A Wavelet-Laplace Variational Technique for Image Deconvolution and Inpainting

Julia A. Dobrosotskaya † and Andrea L. Bertozzi †

***Abstract*—We construct a new variational method for blind deconvolution of images and inpainting, motivated by recent PDE-based techniques involving the Ginzburg-Landau functional, but using more localized wavelet-based methods. We present results for both binary and grayscale images. Comparable speeds are achieved with better sharpness of edges in the reconstruction.**

***Index Terms*—Inpainting, wavelet diffusion, binary image, barcode.**

I. PDEs and wavelets: combining advantages

A. Introduction

This paper presents a general framework for the recovery of piecewise constant signals and images using a combination of variational methods and wavelet analysis. The variational formulation of the problem allows us to build the properties of the recovered signal directly into the analytical machinery. The efficient wavelet representation allows us to capture and preserve sharp features in the signal while it evolves in accordance with the variational laws. We consider two applications of the introduced method: bar code deconvolution and image inpainting. Both lie within the same broad research area of signal recovery from incomplete and corrupted data.

Wavelet analysis has proven invaluable for dealing with a wide class of signals and images with spatially localized features [13]. Wavelets are designed to capture most of the signal energy using a few wavelet coefficients. This property has been leveraged into powerful, spatially adaptive, signal estimation algorithms that are based on simply shrinking the wavelet coefficients of the noisy signal [16]. Discrete (fast) wavelet transforms allow for fast implementation of linear wavelet methods. At the same time, methods utilizing non-linear operations in the wavelet domain accomplish tasks which are not possible for traditional linear/Fourier approaches to such problems as wavelet denoising, linear inverse problems, non-linear wavelet packet approximation and non-linear multiresolution [5]. Overall, the ability to capture sharp discontinuities within relatively sparse data makes wavelets convenient framework for discrete image analysis.

A parallel set of methods are variational, based on nonlinear PDEs. The seminal methods of Rudin-Osher-Fatemi [20] and Perona-Malik [18] for image denoising led to a vast body of work on variational techniques for image understanding [3], [11]. Similarities between image topology issues and phase transitions in material science and fluid mechanics motivated the integration of diffuse interface models into image processing [6], [23]. These techniques are generally based on the phase-field model of Modica and Mortola [14] and allow for topology transition between two states of a physical system. In general, non-linear diffusion models allow the inclusion of a priori knowledge to ensure regularity while preserving important features including edges [15]. Anisotropic diffusion filters provide the advantage of combining image regularization with edge enhancement properties [18],[17]. As a result, diffusion-based methods have been extensively used in image analysis. The diffuse-interface methods appear both as self-contained techniques [6], [1] and as tools for approximating the total variation norm [20].

Wavelet methods appeared in a variational context in [22], [2]. The former discusses the TV(total variation) minimization in the wavelet domain, that successfully reproduces lost coefficients. The latter describes how some wavelet-based image processing algorithms actually produce exact or approximate minimizers of variational problems.

The key idea in this paper combines the basic geometric framework of diffuse interface methods with the non-locality of wavelets. Wavelet based functionals are inherently multiscale and take advantage of simultaneous space and frequency localization. The main points of the algorithm are:

- *Eliminate blurry edges*: The edge sharpness is not restricted by the diffuse interface scale as in the case of PDE methods([1],[8],[7]). Our approach is to take the Ginzburg - Landau energy functional, but replace differential operators with pseudo-differential ones that have eigen-functions that better suit the problem. Wavelet basis functions are natural candidates due to their localized structure and ability to represent images efficiently.
- *Preserve the usage of fast solvers*: one advantage of PDE variational techniques is the usage of FFT. Respectively, in the wavelet-based technique one can employ the discrete wavelet transform to speed up the method.
- *Design a technology that does not require PDE solving* (thus making the technique accessible for non-math specialists). Wavelet techniques are common in signal processing making them a good candidate for implementation by developers unfamiliar with PDEs. There exists a well-developed theory for wavelet decomposition.
- *Create a technique that can be easily extended from binary images to greyscale.*

Our idea is to take a relatively simple PDE model and adapt it to a wavelet basis. The TV seminorm was proven to be a natural and efficient measure of image regularity. To avoid computational challenges related to equations for minimizers of this norm, we reformulate the problem using the phase-field method and approximate the TV functional (in the Γ sense). The Ginzburg-Landau functional,

$$|u|_{GL} = \frac{\epsilon}{2}\int |\nabla u(x)|^2 dx + \frac{1}{4\epsilon}\int W(u)dx, \qquad (1)$$
$$W(u) = (u^2-1)^2.$$

is a diffuse interface approximation to the Total Variational functional $\int |\nabla u| dx$ in the case of binary images. Several efficient algorithms for deconvolution and image inpainting have been proposed, using this functional as the primary regularizer for the solution of an ill-posed problem. In this paper we replace $\int |\nabla u(x)|^2 dx$ with a wavelet-based semi-norm, with the goal of removing the 'fuzzy' diffuse interface features associated with the methods described above. By characterizing signal regularity in terms of the decay of

† UCLA Mathematics Department, Box 951555, Los Angeles, CA 90095-1555 USA (email: juliadob@math.ucla.edu; bertozzi@math.ucla.edu)

July 10, 2007. This work was supported by the Department of Defense and NSF grant ACI-0321917.

wavelet coefficients via a Besov semi-norm (see, for instance, [13]), we are able to construct a method with similar properties to the PDE-based methods but without ϵ-scale blur.

The key new idea of our method is to construct an operator, called the "wavelet Laplacian" which has wavelet basis functions as eigenfunctions, and acts on them in the same "scale - proportional" manner as the Laplace operator does on the Fourier basis. Such a construction diagonalizes the operator in the wavelet basis. While the Laplace equation

$$\Delta u(x) = -4\pi^2 \sum_k k^2 < u(x), e^{-2\pi i k x} > e^{2\pi i k x} = 0$$

can be treated as the Euler-Lagrange equation for the minimizers of

$$\int |\nabla u(x)|^2 dx = 4\pi^2 \sum_{k=1}^{\infty} k^2 | < u(x), e^{-2\pi i k x} > |^2,$$

let us define the "wavelet Laplacian" in a similar context. Given an orthonormal wavelet basis $\{\psi_{j,k}\}$ define for any function $f \in L_2(R)$

$$\Delta_w u := -\sum_{j=1}^{\infty} 2^{2j} \langle u, \psi_{j,k} \rangle \psi_{j\,k}.$$

Then the equation $u_t = \Delta_w u$ is a "gradient descent equation" corresponding to the problem of minimizing the following "weighted energy" functional:

$$E_w(u) := 1/2 \sum_{j=1}^{\infty} |2^j \langle u, \psi_{j,k} \rangle|^2.$$

This expression is a square of the Besov 1-2-2 semi-norm if the wavelet ψ is smooth. For most examples in the paper we choose the compactly supported, piecewise constant Haar wavelet to generate E_w, due to its resemblance to a sharp interface jump. We will denote $\|u\|_w^2 = E_w$ regardless the choice of basis in all formulas. In the end of the paper we show an example with Daubechies filter 4.

II. Blind Deconvolution of Bar Code Signals

A. Prior results on the deconvolution of bar-codes

The idea is inspired by Esedoglu's variational algorithm [8] of recovering the original bar-code by the output of a scanner with (partially) unknown parameters. The process of scanning is modeled as a convolution of an uknown binary (black and white) barcode signal, with some Gaussian kernel:

$$T_{\alpha,\sigma} : u(x) \to \alpha \cdot G_\sigma * u(x), \qquad \alpha > 0,$$

$$G_\sigma(x) = \frac{1}{2\sigma\sqrt{2\pi}} exp(-\frac{x^2}{2\sigma^2}).$$

The amplitude α (distance to the scanned surface) can be treated either as estimated a priori, or unknown; the standard deviation σ (width of the scanning window) is always treated as unknown, as this parameter is much harder to estimate in practice. The known data, i.e. the scanner output, is assumed to be of the form $f = T_{\alpha,\sigma}(u) +$ unknown noise . To deconvolve, the author [8] uses a variational approach by minimizing the energy functional

$$E_{\epsilon,\lambda}(u,\sigma) = |u|_{GL} + \lambda \int (\alpha \cdot G_\sigma * u - f)^2 dx.$$

The first two terms of the energy functional in (II-A) represent a diffuse-interface approximation (in the Γ sense) for the widely used TV(total variation) norm [9]. The algorithm proceeds by computing a minimum of the energy using gradient flow methods (applied to all unknown parameters and the minimizer itself).

B. Blind deconvolution using the wavelet operator

As discussed in section I-A, we revise (II-A) by exchanging the Fourier basis for the Haar wavelet basis:

$$E_{\epsilon,\mu} = \epsilon \|u\|_{1,2,2}^2 + \frac{1}{4\epsilon} \int W(u) dx + \mu \int (\alpha \cdot G_\sigma * u - f)^2 dx.$$

The gradient descent equation for the minimizer u is

$$u_t = 2\epsilon \Delta_w u - \frac{1}{\epsilon} W'(u) - 2\mu\alpha G_\sigma * (\alpha G_\sigma * u - f).$$

We discretize in time using the gradient stable convexity splitting scheme as in [1]:

$$E = (E_{11} - E_{12}) - (E_{21} - E_{22})$$

$$E_{11} = \frac{\epsilon}{2}\|u\|_w^2 + \frac{c_1}{2}\|u\|_{L_2}^2, E_{12} = \frac{c_1}{2}\|u\|_{L_2}^2 - \frac{1}{4\epsilon}\int W(u)dx$$

$$E_{21} = \frac{c_2}{2}\|u\|_{L_2}^2 E_{22} = \frac{c_2}{2}\|u\|_{L_2}^2 - \mu\|(G_\sigma * u^n - f)\|_{L_2}^2$$

Here c_1 and c_2 are chosen so that $(E_{11} - E_{12})$ and $(E_{21} - E_{22})$ are indeed convex. The numerical scheme follows:

$$u^{n+1} - 2(\delta t)\epsilon \Delta_w u^{n+1} + (c_1 + c_2)u^{n+1} =$$

$$u^n - (\delta t)(\frac{1}{\epsilon} W'(u^n) - 2\mu\alpha G_\sigma * f$$

$$-\mu\alpha^2 G_\sigma * G_\sigma * u^n + (c_1 + c_2)u^n)$$

Section III explains the choice of the splitting constants (see Lemma in III-C). The left hand side of the relation is actually a diagonal matrix applied to the vector of wavelet coefficients. Thus u^{n+1} can be easily computed by first explicitly computing the right hand side and then inverting the operator on the left hand side in the wavelet domain.

C. Examples of deconvolved bar-codes

The deconvolution problem was solved numerically using the described scheme. We compare the new method with the algorithm in [8]. Results obtained for the same set of scanning parameters show some practical advantages of the wavelet algorithm: assuming the amplitude is known, our method recovers a sharper barcode. To compare how well both methods perform in pure signal deconvolution (apart from dealing with the unknown parameters), we first consider the barcode recovery from scanning data with known α and σ in the presence of noise - Fig. 1.

We consider the same sample barcode as in [8]. The unknown parameter σ is successfully determined for the same range of values as in [8]. Namely, the blind recovery of the barcode sample used in [8] is successful for $\sigma \leq 0.014$ assuming one uses the same setting: the barcode is supported on the interval $[0, 1]$, the width of the smallest "bar element" is several times less than the "scanning width" σ.

The difference with the original experiment [8] lies in the way one treats the scanned signal: instead of taking the full number of scanned pixels (1024), we split the signal into parts which we recover separately. This simplification helps the algorithm to withstand much higher noise levels (SNR equal to 5-7 dB). However, the method is equally successful when applied to the full length signal for larger SNR . Another advantage of the splitting approach is the ability to determine σ faster by deconvolving only a piece of the signal, and to recover subsequent signal parts with known σ.

In the one-dimensional case we use the standard fast discrete wavelet transform. Our algorithm performance is 0.0125 sec per iteration. Recovering the barcode of length 256 with unknown $\sigma \in (0, 0.0014)$ requires about 250 iterations to determine σ and less than 100 iterations to refine the resulting barcode. In the example

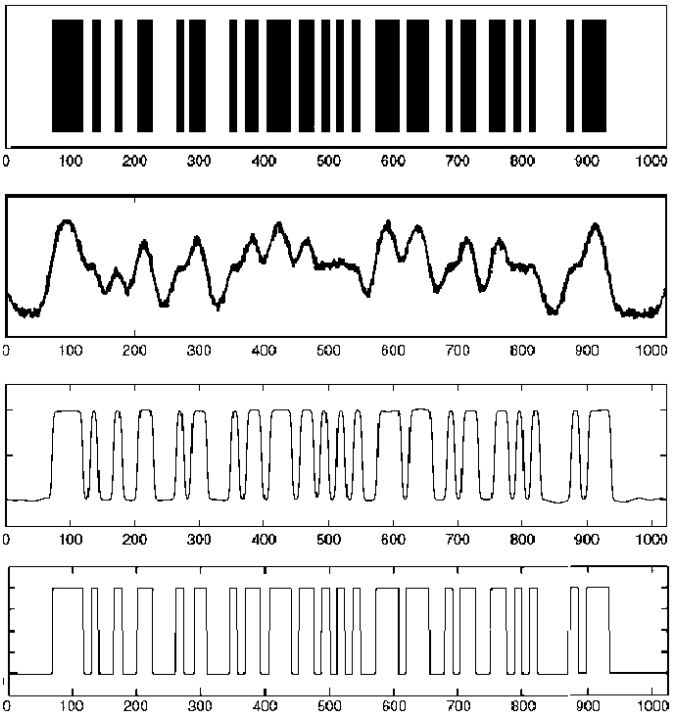

Fig. 1. Comparison of the PDE based and wavelet-based deconvolution (going down): (i) the original barcode, (ii) scanner output, (iii) PDE deconvolution, (iv) wavelet deconvolution.

below (Fig. 2) $\sigma = .012$ was recovered as $\sigma^* = 0.0121$ in the presence of noise with SNR(dB) = 5.9. The reconstructed bar code matches the original one up to the maximum edge shift of 1 pixel.

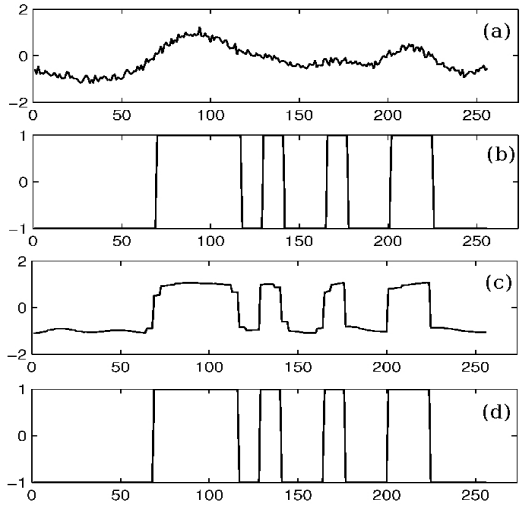

Fig. 2. Enlarged part of a sample barcode:(a) scanner output; (b) (one fourth of) the original (1024-point) barcode; (c) intermediate stage - σ is recovered; (d)final result: the barcode recovered by the wavelet-based blind deconvolution algorithm in the presence of noise with mean square SNR = 5.9 dB.

The process of the bar-code recovery is split into two stages. The first stage starts with a large fidelity coefficient ($\mu = O(10^3 - 10^4)$), the evolution is continued until the parameter σ stabilizes at some value with sufficient precision - see Fig.2 (c). This stopping criterion is convenient since the switching can be performed automatically. At this point the value of σ is assumed to be recovered, and is treated as known. Meanwhile extrema of the recovered signal become easily distinguishable with stable positions.

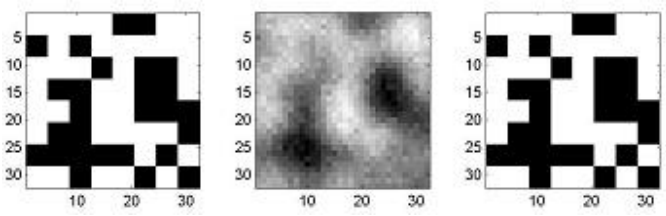

Fig. 3. Wavelet-based deconvolution (left to right): the original barcode, noisy (mean square SNR $\approx$ 12dB) scanner output with $\sigma = .038$, wavelet-based recovery with a substitute value $\sigma = .03$.

The second stage includes evolution of the signal with the known value of σ. The fidelity coefficient μ is taken to be $O(1)$, and the diffuse interface dynamics has the most influence. As a result, the very roughly shaped barcode obtained at the first stage becomes binary - as in Fig. 2 (d).

For the method of [8] one has a serious time restriction coming from the completely explicit numerical scheme for the scanning width σ. Adaptive stepping allows to speed up the process. However, at least 200 iterations are needed to perform the recovery of a signal 256 units long with a known value of σ. Naturally, more iterations are needed for blind deconvolution.

The speed of both algorithms applied to a 256 units long signal are given below (all operations not related to the PDE/wavelet change are performed in an identical manner). Those tests were performed by running the Matlab code on a laptop with CPU specifications: Intel Pentium M, 1.73GHz, 504MB of RAM.

	CPU time per 10 iter-s	# of iter-s	σ
PDE:	0.5	$\gtrsim 200$	known
PDE:	0.6	$\gtrsim 450$	unknown
Wvlt:	0.6	$\gtrsim 150$	known
Wvlt:	0.7	$\gtrsim 200$	unknown

As opposed to the wavelet version, the efficiency of the PDE algorithm depends on the choice of parameters and the starting value of σ. Those factors can influence the length of the recovery process and even its success. For "untuned" parameters the PDE method can produce artifacts, like false bars, or fail to recover the true σ. For the case of the wavelet recovery, the starting value of the scanning width is not crucial, as long as it has the same order of magnitude. Overall, the wavelet method tends to be more tolerant to parameter inconsistencies.

The method directly generalizes to 2 D bar-codes. The important practical difference lies in the necessity to use the shift-invariant stationary wavelet transform instead of the standard non-redundant method in 2D, because of undesirable artifacts the latter produces. Fig.3 shows a random combination of black and white squares of size 4, serving as a 32 x 32 pixels barcode piece, the scanner output with unknown σ (originally $\sigma = .077$), followed by its perfect recovery. This technique is easily adapted to any rectangular-based bar codes by building the rescaling parameters into the convolution kernel. Moreover, it makes sense to remove all elements of the scale smaller than that of the bar-code building element (i.e. to set the Haar basis coefficients of the finest level to be equal to zero). This operation increases the precision of signal recovery and enhances sharp edges in the final output. Without removing the last levels the deconvolution stays perfect for smaller values of sigma and lower noise.

We could recover the value of σ in the same manner as in 1D and with the same precision, but numerical computations take significantly longer due to the need for the redundant wavelet transform.

However, if we use a substitute value of σ, smaller than the true σ_0 but of the same order of magnitude, it is equivalent to adding more noise to the scanner output. If the algorithm works efficiently with much lower signal to noise ratios, then the substitution of σ does not influence the recovery output.

III. Wavelet-based Image Inpainting

A. Prior results in image inpainting.

The notion of digital image inpainting using nonlinear PDEs appeared in [12], where the authors introduced a novel technique based on a 3rd order nonlinear equation. Chan and Shen [4] adapted efficient variational denoising and segmentation models to the inpainting task by localizing the support of the forcing term inside the complement of the inpainting region. Consequent work by Esedoglu and Shen [7] utilizes this idea for converting the Mumford-Shah segmentation model to inpainting technique.

Our work builds on a recent phase-field approach [1]. It describes a model for inpainting binary images using a modified Cahn-Hilliard equation which has a number of applications for high contrast images, including inpainting of degraded text, super-resolution of images, and continuation of roads on aerial photographs. The model has two scales, the diffuse interface scale, on which it can accomplish topological transitions, and the feature scale of the image.

We remark that wavelets appeared in the context of variational image inpainting in [22]. The authors consider missing or damaged regions in the wavelet domain, instead of the pixel domain. This model has shown itself to be very efficient in reconstructing the true signal from lossy wavelet data using a TV norm in the image domain. Nevertheless, it becomes less applicable when the actual damage has occurred in the spatial domain (and the lost and true information is defined spatially). The relation between variational problems and the wavelet-based image processing algorithms is discussed in [2].

Having in mind the advantages of Cahn-Hilliard model, we keep the same general form of the energy functional and replace differential operators by pseudo-differential wavelet operators. Despite the gains made by the CH inpainting method, the presence of the diffuse interface scale introduces a longer timescale into the computation and results in local blurring of edges. Both of these issues can be addressed by the ideas described in the previous section.

B. The essential idea of the wavelet inpainting method

Rather than working with the full CH model, we adapt a simple inpainting model based on the second order Allen-Cahn equation. Adding the inpainting fidelity term completes the energy construction:

$$E(u) = \frac{\epsilon}{2}\|u\|_w^2 + \frac{1}{4\epsilon}\int W(u)dx + \frac{\mu}{2}\|(u-f)\chi_\Omega\|_{L_2}^2, \qquad (E)$$

where $W(u)$ is a double-well potential, Ω is a "known" domain in $[0,1]^2$ and $E_w(f) < \infty$. The corresponding gradient descent equation is

$$u_t = \epsilon\Delta_w u - \frac{1}{\epsilon}W'(u) - \mu(u-f)\chi_\Omega.$$

Evolving this forward in time, and running to steady state gives a solution to the inpainting problem.

C. Numerical implementation

We describe a possible approach to the numerical implementation, which, as in the barcode problem, uses the convexity splitting method analogous to [1]. Our scheme is

$$\frac{u_{n+1}-u_n}{dt} = (\epsilon\Delta_w u_{n+1} - c_1 u_{n+1}) - (-c_1 u_{n+1} + \frac{1}{\epsilon}W'(u^n)) + (-c_2 u_{n+1}) - (-c_2 u_n + \mu(u^n - f)\chi_\Omega)).$$

Lemma (stability) *The numerical schemes introduced in II-B and III-C are unconditionally stable whenever parameters c_i satisfy: $c_2 > \frac{1}{\epsilon}(3C^2\|u\|_\infty + 1)$, $c_2 > \mu$ where $\|u\|_\infty$ is an a priori bound on the L_∞ norm of the solution.*

The discrete version of the wavelet operator Δ_w (denoted by Δ_w^d in the above scheme) can be defined in different ways depending on how one defines the discrete wavelet transform on a finite interval. The classic ([13]) discrete wavelet transform (DWT) allowed to achieve plausible reconstruction only in the case of piece-wise constant images with edges located at some coarse level of the dyadic grid. However, in the general case it can produce energy minima with artifacts previously described in image denoising (see [19]).

The stationary wavelet transform (SWT) gives a perfect solution to this problem. With the use of redundant representation and averaging, effects of being "too discrete" are eliminated, and the dynamics of the solution looks similar to the PDE simulation. The cost of these improvements is more memory and time required. One way to address this is to simplify the SWT algorithm by storing fewer variants of coefficients and using less averaging.

D. Examples of the wavelet inpainting

We compare our results with those of Cahn-Hilliard inpainting. As was described in [1], Cahn-Hilliard inpainting works the best in a 2-step procedure: the first part connects pieces of the picture roughly, using a large interface width $\epsilon = O(1)$, the second part refines the details with $\epsilon = O(.01)$ after the major "information redistribution" was done.

The wavelet algorithm was designed in the same spirit, allowing the interface width parameter ϵ to change from large to small, this change being implemented as gradual "zooming" to details of smaller and smaller scale. Since ϵ puts restrictions on the maximum level of decomposition the minimizer can have, finding a steady state for a coarser scale followed by refinement allows to reconstruct "optimal" details on each subsequent wavelet level.

The following examples (Fig.4) show the typical numerical output for inpainting based on the wavelet energy. The left hand image in each picture shows the given part of the true image, with the grey area covering the missing region, Ω^c, and also serves as the initial condition when running the simulation. The right hand image is the result of inpainting, i.e. the steady state of a respective equation.

Inpainting of simple shapes.

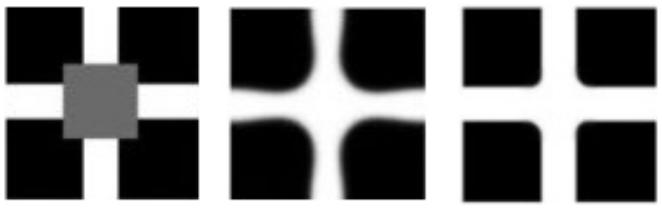

Fig. 4. Cahn-Hilliard vs. wavelet inpainting (left to right): the input image, CH inpainting (for two stages respectively: $\epsilon^{-1} = 1$ and 100, $\mu = 3 \times 10^5$ and 10^7, 500 iterations for each stage), wavelet-based inpainting $\epsilon = 1/32$, $\mu = 100$, 72 iterations.

If the missing region is comprised of relatively thin and "spread out" pieces, there is no need to run the simulation for several values of the interface parameter ϵ, because the length of pieces that need to be connected is comparable to the dominating image scale. In the example shown on Fig 5 $\epsilon = 1/32$ was kept constant until the steady state was reached after about 150 iterations.

Fig. 5. (Left vs. right) the original image with missing gray area, inpainted image. Parameters: $\epsilon = 1/32$, $mu = 500$, $dt = 1$.

The best performance is obtained with a good initial condition. One can be obtained using a standard linear extrapolation inside the unknown area. However, such minor change of initial condition does not have much influence in PDE based inpainting. As a result, the number of iterations needed to inpaint the same image for the wavelet method becomes 5-10 times less than for the PDE technique. Indeed, the difference in the number of iterations performed by each algorithm is noticeable. There is a technical obstacle that prevents the wavelet method from becoming much faster than PDE inpainting: the discrete implementation of the stationary wavelet transform in Matlab requires much more time than the fast Fourier transform. Indeed, the wavelet solver swt implementing "a trois" algorithm for the undecimated wavelet transform performs $O(N^2 \ln N)$ operations for an N-dimensional signal, while fft - standard fast Fourier algorithm needs only $O(N \ln N)$ operations. As a result, the wavelet method needs more time per iteration than the Cahn-Hilliard PDE simulation:

	CPU time per 10 iter-s	Number of iter-s
CH:	0.86	$O(10^3)$
WAC:	12.52	$O(10^1) - O(10^2)$

Nevertheless, the wavelet method overperforms PDEs if it reaches the steady state after a factor of 15 fewer iterations than Cahn-Hilliard simulation. The above table estimates the speed of our method based on the Haar wavelet, which is technically one of the slowest. Choosing the Daubechie 7/9 wavelet instead, we can decrease the CPU time per 10 iterations at least by a factor of two. An additional idea that can speed up computations is the adaptive choice of the level of decomposition. In case the missing region has relatively small scale, it suffices to change only fine and medium details of the image (presumably, the main part of the image is present and the energy balance of the coarsest levels is preserved).

Gray-scale image inpainting

To generalize the solution of the inpainting problem to the gray-scale case, let us split the signal bit-wise into channels:

$$u(x) \rightsquigarrow \sum_{m=0}^{K-1} u_m(x) 2^{-m}$$

here u_m denotes the m-th component ("digit") in the binary representation of the signal, and $u_m(x) \in \{0, 1\}$ for any x. In such a manner, the problem gets reduced to K dyadic inpainting problems, where each fidelity term is obtained as a respective component of the known part of the image, i.e., there is absolutely no interaction between channels after they are separated in the given image and before the inpainted image is synthesized from them. We give two samples of grayscale inpainting. The occlusion removal on Fig.6, where the image has less quality and more distinct color transitions, was performed using Haar wavelet. The tiger image in Fig.7 with 256 gray levels was recovered using Daubechies wavelet 7/9, with higher level of regularity and better "stripe-connecting" properties.

Fig. 6. Grayscale inpainting under occluding text. 5-bit image(16 gray levels) is inpainted after 58 iterations (using Haar wavelet).

Fig. 7. Grayscale inpainting under occluding text. 8-bit image(16 gray levels) is inpainted after 40 iterations (using Daubechies 7/9 wavelet).

Road inpainting

The described technique can be utilized to perform specialized inpainting: for instance, to follow roads in aerial photos with possible defects and occlusions. In our example of the road inpainting (Fig. III-D a) the missing area is determined by thresholding the image and treating all pixels with low intensity as unknown (black area in Fig.III-D b) However, occlusions could be detected by other methods, e.g. using multi-spectral data to identify occluded regions of a photo.

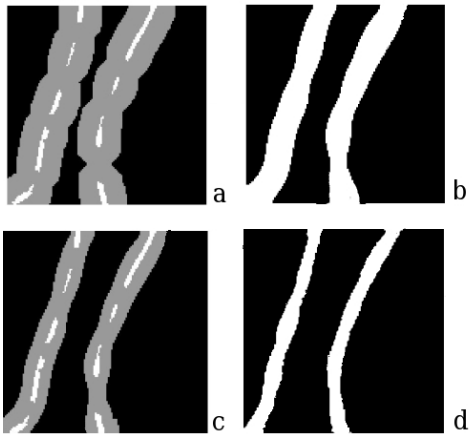

Fig. 8. Leading bit-channel inpainting. a) first stage setting; b) first stage result; c) second stage setting; d) second stage result.

As the objective of inpainting is the road recovery, the first (highest intensity) bit of the image is prioritized as the "leading channel"

and treated separately. To create the predisposition for broken road connection without merging distinct neighboring roads, we choose the maximum distance R allowed between the segments of the same road (12 pixels). Let V be the set of the initially visible road parts (the white area in Fig. 8 a). On the first stage of the leading channel recovery we assume the unknown area Ω_1^c to be the R-pixel neighborhood of V excluding V itself (the gray area in Fig. III-D a). We choose the initial value $u_0 = \chi_V + 1/2\chi_{\Omega_1^c}$. After performing the binary inpainting we get an intermediate solution u such that supp(u) is a connected road covering (white area in Fig. 8 b), nevertheless too wide to determine the precise road location. On the second stage we refine the unknown area $\Omega_2^c = \text{supp}(u) \setminus V$(the gray area in Fig.8 c) and repeat the binary inpainting. The output u represents the recovered leading channel (Fig. 8 d). Other bit channels are recovered by usual binary inpainting with unknown area as in Fig.III-D b.

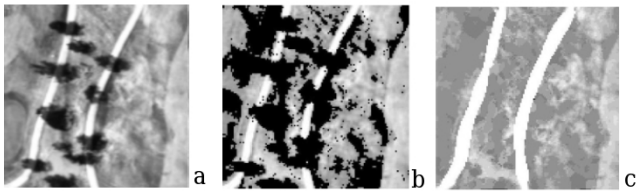

Fig. 9. Multistep road inpainting: a) the original road photo with natural and artificial shadows occluding the road silhouette; b) thresholded image with black area to be inpainted; c) the final output of bitwise inpainting.

ACKNOWLEDGMENT

The authors thank Selim Esedoglu, Jean-Michel Morel and Mauro Maggionini for useful conversations and the NGA for the idea of the road inpainting application.

REFERENCES

[1] S. Esedoglu A. Bertozzi and A. Gillette. Inpainting of binary images using the cahn-hilliard equation. *IEEE Transactions on Image Processing*, 16(1), January 2007.

[2] N. Lee A. Chambolle, R. A. DeVore and B. J. Lucier. Nonlinear wavelet image processing: Variational problems, compression, and noise removal through wavelet shrinkage. *IEEE Transactions on Image Processing*, 7(3):319–333, 1998.

[3] Yves Meyer Ali Haddad. Variational methods in image processing. Cam report, UCLA, September 2004.

[4] T. F. Chan and J. Shen. Mathematical models of local non-texture inpaintings. *SIAM Journal on Applied Mathematics*, 62(3):1019–1043, 2001.

[5] David L. Donoho. Nonlinear wavelet methods for recovery of signals, densities, and spectra from indirect and noisy data. In *Proceedings of Symposia in Applied Mathematics*, 1993.

[6] Heike Emmerich. *Diffuse Interface Approach in Materials Science Thermodynamic Concepts and Applications of Phase-Field Models*. Springer books, 2003.

[7] S. Esedoglu and J. Shen. Digital inpainting based on the mumford-shah-euler image model. *European Journal of Applied Mathematics*, (13):353–370, 2002.

[8] Selim Esedoglu. Blind deconvolution of bar code signals. *Inverse Problems*, (20):121–135, 2004.

[9] E. De Giorgi. Some remarks on gamma-convergence and least squares methods. In G. Dal Maso and G. F DellAntonio, editors, *Composite Media and Homogenization Theory (Proceedings)*, pages 135–142. Birkhauser, 1991.

[10] A. Bertozzi M. Bertalmio and G. Sapiro. Navier-stokes, fluid dynamics, and image and video inpainting. *IEEE Computer Vision and Pattern Recognition (CVPR), Hawaii*, 1(I):355–362, December 2001.

[11] G. Sapiro M. Bertalmio, L. Vese and S. Osher. Simultaneous structure and texture image inpainting. *IEEE Transactions on Image Processing*, 12(8):882–889, 2003.

[12] V. Caselles M. Bertalmio, G. Sapiro and C. Ballester. Image inpainting. In Kurt Akeley, editor, *Siggraph 2000, Computer Graphics Proceedings*, pages 417–424. Wesley Longman, 2000.

[13] S. Mallat. *Wavelet Tour of Signal Processing*. Academic Press, September 1999.

[14] L. Modica and S. Mortola. Un esempio di gamma-convergenza. *Boll. Un. Mat. Ital.*, 5(14:1):285–299, 1977.

[15] S. Zhou M.Y. Wang and H. Ding. Non-linear diffusions in topology optimization. *Structural and Multidisciplinary Optimization*, 28(4):262–276, October 2004.

[16] R. Neelamani, H. Choi, and R. Baraniuk. Wavelet-based deconvolution using optimally regularized inversion for ill-conditioned systems. *Wavelet Applications in Signal and Image Processing VII, Proc. SPIE*, 3813:58–72, July 1999.

[17] Stefano Soatto Luminita A. Vese Paolo Favaro, Stanley Osher. 3d shape from anisotropic diffusion. *IEEE Computer Vision and Pattern Recognition (CVPR)*, 1(I):179–186, 2003.

[18] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, I(12):629–639, 1990.

[19] D.L.Donoho R.R.Coifman. Translation-invariant de-noising. Lecture Notes in Statistics: Wavelets and Statistics, New York: Springer-Verlag, 1995. pp. 125-150.

[20] Fatemi E. Rudin L. I., Osher S. Nonlinear total variation based noise removal algorithms. *Physica D.*, 60:259–268, 1992.

[21] J. Shen T. F. Chan and L. Vese. Variational pde models in image processing. *Notices of the American Mathematical Society*, 50(1):14–26, 2003.

[22] J. Shen T. F. Chan and H.-M. Zhou. Total variation wavelet inpainting. *Journal of Mathematical Imaging and Vision*, 25(1), July 2006.

[23] S.H.Kang Y.M. Jung and J. Shen. Multiphase image segmentation via modica-mortola phase transition. Cam report, UCLA, 2006.