Orientation diffusions

P. Perona

Abstract— **Diffusions are useful for image processing and computer vision because they provide a convenient way of smoothing noisy data, of analyzing images at multiple scales, and of enhancing discontinuities. A number of diffusions of image brightness have been defined and studied so far; they may be applied to scalar and vector-valued quantities that are naturally associated with intervals of either the real line, or other flat manifolds.**

Some quantities of interest in computer vision, and other areas of engineering that deal with images, are defined on curved manifolds; typical examples are orientation and hue that are defined on the circle. Generalizing brightness diffusions to orientation is not straightforward, especially in the case where a discrete implementation is sought. An example of what may go wrong is presented.

A method is proposed to define diffusions of orientation-like quantities. First a definition in the continuum is discussed, then a discrete orientation diffusion is proposed. The behavior of such diffusions is explored both analytically and experimentally. It is shown how such orientation diffusions contain a nonlinearity that is reminiscent of edge-process and anisotropic diffusion. A number of open questions are proposed at the end.

Keywords— **Orientation analysis, texture analysis, diffusions, scale-space**

I. Introduction

A. Why orientation diffusions

Consider the image in Figure 1. It is quite clear that over most of the area of this picture the important information is contained in the orientation of the lines, rather than in the brightness values.

It is not easy, sometimes not possible, to determine the local orientation of the image using local operators. In figure 2 the orientations detected using a simple gradient-based method are shown; such an orientation image is quite noisy. In order to obtain a more precise determination of orientation one may use more sophisticated filtering schemes, for example one may use gabor-like filters that are tuned to the specific frequency that is present in the image (see e.g. [1]), however, by doing so one gives up locality (such filters are larger), generality (such filters need to be tuned for a specific frequency), and one may not be able to detect interesting events such as the two orientation singularities that are present in the image. It would be nice to be able to detect a noisy signal using very local operators, and then 'average' this information locally and adaptively in order to achieve both good noise rejection and localization of interesting singularities.

A similar problem has been faced for a long time by researchers working on brightness images. Low-pass filtering, or, equivalently, diffusing the image, has long been seen as a convenient way of rejecting noise and reconstructing the underlying data. By analogy, it may be a good idea to blur, or diffuse, orientation information in order to obtain more reliable information based on a neighborhood of each point.

California Institute of Technology 136-93, Pasadena, CA 91125 (perona@caltech.edu), and Università di Padova, Italy

Fig. 1. The image of a fingerprint.

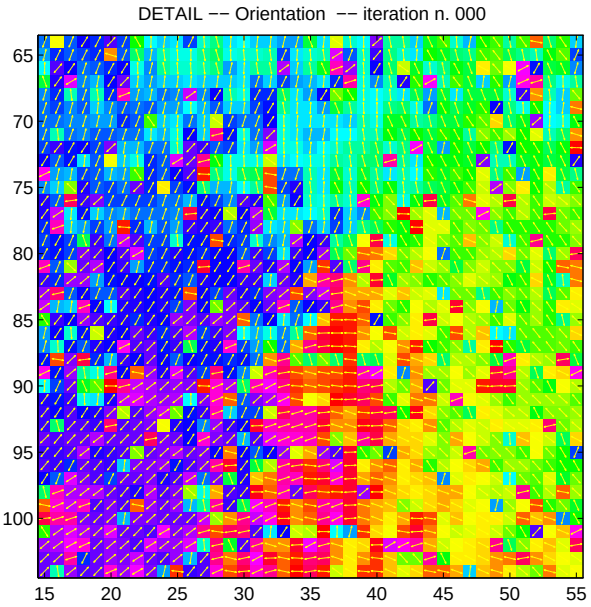

Fig. 2. The orientations detected on a detail of the image in figure 1 by using the gradient of a slightly blurred version of the image.

Another reason for diffusing images is that diffusions naturally produce a 'scale-space', i.e. a fine-to-coarse family of derived images. The main structures present in the image are easier to detect in the coarser scale portions of the scale-space, while the finer scale portions, closer to the original image, are more suitable for localization tasks [2], [3], [4] (for a comprehensive review of work on diffusions see the book edited by B. ter Haar Romeny [5]). It is reasonable to try and define scale-space operators for orientation images

as well.

A number of researchers [6], [7], [8], [9], [10] have proposed methods for diffusing, or low-pass filtering, orientation images. These methods may be, in essence, reduced to a simple three-step process. First, embed orientation in the plane via the map $\theta \to w = [\cos(\theta), \sin(\theta)]$; second, diffuse in the plane for a time t; third, project back to the circle via the map $w_t = [x, y] \to \arctan(y/x)$. Two caveats: In the cited works the norm of the orientation vector w is not restricted to be equal to one, rather it is set to represent image contrast; here we are purely interested in orientation, therefore we will use norm one vectors. Moreover, some of the authors referenced above diffuse an 'orientation tensor', rather than the orientation vector w; as pointed out in Granlund and Knuttson's book [9], page 229, the two are equivalent if the quantity being diffused in an orientation defined in the plane. As it will be shown in section II this process has two effects that may be unwanted: it is not consistent if it is iterated, and it produces large-scale changes in the topology of the orientation image with large jumps in the value of orientation.

B. Representing orientation

Before trying to define orientation diffusions we need to clarify the concept of orientation. Orientations may be associated with points θ on a circle; however, one has to be aware of the fact that there are two main 'types' of orientation. After being rotated by 180^o degrees a brightness step edge looks different; however, a line will look the same. The step edge is 360^o-periodic, while the line is 180^o periodic. Hue is 2π-periodic as well. There may be quantities that are $\frac{\pi}{2}$-periodic, for example cross-junctions as depicted in figure 3.

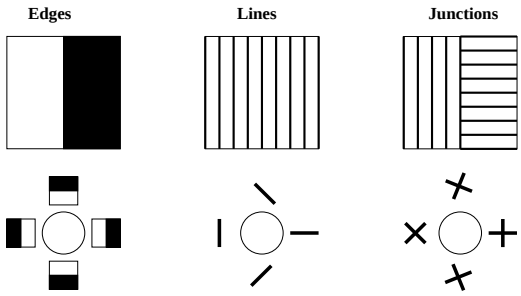

Fig. 3. There are different kinds of oriented structures in images. Some are 2π-periodic, some are π-periodic, others may have even smaller periods. This dishomogeneity is more apparent than real; as Granlund, Knuttson and collaboratorshave suggested, whatever its length the period may be mapped onto the circle thus obtaining a unified representation.

Granlund and Knuttson [9] have proposed to embed each one of these periodic quantities in the $(-\pi, \pi)$ circle by appropriate scaling of the orientation variable. For example, the map $\theta \to 2\theta$ maps line orientation onto the circle, as shown in figure 3 (center). Analogous embedding may be done on the complex plane as $e^{ia\theta}$ by appropriate choice of the constant a. As a result we may think of any orientation in a unified way as a quantity that is defined on the circle. In the rest of this paper we will therefore not need to worry any longer about the problem of which orientation we are using and we will use $\theta \in [-\pi, \pi]$ as our variable representing orientation.

II. Continuous space

The most straightforward idea for extending the idea of brightness diffusions to an orientation image $\theta_0(x, y)$ is to ignore the fact that θ is defined on the circle, rather than an interval of the line, and proceed as in the conventional linear diffusion case:

$$\begin{aligned} \overline{\theta}(x, y, 0) &= \theta_0(x, y) \\ \overline{\theta}_t(x, y, t) &= \triangle \overline{\theta}(x, y, t) \end{aligned} \tag{1}$$

Is this a good idea? In this section it will be shown that this is indeed a good method in the continuous case, and that it is related to the traditional embed-diffuse-project method that was sketched in the introduction. At the end of the section it will be shown, however, that a naive discretization of this method does not work; discrete orientation diffusions will be introduced in section III.

As discussed in the introduction, a popular way of diffusing orientation is to diffuse the components of an orientation vector (equivalently, an orientation tensor) that is embedded in a Euclidean space:

Definition 1: Call $\mathcal{D}'_{[t]}$ the operator that maps the function $\theta_0(x, y)$ into the function $\theta(x, y, t)$ by the sequence of embedding, diffusion and projection:

$$\begin{aligned} w(x, y, 0) &= e^{i\theta_0(x,y)} && \text{embedding} \\ w_t(x, y, t) &= \lambda \triangle w(x, y, t) && \text{diffusion} \\ \theta(x, y, t) &= \text{phase } w(x, y, t) && \text{projection} \end{aligned}$$

where w_t indicates $\frac{\partial w}{\partial t}$ (this notation will be used throughout the paper), and λ is a positive constant determining the velocity of the diffusion.

Observation 1:
This technique does not have the semigroup property. i.e. it is not necessarily true that when $0 < t_1 < t_2$ we have $\theta(x, y, t_2) = \mathcal{D}'_{[t_2 - t_1]}\theta(x, y, t_1)$. See the example in figure 4. If one wants a consistent behavior one has to use both phase and magnitude of $w(x, y, t_1)$ as the initial condition of further diffusion.

Observation 2: This technique may produce results that are 'topologically wrong'. In figure 4, 4th and 5th columns, it is demonstrated that it may 'unwind' an orientation loop while smoothing it. In the process, around location 45 the orientation suddenly jumps by π.

It is possible, however, to define a related diffusion that preserves the semigroup property. The main problem with the traditional 'plane-embedding' way of defining orientation diffusions stems from the fact that the result of the diffusion in the plane is projected back onto the circle only once, at the end of a time interval. As it may be noticed in figure 4 the result of repeating short diffusions and projection cycles is more 'correct' than what is obtained via a single cycle composed of a long diffusion followed by a projection.

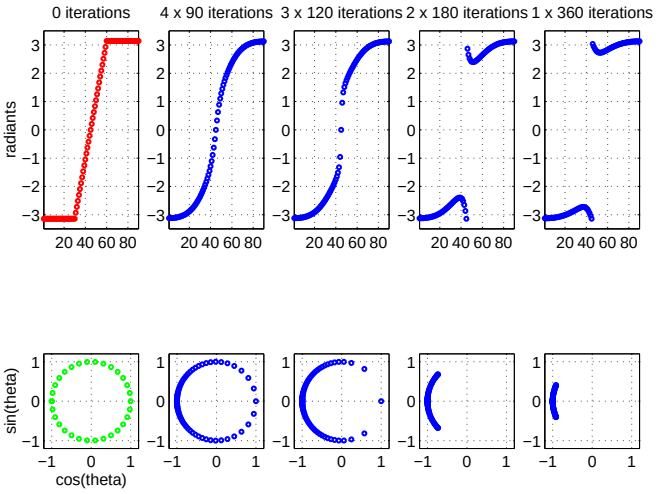

Fig. 4. Example of violation of the semigroup property. The diffusion is performed on the 1D orientation signal represented in the far-left column (each signal is represented twice: as an angle (top row), and as a polar plot (bottom row)). The signal is equal to π (equivalently $-\pi$) at the boundaries, and undergoes a full rotation between the 30th and the 60th sample. The 2nd to 5th columns show the result of applying $\mathcal{D}'_{[90]}$ 4 times, $\mathcal{D}'_{[120]}$ 3 times etc, so that the total diffusion time is constant: $t = 360$ time units. The last two columns present diffusion results that are qualitatively very different from the one in the 2nd column. The 3rd column shows a behavior that is close to critical. Notice that the result of diffusing the signal in the plane (the last column) is 'wrong' because as a result of the diffusion the 2π loop has disappeared, and the signal just oscillates around π. Compare with the result of applying $\mathcal{D}'^4_{[90]}$, where the kink is smoothed while the topology of the loop is preserved.

One might think of taking this last observation to the limit: alternate vanishingly short diffusion phases with projections onto the circle in order to obtain a consistent behaviour.

Definition 2: Call $\mathcal{D}'^k$ the k-th power of $\mathcal{D}'$, i.e. the operation of applying $\mathcal{D}'$ k times. Then we may define $\mathcal{D}$ as the limit obtained applying $\mathcal{D}'$ many times for very short intervals:

$$\mathcal{D}_{[t]} = \lim_{\delta \to 0} \mathcal{D}'^{\frac{t}{\delta}}_{[\delta]}$$

It turns out that this strategy generates a diffusion that has the semigroup property:

Proposition 1: The diffusion

$$\theta(x,y,t) = \mathcal{D}_{[t]}\theta_0(x,y)$$

enjoys the semigroup property.

Proof: We want to prove that $\mathcal{D}_{[t_1+t_2]} = \mathcal{D}_{[t_2]}\mathcal{D}_{[t_1]}$. Starting from the definition of $\mathcal{D}$ and dividing the limit into two limits one may verify this fact. ■

How do we calculate this diffusion in practice? It is possible, of course, to alternate short diffusion cycles with projections, as described above. There is a more natural implementation though. It is possible to prove that one may compute $\mathcal{D}_{[t]}\theta_0(x,y)$ by solving the usual heat diffusion equation:

Proposition 2: Call $\overline{\theta}(x,y,t)$ the solution of the diffusion equation

$$\begin{aligned} \overline{\theta}(x,y,0) &= \theta_0(x,y) \\ \overline{\theta}_t(x,y,t) &= \triangle\overline{\theta}(x,y,t) \end{aligned} \tag{2}$$

and suppose that it is twice differentiable (the symbol $\triangle$ here indicates the 'Laplacian' on the circle, one may calculate it by taking formally the spatial derivatives of θ).

Then $\overline{\theta}(x,y,t) = \mathcal{D}_{[t]}\theta_0(x,y)$.

Proof: Let's start by taking a first order Taylor expansion of $\mathcal{D}'_{[\delta]}\theta(x,y,t)$:

$$\begin{aligned} w(x,y,t) &= e^{i\theta(x,y,t)} \\ \theta(x,y,t+\delta) &= \text{phase } w(x,y,t+\delta) \\ &= \text{Im } \log w(x,y,t+\delta) \\ w(x,y,t+\delta) &= w(x,y,t) + \delta\triangle w(x,y,t) + \mathcal{O}(\delta^2) \end{aligned}$$

(we may take any determination of the complex log in the following calculations, provided that it is always the same).

We may compute aside the laplacian of $w(x,y,t)$:

$$\begin{aligned} \triangle w(x,y,t) &= w_{xx} + w_{yy} \\ &= \left(i\theta_{xx} - \theta_x^2 + i\theta_{yy} - \theta_y^2\right) e^{i\theta} \\ &= \left(i\triangle\theta - |\nabla\theta|^2\right) e^{i\theta} \end{aligned}$$

Let's now expand in Taylor series the complex log of $w(x,y,t+\delta)$ remembering that $\log(1+x) = 1 + x + \mathcal{O}(x^2)$:

$$\begin{aligned} &\log w(x,y,t+\delta) = \\ &= \log\left(w(x,y,t) + \delta\triangle w(x,y,t) + \mathcal{O}(\delta^2)\right) \\ &= \log\left(e^{i\theta}\left[1 + \delta\left(i\triangle\theta - |\nabla\theta|^2\right) + \mathcal{O}(\delta^2)e^{-i\theta}\right]\right) \\ &= i\theta + 1 + \delta\left(i\triangle\theta + |\nabla\theta|^2\right) + \mathcal{O}(\delta^2)e^{-i\theta} + \mathcal{O}(\delta^2) \end{aligned}$$

Therefore:

$$\begin{aligned} \theta(x,y,t+\delta) &= \text{Im } \log w(x,y,t+\delta) \\ &= \theta(x,y,t) + \delta\triangle\theta(x,y,t) + \mathcal{O}(\delta^2) \end{aligned} \tag{3}$$

From this and from the definition of $\mathcal{D}$:

$$\begin{aligned} \mathcal{D}_{[\delta]}\theta &= \lim_{h\to 0} \mathcal{D}'^{\frac{\delta}{h}}_{[h]}\theta \\ &= \theta(x,y,t) + \lim_{h\to 0}\frac{\delta}{h}h\triangle\theta(x,y,t) + \mathcal{O}(\delta^2) \\ &= \theta(x,y,t) + \delta\triangle\theta(x,y,t) + \mathcal{O}(\delta^2) \end{aligned}$$

And therefore:

$$\theta_t = \lim_{\delta\to 0}\frac{\mathcal{D}_{[\delta]}\theta - \theta}{\delta} = \triangle\theta \tag{4}$$

■

A caveat: the derivatives in the Laplacian $\triangle$ should be taken keeping in mind that theta is periodic, and therefore that $0 = 2\pi$.

In order to implement this diffusion on a computer one has to discretize it both in space and time. Unfortunately, is not clear how to discretize this diffusion on a lattice in

X and Y: what is the discrete version of the Laplacian? Plain finite differences will not do for approximating the spatial derivatives of θ.

One could think of using a plain finite-difference approximation of the Laplacian of equation 2 as is commonly done for brightness diffusions (for simplicity of notation a single space coordinate x is used here):

$$\begin{aligned}\theta(x,t+1) &= \theta(x,t) + \lambda\left[\theta(x-1,t) - 2\theta(x,t) + \theta(x+1,t)\right] \\ &= \theta(x,t) + \lambda\left[\theta(x-1,t) - \theta(x,t)\right] + \\ &\quad +\lambda\left[\theta(x+1,t) - \theta(x,t)\right] \end{aligned} \quad (5)$$

This equation is easy to interpret: the variable θ is incremented at the location x by an amount that is proportional to the sum of the difference of $\theta(x)$ with its neighbors $\theta(x-1)$ and $\theta(x+1)$.

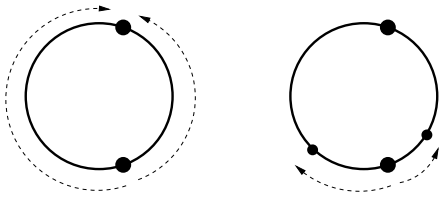

Fig. 5. The difference between two orientations may be expressed by two different notations, one positive and one negative, that denote however the same quantity. When such difference is multiplied by a small scalar we have an ambiguity.

While this is a perfectly natural discretized algorithm when the quantity to be diffused is defined on a flat manifold, it leads to an ambiguity on the circle. Suppose that two adjacent lattice locations $A = x$ and $B = x + 1$ have associated orientations $\theta_A = 0$, and $\theta_B = \pi/2$, say. The difference $\theta_B - \theta_A$ may be represented in two ways: $\pi/2$ and $-3\pi/2$. For now this representation ambiguity is only apparent both notations denote the same element of the circle. However, as soon as we multiply these quantities by a small scalar λ, in order to calculate the increment to be applied to θ as a consequence of the diffusion, we find ourselves in the embarassing position of having to choose between a counterclockwise small increment, and a clockwise larger increment (see figure 5).

Another problem is how to 'bias' the diffusion to keep the solution sufficiently close to some constant 'data' terms θ_d. Also in this case finite differences, or distances, $d(\theta, \theta_d)$ have to be computed. This is another case where the differential definition of 'distance' that is implicit in equation 2 is not helpful.

Also, is not clear how to define the diffusion in general. The hyp. that the solution is twice differentiable makes the differential formulation useless in practice.

III. Discrete space

In the previous section we have found a simple expression for the 'natural' diffusion of orientation on a continous spatial domain. We have seen that it is not possible to obtain the discrete version of this diffusion on a spatial lattice by simply taking spatial finite differences in order to approximate the spatial derivatives of θ.

In order to solve this problem we need to take one step back and derive the notion of diffusion of orientation by revisiting the relationship between diffusions and variational problems.

First of all recall that the common diffusion equation $u_t = \triangle u$, where u is a real-valued function defined on the plane (line), may be obtained as the 'gradient descent' solution of the Euler equation of an energy, or cost, functional $\mathcal{E}(u) = \int_{x,y} |\nabla u|^2 \, dx \, dy$.

The discrete analog (for the sake of convenience we may use one discrete spatial coordinate k) of such a cost function is $\mathcal{E}(u) = \frac{1}{2}\sum_k (u_{k+1} - u_k)^2$. The gradient descent strategy for minimizing such a cost function consists in the algorithm:

$$\frac{\partial}{\partial t} u_k = -\lambda \frac{\partial}{\partial u_k} \mathcal{E}(u) = \lambda\left(u_{k-1} - 2u_k + u_{k+1}\right) \quad (6)$$

where λ is a 'velocity' parameter.

Observe that the quantity $u_{k-1} - 2u_k + u_{k+1}$ is a discretization of $\triangle u$, and $\lim_{\delta \to 0} \frac{u_{x-\delta} - 2u_x + u_{x+\delta}}{\delta^2} = \triangle u$. Moreover, the cost function or 'energy' function is composed of the sum of pairwise terms $(u_{k+1} - u_k)^2$, i.e. it is the sum of the costs or energies associated to each pair of neighboring locations.

By analogy, we may be able to obtain a definition of a diffusion of orientation-like quantities on a discrete lattice if we manage to define the pairwise energy of two orientations. There are many reasonable definitions. If one uses the physical analogy with the energy associated to two superimposed magnets oriented as θ_A and θ_B, then the pairwise energy is $\mathcal{E}(\theta_A, \theta_B) = 1 - cos(\theta_A - \theta_B)$. The same result is obtained from modelling the energy associated to a spring of zero resting length that is attached to the points θ_A and θ_B of a circle, and again by modelling the energy that is associated to a point mass that slides on a circle and that is subject to gravity; this is represented in Fig. 6.

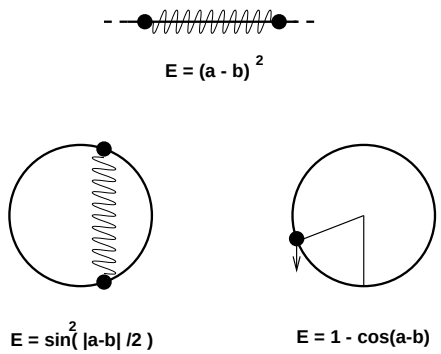

Fig. 6. The energy associated to a pair of locations on the straight line (top), and to a pair of orientations (bottom). Notice that $2\sin^2\left(\frac{\theta_A - \theta_B}{2}\right) = 1 - cos(\theta_A - \theta_B)$, therefore both physical systems are governed by the same energy.

If we adopt this definition then our total cost is: $\mathcal{E}(\theta) = \sum_k 1 - \cos(\theta_{k+1} - \theta_k)$. We may minimize such 'energy' or 'cost' by gradient descent:

$$\theta(k,0) = \theta_0(k)$$

$$\frac{\partial \theta_k}{\partial t} = \lambda [\sin(\theta_{k+1} - \theta_k) + \sin(\theta_{k-1} - \theta_k)]$$

This is the 'right' way to discretize orientation diffusions.

A couple of observations:

Observation 3: The expression $\sin(\theta_{k+1} - \theta_k) + \sin(\theta_{k-1} - \theta_k)$ may be seen as the discretization of the Laplacian for orientation-type variables: $\lim_{\delta \to 0} \frac{\sin(\theta_{k+\delta} - \theta_k) + \sin(\theta_{k-\delta} - \theta_k)}{\delta^2} = (\triangle\theta)(k)$.

Observation 4: Similarly one may think of $\sin(\theta_A - \theta_B)$ as the 'right' way to compute the finite difference on the circle. Observe that, according to this definition, the 'distance' (one may think of it as a force) between 0 and π is 0. This is due to the fact that a small perturbation of the value of either angle produces only a negligible second-order change in the value of the associated energy.

Observation 5: On a square-connected bidimensional lattice the above definition of discrete laplacian generalizes to:

$$\triangle\theta|_{(i,j)} = \sum_{n \in N(i,j)} \sin(\theta_n - \theta_{i,j}) \quad (7)$$

(8)

where $N(i,j)$ is the set of neighbours of the lattice location (i,j). This definition, using a 3x3 neighborhood structure, was used in the experiments presented in this paper. One caveat: this 2D definition is not rotation-invariant.

Observation 6: The energy used here is intuitive because it has a ready physical interpretation. It is by no means the only possible choice. Interesting computational behaviors may be obtained using different energy functions.

Observation 7: A contrast-invariant orientation scale-space, i.e. one that is invariant with respect to any monotonic change of image contrast $I' = f(I)$, may be obtained if the initial conditions of the orientation diffusion are obtained by measuring the orientation of the iso-brightness curves in the image, or, more realistically, by first smoothing the image a little by mean curvature motion [] and then extracting the orientation of the iso-brightness curves. Orientation diffusion is obviously contrast-invatiant.

IV. Biased diffusion

At times it is useful to bias the diffusion with some data term $\theta_d(x,y)$. We may exploit the energy based derivation of finite differences in this case as well:

$$\begin{aligned} \theta(x,y,0) &= \theta_0(x,y) \\ \theta_t &= \lambda \triangle\theta + \beta \sin(\theta_d - \theta) \end{aligned} \quad (9)$$

where $\triangle$ denotes either the Laplacian, or its discrete version as derived above.

V. Nonlinearity in the interaction

Observe the nonlinearity that is intrinsic in this diffusion scheme. While in the brightness diffusion the energy is quadratic, and therefore the flux, i.e. the amount of brightness flowing through a surface in the unit time, is proportional to the difference between neighbors, in the orientation diffusion that we just defined the flux is $\sin(\theta_A - \theta_B)$, and therefore depends nonlinearly on the difference $\theta_A - \theta_B$. This is reminiscent of the nonlinear probability density, energy and diffusion terms proposed by Geman and Geman [11], Blake and Zisserman [12], Perona and Malik [13], and others.

Notice that when $\theta_A - \theta_B = \pi/2$ one has maximum flux, i.e. maximum interaction between neighboring orientations, while when $\theta_A - \theta_B = \pi$ the flux is zero. In this latter case the two nodes of the lattice do not interact, as if a discontinuity process [11], [12] had intervened. As pointed out earlier, the energy function that is used here is by no means the only possibility, although it may be thought of as the simplest, or the most 'natural', one. More aggressive nonlinear behaviors may be obtained using e.g. an energy that is the square of the 'natural' one.

One further observation: if the flux was not zero when neighboring lattice locations differ by π, theflux would be discontinuous there: because of symmetry it needs to change sign when the difference passes from from $\pi - \epsilon$ to $\pi + \epsilon$.

VI. Uncertainty

It is often the case that orientations may be measured with a variable level of accuracy. For example an image may contained strongly oriented structure (edges, wood texture) and unoriented patches as well (sand texture, clouds, flat sky). As a result the data that the diffusion is 'filtering' or 'interpolating' has a variable degree of reliability.

For this reason it is useful at times to incorporate in the diffusion a measure of two kinds of uncertainty: (a) the uncertainty of the data in the bias term; (b) the degree of approximation with which the diffusion has estimated local orientation.

There are a number of different approaches to this problem. We examine here some simple ideas.

A. Uncertainty in the data

At times a model of the uncertainty with which orientation may be measured in the image is available, and this may vary throughout the image. Consider for instance the image in figure 1: in the white background region it is very difficult to make out any orientation, while in the fingerprint area orientation is relatively easy to measure almost everywhere. One may want to incorporate this information in the diffusion process.

One way to represent uncertainty in the data is via a probability density function $f_{\Theta_d}(\theta)$. The 'cost' in the system composed by a deterministic orientation θ and such an uncertain orientation θ_d may be calculated as the expectation of the energy:

$$\begin{aligned} \mathcal{C}(\theta, f_{\Theta_d}) &= \mathrm{E}\mathcal{E}(\theta, f_{\Theta_d}) \\ &= \int_{-\pi}^{\pi} \mathcal{E}(\theta, \phi) f_{\Theta_d}(\phi) d\phi \end{aligned} \quad (10)$$

It is often reasonable to assume that θ_d is distributed uniformly in the interval: $\theta_d \in (\mu_d - \sigma_d, \mu_d + \sigma_d)$, i.e.

$f_{\Theta_d}(\theta) = \frac{1}{2\sigma_d}$ for θ in the interval, and 0 elsewhere. In this case:

$$\begin{aligned} \mathcal{C}(\theta,\theta_d) &= 1 - \frac{1}{2\sigma_d}\int_{\mu_d-\sigma_d}^{\mu_d+\sigma_d} \cos(\phi-\theta)d\phi \qquad (11)\\ &= 1 - \sin(\mu_d+\sigma_d-\theta) + \sin(\mu_d-\sigma_d-\theta)\\ &= 1 - \frac{\sin\sigma_d}{\sigma_d}\cos(\mu_d-\theta) \end{aligned}$$

from this we may derive the expression for the bias term in the diffusion:

$$\begin{aligned} \theta(x,y,0) &= \theta_0(x,y)\\ \theta_t &= \lambda\triangle\theta + \beta\frac{\sin\sigma_d}{\sigma_d}sin(\mu_d-\theta) \qquad (12) \end{aligned}$$

Observe that what we have obtained is very similar to the previous diffusion; we have just added a constant weight term $w(\sigma_d) = \frac{\sin\sigma_d}{\sigma_d}$ that accounts for the certainty of the data. As one would expect when $\sigma_d \to 0$ (no uncertainty) then $w(\sigma_d) \to 1$ and when $\sigma_d = \pi$ (maximal uncertainty) then $w(\sigma_d) \to 0$.

In the case that an uncertainty is also associated with the diffused θ then one obtains an additional multiplicative term $\frac{\sin\sigma}{\sigma}$.

B. Uncertainty in the value of the diffusion

The local energy may be useful to calculate the 'reliability' $\mathcal{E}_F$ of the value of $\theta(x,y,t)$, and the distance $\mathcal{E}_d(x,y,t)$ of θ from the underlying data θ_d:

$$\begin{aligned} \mathcal{E}_F(x,y,t) &= \sum_{n\in\mathcal{N}(x,y)} 1 - cos(\theta_n(t) - \theta(x,y,t))\\ \mathcal{E}_d(x,y,t) &= 1 - \frac{\sin\sigma_d}{\sigma_d}\cos(\mu_d - \theta(x,y,t)) \end{aligned}$$

Where the uniform probability density described in the previous section is assumed for convenience.

The probabilistic interpretation of these energies is not clear. However, they prove to be useful e.g. for detecting and locating the singularities in the orientation as discussed in section VII and as shown in figure 7.

VII. Experiments

A. Smoothing of orientation

The lattice implementation of the orientation diffusion, with the discrete laplacian, as described by equation 8, was tested on the striped portion of the fingerprint image shown in figure 1. The intitial conditions of the diffusion, i.e. the initial orientation, were calculated from the gradient of the image brightness (orientation of the gradient plus 90^o) shown in figure 2. The boundary conditions were adiabatic (i.e. zero-gradient).

The results of running the diffusion may be seen in figure 7. Notice that the diffusion has the desired effect of smoothing orientation. The two singularities of orientation that are present in the fingerprint would be difficult to detect in the original image using a local operator. They are clearly visible in the smoothed orientation image after 50 iterations (second row, second column of figure 7); they may be easily detected and extracted from the corresponding energy image (third row, second column of figure 7).

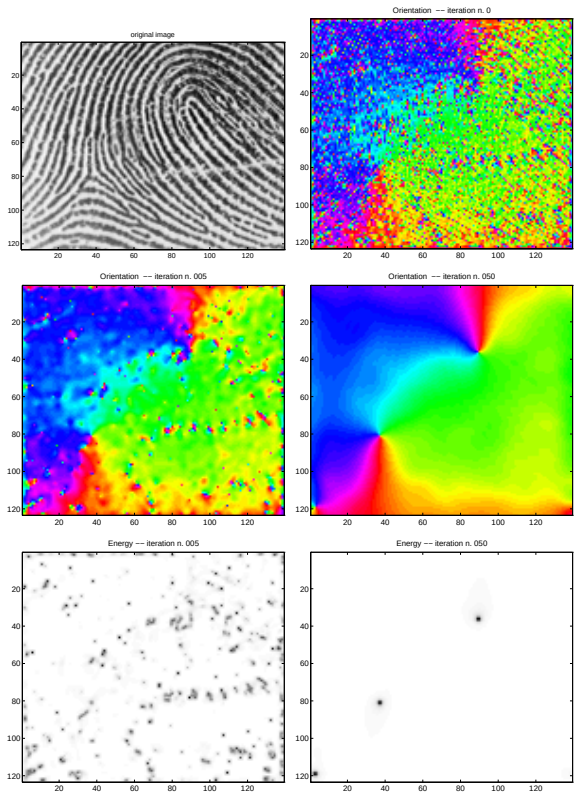

Fig. 7. Results of running orientation diffusion. The original image is shown top-left. The corresponding orientation, computed from the brightness gradient, is shown top-right. The orientation image (top-right) was used as the initial conditions of the diffusion. The second and third rows show the results of running the diffusion for 5, and 50 iterations (left and right columns). The second row shows the orientation and the third the energy.

B. Scale-space and causality

In figure 8 a detail of the previous experiment is shown. The detail corresponds to a portion of the image surrounding the bottom-left orientation singularity of the fingerprint. The original orientation image (top-right) is quite messy, however, after five iterations of the diffusion process (2nd row) most of the picture is already smooth. In the energy image (2nd row, left) a number of high-energy locations are visible. The rightmost one is due to a misaligned orientation which is realigned between iterations 5 and 10. Three more, due to 'fundamental' singularities in orientation, are clearly visible. The bottom-left two singularities display an interesting and commonly observed behavior: they merge and vanish shortly after iteration 10. By iteration 30 only one singularity has remained, the one that is present in the data at a coarse scale, i.e. the delta singularity.

In the scale-space literature this behavior is called 'causal': singularities may disappear, but not appear, as a result of running the diffusion forward in time. When a

singularity is found at a certain time, its 'cause' may always be found at previous times.

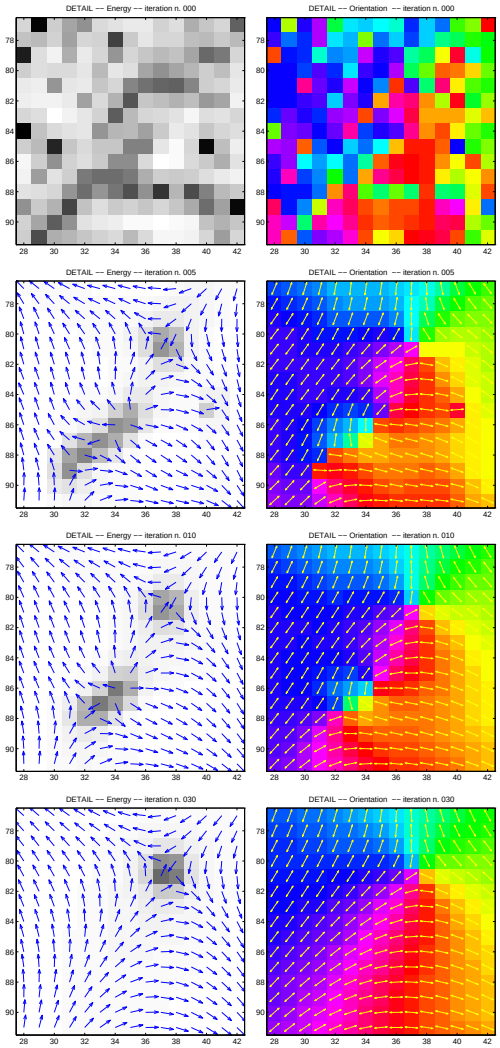

Fig. 8. Detail of the diffusion shown in figure 7. 0 iterations, 5, 10 and 30 are shown in the rows 1 thru 4. Orientation (left column) and energy (right column) are shown. The arrows that are superimposed on the energy and orientation images show orientation using two different conventions. Since in fingerprints the orientation of interest is of the line kind (see figure 3, center), such orientation may be shown via double-headed arrows that are parallel to the lines in the print (right column images), and as single-headed arrows that show the unified representation of orientation of the circle. In this case the single-headed arrows are oriented as 2θ where θ is the orientation of the underlying line. Notice the two singularities that merge and disappear briefly after iteration 10.

In order to better explore wether the causality behavior that one observes in figure 8 is intrinsic to the equation we have run a number of experiments using randomly oriented initial conditions. Violations of causality were never observed. One such experiment is shown in figure 9. A detail of figure 9 showing two singularities merging is shown in figure 10.

VIII. Discussion and conclusions

It was shown that the usual definition of diffusion such as is used on brightness data in computer vision may be extended trivially to the case of orientations in the continuous case; however the discrete case is particularly problematic. A common definition of orientation diffusion, obtained by embedding orientation in the plane, and then diffusing, was shown not to enjoy the semigroup property and to produce large-scale degradations the topology of the solution.

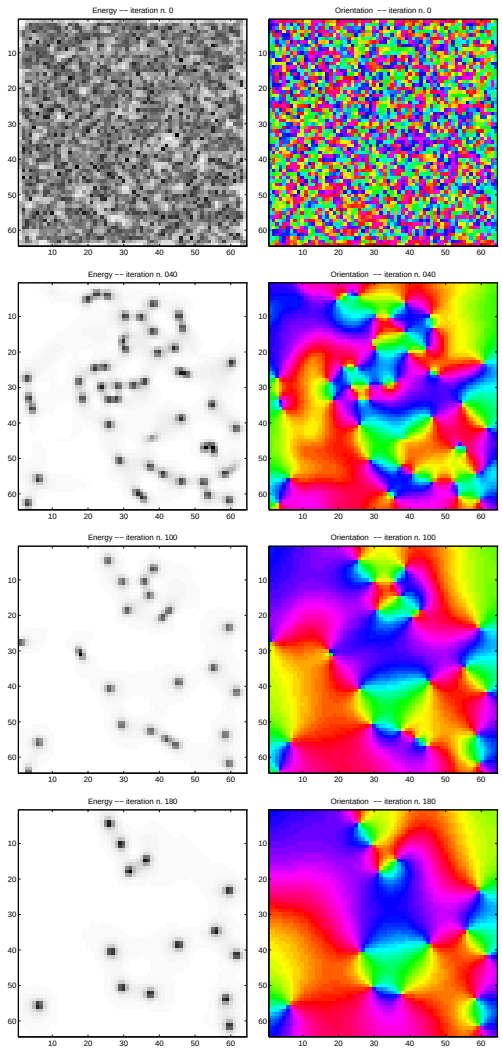

Fig. 9. The causal behavior of orientation diffusion is evident in this experiment where an image composed of random orientations was diffused. Notice that the orientation singularities are points and that they disappear, usually in pairs, as diffusion progresses. No 'birth' of singularities may be observed.

A definition of orientation diffusion was suggested; it is based on gradient descent of an energy function inspired by the model of simple physical systems. It is defined on a lattice and it may be implemented with simple calculations on 2-cliques. Experiments conducted on the image of a fingerprint demonstrate that this diffusion de-noises orientation data effectively.

The scale-space behavior of this equation was explored experimentally. Our empirical observations may be summarized in three conjectures:

1. The generic singularities are points, rather than lines.
2. The diffusion has the causality property, in that point-singularities may disappear, but never appear.

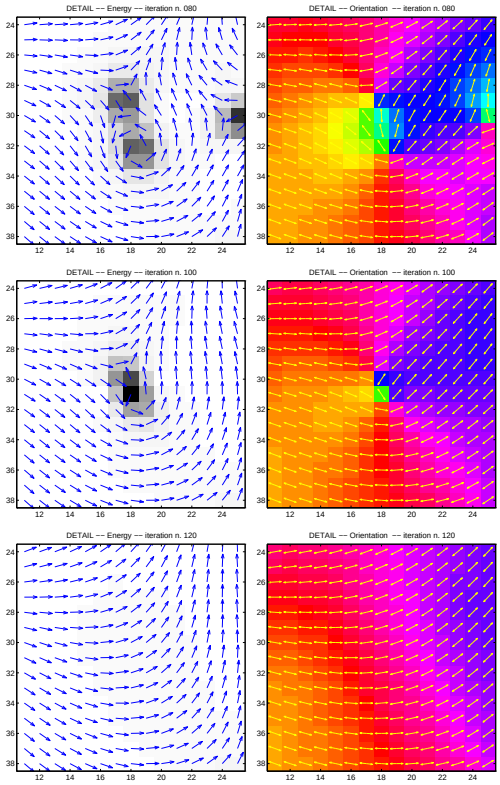

Fig. 10. Detail of the experiment of Fig. 9. The two singularities in the center of the picture merge and disappear.

3. The point-singularities disappear by either merging with other singularities, or by flowing off the boundary of the diffusion.

No proof of these conjectures is known to the author.

A few further questions remain open:

1. The definition of the orientation diffusion in the continuum was given for smooth solutions only. It is quite useless as such. It would be important and probably interesting to try and extend the definition to the case where solutions are allowed to have singularities, i.e. points where the orientation is discontinuous. At such points the energy function diverges, e.g. in the neighborhood of a point from which orientations radiate outwards the energy diverges with the cube of the distance from such point. However, it is not clear that the velocity with which the solution of the diffusion evolves should be infinite. E.g. one would expect that orientations symmetrically radiating from a point would be a stable equilibrium for the equation. Unlike the case of brightness diffusions we would therefore have a number of nontrivial stable equilibria, and some such equilibria might be 'dynamic' in that the corresponding 'orientation flux' would be nonzero.
2. While the definition of singularities in the continuous case is clear, the term 'singularity' only had an intuitive meaning in the discrete case (i.e. when the domain of definition of the orientation function is a lattice). It is unclear wether a rigorous and meaningful definition might be found.
3. Singularities come in different varieties: points from which orientations radiate outwards, points that are encircled by orientations, parabolae etc. Many of these were observed by Granlund and collaborators. It would be interesting to collect a complete taxonomy of such singularities, and relate this taxonomy to the scale-space behavior of pairs of singularities. This would parallel the catastrophe theory that was developed in the case of brightness for linear scale-space [2], [14], [15], [16]. Work on harmonic maps and liquid crystals, e.g. [17], may shed some light on this issue.

Acknowledgements

The second proof was derived in collaboration with Alfonse Salden. Conversations with Serge Belongie, Joachim Weichert, Sanjoy Mitter, and Fabio Fagnani are gratefully acknowledged. Comments from the anonymous reviewers were very helpful in improving a previous version of the paper.

References

[1] P. Perona, "Deformable kernels for early vision," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 17, no. 5, pp. 488–499, 1995.
[2] J. J. Koenderink, "The structure of images," *Biol. Cybern.*, vol. 50, pp. 363–370, 1984.
[3] A.P. Witkin, "Scale-space filtering," in *International Joint Conference on Artificial Intelligence*, 1983, pp. 1019–1021, Karlsruhe.
[4] P. Kube and P. Perona, "Scale-space properties of quadratic feature detectors," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 10, pp. 987–999, 1996.
[5] Bart M. ter Haar Romeny, Ed., *Geometry-Driven Diffusion in Computer Vision*, Kluwer, 1994.
[6] M. Kass and A. Witkin, "Analyzing oriented patterns," *Comp. Vision, Graphics and Image Proc.*, vol. 37, pp. 362–385, 1987.
[7] A. R. Rao and B. G. Schunck, "Computing oriented texture fields," *Comp. Vision, Graphics and Image Proc.*, vol. 53, pp. 157–185, 1991.
[8] J. Bigün, G.H. Granlund, and J. Wiklund, "Multidimensional orientation estimation with applications to texture analysis and optical flow," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 8, pp. 775–790, 1991.
[9] G. H. Granlund and H. Knuttson, *Signal processing for computer vision*, Kluwer, 1995.
[10] J. Weickert, "Foundations and applications of nonlinear anisotropic diffusion filtering.," *Zeitschrift für angewandte Mathematic und Mechanic*, vol. 76, pp. 283–286, 1996.
[11] S. Geman and D. Geman, "Stochastic relaxation, gibbs distributions, and the bayesian restoration of images," *IEEE Transactions on PAMI*, vol. 6, pp. 721–741, November 1984.
[12] A. Blake and A. Zisserman, *Visual reconstruction*, MIT press, 1987.
[13] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 7, pp. 629–639, July 1990.
[14] A. Yuille and T. Poggio, "Scaling theorems for zero crossings," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 8, pp. 15–26, January 1986.
[15] J. Babaud, A. Witkin, Baudin, and R.Duda, "Uniqueness of the gaussian kernel for scale-space filtering," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 8, January 1986.
[16] J. J. Clark, "Authenticating edges produced by zero-crossing algoritms," *IEEE Trans. Pattern Anal. Mach. Intell.*, 1989.
[17] J.-M. Coron, H. Brezis, and E. Lieb, "Harmonic maps with defects," *Comm. Math. Physics*, 1986.