ANALYZING ORIENTED PATTERNS

Michael Kaas
Andrew Witkin

Schlumberger Palo Alto Research
3340 Hillview Ave.
Palo Alto, CA 94304

ABSTRACT

Oriented patterns, such as those produced by propagation, accretion, <r deformation, ore common in nature and therefore an important dass for visual analysis. Our approach to understanding such patterns is to decompose them into two parts: a flow field, describing the direction of anisotropy, and the residual pattern obtained by describing the image in a coordinate system built, from the flow field. We develop a method for the local estimation of anisotropy and a method for combining the estimates to construct a flow coordinate system. Several examples of the use of these methods are presented. These included the use of the (low coordinates to provide preferred directions for edge detection, detection of anomalies, fitting simple models to the straightened pattern, and detecting singularities in the flow field.

I Introduction

A central focus in recent computational vision has been the decomposition of the original intensity image into intrinsic images (Horn 1977, Barrow & Tenenbaum, 1978; Marr, 1982), representing such properties as depth, reflectance, and illuminance. These intrinsic properties are believed to be more meaningful than image intensity because they describe basic independent constituents of the image formation process. Thus, for example, in separating shape from illumination, we can recognize an invaxiance of shape regardless of changing illumination.

The advantages of decomposing what we see into its more-orless independent parts extends beyond the image formation process to the shapes and patterns on which that process operates. For instance, decomposing a bent rod into a straight rod and a bending transformation reveals the similarity between a bent rod and one that hasn't been bent, or some other solid that's been bent the same way (Barr, 1984).

Just as we need to understand the image-forming process to decompose an image into intrinsic images, we need to understand the processes that generate patterns to decompose them into their intrinsic parts. But, whde there is only one image-forming process, a staggering variety of processes shape and color the world around us. Our only hope of dealing with this complexity is to begin with some basic pattern classes that recur in nature, and understand how to decompose and describe them.

One such class are oriented patterns, notably those produced by propagation, accretion, or deformation. To understand an oriented pattern we must be able to say (l) what is propagating, accreting, or deforming, and (2) which way and how much. More precisely, we must estimate everywhere the direction and magnitude of anistropy (which we will call the flow field,) and describe the residual pattern, independent of that field. Why this decomposition leads to simpler, more regular descriptions is best illustrated by example:

- A typical oriented pattern created by propagation is the streaked trail left by a paint brush dipped in variegated paint. The flow field describes the trajectory of the brush, the residual pattern depending only on the distribution of paint on the brush.
- Accretion typically results in laminar structures, such as wood grain. Here, the flow field gives isochrones (the moving accretion boundary,) and the residual pattern describes the change in color or brightness of the accreting material over time.
- If an isotropic body is deformed, the flow field principally describes the bending and stretching it has undergone, while the residual pattern describes the undeformed body.

In all these cases, separate descriptions of the flow field and the residual pattern are appropriate because they describe different processes. The path of propagation for many physical processes is controlled by very different mechanisms than control the coloration of the trail left behind. Similarly, the mechanisms which control the shape of an accretion boundary are frequently unrelated to the processes controlling the color of the accreted material. Finally, the forces which deform a piece of material are often completely unrelated to the process which created the piece of material in the first place By separately describing these processes, we can create descriptions of the whole which are often simpler than is possible without the separation because each of the pieces may have different regularities.

Orientation selective mechanisms have been extensively stud ied by physiologists since Ilubel and Wiesel's (1962) discovery of orientation selective cells in mammalian visual cortex (see Schiller et. al. (1976) for a comprehensive example). There has also been considerable interest among psychologists in the perception of oriented patterns, particularly dot patterns (Glass, 1969). Only recently have the computational issues involved received attention. Stevens (1978) examined the groupmg of tokens in Glass patterns based on orientation. While successful with Glass patterns, his methods were never extended to natural imagery. Zucker (1983) investigated the estimation of orientation by combining the outputs of linear operators. Zucker's estimation method for what he calls "Type II" patterns, while differing in many respects, is quite dose in spirit to our own.

Little progress has been made in using local orientation estimates to interpret patterns, perhaps because reliable estimates have proved difficult to obtain. The key difference between our work and earlier efforts lies in our use of the flow field to build a natural coordinate system for analyzing the pattern.

The remainder of the paper covers the computation of the flow field by local estimation of orientation, the construction of a coordinate system using the flow field, and some examples of analysis and description using flow coordinates.

In Section 2 we develop an estimator for the local flow direction, that direction in which intensity tends to vary most slowly

due to an underlying anisotropic process. The estimator, based on the direction of least spatial variance in the output of an oriented filter, is computed as follows: After initial filtering, the intensity gradient is measured at each point in the image. The gradient angle, 0, is then doubled (by treating the gradient vectors as complex numbers and squaring them) to map directions differing by π into a single direction. The transformed vectors are then summed over a weighted neighborhood around the point of interest. The angle of the summed vector is halved, undoing the previous transformation. This gives an estimate for the direction of *greatest* variance, which is then rotated by π/2 to yield the flow direction.

In section 3, we describe the construction and use of coordinate systems based on the result of local estimation. Integral curves in the flow field are computed numerically, by following the estimated vectors from point to point. A coordinate system is constructed in which the integral curves are parameter lines. Transforming the image into these "flow coordinates" straightens the pattern, removing the effects of changing orientation. We present several examples of analysis and description of the flow field and the straightened pattern.

II Flow Computation

For intensity patterns created by anisotropic processes such as propagation, accretion, or deformation, variation in the flow direction is much slower than variation in the perpendicular direction. Anisotropy in such patterns will be evident in the local power spectrum. The high frequency energy will tend to cluster along the line in the Fourier domain perpendicular to the flow orientation.

A simple way to detect this clustering is to sum the energy in an appropriate region of the power spectrum and examine how the sum is affected by rotations. This can be done by examining the energy in the output of an appropriate orientation-selective linear filter. The orientation at which the energy is maximal can be expected to be perpendicular to the flow orientation.

Selection of the filter involves a number of tradeoffs. Very low spatial frequencies are affected more strongly by illumination effects than surface coloration, so they are inappropriate for measuring textural anisotropy. Very high spatial frequencies are sensitive to noise and aliasing effects so they too are inappropriate. Hence some type of roughly bandpass filtering is required. The orientation specificity of the filter is also quite important If the filter is too orientation-specific then a large spatial neighborhood will be required in order to make a reliable measurement of the energy. Conversely, if the filter responds over a wide range of orientations then it will be difficult to localize the orientation very accurately. Thus there is a trade-off between angular- and spatial- resolution.

One reasonable choice for the frequency response of the filter is

$$F(r,\theta) = [e^{r^2\sigma_1^2} - e^{r^2\sigma_2^2}]2\pi i r\cos(\theta). \qquad (1)$$

The filter is bandpass with passband determined by and a_2. In our experience, ratios of the sigmas in the range of 2.0 to 10.0 work well. The orientation specificity or *tuning curve* is provided by the cosine dependence of the filter on *0.* This appears to strike a reasonable balance between angular- and spatial- resolutions for the range of patterns we have examined. The filter's power spectrum is shown in figure 1

The cosine orientation tuning-curve of the filter has some unusually good properties for computing the filter output at different orientations. The impulse response $S(x_y y)$ of the filter

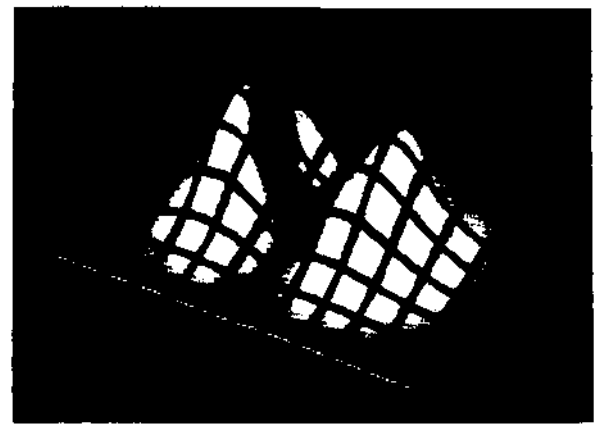

Figure 1: The power spectrum of the filter in equation 1 for $\sigma_2 = 2\sigma_1$.

is

$$S(x,y) = \frac{\partial}{\partial x} H(x,y)$$

where

$$H(x,y) = [\sigma_1{}^2 e^{r^2\cdot\sigma_1^2} - \sigma_2{}^2 e^{r^2\cdot 2\sigma_2^2}]$$

is an isotropic filter. Let $C = H * I$ and let $R_\theta[S]$ denote a counter-clockwise rotation of S by an angle θ. Then the convolution $R_\theta[S] * I$ is just the directional derivative of $H * I$ in the θ direction. The directional derivative can easily be written in terms of the gradient so we have

$$R_\theta[S] * I = (\cos\theta, \sin\theta) \cdot \nabla H * I. \qquad (2)$$

Thus a single convolution suffices for all orientations.

Since the filter S severely attenuates very low frequencies $R_\theta[S] * I$ can be safely regarded as zero-mean. Thus the variance in its output can be estimated by the expression

$$V(\theta) = W * (R_\theta[S] * I)^2$$

where $W(x,y)$ is a local weighting function with unit integral. We use Gaussian weighting functions $W(x,y)$ because approximate Gaussian convolutions can be computed efficiently (Burt 1979).

Using the gradient formulation of the filter output in equation 2, we can write the variance $V(\theta)$ as

$$V(\theta) = W * [\cos(\theta)C_x + \sin(\theta)C_y]^2. \qquad (3)$$

A. Interpretation of Filter Output

There remains the issue of interpreting $V(\theta)$. Assume that there is only one axis of anisotropy. Then $V(\theta)$ will have two extrema π apart corresponding to that axis. Let $V_2(\theta) = V(\theta/2)$. Then $V_2(\theta)$ will have a single extremum in the interval $0 < \theta < 2\pi$. A computationally inexpensive way of estimating the position of this extremum is to consider V_2 as a distribution and compute its mean. Since θ is periodic, V_2 should be considered as a distribution on the unit circle. Hence its mean is the vector integral $(\alpha,\beta) = \int_0^{2\pi} V_2(\theta)(\cos\theta, \sin\theta)d\theta$. The angle $\tan^{-1}(\beta/\alpha)$ is an estimate of the angle of the peak in V_2 and hence twice the angle of the peak in V. Thus the angle ϕ of *greatest* variance can be written

$$\phi = \tan^{-1}\frac{\beta}{\alpha}$$
$$= \tan^{-1}\left(\frac{\int_0^{2\pi} V_x(\theta)\sin(\theta)d\theta}{\int_0^{2\pi} V_x(\theta)\cos(\theta)d\theta}\right)/2 \qquad (4)$$
$$= \tan^{-1}\left(\frac{\int_0^{\pi} V(\theta)\sin(2\theta)d\theta}{\int_0^{\pi} V(\theta)\cos(2\theta)d\theta}\right)/2$$

These integrals are evaluated in Appendix A to show that the angle of anisotropy ϕ can be written

$$\phi = \tan^{-1}\left(\frac{W * 2C_xC_y}{W * (C_x^2 - C_y^2)}\right)/2 \qquad (5)$$

which directly yields a simple algorithm for computing ϕ.

B. Combining Gradient Orientations

Notice that the right hand side of equation 5 can be regarded as the orientation of a locally weighted sum of the vectors of the form $J(x,y) = (C_x^2 - C_y^2, 2C_xC_y)$. These vectors are related in a simple way to the gradient vectors $G(x,y) = (C_x, C_y)$. The magnitude of $J(x,y)$ is just the square of the magnitude of $G(x,y)$ and the angle between $J(x,y)$ and the x-axis is twice the angle between $G(x,y)$ and the x-axis. This follows easily from the observation that $(C_x + C_y i)^2 = C_x^2 - C_y^2 + 2C_xC_y i$.

One might be tempted to believe that smoothing the gradient vectors *G(x,y)* would be nearly as good a measure of anisotropy as smoothing the rotated squared gradient vectors *J(x,y)*. This is emphatically not the case Consider an intensity ridge such as /(x,y) = exp(- x^2). The gradient vectors on the left half-plane all point to the right and the gradient vectors on the right half-plane all point to the left. Adding them together results in cancellation. By contrast, if they are first rotated to form the *J* vectors, they reinforce. The types of oriented patterns we are concerned with often have nearly symmetric distributions of gradient directions around the axis of anisotropy. In such patterns, if the gradients are added together directly, the cancellation is so severe that the result often has little relation to the direction of anisotropy. Thus the difference between rotating the gradient vectors or leaving them be is often the difference between being able or unable to detect the anisotropy. Note also that smoothing the image first and then computing the gradients is exactly the same as computing the gradients and then smoothing. It will not avoid the difficulties of cancellation.

C. Coherence

In addition to finding the direction of anisotropy, it is important to determine how strong an anisotropy there is. If the orientation of the local *J* vectors are nearly uniformly distributed between 0 and 2Π, then the orientation of slight anisotropy is not very meaningful. Conversely, if all the *J* vectors arc pointing the same way then the indication of anisotropy is quite strong and *Φ* is very meaningful. A simple way of measuring the strength of the peak in the distribution of *J* vectors is to look at the ratio x(x,y) = |W *JI/W * IJI which we will call the *coherence* of the flow pattern. If the *J* vectors are dose to uniformly distributed, then the ratio will be nearly zero. If the ./ vectors all point the same way, the ratio will be one. In between, the ratio will increase as the peak gets narrower.

D. Summary

The computation of the flow direction and local coherence can be summarized as follows. First the image /(z, y) is convolved with the isotropic portion $H(x,y)$ of the filter response. The result $C(x,y)$ is then differentiated (by finite differences) to form $C_x(x,y)$ and $C_y(x,y)$. The resulting vectors $(C_x(x,y), C_y(x,y))$ are rotated by computing $J_1(x,y) = 2C_x(x,y)C_y(x,y)$ and $J_2(x,y) = C_x^2(x,y) - C_y^2(x,y)$. The gradient magnitude $J_3(x,y) = |C_x^2(x,y) + C_y^2(x,y)|^{1/2}$ also has to be computed in order to measure the coherence. The next step is to convolve $J_1(x,y)$, $J_2(x,y)$, and $J_3(x,y)$ with the weighting function $W(x,y)$ to obtain $J_1^*(x,y)$, $J_2^*(x,y)$, and $J_3^*(x,y)$. The angle $\phi(x,y)$ of anisotropy and the coherence $\chi(x,y)$ can then be computed from the formulas

$$\phi(x,y) \approx \tan^{-1}(J_1^*(x,y), J_2^*(x,y))/2$$

and

$$\chi(x,y) = (J_1^*(x,y)^2 + J_2^*(x,y)^2)^{1/2}/J_3^*(x,y).$$

An example *of* this computation applied to a picture of a piece of wood is shown in figure 2. The flow direction *Φ(x,y)* + 7r 2 is displayed by the orientation of small needles superimposed on the image. The lengths of the needles is proportional to the coherence x(x,y). Note that the pattern is strongly oriented except near the knot in the middle.

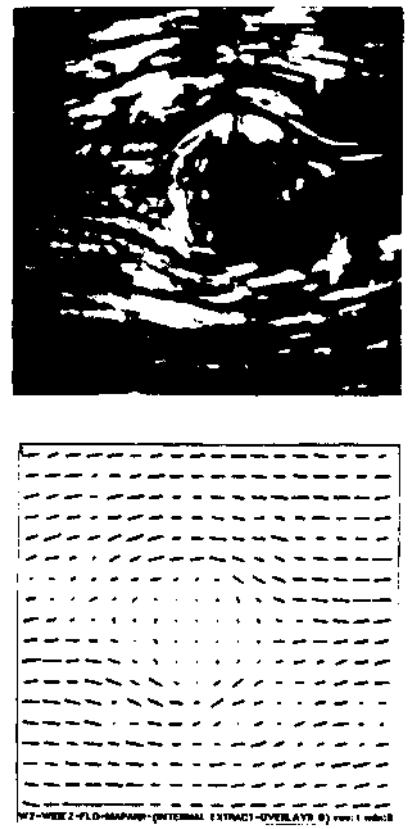

Figure 2: An image of wood grain with its flow field. Estimated flow directions are given by the black needles. The length of the needle encodes coherence. Notice that coherence is low within the knot at the center.

E. Relation To Prior Work

The flow computation just described bears an interesting relation to an early proposal of David Marr that information about local distributions of oriented edge elements be included in the *primal sketch* (Marr 76). If this proposal is combined with his later work with Hildreth on edge detection (Marr & Hildreth, 1980) it results in a special case of the above computation. Marr and Hildreth define edges as zero-crossings in the Laplacian of the Gaussian smoothed image. The natural combination of Marr's proposal with this definition of edge elements calls for examining the local density of zero crossings as a function of orientation. For stationary zero-mean Gaussian processes the square of the oriented zero-crossing density is approximately $V(\theta)$ (see appendix B). Thus in the special case where the point spread function of the filter is $S = (\partial/\partial x)\nabla^2 \exp(-(x^2+y^2)/2\sigma^2)$ our computation can be viewed as computing the direction of minimal edge density in the Marr-Hildreth theory.

Zucker's work on flow (Zucker 1983) is also related to a special case of the above computation. For biological reasons, he prefers to use oriented second derivatives of Gaussians as the initial filters. These have $F(r,\theta) = r^2 \exp(-r^2/2)\cos^2(\theta)$. Instead of looking at the variance of the filter outputs as the orientation is changed, he combines the outputs in a biologically motivated relaxation process. Although quite different in detail, the computation described here has much in common with his technique.

III Flow coordinates.

The orientation field is an abstraction from the anisotropic pattern that defines it. We can, for example, get the same spiral field from a pattern composed of bands, irregular streaks, dot pairs, etc. In addition to measuring the orientation field, it is useful to be able to produce a description of the underlying pattern independent of the changing direction of anisotropy. Such a description would make it possible to recognize, for example, that two very different orientation fields are defined by the same kind of bands or streaks.

A powerful way to remove the effects of changing orientation is to literally "straighten" the image, subjecting it to a deformation that maps the flow lines into straight, parallel lines in a canonical (e.g. horizontal) orientation. Performing this deformation is equivalent to viewing the image in a coordinate system (u,v), with $u = u(x,y)$ and $v = v(x,y)$ that everywhere satisfies

$$\nabla u \cdot (\sin\phi, -\cos\phi) = 0. \tag{6}$$

Equation 6 does not determine a unique coordinate system. An additional constraint may be imposed by choosing lines of constant v orthogonal to those of constant u, i.e,

$$\nabla v \cdot (\cos\phi, \sin\phi) = 0 \tag{7}$$

which has the desirable effect of avoiding the introduction of spurious shear in the deformation.

Even with equation 7, an additional constraint is needed, because we are free to specify arbitrary scaling functions for the u and v axes. In the spirit of equation 7, we want to choose these functions to avoid the introduction of spurious stretch or dilation. Although difficult to do globally (one might minize total stretch,) we will usually want to construct a fairly local coordinate frame around some point of interest. For this purpose, it suffices to take that point as the origin, scaling the axes u — 0 and v = 0 to preserve arc-length along them.

Intuitively, the flow field describes the way the pattern is bent, and viewing the image in these *flow coordinates* straightens the pattern out. Figure 3 shows the flow coordinate grid for the wood-grain image from figure 1. The grid lines were computed by taking steps of fixed length in the direction ($\cos\Phi$, $\sin\Phi$) or ($-\sin\Phi\cos\Phi$) for lines across and along the direction of flow respectively, using bilinear interpolation on the orientation field. Since Φ is always computed between 0 and 7r, we must assume that there are no spurious discontinuities in direction to track smoothly.

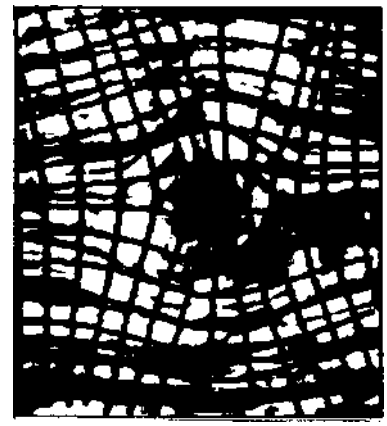

Figure 3: A flow coordinate grid obtained for the image of figure 2.

For many purposes it is unnecessary to compute the deformed image explicitly, but doing so vividly illustrates the flow coordinates' ability to simplify the pattern. Figure 4 shows the deformation from image coordinates to flow coordinates in several stages. As the grain lines straighten the knot shrinks and finally vanishes. The deformed images were anti-aliased using texture-map techniques (Williams, 1983) The deformed linage shows, to a reasonable approximation, what the grain would have looked like had it not been subjected to the deforming influence of the knot.

Thus far, we have separated the image into a flow field, and a pattern derived by viewing the image in flow coordinates. We argued earlier that the advantage of this decomposition, like the decomposition of an image into intrinsic images, is that the components are liable to be simpler and more closely tied to independent parts of the pattern-generating process than is the original image. To exploit the decomposition, we need ways of analyzing, describing, and comparing both the flow field and the straight ened pattern. These are difficult problems. In the remainder of this section, we present several examples illustrating the utility of the decomposition.

A. A coordinate frame for edge detection.

Oriented measurements have been widely used in edge detection. For example Marr & Poggio (1979) employed directional second derivative operators, whose zero-crossings were taken to denote rapid intensity changes. Due to the difficulty in selecting an orientation, Marr & Hildreth (1980 later abandoned this scheme, in favor of zero-crossings of the Laplacian, a non-directional operator.

The flow field provides two meaningful directions along and across the direction of flow — in which to look for edges within an oriented pattern. Zero-crossings in the second directional derivative in the direction of Φ (against the grain) should highlight edges that contribute to defining the flow field, while zero-

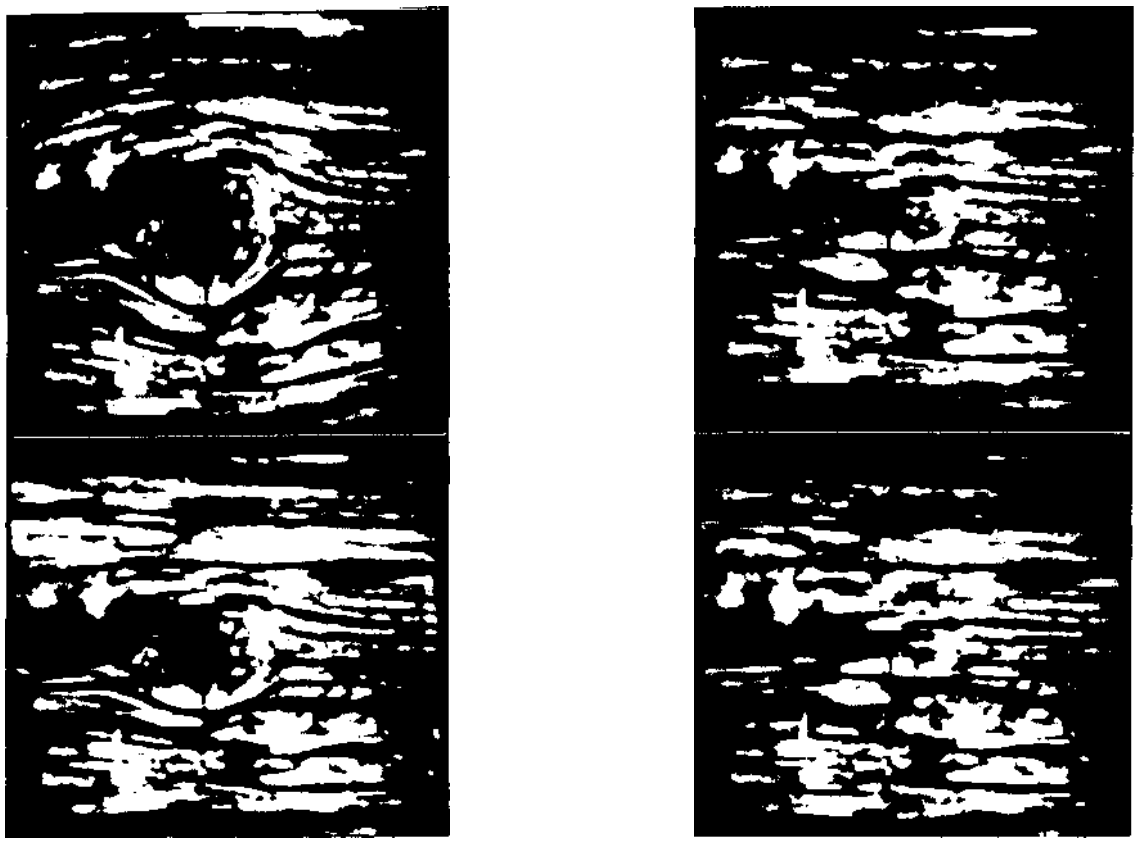

Figure 4: Deformation, in stages, from image coordinates to flow coordinates. Upper left: the original image; Lower left, upper right: two intermediate stages, in which the grain's curvature has diminshed, and the knot compressed; Lower right, the image as seen in flow coordinates: the grain lines are straight and the knot has vanished, showing approximately what the grain would have looked like had it not been deformed by the intrusion of the knot.

crossings in the second derivative perpendicular to Φ (with the grain) should highlight anomalous elements or terminations. The sum of these two derivatives is the Laplacian.

The two directional derivatives of the wood grain image are shown, with the Laplacian, in figure 5. Indeed, the derivative against the grain captures all the elements comprising the grain pattern, while the derivative with the grain does not appear meaningful. The Laplacian confuses these very different signals by adding them together.

The derivative along the grain can also be meaningful, where anomalous elements are present. In addition to being perceptually salient, such anomalies are often physically significant, with origins such as cracks, intrusions, or occlusions, that are distinct from those of the main pattern. In man-made structures, anamolies are often important because they indicate some variety of flaw.

Figure 6 shows a pattern of aligned elements (straw) with some anomalous elements. The directional second derivatives along and across the flow direction are shown, together with their sum (which is just the Laplacian.) Differentiating along the grain highlights anomalous elements, attenuating the rest (thus finding the "needles" in the haystack.) Differentiating across the grain supresses the anomalies. The Laplacian shows both.

A related demonstration is shown in figure 7, in which the anomalous elements have actually been removed by directional median filtering in the flow direction.

B. Singularities.

We have shown several ways in which viewing an oriented pattern in flow coordinates facilitates analysis and description of the pattern. Describing and anlyzing the flow field itself is the other side of the coin. The topology of a flow field, as of any vector field, is determined by the structure of its singularities, those points at which the field vanishes. Identifying and describing singularities is therefore basic to describing the flow field. The singularities provide the framework around which metric, properties, such as curvature, may be described. Singularities arc also perceptually salient (see figure 8.)

A robust basis for identifying singularities is the index or winding number (Spivak, 1979.) Suppose we follow a closed curve on a vector field. As we traverse the circuit, the vector rotates continuously, returning to its original orientation when the circuit is completed. The index or winding number of the curve is the number of revolutions made by the vector in traversing the curve. The index of a point is the index of a small circle as we shrink it around the point:

$$\begin{aligned}\mathbf{ind}(x,y) &= \lim_{\epsilon\to 0} \frac{1}{2\pi}\int_0^{2\pi} \frac{\partial}{\partial\theta}\phi(x+\epsilon\cos\theta, y+\epsilon\sin\theta)\,d\theta \\ &= \lim_{\epsilon\to 0} \frac{1}{2\pi}\int_0^{2\pi} (-\sin\theta,\cos\theta) \\ &\qquad \cdot \nabla\phi(x+\epsilon\cos\theta, y+\epsilon\sin\theta)\,d\theta.\end{aligned}$$

To compute the winding number numerically, we divide the flow field into suitably small rectangles, summing the rotation of Φ around each rectangle. As in computing the flow lines, we assume

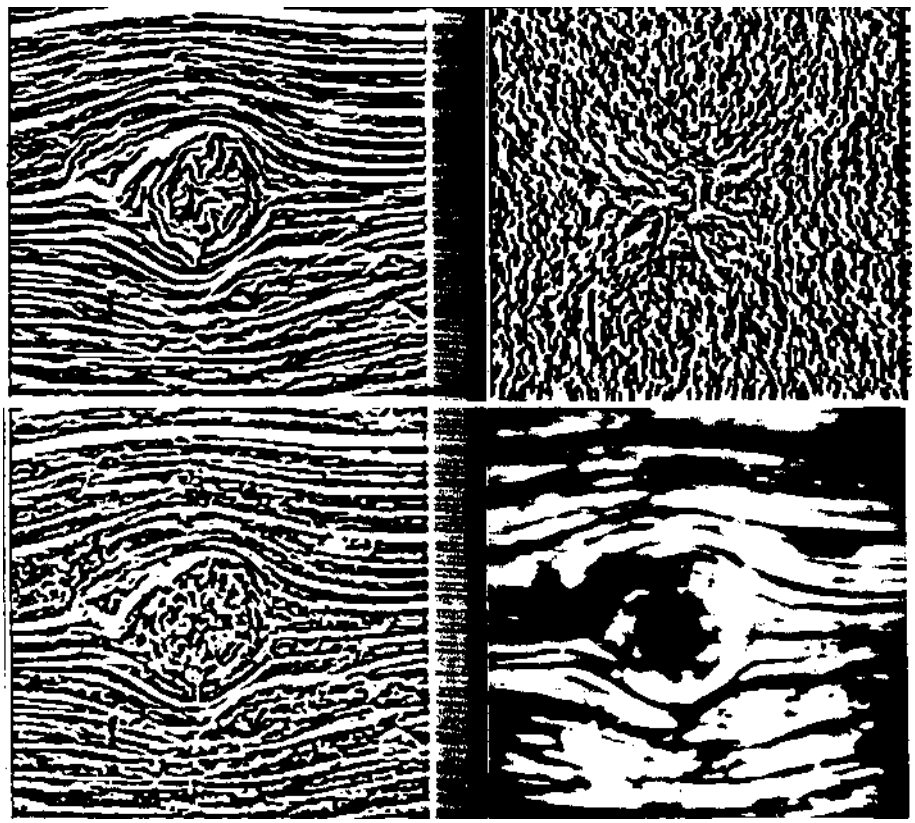

Figure 5: Using flow coordinates for edge detection Upper left: 2nd directional derivative across the How direction. Upper right: 2nd directional derivative along the flow direction. The first of these highhghts the oriented structure, the second supresses it. Lower left: the sum of the directional derivatives is the Laplacian.

that Φ has no spurious discontinuities. Where the result is non-zero, the rectangle surrounds a singularity. Figure 9 shows an example of the detection of singularities using winding number, for a fingerprint. We are currently working on classifying the singularities, and using them to describe the topology of the flow field.

IV Conclusion

We addressed the problem of analyzing oriented patterns by decomposing them into a flow field, describing the direction of anisotropy, and describing the pattern independent of changing flow direction.

A specific computation for estimating the flow direction was proposed The computation can be viewed as a) finding the direction of maximal variance in the output of a linear filter, b) combining gradient directions locally, or c) finding the direction of maximal edge density. The computation has been applied to a number of natural and man-made patterns with consistent success.

The flow field was then used to form a coordinate system in which to view the pattern. Two orthogonal familes of curves—along and across the direction of flow form the coordinate system's parameter lines. Viewing the pattern in these flow coordinates amounts to deforming the pattern so that the flow lines become parallel straight lines. This deformation produces a pattern that is simpler, more regular, and therefore more amenable to analysis and description than the original one.

Several examples of the use of this decomposition were presented. These included the use of the flow coordinates to provide preferred directions for edge detection, detection of anomalies, fitting simple models to the straightened pattern, and detecting singularities in the flow field.

Our ongoing work focuses on the analysis of patterns with multiple axes of anisotropy, statistical modeling and resynthesis of straightened patterns, and richer description of the structure of the flow field.

Appendix A: Derivation of Equation 5

The integrals in equation 4 can be evaluated fairly easily by expansion using equation 3 for V. The numerator of $\tan(2\phi)$ can be written

$$\int_0^\pi V(\theta)\sin(2\theta)d\theta = W * [2C_x^2\int_0^\pi \sin(\theta)\cos^3(\theta)d\theta$$
$$+4C_xC_y\int_0^\pi \sin^2(\theta)\cos^2(\theta)d\theta$$
$$+2C_y^2\int_0^\pi \sin^3(\theta)\cos(\theta)d\theta].$$

The first and third integrals are zero and the middle integral is $\pi/8$. Hence

$$\int_0^\pi V(\theta)\sin(2\theta)d\theta = (\pi/2)W * C_xC_y$$

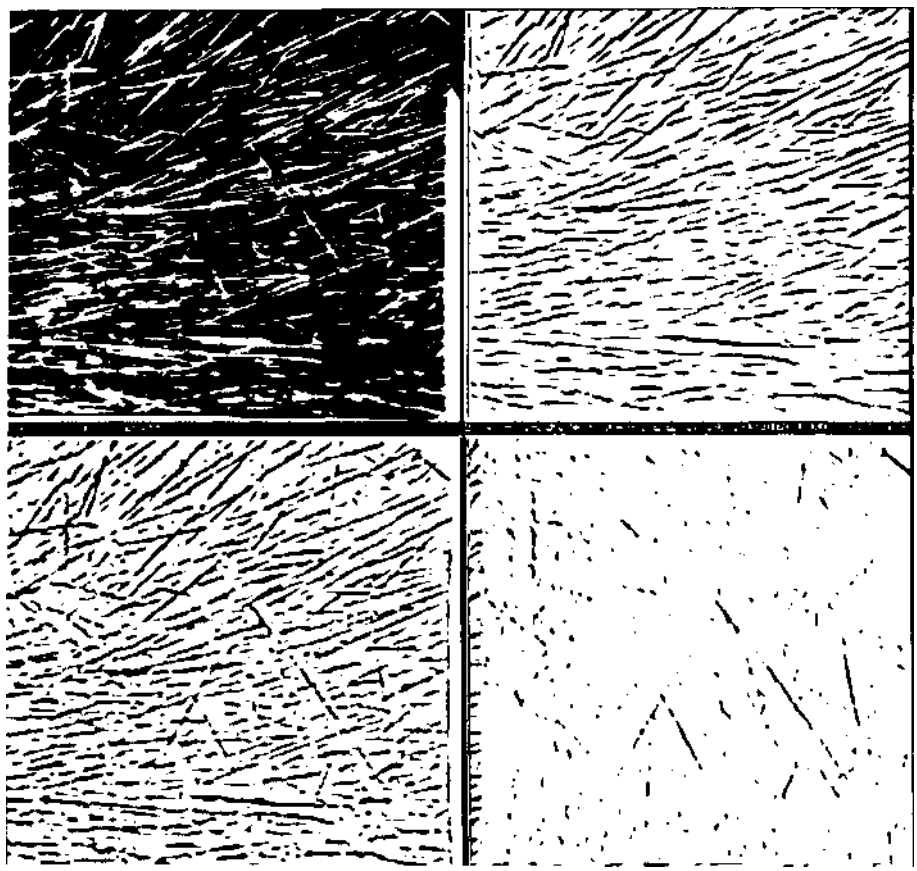

Figure 6: "Finding the needle in the haystack." In this straw pattern, directional derivatives across the flow direction show elements aligned with the pattern (upper right.) Those along the flow direction show anomalous elements (lower right.) The Laplacian (lower left) shows both.

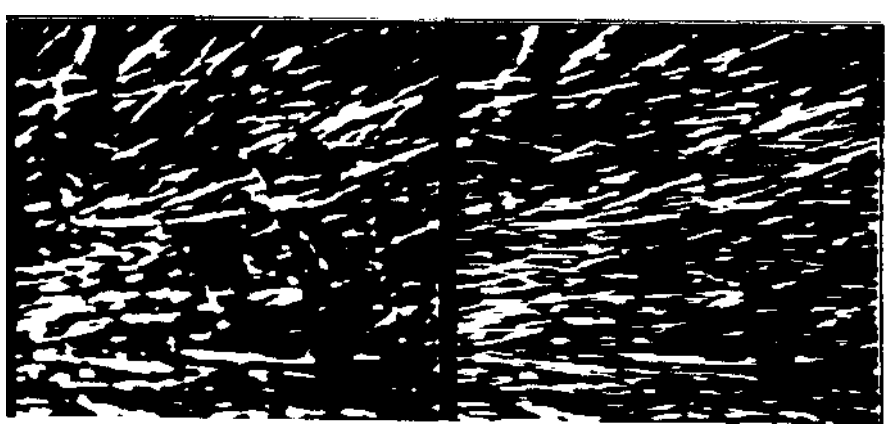

Figure 7: Left: the straw picture from figure 6. Right: the anomalous elements have been removed by directional median filtering in the flow direction. (Following a suggestion by Richard Szeliski.)

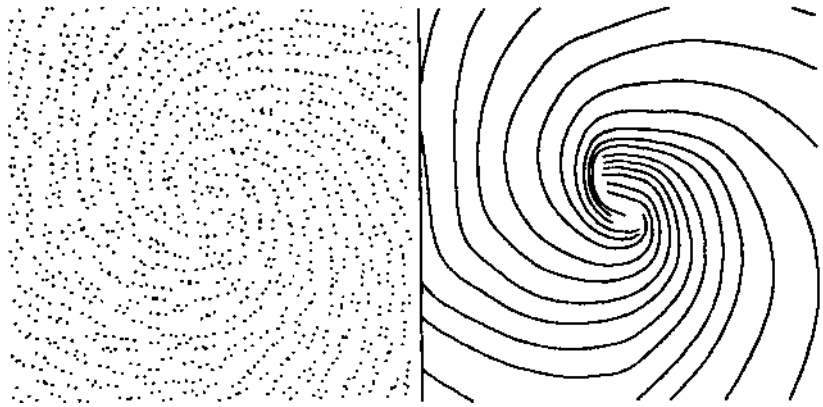

Figure 8: A spiral Glass pattern and its flow lines. The pattern is perceptually dominated by the singularity at the center. Since the flow field vanishes at a singularity, the flow lines obtained by integrating the flow field tend to become ill behaved as they approach one.

The denomenator of $\tan(2\phi)$ similarly is

$$\int_0^\pi V(\theta)\cos(2\theta)d\theta = W * 2C_x^2 \int_0^\pi 2\cos^4(\theta) - \cos^2(\theta)d\theta$$

$$+4C_xC_y \int_0^\pi \sin(4\theta)d\theta$$

$$+2C_y^2 \int_0^\pi 2\cos^2(\theta)\sin^2(\theta) - \sin^2(\theta)d\theta].$$

Here the first integral is $\pi/4$, the second integral is zero and the third is $-\pi/4$. Hence the above expression can be simplified to

$$\int_0^\pi V(\theta)\cos(2\theta)d\theta = (\pi/4)W * (C_x^2 - C_y^2).$$

Substituting for both the numerator and denomenator of $\tan(2\phi)$ in equation 4 yields

$$\phi = \tan^{-1}\left(\frac{W * 2C_xC_y}{W * (C_x^2 - C_y^2)}\right) . 2$$

There is a close connection between the flow computation described in section 2 and the density of edges in the Marr-Hildreth theory of edge detection. Suppose the image $I(x,y)$ is a stationary zero-mean Gaussian process and let $D(x,y)$ be the Laplacian of a Gaussian. The density of zero-crossings in $B_\theta(x,y) = D * R_\theta[I]$ along the x-axis can be approximated by (Papoulis, 1965; Rice, 1944-45)

$$\lambda^2(\theta) \approx \frac{1}{\pi^2}\frac{\int \omega^2 K(\omega)d\omega}{\int K(\omega)d\omega}$$

where $K(\omega)$ is the power spectrum of a slice through B_θ along the x-axis. Using Papoulis' Fourier transform conventions, the numerator can be converted to a spatial integral through the identity

$$\int B_x^2 dx = \frac{1}{2\pi}\int \omega^2 K(\omega)d\omega.$$

Similarly, the denomenator can be converted with the identity

$$\int B^2 dx = \frac{1}{2\pi}\int K(\omega)d\omega.$$

If the integrals are computed locally with the windowing function W, then we have the following estimate for zero crossing density:

$$\lambda^2(\theta) \approx \frac{1}{\pi^2}\frac{W * B_x^2}{W * B^2}$$

If W is radially symmetric, $W * B^2$ will not depend on θ so the maximum zero-crossing density will occur at the maximum of $W * B^2$. By assumption, the mean of the process is zero, so $W * B^2$ is the variance of $D_x * R_\theta[1]$. Thus, for stationary zero-mean Gaussian processes, if S — D_X, $A^2(0)$ $V(\theta)$.

References

Horn, B.K.P "Understanding image intensities," *Aritifical Intelligence,* 8, 1977, 201-231.

Barrow, H., & Tenenbaum, J. M. "Recovering intrinsic scene charactristics from images. In Hanson & Riseman (Eds.), *Computer Vision Systems.* New York: Academic Press, 1978.

Marr, D. Vision, 1982, Freeman, San Fransisco CA.

Barr, A. "Global and local deformation of solid primitives." *Computer Graphics,* 18, pp. 21 30, July 1984.

Brodatz, P. Textures. New York: Dover, 1966

Glass, L. "Moire effect from random dots." *Nature.,* 1969, 243, 578-580.

Hubel, D. H. & Wiesel, T. N. "Receptive fields, binocular interaction and functional architecture in the cat's visual cortex." *J. Physiol., Lond.,* 1962, 166, 106 154

Papoulis, *Probability, Random Variables and Stochastic Processes.* New York: McGraw-Hill, 1965.

Marr, D. "Early processing of visual information." *Proc. Royal Soc,* 1976, B 275, 484-519.

Marr, D., & Hildreth, E. "Theory of edge detection." *Proc. Royal Soc,* 1980, B 207, 187-217.

Marr, D., & Poggio, T. "A computational theory of human stereo vision." Proc. Royal Soc, 1979, 204, 301-328.

Rice, S. O. "Mathematical Analysis of Random Noise." *Bell Sys. Tech. J.,* 23-24, 1944-1945.

Schiller, P. H., Finlay, B. L. & Volman, S. F. "Quantitative studies of single-cell properties in monkey striate cortex. 1J. Orientation specificity and ocular dominance " ./. *Neurophysiology* 39, 1976, 1320-1333.

Spivak, *Differential Geometry*. Berkely, California: Publish or Perish, 1979.

Stevens, K. "Computation of locally parallel structure." *Biological Cybernetics,* 1978, 29 29-26.

Williams, L. "Pyramidal Parametrics" *Computer Graphics,* 17 No. 3, 1983

Zucker, S. "Computational and psychophysical experiments in grouping." In Beck, Hope, Rosenfold (Eds.*),Human and Machine Vision,* New York: Academic Press, 1983.

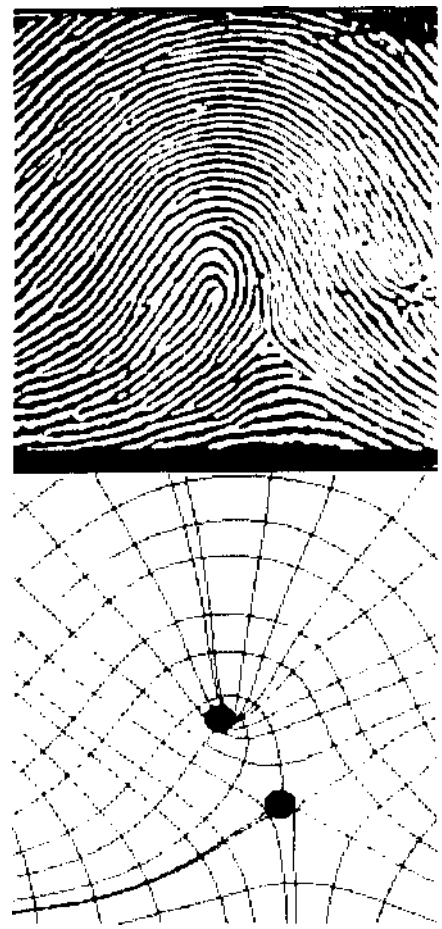

Figure 9. A fingerprint and its flow coordinate grid. The two white circles represent the major singularities, which were detected by measuring the flow field's winding number.