

ASPECTS OF TOTAL VARIATION REGULARIZED L^1 FUNCTION APPROXIMATION*

TONY F. CHAN[†] AND SELIM ESEDOĞLU[†]

Abstract. The total variation–based image denoising model of Rudin, Osher, and Fatemi [*Phys. D*, 60, (1992), pp. 259–268] has been generalized and modified in many ways in the literature; one of these modifications is to use the L^1 -norm as the fidelity term. We study the interesting consequences of this modification, especially from the point of view of geometric properties of its solutions. It turns out to have interesting new implications for data-driven scale selection and multiscale image decomposition.

Key words. total variation, denoising, scale space

AMS subject classifications. 94A08, 65K10

DOI. 10.1137/040604297

1. Introduction. Variational models for image reconstruction have had great success. One of the best known and influential examples is the total variation–based model of Rudin, Osher, and Fatemi (ROF) [22]. This model and its variants have been a very active research topic. The idea behind the model is to exhibit the reconstructed image as the minimizer of the following energy:

$$(1.1) \quad \int_D |\nabla u| + \lambda \int_D (f - u)^2 dx.$$

The functional is to be minimized over all $u \in L^2(D)$. Here D is a domain in \mathbf{R}^N , $N \geq 2$, with Lipschitz boundary; it represents, for example, the computer screen. In this paper, we will work with $D = \mathbf{R}^N$ for convenience. The function $f(x)$ represents the observed and possibly degraded image and is taken to be in $L^2(D)$. The second integral in the functional is the *fidelity* term; it encourages the solution $u(x)$ that is being sought to approximate the *observed image* $f(x)$. The first integral in the functional is the *regularization* term; it is the essential novelty of the ROF model, as it allows for the reconstruction of images with discontinuities across hypersurfaces. Nevertheless, it disfavors oscillations and is responsible for the elimination of noise in applications to noisy images.

The standard ROF model (1.1) is well known to have certain limitations. One important issue is the loss of contrast in solutions even for noise-free observed images. For example, Strong and Chan studied in [25] the case when the observed image $f(x)$ is a disk and showed that the solution to (1.1), for any given λ , is of the form $cf(x)$, where $c \in [0, 1)$ is a constant. We never get $c = 1$, no matter how large the constant λ is chosen. More generally, given any observed image $f(x)$ and $\lambda > (2\|f\|_*)^{-1}$, it can be shown [15] for the corresponding solution $u(x)$ that $\|f - u\|_* = \frac{1}{2\lambda}$. Here, $\|\cdot\|_*$ denotes the dual norm of total variation. (See [15] for definition of the dual norm and

*Received by the editors February 19, 2004; accepted for publication (in revised form) November 19, 2004; published electronically July 26, 2005. This work was supported in part by NSF contract DMS-9973341, NSF contract ACI-0072112, NSF contract DMS-0410085, ONR contract N00014-03-1-0888, and NIH contract P20 MH65166.

<http://www.siam.org/journals/siap/65-5/60429.html>

[†]Mathematics Department, UCLA, Box 951555, Los Angeles, CA 90095 (TonyC@college.ucla.edu, esedoglu@math.ucla.edu).

proofs of the statements just mentioned.) It is in general desirable for image denoising algorithms to have a large class of “noise-free” images that they leave invariant. For the standard ROF model, as these results show, that class consists of only the trivial image $f(x) := 0$.

Recently, work of Meyer inspired research into understanding the role of the fidelity term better. It highlighted the fact that the choice of a suitable fidelity term can have far reaching consequences. For example, following up on Meyer’s ideas, Vese and Osher [27] and then Osher, Sole, and Vese [21] came up with variants of the original model that replace the fidelity term with weaker norms. It is shown in these works that this modification allows for much better separation of the high frequency component of images, such as noise and texture, from the piecewise smooth or “cartoon” part.

In this paper, we ask related but rather different questions. We study a version of the ROF model that uses the L^1 -norm as a measure of fidelity between the observed and denoised images. Given an observed image $f(x) \in L^1(\mathbf{R}^N)$, this model is based on the following variational problem:

$$(1.2) \quad \inf_{u(x) \in BV(\mathbf{R}^N)} \int_{\mathbf{R}^N} |\nabla u| + \lambda \int_{\mathbf{R}^N} |u(x) - f(x)| dx.$$

Our goal in this paper is to explore the consequences of this modest modification on the standard ROF model. In particular, we shall obtain some results that allow us to contrast the modified model (1.2) with the standard one (1.1). Also, the new understanding we develop about the nature of the scale space, lack of uniqueness of solutions, and lack of continuous dependence on data will suggest applications beyond mere removal of noise for the modified model: We will argue that some of these ordinarily undesirable characteristics can be real assets. Indeed, it turns out that the L^1 fidelity-based model has many desirable, and some unexpected, consequences in applications such as multiscale image decomposition and data-driven parameter selection.

Some distinctions between the modified model (1.2) and the standard ROF model (1.1) are immediate:

- The way the fidelity and regularization terms scale with respect to each other in the modified and standard models is different. In particular, unlike the standard model, the modified model is contrast invariant in the following sense: If $u(x)$ is a solution of the modified model for the observed image $f(x)$, then $cu(x)$ is a solution of the modified model for the observed image $cf(x)$.
- The original model is strictly convex, and therefore its solution (the minimizer of the functional) is unique. The modified model is not strictly convex, leading to nonuniqueness of minimizers. This makes the scale space generated by the modified model qualitatively very different—and, as explained in sections 6 and 7, for certain purposes more suitable—than that of the standard ROF model.

We concentrate especially on the scale space and geometric features of the decomposition technique derived from this model. The analytical and numerical results presented in this paper suggest the following major advantages of the L^1 fidelity-based model over the standard one:

- The regularization imposed on solutions by the L^1 model is more geometric. By “more geometric” we mean that the regularization process has less dependence on the contrast of image features than on their shapes. Indeed,

as some of our analytical results show, the L^1 model almost decouples the level sets of the given image from each other and treats them independently of their associated level (grayscale value).

- As distinct from the standard model, small features in the image maintain their contrast even as the fidelity parameter λ is lowered, maintaining good contrast until they suddenly disappear.
- An unexpected consequence of the modification is that it suggests a data-driven scale selection technique: It seems possible to identify certain critical values of the parameter λ at which features at the corresponding scale go through a discontinuous change.

Using the ROF model with L^1 fidelity is a natural idea, and was introduced and studied in the context of image denoising and deblurring by previous authors [1, 3, 4, 16, 17, 18, 8]. Among these, Alliney and Nikolova's works are relevant to ours. Alliney's previous work involves the variational model (1.2) in only one space dimension; moreover, his results are restricted to the discrete versions of the energy. Nevertheless, many of his observations are directly relevant to our results (see, for instance, Proposition 4.2 that we quote from his work), and some of our results (for instance part of Theorem 5.2) can be thought of as continuum analogues of his results in arbitrary dimensions. In [16] Nikolova shows that for certain types of noise the total variation regularization with L^1 fidelity outperforms the standard model. And [17] contains many impressive numerical results that clearly demonstrate the advantages of using the L^1 norm for a fidelity term in some applications. In fact, the analysis presented in [16] applies more generally to fidelity terms that are, like the L^1 fidelity term and unlike the L^2 fidelity term, nondifferentiable at the origin. The techniques of Nikolova also allow her to study certain typical properties of minimizers to the ROF model and its variants with different types of fidelity terms. For example, among the results is a characterization of the staircasing effect. Moreover, she calls attention to the fact that, with L^1 -type fidelity terms, the solution reconstructs the given image exactly at some pixels; this relates to the contrast preserving property we touched on above. However, unlike the focus of this paper, results in [16, 17] mostly concern discrete versions of the denoising energies and depend on the discretization size; continuum analogues are not treated. Our focus in this paper is squarely on the continuum energies so that we can study geometric properties of their minimizers independently of the discretization.

We conclude the introduction with an outline of the remaining sections. Section 2 introduces the notation that is used throughout the paper. Section 3 works out the solution to minimization problems (1.1) and (1.2) in the simple case when the observed image $f(x)$ is the characteristic function of a disk in two dimensions. This illustrates some of the results obtained in subsequent sections for more general types of images. Section 4 consists of a collection of simple but useful facts that follow immediately from the definitions of section 2; these are used in the following sections of the paper. Section 5 deals with properties of minimizers of energy (1.2). In particular, it considers the case where the observed image is the characteristic function of a bounded set. It recalls the known results for the standard ROF model in this case and uses them for comparison. Section 6 elaborates on the differences between the scale spaces generated by the two models given by (1.1) and (1.2); it shows that the model based on L^1 fidelity makes it possible to determine special values of the parameter λ completely from the given observed image. Finally, section 7 presents numerical experiments and gives some implementation details. The numerical results corroborate the overall picture suggested by the analytical results of the previous sections.

2. Notation. In this section we introduce notation that will be used throughout the paper to compare the original ROF model (1.1) with the modified one (1.2) that uses an L^1 fidelity term. First, we recall the standard definitions of total variation of a function and the perimeter of a set [11, 12]. The total variation of a function $u(x) \in L^1_{loc}(\mathbf{R}^N)$ is defined to be

$$\int_{\mathbf{R}^N} |\nabla u(x)| := \sup_{\substack{\phi \in C_c^1(\mathbf{R}^N; \mathbf{R}^N) \\ |\phi(x)| \leq 1 \forall x \in \mathbf{R}^N}} - \int_{\mathbf{R}^N} u(x) \operatorname{div} \phi(x) \, dx.$$

The perimeter of a set $\Sigma \subset \mathbf{R}^N$ is defined in terms of the above definition to be

$$\operatorname{Per}(\Sigma) := \int_{\mathbf{R}^N} |\nabla \mathbf{1}_\Sigma(x)|.$$

For a given possibly noisy image $f(x) \in L^1(\mathbf{R}^N)$, we will denote the energy of the total variation model with L^1 fidelity as $E_1(u, \lambda)$:

$$E_1(u, \lambda) := \int_{\mathbf{R}^N} |\nabla u| + \lambda \int_{\mathbf{R}^N} |f - u| \, dx.$$

It will be compared, for $f \in L^1(\mathbf{R}^N) \cap L^2(\mathbf{R}^N)$, with the energy of the standard ROF model, which we denote by $E_2(u, \lambda)$:

$$E_2(u, \lambda) := \int_{\mathbf{R}^N} |\nabla u| + \lambda \int_{\mathbf{R}^N} (f - u)^2 \, dx.$$

Of particular interest are the minimum values of these energies as a function of the parameter λ :

$$\begin{aligned} \mathcal{E}_1(\lambda) &:= \min_{u \in L^1(\mathbf{R}^N)} E_1(u, \lambda), \\ \mathcal{E}_2(\lambda) &:= \min_{u \in L^2(\mathbf{R}^N)} E_2(u, \lambda). \end{aligned}$$

Minimizers of the standard ROF energy $E_2(\cdot, \lambda)$ for a fixed λ are unique; this is a consequence of the energy's strict convexity. Minimizers of the modified energy $E_1(\cdot, \lambda)$ need not be unique in general. We therefore introduce the following notation to denote the set of minimizers of $E_1(\cdot, \lambda)$ at a given $\lambda \geq 0$:

$$M(\lambda) := \left\{ u \in L^1(\mathbf{R}^N) : E_1(u, \lambda) = \mathcal{E}_1(\lambda) \right\}.$$

For any given $f(x) \in L^1(\mathbf{R}^N)$ and $\lambda \geq 0$, the set $M(\lambda)$ is nonempty: A standard argument shows the existence of minimizers. Because of nonuniqueness, $M(\lambda)$ can have several elements. Different elements of $M(\lambda)$ can stand at different distances from the observed image $f(x)$. This motivates the following notation:

$$\begin{aligned} \mu^+(\lambda) &:= \sup \left\{ \|f - u\|_{L^1(\mathbf{R}^N)} : u \in M(\lambda) \right\}, \\ \mu^-(\lambda) &:= \inf \left\{ \|f - u\|_{L^1(\mathbf{R}^N)} : u \in M(\lambda) \right\}. \end{aligned}$$

The values of the parameter λ at which $M(\lambda)$ contains elements whose distances to the given image $f(x)$ are different turn out to be special. We therefore adopt the following notation to denote this set of special λ values:

$$S(f) := \left\{ \lambda \in \mathbf{R}^+ : \mu^-(\lambda) \neq \mu^+(\lambda) \right\}.$$

To emphasize the dependence of $E_i(\cdot, \lambda)$, $\mathcal{E}_i(\lambda)$, $M(\lambda)$, and $\mu^\pm(\lambda)$ on the observed image $f(x)$ in addition to λ , we will write $E_i(\cdot, \lambda, f)$, $\mathcal{E}_i(\lambda, f)$, $M(\lambda, f)$, and $\mu^\pm(\lambda, f)$ whenever necessary.

3. An example. In this section we consider a very simple but illustrative example. Namely, we work out explicitly the solution to the problem of minimizing the two dimensional version of $E_1(\cdot, \lambda)$ in the case when the observed image $f(x)$ is given by the characteristic function $\mathbf{1}_{B_r(0)}(x)$ of a disk $B_r(0)$ that is centered at the origin and with radius r . It is important to compare the result with the one for the standard ROF model, which—as we noted in the introduction—was calculated in [25].

We start by recalling the calculation of [25]. For $\lambda \geq 0$ and the observed image given by $f(x) = \mathbf{1}_{B_r(0)}(x)$, the unique minimizer $u_\lambda(x)$ of $E_2(\cdot, \lambda)$ is given by

$$u_\lambda(x) \equiv \begin{cases} 0 & \text{if } 0 \leq \lambda \leq \frac{1}{r}, \\ \left(1 - \frac{1}{\lambda r}\right) \mathbf{1}_{B_r(0)}(x) & \text{if } \lambda > \frac{1}{r}. \end{cases}$$

Turning now to the case of $E_1(\cdot, \lambda)$, one can reason (for example with the help of some of the results presented in sections 5 and 6 of this paper) that for each $\lambda \geq 0$ every minimizer has to be of the form $c\mathbf{1}_{B_r(0)}(x)$ for some constant $c \in [0, 1]$. We therefore need to minimize the function

$$E_1(c\mathbf{1}_{B_r(0)}(x), \lambda) = 2\pi r c + \lambda \pi r^2 |1 - c|$$

over $c \in [0, 1]$. We get

$$M(\lambda) = \begin{cases} \{0\} & \text{if } 0 \leq \lambda < \frac{2}{r}, \\ \{c\mathbf{1}_{B_r(0)}(x) : c \in [0, 1]\} & \text{if } \lambda = \frac{2}{r}, \\ \{\mathbf{1}_{B_r(0)}(x)\} & \text{if } \lambda \geq \frac{2}{r}. \end{cases}$$

Thus, we see that the solution is unique for all except one special value of the parameter λ . The special value is related to radius of the disk; for more general images we would expect such special values of the parameter λ to be related to the geometric scale of distinct objects contained in the scene.

The difference between scale spaces generated by the standard ROF model and the one with L^1 fidelity is made abundantly clear by this simple example. When L^1 fidelity is used, unlike in the standard ROF model, the scale space is mostly constant; it only makes a sudden transition at a special value of the scale parameter. This difference can also be manifested by plotting the “fidelity of minimizer” as a function of the parameter λ for each model and comparing the qualitative properties. Figure 1 shows the plots obtained based on the minimizers calculated above.

This example brings out another elementary aspect of using an L^1 fidelity term with total variation regularization. Fix a $\lambda > 0$. Then the unique minimizer of $E_1(\cdot, \lambda)$ with the observed image $f(x) = \mathbf{1}_{B_r(0)}(x)$ is identically 0 if $r < \frac{2}{\lambda}$, but $\mathbf{1}_{B_r(0)}(x)$ if $r > \frac{2}{\lambda}$. Thus the dependence of the solution to the L^1 model on the observed image is not continuous with respect to, say, the L^1 -norm. This is clearly related to the lack of uniqueness in solutions to the model, and is a price to pay for having solutions in




FIG. 1. Left: Plot of $\|u_\lambda(x) - f(x)\|_{L^2}^2$ vs. λ^{-1} for the example of section 3, where $u_\lambda(x)$ denotes the unique minimizer of $E_2(\cdot, \lambda)$. Right: Plot of $\mu^+(\lambda)$ vs. λ^{-1} for the ROF model with L^1 fidelity, using the example of section 3.

which features of interest maintain good contrast until they are completely eliminated. However, sections 6 and 7 explain some applications for which such a discontinuity can actually be desirable, and Proposition 6.4 shows that certain important features of the *scale space* are continuous as a function of observed signal.

4. Basic facts. In this section, we collect a number of elementary facts that follow immediately from the definitions introduced in the previous section. These results will be useful in the subsequent sections.

The following claim shows that the minimum energies $\mathcal{E}_i(\lambda)$ are well-behaved functions of the parameter λ .

CLAIM 1. For any given observed image $f(x) \in L^1(\mathbf{R}^N)$ the function $\mathcal{E}_1(\lambda)$, and for any given observed image $f(x) \in L^2(\mathbf{R}^N)$ the function $\mathcal{E}_2(\lambda)$, satisfy the following properties:

1. $\mathcal{E}_i(\lambda)$ for $i = 1, 2$ are increasing and concave.
2. $\mathcal{E}_i(0) = 0$ for $i = 1, 2$.
3. $0 \leq \mathcal{E}_1(\lambda) \leq \|f\|_{L^1} \lambda$ and $0 \leq \mathcal{E}_2(\lambda) \leq \|f\|_{L^2}^2 \lambda$ for all $\lambda \in [0, \infty)$.
4. $\mathcal{E}_i(\lambda)$ are Lipschitz continuous for $i = 1, 2$.

Proof. $\mathcal{E}_i(\lambda)$ are defined as pointwise infima of a collection of linear functions that are increasing in λ ; this makes them increasing and concave. Statements 2 and 3 follow from the trivial fact that $\mathcal{E}_i(\lambda) \leq E_i(0, \lambda)$ for $i = 1, 2$. Statement 4 now follows from the first three. \square

CLAIM 2. The set $M(\lambda)$ is closed and convex.

Proof. This follows from convexity of the energy E_1 . \square

The following claim, which must be a well-known fact, shows that the fidelity of the minimizer to the original ROF model varies continuously as a function of λ . This should be contrasted with the results for the L^1 model that are obtained in the subsequent sections. We include its proof for completeness.

CLAIM 3. Given $f(x) \in L^2(\mathbf{R}^N)$, for each $\lambda \geq 0$ let $u_\lambda(x)$ denote the unique minimizer of $E_2(\cdot, \lambda)$. Then the function $\lambda \rightarrow \|f - u_\lambda\|_{L^2}$ is continuous.

Proof. Fix $\lambda_* \geq 0$ and let $u_{\lambda_*}(x)$ be the unique minimizer of $E_2(\cdot, \lambda_*)$. Let $\{\lambda_j\}_j^\infty \subset \mathbf{R}^+$ converge to λ_* . Consider the sequence of corresponding minimizers: $\{u_{\lambda_j}\}$. The obvious relation $E_2(u_{\lambda_j}, \lambda_j) \leq E_2(0, \lambda_j) = \lambda_j \|f\|_{L^2}^2$ implies that the sequence has uniformly bounded total variation and L^2 -norm. It also implies that $\|u_\lambda - f\|_{L^2} \leq \|f\|_{L^2}$ for every $\lambda \geq 0$. Applying the standard compactness prop-

erty (for functions with uniformly bounded total variation) on compact sets, we can find a subsequence, also denoted $\{u_{\lambda_j}\}$, such that $u_{\lambda_j}(x) \rightarrow v(x) \in L^1_{loc}(\mathbf{R}^N)$ in L^1 on any bounded set. We may then pass to another subsequence to make sure that $u_{\lambda_j}(x) \rightarrow v(x)$ pointwise a.e. as well. Fatou's lemma then shows that $\|v - f\|_{L^2} \leq \liminf_{j \rightarrow \infty} \|u_{\lambda_j} - f\|_{L^2}$, so that in fact $v \in L^2(\mathbf{R}^N)$. Also, the standard lower semi-continuity result for total variation implies that $\int |\nabla v| \leq \liminf_{j \rightarrow \infty} \int |\nabla u_{\lambda_j}|$. Hence we get that $E_2(v, \lambda_*) \leq \liminf_{j \rightarrow \infty} E_2(u_{\lambda_j}, \lambda_j)$.

On the other hand, $E_2(u_{\lambda_*}, \lambda_*) \geq \limsup_{j \rightarrow \infty} E_2(u_{\lambda_j}, \lambda_j)$. To see this, suppose not. Then there is $\varepsilon > 0$ and arbitrarily large j such that $E_2(u_{\lambda_*}, \lambda_*) \leq E_2(u_{\lambda_j}, \lambda_j) - \varepsilon$. But also, $\lim_{j \rightarrow \infty} E_2(u_{\lambda_*}, \lambda_j) = E_2(u_{\lambda_*}, \lambda_*)$. These two statements mean $E_2(u_{\lambda_*}, \lambda_j) < E_2(u_{\lambda_j}, \lambda_j)$ for some large j , which is a contradiction, since u_{λ_j} are supposed to be minimizers of $E_2(\cdot, \lambda_j)$. This, along with the remarks of the previous paragraph, adds up to the following conclusion:

$$\limsup_{j \rightarrow \infty} E_2(u_{\lambda_j}, \lambda_j) \leq E_2(u_{\lambda_*}, \lambda_*) \leq E_2(v, \lambda_*) \leq \liminf_{j \rightarrow \infty} E_2(u_{\lambda_j}, \lambda_j).$$

We thus see that v is a minimizer of $E_2(\cdot, \lambda_*)$; by uniqueness of minimizers of $E_2(\cdot, \lambda_*)$, we get that $v = u_{\lambda_*}$.

If $\lambda_* = 0$, then $u_{\lambda_*} = 0$ and so $\|u_{\lambda_*} - f\|_{L^2} = \|f\|_{L^2}$. Recalling from above that $\|u_{\lambda} - f\|_{L^2} \leq \|f\|_{L^2}$ for all λ , we see that in this case

$$\limsup_{j \rightarrow \infty} \|u_{\lambda_j} - f\|_{L^2} \leq \|u_{\lambda_*} - f\|_{L^2} \leq \liminf_{j \rightarrow \infty} \|u_{\lambda_j} - f\|_{L^2},$$

which establishes continuity of the map in question at $\lambda = 0$.

If $\lambda_* > 0$, we reason as follows: We must once again have $\limsup_{j \rightarrow \infty} \|u_{\lambda_j} - f\|_{L^2} \leq \|u_{\lambda_*} - f\|_{L^2}$, which immediately leads to the conclusion of the claim. To see this, we suppose that it is false and proceed as we did in the previous paragraphs. There is then arbitrarily large j and an $\varepsilon > 0$ such that $\|u_{\lambda_*} - f\|_{L^2} \leq \|u_{\lambda_j} - f\|_{L^2} - \varepsilon$. But then

$$E_2(u_{\lambda_*}, \lambda_*) \leq \liminf_{j \rightarrow \infty} E_2(u_{\lambda_j}, \lambda_j) - \varepsilon \lambda_j.$$

Also, $E_2(u_{\lambda_*}, \lambda_j) \rightarrow E_2(u_{\lambda_*}, \lambda_*)$ as $j \rightarrow \infty$. These last two statements lead as before to the contradictory statement that $E_2(u_{\lambda_*}, \lambda_j) < E_2(u_{\lambda_j}, \lambda_j)$. \square

We will see whether the analogue of Claim 3 holds for E_1 . In that regard, we first make the following basic observation.

CLAIM 4. *Let $\lambda_2 > \lambda_1 \geq 0$, and assume that u_{λ_1} and u_{λ_2} are any two minimizers of $E_1(\cdot, \lambda_1)$ and $E_1(\cdot, \lambda_2)$, respectively. Then*

$$\|u_{\lambda_1} - f\|_{L^1(\mathbf{R}^N)} \geq \|u_{\lambda_2} - f\|_{L^1(\mathbf{R}^N)}.$$

Proof. Suppose $\|u_{\lambda_2} - f\|_{L^1} > \|u_{\lambda_1} - f\|_{L^1}$. Then, since $u_{\lambda_1} \in M(\lambda_1)$, we have $E_1(u_{\lambda_1}, \lambda_1) \leq E_1(u_{\lambda_2}, \lambda_1)$. We then have

$$\begin{aligned} E_1(u_{\lambda_1}, \lambda_2) &= E_1(u_{\lambda_1}, \lambda_1) + (\lambda_2 - \lambda_1)\|u_{\lambda_1} - f\|_{L^1} \\ &\leq E_1(u_{\lambda_2}, \lambda_1) + (\lambda_2 - \lambda_1)\|u_{\lambda_1} - f\|_{L^1} \\ &< E_1(u_{\lambda_2}, \lambda_1) + (\lambda_2 - \lambda_1)\|u_{\lambda_2} - f\|_{L^1} \\ &= E_1(u_{\lambda_2}, \lambda_2), \end{aligned}$$

which is a contradiction, since $u_{\lambda_2} \in M(\lambda_2)$ by hypothesis. \square

COROLLARY 4.1. *The functions $\mu^\pm(\lambda)$ are decreasing. In fact,*

$$\mu^-(\lambda_1) \leq \mu^+(\lambda_1) \leq \mu^-(\lambda_2) \leq \mu^+(\lambda_2)$$

whenever $\lambda_1 > \lambda_2 \geq 0$.

The functions $\mu^\pm(\lambda)$ are the analogue for E_1 of $\|u_\lambda - f\|_{L^2}$ in Claim 3. These functions in general can be discontinuous; in fact, their set of discontinuity is precisely $S(f)$ according to our notation. The corollary above allows us to make the following simple statement about the discontinuities of these functions.

CLAIM 5. *For any given $f \in L^1(\mathbf{R}^N)$, the set $S(f)$ is at most countable.*

Proof. If $\lambda \in S(f)$, then $\mu^-(\lambda) < \mu^+(\lambda)$. By the corollary above, at such a λ both μ^- and μ^+ have a jump discontinuity. The set of discontinuities of a monotone function are at most countable. \square

Finally, for completeness let us state the following rather obvious fact about the asymptotic value of the functions $\mu^\pm(\lambda)$ as $\lambda \rightarrow \infty$.

CLAIM 6. *Given $f(x) \in L^1(\mathbf{R}^N)$, we have $\lim_{\lambda \rightarrow \infty} \mu^\pm(\lambda) = 0$.*

Proof. Given $\varepsilon > 0$, we can find $f_\varepsilon(x) \in BV(\mathbf{R}^N)$ such that $\|f_\varepsilon - f\|_{L^1} \leq \frac{\varepsilon}{2}$. If $u_\lambda(x) \in M(\lambda)$ with $\mu^+(\lambda) = \|u_\lambda - f\|_{L^1}$, then

$$\mu^-(\lambda) \leq \mu^+(\lambda) \leq \frac{1}{\lambda} E_1(u_\lambda, \lambda) \leq \frac{1}{\lambda} E_1(f_\varepsilon, \lambda) \leq \frac{1}{\lambda} \int |\nabla f_\varepsilon| + \frac{\varepsilon}{2}.$$

Hence, for all large enough λ we have $\mu^\pm(\lambda) \leq \varepsilon$. \square

The following fact is taken directly from [3]. It says that any image $u_*(x)$ which arises as the solution to model (1.2) for some observed image $f(x)$ is in fact also the solution to model (1.2) with observed image $f(x)$ taken to be $u_*(x)$ itself, provided that the parameter λ is taken large enough. We include it as a good way to emphasize the difference of model (1.2) from (1.1) in regard to the loss of contrast in solutions.

PROPOSITION 4.2. *Let $\lambda_* \geq 0$, $f(x) \in L^1(\mathbf{R}^N)$, and $u_*(x) \in M(\lambda_*, f)$. Then for every $\lambda \geq \lambda_*$ we have $u_*(x) \in M(\lambda, u_*)$.*

Proof. For the proof of this claim, see [3].

5. Minimizers of E_1 . In this section, we study the behavior of the ROF model with L^1 fidelity on simple images. Our motivation is twofold. First, studying the behavior of image denoising models on simple images is a first step towards understanding the type of images they can successfully process. Second, this type of question allows us to compare different models. In fact, we will stress the difference of these results from the analogous ones obtained for the standard ROF model by previous authors. In particular, our results will bolster the intuitive observation that the L^1 fidelity term leads to more geometric regularizations.

The following proposition constitutes our starting point. It shows that the ROF model with L^1 fidelity term almost decouples the level sets of the given image from each other; it almost becomes a geometry problem for each level set, independent of the level. This idea of writing total variation-based optimization problems in terms of level sets appears previously in the works [23, 24] of Strang, and is used to show the existence of binary solutions, as we do in Theorem 5.2.

PROPOSITION 5.1. *The energy $E_1(u, \lambda)$ can be rewritten as follows:*

$$(5.1) \quad E_1(u, \lambda) = \int_{-\infty}^{\infty} \text{Per}(\{x : u(x) > \gamma\}) + \lambda |\{x : u(x) > \gamma\} \Delta \{x : f(x) > \gamma\}| \, d\gamma.$$

Proof. Recall the coarea formula for functions of bounded variation (see [12] or [11]):

$$(5.2) \quad \int_{\mathbf{R}^N} |\nabla u| = \int_{-\infty}^{\infty} \text{Per}(\{x : u(x) > \gamma\}) d\gamma.$$

Also, there is the following “layer cake” formula:

$$\begin{aligned} \int_{\mathbf{R}^N} |u - f| dx &= \int_{\{u>f\}} |u - f| dx + \int_{\{f>u\}} |u - f| dx \\ &= \int_{\{u>f\}} \int_{f(x)}^{u(x)} d\gamma dx + \int_{\{f>u\}} \int_{u(x)}^{f(x)} d\gamma dx \\ &= \int_{\mathbf{R}^N} \int_{\mathbf{R}} \mathbf{1}_{\{u>f\}}(x) \mathbf{1}_{[f(x),u(x))}(\gamma) + \mathbf{1}_{\{f>u\}}(x) \mathbf{1}_{[u(x),f(x))}(\gamma) d\gamma dx \\ &= \int_{\mathbf{R}} \int_{\mathbf{R}^N} \mathbf{1}_{\{u>f\}}(x) \mathbf{1}_{[f(x),u(x))}(\gamma) + \mathbf{1}_{\{f>u\}}(x) \mathbf{1}_{[u(x),f(x))}(\gamma) dx d\gamma, \end{aligned}$$

where we simply changed the order of integration in the last step. But now we have

$$\mathbf{1}_{\{u>f\}}(x) \mathbf{1}_{[f(x),u(x))}(\gamma) = 1 \quad \text{iff } x \in \{u > f\} \cap \{u > \gamma\} \cap \{f > \gamma\}^c$$

and 0 otherwise, and

$$\mathbf{1}_{\{f>u\}}(x) \mathbf{1}_{[u(x),f(x))}(\gamma) = 1 \quad \text{iff } x \in \{f > u\} \cap \{u > \gamma\}^c \cap \{f > \gamma\}$$

and 0 otherwise. That means

$$\mathbf{1}_{\{u>f\}}(x) \mathbf{1}_{[f(x),u(x))}(\gamma) + \mathbf{1}_{\{f>u\}}(x) \mathbf{1}_{[u(x),f(x))}(\gamma) = \mathbf{1}_{\{u>\gamma\} \Delta \{f>\gamma\}}(x).$$

Therefore

$$\int_{\mathbf{R}^N} |u - f| dx = \int_{-\infty}^{\infty} |\{x : u(x) > \gamma\} \Delta \{x : f(x) > \gamma\}| d\gamma.$$

Putting these formulas together gives the one in the statement of the claim. \square

We now explore some consequences of Proposition 5.1. First, we consider what happens when the observed image is binary. In other words, we assume that $f(x)$ is the characteristic function of a domain. We assume that the domain is bounded, but for now make no assumptions about the boundary of the domain.

THEOREM 5.2. *If the observed image $f(x)$ is the characteristic function of a bounded domain $\Omega \subset \mathbf{R}^N$, then for any $\lambda \geq 0$ there is a minimizer of $E_1(\cdot, \lambda)$ that is also the characteristic function of a (possibly different) domain. In other words, when the observed image is binary, then for each $\lambda \geq 0$ there is at least one $u(x) \in M(\lambda)$ which is also binary.*

In fact, if $u_\lambda(x) \in M(\lambda)$ is any minimizer of $E_1(\cdot, \lambda)$, then for almost every $\gamma \in [0, 1]$ we have that the binary function

$$\mathbf{1}_{\{x:u_\lambda>\gamma\}}(x)$$

is also a minimizer of $E_1(\cdot, \lambda)$.

Proof. Let $f(x) := \mathbf{1}_\Omega(x)$, where Ω is a bounded domain in \mathbf{R}^N . It can be easily seen that any minimizer $u(x)$ of E_1 satisfies $u(x) \in [0, 1]$ for almost every $x \in \mathbf{R}^N$. Formula (5.1) of Proposition 5.1 above becomes, in this case,

$$E_1(u, \lambda) = \int_0^1 \text{Per}(\{x : u(x) > \gamma\}) + \lambda |\{x : u(x) > \gamma\} \Delta \Omega| \, d\gamma.$$

This suggests that we consider for each level set of $u(x)$ the following geometry problem:

$$(5.3) \quad \min_{\Sigma \subset \mathbf{R}^N} \left(\text{Per}(\Sigma) + \lambda |\Sigma \Delta \Omega| \right).$$

Standard compactness and lower semicontinuity facts show the existence of minimizers; let $\Sigma_* \subset \mathbf{R}^N$ be one of them. Let $u_\lambda(x)$ be any minimizer of $E_1(\cdot, \lambda)$, i.e., $u_\lambda(x) \in M(\lambda)$. Set

$$\Sigma(\gamma) := \{x : u_\lambda(x) > \gamma\}.$$

Then

$$(5.4) \quad \text{Per}(\Sigma(\gamma)) + \lambda |\Sigma(\gamma) \Delta \Omega| \geq \text{Per}(\Sigma_*) + \lambda |\Sigma_* \Delta \Omega|$$

for almost every $\gamma \in [0, \infty)$. This now immediately implies that

$$E_1(u_\lambda(x), \lambda) \geq E_1(\mathbf{1}_{\Sigma_*}(x), \lambda),$$

which means that $\mathbf{1}_{\Sigma_*}(x)$ is also a minimizer of $E(\cdot, \lambda)$.

Furthermore, since $u_\lambda(x)$ is a minimizer, the inequality of (5.4) is in fact an equality for almost every $\gamma \in [0, 1]$. Thus, $\Sigma(\gamma)$ is a minimizer of the geometry problem (5.3), and $\mathbf{1}_{\Sigma(\gamma)}(x)$ is a minimizer of $E_1(\cdot, \lambda)$ for almost every γ . \square

Remark. A version of the first statement of Theorem 5.2 was obtained for the discrete analogue of model (1.2) in one space dimension by Alliney in [4]. \square

Remark. The claim leaves open the possibility that for a given $\lambda \geq 0$ there might be a $u \in M(\lambda)$ that takes more than two values.

Remark. The conclusion of Theorem 5.2 is interesting because it establishes the equivalence of a nonconvex problem (the geometry problem of minimizing over only binary images, which is encountered in many applications such as improving the appearance of fax documents) to a convex problem (minimizing over all images). Indeed, it follows from the corollary that to obtain a solution to (5.3), one can first minimize $E_1(\cdot, \lambda)$, taking $f(x) = \mathbf{1}_\Omega(x)$ as the observed image, and then look at a level set of the solution obtained. The resulting algorithm would be very different from the standard level set method of Osher and Sethian [19, 20]. Whether this observation can be turned into a useful computational tool needs to be explored, but this question will not be pursued any further here.

Theorem 5.2 highlights an important *qualitative difference* of the L^1 model from the standard ROF model. In contrast to the content of these claims, it is easy to show that for certain types of binary images (even with smooth edge sets) the minimizer of the standard ROF model takes more than two values for every large enough choice of the parameter λ .

We do not know whether the following comparison principle holds for the geometry problem (5.3): If $\Omega_1 \subset \Omega_2$ and Σ_1, Σ_2 are minimizers of (5.3) with $\Omega = \Omega_1$ and

$\Omega = \Omega_2$, respectively, then do we necessarily have $\Sigma_1 \subset \Sigma_2$? If true, this would imply, in particular, uniqueness for solutions of (5.3). In any case, we can make the following statement.

COROLLARY 5.3. *If the observed image $f(x)$ is the characteristic function of a bounded convex domain $\Omega \subset \mathbf{R}^N$, then for almost every $\lambda \geq 0$ the minimizer of $E_1(\cdot, \lambda)$ is unique and is the characteristic function of a set contained in Ω .*

Proof. Let $\lambda \in [0, \infty) \setminus S(f)$, and let $u_\lambda(x) \in M(\lambda)$. We recall from the proof of Theorem 5.2 that, using the same notation as in that proof, the set $\Sigma(\gamma)$ minimizes the geometry problem (5.3) for almost every $\gamma \in [0, 1]$. Let $1 \geq \gamma_1 > \gamma_2 \geq 0$, and assume that $\Sigma(\gamma_1) \neq \Sigma(\gamma_2)$ both minimize the geometry problem. By definition, we have $\Sigma(\gamma_1) \subset \Sigma(\gamma_2)$. Furthermore, convexity of Ω implies that

$$\text{Per}(\Sigma(\gamma_i) \cap \Omega) \leq \text{Per}(\Sigma(\gamma_i)) \text{ for } i = 1, 2.$$

Since $\mathbf{1}_{\Sigma(\gamma_1)}(x)$ and $\mathbf{1}_{\Sigma(\gamma_2)}(x)$ are minimizers, it follows that $\Sigma(\gamma_1) \subset \Sigma(\gamma_2) \subseteq \Omega$. Hence, $|\Sigma(\gamma_1) \Delta \Omega| \neq |\Sigma(\gamma_2) \Delta \Omega|$. But then $\lambda \in S(f)$, which is a contradiction. We have thus reached the conclusion that if $\lambda \in [0, \infty) \setminus S(f)$, then any minimizer of $E_1(\cdot, \lambda)$ is necessarily binary (i.e., the characteristic function of a set). Now suppose that $u_1(x)$ and $u_2(x)$ are two binary minimizers of $E_1(\cdot, \lambda)$. By convexity of $E_1(\cdot, \lambda)$, we then have that $\frac{1}{2}(u_1(x) + u_2(x))$ is also a minimizer, and thus binary. But the average of two binary functions is binary only if the two functions are identical.

Thus, whenever $\lambda \in [0, \infty) \setminus S(f)$, the minimizer of $E_1(\cdot, \lambda)$ is unique and is binary: It is of the form $\mathbf{1}_\Sigma(x)$ for some set Σ . The argument above shows that $\Sigma \subseteq \Omega$. And Claim 5 says that $S(f)$ is at most countable and thus negligible. That proves the claim. \square

As an aside, we note the following result about problem (5.3) that follows immediately from the previous corollary (perhaps it can be obtained also in a less roundabout way).

COROLLARY 5.4. *Let Ω be a bounded convex domain in \mathbf{R}^N . Then, for almost every $\lambda \geq 0$, the solution of problem (5.3) is unique.*

Proof. If Σ_1 and Σ_2 are solutions to (5.3), then $\mathbf{1}_{\Sigma_1}(x)$ and $\mathbf{1}_{\Sigma_2}(x)$ are minimizers of $E_1(\cdot, \lambda)$ with the observed image given by $f(x) = \mathbf{1}_\Omega(x)$. Conditions on Ω imply that Corollary 5.3 applies so that $\Sigma_1 = \Sigma_2$. That proves the claim. \square

We will next consider some simple images $f(x)$ for which the minimizer of $E_1(\cdot, \lambda)$ turns out to be precisely the image $f(x)$ itself for every large enough λ . In section 1, we recalled a result from Meyer’s lecture notes [15] which says that for the standard ROF model given by $E_2(\cdot, \lambda)$ the only such image is $f(x) := 0$. For E_1 , however, there are many such images, as shown by Proposition 4.2, which we quoted in section 4 from [3]. The following lemma will be instrumental in establishing whether certain simple observed images $f(x)$ have this property.

LEMMA 5.5. *Given an observed image $f(x) \in BV(\mathbf{R}^N)$, assume that there is a vector field $\phi(x)$ with the following properties:*

1. $\phi(x) \in C_c^1(\mathbf{R}^N; \mathbf{R}^N)$,
2. $|\phi(x)| \leq 1$ for all $x \in \mathbf{R}^N$,
3. $\int_{\mathbf{R}^N} f(x) \text{div } \phi(x) \, dx = \int_{\mathbf{R}^N} |\nabla f|$.

Then there exists a threshold $\lambda_ \geq 0$ such that $M(\lambda) = \{f(x)\}$ for all $\lambda > \lambda_*$. In other words, the unique minimizer of $E_1(\cdot, \lambda)$ is given by the observed image $f(x)$.*

Proof. Set $\lambda_* := \max_{x \in \mathbf{R}^N} |\text{div } \phi(x)|$. Take any $\lambda > \lambda_*$. Then, given any

$u(x) \in BV(\mathbf{R}^N)$, we have

$$\begin{aligned} E_1(u, \lambda) &= \int |\nabla u| + \lambda \int |u - f| dx \\ &\geq \int u \operatorname{div} \phi dx + \lambda \int |u - f| dx \\ &= \int f \operatorname{div} \phi dx + \lambda \int |u - f| dx + \int (u - f) \operatorname{div} \phi dx \\ &\geq E_1(f, \lambda) + \left(\lambda - \max_{x \in \mathbf{R}^N} |\operatorname{div} \phi(x)| \right) \int |u - f| dx. \end{aligned}$$

Since $\lambda > \lambda_* := \max |\operatorname{div} \phi(x)|$, the last inequality shows that $E_1(u, \lambda) > E_1(f, \lambda)$ unless $u \equiv f$. Since u is a minimizer, it must in fact be the case that $u \equiv f$. \square

Lemma 5.5 can now be applied, for example, to binary images to obtain an important class of exact solutions. This requires making some smoothness assumption about the interface between the two values of the binary function.

THEOREM 5.6. *Let $\Omega \subset \mathbf{R}^N$ be a bounded domain with C^2 boundary. Let the observed image $f(x)$ be given by $f(x) = \mathbf{1}_\Omega(x)$. Then there exists a threshold $\lambda_* \geq 0$ such that whenever $\lambda > \lambda_*$, the unique minimizer of $E_1(\cdot, \lambda)$ is the observed image $f(x) = \mathbf{1}_\Omega(x)$ itself.*

Proof. Since the boundary $\partial\Omega$ of the bounded domain Ω is assumed to be C^2 , the outward unit normal vector field $n(x) : \partial\Omega \rightarrow \mathbf{S}^{N-1}$ of $\partial\Omega$ can be extended in a C^1 manner to a tubular neighborhood of $\partial\Omega$, so that one gets a vector field $\phi(x) \in C_c^1(\mathbf{R}^N; \mathbf{R}^N)$ such that $\phi(x)|_{x \in \partial\Omega} = n(x)$, and $|\phi(x)| \leq 1$ for all $x \in \mathbf{R}^N$. But then

$$\begin{aligned} \int_{\mathbf{R}^N} f \operatorname{div} \phi dx &= \int_{\Omega} \operatorname{div} \phi(x) dx = \int_{\partial\Omega} \phi(x) \cdot n(x) d\sigma \\ &= \operatorname{Per}(\partial\Omega) = \int_{\mathbf{R}^N} |\nabla f| dx. \end{aligned}$$

Hence, the vector field $\phi(x)$ satisfies all the requirements of Lemma 5.5, from which the conclusion of the present claim follows. \square

At this point it is worth recalling the behavior of the standard ROF model on binary images of the form $f(x) = \mathbf{1}_\Omega(x)$. As we noted above, simple considerations show that the minimizer of the standard ROF model almost never turns out to be $u(x) = f(x) = \mathbf{1}_\Omega(x)$. A related question is whether the solution $u(x)$ has at least the correct “set of edges”; see [10]. In case Ω is a ball, one can calculate the minimizer explicitly [25]; it turns out to be $u(x) = c\mathbf{1}_\Omega(x)$, where $c = 1 - \frac{\operatorname{Per}(\Omega)}{2\lambda|\Omega|}$. In particular, $u(x)$ has the same set of edges as $f(x)$. The results of [5] generalize the results of [25] but also show that the class of binary images that have this weaker property (i.e., images for which the solution to the standard ROF model turns out to be a constant multiple of the observed image) is still rather limited; for example, there are smooth but nonconvex shapes that lack this property.

Remark. Theorem 5.6 can easily be extended to images of a more general form. Indeed, if the level sets $\{x : f(x) = \gamma\}$ of the given image $f(x)$ are smooth and vary smoothly with respect to γ , the same conclusion holds. We also see, among other things, that such an image $f(x)$ cannot have strict local extrema, for at a strict local extremum the level sets shrink to a point. Moreover, there are also binary images that lack this property (i.e., which are not exactly recovered for any $\lambda \geq 0$, no matter how

large). In fact, a repetition of some of the arguments of Meyer given in his lecture notes [15] on the standard ROF model show that the characteristic function of, say, a square cannot arise as the solution to the ROF model with L^1 fidelity either, no matter what the observed image $f(x) \in L^1$ is, and no matter how large the parameter λ is chosen.

Remark. A discrete version of Theorem 5.6 is proved in [17] for denoising models with nonsmooth (including L^1) fidelity terms and smooth regularization terms. In those results, unlike ours, the threshold value for the parameter λ necessarily involves the grid size.

The last few claims dealt with the behavior of the L^1 fidelity-based model for large values of the parameter λ . Next, we consider what happens when $\lambda \geq 0$ is small enough. The following claim is a very simple application of the isoperimetric inequality.

PROPOSITION 5.7. *Let $R > 0$. Then there exists a threshold $\lambda_* = \lambda_*(R, N)$ such that if $f \in L^1(\mathbf{R}^N)$ with $\text{supp}(f) \subset B_R(0)$, then $M(\lambda) = \{0\}$ for any $\lambda < \lambda_*$. In other words, the unique minimizer of $E_1(\cdot, \lambda)$ is given by $u(x) \equiv 0$.*

Proof. Let $C = C(N)$ be the isoperimetric constant

$$\int_{\mathbf{R}^N} |\nabla u| \geq C(N) \|u\|_{L^{\frac{N}{N-1}}(\mathbf{R}^N)} \quad \text{for all } u \in BV(\mathbf{R}^N).$$

Then we set

$$\lambda_*(R, N) := \frac{C(N)}{R\omega_N^{\frac{1}{N}}},$$

where ω_N is the volume of the unit ball in \mathbf{R}^N . Take a $\lambda > \lambda_*$ and let $u(x) \in M(\lambda)$. Then $E_1(u, \lambda) \leq E_1(0, \lambda)$. By the isoperimetric inequality, that means

$$C(N) \|u\|_{L^{\frac{N}{N-1}}(\mathbf{R}^N)} + \lambda \|u - f\|_{L^1(\mathbf{R}^N)} \leq \lambda \|f\|_{L^1(\mathbf{R}^N)} = \lambda \|f\|_{L^1(B_R(0))}.$$

We apply Holder's inequality to the first term on the left-hand side after splitting it into integrations over $B_R(0)$ and $B_R^c(0)$. That gives

$$\frac{C(N)}{R\omega_N^{\frac{1}{N}}} \|u\|_{L^1(B_R(0))} + \lambda \|u - f\|_{L^1(B_R(0))} + C(N) \|u\|_{L^{\frac{N}{N-1}}(B_R^c(0))} \leq \lambda \|f\|_{L^1(B_R(0))},$$

which shows that if $\lambda < C(N)/R\omega_N^{\frac{1}{N}} = \lambda_*$, then

$$\|u\|_{L^1(B_R(0))} = \|u\|_{L^{\frac{N}{N-1}}(B_R^c(0))} = 0.$$

In other words, $u \equiv 0$. □

Remark. This behavior of the L^1 model is to be expected, based on its contrast invariance, as we have already noted in the introduction. It differs from the behavior at small λ values of the standard ROF model, which, according to [15], entails not just the support of a given compactly supported image $f(x)$ but its $\|\cdot\|_*$ -norm.

6. Scale space and the set $S(f)$. The set $S(f)$ of discontinuities of the functions μ^\pm play a distinguished role in the scale space generated by varying the parameter λ in the L^1 model. As the value of λ is gradually decreased, minimizers of the image models become coarser as small scale objects in the image merge to form larger

scale structures. Intuitively, for the L^1 model we can expect the values of $\lambda \in S(f)$ to correspond to scales of distinct objects that make up the image. These are the values of λ at which the scale space makes a rapid and drastic transition.

We would first like to prove that the set $S(f)$ is nonempty for the kind of images we have been considering in the previous sections, namely images of the form $f(x) = \mathbf{1}_\Omega(x)$, where Ω is a bounded domain. Our arguments are based on verifying this claim for the special case where the given image is the characteristic function of a ball, and then generalizing the result to $f(x) = \mathbf{1}_\Omega(x)$ by comparing Ω with a ball that is contained in Ω .

LEMMA 6.1. *Let Ω be a bounded domain in \mathbf{R}^2 , and assume that $B_R(p) \subset \Omega$. Consider the observed image given by $f(x) = \mathbf{1}_\Omega(x)$. Then for any $\lambda \geq 0$ and $r \in (0, R)$ we have*

$$E_1(\mathbf{1}_{B_r(p)}(x), \lambda) > \min \left\{ E_1(0, \lambda), E_1(\mathbf{1}_{B_R(p)}(x), \lambda) \right\}.$$

Proof. Since $B_r(p) \subset B_R(p) \subset \Omega$ for each $r \in (0, R)$, we have

$$\|\mathbf{1}_\Omega(x) - \mathbf{1}_{B_r(p)}(x)\|_{L^1(\mathbf{R}^2)} = |\Omega| - \pi r^2.$$

That means

$$E_1(\mathbf{1}_{B_r(p)}(x), \lambda) = \lambda(|\Omega| - \pi r^2) + 2\pi r.$$

Considering $E_1(\mathbf{1}_{B_r(p)}(x), \lambda)$ as a function of r , we see that it achieves its minimum on $[0, R]$ strictly at the end points of the interval. \square

In order to show that $\mu^\pm(\lambda)$ is a discontinuous function, we will show that its range omits a full interval of values but does include certain values on either side of that interval. The next claim exhibits such an omitted interval.

LEMMA 6.2. *Let Ω be a bounded domain in \mathbf{R}^2 , and let $B_R(0) \subset \Omega$. Consider the observed image given by $f(x) = \mathbf{1}_\Omega(x)$. There is no $\lambda \in \mathbf{R}^+$ such that*

$$|\Omega| - \pi R^2 < \mu^+(\lambda) < |\Omega|.$$

Proof. Suppose there is a $\lambda \geq 0$ such that $|\Omega| - \pi R^2 < \mu^+(\lambda) < |\Omega|$. There exists $u(x)$ such that $u(x) \in M(\lambda)$ and $\|u - f\|_{L^1(\mathbf{R}^N)} = \mu^+(\lambda)$. As before, let $\Sigma(\gamma) := \{x : u(x) > \gamma\}$. By Proposition 5.1, we have $\mathbf{1}_{\Sigma(\gamma)}(x) \in M(\lambda)$ for almost every $\gamma \in (0, 1)$. Therefore, for almost every γ we have

$$\|\mathbf{1}_{\Sigma(\gamma)}(x) - f\|_{L^1(\mathbf{R}^2)} < |\Omega|$$

(otherwise $\mu^+(\lambda) \geq |\Omega|$). It also cannot be the case that $|\Sigma(\gamma) \Delta \Omega| \leq |\Omega| - \pi R^2$ for almost every $\gamma \in (0, 1)$ since we know that

$$\int_0^1 |\Sigma(\gamma) \Delta \Omega| d\gamma = \|u - f\|_{L^1(\mathbf{R}^2)} = \mu^+(\lambda) > |\Omega| - \pi R^2.$$

Thus, there exists $\gamma_* \in (0, 1)$ such that

$$\mathbf{1}_{\Sigma(\gamma_*)}(x) \in M(\lambda) \quad \text{and} \quad |\Omega| - \pi R^2 < \|\mathbf{1}_{\Sigma(\gamma_*)}(x) - f(x)\|_{L^1(\mathbf{R}^2)} < |\Omega|.$$

Case 1. $|\Sigma(\gamma_*)| \geq \pi R^2$. But then $\text{Per}(B_R(0)) \leq \text{Per}(\Sigma(\gamma_*))$, and

$$|\Omega \Delta B_R(0)| = |\Omega| - \pi R^2 < \|\mathbf{1}_{\Sigma(\gamma_*)}(x) - f(x)\|_{L^1}.$$

Hence, $E_1(\mathbf{1}_{B_R(0)}(x), \lambda) < E_1(\mathbf{1}_{\Sigma(\gamma_*)}(x), \lambda)$. This is a contradiction, since $\mathbf{1}_{\Sigma(\gamma_*)}(x)$ was supposed to be a minimizer.

Case 2. $|\Sigma(\gamma_*)| < \pi R^2$. In this case, take $r = \frac{1}{\sqrt{\pi}}|\Sigma(\gamma_*)|^{\frac{1}{2}}$. Since $r \in (0, R)$, we have that $B_r(0) \subset \Omega$. This implies

$$\|\mathbf{1}_{B_r(0)}(x) - f(x)\|_{L^1(\mathbf{R}^2)} \leq \|\mathbf{1}_{\Sigma(\gamma_*)}(x) - f(x)\|_{L^1(\mathbf{R}^2)}.$$

Moreover, as before, $\text{Per}(B_R(0)) \leq \text{Per}(\Sigma(\gamma_*))$. Therefore,

$$E_1(\mathbf{1}_{B_r(0)}(x), \lambda) \leq E_1(\mathbf{1}_{\Sigma(\gamma_*)}(x), \lambda) = E_1(u(x), \lambda).$$

On the other hand, by Lemma 6.1 we have

$$E_1(\mathbf{1}_{B_r(0)}(x), \lambda) > \min \left\{ E_1(0, \lambda), E_1(\mathbf{1}_{B_R(0)}(x), \lambda) \right\}.$$

This is a contradiction, since $u(x) \in M(\lambda)$.

THEOREM 6.3. *Let Ω be a nonempty bounded domain in \mathbf{R}^2 . Consider the observed image given by $f(x) = \mathbf{1}_\Omega(x)$. Then the functions $\mu^\pm(\lambda)$ are discontinuous.*

Proof. By Proposition 5.7, we have that $\mu^+(\lambda) = \|f\|_{L^1} = |\Omega|$ for all small enough λ . On the other hand, by Claim 6 we have that $\mu^\pm(\lambda) \rightarrow 0$ as $\lambda \rightarrow \infty$. However, by Lemma 6.2, there is a range of values near $|\Omega|$ that the function μ^+ cannot take. It therefore has to be discontinuous. Discontinuity of μ^- follows from that of μ^+ via Claim 4. \square

Remark. This should be contrasted with the situation for the standard total variation model (with L^2 fidelity), which is explained in Claim 3.

We thus see that the scale spaces generated by the two models, the standard ROF model and the one with L^1 fidelity, are very different. With the standard ROF model, pronounced objects of distinct scale with sharp edges in the image gradually lose their contrast and merge with their neighbors as the parameter λ is lowered. With the L^1 model, such objects maintain their contrast with respect to their neighbors—however, their boundaries might be gradually smoothed out. This goes on until a critical value of λ is reached—one that belongs to the set $S(f)$, at which point the object suddenly merges with a neighboring one.

At this point, it is also worth comparing the scale space generated by the L^1 model with that generated by anisotropic diffusion via motion by mean curvature of level sets. The two are drastically different. This can be seen most easily in the case when $f(x)$ is the characteristic function of a disk. The scale space generated by motion by curvature consists of a family of concentric disks shrinking gradually to a point. Hence the same feature, i.e., the original disk, appears at many intermediate scales, albeit in different sizes. On the other hand, the scale space generated by the total variation model with L^1 fidelity term consists of either the original disk or the constant background at any given scale.

Finally, we return to the topic of continuous dependence on the observed signal for the L^1 model. Despite our remarks in section 3, we show in the next claim that the fidelity of minimizer versus λ graph depends on the observed image continuously.

PROPOSITION 6.4. *Let $\{f_j(x)\}_{j=1}^\infty$ be a sequence in $L^1(\mathbf{R}^N)$ that converges to $f(x)$ in the L^1 -norm. Then, for almost all $\lambda \geq 0$, $\mu^\pm(\lambda, f_j)$ converges to $\mu^\pm(\lambda, f)$ as $j \rightarrow \infty$.*

Proof. Let $\mathcal{S} := S(f) \cup (\cup_{j=1}^\infty S(f_j))$. According to Claim 5, $S(f)$ and each $S(f_j)$ are countable. Therefore, \mathcal{S} is countable and thus negligible. Fix $\lambda \in [0, \infty) \setminus \mathcal{S}$. For

each j , take $u_j \in M(\lambda, f_j)$. The sequence $\{u_j\}_{j=1}^\infty$ is bounded in total variation norm and hence is precompact in L^1 on compact sets. Passing to a subsequence if necessary, we may assume that $u_j \rightarrow u_\infty$ pointwise a.e. as $j \rightarrow \infty$.

We must have $u_\infty \in M(\lambda, f)$. To see this, assume otherwise. $M(\lambda, f)$ is nonempty, so take a $u \in M(\lambda, f)$. By lower semicontinuity we have

$$E_1(u, \lambda, f) < E_1(u_\infty, \lambda, f) \leq \liminf_{j \rightarrow \infty} E_1(u_j, \lambda, f_j).$$

However, $E_1(u, \lambda, f_j) \rightarrow E_1(u, \lambda, f)$ as $j \rightarrow \infty$. Therefore, for large enough j , we get $E_1(u, \lambda, f_j) < E_1(u_j, \lambda, f_j)$. This gives a contradiction, since $u_j \in M(\lambda, f_j)$.

Now that we know $u_\infty \in M(\lambda, f)$, recall next that $\lambda \notin \mathcal{S}$. Therefore,

$$\mu^\pm(\lambda, f) = \|u_\infty - f\|_{L^1} = \lim_{j \rightarrow \infty} \|u_j - f_j\|_{L^1} = \lim_{j \rightarrow \infty} \mu^\pm(\lambda, f_j).$$

That proves the claim. □

7. Computation. In this section, we show numerical examples that bring out unique features of the total variation-based denoising model with L^1 fidelity term. We also give some details on the numerical schemes used to obtain these results.

Our computations are based on gradient descent schemes for decreasing the energies involved. The nondifferentiability of the terms involved in the energies calls for some sort of regularization. The regularized versions of energies $E_1(\cdot, \lambda)$ and $E_2(\cdot, \lambda)$ used in our numerical experiments are the following:

$$\begin{aligned} E_1^{\varepsilon, \delta}(u, \lambda) &:= \int_{\mathbf{R}^N} \sqrt{|\nabla u|^2 + \varepsilon} + \lambda \int_{\mathbf{R}^N} \sqrt{(f - u)^2 + \delta} \, dx, \\ E_2^\varepsilon(u, \lambda) &:= \int_{\mathbf{R}^N} \sqrt{|\nabla u|^2 + \varepsilon} + \lambda \int_{\mathbf{R}^N} (f - u)^2 \, dx. \end{aligned}$$

This type of approximation to total variation-based models is very standard. The discrete versions of these energies lead to the following equally standard explicit gradient descent schemes in two space dimensions:

$$\begin{aligned} \frac{u_{i,j}^{n+1} - u_{i,j}^n}{\delta t} &= D_x^- \left(\frac{D_x^+ u_{i,j}^n}{\sqrt{(D_x^+ u_{i,j}^n)^2 + (D_y^+ u_{i,j}^n)^2 + \varepsilon}} \right) \\ &\quad + D_y^- \left(\frac{D_y^+ u_{i,j}^n}{\sqrt{(D_x^+ u_{i,j}^n)^2 + (D_y^+ u_{i,j}^n)^2 + \varepsilon}} \right) + \lambda \frac{(f - u_{i,j}^n)}{((f - u_{i,j}^n)^2 + \delta)^\alpha}, \end{aligned}$$

where $\alpha = \frac{1}{2}$ for $E_1^{\varepsilon, \delta}$ and $\alpha = 0$ for E_2^ε . Here, D^+ and D^- denote forward and backward difference quotients, respectively, in the direction of their subscript.

We note that efficient numerical minimization of energies considered in this work is a topic unto itself; no doubt there are better ways to do it than the gradient descent approach taken and the specific choice of scheme made above. In particular, it is better to use algorithms that do not need to regularize the nondifferentiable terms appearing in the energy. Such an algorithm is presented by Alliney in [2] with applications to one dimensional signals in the context of an objective functional with mixed l^1, l^2 norms. Also, Chambolle recently developed an efficient algorithm for minimizing the standard ROF model for images without regularizing the total variation term




FIG. 2. Example of scale space generated by the standard total variation model. Compare with the same example for the model with L^1 fidelity, shown in Figure 3.

[7]. Further alternative numerical approaches to total variation–based models can be found in [6, 9]. Whether these algorithms can be adapted to our setting is a very interesting question that will be explored elsewhere.

An important point that we need to clarify is the following. Although, as we already noted several times, the energy $E_1(\cdot, \lambda)$ is not strictly convex and its minimizers in general lack uniqueness, for any given $\delta > 0$ the approximate energy $E_1^{\varepsilon, \delta}(\cdot, \lambda)$ is strictly convex so that its minimizers enjoy uniqueness. It is these minimizers that we have computed. Moreover, it is a very routine matter to verify that a sequence of minimizers of $E_1^{\varepsilon, \delta}(\cdot, \lambda)$ converges to the set of minimizers $M(\lambda)$ of $E_1(\cdot, \lambda)$ as $\varepsilon, \delta \rightarrow 0^+$. The analogous convergence statement is, of course, true also for a sequence of minimizers of $E_2^{\varepsilon}(\cdot, \lambda)$.

Figures 2 and 3 compare the scale spaces generated by the standard total variation model and the one with L^1 fidelity on a synthetic image. This experiment makes the more geometric nature of the L^1 model abundantly clear. The observed image consists of squares of various sizes and gray levels. In the scale space generated by the standard total variation model, the squares gradually lose their contrast (while at the same time their geometries get regularized) and gradually disappear. Moreover, some large squares with low contrast against the background—namely the square near the upper right corner—disappear before some smaller squares that have higher contrast against the background—namely the two intermediate sized squares along the diagonal. On the other hand, in the scale space generated by the model with L^1 fidelity, the squares get processed only in terms of their geometry: They preserve their contrast very well until all of a sudden they disappear. (They should, in fact, preserve their contrast perfectly, but because our numerical scheme regularizes the L^1 fidelity term to make it differentiable, in practice there is some loss of contrast.) In principle, the contrast of the squares plays no role in determining the order in which they are removed; that order is determined completely in terms of the geometry of the features.

Figure 4 shows the graph of the fidelity of the minimizer versus λ for the standard




FIG. 3. Example of scale space generated by the total variation model with L^1 fidelity. Compare with Figure 2.




FIG. 4. Plot of the fidelity of minimizer (i.e., $\|u_\lambda(x) - f(x)\|_{L^2}^2$) versus λ^{-1} for the standard ROF model (top graph) and of the fidelity of minimizer (i.e., $\|u_\lambda(x) - f(x)\|_{L^1}$) versus λ^{-1} for the ROF model with L^1 fidelity (bottom graph).

total variation model and for the model with L^1 fidelity. An important ambiguity that we need to resolve is how the nonuniqueness of minimizers of $E_1(\cdot, \lambda)$ affects the fidelity-versus- λ plot for $E_1(\cdot, \lambda)$. To answer this question, recall that the fidelity of various minimizers of $E_1(\cdot, \lambda)$ differs from each other at only countably many values of λ . In particular, all ways of obtaining the second graph in Figure 4 yield plots that are identical up to a set of measure 0. Hence, there is no ambiguity in the results shown.

Discontinuities in the minimizer’s fidelity-versus- λ graph for the L^1 model correspond to distinguished values of the parameter λ . As can be seen from the results, these are the values of λ at which a drastic change in the scale space takes place. Namely, at such values of λ one of the “features” (squares in this example) gets eliminated. There is no such distinguished value of λ in the plot for the standard ROF model at which the graph becomes discontinuous (as shown both by our theoretical results and by the numerical example shown). However, the graph in that case might have kinks, which are of course harder to detect than discontinuities. Thus, unlike

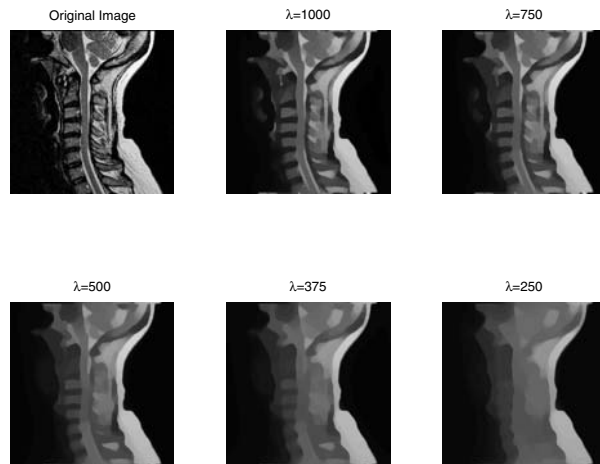


FIG. 5. Scale space generated by the standard ROF model.

the standard total variation model, the model with L^1 fidelity thus suggests a method for *data-driven parameter selection*.

For those familiar with the notion of an L-curve [13, 14] (which is a technique for choosing regularization parameters in ill-posed inverse problems), let us point out that from the point of view of this paper there is no apparent useful connection between the fidelity-versus- λ graph and the L-curve. According to the L-curve method, to determine a distinguished value of the regularization parameter λ , one should find the corner (point of maximum curvature) in the $\int |\nabla u_\lambda|$ versus $\|u_\lambda - f\|_{L^1}$ graph. However, for instance in the case of the example of section 3 (i.e., with $f(x) = \mathbf{1}_{B_r(0)}(x)$), the curvature of this graph is easily seen to be independent of the radius r ; thus, the L-curve method does not yield any scale information.

The special values of parameter λ obtained from the fidelity-of-minimizer graph via the L^1 model can be used in many ways. For example, denoising models are sometimes used for generating multiscale decomposition of images, as in [26]. In such applications, it is necessary to select a *schedule* for the parameter λ a priori. In [26], this schedule is chosen in the form $\lambda = 2^j \lambda_0$, with $j = 1, 2, 3, \dots$, and the initial value λ_0 is arbitrarily chosen by the user. The L^1 scale space suggests a more natural data-driven way to select these parameters using the discontinuities in the fidelity-of-minimizers graph. Moreover, even if one opts to use a λ -schedule of the form used in [26], the theoretical results and preliminary numerical examples of this paper suggest that one might obtain a much cleaner decomposition using the ROF model with L^1 fidelity in place of the standard ROF model. All these ideas pertaining to multiscale decomposition of images using the L^1 fidelity-based model will be explored elsewhere.

Finally, Figures 5 and 6 illustrate the differences between the standard ROF model and the one with L^1 fidelity on a real medical image. In this example also, one can see that the small scale features in the observed image, such as these indicated by the arrow on the lower-left-hand-side image of Figure 6, maintain their contrast much better in the L^1 fidelity model than in the standard ROF model, even as the parameter λ is gradually decreased to very low values.

8. Conclusion. We have considered the total variation-based image denoising model of Rudin, Osher, and Fatemi with the L^1 -norm as the fidelity term. Our

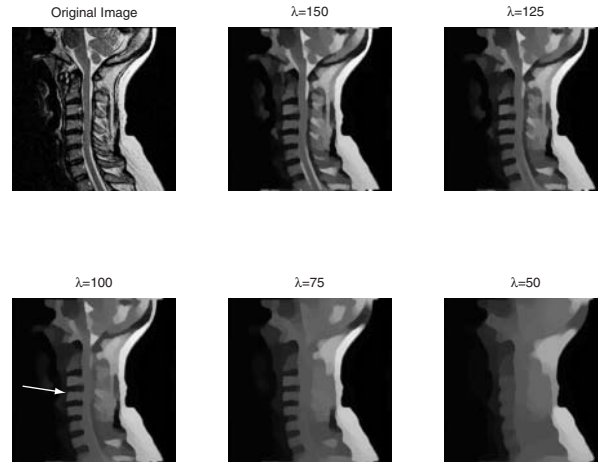


FIG. 6. Scale space generated by the ROF model with L^1 fidelity term.

results highlight that this modification leads to many interesting qualitative differences in the behavior of the modified model from the standard one. These differences have important consequences for image denoising. They also suggest interesting new research directions into applications to data-driven parameter selection and multiscale image decomposition.

Acknowledgments. The authors would like to thank Antonin Chambolle, Mila Nikolova, Stanley Osher, and Luminita Vese for helpful discussions.

REFERENCES

- [1] S. ALLINEY, *Digital filters as absolute norm regularizers*, IEEE Trans. Signal Process., 40 (1992), pp. 1548–1562.
- [2] S. ALLINEY AND S. A. RUZINSKY, *An algorithm for the minimization of mixed l_1 and l_2 norms with applications to Bayesian estimation*, IEEE Trans. Signal Process., 42 (1994), pp. 618–627.
- [3] S. ALLINEY, *Recursive median filters of increasing order: A variational approach*, IEEE Trans. Signal Process., 44 (1996), pp. 1346–1354.
- [4] S. ALLINEY, *A property of the minimum vectors of a regularizing functional defined by means of the absolute norm*, IEEE Trans. Signal Process., 45 (1997), pp. 913–917.
- [5] G. BELLETTINI, V. CASELLES, AND M. NOVAGA, *Total variation flow in \mathbf{R}^N* , J. Differential Equations, 184 (2002), pp. 475–525.
- [6] F. CATTÉ, F. DIBOS, AND G. KOEPFLER, *A morphological scheme for mean curvature motion and applications to anisotropic diffusion and motion of level sets*, SIAM J. Numer. Anal., 32 (1995), pp. 1895–1909.
- [7] A. CHAMBOLLE, *An algorithm for total variation minimization and applications*, J. Math. Imaging and Vision, 20 (2004), pp. 89–97.
- [8] E. CHEON AND A. PARANJPYE, *Noise Removal Project by Total Variation Minimization*, Math 199 Project Report, S. Osher and L. Vese, advisors, UCLA Mathematics Department, Los Angeles, 2002, <http://www.math.ucla.edu/~lvese/MATH199/index.html>.
- [9] F. DIBOS AND G. KOEPFLER, *Global total variation minimization*, SIAM J. Numer. Anal., 37 (2000), pp. 646–664.
- [10] D. C. DOBSON AND F. SANTOSA, *Recovery of blocky images from noisy and blurred data*, SIAM J. Appl. Math., 56 (1996), pp. 1181–1198.
- [11] L. C. EVANS AND R. F. GARIEPY, *Measure Theory and Fine Properties of Functions*, Stud. Adv. Math., CRC Press, Boca Raton, FL, 1992.
- [12] E. GIUSTI, *Minimal Surfaces and Functions of Bounded Variation*, Monogr. Math. 80, Birkhäuser-Verlag, Basel, 1984.

- [13] P. C. HANSEN, *Analysis of discrete ill-posed problems by means of the L-curve*, SIAM Rev., 34 (1992), pp. 561–580.
- [14] P. C. HANSEN AND D. P. O’LEARY, *The use of the L-curve in the regularization of discrete ill-posed problems*, SIAM J. Sci. Comput., 14 (1993), pp. 1487–1503.
- [15] Y. MEYER, *Oscillating patterns in image processing and nonlinear evolution equations*, AMS University Lecture Series 22, AMS, Providence, RI, 2002.
- [16] M. NIKOLOVA, *Minimizers of cost-functions involving nonsmooth data-fidelity terms. Application to the processing of outliers*, SIAM J. Numer. Anal., 40 (2002), pp. 965–994.
- [17] M. NIKOLOVA, *A variational approach to remove outliers and impulse noise*, J. Math. Imaging and Vision, 20 (2004), pp. 99–120.
- [18] M. NIKOLOVA, *Weakly constrained minimization. Application to the estimation of images and signals involving constant regions*, J. Math. Imaging and Vision, 21 (2004), pp. 155–175.
- [19] S. OSHER AND R. FEDKIW, *Level Set Methods and Dynamic Implicit Surfaces*, Appl. Math. Sci. 153, Springer-Verlag, New York, 2003.
- [20] J. SETHIAN AND S. OSHER, *Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton–Jacobi formulations*, J. Comput. Phys., 79 (1988), pp. 12–49.
- [21] S. OSHER, A. SOLÉ, AND L. VESE, *Image decomposition and restoration using total variation minimization and the H^{-1} norm*, Multiscale Model. Simul., 1 (2003), pp. 349–370.
- [22] L. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Phys. D, 60 (1992), pp. 259–268.
- [23] G. STRANG, *L^1 and L^∞ approximation of vector fields in the plane*, in Nonlinear Partial Differential Equations in Applied Science (Tokyo, 1982), North–Holland Math. Stud. 81, North–Holland, Amsterdam, 1983, pp. 273–288.
- [24] G. STRANG, *Maximal flow through a domain*, Math. Programming, 26 (1983), pp. 123–143.
- [25] D. STRONG AND T. F. CHAN, *Edge-preserving and scale-dependent properties of total variation regularization*, Inverse Problems, 19 (2003), pp. S165–S187.
- [26] E. TADMOR, S. NEZZAR, AND L. VESE, *A multiscale image representation using hierarchical (BV, L^2) decompositions*, Multiscale Model. Simul., 2 (2004), pp. 554–579.
- [27] L. VESE AND S. OSHER, *Modeling textures with total variation minimization and oscillating patterns in image processing*, J. Sci. Comput., 19 (2003), pp. 553–572.