Full Professor and a Director of the Biomedical Imaging Group, EPFL. He is the author of 200 journal papers. His main research interests include biomedical image processing, sampling theories, multiresolution algorithms, wavelets, and the use of splines for image processing.

Dr. Unser is a EURASIP Fellow and a member of the Swiss Academy of Engineering Sciences. He was the Associate Editor for the *IEEE Signal Processing Letters* in 1994–1998, for the *IEEE Transactions on Image Processing* in 1992–1995, and for the *IEEE Transactions on Medical Imaging* in 1999–2002 and 2006–2007. He has also held the position of associate Editor-in-Chief in 2003–2005 for the *IEEE Transactions on Medical Imaging*. He is currently a member of the editorial boards of *Foundations and Trends in Signal Processing* and *Sampling Theory in Signal and Image Processing*. He coorganized the first *IEEE International Symposium on Biomedical Imaging* and was the Founding Chair of the technical committee of the *IEEE-SP Society on Bio Imaging and Signal Processing*. He was the recipient of the 1995 and 2003 Best Paper Awards, the 2000 Magazine Award, and two *IEEE Technical Achievement Awards* (2008 SPS and 2010 EMBS). He is one of ISIs Highly Cited Authors in Engineering (<http://isihighlycited.com>).