Images as Embedded Maps and Minimal Surfaces: Movies, Color, Texture, and Volumetric Medical Images *

R. Kimmel,[1] R. Malladi,[2] and N. Sochen[3]

[1] Computer Science Department, Technion, Haifa 32000, Israel.

[2] Lawrence Berkeley National Laboratory
University of California, Berkeley, CA 94720.

[3] Electrical Engineering Department Technion, Haifa 32000, Israel.

Abstract

We extend the geometric framework introduced in [46] for image enhancement. We analyze and propose enhancement techniques that selectively smooth images while preserving either the multi-channel edges or the orientation-dependent texture features in them. Images are treated as manifolds in a feature-space. This geometrical interpretation lead to a general way for grey level, color, movies, volumetric medical data, and color-texture image enhancement.

We first review our framework in which the Polyakov action from high-energy physics is used to develop a minimization procedure through a geometric flow for images. Here we show that the geometric flow, based on manifold volume minimization, yields a novel enhancement procedure for color images. We apply the geometric framework and the general Beltrami flow to feature-preserving denoising of images in various spaces.

Next, we introduce a new method for color and texture enhancement. Motivated by Gabor's geometric image sharpening method [14], we present a geometric sharpening procedure for color images with texture. It is based on inverse diffusion across the multi-channel edge, and diffusion along the edge.

Keywords:*Scale-space, minimal surfaces, PDE based non-linear image diffusion, selective smoothing, color processing, texture enhancement, movies and volumetric medical data.*

*This work is supported in part by the US-Israel Binational Science Foundation, in part by the Applied Mathematics Subprogram of the Office of Energy Research under DE-AC03-76SFOOO98, ONR grant under NOOO14-96-1-0381, and partially by the National Science Foundation under grant PHY-90-21139.

1 Introduction

We extend in this paper the geometric framework presented in [45] both in scope and in applications. It is applied to volumetric images, movies, texture analysis and color images. We study in detail the structure of the different norms suggested for color processing and show that our area norm satisfies basic requirements of Lambertian color images. We show further that recent enhancement techniques of [52, 54] for grey level and color images are another example of a PDE technique that fits our framework and suggest a new enhancement technique that transcends Riemannian geometry.

The framework is based on geometrical ideas borrowed from high-energy physics. The essence of the method is summarized in two steps: (a) Representing an image as a Riemannian manifold embedded in a higher dimensional spatial-feature manifold. E.g. a three-dimensional manifold embedded in a four-dimensional space for volumetric medical images and movies. (b) A non-linear scale-space equation applied to images, derived as a gradient descent of a norm functional (the Polyakov action) that weights embedding maps in a geometrical way.

The explicit form of the scale-space PDE (or the coupled PDEs) depends on the choice of dynamic coordinates and the geometry of the image manifold, i.e., its metric. We work in this paper with Euclidean space-feature manifold. One can also choose to work with a non-Euclidean embedding space, see [47].

The importance of edges in scale space construction is obvious. We would like to build our minimization schemes such that boundaries between objects survive along the scale space, while homogeneous regions are simplified and flattened in a more rapid way. An important question, for which there are only partial answers, is how to treat multi-valued images in a geometric way. A color image is a good example since one actually considers 3 images (Red, Green, Blue) that are composed into one. The geometric framework attempts to answer this question. An edge-preserving enhancement procedure is a result of minimizing the Polyakov action norm with respect to the feature coordinates with the induced metric and is expressed via a geometric flow for images that we named *Beltrami flow*.

Texture plays an important role in the understanding process of many images, specially those that involve natural scenes. Therefore, it became an important research subject in the fields of psychophysics and computer vision. The study of texture starts from the pre-image that describes the physics and optics that transform the $3D$ world into an image. It tracks human perception from the image formation on the retina and its interpretation at the first perception steps in the brain.

Preserving the orientation information while diffusing a given texture image is important in certain cases, say in denoising a fingerprint image. We imagine a procedure that preserves domains of constant/homogeneous texture, enhances the texture in each domain, and thereby enhances the boundaries between neighboring domains with different textures. In this paper, we apply the geometrical framework to improve and enhance color and texture based images.

A popular method for texture analysis is to decompose a given image into a set of sub-band images using the $2D$ Gabor/Morlet-wavelet transform. Some nice mathematical properties and the relation of this transform to the physiological behavior were studied in [20, 33]. This model was later used for the segmentation, interpretation and analysis of texture [5, 21], and for texture-based browsing [26]. In Section 4, we use the Gabor/Morlet-wavelet transform to split a given image into a set of sub-band images. We then show that an enhancement procedure can be constructed based on a flow in the transformed space, i.e. the transform coefficients are treated as higher dimensional manifolds. Other flows in similar feature spaces were recently proposed in [39, 37, 7, 41, 57]; see also [54] for orientation preserving flows. These approaches begin with a flat metric, or singular structures [10], that do not necessarily yield a meaningful minimization process when going to more than one channel. The main difference between these schemes and the one we propose is the geometric interpretation of the information as a manifold flowing in a special way so as to minimize its volume. Our geometric perspective of a color image as a surface embedded in a higher dimensional space enabled us to define coupling in the multi-channel color space. Other schemes have also considered image as a surface [2, 13, 12, 59, 25], some even used the image information to build a Riemannian metric for segmentation [6, 15, 43]. However, these methods were not generalized to feature space or any co-dimension higher than one.

The geometric framework has the following properties: (1) It is a general way of writing the geometrical scale-space and enhancement algorithms for grey-scale, color, volumetric, time-varying, and texture images, (2) it unifies many existing partial differential equation based schemes for image processing, (3) it leads to feature-preserving schemes that are suitable for enhancement and segmentation tasks, and (4) it offers a general coupling between channels in a multi-channel image processing.

The remainder of this paper is organized as follows: We start with a short review of the geometrical framework in order to be as self contained as possible and to fix ideas and notations. The next sections are ordered according to increasing codimension. Section 3 deals with the simple cases of grey-level images, 3D images and movies which can be described as hypersurfaces i.e. embedding map with codimension 1. Texture analysis is the focus of Section 4 where the problem is formulated as parallel processing of different sub-bands each of which described as a 3-dimensional manifold embedded in a 5-dimensional space, i.e. a codimension 2 problem. We next move to a codimension 3 problem: Color. The understanding, processing and enhancement of colored images are an active fields of research for more then a century. We motivate the metric and the resulting edge-enhancing flow for color images via a simple color image formation model. In Section 6 we link the Beltrami flow to some recent coherence enhancing flows. Finally, we extend our geometric framework and introduce a new sharpening flow that is based on inverse diffusion across the edge. We finish with our concluding remarks.

2 The geometric framework

Let us first review the geometrical framework in which images are considered as Riemannian manifolds. We limit our discussion to variational methods in non-linear scale space image processing, and while our analysis in the following sections assumes Euclidean embedding space, we leave the discussion in this section in its full generality.

Suppose we have an n-dimensional manifold Σ with coordinates $\sigma^1, \sigma^2, \ldots, \sigma^n$ embedded in an m-dimensional manifold M with coordinates $X^1, X^2, \cdots, X^m$, where $m > n$. The embedding map $X : \Sigma \to M$ is given explicitly by the m functions of n variables

$$X : (\sigma^1, \ldots, \sigma^n) \to (X^1(\sigma^1, \ldots, \sigma^n), \ldots, X^m(\sigma^1, \ldots, \sigma^n)).$$

This map is an embedding if the map is an injection and the rank of the Jacobian is n.

If we denote the length of an image by σ^1 and the width by σ^2 then a possible embedding map of a grey-level image is

$$\left(X^1(\sigma^1,\sigma^2) = \sigma^1, X^2(\sigma^1,\sigma^2) = \sigma^2, X^3(\sigma^1,\sigma^2) = I(\sigma^1,\sigma^2)\right)$$

where $I(\sigma^1, \sigma^2)$ is the intensity. If we further denote $X^1 \equiv x$ and $X^2 \equiv y$ then it can be written with a slight abuse of notations as $(x, y, I(x, y))$.

Up to now we discussed coordinates only with no mention of the geometry of the manifolds. In order to do that we introduce Riemannian structure i.e. a metric (in some cases pseudo-Riemannian structure is needed). The metric at a given point on the manifold describes the way we measure distances without being dependent on the coordinates, i.e. the metric on Σ measures locally the distances at a point as follows

$$ds^2 = g_{\mu\nu} d\sigma^\mu d\sigma^\nu \qquad \mu, \nu \in \{1, \ldots, n\}$$

and summation is implied on identical indices. Similarly on M

$$ds^2 = h_{ij} dX^i dX^j \qquad i, j \in \{1, \ldots, m\}.$$

In an isometric embedding, i.e. one that preserves length these two line elements are equal, applying the chain rule $dX^i = \partial_\mu X^i d\sigma^\mu$, where $\partial_\mu \equiv \frac{\partial}{\partial \sigma^\mu}$ and there is a sum over μ, gives the induced metric formula

$$g_{\mu\nu} = h_{ij} \partial_\mu X^i \partial_\nu X^j.$$

For the embedding of a grey level image in a Euclidean 3-dimensional space we obtain the following metric

$$g_{\mu\nu} = \begin{pmatrix} 1 + I_x^2 & I_x I_y \\ I_x I_y & 1 + I_y^2 \end{pmatrix}.$$

Denote by (Σ, g) the image manifold and its metric and by (M, h) the space-feature manifold and its metric, then the Polyakov action [32] provides a convenient measure on the space of embedding maps $\mathbf{X} : \Sigma \to M$. It reads as follows

$$S[X^i, g_{\mu\nu}, h_{ij}] = \int d^m\sigma \sqrt{g} g^{\mu\nu} \partial_\mu X^i \partial_\nu X^j h_{ij}(\mathbf{X}), \qquad (1)$$

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. For more details see [46]. This is a natural generalization of the L2 norm to manifolds.

Many scale-space methods, linear and non-linear can be shown to be gradient descent flows of this functional with appropriately chosen metric of the image manifold. The gradient descent equation is

$$X_t^i = -\frac{1}{\sqrt{g}} \frac{\delta S}{\delta X^i},$$

where we limit ourselves above, and from now on to Euclidean embedding space.

The metric is a free "parameter" of the framework and different choices lead to different scale-space schemes as shown in [45]. For the choice of the metric as the induced metric the norm becomes simply the area or the volume of the image manifold, and the flow is towards a minimal surface solution. Minimizing the area action with respect to the feature coordinate (fixing the x and y coordinates), we obtain the area minimization direction given by applying the *second order differential operator of Beltrami* on the feature coordinates. Filtering the image based on this result, yields an efficient geometric flow for smoothing the image while preserving the edges. It is written as

$$\mathbf{I}_t = \Delta_g \mathbf{I}, \tag{2}$$

where, for color $\mathbf{I} = (R, G, B)$. The Beltrami operator, denoted by Δ_g, that is acting on $\mathbf{I}$ is a generalization of the Laplacian from flat spaces. It is defined by [1]

$$\Delta_g \mathbf{I} \equiv \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu \mathbf{I}). \tag{3}$$

For $2D$ images grey-level or color, the flow is given by

$$I_t^i = \frac{1}{g}\left(\partial_x p^i + \partial_y q^i\right) - \frac{1}{2g^2}\left(p^i(\partial_x g) + q^i(\partial_y g)\right) \tag{4}$$

where $g_{\mu\nu} = \delta_{\mu\nu} + \sum_i (\partial_\mu I^i)(\partial_\nu I^i)$, $g = g_{11}g_{22} - g_{12}^2$, and

$$p^i = g_{22}\partial_x I^i - g_{12}\partial_y I^i, \qquad \text{and} \qquad q^i = -g_{12}\partial_x I^i + g_{11}\partial_y I^i. \tag{5}$$

For the gray level case, the above evolution equation is the mean curvature flow of the image surface divided by the induced metric $g = \det(g_{\mu\nu})$. It is the evolution via the $\mathbf{I}$ components of the mean curvature vector $\mathbf{H}$. I.e. for the surface $(\mathbf{x}(\sigma_1, \sigma_2), \mathbf{I}(\sigma_1, \sigma_2))$ in the Euclidean space $(\mathbf{x}, \mathbf{I})$, the curvature vector is given by $\mathbf{H} = \Delta_g(\mathbf{x}(\sigma_1, \sigma_2), \mathbf{I}(\boldsymbol{\sigma_1}, \boldsymbol{\sigma_2}))$. If we identify $\mathbf{x}$ with σ then $\Delta_g I^i(\mathbf{x}) = \mathbf{H} \cdot \hat{I}^i$. Where, this direct computation applies for co-dimensions > 1. The determinant of the induced metric matrix g =det$(g_{\mu\nu})$ may be considered as a generalized form of an edge indicator. Therefore, the flow (2) is a selective smoothing mechanism that preserves edges and can be generalized to any dimension. In [46, 19], methods for constraining the evolution and the construction of convergent schemes based on the knowledge of the noise variance, were reported. Let us consider the simple gray level case, in which the image is considered as the surface $(x, y, I(x, y))$ in the (x, y, I) Euclidean space. If we set the aspect ratio such that $dI \gg dx$, then the principle curvatures can be approximated by the iso-contour curvature, and the flow line curvature which vanishes for this selection of the aspect ratio. Then, the mean curvature of the surface which is a sum of the principle curvatures may be approximated by the iso-contour curvature, and the Beltrami flow becomes the TV flow [38] up to a factor.

[1] In case the embedding space is chosen non-Euclidean there is an extra term. See [46].

3 Movies and Volumetric Medical Images

Traditionally, MRI volumetric data is referred to as 3D medical image. Following our framework, a more appropriate definition is of a 3D surface in 4D (x, y, z, I). In a very similar manner we will consider gray level movies as a 3D surfaces in 4D, where all we need to do is the mental exercise of replacing z of the volumetric medical images by the sequence (time) axis. In Fig. 1, the first row shows images at different z locations and the second row shows the corresponding denoised images. This is a relatively simple case, since now we have co-dimension equal to one.

The line element is

$$ds^2 = dx^2 + dy^2 + dz^2 + dI^2.$$

The induced metric in this case is given by

$$(g_{\mu\nu}) = \begin{pmatrix} 1+I_x^2 & I_xI_y & I_xI_z \\ I_xI_y & 1+I_y^2 & I_yI_z \\ I_xI_z & I_yI_z & 1+I_z^2 \end{pmatrix}, \tag{6}$$

and the Beltrami flow is:

$$I_t = \frac{1}{\sqrt{g}}\mathrm{div}\left(\frac{\nabla I}{\sqrt{g}}\right), \tag{7}$$

where now $\nabla I \equiv (I_x, I_y, I_z)$ and $g = 1 + I_x^2 + I_y^2 + I_z^2$.

The meaning of edge preserving in movies is as follows: In a shot where things stay more or less in the same place the algorithm will tend to flatten the boundaries i.e. it is an "anti shake" or "steady shot" filter. Yet it does not have an impact on an adjacent different scene since it preserves sharp changes along the time axis.

4 2D Gabor/Morlet-Wavelets as a Space for Texture Images

In this section, we apply the Beltrami flow in a decomposition space for the enhancement of texture images. In [20] Lee argues that the 2D Gabor/Morlet wavelet transform with specific coefficients is an appropriate mathematical description for images. He motivated his model by recent neurophysiological evidence based on experiments on the visual cortex of mammalian brains. These experiments indicate that a good model for the filter response of simple cells are self-similar 2D Gabor/Morlet wavelets, see also [29]. We refer the interested reader to [27] for implementation considerations, and to the rich literature on wavelet theory, e.g. [9].

Following Lee [20], we briefly describe the 2D Gabor/Morlet wavelets that model the simple cells while satisfying Daubechies' wavelet theory [9]. The 2D wavelet transform on an image $I(x, y)$, is defined as

$$(T^{wav}I)(x_0, y_0, \theta, a) = \|a\|^{-1}\int\int dxdyI(x,y)\psi_\theta\left(\frac{x-x_0}{a}, \frac{y-y_0}{a}\right), \tag{8}$$

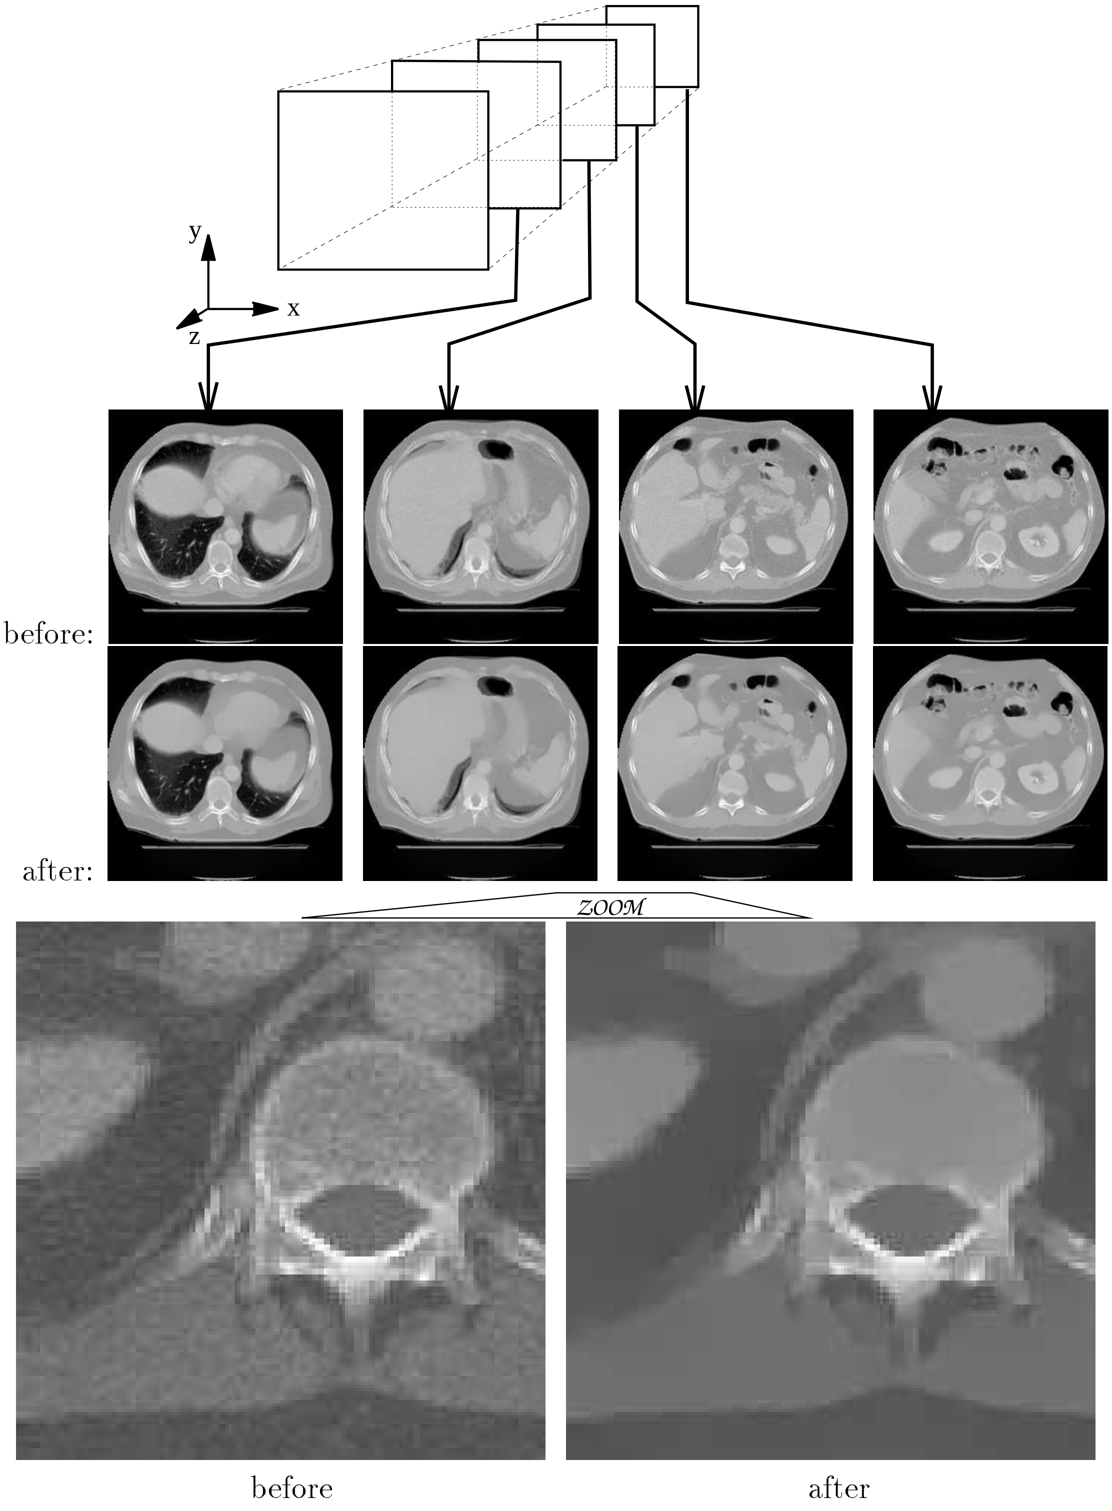

Figure 1: Movie or volumetric data; see text.

where a is a dilation parameter, x_0 and y_0 are the spatial translations, and θ is the wavelet orientation parameter.

$$\psi(x, y, x_0, y_0, \theta, a) = \|a\|^{-1}\psi_\theta(\frac{x - x_0}{a}, \frac{y - y_0}{a}), \tag{9}$$

is the 2D elementary wavelet function rotated by θ. Based on neurophysiological experiments, a specific Gabor elementary function is used as the *mother wavelet* to generate the 2D Gabor/Morlet wavelet family by convolving the image with

$$\psi(x, y) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{8}(4x^2+y^2)}(e^{ikx} - e^{-\frac{k^2}{2}}), \tag{10}$$

and $\psi_\theta(x, y) = \psi(\tilde{x}, \tilde{y})$ is defined by rotation of (x, y) via

$$\left\{ \begin{array}{rcl} \tilde{x} & = & x\cos\theta + y\sin\theta \\ \tilde{y} & = & -x\sin\theta + y\cos\theta. \end{array} \right. \tag{11}$$

The discretization of Eq. (8) is given by

$$W_{p,q,l,m} = (T^{wav}_{p,q,l,m} I) = a^{-m} \int\int dxdy I(x, y)\psi_{l\Delta\theta}(a^{-m}(x - p\Delta x), a^{-m}(y - q\Delta x)), \tag{12}$$

where Δx is the basic sampling interval, and the angles are given by $\Delta\theta = 2\pi l/L$, $l = 0, ..., L - 1$, L being the total number of orientations; p, q and m are integers determining the position and scaling. Note that as m increases the sample intervals get larger forming a pyramidal structure. Eq. 12 can be interpreted as a projection onto a discrete set of basis functions, namely

$$W_{p,q,l,m} = \langle I, \psi_{p,q,l,m} \rangle. \tag{13}$$

The real number k determines the frequency bandwidth of the filters in octaves via the approximation

$$k = \frac{a^\phi + 1}{a^\phi - 1}\sqrt{2\ln 2}, \tag{14}$$

where ϕ is the bandwidth in octaves, e.g. for $a = 2$ and $\phi = 1.5$ we get $k \approx 2.5$. In the above approximation the DC normalization term $e^{-k^2/2}$ that is required to make a wavelet basis out of the Gabor basis is ignored and we consider $a = k/\omega_0$. So the peaks of the scaled mother wavelets in the frequency domain are (approximately) at the locations $a^{-m}\omega_0$.

For our application we have chosen $L = 16$ (16 orientations), $a = 2$, $\Delta x = 1$, $k = 2.5$, and 5 scales, i.e. $m \in \{0, .., 4\}$. This selection results in a 'tight frame' [11] that allows simple summation reconstruction.

4.1 Beltrami Flow for Texture Enhancement

We denote the 2D Gabor/Morlet-wavelet transform as $W(x, y, \theta, \sigma)$, where for the discrete case $\sigma = a^m$ and $\theta = l\Delta\theta$. Let $R = \mathrm{Real}(W)$ and $J = \mathrm{Imag}(W)$ be its real and imaginary part. The response of a simple cell is then modeled by the projection of the image onto a specific Gabor/Morlet wavelet.

The Gabor/Morlet-wavelet transform of an image in our framework is a mapping $W : (x, y, \theta, \sigma) \to (x, y, \theta, \sigma, R, J)$, i.e. a $4D$ manifold embedded in $6D$. The Beltrami operator is not limited to act on gray level images (2D surfaces in 3D) as we show in Section 5 for color. First, the metric $g_{\mu\nu}$ is "pulled back" from the relevant arclength definition in the spatial-orientation complex space, namely

$$ds^2 = dx^2 + dy^2 + d\theta^2 + d\sigma^2 + dJ^2 + dR^2.$$

For practical implementation we consider each scale as a separate space. This is in contrast to writing the arclength for the full transform. Therefore, the arclength for a given scale σ is $ds^2 = dx^2 + dy^2 + d\theta^2 + dJ^2 + dR^2$, and the induced metric for each scale is given by

$$(g_{\mu\nu}) = \begin{pmatrix} 1 + R_x^2 + J_x^2 & R_x R_y + J_x J_y & R_x R_\theta + J_x J_\theta \\ R_x R_y + J_x J_y & 1 + R_y^2 + J_y^2 & R_y R_\theta + J_y J_\theta \\ R_x R_\theta + J_x J_\theta & R_y R_\theta + J_y J_\theta & 1 + R_\theta^2 + J_\theta^2 \end{pmatrix}. \tag{15}$$

As we have seen before, the above result can be understood from the arclength definition and applying the chain rule $dR = R_x dx + R_y dy + R_\theta d\theta$, and similarly for dJ to obtain the desired bilinear structure.

Finally, the area-minimizing and feature-preserving *Beltrami flow* that operates on the Gabor/Morlet-wavelet transform of a texture image can be compactly written as

$$\begin{aligned} R_t &= \Delta_g R \\ J_t &= \Delta_g J. \end{aligned} \tag{16}$$

The main difference from nonlinear diffusion in the image plane, is the freedom to work separately on each scale subspace. There is also a new meaning the term 'edge' in the decomposion space, since edge is now defined as the spatial difference in orientation rather than an explicit change in the gray level.

As a by product of the wavelet decomposition, at each scale σ we now have the complex function $W_\sigma(x, y, \theta) = R_\sigma(x, y, \theta) + iJ_\sigma(x, y, \theta)$. It defines a 3D manifold in the 5D space $(x, y, \theta, R_\sigma, J_\sigma)$. The extra coordinate θ that describes the behavior of the image along a specific direction enables us to smooth the image while keeping the meaningful orientation structure of the texture. Moreover, we have the freedom to apply different filters to the different scales. This enables us to preserve the nature of texture images by processing them only at significant scales. In other words, we can sharpen a specific scale without effecting the rest of the sub-band images. The first row is Fig. 2 presents the original image and the result of applying the Beltrami flow in the decomposition space to filter out non-oriented structures in a gray level image. More examples are shown in the second and third rows of Fig. 2.

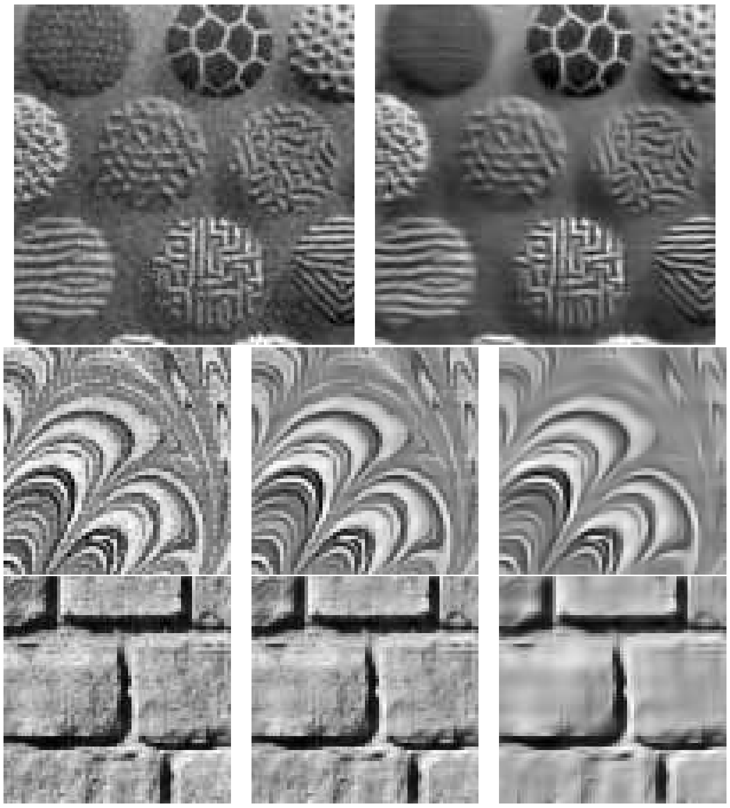

Figure 2: Top Row: Original image 128×128 is on the left, Result of Beltrami flow for 70 numerical iterations of each sub-scale in the decomposition space is on the right. Second and Third Rows: Two steps along the evolution for two different texture images, Left is the original image 64×64.

5 Color Processing

We show in this section that the geometric framework results in a meaningful operator for enhancing color images, following the presentation in [16]. The area functional, or "norm", captures the way we would like the smoothing process to act on the different color channels while exploring the coupling between them. Next, the steepest descent flow associated with the first variation of this functional is shown to be a proper selective smoothing filter for the color case. In this section we briefly review the geometric framework and justify the usage of the area norm and the Beltrami steepest descent flow in the color case. We list the requirements, compare to other recent norms, and relate to line element theories in color.

In [45, 46, 19], the geometrical framework for image diffusion was introduced. Minimizing the area of the image surface was claimed to yield a proper filter for color image enhancement. The area norm may serve for intermediate asymptotic analysis in low level vision, that is referred to as scale space in the computer vision community [36]. The norm may be coupled with variance constraints that are implemented via projection methods that were used for convergence based denoising [38] for image processing. Another popular option is to combine the norm with lower dimensional measures to create variational segmentation procedures, like the Mumford-Shah [35, 28, 18]. In this section we justify the usage of the area norm for color images obtained by the geometric framework and the Beltrami flow as an appropriate scale-space.

Here we limit our discussion to variational methods in non-linear scale space image processing, and to Euclidean color space. Given other significant groups of transformations in color, one could design the invariant flow with respect to that group based on the philosophy of images as surfaces in the hybrid space (x, y, R, G, B) through an arclength definition.

Let us first briefly review the geometric framework and the Beltrami flow and explore its relation to line element theory in color. Next we list the coupling requirements for the color case. A simple 'color image formation' model defines a 'proper' order of events for a desired enhancement. It is shown that this sequence of events is captured by the area norm.

5.1 The Geometric Framework and Color Processing

According to the geometric framework, images are considered as surfaces rather than functions. The area of the image surface minimized in a special way yields filters for texture, volume, movie, and color image enhancement;

Usually, a color image is considered as 3 images Red, Green, and Blue, that are composed into one. How should we treat such a composition? To answer this question, we view color images as *embedding maps*, that flow towards *minimal surfaces*. See [60] for a non-variational related effort.

At this point we would like to go back more than a hundred years, when physicists started to describe the human color perception as simple geometric space. Helmholtz [50] was the first to define a 'line element' (arclength) in color space. He first used a Euclidean R, G, B space defined by the arclength

$$ds^2 = (d\log R)^2 + (d\log G)^2 + (d\log B)^2. \tag{17}$$

His first model failed to represent empirical data of human color perception. Schrödinger [42] tried to improve Helmholtz's model by introducing the arclength

$$ds^2 = \frac{1}{l_R R + l_G G + l_B B}\left(\frac{l_R (dR)^2}{R} + \frac{l_G (dG)^2}{G} + \frac{l_B (dB)^2}{B}\right), \tag{18}$$

where l_R, l_G, l_B are constants. Schrödinger's model was later found to be inconsistent with findings on threshold data of color discrimination.

If we summarize the existing models for color space, we have two main cases: 1. The *inductive* line elements that derive the arclength by simple assumptions on the visual response mechanisms. For example, we can assume that the color space can be simplified and represented as a Riemannian space with zero Gaussian curvature, e.g. Helmholtz [50] or Stiles [49, 58] models. Another possibility for inductive line elements is to consider color arclengths like Schrödinger, or Vos-Walraven [51]. These models define color spaces with non-zero curvature ('effective' arclength). 2. The *empirical* line elements, in which the metric coefficients are determined to fit empirical data. Some of these models describe a Euclidean space like the CIELAB (CIE 1976 $(L^*a^*b^*)$) [58], recently used in [41]. Others, like MacAdam [23, 24], are based on an effective arclength.

The geometric framework is not limited to zero curvature spaces, and can incorporate any inductive or empirical color line element. See for example [47].

In case we want to perform any meaningful processing operation on a given image, we need to define a spatial relation between the points in the image plane $\mathbf{x}$. As a first step define the image plane to be Euclidean, which is a straightforward assumption for 2D images, that is:

$$ds_{\mathbf{x}}^2 = dx^2 + dy^2. \tag{19}$$

In order to construct a valuable geometric measure for color images we need to combine the spatial and color measures. The simplest combination of this hybrid spatial-color space is given by:

$$ds^2 = ds_{\mathbf{x}}^2 + \beta^2 ds_c^2. \tag{20}$$

The parameter β has dimensions [distance/intensity] and fixes the relative scale between the intensity of colors and the spatial distances. For a large β it defines a regularization of the color space.

Given the above arclength for color images, we pose the following question: How should a given image be simplified? In other words: What is the measure/norm/functional that is meaningful? What kind of variational method should be applied in this case?

The next geometrical measure after arclength is area. Minimization of area is a well known and studied physical phenomena. We will show, that for the right aspect ratio β, the area is a meaningful measure for our color case. Once the minimization measure is determined, one still needs to determine the parameterization for the steepest decent flow. A geometric flow for area minimization, that preserves edges is given by the Beltrami flow.

Let x and y be the *spatial* coordinates and the intensity R, G, B the *feature* coordinates, and describe color images as 2D surfaces in the 5D (x, y, R, G, B) space. The arclength is given by

$$ds^2 = dx^2 + dy^2 + dR^2 + dG^2 + dB^2. \tag{21}$$

As an introduction we have chosen the over simplified Euclidean color space, and for the time being assume $\beta = 1$. Next, we *pull back* the image surface *induced metric* from the arclength definition. By applying the chain rule $dR = R_x dx + R_y dy$, and rearranging terms, we obtain a distance measure on the surface defined via

$$ds^2 = g_{11}dx^2 + 2g_{12}dxdy + g_{22}dy^2,$$

where $g_{\mu\nu} = \delta_{\mu\nu} + \sum_i (\partial_\mu I^i)(\partial_\nu I^i)$ are the induced metric coefficients, $i \in \{1,2,3\}$ indicates the different color channels: $I^1 = R$, $I^2 = G$ and $I^3 = B$.

For the Euclidean color case with the induced metric, the norm is the area $\int d^2\sigma\sqrt{g}$. Here g is the determinant of the metric matrix $g = \det(g_{ij}) = g_{11}g_{22} - g_{12}^2$ given by its components $g_{\mu\nu} = \delta_{\mu\nu} + \sum_i (\partial_\mu I^i)(\partial_\nu I^i)$. If we multiply the intensities by a constant β, this action functional is given explicitly by

$$S = \int \sqrt{1 + \beta^2 \sum_i |\nabla I^i|^2 + \beta^4 \frac{1}{2} \sum_{ij} (\nabla I^i, \nabla I^j)^2} dxdy. \tag{22}$$

where $(\nabla R, \nabla G) \equiv R_x G_y - R_y G_x$ is the magnitude of the cross product of the vectors ∇R and ∇G. The action in Eq. (22) is the area of the image as a surface.

This functional obviously depends on the scalar β. For $\beta \gg 1$ it practically means mapping the intensity values that usually range between 0 and 255 to, let us say, $[0, 1000]$. Roughly speaking, for this limit of β, the order of events along the scale of the flow is as follows: First the different colors align together, then starts the selective smoothing geometric flow (similar to the single channel TV-L_1). On the other limit, where $\beta^2 \ll 1$, the smoothing tends to occur uniformly as a multi-channel heat equation (L_2).

5.2 Color Image Formation and Coupling Requirements

Let us elaborate on the selection of area as a proper measure for color images. The question we try to answer is how should we link between the different spectral channels. Let us assume that each color is 'equally important' and thus the measure we define should be symmetric. Within the scale space philosophy, we want the different spectral channels to get smoother in scale. This requirement leads to the minimization of the different color channels' gradient magnitudes combined in one way or another.

Next we argue that an important demand for color image processing is the alignment requirement of the different color channels. That is, we want the color channels to align together as they become smoother in scale. Figure 3 shows one level set of the Red and Green colors and their corresponding gradient vectors at one point along the level set. The requirement that the color channels align together as they evolve, amounts to minimizing the cross products between their gradient vectors.

A simplified color image formation model is a result of viewing Lambertian surface patches (not necessarily flat). Such a scene is a generalization of a 'Mondriaan world'. Each channel is considered as the projection of the real 3D world surface normal $\hat{\mathbf{N}}(\mathbf{x})$ onto the light source direction $\vec{l}$, multiplied by the albedo $\rho(x, y)$. The albedo captures the characteristics of the

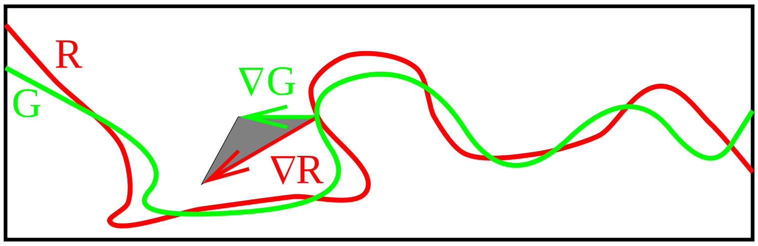

Figure 3: The cross product between ∇R and ∇G, $\frac{(\nabla G, \nabla R)}{2}$ displayed as the area of the gray triangle, measures the alignment between them.

3D object's material, and is different for each spectral channel. The 3 color channels may then be written as

$$I^i(\mathbf{x}) \;=\; \rho_i(\mathbf{x})\hat{\mathbf{N}}(\mathbf{x}) \cdot \vec{l}, \tag{23}$$

see Figure 4. Which means that the different colors capture the change in material via the

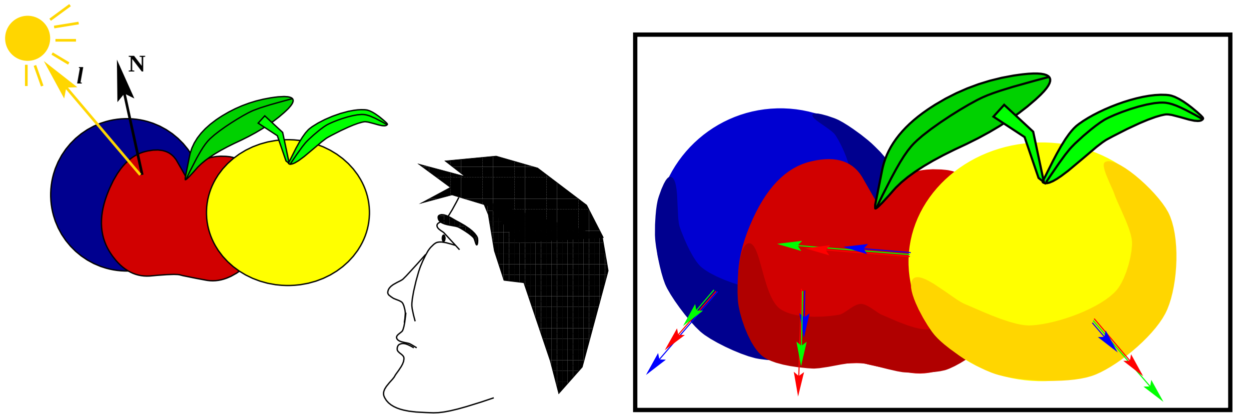

Figure 4: A simplified Lambertian color image formation model (left), leads to spectral channel alignment (right), see text.

albedo that multiplies the normalized shading image $\tilde{I}(\mathbf{x}) = \hat{\mathbf{N}}(\mathbf{x}) \cdot \vec{l}$.

Let us also assume that the material, and therefore the albedo, are the same within a given object in the image, e.g. $\rho_i(\mathbf{x}) = c_i$, where c_i is a given constant. The intensity gradient for each channel within a given object is then given by

$$\begin{aligned} \nabla I^i(\mathbf{x}) &= \tilde{I}(\mathbf{x})\nabla\rho_i(\mathbf{x}) + \rho_i(\mathbf{x})\nabla\tilde{I}(\mathbf{x}) \\ &= \tilde{I}(\mathbf{x})\nabla c_i + c_i\nabla\tilde{I}(\mathbf{x}) \\ &= c_i\nabla\tilde{I}(\mathbf{x}). \end{aligned} \tag{24}$$

Under the above assumptions, all color channels should have the same gradient direction within a given object. Moreover, the gradient direction should be orthogonal to the boundary for each color, since both the normalized shading image $\tilde{I}$ and the albedo ρ_i change across the boundaries. Our Lambertian world, without inter-reflections and specularities, is an oversimplified model for color image formation. Yet, its simplicity and locality direct us to a proper order of events we expect our local differential filters to follow. The Lambertian shading model is indeed the simplest image synthesis method in computer graphics. We can thus conclude that a first step in color processing should be the alignment of the colors so that their gradients agree. Only next should come the diffusion of all the colors simultaneously. As we show, the area minimization approach takes care for the gradients alignment, and solves the undesired twist between the channels.

For a large enough β, Eq. (22) follows exactly these requirements and the area norm is a regularization form of

$$\int \sqrt{\sum_i |\nabla I^i|^2 + \beta^2 \sum_{ij} (\nabla I^i, \nabla I^j)^2} dx dy, \tag{25}$$

that captures the order of events described above. For an even larger β it can be considered as a regularization of the *affine invariant* norm

$$\int \sqrt{\sum_{ij} (\nabla I^i, \nabla I^j)^2} dx dy. \tag{26}$$

If we also add the demand that edges should be preserved and search for the simplest geometric parameterization for the flow, we end up with the Beltrami flow as an appropriate selection.

Figure 5 shows snapshots from the Beltrami scale space in color for 3 images. Next, the flow is used to selectively smooth the JPEG compression distortions in Figure 6. Observe how the color perturbations are smoothed: The cross correlation between the colors holds the edges while selectively smoothing the uncorrelated data. In Figure 7, we deal with multiplicative noise. The Beltrami flow is used again as a denoising filter, now in the log domain to selectively smooth the image. The L_2 difference between the noisy and filtered images, is assumed to be known, and serves as a stopping criteria.

5.3 Previous norms for color Images

Let us review recent norms suggested for color processing. We start with two non-variational methods that will lead us to the variational norms: Chambolle [7], suggested a flow by the second derivative in the direction of minimal change with respect to the spectral channel with the largest gradient. Sapiro and Ringach [41] considered geometric diffusion in the direction of maximal L_2 change, see [52, 54] for a related effort. They used the eigenvalues, $\lambda_\pm$, of the matrix (though not a metric) $g_{\mu\nu} = \sum_i (\partial_\mu I^i)(\partial_\nu I^i)$ as a generalized edge detector to preserve edges.

In [39] Sapiro suggested to consider the variational method of the general form $\int f(\lambda_-, \lambda_+)$. Blomgren and Chan reported in [3, 4], that from the class of all possible norms of the form $f(\lambda_+, \lambda_-)$, the $f(\lambda_+ + \lambda_-)$ is the most natural one. This brings us to Shah's multi-channel model [44], that is based on the norm $\int \sqrt{\sum_{i=1} |\nabla I^i|^2}$ as part of a generalized Mumford-Shah functional.

Figure 5: Three snapshots along the scale space (left most is the original image). [This is a color figure]

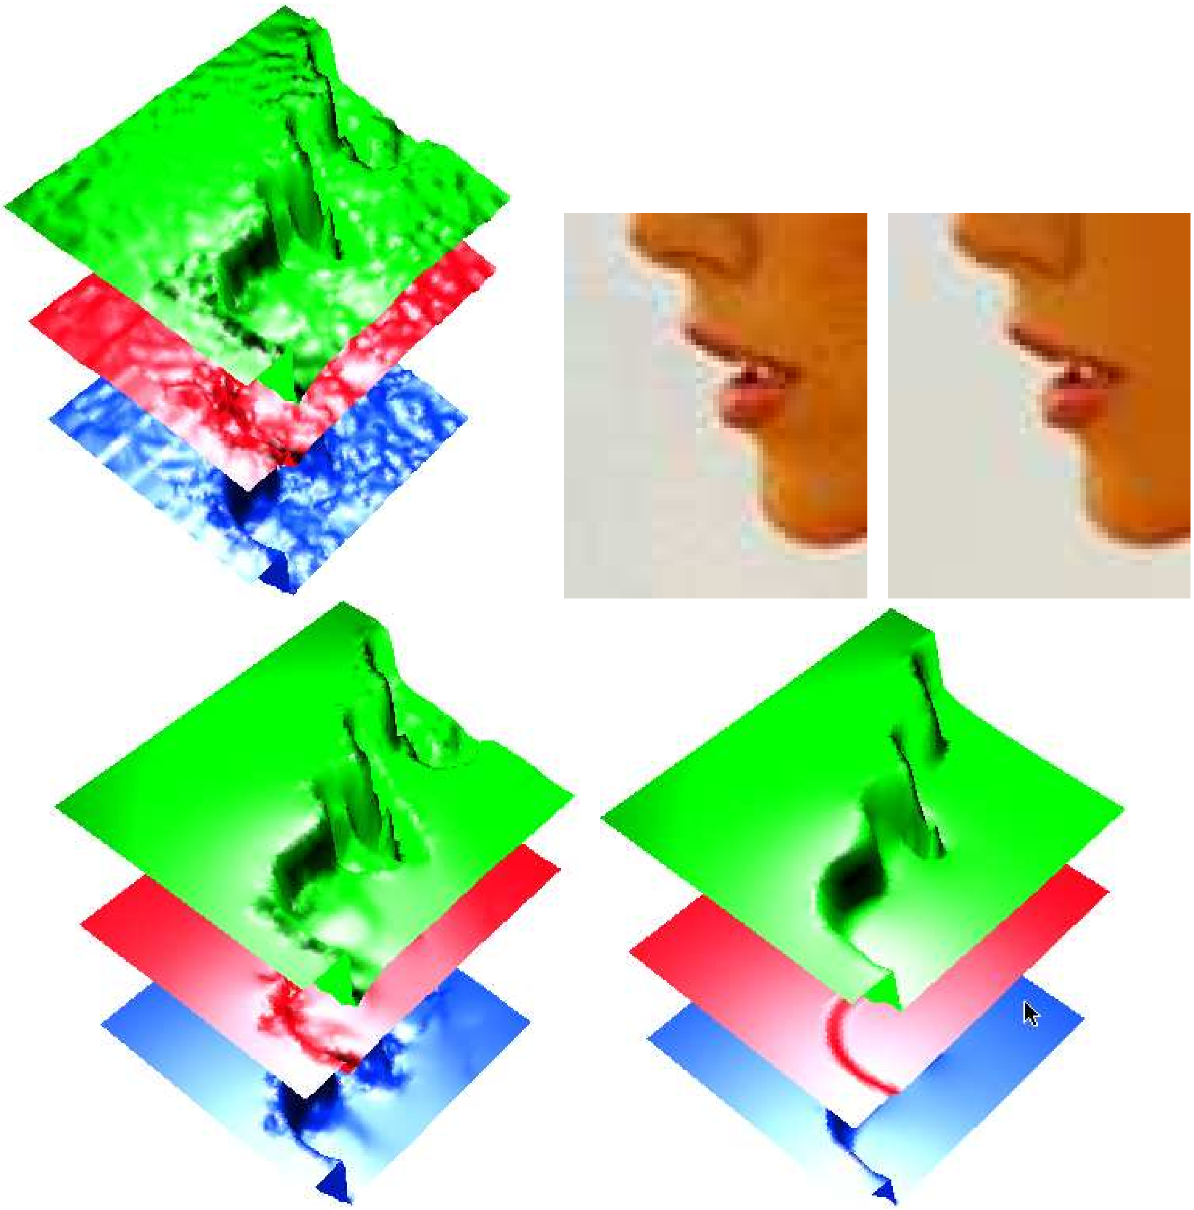

Figure 6: Three snapshots along the scale space for selectively smoothing JPEG lossy effects. The three channels are rendered as surfaces. The original image is on the left. [This is a color figure]

Figure 7: Original benchmark image 128×128 is on the left. A random noise, n, with zero mean, uniformly distributed over 10% of the $\log_2$ intensity scale is added to the image $\log_2$ intensity, $\log_2 \tilde{I} = \log_2 I + n$, which is a multiplicative noise. Result of Beltrami flow as a selective smoothing denoising in color is on the right.

Blomgren and Chan [3] defined a different "color TV" norm

$$\mathrm{TV}_m = \sqrt{\sum_{i=1}^{m} \left(\int |\nabla I^i| \right)^2},$$

with a constraint. In this case the coupling between the colors is only by the constraint. Actually, without the constraint the minimization yields a channel by channel curvature flow.

In order to preserve the edge and resolve color fluctuations one needs to use the cross alignment within the definition of the norm. While none of the previous norms included the cross-alignment terms in a proper way, the geometric framework of images as surfaces lead us to the norm that resolves the twist (torsion) between the channels via the cross-alignment term. We have thereby shown that the geometric framework yields a proper norm with respect to recent norms, and with respect to a list of objective requirements and considerations of color image formation. Next we apply the Beltrami operator to construct an orientation-preserving flow for texture images.

6 The Metric as a Structure Tensor

In [14, 22], Gabor considered an image enhancement procedure based on a single numerical step along a directional flow. It is based on the anisotropic flow via the <u>inverse</u> second directional derivative in the 'edge' direction (∇I direction) and the geometric heat equation (second derivative in the direction parallel to the edge). The same idea of steering the diffusion direction motivated many recent works[2]. Cottet and Germain [8] used a smoothed

[2]This definition of anisotropic flow differs from the Perona-Malik [31] framework, that is locally isotropic. See [34] for many interesting extensions and applications of the locally isotropic flow.

version of the image to direct the diffusion, while Weickert [55, 53] smoothed also the structure tensor $\nabla I \nabla I^T$ and then manipulated its eigenvalues to steer the smoothing direction. Eliminating one eigenvalue from a structure tensor, first proposed as a color tensor in [10], was used in in [41, 40], in which the tensors are not necessarily positive definite. While in [52, 54], the eigenvalues are manipulated to result in a positive definite tensor. See also [7], where the diffusion is in the direction perpendicular to the maximal gradient of the three color channels (this direction is different than that of [41]).

6.1 Relation to Weickert's coherence enhancement scheme

Motivated by these results we follow [17] and first link the anisotropic orientation diffusion (coherence enhancement [55]) to the geometric framework, and then invert the diffusion direction across the edge. Let us first show that the diffusion directions can be deduced from the smoothed metric coefficients $g_{\mu\nu}$ and may thus be included within the Beltrami framework under the right choice of directional diffusion coefficients.

The induced metric $(g_{\mu\nu})$ is a symmetric uniformly positive definite matrix that captures the geometry of the image surface. Let λ_1 and λ_2 be the largest and the smallest eigenvalues of $(g_{\mu\nu})$, respectively. Since $(g_{\mu\nu})$ is a symmetric positive matrix its corresponding eigenvectors u_1 and u_2 can be chosen orthonormal. The diagonalizing matrix be $U \equiv (u_1|u_2)$ is Hermitian, and $\Lambda \equiv \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}$, then we readily have the equality

$$(g_{\mu\nu}) = U \Lambda U^T. \tag{27}$$

Note also that

$$(g^{\mu\nu}) \equiv (g_{\mu\nu})^{-1} = U \Lambda^{-1} U^T = U \begin{pmatrix} 1/\lambda_1 & 0 \\ 0 & 1/\lambda_2 \end{pmatrix} U^T, \tag{28}$$

and that

$$g \equiv \det(g_{\mu\nu}) = \lambda_1 \lambda_2. \tag{29}$$

We will use the image metric as a control on the amount and direction of the diffusion, i.e., as a structure tensor. The coherence enhancement Beltrami flow $\mathbf{I}_t = \Delta_{\hat{g}} \mathbf{I}$ for color-texture images is then given as follows:

1. Compute the metric coefficients $g_{\mu\nu}$. For the N channel case (for color $N = 3$) we have

$$g_{\mu\nu} = \delta_{\mu\nu} + \sum_{k=1}^{N} (\partial_\mu I^k)(\partial_\nu I^k). \tag{30}$$

2. Diffuse the $g_{\mu\nu}$ coefficients by convolving with a Gaussian of variance ρ, thereby

$$\tilde{g}_{\mu\nu} = G_\rho * g_{\mu\nu}. \tag{31}$$

For $2D$ images $G_\rho = \frac{1}{\pi\rho^2} e^{-(x^2+y^2)/\rho^2}$.

3. Change the eigenvalues, λ_1, λ_2, $\lambda_1 > \lambda_2$, of $(\tilde{g}_{\mu\nu})$ so that $\lambda_1 = \alpha^{-1}$ and $\lambda_2 = \alpha$, for some given positive scalar $\alpha \ll 1$. This yields a new metric $\hat{g}_{\mu\nu}$ that is given by:

$$(\hat{g}_{\mu\nu}) = \tilde{U} \begin{pmatrix} \alpha^{-1} & 0 \\ 0 & \alpha \end{pmatrix} \tilde{U}^T = \tilde{U} \Lambda_\alpha \tilde{U}^T. \tag{32}$$

4. Evolve the k-th channel via Beltrami flow, that by the selection $\hat{g} \equiv \det(\hat{g}_{\mu\nu}) = \lambda_1 \lambda_2 = \alpha^{-1}\alpha = 1$ now reads

$$\begin{aligned} I^k_t &= \Delta_{\hat{g}} I^k \equiv \frac{1}{\sqrt{\hat{g}}} \partial_\mu \sqrt{\hat{g}} \hat{g}^{\mu\nu} \partial_\nu I^k = \partial_\mu \hat{g}^{\mu\nu} \partial_\nu I^k \\ &= \operatorname{div} \left(\tilde{U} \begin{pmatrix} \alpha & 0 \\ 0 & \alpha^{-1} \end{pmatrix} \tilde{U}^T \nabla I^k \right) = \operatorname{div} \left(\tilde{U} \Lambda_\alpha \tilde{U}^T \nabla I^k \right). \end{aligned} \tag{33}$$

Note again that both for gray level and color images the above flow is similar to the coherence-enhancing anisotropic diffusion with the important property of a uniformly positive definite diffusion tensor. For color images, $(g_{\mu\nu}) = \mathcal{I} + \sum_i \nabla I^i \nabla {I^i}^T$, where $\mathcal{I}$ is the identity matrix, and I^i are the color channels $((I^r, I^g, I^b) \equiv (I^1, I^2, I^3))$. In this case all that is done is the identity added to the structure tensors $\nabla I \nabla I^T$ for gray and $\sum_i \nabla I^i \nabla {I^i}^T$ for color. This addition does not change the eigenvectors and thus the above flow is equivalent to Weickert schemes [55, 53, 52, 54]. Next, we introduce a new inverse/direct diffusion model.

6.2 Inverse Diffusion Across the Edge

Let us take one step further, and exit our Riemannian framework by defining $(g_{\mu\nu})$ to be a non-singular symmetric matrix with one positive and one negative eigenvalues, i.e., a pseudo-Riemannian metric. That is, instead of a small diffusion we introduce an inverse diffusion across the edge. Here we extend Gabor's idea [14, 22] of inverting the diffusion along the gradient direction.

Inverting the heat equation is an inherently unstable process, see for example [48]. If we keep smoothing the metric coefficients, and apply the heat operator in the perpendicular direction we get a coherence-enhancing flow with sharper edges that is stable for a short duration of time. Obviously, the stability here is an empirical observation, since an inverse second order sharpening of a given sharp image yields visible artifacts. Nevertheless, a one numerical step of inverse diffusion is the classical sharpening technique in image processing.

Our idea is to change the sign of one of the modified eigenvalues in the algorithm described in the previous section, see Figure 8.

We change steps 3 and 4 of the previous scheme that now reads:

1. Compute the metric coefficients $g_{\mu\nu} = \delta_{\mu\nu} + \sum_{k=1}^{N} (\partial_\mu I^k)(\partial_\nu I^k)$.

2. Diffuse the $g_{\mu\nu}$ coefficients by convolving with a Gaussian of variance ρ.

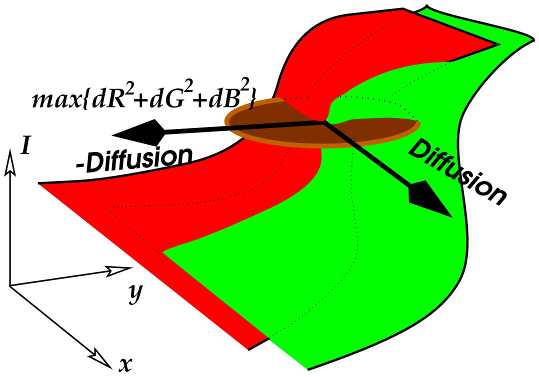

Figure 8: Motivated by the geometric framework and Gabor's sharpening algorithm we steer the diffusion directions and invert the diffusion direction across the edge. The edge direction is extracted by 'sensing' the multi-channel structure after smoothing the metric.

3. Change the eigenvalues of $(\tilde{g}_{\mu\nu})$ such that the largest eigenvalue λ_1 is now $\lambda_1 = -\alpha^{-1}$ and $\lambda_2 = \alpha$, for some given positive scalar $\alpha < 1$. This yields a new matrix $\hat{g}_{\mu\nu}$ that is given by:

$$(\hat{g}_{\mu\nu}) = \tilde{U}\begin{pmatrix} -\alpha^{-1} & 0 \\ 0 & \alpha \end{pmatrix}\tilde{U}^T = \tilde{U}\Lambda_\alpha\tilde{U}^T. \tag{34}$$

   We have used a single scalar α for simplicity of the presentation. Different eigenvalues can be chosen, one example are eigenvalues that depend on the original ones and bring us back to the Beltrami flow. By manipulating the eigenvalues we control the direction as well the intensity of the diffusion that can just as well be edge dependent. In this application the key idea is to modify the largest eigenvalue to be negative. This modification inverts the diffusion direction across the multi-spectral edge and thereby enhance it.

4. Evolve the k-th channel via the flow, that by the selection $|\hat{g}| \equiv |\det(\hat{g}_{\mu\nu})| = |\lambda_1\lambda_2| = |-\alpha^{-1}\alpha| = 1$, reads

$$\begin{aligned} I_t^k = \frac{1}{\sqrt{|\hat{g}|}}\partial_\mu\sqrt{|\hat{g}|}\hat{g}^{\mu\nu}\partial_\nu I^k &= \partial_\mu\hat{g}^{\mu\nu}\partial_\nu I^k \\ &= \mathrm{div}\left(\tilde{U}\begin{pmatrix} -\alpha & 0 \\ 0 & \alpha^{-1} \end{pmatrix}\tilde{U}^T\nabla I^k\right). \end{aligned} \tag{35}$$

For the gray level case with $\rho = 0$ it simplifies to highly unstable inverse heat equation. However, as ρ increases the smoothing along the edges becomes fundamental and the scheme is similar in its spirit to that of [14]. Different control methods can be applied. One example is an additional regularization term $\frac{\lambda}{2}(I(0) - I(t))^2$ that penalizes the departure from the original image, similar to the ideas explored in [48]. Another example is location dependent modification of the eigenvalues with local sensitivity to the image metric.

Gabor's [14] comment on the inverse diffusion operation in the gradient direction is that *'It is very similar to the operation which the human eye carries out automatically, and it is not surprising that even the first steps in imitating the human eye by mechanical means lead to rather complicated operations'.* It is important to note that the idea of stabilizing the inverse heat equation is extensively used in image processing. Exploring this area is beyond the scope of this paper. However, we like to refer the reader to the 'shock filters' introduced by Osher and Rudin in [30] for gray level images, and the extension of Alvarez and Mazorra [1] who apply geometrical inverse diffusion in the gradient direction combined with a directional smoothing in the orthogonal direction for gray level images.

6.3 Color Orientation-Enhancing Results

For completeness of the exposition we first repeat the gray level case as in [55, 53] and present an example of a fingerprint enhancement in gray level in Fig. 9.

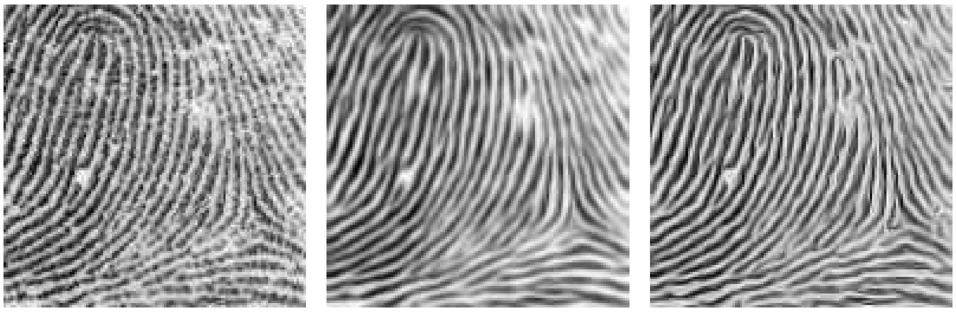

Figure 9: Left: Original fingerprint image 128×128. Middle: Result of the diffusion flow with smoothed metric ($\rho = 6$) and steered eigenvalues ($\alpha = 10^{-5}$), after 4 numerical iterations. Right: Result of the inverse/direct diffusion flow with smoothed metric ($\rho = 2$), steered eigenvalues and negative eigenvalue in the gradient direction (inverse/direct diffusion) after 4 numerical iterations ($\alpha = 0.55$).

In [56] the coherence enhancement flow was applied on several color masterpieces by van Gogh, which resulted in a 'coherence enhancement of expressionism'. In the next example we have chosen to 'enhance and sharpen impressionism'. We apply first the anisotropic oriented diffusion flow and then the new oriented diffusion along/inverse diffusion across the edge on a color painting by Claude Monet, see Fig. 10.

Next, we apply the color-oriented diffusion, and the oriented inverse/direct diffusion algorithms to a standard color-texture test image. Figure 11 compares again the flow with and without the inverse heat operator across the edge direction.

In the last example, we use the standard 'mandrill' color test image Fig. 12. Again, for comparison, the second row presents two steps along the color oriented diffusion flow.

Figure 10: Top: Original picture "Femme à l'ombrelle tournée vers la gauche," by Claude Monet 1875 ("woman with umbrella turning left") 521×784 (left), and the result of the inverse/direct diffusion flow ($\rho = 4$) after 8 numerical iterations (right). Bottom: Orientation preserving diffusion for 8 (left) and 128 (right) iterations.

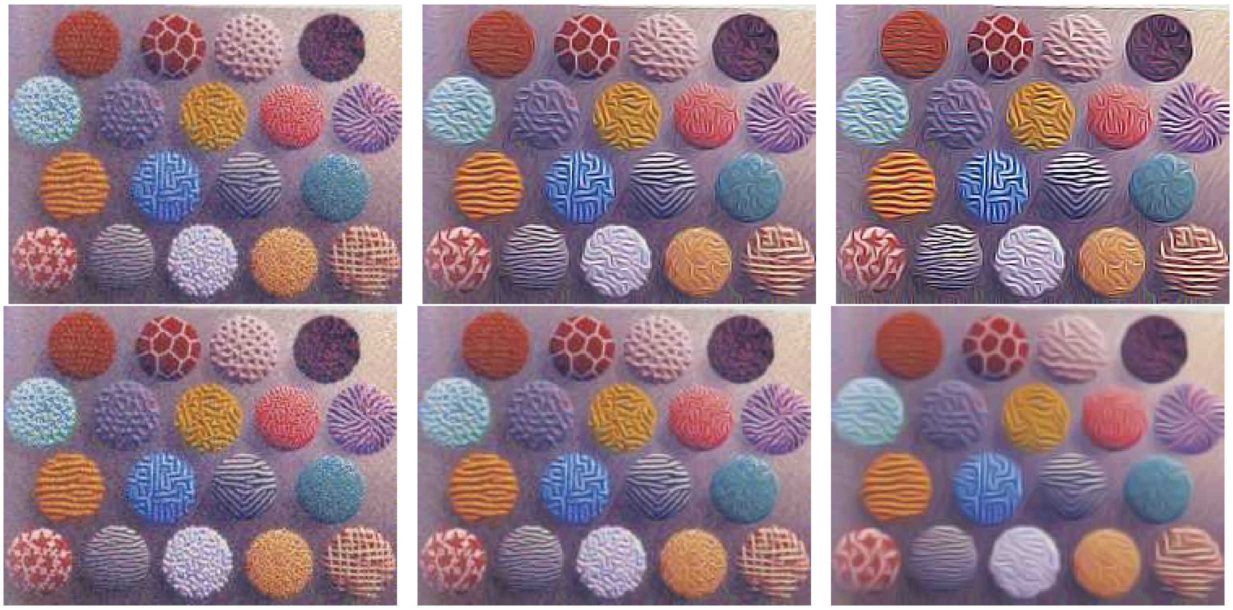

Figure 11: Color and texture: Top Row: Original 'Shells' image 242×184 (left), and the result of the color and texture inverse/direct diffusion. Flow for 4 (middle) and 8 (right) numerical iterations, $\alpha = 0.55$. Second Row: Original (left), and the result of the orientation preserving diffusion flow (smoothed metric and steered eigenvalues $\alpha = 10^{-5}$, $\rho = 2$) for 2 (middle) and 16 (right) numerical iterations.

7 Concluding Remarks

We applied the geometric framework and used it to design novel procedures for enhancement of color and texture images. These procedures are based on the interpretation of the image as a surface and a heat flow with respect to a given metric (Beltrami operator) as a filter.

We dealt with image enhancement and reconstruction of color and orientation based texture. These two different spaces were linked by a geometrical measure. The proposed filters align the color channels without un-coupling disturbances while enhancing the orientation based texture features and/or preserving the edges. Lee's [20] decomposition space was used for texture processing via the geometric framework.

We linked the geometric framework to recent color and texture enhancement algorithms and introduced a new sharpening procedure that extends the geometric framework. It is based on inverse diffusion across the edge for better sharpening results.

A direct application of the proposed methods is to enhance, selectively smooth, or sharpen color-texture and volumetric images. It can also be used to reduce the image entropy prior to compression and enhance its coherence in the reconstruction process, e.g. image restoration and denoising of lossy compression effects. It was shown that the geometrical framework can be applied to color, movies, and volumetric medical data, as well as non-trivial decomposition spaces.

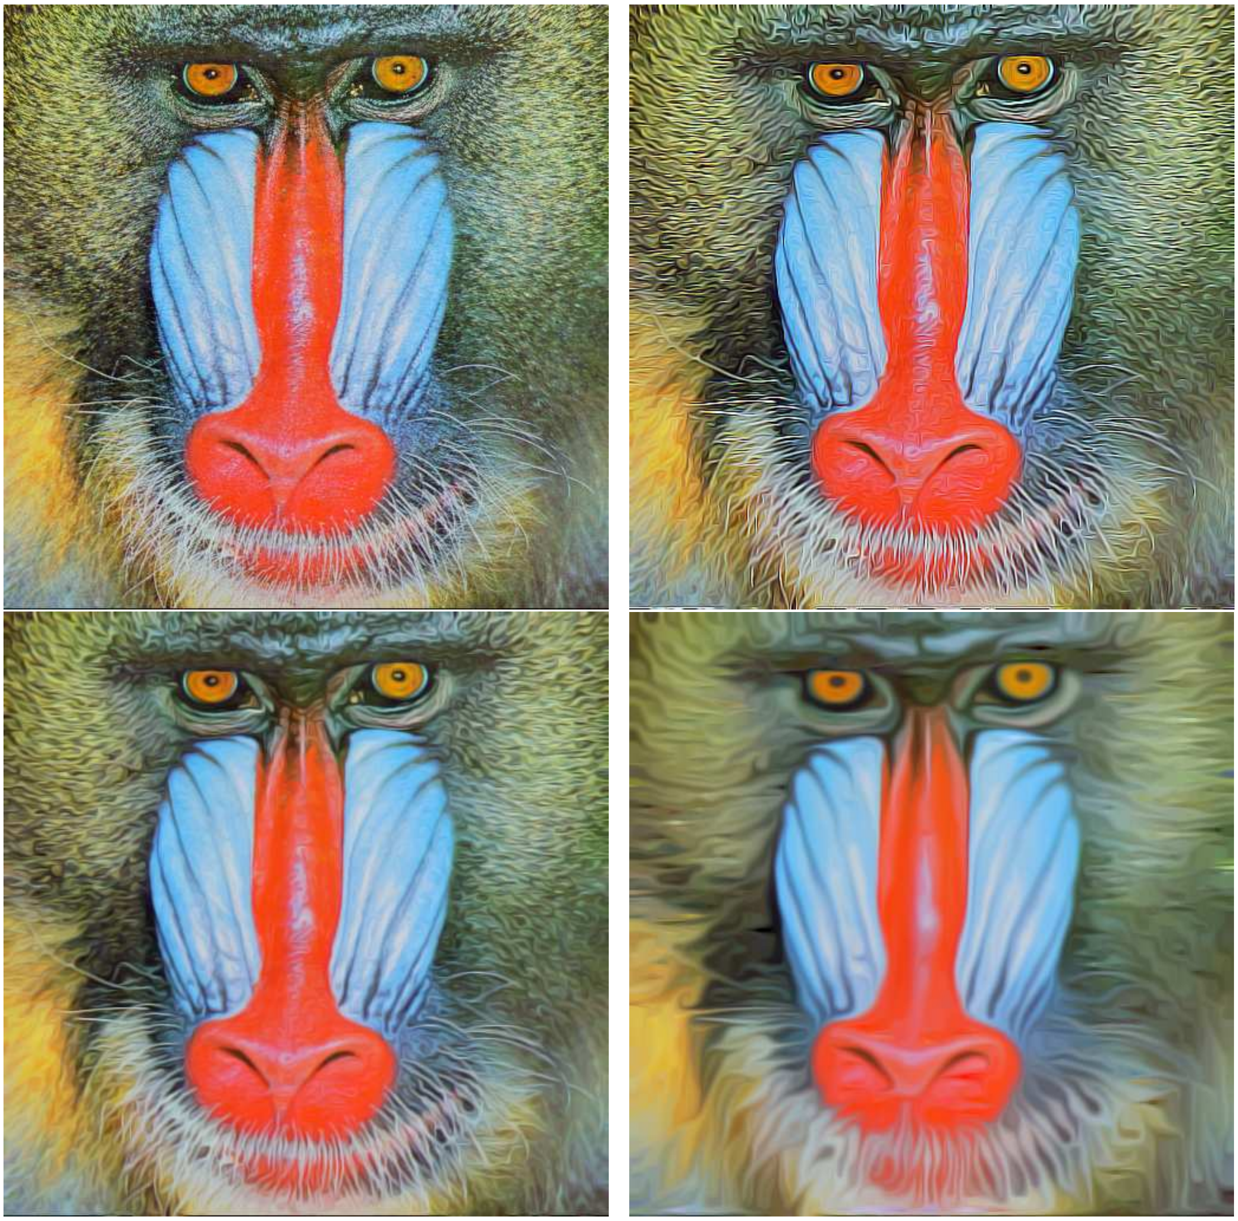

Figure 12: Color and texture: Diffusion flow with smoothed metric and steered eigenvalues ($\alpha = 10^{-5}$); Top Row: Original 'mandrill' image 512×512 (left), and the result of orientation-preserving flow and negative eigenvalue (inverse diffusion) in gradient direction, $\alpha = 0.39$. Second Row: Two steps along an orientation-preserving diffusion flow.

Acknowledgments

We thank Dr. David Adalsteinsson for his help in manipulating and rendering Figure 6 on his powerful Mac, Dr. Sherif Makram-Ebeid for interesting discussions in person and over the net, and Dr. Yacov Hel-Or for his benchmark images. We also thank the anonymous reviewers for their detailed comments that helped us improve the presentation.

References

[1] L Alvarez and L Mazorra. Signal and image restoration using shock filters and anisotropic diffusion. *SIAM J. Numer. Anal*, 31:590–605, 1994.

[2] A Blake and A Zisserman. *Visual Reconstruction.* MIT Press, Cambridge, Massachusetts, 1987.

[3] P Blomgren and T F Chan. Color TV: Total variation methods for restoration of vector valued images. cam TR, UCLA, 1996.

[4] P Blomgren and T F Chan. Color TV: Total variation methods for restoration of vector valued images. *IEEE Trans. on Image Processing*, 7(3):304–309, 1998.

[5] A C Bovik, M Clark, and W S Geisler. Multichannel texture analysis using localized spatial filters. *IEEE Trans. on PAMI*, 12(1):55–73, 1990.

[6] V Caselles, R Kimmel, and G Sapiro. Geodesic active contours. In *Proceedings ICCV'95*, pages 694–699, Boston, Massachusetts, June 1995.

[7] A Chambolle. Partial differential equations and image processing. In *Proceedings IEEE ICIP*, volume 1, pages 16–20, Austin, Texas, November 1994.

[8] G H Cottet and L Germain. Image processing through reaction combined with nonlinear diffusion. *Math. Comp.*, 61:659–673, 1993.

[9] I Daubechies. The wavelet transform, time frequency localization and signal analysis. *IEEE Trans. Information Theory*, 36(5):961–1004, 1990.

[10] S Di Zenzo. A note on the gradient of a multi image. *Computer Vision, Graphics, and Image Processing*, 33:116–125, 1986.

[11] R J Duffin and A C Schaeffer. A class of nonharmonic Fourier series. *Trans. Am. Math. Soc.*, 72:341–366, 1952.

[12] A I El-Fallah and G E Ford. On mean curvature diffusion in nonlinear image filtering. *Pattern Rec. Letters*, 19:433–437, 1998.

[13] 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.

[14] D Gabor. Information theory in electron microscopy. *Laboratory Investigation*, 14(6):801–807, 1965.

[15] S Kichenassamy, A Kumar, P Olver, A Tannenbaum, and A Yezzi. Gradient flows and geometric active contour models. In *Proceedings ICCV'95*, pages 810–815, Boston, Massachusetts, June 1995.

[16] R Kimmel. A natural norm for color processing. In *Proc. of 3-rd Asian Conf. on Computer Vision*, pages 88–95, Hong Kong, January 1998. Springer-Verlag, LNCS 1351.

[17] R Kimmel, R Malladi, and N Sochen. Image processing via the beltrami operator. In *Proc. of 3-rd Asian Conf. on Computer Vision*, pages 574–581, Hong Kong, January 1998. Springer-Verlag, LNCS 1351.

[18] R Kimmel and N Sochen. Geometric-variational approach for color image enhancement and segmentation. In *Proc. 2nd Int. Conference on Scale-Space Theories in Computer Vision Geometric Image Flows, Nonlinear Diffusion, Functional Minimisation, and Linear Scale-Space*, Corfu, Sept. 1999.

[19] R Kimmel, N Sochen, and R Malladi. From high energy physics to low level vision. In *Lecture Notes In Computer Science: First International Conference on Scale-Space Theory in Computer Vision*, volume 1252, pages 236–247. Springer-Verlag, 1997.

[20] T S Lee. Image representation using 2D Gabor-wavelets. *IEEE Trans. on PAMI*, 18(10):959–971, 1996.

[21] T S Lee, D Mumford, and A L Yuille. Texture segmentation by minimizing vector valued energy functionals: the couple-membrane model. In G Sandini, editor, *Lecture Notes in Computer Science, 588, Computer Vision: ECCV'92*, pages 165–173. Springer-Verlag, 1992.

[22] M Lindenbaum, M Fischer, and A M Bruckstein. On Gabor's contribution to image enhancement. *Pattern Recognition*, 27(1):1–8, 1994.

[23] D L MacAdam. Visual sensitivity to color differences in daylight. *J. Opt. Soc. Am.*, 32:247, 1942.

[24] D L MacAdam. Specification of small chromaticity differences. *J. Opt. Soc. Am.*, 33:18, 1943.

[25] 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.

[26] B S Manjunath and W Y Ma. Texture features for browsing and retrieval of image data. *IEEE Trans. on PAMI*, 18(8):837–841, 1996.

[27] T Masters. *Signal and image processing with neural networks: A C++ Sourcebook.* Wiley, New York, 1996.

[28] D Mumford and J Shah. Boundary detection by minimizing functionals. In *Proceedings of CVPR, Computer Vision and Pattern Recognition*, pages 22–26, San Francisco, 1985.

[29] B A Olshausen and D J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381:607–609, 1996.

[30] S J Osher and L I Rudin. Feature–oriented image enhancement using shock filters. *SIAM J. Numer. Analy.*, 27(4):919–940, August 1990.

[31] P Perona and J Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE-PAMI*, 12:629–639, 1990.

[32] A M Polyakov. Quantum geometry of bosonic strings. *Physics Letters B*, 103B(3):207–210, 1981.

[33] M Porat and Y Y Zeevi. The generalized Gabor scheme of image representation in biological and machine vision. *IEEE Trans. on PAMI*, 10(4):452–468, 1988.

[34] M Proesmans, E Pauwels, and L van Gool. Coupled geometry-driven diffusion equations for low level vision. In B M ter Haar Romeny, editor, *Geometric–Driven Diffusion in Computer Vision*, pages 191–228. Kluwer Academic Publishers, The Netherlands, 1994.

[35] 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.

[36] In B M ter Haar Romeny, editor, *Geometric–Driven Diffusion in Computer Vision*. Kluwer Academic Publishers, The Netherlands, 1994.

[37] Y Rubner and C Tomasi. Coalescing texture descriptors. In *Proceedings of the ARPA Image Understanding Workshop*, pages 927–936, February 1996.

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

[39] G Sapiro. Vector-valued active contours. In *Proceedings IEEE CVPR'96*, pages 680–685, 1996.

[40] G Sapiro and D Ringach. Anisotropic diffusion of multivalued images. In *12th Int. Conf. on Analysis and Optimization of Systems: Images, Wavelets and PDE'S*, volume 219 of *Lecture Notes in Control and Information Sciences*, pages 134–140, London, 1996. Springer.

[41] G Sapiro and D L Ringach. Anisotropic diffusion of multivalued images with applications to color filtering. *IEEE Trans. Image Proc.*, 5:1582–1586, 1996.

[42] E Schrödinger. Grundlinien einer theorie der farbenmetrik in tagessehen. *Ann. Physik*, 63:481, 1920.

[43] J Shah. A common framework for curve evolution, segmentation and anisotropic diffusion. In *Proceedings IEEE CVPR'96*, pages 136–142, 1996.

[44] J Shah. Curve evolution and segmentation functionals: Application to color images. In *Proceedings IEEE ICIP'96*, pages 461–464, 1996.

[45] 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. Presented in ONR workshop, UCLA, Sept. 5 1996.

[46] N Sochen, R Kimmel, and R Malladi. A general framework for low level vision. *IEEE Trans. on Image Processing*, 7(3):310–318, 1998.

[47] N Sochen and Y Y Zeevi. Images as manifolds embedded in a spatial feature non-Euclidean space. In *IEEE ICIP'98*, pages 166–170, Chicago, IL, 1998.

[48] A Steiner, R Kimmel, and A M Bruckstein. Shape enhancement and exaggeration. *Graphical Models and Image Processing*, 60(2):112–124, 1998.

[49] W S Stiles. A modified Helmholtz line element in brightness-colour space. *Proc. Phys. Soc. (London)*, 58:41, 1946.

[50] H Helmholtz von. *Handbuch der Psychologischen Optik*. Voss, Hamburg, 1896.

[51] J J Vos and P L Walraven. An analytical desription of the line element in the zone-fluctuation model of colour vision II. The derivative of the line element. *Vision Research*, 12:1345–1365, 1972.

[52] J Weickert. Scale-space properties of nonlinear diffusion filtering with diffusion tensor. Report no. 110, laboratory of technomathematics, University of Kaiserslautern, P.O. Box 3049, 67653 Kaiserslautern, Germany, 1994.

[53] J Weickert. Multiscale texture enhancement. In *Computer analysis of images and patterns; Lecture Notes in Computer Science*, volume 970, pages 230–237. Springer, 1995.

[54] J Weickert. Coherence-enhancing diffusion of colour images. In *Proc. VII National Symposium on Pattern Rec. and Image Analysis*, volume 1, pages 239–244, Barcelona, 1997.

[55] J Weickert. *Anisotropic Diffusion in Image Processing*. Teubner Stuttgart, 1998. ISBN 3-519-02606-6.

[56] J Weickert. Coherence-enhancing diffusion in color images. *Image and Vision Computing*, 17(3/4):199–210, 1999.

[57] R Whitaker and G Gerig. Vector-valued diffusion. In B M ter Haar Romeny, editor, *Geometric–Driven Diffusion in Computer Vision*, pages 93–134. Kluwer Academic Publishers, The Netherlands, 1994.

[58] G Wyszecki and W S Stiles. *Color Science: Concepts and Methods, Qualitative Data and Formulae, (2nd edition)*. John Wiley & Sons, 1982.

[59] S D Yanowitz and A M Bruckstein. A new method for image segmentation. *Computer Vision, Graphics, and Image Processing*, 46:82–95, 1989.

[60] A. Yezzi. Modified curvature motion for image smoothing and enhancement. *IEEE Trans. IP*, 7(3):345–352, 1998.