

AN UPWIND FINITE-DIFFERENCE METHOD FOR TOTAL VARIATION-BASED IMAGE SMOOTHING

ANTONIN CHAMBOLLE*, STACEY E. LEVINE†, AND BRADLEY J. LUCIER‡

Abstract. In this paper we study finite-difference approximations to the variational problem using the BV smoothness penalty that was introduced in an image smoothing context by Rudin, Osher, and Fatemi. We give a dual formulation for an upwind finite-difference approximation for the BV seminorm; this formulation is in the same spirit as one popularized by the first author for a simpler, less isotropic, finite-difference approximation to the (isotropic) BV seminorm. We introduce a multiscale method for speeding the approximation of both Chambolle’s original method and of the new formulation of the upwind scheme. We demonstrate numerically that the multiscale method is effective, and we provide numerical examples that illustrate both the qualitative and quantitative behavior of the solutions of the numerical formulations.

AMS subject classifications. 15A15, 15A09, 15A23

1. Introduction. In an influential paper, Rudin, Osher, and Fatemi [28] suggested using the bounded variation seminorm to smooth images. The functional proposed in their work has since found use in a wide array of problems (see, e.g., [9]), both in image processing and other applications. The novelty of the work was to introduce a method that preserves discontinuities while removing noise and other artifacts.

In the continuous setting, the behavior of the solutions of the model proposed in [28] is well understood (see, e.g., [11], [12], and [27]). The qualitative properties of solutions of its discrete versions are not, perhaps, as well known or understood. In this work we study the behavior of solutions of the discrete approach used in, e.g., [13] as well as a new “upwind” variant of this model that better preserves edges and “isotropic” features. We also introduce a multiscale method for improving the initial guess of certain iterative methods for solving discrete variational problems based on the BV variational model.

We begin by giving some background on the work in [28]. Working on the unit square $I = [0, 1]^2$, where the (isotropic) bounded variation seminorm is defined as

$$\begin{aligned} |g|_{\text{BV}(I)} &:= \int_I |Dg(x)| \, dx \\ &:= \sup \left\{ \int_I g \nabla \cdot p \mid p : I \rightarrow \mathbb{R}^2, p \in C_0^1(I), |p(x)| \leq 1 \text{ for all } x \in I \right\}. \end{aligned} \tag{1.1}$$

The formulation of BV smoothing on which we depend is as follows: Given a non-constant function f and a number $\lambda > 0$, find the function \tilde{f} that minimizes over

*CMAP, Ecole Polytechnique, CNRS, 91128 Palaiseau, France. The work of this author was partially supported by the ANR “MICA” project, grant ANR-08-BLAN-0082 (antonin.chambolle@polytechnique.fr).

†Department of Mathematics and Computer Science, Duquesne University, 440 College Hall, Pittsburgh, PA 15282, USA. The work of this author was partially supported by NSF-DMS grant #0505729 and the Institute for Mathematics and its Applications with funds provided by the National Science Foundation (sel@mathcs.duq.edu).

‡Department of Mathematics, Purdue University, 150 N. University St., West Lafayette, IN 47907, USA. The work of this author was partially supported by the Office of Naval Research, Contract N00014-91-J-1152, and by the Institute for Mathematics and its Applications with funds provided by the National Science Foundation (lucier@math.purdue.edu).

all g

$$\frac{1}{2}\|f - g\|_{L_2(I)}^2 + \lambda|g|_{\text{BV}(I)}. \quad (1.2)$$

Fix f ; for each λ this problem has a unique solution that satisfies

$$\frac{\tilde{f} - f}{\lambda} + \partial_X|\tilde{f}|_{\text{BV}(I)} \ni 0, \quad (1.3)$$

where $\partial_X\varphi(g)$ is the subdifferential of the convex, lower-semicontinuous map $\varphi : L_2(I) \rightarrow \mathbb{R}$. (See [21] for definitions and the basic results we quote here.) Furthermore, if we set

$$\|\tilde{f} - f\|_{L_2(I)}^2 = \sigma^2,$$

then σ^2 is a continuous, one-to-one, increasing function of λ and

$$\sigma^2 < \left\| f - \int_I f \right\|_{L_2(I)}^2. \quad (1.4)$$

Thus, given σ^2 that satisfies (1.4) there is a unique λ such that the solution \tilde{f} of (1.2) given above is also the solution of the problem: “Find \tilde{f} that minimizes the functional, $|g|_{\text{BV}(I)}$, over all g with $\|f - g\|_{L_2(I)}^2 = \sigma^2$.” Rudin, Osher, and Fatemi [28] introduced BV image smoothing in the latter form, but these two formulations are equivalent [15]. About the same time, Bouman and Sauer [7] proposed a discrete version of (1.2) in the context of tomography. More recently, the BV smoothing technique is used as just one step in the so-called inverse scale space approach [25, 10].

The main purpose of this work is twofold. First we propose a new formulation of an upwind finite-difference approximation to the bounded variation seminorm based on work by Osher and Sethian [26]. We then study qualitative properties of the solutions given by this new method and the discrete approximation of (1.2) given in [13].

The paper is organized as follows. In Section 2, we discuss the algorithm introduced in [13] for minimizing a finite-difference approximation to the ROF functional. In Section 3, we present the new upwind formulation, and show how this can be used to minimize the ROF functional in a manner similar to [13]. In Section 4, we introduce a multiscale algorithm that greatly reduces the run time of both methods.

Section 5 contains numerical examples that demonstrate the qualitative properties of the algorithm and observed rates of convergence for two special problems with known solutions. We also illustrate the improved isotropy of the upwind method when removing noise from a synthetic image; we finish by computing “cartoon/texture/noise” decompositions of two natural images to show that the upwind method keeps more details of the images.

2. Discrete BV Variational Smoothing. To begin we consider discretizations of (1.2). A standard approach, as first suggested in [28], is to regularize the BV seminorm and consider the problem of minimizing, with $\epsilon > 0$,

$$\frac{1}{2}\|f - g\|_{L_2(I)}^2 + \lambda \int_I \sqrt{|\nabla g|^2 + \epsilon^2}.$$

This functional is differentiable in g , and one can follow the flow of its associated Euler-Lagrange equation; numerical methods approximate this flow. This is sufficient

for some applications, but the solution now depends on the regularization parameter ϵ .

In this work we consider a discrete analog of (1.2) and follow the dual approach proposed in [13]. The material in this section is classical; we refer the reader to [13] for a more extended treatment, and to [17] which puts [13] into some historical context. Another approach to solving the discrete problem that arises here is through second-order cone programming [22]. Additionally, if one considers an anisotropic variant e.g.,

$$|g| = \int_I |g_x| + |g_y| \quad (\text{interpreted appropriately})$$

to the usual isotropic BV seminorm, then one can apply methods involving graph cuts [18].

Given $N > 1$, we let $h = 1/N$ and consider discrete functions

$$g_i, \quad i = (i_1, i_2), \quad 0 \leq i_1, i_2 < N.$$

A discrete $L_2^h(I)$ norm of g is defined by

$$\|g\|_{L_2^h(I)}^2 = \sum_{0 \leq i_1, i_2 < N} |g_i|^2 h^2.$$

One way to compute a discrete gradient of a discrete scalar function f_i is given by

$$\nabla_h g_i := \left(\frac{g_{i+(1,0)} - g_i}{h}, \frac{g_{i+(0,1)} - g_i}{h} \right)^T. \quad (2.1)$$

When we use this discrete gradient in the general formula (2.2), we say that the resulting discrete variational functional $|g|_{\text{BV}^h(I)}$ is an “anisotropic” approximation to the *isotropic* seminorm $|g|_{\text{BV}(I)}$.

For any discrete gradient operator, one needs to specify the value of g_i for some values of i outside $[0, N]^2$; to do so, we need to specify boundary conditions. We assume that scalar discrete functions g, f , etc., are either periodic (with period N) or satisfy Neumann or Dirichlet boundary conditions. For Dirichlet boundary conditions, g_i is zero for i outside $[0, N]^2$; for Neumann conditions, we consider g to be reflected across the lines $i_1 = N - 1/2$ and $i_2 = N - 1/2$ and then extended periodically across the plane with period $2N$.

Given the discrete gradient (2.1) (or indeed *any* discrete gradient), we can define an associated discrete BV seminorm

$$|g|_{\text{BV}^h(I)} = \sum_{0 \leq i_1, i_2 < N} |\nabla_h g_i| h^2 \quad (2.2)$$

and then in turn a discrete analogue to (1.2),

$$\frac{1}{2} \|f - g\|_{L_2^h(I)}^2 + \lambda |g|_{\text{BV}^h(I)}. \quad (2.3)$$

For any discrete gradient, we define an associated discrete divergence of vector functions $p_i = (p_i^{(1)}, p_i^{(2)})$; given the discrete gradient (2.1), we define the associated divergence by

$$\nabla_h \cdot p_i = \frac{p_i^{(1)} - p_{i-(1,0)}^{(1)}}{h} + \frac{p_i^{(2)} - p_{i-(0,1)}^{(2)}}{h}. \quad (2.4)$$

As for boundary conditions, we note that in the following we compute discrete divergences only of discrete vector fields that are themselves discrete gradients; therefore we compute p_i for i outside of $[0, N]^2$ in a manner consistent with whatever boundary condition we have chosen for discrete scalar functions f_i .

Because of how the discrete gradient and divergence are related, we have

$$\begin{aligned} |g|_{\text{BV}^h(I)} &= \sum_{0 \leq i_1, i_2 < N} |\nabla_h g_i| h^2 \\ &= \sup_{|p_i| \leq 1} \sum_{0 \leq i_1, i_2 < N} (-\nabla_h) g_i \cdot p_i h^2 \\ &= \sup_{|p_i| \leq 1} \sum_{0 \leq i_1, i_2 < N} g_i (\nabla_h \cdot p_i) h^2. \end{aligned} \quad (2.5)$$

The first equality is obvious; the second (which can be interpreted as “the adjoint of the discrete divergence is the negative of the discrete gradient”) follows by summation by parts.

Thus, if the symmetric convex set \overline{K} is defined by

$$\overline{K} = \{u_i = \nabla_h \cdot p_i \mid |p_i| \leq 1, p_i = (p_i^{(1)}, p_i^{(2)})\}, \quad (2.6)$$

then

$$|g|_{\text{BV}^h(I)} = \sup_{u \in \overline{K}} \sum_{0 \leq i_1, i_2 < N} g_i u_i h^2 =: \langle g, u \rangle.$$

Using classical convexity arguments, it is shown in [13] that the minimizer \tilde{f} of (2.3) is $f - \lambda \nabla_h \cdot \bar{p}$ with \bar{p} a minimizer of

$$F(p) := \left\| \nabla_h \cdot p - \frac{f}{\lambda} \right\|_{L_2(I)}^2 \quad (2.7)$$

subject to the constraint

$$p \in K := \{p : [0, N]^2 \rightarrow \mathbb{R}^2 \mid |p_i| \leq 1 \text{ for all } i\}. \quad (2.8)$$

In other words,

$$\tilde{f} = f - \pi_{\lambda \overline{K}} f, \quad (2.9)$$

where $\pi_{\lambda \overline{K}}$ is the orthogonal projector in $L_2^h(I)$ of f onto the convex set $\lambda \overline{K}$.

A specific iterative algorithm for finding a discrete vector field p that minimizes (2.7) is also given in [13]. One sets $p_i^0 = 0$ for all i (so that p^0 is obviously in K) and then calculates

$$p_i^{n+1} := \frac{p_i^n - \tau(-\nabla_h)(\nabla_h \cdot p^n - f/\lambda)_i}{1 + \tau|(-\nabla_h)(\nabla_h \cdot p^n - f/\lambda)_i|}. \quad (2.10)$$

He notes that $p^n \in K$ for all n and shows, using the Karush-Kuhn-Tucker theorem, that if $0 < \tau \leq h^2/8$ the limit $\lim_{n \rightarrow \infty} p^n$ exists and gives a minimizer of (2.7) over all p with $|p_i| \leq 1$. Procedure (2.10) can be written as the two-step process

$$p_i^{n+1/2} := p_i^n - \tau(-\nabla_h)(\nabla_h \cdot p^n - f/\lambda)_i, \quad (2.11)$$

$$p_i^{n+1} := \frac{p_i^{n+1/2}}{1 + |p_i^{n+1/2} - p_i^n|}, \quad (2.12)$$

where the first formula is just gradient descent of the functional $F(p)$ with step τ , and the second is a nonlinear projector that ensures that $p^{n+1} \in K$ if $p_i^n \in K$.

In [14], the same author speculates whether the two-step procedure

$$p_i^{n+1/2} := p_i^n - \tau(-\nabla_h)(\nabla_h \cdot p^n - f/\lambda)_i, \quad (2.13)$$

$$p_i^{n+1} := \frac{p_i^{n+1/2}}{\max(1, |p_i^{n+1/2}|)} \quad (2.14)$$

yields vector fields p^n such that $\nabla_h \cdot p^n$ converges to the projection $\pi_{\overline{K}}(f/\lambda)$. Again, the first half-step is gradient descent along the functional $F(p)$ given in (3.10), while the second is the L_2^h -projection of p onto the set K . We note at the end of the next section that this iteration is a special case of a more general minimization algorithm developed originally by Eicke [20], which itself has been generalized by Combettes [16, 17]. There is recent work related to algorithms of this type by Aujol [4].

3. Upwind BV Smoothing. One might think that, given $g \in \text{BV}(I)$, the $L_2(I)$ projection of g on a grid with sidelength h ,

$$g_i = \frac{1}{h^2} \int_{I_i} g, \quad I_i = h(I + i),$$

or the (multi-valued) $L_1(I)$ projection of g on the same grid

$$\begin{aligned} g_i &= \text{any } m \text{ such that } |\{x \in I_i \mid g(x) \geq m\}| \geq 1/2 \\ &\quad \text{and } |\{x \in I_i \mid g(x) \leq m\}| \geq 1/2 \end{aligned} \quad (3.1)$$

would satisfy

$$\lim_{h \rightarrow 0} |g|_{\text{BV}^h(I)} = |g|_{\text{BV}(I)}, \quad (3.2)$$

but this is not true in general. If g is C^1 , then integration by parts in (1.1) shows that (3.2) holds. Again, if g is the characteristic function of the set $\{x_1 < 1/2\}$ or $\{x_2 < 1/2\}$, or even $\{x_1 + x_2 < 1\}$, then (3.2) holds, but if g is the characteristic function of $\{x_1 < x_2\}$ and we use the projection (3) then a calculation shows that

$$2 = \lim_{h \rightarrow 0} |g|_{\text{BV}^h(I)} \neq |g|_{\text{BV}(I)} = \sqrt{2}. \quad (3.3)$$

In fact, this is not so much of an issue as far as *minimization problems* are concerned, since it is well known in this case that the correct notion of convergence is Γ -convergence [8], which ensures convergence of the minimizers of variational problems. It can be shown without much difficulty that the semi-norms $|\cdot|_{\text{BV}^h}$ Γ -converge as $h \rightarrow 0$ to the BV seminorm. However, it follows from the inequality (3.3) that the approximation of $\chi_{\{x_1 < x_2\}}$, as a possible solution of a minimization problem, will be possible only after some smoothing of the discontinuity, so that the output of a discrete minimization will usually not be as sharp as one could hope.

This anisotropy motivates us to define an “upwind” discrete $\text{BV}(I)$ norm of a discrete scalar function g_i given by

$$\|g\|_{\text{BV}^h(I)} := \sum_i |(-\nabla_h)g_i \vee 0| h^2, \quad (3.4)$$

where we have defined the discrete gradient

$$(-\nabla_h)g_i = \left(\frac{g_i - g_{i+(1,0)}}{h}, \frac{g_i - g_{i-(1,0)}}{h}, \frac{g_i - g_{i+(0,1)}}{h}, \frac{g_i - g_{i-(0,1)}}{h} \right)^T. \quad (3.5)$$

and we denote by $p \vee q$ and $p \wedge q$ the componentwise maximum and minimum, respectively, of the vectors p and q . (Similarly, if we write an inequality between vectors, $p \leq q$, then we mean that this inequality holds componentwise.) This type of operator is based on the classical first-order upwind finite-difference scheme used to solve hyperbolic partial differential equations; upwind methods have found important applications in level set methods [26].

In other words, we include a difference in the vector norm of the i th term in (3.4) only if g^h is increasing into g_i^h . Nothing changes in the following calculations (and one sees little change in the images themselves) if we change componentwise maximum (\vee) to componentwise minimum (\wedge) in (3.4). In their paper, Osher and Sethian [26] were solving Hamilton-Jacobi equations where this substitution could not be made: their problem, unlike ours, has a true notion of “wind”.

In the present case, the problem we are solving is degenerate elliptic and there is, strictly speaking, no direction of “wind.” However it still produces the desired effect, which is to preserve some discontinuities better than standard discretizations. We may refer, for a similar idea, to Appleton and Talbot [3] who recently proposed to compute minimal surfaces by solving some hyperbolic system, discretized with an upwind scheme. Also in their case the “speed” and “wind” can be reversed, still, their approach produces sharp discontinuities as desired — on the other hand, it does not really correspond to the minimization of a convex discrete functional such as our upwind TV.

In (3.4) we include a difference in the sum only if the difference is *positive*, i.e., the discrete function g_i is *increasing* as it goes to g_i from the given direction. Note that for smooth $g(x)$ this is a convergent approximation to the BV seminorm of f and for jumps across vertical, horizontal, or diagonal lines you get the correct value of the BV seminorm; that is, $\lim_{h \rightarrow 0} |g|_{\text{BV}^h(I)} = |g|_{\text{BV}(I)}$.

We then write this “upwind” seminorm as

$$\sum_i |(-\nabla_h)g_i \vee 0| h^2 = \sup_{|p_i| \leq 1, p_i \geq 0} \sum_i (-\nabla_h)g_i \cdot p_i h^2.$$

Here we require not only that the Euclidean norm $|p_i|$ of $p_i \in \mathbb{R}^4$ be no larger than one, but also that each coordinate of p_i be non-negative, so that p is in the set

$$K := \{p : [0, N]^2 \rightarrow \mathbb{R}^4 \mid |p_i| \leq 1 \text{ and } p_i \geq 0\}. \quad (3.6)$$

Thus we have dealt with the “extra” nonlinearity of (3.4) by incorporating it into the convex set that contains p .

If we now define the discrete divergence that is the adjoint of the discrete gradient (3.5),

$$\nabla_h \cdot \xi_i = \frac{\xi_i^{(1)} - \xi_{i-(1,0)}^{(1)}}{h} + \frac{\xi_i^{(2)} - \xi_{i+(1,0)}^{(2)}}{h} + \frac{\xi_i^{(3)} - \xi_{i-(0,1)}^{(3)}}{h} + \frac{\xi_i^{(4)} - \xi_{i+(0,1)}^{(4)}}{h}, \quad (3.7)$$

and again apply summation by parts, we see that this new discrete seminorm is equal to

$$\sup_u \sum_i g_i u_i h^2 = \sup_u \langle g, u \rangle$$

with u in the convex set

$$\overline{K} := \{\nabla_h \cdot p \mid p : [0, N)^2 \rightarrow \mathbb{R}^4, |p_i| \leq 1, p_i \geq 0\}. \quad (3.8)$$

Thus as in [13], the minimizer over discrete scalar functions g_i of

$$\frac{1}{2} \|f - g\|^2 + \lambda \|(-\nabla_h)g \vee 0\|_1 \quad (3.9)$$

can be written as the difference between f and the unique projection of f onto the convex set $\lambda \overline{K}$.

In other words, the minimizer of (3.9) is $f - \lambda \nabla_h \cdot \bar{p}$ where \bar{p} is any minimizer of the functional

$$F(p) := \left\| \nabla_h \cdot p - \frac{f}{\lambda} \right\|^2 \quad (3.10)$$

subject to the constraint that $p \in K$, where K is defined by (3.6). An iterative method to compute a p that minimizes (3.10) is given by

$$p_i^{n+1/3} := p_i^n - \tau(-\nabla_h)(\nabla_h \cdot p^n - f/\lambda)_i, \quad (3.11)$$

$$p_i^{n+2/3} := p_i^{n+1/3} \vee 0, \quad (3.12)$$

$$p_i^{n+1} := \frac{p_i^{n+2/3}}{\max(1, |p_i^{n+2/3}|)}, \quad (3.13)$$

where p_i^0 is chosen arbitrarily in K . The computation of $p_i^{n+1/3}$ is simply gradient descent of (3.10), while the next two steps compute the projection of $p^{n+1/3}$ onto K .

We remark that [29] bounds the difference in $L_2(I)$ between discrete minimizers of (3.9) and the minimizer of (1.2) as the mesh size $h \rightarrow 0$.

The notational similarity of (3.10) and (2.7) is deliberate; both can be formulated as minimizing over all x in a closed convex set K

$$F(x) = \|Ax - b\|^2 \quad (3.14)$$

for some bounded linear operator A and vector b . In our cases we have $x = p$, $Ax = \nabla_h \cdot p$, $b = f/\lambda$, and K is given by (2.8) or (3.6).

We consider the general iteration

$$x^{n+1} = \pi_K(x^n - \tau(A^*(Ax^n - b))), \quad (3.15)$$

where π_K is the orthogonal projection onto the set K , A^* is the adjoint of the operator A , and τ is suitably small. In other words, we first perform gradient descent on the functional $F(x)$ and then project the intermediate result onto the convex set K .

The convergence of this algorithm was studied by Eicke [20], Theorem 3.2, and is a special case of a general theory developed later by Combettes and his collaborators [16, 17]. For pedagogical purposes we recommend the analysis in [20], which is particularly short and self contained. The result applied to our (finite-dimensional) problem gives the following: if $0 < \tau < 2\|A\|^{-2}$ then x^n converges to a minimizer of $F(x)$ on K . (Part of the result goes back to Opial [24].) In our case this means that the method converges if $0 < \tau < h^2/8$ for the discrete divergence (2.4) [13]; a similar argument shows that for the discrete divergence (3.7) we obtain convergence when $0 < \tau < h^2/16$.

The iteration (3.15) is efficient only if π_K , A , and A^* can be calculated quickly. In our case, it takes $O(N^2)$ operations to calculate $\nabla_n \cdot p$ or $-\nabla_h f$ on an $N \times N$ image. For the set K defined by (2.8), we have simply

$$(\pi_K p)_i = \frac{p_i}{\max(1, |p_i|)}.$$

For K defined by (3.6) we set

$$(\pi_K p)_i = \frac{\bar{p}_i}{\max(1, |\bar{p}_i|)} \text{ where } \bar{p}_i = p_i \vee 0.$$

So with either (2.8) or (3.6) we can calculate $\pi_K p$ on an $N \times N$ image in $O(N^2)$ operations. Practically, each step of the upwind scheme (3.11)–(3.13) takes about four times as many operations as one step of the scheme (2.13)–(2.14).

4. A Multiscale Algorithm. The characterization of the minimizer (2.9) allows us the following observation: We need only construct a vector field p that minimizes $F(p)$ over all $p \in K$. Any general iteration of the form (3.15) converges as long as the initial data is in the set K . We propose here to use a multiscale technique to get a good approximation $p^0 \in K$ for our iterations.

We consider two grids in I , one with grid spacing $2h$ and one with grid spacing h . We will construct a scalar injector from the $2h$ -grid to the h -grid (called I_{2h}^h) and a scalar projector from the h grid to the $2h$ grid (called I_h^{2h}). Similarly, we will have an operator-dependent injector \tilde{I}_{2h}^h on vector fields. Our general approach will then be as follows.

Given data f_h on a grid with spacing h , we calculate data $f_{2h} = I_h^{2h} f_h$ on a grid with spacing $2h$. We then calculate the minimizer p_{2h} of (3.10) with data f_{2h} using our iterative algorithm (not yet specifying the initial value p^0). Next, we begin the iteration solving (3.10) with data f_h with the initial vector field $p^0 = \tilde{I}_{2h}^h p_{2h}$.

We now explain our choice of I_h^{2h} and \tilde{I}_{2h}^h . Assume that N is even with $h = 1/N$, and define the discrete inner product on $N \times N$ arrays

$$\langle u, v \rangle_h = \sum_i u_i v_i h^2$$

where the sum is taken over all $i = (i_1, i_2)$ with $0 \leq i_1, i_2 < N$. A similar inner product can be defined on $N/2 \times N/2$ arrays with grid spacing $2h$:

$$\langle u, v \rangle_{2h} = \sum_i u_i v_i (2h)^2,$$

where the sum now is over all $i = (i_1, i_2)$ with $0 \leq i_1, i_2 < N/2$.

Our injector I_{2h}^h will simply be the constant injector on 2×2 squares:

$$(I_{2h}^h u)_i = u_{\lfloor i/2 \rfloor},$$

where $\lfloor i/2 \rfloor = (\lfloor i_1/2 \rfloor, \lfloor i_2/2 \rfloor)$ and $\lfloor x \rfloor$ is the largest integer no greater than x . The corresponding projector I_h^{2h} is defined as the adjoint of I_{2h}^h with respect to the h - and $2h$ -inner products, i.e., for an $N \times N$ grid function v and a $N/2 \times N/2$ grid function u

$$\langle I_{2h}^h u, v \rangle_h = \langle u, I_h^{2h} v \rangle_{2h}.$$

A direct calculation shows that

$$(I_h^{2h}v)_i = \frac{1}{4}(v_{2i} + v_{2i+(1,0)} + v_{2i+(0,1)} + v_{2i+(1,1)}),$$

i.e., it is simply the average of the values of v on 2×2 subgrids. (The factor $1/4$ comes in because of the different weights in the two inner products.)

After we calculate the minimizer p_{2h} of (3.10) over all $p \in K$ on the grid with spacing $2h$, we start the iteration on the grid with spacing h with

$$p^0 = \bar{I}_{2h}^h p_{2h}.$$

Here $\bar{I}_{h,2}^h$ is an injector that satisfies

$$\nabla_h \cdot \bar{I}_{2h}^h p_{2h} = I_{2h}^h \nabla_{2h} \cdot p_{2h}. \quad (4.1)$$

Specifically, for the anisotropic operators (2.4) and (2.1) we have

$$(\bar{I}_{2h}^h p)_i = \frac{1}{2} \left(p_{\lfloor i/2 \rfloor}^{(1)} + p_{\lfloor (i-(1,0))/2 \rfloor}^{(1)}, p_{\lfloor i/2 \rfloor}^{(2)} + p_{\lfloor (i-(0,1))/2 \rfloor}^{(2)} \right)^T$$

and for the upwind operators (3.5) and (3.7) we have

$$(\bar{I}_{2h}^h p)_i = \frac{1}{2} \begin{pmatrix} p_{\lfloor i/2 \rfloor}^{(1)} + p_{\lfloor (i-(1,0))/2 \rfloor}^{(1)} \\ p_{\lfloor i/2 \rfloor}^{(2)} + p_{\lfloor (i+(1,0))/2 \rfloor}^{(2)} \\ p_{\lfloor i/2 \rfloor}^{(3)} + p_{\lfloor (i-(0,1))/2 \rfloor}^{(3)} \\ p_{\lfloor i/2 \rfloor}^{(4)} + p_{\lfloor (i+(0,1))/2 \rfloor}^{(4)} \end{pmatrix}.$$

The next section contains some computational experiments that show that this multiscale technique is effective.

5. Experimental Results. We did a series of experiments (a) to measure the effectiveness of the multiscale predictor for the initial vector field p , (b) to examine the experimental convergence rates for two sets of initial data for the continuous BV problem with known analytic solutions, and (c) to illustrate some of the qualitative properties of the solutions to the discrete problems.

5.1. Effectiveness of the multiscale predictor. All of our algorithms compute only an approximate minimizer of discrete BV problems—we must decide when to stop the iteration (3.15). In [14] one finds a simple error bound; first we let

$$\tilde{f}^n = f - \lambda \nabla_h \cdot p^n.$$

Then we have

$$\|\tilde{f} - \tilde{f}^n\|_{L_2^h(I)}^2 \leq \lambda(|\tilde{f}^n|_{\text{BV}^h(I)} - \langle \nabla_h f^n, p^n \rangle) = \epsilon(p^n)^2. \quad (5.1)$$

So we can ensure that we've computed \tilde{f} to within an error of ϵ , i.e.,

$$\|\tilde{f} - \tilde{f}^n\|_{L_2^h(I)} \leq \epsilon,$$

if we iterate until $\epsilon(p^n) \leq \epsilon$. It is possible that the true error is less (perhaps *much* less) than the bound.

In our computations the initial data is 255 times the characteristic function of a disk or of a square, and we set $\epsilon(p^n) = 1/4$, so we ensure that the $L_2^h(I)$ error between our computed discrete minimizer and the exact discrete minimizer is no more than $1/4$ greyscales.

We begin by discussing the efficiency improvements we observed by our multiscale method. As mentioned in the previous paragraph, we do not compute exact discrete minimizers, but only approximate discrete minimizers. Our final goal is to compute \tilde{f}_h^n such that

$$\|\tilde{f}_h - \tilde{f}_h^n\|_{L_2^h(I)} \leq \epsilon_h$$

for some ϵ_h . We begin the iterative algorithm with some initial guess for p_h^0 . Without our multiscale algorithm we take $p_h^0 = 0$. Our multiscale algorithm says that we should iteratively compute p_{2h}^n on a grid with grid size $2h$ until

$$\|\tilde{f}_{2h} - \tilde{f}_{2h}^n\|_{L_2^{2h}(I)} \leq \epsilon_{2h}$$

and then set

$$p_h^0 = \pi_{K_h} \tilde{f}_{2h}^n p_{2h}^n. \quad (5.2)$$

In our computations we take $\epsilon_{2h} = \epsilon_h = 1/4$. It is possible that, given $\epsilon_h = 1/4$, a better choice of ϵ_{2h} can improve our initial vector field p_h^0 .

We compare the computational effort needed by our multiscale method and the iterative method operating solely on the grid with grid spacing h and $p_h^0 = 0$. We note that each iteration $p_{2h}^n \rightarrow p_{2h}^{n+1}$ on a grid with mesh size $2h$ takes (roughly) $1/4$ as many operations as one iteration on the grid of size h . Thus if the number of iterations on grids with spacing $h, 2h, 4h$, etc., are N_h, N_{2h}, N_{4h} , respectively, we report the number of equivalent iterations on the finest grid,

$$N_h + \frac{1}{4}N_{2h} + \frac{1}{16}N_{4h} + \cdots. \quad (5.3)$$

In our examples we take the data to be $f = 255\chi_{[\frac{1}{4}, \frac{3}{4}]^2}$, the characteristic function of a subsquare with sidelength $1/2$ inside the computational domain of $[0, 1]^2$. We use Dirichlet boundary conditions. We computed numerical solutions on grids with $h = 1/128, 1/256$, and $1/512$. We chose three values of λ for which the $L_2(I)$ distance between the solution of the continuous problem (1.2) and the initial data f is 16, 32, and 64. For this purpose we use the characterization of the exact solutions (see [6], Section 4 in [2], or Appendix 1 in [1]), and found that the corresponding values of λ are 3.771636443, 7.820179629, and 16.26268646, respectively.

A summary of the iteration count to solve both problems (2.3) and (3.9) such that the error bound (5.1) is $< 1/4$, both with and without the multiscale approach, are summarized in Tables 5.1 and 5.2. The iteration count with the multiscale approach is reported as the equivalent number of iterations at the finest resolutions as described above.

We see that the multiscale method of choosing p_h^0 speeds up the computation, and greatly so in some cases.

5.2. Experimental Error Estimates. Next we discuss the observed error between the minimizer \tilde{f} of the continuous problem (1.2) and the approximate minimizers \tilde{f}_h^n (with an error in $L_2^h(I)$ of < 0.25) of the two discrete problems (2.3) and (3.9).

TABLE 5.1

Iteration count with $p^0 = 0$ on grids of size 128, 256, and 512, and differences $\sigma = \|f - \tilde{f}\|_{L_2(I)}$ of 16, 32, and 64, with initial data (5.4); columns 1–3 are the result of iterating (2.13)–(2.14), while columns 4–6 are the result of iterating (3.11)–(3.13). In all cases, iterations were stopped when the error bound (5.1) satisfied $\epsilon(p^n) \leq 1/4$.

	Number of iterations with $p^0 = 0$					
	“anisotropic” $ g _{\text{BV}^h(I)}$			“upwind” $ g _{\text{BV}^h(I)}$		
resolution	$\sigma = 16$	$\sigma = 32$	$\sigma = 64$	$\sigma = 16$	$\sigma = 32$	$\sigma = 64$
128×128	4,815	21,772	119,468	4,293	5,414	13,049
256×256	10,466	45,744	255,096	17,173	21,162	33,158
512×512	36,653	103,817	514,060	68,324	83,908	113,843

TABLE 5.2

Equivalent iteration count (5.3) using the multiscale estimate (5.2) for p_h^0 on grids of size 128, 256, and 512, and differences $\sigma = \|f - \tilde{f}\|_{L_2(I)}$ of 16, 32, and 64, with initial data (5.4); columns 1–3 are the result of iterating (2.13)–(2.14), while columns 4–6 are the result of iterating (3.11)–(3.13). In all cases, iterations were stopped when the error bound (5.1) satisfied $\epsilon(p^n) \leq 1/4$.

	Equivalent iterations with the multiscale estimate for p^0					
	“anisotropic” $ g _{\text{BV}^h(I)}$			“upwind” $ g _{\text{BV}^h(I)}$		
resolution	$\sigma = 16$	$\sigma = 32$	$\sigma = 64$	$\sigma = 16$	$\sigma = 32$	$\sigma = 64$
128×128	1,393	2,358	10,047	1,694	2,574	3,476
256×256	4,525	6,722	12,250	5,460	8,851	12,484
512×512	14,615	22,328	33,115	17,197	30,676	44,289

For the three problems with

$$f = 255\chi_{[\frac{1}{4}, \frac{3}{4}]^2}, \quad (5.4)$$

and λ equal to 3.771636443, 7.820179629, and 16.26268646 (corresponding to $\|f - \tilde{f}\|_{L_2(I)} = 16, 32$, and 64, respectively) we computed numerical approximations to the exact solutions on a grid of size 2048×2048 ; the value of the numerical approximation on the subsquare $\frac{1}{2048}(I + i)$ is taken to be

$$\tilde{f}\left(\frac{1}{2048}\left(i + \left(\frac{1}{2}, \frac{1}{2}\right)\right)\right). \quad (5.5)$$

A simple geometric argument shows that this approximation is a near-best piecewise constant projection in $L_2(I)$ of \tilde{f} onto a $2,048 \times 2,048$ grid (just follow the argument for Example 2 in Section III.E of [19], as the measure of the subgrid square where \tilde{f} is less than the value (5.5) is no less than $1/4$, and no greater than $3/4$, times the measure of the subgrid square).

We compute piecewise constant approximations f_h^n of the minimizers of (2.3) and (3.9) on grids of size 128×128 , 256×256 , and 512×512 , and measure the $L_2(I)$ distance between these approximations f_h^n and the projection onto a 2048×2048 grid of \tilde{f} given above. These differences are reported in Table 5.3.

We also computed approximate solutions f_h^n when f is 255 times the characteristic function of the disk

$$f = 255\chi_{|x - (\frac{1}{2}, \frac{1}{2})| \leq \frac{1}{4}}, \quad (5.6)$$

with $\|f - \tilde{f}\|_{L_2(I)} = 16, 32, 64$. Again we used Dirichlet boundary conditions and the discrete minimization problems (2.3) and (3.9); again we ensured that $\epsilon(p^n) \leq 1/4$, so we know that $\|\tilde{f}_h - f_h^n\|_{L_2(I)} \leq 1/4$. The exact solution \tilde{f} is simply a multiple of the

TABLE 5.3

$L_2(I)$ errors on grids of size 128, 256, and 512, and differences $\sigma = \|f - \tilde{f}\|_{L_2(I)}$ of 16, 32, and 64, with initial data a multiple of the characteristic function of a square (5.4); columns 1–3 are the result of (2.3); columns 4–6 are the result of (3.9); α is the estimated order of convergence, $\|\tilde{f} - \tilde{f}_h\|_{L_2} \approx Ch^\alpha$.

	$L_2(I)$ difference between continuous and discrete solutions					
	“anisotropic” $ g _{BV^h(I)}$			“upwind” $ g _{BV^h(I)}$		
resolution	$\sigma = 16$	$\sigma = 32$	$\sigma = 64$	$\sigma = 16$	$\sigma = 32$	$\sigma = 64$
128×128	1.613	1.889	2.113	1.533	1.813	2.045
256×256	0.962	1.134	1.249	0.900	1.041	1.145
512×512	0.554	0.654	0.733	0.508	0.578	0.639
α	0.772	0.765	0.764	0.796	0.824	0.839

TABLE 5.4

$L_2(I)$ errors on grids of size 128, 256, and 512, and differences $\sigma = \|f - \tilde{f}\|_{L_2(I)}$ of 16, 32, and 64, with initial data a multiple of the characteristic function of a disk (5.6); columns 1–3 are the result of (2.3); columns 4–6 are the result of (3.9); α is the estimated order of convergence, $\|\tilde{f} - \tilde{f}_h\|_{L_2} \approx Ch^\alpha$.

	$L_2(I)$ difference between continuous and discrete solutions					
	“anisotropic” $ g _{BV^h(I)}$			“upwind” $ g _{BV^h(I)}$		
resolution	$\sigma = 16$	$\sigma = 32$	$\sigma = 64$	$\sigma = 16$	$\sigma = 32$	$\sigma = 64$
128×128	10.637	9.223	6.004	9.925	8.312	5.143
256×256	7.929	6.981	4.542	7.061	6.051	3.795
512×512	6.029	5.360	3.495	5.185	4.503	2.852
α	0.410	0.392	0.390	0.468	0.442	0.425

characteristic function of same disk such that $\|f - \tilde{f}\|_{L_2(I)}$ is the correct value. The values of $\|f - \tilde{f}\|_{L_2(I)} = 16, 32$, and 64 correspond to $\lambda = 4.5134516668, 9.02703337$, and 18.05406674 , respectively. We compared the same piecewise constant approximation (5.5) to \tilde{f} on a 2048×2048 grid to piecewise constant approximations \tilde{f}_h^n the discrete solutions of (2.3) and (3.9) on grids of size $128 \times 128, 256 \times 256$, and 512×512 . Here the piecewise constant, discrete initial data \tilde{f}_h is not equal to the true initial data \tilde{f} ; we used the same interpolation method (5.5) to compute our discrete data \tilde{f}_h on grids of size 128×128 , etc. The results are reported in Table 5.4.

With both the characteristic function of the disk and the square as data, the solution is in the Sobolev space $W^{\beta,2}$ for β at most $1/2$, so one might suspect that the maximum possible rate of approximation in $L_2(I)$ by piecewise constants is $O(h^\beta)$, and this is, roughly, what one observes when the data is a multiple of the characteristic function of a disk. When the data is a multiple of the characteristic function of the square, however, the discontinuities in the solution are aligned with the computational grid, and the convergence rate, which is at most one for piecewise constant approximations, is limited by the smoothness of \tilde{f} inside only the subsquare $[\frac{1}{4}, \frac{3}{4}]^2$, and the experimental convergence rate is clearly above $1/2$.

5.3. Qualitative properties of the numerical solutions. The anisotropy of the operator (2.4) was briefly noted in the previous section. The operator (3.7) was offered as a “more isotropic” operator, but it, too, is fundamentally anisotropic—if a function has discontinuities across curves that are not vertical, horizontal, or diagonal lines then indeed

$$\lim_{h \rightarrow 0} |g|_{BV^h(I)} \neq |g|_{BV}$$




FIG. 5.1. The first row contains, on the left, the solution of the discrete BV problem using the anisotropic definition (2.5); and on the right, the solution of the discrete BV problem using the upwind definition (3.4). The initial data is a multiple of the characteristic function of a disk (5.6) projected onto a 128×128 grid. The second row contains the same solutions, but now in “false color.”

even for the “upwind” discrete BV seminorm. Here we briefly give two examples that illustrate the effect of this anisotropy.

As usual, we work only with Dirichlet boundary conditions. We begin with $f = 255 \chi_D$, where D is the disk with center $(\frac{1}{2}, \frac{1}{2})$ and radius $\frac{1}{4}$. Using the iteration (3.15), we then find the approximate minimizer over all discrete g with $\|f - g\|_{L_2^h(I)} \leq 64$ of $\|g\|_{\text{BV}^h(I)}$ for both the anisotropic definition (2.5) and the “upwind” definition (3.4) of the discrete BV seminorm. The exact minimizer of the continuous problem in this case is known [27] to be $(255 - \frac{2\lambda}{r}) \chi_D$, where $r = \frac{1}{4}$ is the radius of the disk and λ (satisfying $\lambda > \frac{1}{r}$) is chosen so that

$$\|255 \chi_D - (255 - \frac{2\lambda}{r}) \chi_D\|_{L_2(I)} = 64.$$




FIG. 5.2. The first row contains, on the left, the solution of the discrete BV problem using the anisotropic definition (2.5); and on the right, the solution of the discrete BV problem using the upwind definition (3.4). The initial data is a multiple of the characteristic function of a square (5.4) projected onto a 128×128 grid. The second row contains the same solutions, but now in “false color.”

We use the error bound $\epsilon(p^n)$ to ensure that the $L_2^h(I)$ errors are no greater than $1/4$. In our experiments we set $h = 1/128$. The results are shown in Figure 5.1 for both the anisotropic discrete BV norm (2.5) and the upwind discrete BV norm (3.4). If one looks closely, one sees that the “northeast” and “southwest” borders of the solution disk in the anisotropic solution are more smoothed than other parts of the border, and the “upwind” solution has generally a sharper border everywhere.

To illustrate this phenomenon more clearly, we include “false color” images of those in Figure 5.1. Here each grayscale was assigned an arbitrary color to show how much the borders are smoothed in each of the solutions. For example, the pixels with a grayscale of 0 were colored with the terra-cotta-like color; the same mapping of greyscales to colors was used in both images.

Some things stick out immediately from the false-color images. First, the grayscale

values of the central plateau of the discrete solutions are different, but their difference of only one greyscale value (113 in the anisotropic image, 112 in the greyscale image) could be explained by numerical error.

Second, the smoothing of the border of the solution truly *is* anisotropic in the “anisotropic” image, and it is smoothed over a distance of about 9 pixels in the northeast and southwest directions, which is significant (and which cannot be due to the discrete error, which as noted above is no more than 0.25 RMS greyscales). The smoothing of the upwind solution is spread over a noticeably smaller distance. The notion that “BV preserves edges”, while true in the continuous setting, clearly needs some qualification in the discrete setting.

There are precisely three places where there is a one-pixel jump from the plateau of the disk to the background color in the anisotropic image—at the right, the bottom, and the “northwest” corner. In the upwind image, there are four such places, at the left and right and top and bottom edges of the plateau.

We now consider the example $f = 255\chi_S$ where S is the square $[\frac{1}{4}, \frac{3}{4}] \times [\frac{1}{4}, \frac{3}{4}]$. As in the last example, we found the approximate minimizer of $\|g\|_{BV}$ satisfying $\|f - g\|_{L_h^2(I)} \leq 64$. Allard gives the exact minimizer of the continuous problem in an appendix to [1].

Both the original and upwind solutions preserve sharp jumps at the sides of the squares (see Figure 5.2), however, we see some significant differences near the corners. As expected, the upwind scheme yields the same type of smoothing at each of the four corners, which seems to match the true solution. On the other hand, the anisotropy of the original scheme is evident in the behavior of the four corners. It is also curious to note that the right and bottom sides of the square are much better preserved than the top and left ones. Furthermore, the truncated top left corner appears to be favoring a sharp jump over what should be a more gradual ramp. Finally, we note the magnified anisotropy of the ramps in the northwest and southeast corners.

To see the effect of the upwind BV algorithm on noisy images we compared the anisotropic and upwind models on a piecewise smooth image with additive Gaussian noise (see Figure 5.3). Here we used Neumann boundary conditions and again an upper bound on the error of $1/4$. The result illustrates the staircasing property of solutions of (1.2) (this is often seen as a drawback, but none-the-less it is the “correct” solution). The northeast quadrant of the original image is flat, and both the anisotropic or upwind methods yield flat solutions in this region. The northwest and southeast quadrants of the image have ramps at the proper orientation so that both approaches yield the same staircased result. However, the southwest corner of the image demonstrates the difference in the two algorithms. The anisotropic method over-smooths at this orientation and thus the staircasing is barely present, whereas the staircasing is at the expected level when using the upwind approach.

Finally, we tested the algorithms on real images. Since real images have both cartoon features (which are well modeled using the total variation norm) as well as textures, we applied this approach to the BV based cartoon + texture decomposition model in [23] which is formulated as follows: Given a nonconstant function f and a number $\lambda > 0$, find functions \tilde{u} (representing the cartoon part of the image) and \tilde{v} (representing the textures) satisfying $\tilde{u} + \tilde{v} = f$ that minimize over all u and v

$$\frac{1}{2}\|v\|_{G(I)} + \lambda|u|_{BV(I)} \quad (5.7)$$

where $G(I) = \{v = \nabla_h \cdot g \mid g \in L_\infty(I \times I)\}$. The authors in [5] proposed a discrete version of (5.7) that is an effective yet simple variant of (2.3). They search for discrete




FIG. 5.3. These images demonstrate the effect of the upwind algorithm on a piecewise smooth, noisy image. Top left: the original image data; top right: the noisy data; bottom left: the reconstruction using (2.5); bottom right: the reconstruction using (3.4).

minimizers $\tilde{u} \in BV(I)$ and $\tilde{v} \in G(I)$ of

$$\frac{1}{2} \|f - u - v\|_{L_2^h(I)} + \lambda |u|_{BV^h(I)} \quad (5.8)$$

where $v \in G(I)$ satisfies

$$\inf_{g_i=(g_i^{(1)}, g_i^{(2)})} \left\{ \|g_i\|_\infty \mid v_i = \nabla_h \cdot g_i \text{ and } |g_i| = \sqrt{(g_i^{(1)})^2 + (g_i^{(2)})^2} \right\} \leq \mu$$

for some appropriately set $\mu > 0$. They show (5.8) can be solved using a straightforward extension of the algorithm proposed in [13] (and used here) by computing u and v using the iterative process:

$$v^{n+1} = \pi_{\mu\overline{K}}(f - u^n), \quad (5.9)$$

$$u^{n+1} = f - v^{n+1} - \pi_{\lambda\overline{K}}(f - v^{n+1}). \quad (5.10)$$



FIG. 5.4. *Left: original image. Middle: reconstruction using (2.5). Right: reconstruction using (3.4). All images are 512×512 .*

\overline{K} is defined as in (2.6) and $\pi_{\lambda\overline{K}}(f)$ is still the limit of $\lambda\nabla_h \cdot \overline{p}$ where \overline{p} is a solution of (2.11)–(2.12) or (2.13)–(2.14) for some appropriately defined discrete gradient and discrete divergence. The choices for these discrete operators in [5] are (2.1) and (2.4).

We tested the algorithm (5.9)–(5.10) using the proposed upwind discrete gradient (3.5) and divergence (3.7) and compared it with the original method. For fine features and diagonal lines, this simple adjustment yields quite noticeable results.

Figures 5.4–5.7 contain the cartoon part, \tilde{u} , of the solutions of (5.8) with real images as initial data. We took $\lambda = 0.0002$ and $\mu = 0.1$; we used Neumann boundary conditions. The stopping criteria were somewhat complicated—we required that

$$\|v^{n+1} - v^n\|_{L_2^h(I)} \leq \min\left(\frac{1}{4}, \frac{1}{8}\|\pi_{\mu\overline{K}}(f - u^n)\|_{L_2^h(I)}\right)$$

(we calculated $\pi_{\mu\overline{K}}(f - u^n)$ to the same accuracy) and

$$\|u^{n+1} - u^n\|_{L_2^h(I)} \leq \min\left(\frac{1}{4}, \frac{1}{8}\|\pi_{\lambda\overline{K}}(f - v^{n+1})\|_{L_2^h(I)}\right)$$

(and again we calculated $\pi_{\lambda\overline{K}}(f - v^{n+1})$ to the same accuracy).

At first glance, the results in Figure 5.4 look quite similar. However, if one zooms in on the fine features in the image (in this case, the bicycle spokes) the difference in the reconstructions of the spokes is evident. Both main spokes in the top images of Figure 5.5 are blurred by (2.5). One can at least see a trace of the very bottom spoke when using (3.4) while it is completely lost when using (2.5). The bottom images show significant blurring when using (2.5) in the main diagonal spoke that spans from the top left to bottom right corner. Furthermore, the low contrast section of the spoke directly above it is completely lost when using (2.5).

Figures 5.6 and 5.7 contain cartoon reconstructions of both the clean and noisy image using (2.5) and (3.4). The original images were 512×512 , and again we took $\lambda = 0.0002$ and $\mu = 0.1$. The model (3.4) better preserves the top string on the boat rigging in the reconstruction of the clean image, and keeps a much more coherent line in the reconstruction of the noisy image. Furthermore, the part of this same string that has the water as its background (thus is more low contrast) is still visible in the reconstruction using (3.4) while it is essentially lost using (2.5).




FIG. 5.5. Close ups of the reconstructions in Figure 5.4. The left column are reconstructions using (2.5) and the right columns are reconstructions using (3.4). These sub-images are 131×131 .

REFERENCES

- [1] W. K. ALLARD, *Total variation regularization for image denoising. I. Geometric theory*, SIAM J. Math. Anal., 39 (2007/08), pp. 1150–1190.
- [2] F. ALTER, V. CASELLES, AND A. CHAMBOLLE, *Evolution of characteristic functions of convex sets in the plane by the minimizing total variation flow*, Interfaces Free Bound., 7 (2005), pp. 29–53.
- [3] B. APPLETON AND H. TALBOT, *Globally optimal geodesic active contours*, J. Math. Imaging Vision, 23 (2005), pp. 67–86.
- [4] J.-F. AUJOL, *Some algorithms for total variation based image restoration*, CMLA Preprint 2008-05.
- [5] J.-F. AUJOL, G. AUBERT, L. BLANC-FÉRAUD, AND A. CHAMBOLLE, *Image decomposition into a bounded variation component and an oscillating component*, J. Math. Imaging Vision, 22 (2005), pp. 71–88.
- [6] G. BELLETTINI, V. CASELLES, AND M. NOVAGA, *The total variation flow in \mathbb{R}^N* , J. Differential Equations, 184 (2002), pp. 475–525.
- [7] C. BOUMAN AND K. SAUER, *Bayesian estimation of transmission tomograms using segmentation based optimization*, IEEE Trans. Nuclear Sci., 39 (1992), pp. 1144–1152.




FIG. 5.6. Top left: original 512×512 image; top right: original 512×512 image with additive Gaussian noise. Bottom left: 131×131 sub-image of original image; bottom right: 131×131 sub-image of noisy image.

- [8] A. BRAIDES, *Γ -convergence for beginners*, vol. 22 of Oxford Lecture Series in Mathematics and its Applications, Oxford University Press, Oxford, 2002.
- [9] X. BRESSON AND T. CHAN, *Fast minimization of the vectorial total variation norm and applications to color image processing*, UCLA CAM Report 07-25, (2007).
- [10] M. BURGER, G. GILBOA, S. OSHER, AND J. XU, *Nonlinear inverse scale space methods*, Commun. Math. Sci., 4 (2006), pp. 179–212.
- [11] V. CASELLES, A. CHAMBOLLE, AND M. NOVAGA, *The discontinuity set of solutions of the TV denoising problem and some extensions*, Multiscale Model. Simul., 6 (2007), pp. 879–894.
- [12] ———, *Regularity for solutions of the total variation denoising problem*, Rev. Mat. Iberoamericana, (2010).
- [13] A. CHAMBOLLE, *An algorithm for total variation minimization and applications*, J. Math. Imaging Vision, 20 (2004), pp. 89–97. Special issue on mathematics and image analysis.
- [14] ———, *Total variation minimization and a class of binary mrf models*, in Energy Minimization Methods in Computer Vision and Pattern Recognition: 5th International Workshop, EMM-CVPR 2005, St. Augustine, FL, USA, November 9-11, 2005: proceedings/ Anand Rangarajan, Baba Vemuri, Alan L. Yuille (eds.), LNCS 3757, Springer-Verlag, 2005, pp. 136–152.
- [15] A. CHAMBOLLE AND P.-L. LIONS, *Image recovery via total variation minimization and related*




FIG. 5.7. Top left: reconstruction of the clean image in Figure 5.6 using (2.5). Top right: reconstruction of the clean image in Figure 5.6 using (3.4). Bottom left: reconstruction of the noisy image in Figure 5.6 using (2.5). Bottom right: reconstruction of the noisy image in Figure 5.6 using (3.4). These sub-images are 131×131 .

- problems, Numer. Math., 76 (1997), pp. 167–188.
- [16] P. L. COMBETTES, *Solving monotone inclusions via compositions of nonexpansive averaged operators*, Optimization, 53 (2004), pp. 475–504.
 - [17] P. L. COMBETTES AND V. R. WAJS, *Signal recovery by proximal forward-backward splitting*, Multiscale Model. Simul., 4 (2005), pp. 1168–1200 (electronic).
 - [18] J. DARBON AND M. SIGELLE, *Exact optimization of discrete constrained total variation minimization problems*, in Combinatorial image analysis, vol. 3322 of Lecture Notes in Comput. Sci., Springer, Berlin, 2004, pp. 548–557.
 - [19] R. A. DEVORE, B. JAWERTH, AND B. J. LUCIER, *Image compression through wavelet transform coding*, IEEE Trans. Inform. Theory, 38 (1992), pp. 719–746.
 - [20] B. EICKE, *Iteration methods for convexly constrained ill-posed problems in Hilbert space*, Numer. Funct. Anal. Optim., 13 (1992), pp. 413–429.
 - [21] I. EKELAND AND R. TEMAM, *Convex analysis and variational problems*, Translated from the French, Studies in Mathematics and its Applications, Vol. 1, North-Holland Publishing Co., Amsterdam, 1976.
 - [22] D. GOLDFARB AND W. YIN, *Second-order cone programming methods for total variation-based*

- image restoration*, SIAM J. Sci. Comput., 27 (2005), pp. 622–645 (electronic).
- [23] Y. MEYER, *Oscillating patterns in image processing and nonlinear evolution equations*, vol. 22 of University Lecture Series, American Mathematical Society, Providence, RI, 2001. The fifteenth Dean Jacqueline B. Lewis memorial lectures.
 - [24] Z. OPIAL, *Weak convergence of the sequence of successive approximations for nonexpansive mappings*, Bull. Amer. Math. Soc., 73 (1967), pp. 591–597.
 - [25] S. OSHER, M. BURGER, D. GOLDFARB, J. XU, AND W. YIN, *An iterative regularization method for total variation-based image restoration*, Multiscale Model. Simul., 4 (2005), pp. 460–489 (electronic).
 - [26] S. OSHER AND J. A. SETHIAN, *Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations*, J. Comput. Phys., 79 (1988), pp. 12–49.
 - [27] W. RING, *Structural properties of solutions to total variation regularization problems*, M2AN Math. Model. Numer. Anal., 34 (2000), pp. 799–810.
 - [28] L. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Physica D, 60 (1992), pp. 259–268.
 - [29] J. WANG AND B. J. LUCIER, *Error bounds for finite difference methods for Rudin-Osher-Fatemi image smoothing*, to appear.