From High Energy Physics to Low Level Vision

Ron Kimmel[1] and Nir Sochen[2] and Ravi Malladi[1]

[1] Lawrence Berkeley National Laboratory, and Dept. of Mathematics
University of California, Berkeley, CA 94720.
[2] Raymond and Beverly Sackler Faculty of Exact Sciences Tel-Aviv University, Israel.

Abstract. A geometric framework for image scale space, enhancement, and segmentation is presented. We consider intensity images as surfaces in the $(\mathbf{x}, \mathbf{I})$ space. The image is thereby a 2D surface in 3D space for gray level images, and a 2D surface in 5D for color images. The new formulation unifies many classical schemes and algorithms via a simple scaling of the intensity contrast, and results in new and efficient schemes. Extensions to multi dimensional signals become natural and lead to powerful denoising and scale space algorithms. Here, we demonstrate the proposed framework by applying it to denoise and improve gray level and color images.

1 Introduction: A philosophical point of view

In this paper we adopt an action potential that was recently introduced in physics and use it to produce a natural scale space for images as surfaces. It will lead us to the construction of image enhancement procedures for gray and color images. This model also integrates many existing segmentation and scale space procedures by a change of a single parameter that switches between the L_1 and L_2 Euclidean norms.

Let the input to the low level vision process be a map $\mathbf{X} : \Sigma \rightarrow M$ where Σ is a one, two, or three dimensional manifold and $\mathbf{X}$ is the embedding of this manifold in a space which is a hybrid space of spatial coordinates and feature coordinates, the "space-feature". For example, the most common map is from a two dimensional surface to $\mathbb{R}^3$ where we have at each point of the plane an intensity $I(x, y)$. The $\mathbb{R}^3$ space-feature has Cartesian coordinates (x, y, I) where x and y are the spatial coordinates and I is the feature coordinate [3]. Higher dimensions of the embedding space are encountered for example in color images. Three dimensional manifolds Σ occur in movie analysis and in volumetric medical images [13]. The output of the low level process in most models consists of a 'simplified', 'denoised', 'deblurred', 'segmented', or 'cleaned' image, for further analysis and processing.

The importance of the dynamics of the image geometry in the perception and understanding of images is by now well established in computer vision. There are many definitions for scale space of images aiming to arrive at a coherent

[3] While in this paper the feature coordinate is simply the zeroth jet space $j^0 I$, we use the term feature space to leave room for a more general cases like texture [12] etc.

framework that unifies many requirements. One such requirement is that *"only isophotes matter"*, or equivalently assume the importance of the *morphological assumption* of the scale space to be *contrast invariant*. We argue that this assumption, though leading to many interesting results, seems to fail in many other natural cases. Let us demonstrate it with a very simple example: Consider the intensity image of a dark object in a white background. At this point the boundary of the object is closely related to one of the isophotes of the gray level image. Consider the intensity image as a function, and add to this function a new smooth function (e.g. a tilted plane). This additional smooth function might be the result of non-uniform lighting conditions [26]. It is clear that in the new intensity image the isophotes play only a minor role in the understanding process.

The importance of edges in scale space construction is obvious: Boundaries between objects should survive as long as possible along the scale space, while homogeneous regions should be simplified and flattened in a more rapid way. We also want to preserve the geometrical beauty that results in some interesting non-linear 'scale spaces' as a result of the morphological assumption. Among these are the Euclidean and affine invariant flows [1, 2, 22]. Moreover, we want our framework to handle multi channel images. A color image is a good example since we actually talk about 3 images (Red, Green, Blue) that are composed into one.

We propose to view images as embedding maps, that flow towards minimal surfaces. We go two dimensions higher than most of the classical schemes [8], and instead of dealing with isophotes as planar curves we deal with the whole image as a surface. For example, we consider a color image as a 2D surface in 5D (x,y,R,G,B).

Section 2 introduces the arclength and the definition of a metric on a surface. Next, Section 3 presents the "action" that we borrowed from high energy physics and the way it produces a general framework for non-linear diffusion in computer vision. In Section 4 we introduce a new flow that we have chosen to name *Beltrami flow*, and present a geometric interpretation in the simplest 3D case. Next, in Section 5 we present the resulting flow for multi channel (color) images and its advantages over previous models, and show some experiments with color images with and without constraints. We refer the interested reader to [24] for further details and more examples.

2 Measuring Distances on Surfaces: The Induced Metric

The basic concept of Riemannian differential geometry is distance. The natural question is how should we measure distances? We will first take the important example $\mathbf{X} : \Sigma \rightarrow \mathbb{R}^3$. We denote the local coordinates on the two dimensional manifold Σ by (σ^1, σ^2), see Fig. 1. The map $\mathbf{X}$ is explicitly given by $(X^1(\sigma^1, \sigma^2), X^2(\sigma^1, \sigma^2), X^3(\sigma^1, \sigma^2))$. Since the local coordinates σ^i are curvilinear, and not orthogonal in general, the distance square between two close points on Σ, $p = (\sigma^1, \sigma^2)$ and $p + (d\sigma^1, d\sigma^2)$ is not $ds^2 = d\sigma_1^2 + d\sigma_2^2$. In fact, the squared

distance is given by a positive definite symmetric bilinear form called the metric whose components we denote by $g_{\mu\nu}(\sigma^1, \sigma^2)$:

$$ds^2 = g_{\mu\nu} d\sigma^\mu d\sigma^\nu \equiv g_{11}(d\sigma^1)^2 + 2g_{12} d\sigma^1 d\sigma^2 + g_{22}(d\sigma^2)^2, \tag{1}$$

where we used Einstein summation convention in the second equality; identical indices that appear one up and one down are summed over. We will denote the inverse of the metric by $g^{\mu\nu}$, so that $g^{\mu\nu} g_{\nu\gamma} = \delta^\mu_\gamma$, where δ^μ_γ is the Kronecker delta.

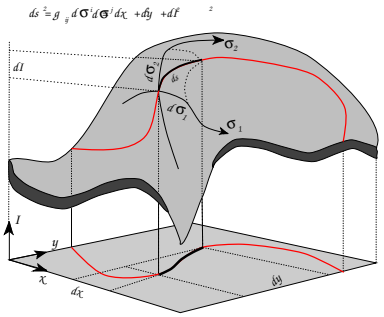

Fig. 1. Length element of a surface curve ds, may be defined either as a function of a local metric defined on the surface $(\sigma_1, \sigma_2; (g_{ij}))$, or as a function of the coordinates of the space in which the surface is embedded (x, y, I).

Let $\mathbf{X} : \Sigma \to M$ be an embedding of (Σ, g) in (M, h), where Σ and M are Riemannian manifold and g and h are their metrics respectively. We can use the knowledge of the metric on M and the map $\mathbf{X}$ to construct the metric on Σ. This procedure, which is denoted formally as $(g_{\mu\nu})_\Sigma = \mathbf{X}^*(h_{ij})_M$, is called the *pullback* and is given explicitly as:

$$g_{\mu\nu}(\sigma^1, \sigma^2) = h_{ij}(\mathbf{X}) \partial_\mu X^i \partial_\nu X^j, \tag{2}$$

where $i, j = 1, ..., \dim M$ are being summed over, and in short we use $\partial_\mu X^i \equiv \partial X^i(\sigma^1, \sigma^2)/\partial\sigma^\mu$.

An example, often used in computer vision, is the embedding of a surface described as a graph in $\mathbb{R}^3$: $\mathbf{X} : (\sigma^1, \sigma^2) \to (\sigma^1, \sigma^2, I(\sigma^1, \sigma^2))$.Using Eq. (2) we get

$$(g_{\mu\nu}) = \begin{pmatrix} 1 + I_x^2 & I_x I_y \\ I_x I_y & 1 + I_y^2 \end{pmatrix} \tag{3}$$

where we used the identification $x \equiv \sigma^1$ and $y \equiv \sigma^2$ in the map $\mathbf{X} = (x, y, I)$.

Actually we can understand this result in an intuitive way: Eq. (2) means that the distance measured on the surface by the local coordinates is equal to the

distance measured in the embedding coordinates, see Fig. 1. Under the above identification, we can write $ds^2 = dx^2 + dy^2 + dI^2 = dx^2 + dy^2 + (I_x dx + I_y dy)^2 = (1 + I_x^2)dx^2 + 2I_x I_y dx dy + (1 + I_y^2)dy^2$.

Next we provide a measure on the space of these maps.

3 Polyakov Action and Harmonic Maps

In this section, we present a general framework for non-linear diffusion in computer vision. We have shown in [24] that many of the known methods fall naturally into this framework. Here we show how to derive new ones. The equations will be derived by a minimization problem from an action functional. The functional in question depends on *both* the image manifold and the embedding space. Denote by (Σ, g) the image manifold and its metric and by (M, h) the space-feature manifold and its metric, then the map $\mathbf{X} : \Sigma \to M$ has the following weight

$$S[X^i, g_{\mu\nu}, h_{ij}] = \int d^m\sigma \sqrt{g} g^{\mu\nu} \partial_\mu X^i \partial_\nu X^j h_{ij}(\mathbf{X}), \tag{4}$$

where m is the dimension of Σ, g is the determinant of the image metric, $g^{\mu\nu}$ is the inverse of the image metric, the range of indices is $\mu, \nu = 1, \ldots, \dim \Sigma$, and $i, j = 1, \ldots, \dim M$, and h_{ij} is the metric of the embedding space.

This functional, for $m = 2$, was first proposed by Polyakov [18] in the context of high energy physics. Given the above functional, we have to choose the minimization. We may choose for example to minimize with respect to the embedding alone. In this case the metric $g_{\mu\nu}$ is treated as a parameter of the theory and may be fixed by hand. Another choice is to vary only with respect to the feature coordinates of the embedding space, or we may choose to vary with respect to the image metric as well. We have shown that different choices yield different flows. Some flows are recognized as existing methods, other choices are new and some will be described below.

Using standard methods in variational calculus, the Euler-Lagrange equations with respect to the embedding are (see [24] for a derivation):

$$-\frac{1}{2\sqrt{g}} h^{il} \frac{\delta S}{\delta X^l} = \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu X^i) + \Gamma^i_{jk} \partial_\mu X^j \partial_\nu X^k g^{\mu\nu} = 0. \tag{5}$$

where Γ^i_{jk} are the Levi-Civita connection coefficients with respect to the metric h_{ij} that describes the geometry of the embedding space (see [24, 25] for a definition of the Levi-Civita connection).

Our proposal is to view scale-space as the gradient descent flow:

$$X^i_t \equiv \frac{\partial X^i}{\partial t} = -\frac{1}{2\sqrt{g}} h^{il} \frac{\delta S}{\delta X^l} \tag{6}$$

Few remarks are in order. First notice that we used our freedom to multiply the Euler-Lagrange equations by a strictly positive function and a positive definite matrix. This factor is the simplest one that does not change the minimization

solution while giving a reparametrization invariant expression i.e. invariant under $\sigma^\mu \rightarrow \tilde{\sigma}^\mu(\sigma^1, \sigma^2)$. The operator that is acting on X^i in the first term is the natural generalization of the Laplacian from flat spaces to manifolds and is called *the second order differential parameter of Beltrami* [14], or for short *Beltrami operator*, and we will denote it by Δ_g.

When the embedding is in an Euclidean space with Cartesian coordinate system the connection elements are zero. In this case $\Delta_g \mathbf{X}$, for grey-level image, is the usual mean curvature vector. This simple definition for the general mean curvature vector provides a straightforward calculation procedure, extends to higher dimensions, to cases with co-dimension greater than one, and to non-trivial geometries and coordinate systems; e.g. we consider a color image as a 2D surface in 5D, in which case the co-dimension is 3.

The Beltrami operator with a metric that corresponds to the plane with non-Cartesian coordinate system was discussed in Florac *et al.* [10]. Our approach is a generalization in two ways, one is the choice of a metric with non-trivial Riemann tensor (or equivalently for surfaces, the Gaussian curvature of the image manifold is different from zero), the other is the possibility to deal with non-trivial embedding. We also have here a framework that can treat curves, surfaces, and higher dimensional image data embedded in gray, color and higher dimensional and geometrically non-trivial embedding spaces.

Evolving a surface according to its curvature vector $H\boldsymbol{\mathcal{N}} = \Delta_g \mathbf{X}$ is the steepest descent flow towards a minimal surface, and may be written as

$$\mathbf{X}_t = H\boldsymbol{\mathcal{N}}, \tag{7}$$

where H is the mean curvature, $\boldsymbol{\mathcal{N}}$ is the normal to the surface:[4] For co-dimension 1:

$$\boldsymbol{\mathcal{N}} = \frac{1}{\sqrt{g}}(-\nabla I, 1)^T \tag{8}$$

where $g = 1 + |\nabla I|^2$. This is mean curvature flow! This should not be a surprise, since if we check the action functional, we notice that, for the choice of the induced metric Eq. (3) as the image metric $g_{\mu\nu}$, we are left with $S = \int d^2\sigma\sqrt{g} = \int d^2\sigma\sqrt{\det(\partial_\mu X^i \partial_\nu X_i)}$,which is the Euler functional that describes the area of the surface (also known in high energy physics as the Nambu action).

In [24] we survey different choices for the dynamic and parametric degrees of freedom in the action functional that lead to known methods. These include the reparameterization invariant linear scale-space by Florac *et al.* [10], Perona-Malik anisotropic diffusion [17], geodesic active contour models for segmentation, Mumford-Shah segmentation models that are based on the L_2 ($\int |\nabla I|^2$) norm [19, 16], Rudin-Osher-Fatemi total variation (TV) method [20] for image enhancement based on the L_1 ($\int |\nabla I|$) norm, and the different Blake-Zisserman membrane models [3]. We actually show that by varying the aspect ratio between the I axis and the xy axes, we can switch between the L_1 and the L_2

[4] In what follows, we denote by g the determinant of the metric, $g \equiv \det(g)$, the metric itself will be denoted as (g_{ij}). Note also that some definitions of the mean curvature include a factor of 2 that we omit in our definition.

norms for image processing. In fact, we can approach the L_1 norm, which is practically regularized in most application to avoid zero denominator. The regularized functional $\int \sqrt{\beta^2 + |\nabla I|^2}$ may be viewed as an area minimization, that approaches the L_1 total variation norm as $\beta \to 0$. For gray level images it is just a mathematical exercise, however, when we deal with more complicated cases like multi channel images or color images, we have a very natural extension of the total variation method. We will show why this extension is better and more natural than previous multi channel norms.

For images which are maps from an m dimensional manifold to n dimensional embedding space with $n-m > 1$, the normals to the image span an $n-m$ normal space. The way the mean curvature is generalized to these maps (for Euclidean embedding spaces with Cartesian Coordinate system) is via the Beltrami operator. This operator is built from the metric only, and there is no need for any extrinsic information to express it. It acts on the embedding coordinates and coincides with the usual definition of the mean curvature for hypersurfaces. When the metric of the embedding space is not trivial the (generalized) mean curvature flow is obtained by the more general Eq. (6).

Note that this flow is the (generalized) mean curvature flow only if we move all the coordinates X^i simultaneously. Below we concentrate on another possibility. We flow only the feature coordinate(s). We call this generalized flow the *Beltrami flow* and discuss its characteristics in the next section.

4 The Beltrami flow

Let the image be an embedding map $\mathbf{X} : \Sigma \to \mathbb{R}^3$, where Σ is a two dimensional manifold, and the flow is natural in the sense that it minimizes the action functional with respect to I and (g_{ij}), while being reparametrization invariant. The coordinates x and y are parameters from this view point and are identified as above with σ^1 and σ^2 respectively. The result of the minimization is the Beltrami operator acting on I:

$$I_t = \Delta_g I \equiv \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu I) = H \boldsymbol{\mathcal{N}}_{\hat{I}} \tag{9}$$

where the metric is the induced one given in Eq. (2), and $\hat{I}$ is the unit vector in the I direction.[5]

[5] The mean curvature flow can be written as $\frac{\partial}{\partial t}\begin{pmatrix}\mathbf{x}\\ \mathbf{I}\end{pmatrix} = \Delta_g \begin{pmatrix}\mathbf{x}\\ \mathbf{I}\end{pmatrix}$. Fixing the xy coordinates amounts to moving the surface via its mean curvature feature components $\mathbf{I}$, thereby preserving the edges at which these component are small: Along the edges (cliffs), the surface normal is almost parallel to the x-y plane. Thus, $I(x,y)$ hardly evolves along the edges while the flow drives other regions of the image towards a minimal surface at a more rapid rate.

4.1 Geometric Flows Towards Minimal Surfaces

A minimal surface is the surface with least area that satisfies given boundary conditions. It has nice geometrical properties, and is often used as a natural model of various physical phenomena, e.g. soap bubbles "Plateau's problem", in computer aided design, in architecture (structural design), and recently even for medical imaging [5]. It was realized by J. L. Lagrange in 1762, that the mean curvature equal to zero is the Euler Lagrange equation for area minimization. Hence, the mean curvature flow is the most efficient flow towards a minimal surface; see Fig. 2 (left).

We refer to [7] for the derivation of H for a graph surface $\mathcal{S} = (x, y, I(x, y))$ (as D.L. Chopp summarizes the original derivation by J.L. Lagrange from 1762)

$$H = \operatorname{div}\left(\frac{\nabla I}{\sqrt{1+|\nabla I|^2}}\right) = \frac{(1+I_y^2)I_{xx} - 2I_x I_y I_{xy} + (1+I_x^2)I_{yy}}{(1+I_x^2+I_y^2)^{3/2}}. \tag{10}$$

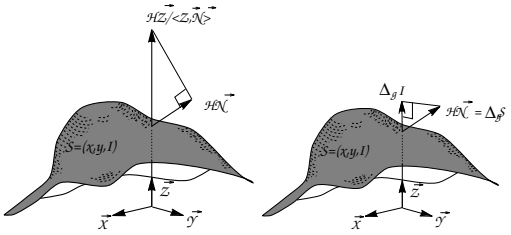

Fig. 2. Left: Mean curvature flow. Consider the surface mean curvature flow $\mathcal{S}_t = H\mathcal{N}$. A geometrically equivalent flow is the graph flow $I_t = H(1+|\nabla I|^2)^{1/2}$ which yields the mean curvature flow when projected onto the normal. Right: Beltrami flow. Now, consider the mean curvature vector $H\mathcal{N}$. It can also be expressed as $H\mathcal{N} = \Delta_g \mathcal{S}$. Beltrami operator that operates on I: $\Delta_g I$, is the third component of this vector: Projection onto the I direction.

The mean curvature for a graph (the image) is given by the following evolution equation

$$I_t = \frac{(1+I_y^2)I_{xx} - 2I_x I_y I_{xy} + (1+I_x^2)I_{yy}}{1+I_x^2+I_y^2}, \tag{11}$$

with the image itself as initial condition $I(x, y, 0) = I(x, y)$. Using Beltrami second order operator Δ_g, Eq. (11) may be read as $I_t = g\Delta_g I$. The Beltrami flow (selective mean curvature flow) on the other hand, $I_t = \Delta_g I$, is given explicitly for the gray level case as

$$I_t = \frac{(1+I_y^2)I_{xx} - 2I_x I_y I_{xy} + (1+I_x^2)I_{yy}}{(1+I_x^2+I_y^2)^2}, \tag{12}$$

see Fig. 2 (right). In other words, the Beltrami flow can be viewed as $\mathcal{S}_t = \frac{1}{g} H \boldsymbol{\mathcal{N}}$, where $g = \sqrt{1+|\nabla I|^2}$ is basically an edge indicator. Other methods that consider gray level images as surfaces are Yanowitz and Bruckstein [26], El-Fallah et al. [9], and Malladi and Sethian [15].

Fig. 3 compares the results of the Beltrami flow and the mean curvature flow both applied to a digital subtraction angiogram (DSA). It demonstrates the edge preserving property of the Beltrami flow.

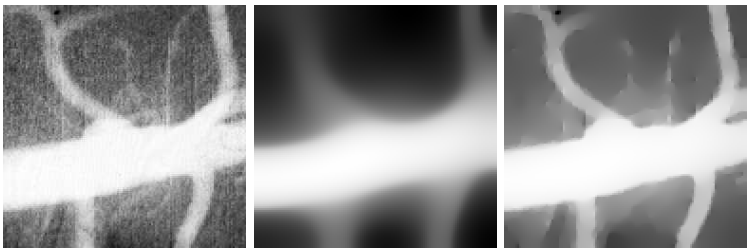

Fig. 3. On the left is the original medical image. In the middle is the result of smoothing via the mean curvature flow, and on the right is the result of the Beltrami flow.

5 Color

We generalize the Beltrami flow to the 5 dimensional space-feature needed in color images. The embedding space-feature space is taken to be Euclidean with Cartesian coordinate system. The image, thus, is the map $f : \Sigma \rightarrow \mathbb{R}^5$ where Σ is a two dimensional manifold.
Explicitly the map is $f = \left(x(\sigma^1,\sigma^2), y(\sigma^1,\sigma^2), I^r(\sigma^1,\sigma^2), I^g(\sigma^1,\sigma^2), I^b(\sigma^1,\sigma^2)\right)$.

We minimize our action (4) with respect to the metric and with respect to (I^r, I^g, I^b). For convenience we denote (r,g,b) by $(1,2,3)$, or in general notation i. Minimizing the action with respect to the metric gives, as usual, the induced metric which is now given by:

$$g_{\mu\nu} = \delta_{\mu\nu} + \sum_i \left[(\partial_\mu I^i)(\partial_\nu I^i)\right],$$

where $\delta_{\mu\nu}$ is the Kronecker delta. The determinant is $g = \det(g_{ij}) = g_{11}g_{22} - g_{12}^2$. Note that this metric differs from the Di Zenzo metric [8] by the addition of 1 to g_{11} and g_{22}. The source of the difference lies is the map used to describe the image; Di Zenzo used $f : \Sigma \rightarrow \mathbb{R}^3$ while we use $f : \Sigma \rightarrow \mathbb{R}^5$.

The action functional under this choice of the metric is the Euler functional $S = \int d^2\sigma\sqrt{g}$,where the *generalized surface area element* $\sqrt{g}$ is defined by

$$g = 1 + \sum_i |\nabla I^i|^2 + \frac{1}{2}\sum_{ij}(\nabla I^i, \nabla I^j)^2, \tag{13}$$

where $(\nabla I^i, \nabla I^j)$ stand for the magnitude of the vector product of the vectors ∇I^i and ∇I^j.

The action is simply the area of the image surface. Minimization with respect to I^i gives the Beltrami flow

$$I^i_t = \Delta_g I^i, \tag{14}$$

where $\Delta_g I^i = \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu I^i)$. Again, this is a flow towards a minimal surface.

5.1 Relation to other color diffusion methods

Chambolle [6], and Sapiro and Ringach [21], generalized the idea of smoothing a single valued function via a second directional derivative in the direction of minimal change, i.e. isophotes curvature flow, into a multi valued function. These are non-variational flows.

As pointed out in [4], for image segmentation, edge preserving and selective smoothing purposes, this is a result of a weakly coupled definition in color space. Blomgren and Chan [4] try to improve these results and defined the color TV norm as $\sqrt{\sum_{i=1} \left(\int |\nabla I^i| \right)^2}$, with a constraint. Observe that in this case the coupling between the channels is only by the constraint. Actually, without the constraint the minimization yields a channel by channel curvature flow.

A different norm was used by Shah in [23]: $\int \sqrt{\sum_{i=1} |\nabla I^i|^2}$. As in all the previous norms this norm is simplified to the L_1 (TV) norm for the single channel case.

We notice that the proposed area $\int \sqrt{g}$ norm, Eq. (13), includes a new term that does not appear in previous norms. The term $\sum_{ij} (\nabla I^i, \nabla I^j)^2$ measures the directional difference of the gradient between different channels. The minimization of a norm that includes this term, directs different channels to align together as they become smoother and simpler in scale. One should recognize this cross correlation of orientation between the channels as a very important feature, e.g. overcoming the color fluctuations along edges as a result of a lossy JPEG compression.

If we now tune the regularization ratio β to small values, we approach an extension of the TV norm that still includes the alignment term and serves as a natural coupling between the channels, see [11] for further details.

5.2 Experimental Results

We now present some results of denoising color images using our model. Spatial derivatives are approximated using central differences and an explicit Euler step is employed to reach the solution. We have tested the Beltrami flow with and without constraints[6] on color images. See Fig. 4.

[6] It is possible to impose a meaningful convergence on the Beltrami flow through the right constraints on the action functional. As a simple example we derive a variance constraint similar to the TV method [20] for image denoising with convergence.

Fig. 4. Upper row: The Beltrami flow as an edge preserving scale space in color. Second row: Reconstructing a color image with noise artifacts introduced by wavelet lossy compression. The noisy image is on the left, the next image is the result of applying the Beltrami flow without constraints. Next is a reconstruction with noise artifacts introduced by JPEG compression. Again the noisy image is on the left and the reconstruction on the right Third row: Reconstruction of a color image corrupted with Gaussian noise; the second image is the result of flow with constraints (convergent scheme), while the third image is the result of the flow without constraints after the same number of numerical iterations. [This is a color image].

Note that since we have a powerful selective smoothing operator, good results may be obtained even without invoking these constraints. Without the

---

Given the variance for every channel, i.e. $\int (I^i - I_0^i)^2 dxdy = \sigma_i^2$ where σ_i is the given noise variance for channel i. The Euler Lagrange is $\sum_{\alpha=x,y} \frac{1}{2}\partial_\alpha \left(\frac{1}{\sqrt{g}} \frac{\partial g}{\partial I^i_\alpha} \right) - \lambda^i (I^i - I_0^i) = 0$. Again, using the freedom of parametrization (multiplying by $g^{-1/2}$), yields the flow $I_t^i = \Delta_g I^i - \frac{1}{\sqrt{g}} \lambda^i (I^i - I_0^i)$, where λ^i is computed via $\lambda^i = -\frac{1}{2\sigma_i^2} \int \sum_{\alpha=x,y} \frac{1}{\sqrt{g}} \frac{\partial g}{\partial I^i_\alpha} (I^i_\alpha - I^i_{0\alpha}) dxdy$.

We used the notation $\partial g / \partial I^i_\alpha$, that for the multi channel (color) case simplifies to $\frac{\partial g}{\partial I^i_x} = 2I^i_x g_{22} - 2I^i_y g_{12}$, and $\frac{\partial g}{\partial I^i_y} = 2I^i_y g_{11} - 2I^i_x g_{12}$.

constraints, the time we run the evolution is related to the noise variance.

6 Concluding Remarks

Inventing a perceptually good segmentation process, and formulating a meaningful scale space for images is not an easy task. Here we tried to address these questions and to come up with a new framework that unifies many previous results and introduces new procedures. There are still many open questions to be asked, like what is the right aspect ratio between the intensity and the image plane? An even deeper question to answer is what is the 'right' embedding space h_{ij}?

The question of what is the 'right norm' when dealing with images is indeed not trivial, and the right answer probably depends on the application. For example, the answer for the 'right' color metric h_{ij} is the consequence of empirical results, experimental data, and the application. Here we covered some of the gaps between the two classical norms in a geometrical way and proposed a new approach to deal with multi dimensional images. We used recent results from high energy physics that yield promising algorithms for enhancement, segmentation and scale space.

Acknowledgments

We thank David Adalsteinsson for his comments on embedding spaces, Korkut Bardakçi for discussions on extrinsic properties of surface embedding, and David Marimont for supplying the color images. This work is supported in part by the Applied Math. Subprogram of the OER under DE-AC03-76SFOOO98, ONR grant under NOOO14-96-1-0381, and NSF under grant PHY-90-21139. All calculations were performed at the Lawrence Berkeley National Laboratory, University of California, Berkeley.

References

1. L Alvarez, F Guichard, P L Lions, and J M Morel. Axioms and fundamental equations of image processing. *Arch. Rational Mechanics*, 123, 1993.
2. L Alvarez, P L Lions, and J M Morel. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Numer. Anal*, 29:845–866, 1992.
3. A Blake and A Zisserman. *Visual Reconstruction*. MIT Press, Cambridge, Massachusetts, 1987.
4. P Blomgren and T F Chan. Color TV: Total variation methods for restoration of vector valued images. cam TR, UCLA, 1996.
5. V Caselles, R Kimmel, G Sapiro, and C Sbert. Minimal surfaces: A geometric three dimensional segmentation approach. *Numerische Mathematik*, to appear, 1996.
6. A Chambolle. Partial differential equations and image processing. In *Proceedings IEEE ICIP*, Austin, Texas, November 1994.
7. D L Chopp. Computing minimal surfaces via level set curvature flow. *J. of Computational Physics*, 106(1):77–91, May 1993.

8. S Di Zenzo. A note on the gradient of a multi image. *Computer Vision, Graphics, and Image Processing*, 33:116–125, 1986.
9. A I El-Fallah, G E Ford, V R Algazi, and R R Estes. The invariance of edges and corners under mean curvature diffusions of images. In *Processing III SPIE*, volume 2421, pages 2–14, 1994.
10. L M J Florack, A H Salden, , B M ter Haar Romeny, J J Koendrink, and M A Viergever. Nonlinear scale-space. In B M ter Haar Romeny, editor, *Geometric–Driven Diffusion in Computer Vision*. Kluwer Academic Publishers, The Netherlands, 1994.
11. R Kimmel. What is a natural norm for multi channel image processing. LBNL report, Berkeley Labs. UC, CA 94720, March 1997.
12. R Kimmel, N Sochen, and R Malladi. On the geometry of texture. Report LBNL-39640, UC-405, Berkeley Labs. UC, CA 94720, November 1996.
13. R Kimmel, N Sochen, and R Malladi. Images as embedding maps and minimal surfaces: Movies, color, and volumetric medical images. In *Proc. of IEEE CVPR'97*, Puerto Rico, June 1997.
14. E Kreyszing. *Differential Geometry*. Dover Publications, Inc., New York, 1991.
15. R Malladi and J A Sethian. Image processing: Flows under min/max curvature and mean curvature. *Graphical Models and Image Processing*, 58(2):127–141, March 1996.
16. D Mumford and J Shah. Boundary detection by minimizing functionals. In *Proceedings of CVPR, Computer Vision and Pattern Recognition*, San Francisco, 1985.
17. P Perona and J Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE-PAMI*, 12:629–639, 1990.
18. A M Polyakov. *Physics Letters*, 103B:207, 1981.
19. T Richardson and S Mitter. Approximation, computation, and distoration in the variational formulation. In B M ter Haar Romeny, editor, *Geometric–Driven Diffusion in Computer Vision*. Kluwer Academic Publishers, The Netherlands, 1994.
20. L Rudin, S Osher, and E Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
21. G Sapiro and D L Ringach. Anisotropic diffusion in color space. *IEEE Trans. Image Proc.*, to appear, 1996.
22. G Sapiro and A Tannenbaum. Affine invariant scale–space. *International Journal of Computer Vision*, 11(1):25–44, 1993.
23. J Shah. Curve evolution and segmentation functionals: Application to color images. In *Proceedings IEEE ICIP'96*, pages 461–464, 1996.
24. N Sochen, R Kimmel, and R Malladi. From high energy physics to low level vision. Report LBNL 39243, LBNL, UC Berkeley, CA 94720, August 1996. *http://www.lbl.gov/∼ron/belt−html.html*.
25. M Spivak. *A Comprehensive Introduction to Differential Geometry*. Publish or Perish, Inc., Berkeley, 1979.
26. S D Yanowitz and A M Bruckstein. A new method for image segmentation. *Computer Vision, Graphics, and Image Processing*, 46:82–95, 1989.

This article was processed using the LaTeX macro package with LLNCS style