

Orientation Diffusion or How to Comb a Porcupine

Ron Kimmel

Computer Science Department, Technion—Israel Institute of Technology, Haifa 32000, Israel

and

Nir Sochen

Department of Applied Mathematics, University of Tel-Aviv, Tel-Aviv 69978, Israel

Received February 21, 2000; accepted November 2, 2001

This paper addresses the problem of feature enhancement in noisy images, when the feature is known to be constrained to a manifold. As an example, we approach the orientation denoising problem via the geometric Beltrami framework for image processing. The feature (orientation) field is represented accordingly as the embedding of a two dimensional surface in the spatial-feature manifold. The resulted Beltrami flow is a selective smoothing process that respects the feature constraint. Orientation diffusion is treated as a canonical example where the feature (orientation in this case) space is the unit circle S^1 . Applications to color analysis are discussed and numerical experiments demonstrate again the power of the Beltrami framework for nontrivial geometries in image processing. © 2002 Elsevier Science (USA)

1. INTRODUCTION

Feature enhancement is part of algorithms that play a major role in many applications. We are witnessing the emergence of a variety of methods for feature denoising that include and generalize the traditional image denoising. Here, we show how harmonic map methods, defined via the Beltrami framework, can be used to perform adaptive feature denoising.

We concentrate on the example of direction diffusion which gives us information on the preferred direction at every pixel. This example is useful in applications like texture and color analysis, and incorporates all the problems and characteristics of more involved feature spaces.

The input to the denoising process is a vector field on the image. The values of this vector field are in the unit circle S^1 . The given vector field is noisy and we wish to denoise it under the assumption that, for the class of images we are interested in, this vector field is piecewise smooth.¹

¹ See [2] for higher dimensional spheres S^n .

Two approaches for this problem are known: in a paper that first addresses directly this issue, Perona [6] uses a single parameter θ as an internal coordinate in S^1 , the second by Tang *et al.* [14, 15] embeds the unit circle S^1 in \mathbb{R}^2 (the sphere S^2 in \mathbb{R}^3) and work with the external coordinates. See also [16] for a related effort. The first approach is problematic because of the periodicity of S^1 . Averaging small angles around zero such as $\theta = \epsilon$ and $\theta = 2\pi - \epsilon$ leads to the erroneous conclusion that the average angle is $\theta = \pi$. Perona solved this problem by exponentiating the angle such that $V = e^{i\theta}$. This is actually the embedding of S^1 in \mathbb{C} which is isometric to \mathbb{R}^2 . This method is specific to two-dimensional embedding space where complex numbers can be used. The problem in using only one internal coordinate manifests itself in the numerical implementation of the PDE through the braking of rotation invariance. In the second approach we must make sure that we stay always in S^1 along the flow. This problem is known as the projection problem. It is solved in the continuum by adding a projection term. Chan and Shen [1] also use external coordinates with a projection term but suggest to use a Total Variation (TV) measure [8] in order to better preserve discontinuities in the vector field. This works well for the case where the codimension is one, like a circle in \mathbb{R}^2 . Yet it is difficult to generalize to higher codimensions like the sphere in \mathbb{R}^5 . Moreover, the flow of the external coordinates is difficult to control numerically since numerical errors should be projected on S^1 and no well-defined projection exist.

We propose a solution to these problems and introduce an adaptive smoothing process, which preserves orientation discontinuities. The proposed solutions work for all dimensions and codimensions, and overcome possible parameterization singularities by introducing several internal coordinates on different patches (charts) such that the union of the patches is S^n . Adaptive smoothness is achieved by the description of the vector field as a two-dimensional surface embedded in three- and four-dimensional spatial-feature manifold for the S^1 and S^2 cases, respectively. We treat here the S^1 case only due to space limitations.

The problem is formulated, in the Beltrami framework [3, 12] in terms of the embedding map

$$Y : (\Sigma, g) \rightarrow (M, h),$$

where Σ is the two-dimensional image manifold, and M , in the case of interest, is $\mathbb{R}^n \times S^1$ with $n = 2$ ($n = 4$) for gray-level (color) images. The key point is the choice of local coordinate systems for both manifolds²: The image manifold Σ (with metric g), and the embedding manifold M (with metric h). At the same time we should verify that the geometric filter does not depend on the specific choice of coordinates we make.

Once a local coordinate system is chosen for the embedding space and the optimization is done directly in these coordinates we can never leave M and avoid the problem of projection. The difficulty represented in the problem of projection is transformed to the problem of the choice of local coordinate system as we describe below. Other examples of enhancement by the Beltrami framework of non-flat feature spaces, like the color perceptual space and the derivatives vector field, can be found in [10, 13].

This paper is organized as follows: We review the Beltrami framework and point to the relation with harmonic maps in Section 2. We analyze the case S^1 in Section 3. Application for the enhancement of color images using perceptual color model is presented in Section 4. Section 5 presents experimental results and we offer our concluding remarks in Section 6.

² Note the difference with [1, 14, 15], where the image metric is flat.

2. THE BELTRAMI FRAMEWORK

Let us briefly review the Beltrami geometric framework for nonlinear diffusion in computer vision [12].

2.1. Representation and Riemannian Structure

We represent an image and other local features as an embedding map of a Riemannian manifold in a higher dimensional space. The simplest example is the image itself which is represented as a 2D surface embedded in \mathbb{R}^3 . We denote the map by $Y : \Sigma \rightarrow \mathbb{R}^3$. Where Σ is a two-dimensional surface, and we denote the local coordinates on it by (x^1, x^2) . The map Y is given in general by $(Y^1(x^1, x^2), Y^2(x^1, x^2), Y^3(x^1, x^2))$. We choose on this surface a Riemannian structure, namely, a metric. The metric is a positive definite and a symmetric 2-tensor that may be defined through the local distance measurements

$$ds^2 = g_{11}(dx^1)^2 + 2g_{12} dx^1 dx^2 + g_{22}(dx^2)^2. \quad (1)$$

We use below the Einstein summation convention in which the above equation reads $ds^2 = g_{\mu\nu} dx^\mu dx^\nu$ where repeated indices are summed over. We denote the inverse of the metric by $g^{\mu\nu}$.

2.2. Image Metric Selection: The Induced Metric

A reasonable assumption is that distances we measure in the embedding spatial-feature space, such as distances between pixels and difference between gray-levels, correspond directly to distances measured on the image manifold. This is the assumption of isometric embedding under which we can calculate the image metric in terms of the embedding maps Y^i and the embedding space metric h_{ij} . This follows directly from the fact that the length of infinitesimal distances on the manifold can be calculated in the manifold and in the embedding space with the same result. Formally, $ds^2 = g_{\mu\nu} dx^\mu dx^\nu = h_{ij} dY^i dY^j$. By the chain rule, $dY^i = \partial_\mu Y^i dx^\mu$, we get $ds^2 = g_{\mu\nu} dx^\mu dx^\nu = h_{ij} \partial_\mu Y^i \partial_\nu Y^j dx^\mu dx^\nu$. From which we have

$$g_{\mu\nu} = h_{ij} \partial_\mu Y^i \partial_\nu Y^j. \quad (2)$$

As an example we take the gray-level image as a two-dimensional image manifold embedded in the three dimensional Euclidean space \mathbb{R}^3 . The embedding maps are

$$(Y^1(x^1, x^2) = x^1, Y^2(x^1, x^2) = x^2, Y^3(x^1, x^2) = I(x^1, x^2)). \quad (3)$$

We choose to parameterize the image manifold by the canonical coordinate system $x^1 = x$ and $x^2 = y$. The embedding, by abuse of notation, is $(x, y, I(x, y))$. The induced metric g_{11} element is calculated as follows

$$g_{11} = h_{ij} \partial_{x^1} Y^i \partial_{x^1} Y^j = \delta_{ij} \partial_x Y^i \partial_x Y^j = \partial_x x \partial_x x + \partial_x y \partial_x y + \partial_x I \partial_x I = 1 + I_x^2. \quad (4)$$

Other elements are calculated in the same manner. The result is

$$(g_{\mu\nu}) = \begin{pmatrix} 1 + I_x^2 & I_x I_y \\ I_x I_y & 1 + I_y^2 \end{pmatrix}. \quad (5)$$

2.3. Polyakov Action: A Measure on the Space of Embedding Maps

Denote by (Σ, g) the image manifold and its metric and by (M, h) the space-feature manifold and its metric. Then the functional $S[\cdot, \cdot, \cdot]$ attaches a real number to a map $Y : \Sigma \rightarrow M$

$$S[Y^i, g_{\mu\nu}, h_{ij}] = \int dV \langle \nabla Y^i, \nabla Y^j \rangle_g h_{ij}, \quad (6)$$

where $dV = dx^1 dx^2 \cdots dx^m \sqrt{g}$ is a volume element and the scalar product $\langle \cdot, \cdot \rangle_g$ is defined with respect to the image metric, i.e., $\langle \nabla Y^1, \nabla Y^2 \rangle_g = g^{\mu\nu} \partial_\mu Y^1 \partial_\nu Y^2$. This functional, for $m = 2$ (two-dimensional image manifold) and $h_{ij} = \delta_{ij}$, was first proposed by Polyakov [5] in the context of high energy physics, and the theory known as *string theory*. Note that the image metric and the feature coordinates, i.e., intensity, color, orientation etc., are independent variables. The minimization of the functional with respect to the image metric can be solved analytically in the two-dimensional case (see for example [11]). The minimizer is the induced metric. If we choose, a-priory, the image metric induced from the metric of the embedding spatial-feature space M then the Polyakov action is reduced to an area (volume) of the image manifold:

$$S[Y^i, h_{ij}] = 2 \int dV = \int dx^1 dx^2 \sqrt{g} = 2 \int dx^1 dx^2 \sqrt{\det(\partial_\mu Y^i \partial_\nu Y^j h_{ij})}. \quad (7)$$

This follows from the form of the induced metric:

$$\langle \nabla Y^i, \nabla Y^j \rangle_g h_{ij} = g^{\mu\nu} \partial_\mu Y^i \partial_\nu Y^j h_{ij} = g^{\mu\nu} g_{\mu\nu},$$

and the identity

$$g^{\mu\nu} g_{\mu\nu} = \text{trace}(G^{-1}G^T) = \text{trace}(G^{-1}G) = \text{trace}(I_2) = 2. \quad (8)$$

Using standard methods in the calculus of variations (see [11]), the Euler–Lagrange equations with respect to the embedding are

$$-\frac{1}{2\sqrt{g}} h^{il} \frac{\delta S}{\delta Y^l} = \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu Y^i) + \Gamma_{jk}^i \langle \nabla Y^j, \nabla Y^k \rangle_g. \quad (9)$$

Since $(g_{\mu\nu})$ is positive definite, $g \equiv \det(g_{\mu\nu}) > 0$ for all x^μ . This factor is the simplest one that does not change the minimization solution while giving a reparameterization invariant expression. The operator that is acting on Y^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* [4], or for short *Beltrami operator*, and is denoted by Δ_g . The second term involves the Levi–Civita connection whose coefficients are given in terms of the metric

of the embedding space

$$\Gamma_{jk}^i = \frac{1}{2} h^{il} (\partial_j h_{lk} + \partial_k h_{jl} - \partial_l h_{jk}). \quad (10)$$

This is the term that takes into account the fact that the image surface flows in a non-Euclidean manifold and not in \mathbb{R}^n .

A map that satisfies the Euler–Lagrange equations

$$-\frac{1}{2\sqrt{g}} h^{il} \frac{\delta S}{\delta Y^l} = 0$$

is a *harmonic map*. The one- and two-dimensional examples are a geodesic curve on a manifold and a minimal surface.

The nonlinear diffusion or scale-space equation emerges as the gradient descent minimization flow

$$Y_t^i = \frac{\partial}{\partial t} Y^i = -\frac{1}{2\sqrt{g}} h^{il} \frac{\delta S}{\delta Y^l} = \Delta_g Y^i + \Gamma_{jk}^i \langle \nabla Y^j, \nabla Y^k \rangle_g. \quad (11)$$

This flow evolves a given surface toward a minimal surface, and in general it changes continuously a map towards a harmonic map.

Few major differences between this flow and those suggested in [1, 6, 14] are

- We use the induced metric while in other flows the image metric is flat.
- The coordinates Y^i are the local coordinates of the feature space while in the above mentioned flows they are coordinates of a THIRD manifold, i.e., \mathbb{R}^{n+1} in which the feature space S^n is embedded.
- The projection term is given as $A^i(\vec{Y})(h_{jk} \langle \nabla Y^j, \nabla Y^k \rangle_g)^{\frac{p}{2}}$, with $p = 1$ in [1] and $p = 2$ in [14]. Comparison with Eq. (11) show that for this to hold in the general case imposes a special form of the Levi–Civita connection, i.e., $\Gamma_{jk}^i = A^i h_{ij}$. This factorization is taking place in special cases such as S^n (see below and in [2] for a direct calculation of Γ_{jk}^i in this case). For the general case it doesn't hold and we have to use Eq. (10), as done for example in [10, 13].
- The flow Eq. (11) has a clear geometric meaning. It is a mean curvature flow projected on the feature space. This projection is an edge preserving operation [12].

3. BELTRAMI S^1 DIRECTION DIFFUSION

We are interested in attaching a unit vector field to an image. More precisely we would like to construct a nonlinear diffusion process that will enhance a given noisy vector field of this form while preserving the unit magnitude of each vector.

3.1. The Embedding Space Geometry

Denote the vector field by two components $(U, V)^T$ such that $U^2 + V^2 = 1$. This description is actually an extrinsic description. The unit circle S^1 is a one-dimensional curve and one parameter should suffice as an internal description. Since S^1 is a simply connected and compact manifold without boundaries, we need at least two coordinate systems to cover all the points of S^1 such that the transition function between the two patches is infinitely

differentiable at all points that belong to the intersection of the two patches. We define the two patches as follows: The coordinate system on $S^1 - \{(\pm 1, 0)\}$ is U , with the induced metric

$$ds_{S^1}^2 = dU^2 + dV^2 = dU^2 + (d(\sqrt{1 - U^2}))^2 = \frac{1}{1 - U^2} dU^2. \tag{12}$$

The coordinate system on $S^1 - \{(0, \pm 1)\}$ is V with the induced metric

$$ds_{S^1}^2 = dU^2 + dV^2 = (d(\sqrt{1 - V^2}))^2 + dV^2 = \frac{1}{1 - V^2} dV^2. \tag{13}$$

It is clear the the transformations $V(U) = \sqrt{1 - U^2}$ and $U(V)$ are differentiable anywhere on the intersection $S^1 - \{(\pm 1, 0), (0, \pm 1)\}$.

The embedding is of the two-dimensional image manifold in the three-dimensional space $\mathbb{R}^2 \times S^1$. The canonical embedding for the first patch is $(Y^1(x, y) = x, Y^2(x, y) = y, Y^3(x, y) = U(x, y))$, and for the second patch is $(Y^1(x, y) = x, Y^2(x, y) = y, Y^3(x, y) = V(x, y))$.

3.2. The S^1 Beltrami Operator

The line element on the image manifold is

$$ds^2 = ds_{\mathbb{R}^2}^2 + ds_{S^1}^2 = dx^2 + dy^2 + \frac{1}{1 - U^2} dU^2, \tag{14}$$

and by the chain rule

$$ds^2 = (1 + A(U)U_x^2) dx^2 + 2A(U)U_xU_y dx dy + (1 + A(U)U_y^2) dy^2, \tag{15}$$

where $A(U) = 1/(1 - U^2)$, and similarly for V .

The induced metric is therefore

$$(g_{\mu\nu}) = \begin{pmatrix} 1 + A(U)U_x^2 & A(U)U_xU_y \\ A(U)U_xU_y & 1 + A(U)U_y^2 \end{pmatrix}, \tag{16}$$

and the Beltrami operator acting on U is $\Delta_g U = \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu U)$, where $g = 1 + A(U)(U_x^2 + U_y^2)$ is the determinant of $(g_{\mu\nu})$, and $(g^{\mu\nu})$ is the inverse matrix of $(g_{\mu\nu})$.

3.3. The Levi-Civita Connection

Since the embedding space is non-Euclidean we have to calculate the Levi-Civita connection. Remember that the metric of the embedding space is

$$(h_{ij}) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & A(U) \end{pmatrix}. \tag{17}$$

The Levi–Civita connection coefficients are given by the fundamental theorem of Riemannian geometry in the following formula $\Gamma_{jk}^i = \frac{1}{2}h^{il}(\partial_j h_{lk} + \partial_k h_{jl} - \partial_l h_{jk})$, where the derivatives are taken with respect to Y^i for $i = 1, 2, 3$.

The only nonvanishing term is Γ_{33}^3 that reads $\Gamma_{33}^3 = \frac{1}{2A(U)}\partial_U(A(U)) = \frac{U}{1-U^2} = Uh_{33}$.

The second term in the EL equations in this case reads $Uh_{33}\|\nabla U\|_g^2$. We can rewrite this expression using the following identities

$$\begin{aligned} h_{33}\|\nabla U\|_g^2 &= (h_{11}g^{11} + h_{22}g^{22} + h_{33}\partial_\mu U \partial_\nu U g^{\mu\nu}) - (h_{11}g^{11} + h_{22}g^{22}) \\ &= g_{\mu\nu}g^{\mu\nu} - (g^{11} + g^{22}) = 2 - \frac{1}{g}(g_{11} + g_{22}) = 2 - \frac{1}{g}(1 + g), \end{aligned} \quad (18)$$

where we used the induced metric identity Eq. (2), and the identity Eq. (8), in order to rewrite

$$2 = \text{trace} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = g_{\mu\nu}g^{\mu\nu} = h_{11}g^{11} + h_{22}g^{22} + h_{33}\partial_\mu U \partial_\nu U g^{\mu\nu}. \quad (19)$$

3.4. The Flow and the Switches

The Beltrami flow is

$$Y_t^i = \Delta_g Y^i + \Gamma_{jk}^i(Y^1, Y^2, Y^3)\langle \nabla Y^j, \nabla Y^k \rangle_g, \quad (20)$$

for $i = 3$, and similarly for the other coordinate charts. Gathering together all the pieces we finally get the Beltrami flow

$$\begin{aligned} U_t &= \Delta_g U + U \frac{g-1}{g} \\ V_t &= \Delta_g V + V \frac{g-1}{g}. \end{aligned} \quad (21)$$

In the implementation we compute the diffusion for U and V simultaneously and take the values $(U, \text{sign}(V)\sqrt{1-U^2})$ for the range $U^2 \leq V^2$, and the values $(\text{sign}(U)\sqrt{1-V^2}, V)$ for the range $V^2 \leq U^2$.

4. COLOR DIFFUSION

There are many coordinate systems and models of color space which try to be closer to human color perception. One of the popular coordinate systems is the HSV system [9]. In this system, color is characterized by the Hue, Saturation, and Value. The Saturation and Value take their value in \mathbb{R}_+ , while the Hue is an angle that parameterizes S^1 .

In order to denoise and enhance color images by a nonlinear diffusion process which is adapted to human perception we use here the HSV system. We need to have special treatment of the Hue coordinate in the lines of Section 3.

Let us represent the image as a mapping $\mathbf{Y}: \Sigma \rightarrow \mathbb{R}^4 \times S^1$ where Σ is the two-dimensional image surface and $\mathbb{R}^4 \times S^1$ is parameterized by the coordinates (x, y, H, S, V) . As mentioned above, a diffusion process in this coordinate system is problematic. We define therefore two coordinates

$$U = \cos H; \quad W = \sin H$$

and continue in a similar way to Section 3. The metric of $\mathbb{R}^4 \times S^1$ on the patch where U parameterizes S^1 and $W(U)$ is nonsingular is

$$h_{ij} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & A(U) & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \tag{22}$$

where $A(U) = 1/(1 - U^2)$.

The induced metric is therefore

$$\begin{aligned} ds^2 &= dx^2 + dy^2 + A(U) dU^2 + dS^2 + dV^2 \\ &= dx^2 + dy^2 + A(U)(U_x dx + U_y dy)^2 + (S_x dx + S_y dy)^2 + (V_x dx + V_y dy)^2 \\ &= (1 + A(U)U_x^2 + S_x^2 + V_x^2) dx^2 + 2(A(U)U_x U_y + S_x S_y + V_x V_y) dx dy \\ &\quad + (1 + A(U)U_y^2 + S_y^2 + V_y^2) dy^2. \end{aligned} \tag{23}$$

Similar expressions are obtained on the other dual patch.

The only nonvanishing Levi–Civita connection’s coefficient is $\Gamma_{33}^3 = U h_{33}$. The resulting flow is

$$\begin{aligned} U_t &= \Delta_g U + 2U - U(g^{11} + g^{22}) \\ W_t &= \Delta_g W + 2W - W(g^{11} + g^{22}) \\ S_t &= \Delta_g S \\ V_t &= \Delta_g V. \end{aligned} \tag{24}$$

Note that the switch between U and W should be applied not only to the U and W equations but also to the S and V evolution equations where, at each point, one needs to work with the metric that is defined on one of the patches.

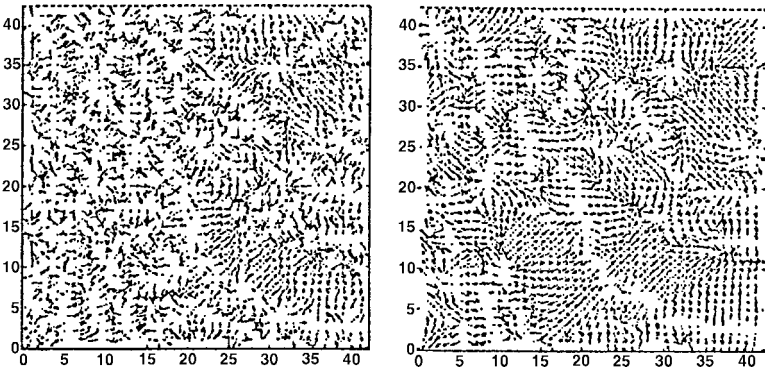


FIG. 1. Two vector fields (left) before and (right) after the flow on S^1 .




FIG. 2. The HSV color model captures human color perception better than the RGB model which is the common way our machines represent colors. The original image (left), the noisy image (middle) and the filtered image (right) demonstrate the effect of the flow as a denoising filter in the HSV color space. For further examples follow the links at <http://www.cs.technion.ac.il/~ron/pub.html>.

5. EXPERIMENTAL RESULTS

Our first example deals with the gradient direction flow via the Beltrami framework. Figure 1 shows a vector field before and after the application of the flow for a given evolution time. The normalized gradient vector field extracted from an image is presented before and after the flow and show the way the field flows into a new smooth orientation transactions field.

Next, we explore a popular model that captures some of our color perception. The HSV (hue, saturation, value) model proposed in [9] is often used as a “user oriented” color model, rather than the RGB “machine-oriented” model.

Figure 2 shows the classical representation of the HSV color space, in which the hue is measured as an angle, while the value (sometimes referred to as brightness) and the color saturation are mapped onto finite non-periodic intervals. This model lands itself into a filter that operates on the spatial x, y coordinates, the value and saturation coordinates, and the hue periodic variable. Our image is now embedded in $\mathbb{R}^4 \times S^1$.

6. CONCLUDING REMARKS

There are two important issues in the process of denoising a constrained feature field. The first is to make the process compatible with the constraint in such a way that the latter is never violated along the flow. The second is the type of regularization which is applied in order to preserve significant discontinuities of the feature field while removing noise.

These two issues are treated in this paper via the Beltrami framework. First a Riemannian structure, i.e., a metric, is introduced on the feature manifold and a several local coordinate systems are chosen to describe *intrinsically* the constrained feature manifold. The diffusion process acts on these coordinates and the compatibility with the constraint is achieved through the intrinsic nature of the coordinate system. The difficulty in working on a

non-Euclidean space transforms itself to the need to locally choose the best coordinate system to work with.

A preservation of significant discontinuities is dealt with by using the induced metric and the corresponding Laplace–Beltrami operator acting on feature coordinates only. This operation is in fact a projection of the mean curvature, in the normal(s) direction(s) to the surface, to the feature direction(s). This projection slows the diffusion process along significant (supported) discontinuities while letting the process in the homogeneous regions proceed in a normal speed.

The result of this algorithm is an adaptive smoothing process for a constrained feature space in every dimension and codimension. As examples we showed how our geometric model coupled with a proper choice of charts handles the orientation diffusion problem. This is a new application of the Beltrami framework proposed in [11]. We tested the new model on vector fields restricted to the unit circle S^1 , and hybrid spaces like the HSV color space. The integration of the spatial coordinates with the color coordinates yield a selective smoothing filter for images in which some of the coordinates are restricted to a circle.

ACKNOWLEDGMENTS

We thank Alfred Bruckstein from the Technion Israel for stimulating discussions on diffusion and averaging and color analysis. We also thank Guillermo Sapiro from University of Minnesota for sharing with us his ideas and results on direction diffusion.

REFERENCES

1. T. Chan and J. Shen, Variational restoration of non-flat image features: Models and algorithms, *SIAM J. Appl. Math.* **61**, 2000, 1338–1361.
2. R. Kimmel and N. Sochen, *How to Comb a Porcupine?*, ISL technical report, Technion Tel-Aviv University, Israel, June 2000.
3. R. Kimmel, R. Malladi, and N. Sochen, Images as embedding maps and minimal surfaces: Movies, color, texture, and volumetric medical images, *Int. J. Comput. Vision* **39**, 2000, 111–119.
4. E. Kreyszing, *Differential Geometry*, Dover, New York, 1991.
5. A. M. Polyakov, Quantum geometry of bosonic strings, *Phys. Lett. B* **103**, 1981, 207–210.
6. P. Perona, Orientation diffusion, *IEEE Trans. Image Process.* **7**, 1998, 457–467.
7. B. M. ter Haar Romeny (Ed.), *Geometry Driven Diffusion in Computer Vision*, Kluwer Academic, Dordrecht/Norwell, MA, 1994.
8. L. Rudin, S. Osher, and E. Fatemi, Nonlinear total variation based noise removal algorithms, *Physica D* **60**, 1991, 259–268.
9. A. R. Smith, Color gamut transform pairs, *SIGGRAPH'79*, pp. 12–19.
10. N. Sochen, R. M. Haralick, and Y. Y. Zeevi, A geometric functional for derivatives approximation, *Proceedings of Scale-Space Theories in Computer Vision, Corfu, 1999*, pp. 507–512.
11. N. Sochen, R. Kimmel, and R. Malladi, *From High Energy Physics to Low Level Vision*, Report, LBNL, UC Berkeley, LBNL 39243, August, Presented in ONR workshop, UCLA, Sept. 5, 1996.
12. N. Sochen, R. Kimmel, and R. Malladi, A general framework for low level vision, *IEEE Trans. Image Process.* **7**, 1998, 310–318.
13. N. Sochen and Y. Y. Zeevi, Representation of colored images by manifolds embedded in higher dimensional non-Euclidean space, *IEEE ICIP'98*, Chicago, 1998.
14. B. Tang, G. Sapiro, and V. Caselles, Direction diffusion, *International Conference on Computer Vision*, 1999.
15. B. Tang, G. Sapiro, and V. Caselles, *Color Image Enhancement via Chromaticity Diffusion*, Technical report, ECE-University of Minnesota, 1999.

16. J. Weickert, Coherence-enhancing diffusion in colour images, *Proc. of VII National Symposium on Pattern Rec. and Image Analysis, Barcelona, Vol. 1*, pp. 239–244, 1997.
-

RON KIMMEL received the B.Sc. (with honors) in computer engineering in 1986, the M.S. in 1993 in electrical engineering, and the D.Sc. in 1995 from the Technion—Israel Institute of Technology. During the years 1986–1991 he served as an R&D officer in the Israeli Air Force. During the years 1995–1998 he has been a postdoctoral fellow at the Computer Science Division of Berkeley Labs and the Mathematics Department, University of California, Berkeley. Since 1998, he has been a faculty member of the Computer Science Department at the Technion, Israel, where he currently is an associate professor. His research interests are in computational methods and their applications in differential geometry, numerical analysis, image processing and analysis, computer aided design, robotic navigation, and computer graphics. Dr. Kimmel was awarded the Hershel Rich Technion Innovation Award, the Henry Taub Prize for excellence in research, the Alon Fellowship, the HTI Postdoctoral Fellowship, and the Wolf, Gutwirth, Ollendorff, and Jury fellowships. He has been a consultant at the HP research Lab in image processing and analysis during the years 1998–2000, and to the Net2Wireless/Jigami research group during 2000–2001.