

The Perceptual Organization of Texture Flow: A Contextual Inference Approach

Ohad Ben-Shahar* and Steven W. Zucker†

Abstract

Locally parallel dense patterns - sometimes called *texture flows* - define a perceptually coherent structure of particular significance to perceptual organization. We argue that with applications ranging from image segmentation and edge classification to shading analysis and shape interpretation, texture flows deserve attention equal to edge segment grouping and curve completion.

This paper develops the notion of texture flow from a geometrical point of view to argue that local measurements of such structures must incorporate two curvatures. We show how basic theoretical considerations lead to a unique model for the local behavior of the flow and to a notion of texture flow “good continuation”. This, in turn, translates to a specification of consistency constraints between nearby flow measurements which we use for the computation of *globally* (piecewise) coherent structure through the contextual framework of relaxation labeling. We demonstrate the results on synthetic and natural images.

Keywords

Texture flow, Perceptual organization, Social conformity of a line, Good continuation, Texture segmentation, Line discontinuities, Point singularities, Shading flow, Local parallelism, Orientation diffusion, Tangential curvature, Normal curvature, Relaxation labeling.

*O. Ben-shahar is with the Computer Science department, Yale university, New Haven, Connecticut. Email: ben-shahahr@cs.yale.edu (corresponding author).

†S.W. Zucker is with the Computer Science and Electrical Engineering departments, Yale university, New Haven, Connecticut. Email: zucker@cs.yale.edu

1 Background

1.1 Texture flow and perceptual organization

A bear’s fur, a zebra’s stripes, and a field of wheat all define a visual structure, sometimes called *texture flow*, whose organization into coherent parts is fundamental to many aspects of computer vision. Informally, texture flows are defined by their orientation content - a dense visual percept characterized by local parallelism and slowly varying dominant local orientation (almost everywhere). This class of patterns is common in both natural and man-made objects (Fig. 1) and for centuries it has been used by artists as a tool to convey both the shape and shading of smoothly varying surfaces and their discontinuities.

In computer vision, parallel structure has been discussed mainly through the *non accidentalness* argument [62, 30], shape recovery [51, 31], or in relation to pattern formation processes [21]. In computational perceptual organization, though, it has played a secondary role despite ample psychological evidence demonstrating that the human visual system organizes and groups parallel structure into coherent units. Examples include the parallel structure that emerges from Glass patterns [10] and what Kanizsa dubbed the “social conformity of a line” [20]. In all cases, this organization dramatically affects our interpretation of both 2D and 3D scenes (Fig. 2).

Texture flows also appear to possess the same “gap completion” property that has been investigated so extensively for curves (e.g., [23, 60]). However, texture flow completion cannot be explained as multiple completions of individual curves. The problem is not only the combinatorics of correspondence between segment pairs along the occluder boundary, but also the fact that such a one-to-one correspondence does not necessarily exist (Fig. 2D). Similarly, it can be argued that perceptual filling-in of texture flow is insensitive to the exact locations at which the integral lines meet the occluder, an observation that stands in contrast to the qualitative changes (e.g., an emergence of inflection points) that such small perturbations may introduce to curve completions.

These observations strongly suggest that texture flow representations should be dense (i.e., continuous in the image plane). They further question what “good continuation” [59] means in the context of



Figure 1: Instances of texture flows can be found in a variety of visual stimuli, both natural and artificial.

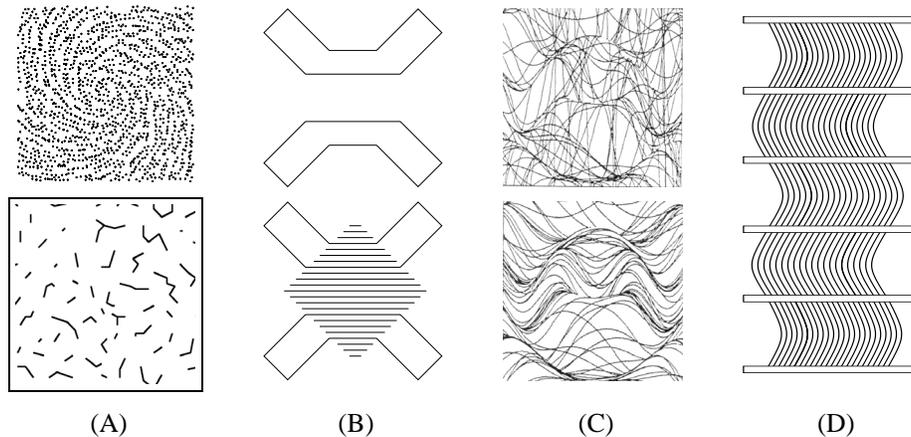


Figure 2: Psychological observations on texture flow. (A) Organization of locally parallel structure in Glass patterns occurs despite the fact that proximity local cues, depicted as a nearest neighbor graph on a zoomed portion of the pattern, convey no such structure. (B) *social conformity of a line* [20] drastically affects our interpretation of scenes. While the top is naturally perceived as two polygons, its superposition with a collection of parallel lines (bottom) is normally perceived as a diamond occluding a cross. The texture flow grouping thus dominates the continuity of the polygons. (C) Local parallelism facilitates 3D shape perception. Although both images depict the same surface, its shape is apparent only when the set of surface contours possess some degree of parallelism (reproduced from [53]). (D) This flow is amodally completed under the occluders and appears uniform and coherent. It requires a great deal of scrutiny to realize that each of its unoccluded segments contains a different number of lines, which precludes the completion of texture flow gaps by a line-by-line process.

texture flow, a better understanding of which is essential for segmentation. Flow patterns are commonly segmented in the neighborhood of *singularities*, which is sometimes correct (e.g., the folds in Dürer’s curtains, Fig. 1) but oftentimes not (Fig. 3). In particular, point singularities such as the one in the zebra’s scapular stripes are intrinsic characteristic of *uniform* processes of biological pattern formation [34] and do not represent object boundaries. Thus, in formalizing good continuation for texture flow one should account for the fact that *coherent* flows may contain singularities, and thus a wide range of orientations within a small region. As we will see, this can be done through a notion of curvature.

Psychophysical data also have been gathered about orientation sensitivity and orientation-based texture segmentation. The data clearly indicates that segmentation performance improves with the amount of spatial information that is available for the inference of a oriented structure. Much of the research has concentrated on the effect of elements’ density and length [36], and orientation gradient [36, 28]. Unfortunately, most psychophysical studies of texture flow have concentrated on flows of piecewise constant orientation, thus constraining themselves to the absolutely minimal model. This trend of ignoring the *varying* aspect of texture flow still prevails in recent studies (e.g., [44]). Although a preliminary work on texture flow segmentation in the presence of *everywhere changing* local orientation was discussed by Nothdurft [37], most of the psychophysics of texture flow remains unexplored.

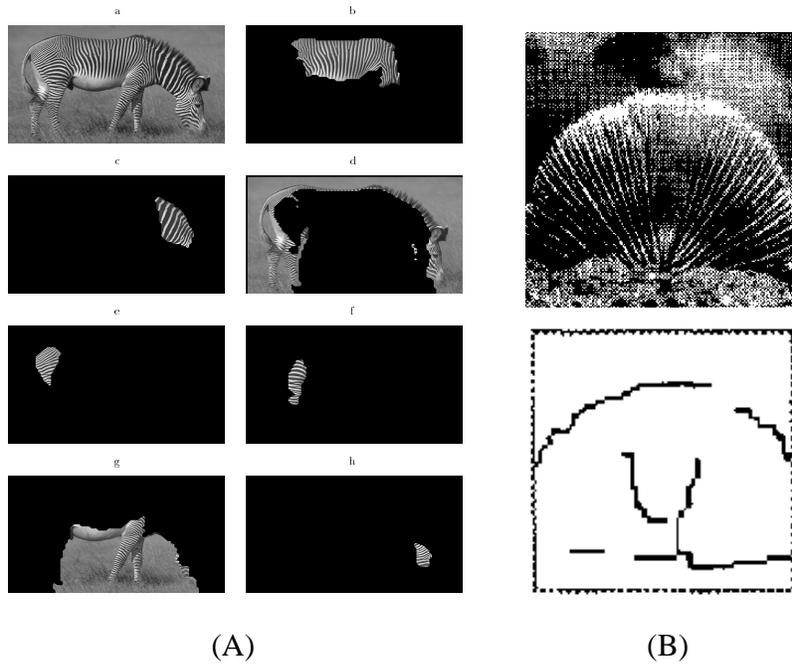


Figure 3: Segmentation of intensity images often produces undesired boundaries in the neighborhood of flow singularities. (A) A segmented zebra reproduced from [46, page 736]. (B) A segmented shell reproduced from [29, page 2651].

1.2 From texture to other flows

The notion of texture flow extends beyond textured objects to general imagery through the observation that other image features also have a dense, locally parallel structure of slowly varying dominant local orientation (a.e.). Prominent among these features is the *shading flow field* - the vector field that underlies the iso-brightness contours (or intensity level sets) of the gray level image of smooth surfaces [5]. The geometry of this field is a precursor for shape, and its interaction with edge geometry provides useful information for edge classification [17]. When applied to shading flow fields instead of textures, a successful organization of coherent flow structure implies the existence of a smoothly curved, untextured surface in the scene (Fig. 4). Since shading flow singularities relate to shape discontinuities, a reliable recovery of their geometry, while preserving their discontinuities and singularities, is important.



Figure 4: The shading flow in these two regions of interest (ROIs) is fundamentally different in terms of their smoothness and coherence. Thus, A method that can organize coherent flows may signal the existence of a smooth surfaces in the scene.

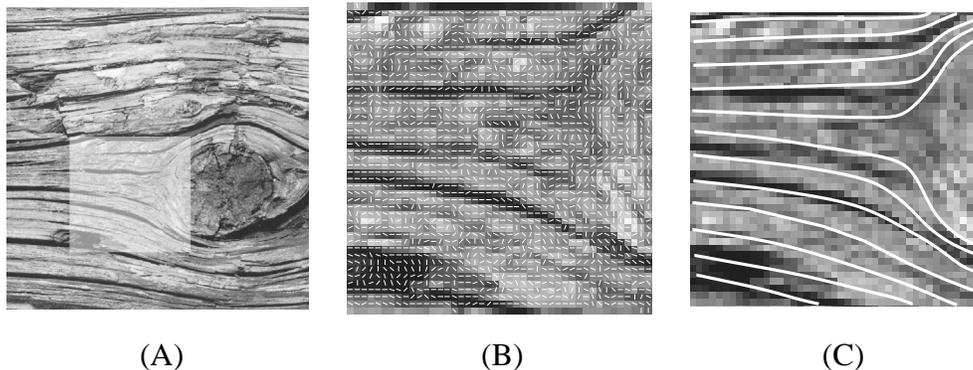


Figure 5: Example of the computational gap between the noisy, locally measured data and a desired coherent representation. (A) A tree stump image [6] and a ROI. (B) The field of dominant directions which was computed with a gradient based filter. (C) The underlying perceptual structure (drawn manually) is hardly evident in the noisy signal from the filter.

For similar reasons, the organization of texture flow into coherent units is equally important for the robust computation of optical flow and segmentation based on common fate [59]. Consequently, having a biologically plausible approach for the organization of texture flow like the one presented in this paper may suggest that perceptual organization-like processes participate in tasks that traditionally were not examined from this point of view.

1.3 The computational gap

To exploit the information conveyed by texture flows, one needs to analyse intrinsically noisy data to extract a piecewise coherent representation (Fig. 5). Thus, any analysis of texture flow needs to provide answers to questions such as (1) what should be measured locally; and (2) how can spurious measurements be refined into globally coherent structures while preserving and detecting “interesting events”, namely discontinuities and singularities. Since texture flows are perceptually dense even when the raw measurements are sparse, an adequate analysis must *fill in* holes and missing measurements. Since texture flows may overlap [56, 22], *multiple orientations* at a point should be allowed. Finally, non flow structures must be rejected rather than transformed into a meaningless flow structure. To our knowledge, no work to date possesses all these features.

1.4 Texture flow representation

Informally, texture flow may be defined as a two dimensional structure characterized by *local parallelism* and *slowly varying* dominant local orientation. As we argued above, psychological observations suggest that abstract representations of texture flow should be continuous (as opposed to discrete).

More formally, the notion of texture flow can be abstracted naturally in different ways:

- An orientation function $\theta(x, y) : \mathbb{R}^2 \rightarrow \mathcal{S}^1$ that defines a dominant orientation at each point (x, y) .
- A 2D submanifold of $\mathbb{R}^2 \times \mathcal{S}^1$. Using local charts to cover this manifold, or a global parametrization and an adequate periodic function of orientation (c.f. [48]), this representation can be embedded as a surface in \mathbb{R}^3 whose Z axis represents orientation (in this paper we label this space by $XY\theta$). A single valued texture flow patch then takes the form $s(x, y) = (x, y, \theta(x, y))$, i.e., up to 2π -periodicity in the θ direction it is a graph of a scalar function over \mathbb{R}^2 .
- A unit length vector field $I(x, y) : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, $I(x, y) = (I_1, I_2)$ with $I_1^2 + I_2^2 = 1$.
- A 2D *frame field* [38] $E(q) = \{E_T(q), E_N(q)\}$ over the image plane. Naturally, one would choose $E_T(q)$ as the unit vector field tangent to the flow, and $E_N(q)$ as normal to it.

Although all these representations¹ are basically equivalent, specifying the representation makes explicit different properties [31]. We make use of this observation below.

1.5 Related computational work

Locally parallel structures have been of interest to computer vision researchers for at least two decades. Following Glass’s work [10], Stevens [50] computed locally parallel structure by a histogram based approach on a discrete and sparse representation of the flow. Zucker [63] argued for a dense, vector field-based representation, emphasized the role of locally parallel structure in the context of perceptual organization, and linked the flow’s local behavior to its *global structure* via partial differential equations.

Few studies of oriented patterns [21, 43] used filtering methods to extract the flow’s local orientation and obtained a more reliable estimation of dominant orientation by appropriately averaging the filter’s response over a neighborhood. Alternatively, the flow’s orientation was also computed by fitting the filter responses to a first order model [42]. This and similar approaches yield a global description; they allow a classification of the flow pattern via its differential properties (c.f. [27]); and they are able to localize point singularities. However, they cannot handle multi-valued data (overlapping flows or textures of multiple dominant orientations) and they behave poorly along line discontinuities.

Following work on scale space [61, 26] and nonlinear diffusion of scalar [41, 57, 49, 3] and vector valued images [45, 49, 4], the trend in the analysis of oriented patterns has shifted from perceptual

¹Representations of *orientation fields* typically relate to 2π -periodicity. However, a simple mapping [40, 11] allows to represent the equally important *fields of directions* [8] and π -periodic signals. Hence, we use all these terms interchangeably.

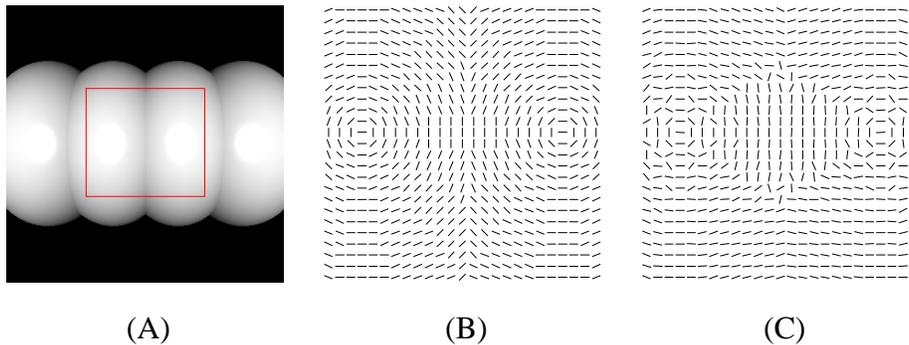


Figure 6: Example of a qualitative structure distortion by orientation diffusion. (A) A gray level image of a worm like object and a ROI. (B) The shading flow field of the object in the ROI. (for clarity we show a noiseless flow). It is clear that the vertical discontinuity, which corresponds to the surface discontinuity, must be preserved in any processing of the field, or else a misinterpretation of the shape is unavoidable. (C) The result of 100 iterations of anisotropic orientation diffusion (Tang *et al.* [52], $P = 1$). Note the qualitative distortion in the structure, the elimination of the vertical discontinuity, and the emergence of new singularities/discontinuities.

organization to formal scale space models of orientation diffusion. The aim of orientation diffusion is to denoise orientation fields and to obtain a multi scale representation of the data set. It is commonly acknowledged that this process should preserve “interesting” events (i.e., singularities in the structure) “as long as possible” in the scale space (e.g.,[49, p. 310]). However, since the signal lies in a non-Euclidean space, the periodic behavior of directional data entails special care in the definition of energy functionals and in operations like averaging and differentiation [40]. Tang *et al.* [52] used the theory of harmonic maps in liquid crystals to minimize the harmonic energy of a pattern by flowing toward its critical points via gradient descent. Using different norms one can change the sensitivity of this diffusion process to large scale discontinuities and hence the rate at which such structures in the orientation field are smoothed away. Related ideas were studied by Chan and Shen [7] and Kimmel and Sochen *et al.* [24].

Orientation diffusion methods are a powerful tool for progressively transforming a noisy set of measurements into a piecewise smooth representation. Their inhomogeneous nature allows them to adaptively damp their influence in areas “suspected” as structural edges and thus to preserve the prominent structure to deeper levels of the scale space. But, orientation diffusion methods are not designed to deal with “holes” or missing measurements, which undermines both occlusion and the more general fact that measurements may be intrinsically sparse. In addition, diffusion schemes do not deal with data containing *multiple* dominant orientations at a point. This prevents the analysis of multi directional textures (see [15, 6] for examples) or overlapping texture flows due to transparency [56]. Finally, orientation diffusion may distort certain kinds of perceptual discontinuities which do not conform to the edge-like discontinuity structure assumed by their edge stopping function (Fig. 6). We study this problem in detail elsewhere [2].

1.6 Overview of our approach

We propose an approach to the perceptual organization of texture flow that can handle sparse data sets and multi-valued inputs. It can fill in missing measurements over short distances and it can reject large scale non-flow structures. It is able to preserve both singularities and line discontinuities, including those which orientation diffusion is prone to distort.

Approaching the problem from a perceptual organization point of view, we first analyse texture flows geometrically. This associates meaning to “good continuation” and results in two new parameters that govern it - a *tangential* curvature and a *normal* curvature. Then, through the formalization of the intuitive definition of texture flow (Section 1.4), we derive local models and reveal a surprising formal link to orientation diffusion. Finally, we show how one of these models, which in $XY\theta$ takes the form of a *right helicoid*, is superior to the others in terms of exploitable invariances on the flow’s curvatures.

The advantage of having a model for the local behavior of a coherent texture flow lies in the ability to assess the degree to which a particular measurement is compatible, or consistent, with the context in which it is embedded, and whether or not that context is part of a single coherent object. Motivated by the columnar architecture of the primary visual cortex [16], we apply this contextual computation through a relaxation labeling network. We construct the compatibilities for our network from the right helicoidal model and show how it may be made stable along discontinuities. Finally, we demonstrate the result of this process both on synthetic and natural images, including its application to shading analysis.

2 The geometry of texture flow

A first step toward the perceptual organization of texture flow is understanding the parameters that govern its behavior. Since these objects are geometrical and continuous, and their characteristic property (i.e., orientation) is local, a differential analysis is appropriate. In this section we discuss how such an analysis necessarily results in a notion of texture flow curvature. We then turn to the issue of coherent flow and derive a model of “good continuation” from a formalization of the intuition in Section 1.4. The theoretical aspects developed here are then used for the computational system in Section 3.

2.1 Texture flow local behavior and measurements

Recall that informally, texture flow may be defined as a two dimensional dense (i.e., continuous) structure characterized by *local parallelism* and *slowly varying* dominant local orientation. A key aspect in this

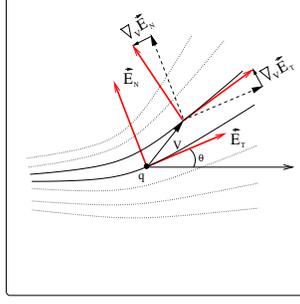


Figure 7: A texture flow can be represented as a differentiable frame field which is everywhere tangent (and normal) to the streamlines of the flow. An infinitesimal translation of the frame in a direction V rotates it by some angle determined by the connection form of the frame field. This rotation is fully characterized by two scalars obtained by the projection of the form on two independent directions. Using the frame itself this yields the two curvatures κ_T and κ_N .

intuition is that texture flow may not be constant. Although this may sound beyond dispute, there are areas of vision sciences, particularly psychophysics, where constant flows constitute a main body of research (e.g., [36, 28]). Thus we first try to formalize the notion of a *varying* dominant local orientation along with the parameters and the constraints that govern it.

Represented as a scalar orientation function $\theta(x, y)$, the local behavior of the flow is governed to first order by its gradient. Unfortunately, this quantity does not have an intuitive relationship to the flow's geometry. For this purpose, the frame field representation provides an insightful alternative.

Taken as a frame field over the image plane, the local behavior of the flow is characterized by the differential behavior of the frame as it moves over the image plane. In particular, one may seek a representation of this behavior in terms of the frame *itself*, thus achieving an object-centered view which is invariant to Euclidean transformations (c.f. the Frenet frame and the Frenet equation for curves [8]). This is captured by Cartan's *connection equations* [38]:

$$\begin{pmatrix} \nabla_V E_T \\ \nabla_V E_N \end{pmatrix} = \begin{bmatrix} 0 & w_{12}(V) \\ -w_{12}(V) & 0 \end{bmatrix} \begin{pmatrix} E_T \\ E_N \end{pmatrix} \quad (1)$$

where ∇_V is the *covariant derivative* in the direction V . The coefficient $w_{12}(V)$ is a function of the tangent vector V , which implies that the local behavior of the flow depends on the direction along which it is measured (Fig. 7). Fortunately, $w_{12}(V)$ is a *1-form* and thus linear. This allows us to fully represent it with two scalars at each point since $w_{12}(V) = w_{12}(a E_1 + b E_2) = a w_{12}(E_1) + b w_{12}(E_2)$. The freedom lies in choosing the basis to represent tangent vectors. In the flow-centered view, the natural choice is to set $E_1 = E_T$ and $E_2 = E_N$, defining the following two scalars :

$$\begin{aligned} \kappa_T &\triangleq w_{12}(E_T) \\ \kappa_N &\triangleq w_{12}(E_N) . \end{aligned} \quad (2)$$

We call κ_T the *tangential curvature* and κ_N the *normal curvature* - they represent the rate of change of the flow's dominant orientation in the tangential and normal directions, respectively. Since the differential of θ is also a 1-form, we can represent $w_{12}(V)$, κ_T , and κ_N all in terms of the appropriate projection:

$$\begin{aligned}\kappa_T &= d\theta(E_T) = \nabla\theta \cdot E_T = \nabla\theta \cdot (\cos\theta, \sin\theta) \\ \kappa_N &= d\theta(E_N) = \nabla\theta \cdot E_N = \nabla\theta \cdot (-\sin\theta, \cos\theta).\end{aligned}\tag{3}$$

In the language of frame fields, κ_T and κ_N are just the *coordinate functions* of $d\theta$ with respect to $\{E_T, E_N\}$. Since E_T and E_N are rigidly coupled, we can rewrite Eq. 3 in terms of E_T only using the standard curl ($\nabla \times$) and divergence ($\nabla \cdot$) operators:

$$\begin{aligned}\kappa_T &= \|\nabla \times E_T\| \\ \kappa_N &= \nabla \cdot E_T\end{aligned}\tag{4}$$

The representation of the differential behavior of the flow with (κ_T, κ_N) is not only more natural than its gradient but it also makes explicit the fact that oriented patterns change not only along their streamlines but across them as well. This stands in contrast to compromises on translational or rotational invariant models (e.g., [54]) where the normal curvature is completely omitted.

2.2 A Gestalt quality of texture flow

If κ_T and κ_N were known functions of position $q = (x, y)$, Eq. 3 could be viewed as a PDE and solved for $\theta(q)$. This raises the question of the degree to which κ_T and κ_N are independent. Are they completely independent, as curvature and torsion are for space curves, or are they dependent in the way the Gaussian and mean curvatures of surfaces are? We observe the following:

Proposition 1 : *Unless κ_T and κ_N both equal zero, they cannot be constant simultaneously in a neighborhood around q , however small, or else the induced flow is nonintegrable.*

This observation, a direct result of the more general Proposition 2, has an important implication: *in general, at least one of the curvatures must vary, or the two curvatures need to **covary** in any neighborhood of the texture flow.* Since κ_T is an ‘‘intra-streamline’’ property while κ_N is an ‘‘inter-streamline’’ property, Proposition 1 can be interpreted as a constraint that the flow as a whole imposes on the collection of streamlines. This Gestalt quality [58] is formally characterized as (see Appendix for the proof):

Proposition 2 : *Given any texture flow $\{E_T, E_N\}$, its curvature functions κ_T and κ_N must satisfy*

$$\nabla\kappa_T \cdot E_N - \nabla\kappa_N \cdot E_T = \kappa_T^2 + \kappa_N^2 .$$

2.3 Texture flow coherence

Since the local behavior of the flow is characterized (up to Euclidean transformation) by a pair of curvatures, it is natural to conclude that nearby local measurements of texture flow orientation should relate to each other based on these curvatures. Put differently, measuring a particular curvature pair $(\kappa_T(q), \kappa_N(q)) = (K_T, K_N)$ at a point q should induce a field of coherent measurements – i.e., an orientation function $\theta(x, y)$ – in the neighborhood of q . Clearly, that field, which we call the texture flow *osculating object*, should depend upon (K_T, K_N) . Coherence of texture flow measurements can then be determined in a manner analogous to cocircularity for tangents to a curve via the osculating circle [39].

Solving Eq. 3 with initial data at a single point q is equivalent to constructing the texture flow surface $s(x, y) = (x, y, \theta(x, y))$ in $XY\theta$ given only a single tangent plane at q . The set of possible solutions is infinite and the problem is underdetermined. We now make it well-posed by including the notion of “slowly varying dominant local orientation”. We formalize this by introducing minimization constraints on the behavior of $\theta(x, y)$. But which one? At this point the choice of representation plays a crucial role. The view of the flow as a 2D scalar function $\theta(x, y)$ suggests looking for the critical points of the harmonic energy

$$\int \int \|\nabla\theta\|^2 dx dy \quad (5)$$

in the spirit of Tang *et al.* [52]. On the other hand, the view of the flow as a surface suggests minimizing the surface tension and looking for critical points of the area functional

$$\int \int \sqrt{1 + \theta_x^2 + \theta_y^2} dx dy. \quad (6)$$

Even worse, not only does the choice of representation affect the choice of minimization and thus the possible solutions, but realizing again the kind of initial data we may use (i.e., orientation and curvature at a *single* point), either choice still leaves us with an infinite number of solutions.

Avoiding making this decision, and requiring that our sought after solution satisfies *both* constraints, results in a surprising conclusion:

Proposition 3 : *Assume (w.l.o.g.) that $q = (0, 0)$ and $\theta(0, 0) = 0$. A function $\theta(x, y)$ which satisfies $(\kappa_T(q), \kappa_N(q)) = (K_T, K_N)$ and is a critical point of both functionals (5) and (6) is either*

$$\begin{aligned} \text{a plane} & \quad \theta(x, y) = K_T x + K_N y, \\ \text{a right Helicoid} & \quad \theta(x, y) = \tan^{-1}\left(\frac{K_T x + K_N y}{1 + K_N x - K_T y}\right), \\ \text{or a left Helicoid} & \quad \theta(x, y) = \tan^{-1}\left(\frac{K_T x + K_N y}{1 - K_N x + K_T y}\right). \end{aligned}$$

The proof (see Appendix) is based on minimal surface theory [35] and on Hamel’s theorem for harmonic minimal surfaces [13, 12, 35]. The three solutions are visualized in Fig. 8. Naturally, all have the same tangent plane at q and one can show that the two helicoids even have identical local shape at q (i.e., the principle curvatures and directions of their surface representation coincide at q). Of the three, the plane is not intrinsically 2π -periodic but since it can be trivially duplicated to become one, we consider it a legitimate solution. Interestingly enough, the three solutions, henceforth called the *harmonic minimal solutions*, uniquely share another property:

Proposition 4 : *The harmonic minimal solutions are the only functions (that satisfy the initial data and) whose p -Laplacian $\Delta_p(\theta)$ vanishes simultaneously for all $p \geq 0$.*

This proposition (see Appendix for the proof and definition of $\Delta_p(\theta)$) reveals a formal relationship to [52] since it implies that each of the harmonic minimal solutions is a stationary point of a gradient flow of the p -harmonic energy of the orientation function $\theta(x, y)$, regardless of the chosen metric p .

2.4 Curvatures covariation revisited and texture flow invariances

The discussion so far suggests two main observations with regard to the modeling and computation of coherent texture flow: (1) that texture flow measurements should include two curvature measurements (in

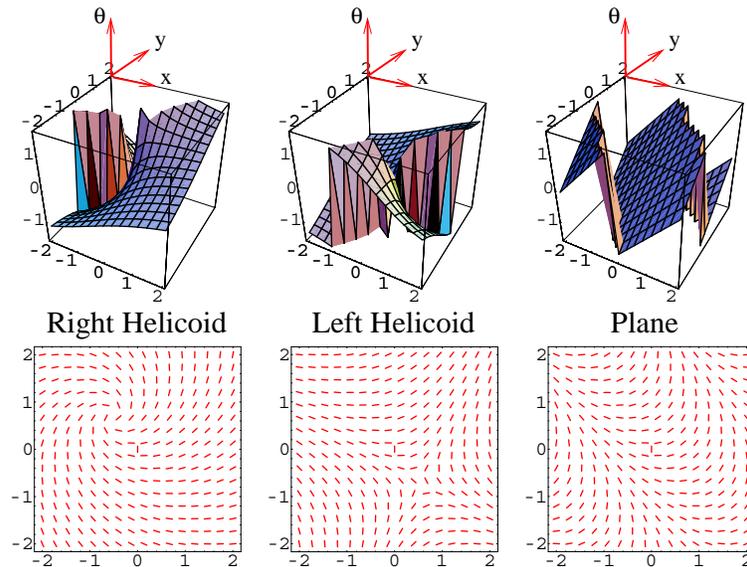


Figure 8: The only three solutions of Eq. 3 that are simultaneous critical points of the energy expressions (5) and (6), and that satisfy the initial conditions at q . We illustrate these three solutions for $K_T = \kappa_T(q) = 0.9$ and $K_N = \kappa_N(q) = 0.7$ both as surfaces in $XY\theta$ and as flows in the image plane. To set a visual context, we plot these solutions in a large neighborhood to show the singular point of the helicoids and the periodicity over the range $(-\frac{\pi}{2}, \frac{\pi}{2}]$.

addition to position and orientation), and (2) that nearby measurements should be considered coherent if they are part of, or “close” to, the same osculating object. A question remains whether or not one of the three candidate functions from Proposition 3 is a better choice than the others.

As mentioned in Sec. 2.2, any nontrivial texture flow, and thus any osculating object, must incorporate some covariation of κ_T and κ_N around the point of measurement q . Representational and computational considerations dictate *simple* and *unbiased* covariation; thus it is striking that

Proposition 5 : *Of the three solutions of Proposition 3 (and of all orientation functions in general), the right helicoid has the “simplest” possible covariation of κ_T and κ_N since it uniquely satisfies*

$$\frac{\kappa_T(x, y)}{\kappa_N(x, y)} = \text{const} = \frac{\kappa_T(q)}{\kappa_N(q)} = \frac{K_T}{K_N} \quad \forall (x, y) \in N(q).$$

It is immediate to verify that the right helicoid has a constant ratio of curvatures by substituting its orientation function into Eq. 3. The fact that this behavior is unique to the right helicoid can be derived from the same equation by imposing constant covariation on κ_T and κ_N and reducing the system to a single quasi-linear differential equation $\vec{n} \cdot (\kappa_N(q) \cos \theta + \kappa_T(q) \sin \theta, \kappa_N(q) \sin \theta - \kappa_T(q) \cos \theta, 0) = 0$ where \vec{n} denotes a normal to the solution surface in $XY\theta$. One can then show that the characteristic curves of this PDE in $XY\theta$ constitute a ruling that defines a right helicoid.

The constant covariation of κ_T and κ_N puts both curvatures on equal ground and has an important computational advantage since for most purposes only one curvature (say, κ_T) needs to be explicitly maintained. Differentiating this constraint reveals the conservation of covariation at much deeper levels:

Proposition 6 : *Let $\kappa_T(x, y)$ and $\kappa_N(x, y)$ be the flow curvature functions induced by the right helicoid through Eq. 3. Let D^{ij} be the differential operator $D^{ij} \triangleq \frac{\partial^i}{\partial x^i} \frac{\partial^j}{\partial y^j}$. Then*

$$\frac{D^{ij} \kappa_T(x, y)}{D^{ij} \kappa_N(x, y)} = \text{const} = \frac{K_T}{K_N} \quad \forall i, j \geq 0 .$$

Another direct consequence of the constant ratio of curvatures is an object-centered flow invariance; everywhere on the flow the direction along which it changes the most is constant relative to the direction of the flow itself. More formally:

Proposition 7 : *If $\theta(x, y)$ is a texture flow induced by the right helicoid (Proposition 3), then*

$$\cos \alpha(x, y) = \frac{\nabla \theta}{\|\nabla \theta\|} \cdot E_T = \text{const}$$

where $\alpha(x, y)$ is the angle that $\nabla \theta$ makes with the flow’s frame.

Finally, it is worth noting that the helicoid, as an osculating object, bears an interesting relationship to another osculating object used for curve detection. It has been proposed [39, 47] that given the orientation and curvature of a curve at a point, a preferred model for its behavior around that point is the *osculating circle* [8]. In $XY\theta$, a circle becomes a helix of pitch 2π . More than the fact that a helicoid is a collection of helices, the interesting link lies in the behavior of curvatures, since a helix is the only curve whose ratio of “curvatures” (curvature and torsion) is constant [9].

3 Analysis via relaxation labeling

The advantage of having a model for the local behavior of a coherent texture flow lies in the ability to assess the degree to which a particular measurement is compatible, or consistent, with the context in which it is embedded. This, in turn, can be used to refine noisy measurements, remove spurious ones, and fill in “holes” so that local ambiguity is reduced and global structures become coherent.

A natural framework with which one can pursue this task while maximizing the average consistency over a domain of interest is *relaxation labeling*. We developed such a relaxation network for the organization of coherent texture flows and derived the compatibility function which governs its behavior from the right helicoidal model. This section describes this system in detail.

3.1 What is relaxation labeling

Following early studies on contextual constraints in the interpretation of line drawings [55], it has become widely acknowledged that interpretation of sensory data is highly unreliable unless carried out in some non-local context. Relaxation labeling [18, 25] is a formal computational framework for doing exactly that. Closely related to popular models of neural networks [14] and polymatrix games [32], it involves representing the interpretation problem as the assignment of labels to nodes in a graph whose edges represent the contextual structure.

Let $I = \{i \mid i = 1..n\}$ be a set of nodes, each of which may take any label λ from the set Λ . Let $p_i(\lambda)$ denote the probability, or confidence in the assignment of label λ to node i . In general, this measure needs to satisfy two constraints:

$$p_i(\lambda) \geq 0 \quad \forall i, \lambda \quad \text{and} \quad \sum_{\lambda \in \Lambda} p_i(\lambda) = 1 \quad \forall i.$$

Of all possible assignments (henceforth denoted as the space \mathbb{K}), those which satisfy $p_i(\lambda) \in \{0, 1\} \quad \forall i \in I, \forall \lambda \in \Lambda$ are called *unambiguous* since they assign a unique label to each node. The role of relaxation

labeling is to start from a given, typically ambiguous labeling assignment, and iteratively change it toward a “better” one, where better refers to the degrees of ambiguity and consistency of the assignment of labels at different nodes.

The fundamental mechanism by which contextual information is propagated in a relaxation labeling network is a *compatibility function* $r_{ij}(\lambda, \lambda')$ which quantifies the contextual information conveyed by label λ' at node j about label λ at node i . In conjunction with the confidence measure p_i , the *contextual support* that a label λ at node i received from its neighborhood is (almost exclusively) defined to be:

$$s_i(\lambda) = s(i, \lambda; \bar{r}, \bar{p}) = \sum_{j=1}^n \sum_{\lambda'=1}^m r_{ij}(\lambda, \lambda') p_j(\lambda'), \quad (7)$$

which can be viewed as a sum of all compatibilities weighted by the confidences. In the same spirit, we can average the supports themselves to yield a scalar measure of consistency across the entire network, or what is traditionally called the *average local consistency*

$$A(\bar{p}) = \sum_i \sum_{\lambda} p_i(\lambda) s_i(\lambda) = \sum_i \sum_{\lambda} \sum_j \sum_{\lambda'} p_i(\lambda) r_{ij}(\lambda, \lambda') p_j(\lambda'). \quad (8)$$

Finally, it remains to clarify when an assignment \bar{p} is *consistent*. Intuitively, this happens when the labels assigned at each node maximally agree with their context. Had we dealt with unambiguous assignments only, this could have been expressed as a maximization of support at all nodes simultaneously, i.e.,

$$s_i(\lambda) \geq s_i(\lambda') \quad \forall i \in I.$$

For ambiguous assignments, this is expressed via the following set of inequalities:

$$\sum_{\lambda} p_i(\lambda) s_i(\lambda; \bar{p}) \geq \sum_{\lambda} \tilde{p}_i(\lambda) s_i(\lambda; \bar{p}) \quad \forall \tilde{p}_i \in \mathbb{K}, \quad \forall i \in I.$$

Using all this machinery, relaxation labeling maps inconsistent labeling to consistent ones via an iterative process. Here we follow the algorithm by Hummel and Zucker [18] which is described by the following confidence update rule:

$$p_i^{t+1}(\lambda) \leftarrow \Pi_{\mathbb{K}} [p_i^t(\lambda) + \delta s_i^t(\lambda)], \quad (9)$$

where $\Pi_{\mathbb{K}}$ is a projection operator that projects its argument onto \mathbb{K} , and δ is a constant step size. A fundamental result from the theory of relaxation labeling relates it to polymatrix game theory, and states that this algorithm converges to consistent labelings [18] related to the Nash equilibria of the polymatrix game [32]. Furthermore, if the compatibility function is symmetric $r_{ij}(\lambda, \lambda') = r_{ji}(\lambda', \lambda)$, the update rule constitutes a gradient ascent on $A(\bar{p})$ that maximizes it locally [18].

3.2 A relaxation network for texture flow

A direct abstraction of the relaxation process for texture flow should involve a 2D image-based network of nodes $i = (x, y)$ (i.e., pixels) whose labels are drawn from the set

$$\Lambda = \{(\theta, \kappa_T, \kappa_N) \mid \theta \in (-\frac{\pi}{2}, \frac{\pi}{2}], \kappa_T, \kappa_N \in [K_{min}, K_{max}]\} \cup \{\text{no-flow}\}$$

after it has been quantized appropriately. Motivated by the columnar architecture of the visual cortex [16] and in order to accommodate for either “no-flow” or multiple flows at a pixel, we replace this abstraction with a 5D network of nodes $i = (x, y, \theta, \kappa_T, \kappa_N)$ whose labels are either T (*TRUE*) or F (*FALSE*). This two-label paradigm was first suggested by Parent and Zucker [39] for their curve inference relaxation labeling algorithm. For each node i , $p_i(T)$ denotes the confidence that a texture flow of orientation θ and curvatures κ_T, κ_N passes through pixel (x, y) . Since by definition of the confidence measure $p_i(F) = 1 - p_i(T)$, we need to maintain and update the confidence of only one label at a node.

Under the two-label relaxation labeling, and providing we make one more intuitive design decision, most of the machinery reviewed in Section 3.1 above takes a particularly convenient form. In particular, we impose the following behavior on the compatibility structure between any two nodes:

$$\begin{aligned} r_{ij}(F, T) &= -r_{ij}(T, T) \\ r_{ij}(F, F) &= +\beta r_{ij}(T, T) \\ r_{ij}(T, F) &= -\beta r_{ij}(T, T), \end{aligned} \tag{10}$$

where $0 \leq \beta \leq 1$ is a constant. These constraints formalize the intuition that if an evidence at node j is *compatible* with a flow at node i , it is equally *incompatible* with the *lack* of flow at node i . In addition, they imply that the lack of evidence is β times less informative than its existence. With these general guidelines the support function at each node now becomes

$$s_i(T) = \sum_{j=1}^n r_{ij}(T, T) [(1 + \beta)p_j(T) - \beta] = -s_i(F).$$

The relationship $s_i(T) = -s_i(F)$ is particularly useful from the point of view of the gradient ascent rule (Eq. 9). It implies that except for the boundary of \mathbb{K} , the gradient ascent procedure stays always inside the space and thus requires no projection. In other words, the projection operation $\Pi_{\mathbb{K}}$, which otherwise may become much more complicated [33], reduces to

$$\Pi_{\mathbb{K}}(x) = \Pi_0^1(x) \triangleq \min(1, \max(0, x)). \tag{11}$$

As a result of the two-label paradigm and the compatibility constraints (10), only quantities that depend solely on the T label need to be represented explicitly. Thus we use p_i, s_i , and r_{ij} to refer to

$p_i(T)$, $s_i(T)$, and $r_{ij}(T, T)$, respectively, and rewrite the relaxation rule (Eq. 9) as

$$p_i^{t+1} \leftarrow \Pi_0^1 [p_i^t + \delta s_i^t] = \Pi_0^1 \left[p_i^t + \delta \sum_j r_{ij} ((1 + \beta)p_j^t - \beta) \right] \quad (12)$$

with t being the iteration number. Note that choosing $\beta = 0$ to reflect the assignment of zero information to lack of evidence further simplifies the relaxation rule to

$$p_i^{t+1} \leftarrow \Pi_0^1 \left[p_i^t + \delta \sum_j r_{ij} p_j^t \right]. \quad (13)$$

3.3 Geometric compatibilities for texture flow

It remains to specify a discrete compatibility function r_{ij} in order to complete the design of the relaxation network. Following our discussion in Section 2, we derive these compatibilities from the right helicoidal model, with an emphasis on two issues; the effect of quantization and boundary stability.

Measurement quantization dictates that every possible node i represents an *equivalence class* C_i of measurements, each of which induces a helicoidal field of compatible labels in the neighborhood of i . In the continuum, the union of all these helicoidal fields that correspond to points in C_i forms a consistent 5D “volume” V_i . After quantization, this volume results in a set of consistent, and thus excitatory labels. It is important to mention that since each node i corresponds to a different set of position, orientation, and curvatures, the shape of the 5D consistency volume V_i around i will vary with i .

While the consistency volume V_i determines *which* of the nodes are compatible with a given node i , it does not determine the *values* of r_{ij} which represent the likelihood that node j shares the same local model (i.e., a helicoid) with node i . Thus, we set r_{ij} to be the probability that C_j intersects the helicoid of a randomly selected point in C_i . In other words, if P_i is a random variable uniformly distributed in C_i , and $H(P_i)$ is the helicoid associated with it, we set

$$r_{ij} = \text{Prob} [H(P_i) \cap C_j \neq \emptyset].$$

In practice, we compute these values directly by sampling. We also surround the excitatory region with an inhibitory surround to compensate for the lack of intrinsic normalization of labels between all the nodes in the 5D network that share the same spatial position². Example of one such compatibility field, for a particular value of K_T and K_N , is illustrated in Fig. 9.

²The relaxation worked well even without such an inhibition but there was the occasional tendency to “leak” or diffuse labels along orientation columns if many iterations were needed.

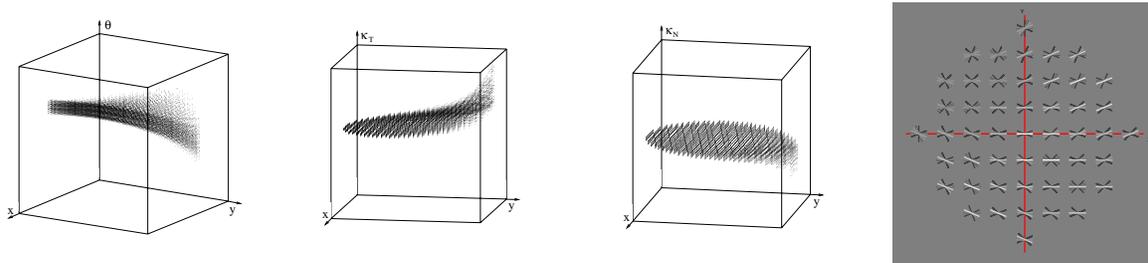


Figure 9: A visualization of the 5D consistency volume corresponding to particular measurements of $\theta = 0$, $\kappa_T = 0.3$ and $\kappa_N = -0.1$. Shown are projections of the volume on the 3D subspaces $XY\theta$, $XY\kappa_T$, and $XY\kappa_N$. The 2D image depicts the same object but now as oriented segments on the image plane. Brightness represents degree of compatibility and the inhibitory surround is shown in black (see text).

3.4 Support normalization for boundary stability

Texture flow *line discontinuities* are curves in the image plane along which the flow undergoes a discontinuity or termination. An important part of the computational task is to preserve these structures. As is implied by our geometrical analysis, and exemplified in Fig. 6, unless we attempt to treat discontinuities not only in orientation but also in curvature, the result might be a distortion of structure. Due to the high dimensionality of the space in which the relaxation labeling process operates, Most line discontinuities, either in orientation or curvature, can be viewed as texture flow boundaries. This is simply because in the 5D assignment space, each of the structures that meet along a discontinuity lies outside the context of its counterpart, and thus both are mutually uninfluential. Consequently, what the relaxation process is required to exhibit is stability along *boundaries* in the assignment space.

One way to achieve such stability is to introduce non-linearities in the support function $s_i(\lambda)$. This

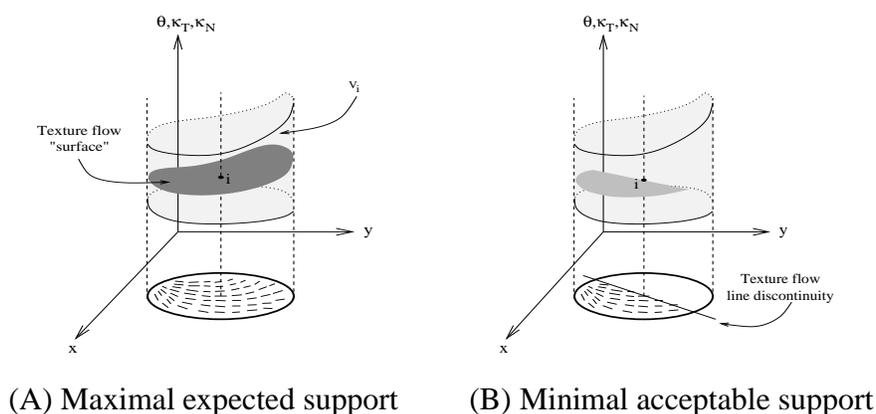


Figure 10: Practical stability of the relaxation labeling process to line discontinuities in the flow can be achieved though the normalization of the support function from the interval $[s_{min}, s_{max}]$ to $[0, 1]$. (A) At each node i , s_{max} is determined by integrating the support gathered from a full confidence, compatible flow that transverse the entire consistency volume V_i . (B) The minimal accepted support s_{min} reflects a flow of some minimally accepted confidence $\rho_{min} < 1$ (depicted here by the brighter surface intensity) that terminates along a line that intersects i .

may allow for great flexibility in the type of discontinuities we can handle, but it may significantly complicate the design and it may also affect the applicability of the fundamental results of relaxation labeling theory. Another approach, which is attractive mostly for curved discontinuities, is to endow the assignment space with labels that correspond to the discontinuity itself, e.g., its local orientation and curvature. This, however, has the disadvantage of further increasing the dimensionality of the assignment space and raising the computation time significantly. To stay within the simpler but better understood framework of linear support (Eq. 7) and to avoid an assignment space of dimensionality higher than 5, we equip our network with some degree of boundary stability by considering the minimum and maximum support that is expected from a coherent context and then normalizing the support function accordingly.

Given the compatibility volume V_i which corresponds to a particular texture flow node i , one can compute the integral of the compatibility coefficients assuming that a single, full confidence flow traverses through V_i . We denote it by s_{max} (Fig. 10). It is clear that the closer i is to a discontinuity or a boundary in the data, less context supports it. At the boundary, the process should remain stable, i.e., neither grow, nor shrink. Thus, one can define the level of support for which no change in confidence is desired. We denote this level by s_{min} and observe that it depends on both the geometry of the discontinuity and the minimally accepted confidence of the existing structure. Here we assume the simplest model of discontinuity, namely one that locally occurs along a straight line. The support one can gather from such a structure of minimally accepted average confidence ρ_{min} (Fig. 10) can then be approximated by $s_{min} = \frac{\rho_{min}s_{max}}{2}$. Normally, ρ_{min} would be set to 0.5, which is the minimal confidence that cannot be disambiguated as the *TRUE* label. Since noise may decrease the number of measurements that contribute to the minimally accepted support, decreasing ρ_{min} accordingly can provide stability around boundaries even in the presence of noise.

In the context of the two-label relaxation labeling paradigm and the gradient ascent update rule (Eq. 13), decrease in the confidence of a label occurs only if $s_i < 0$. Since we seek stability for all supports in the range $[s_{min}, s_{max}]$, it remains to normalize the support values by mapping this interval to the unit interval $[0, 1]$ before submitting it to the update rule. This can be done via the transformation

$$s_i \leftarrow \frac{s_i - s_{min}}{s_{max} - s_{min}} .$$

3.5 Analysis of multi-oriented texture patterns

A unique feature of the two-label relaxation labeling network is the ability to handle data sets that have more than one dominant orientation at a point. Such data sets can emerge from textures that are inherently

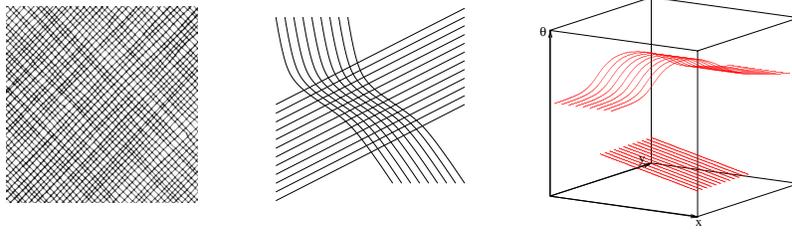


Figure 11: Overlapping texture flows. (A) An example of a multi-directional texture (from [15]). (B) An example of texture laciness [56]. (C) In both cases, the overlapping flows are represented as separated flows whose projection to the image plane share a common region. In terms of the relaxation, though, the different coherent parts maintain a big separation in the assignment space thus they remain mutually uninfluential. Shown here is the $XY\theta$ representation of the stimulus in B.

multi-directional (Fig. 11A. See also [15, 6] for many other examples) or from texture flows that lie on transparent and overlapping surfaces. The latter phenomenon has been termed *texture laciness* [56]. Given such a stimulus, multi-valued texture flows can be extracted from it via a population of oriented filters. In our experiments, we used the Logical/Linear edge/line detector [19], which by design can return multiple edgels at a point.

Once the relaxation labeling network has been initialized by a process that can deal with multi-orientation patterns, the computational process cannot, and need not, distinguish between a single texture of multiple orientations and an overlapping arrangement of *distinct* texture flows. In fact, with only the geometrical (i.e., orientation) information in hand, such a distinction can be decided only by a more global process similar to the way transparency is decided for a collection of intensity patches [20]. From the point of view of the relaxation network, every multi-directional texture is perceived, and thus analyzed, as a collection of *separate* texture flows that happen to share the same spatial region. This separation is achieved at no extra cost as a result of the high dimensionality of the space in which the relaxation operates. In particular, overlapping texture flows are likely to have different orientations wherever they overlap, thus the representations of these flows in the 5D assignment space will be mutually uninfluential (Fig. 11B,C).

4 Experimental results

We tested the proposed model and relaxation system on a variety of inputs, both synthetic and natural. In all cases we defined our network and compatibility fields to be π -periodic (i.e., to relax fields of directions, rather than orientations) with 8 equivalence classes, and to support curvatures in the range $[-0.2, 0.2]$ with 5 equivalence classes. This section discusses our results.

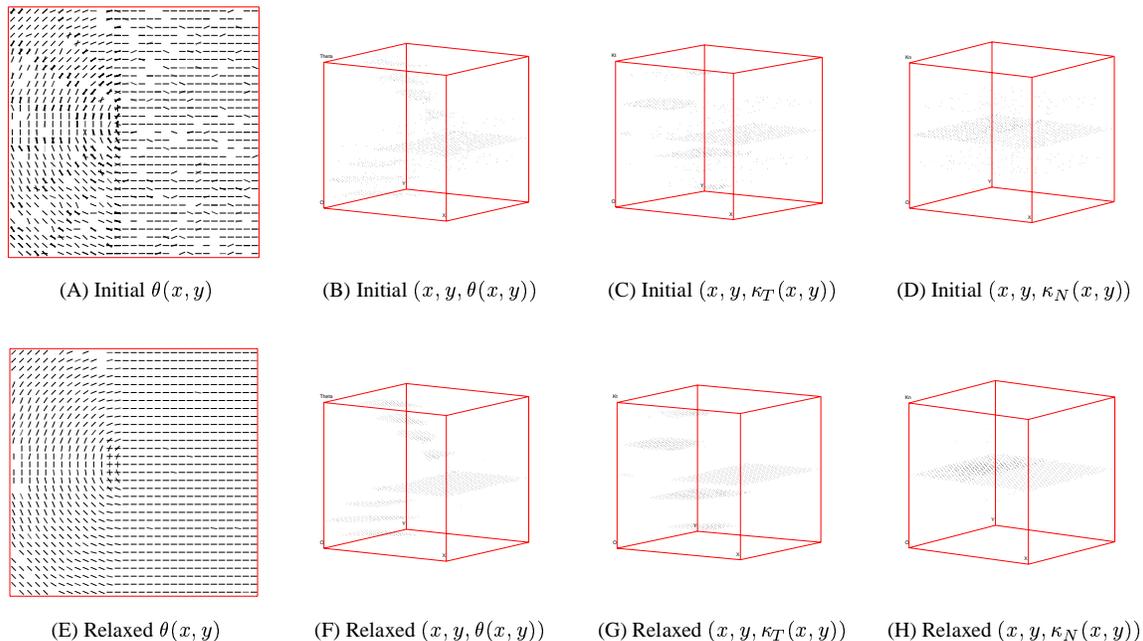


Figure 12: The organization of texture flow via relaxation labeling based on right helicoidal compatibilities. (A) A noisy synthetic texture flow composed of two parts. (B-D) Initial distribution of orientation and curvatures measurements. Note that the Z axes represents orientation, tangential curvature, and normal curvature, respectively. (E-H) Relaxed distributions ($\delta = 0.25$, 10 iterations). Missing measurements were filled in while noisy measurements were filtered out, making the entire structure globally coherent while keeping its discontinuity intact.

Fig. 12 shows an example of a synthetic texture flow composed of two parts. For such flows we computed exact initial measurements and then corrupted them along the orientation and curvature dimensions with additive and/or salt-and-pepper noise³. Although the pattern in this figure is mostly continuous in terms of orientation, there is a vertical line of discontinuity in the tangential curvature. As we discussed in Section 1.2, such discontinuities often occur in shading flow fields and their preservation is crucial for accurate shape from shading analysis and part decomposition. As can be seen in the figure, the relaxation labeling process is able to eliminate the noise, to fill in holes, and to converge to a globally coherent structure while preserving the discontinuity. Fig. 12 also provides a visualization of the change in the distribution of labels in the 5D assignment space (shown by its 3D projections). For lack of space we include this kind of visualization for one case only.

To examine the robustness of our approach in the presence of noise, we conducted an exhaustive test by progressively adding noise to a constant flow and checking the behavior of the relaxation network for a fixed number of iterations. The results are reported in Fig. 13. Observe the large amount of noise

³Note that the figure illustrates the noisy signal *after* it was quantized to fit the quantization level of the network and to reflect the possible flow labels that the relaxation labeling process can manipulate.

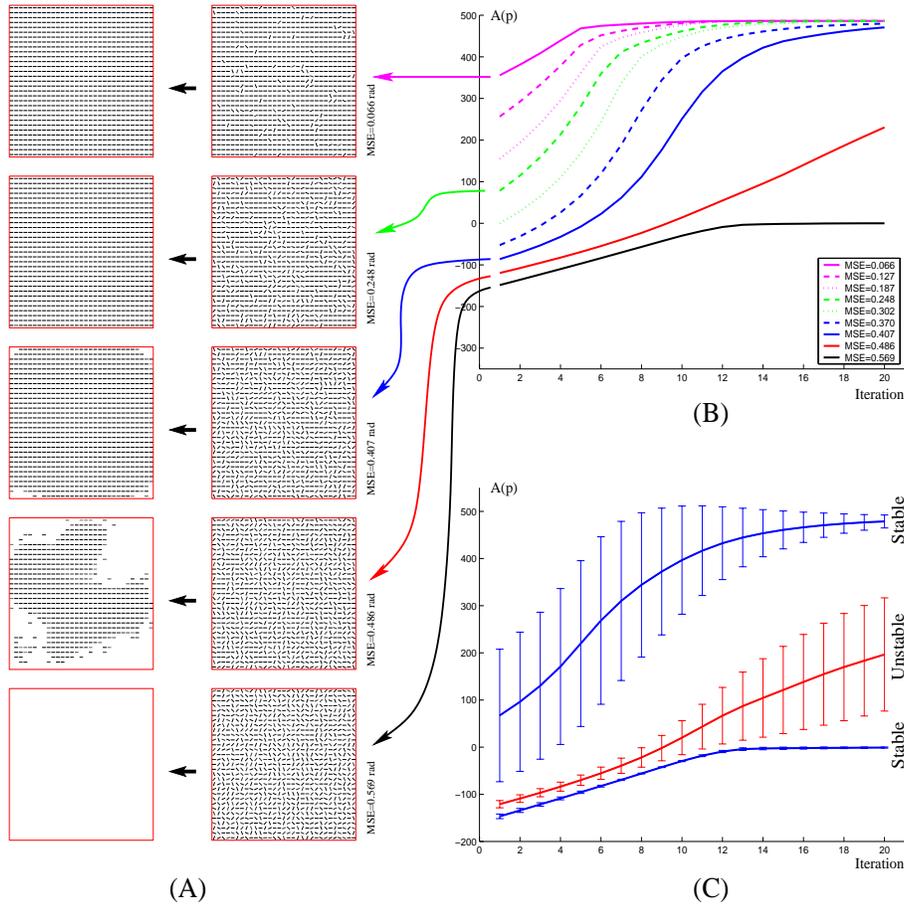


Figure 13: Robustness to noise in texture flow computation. (A) Each row illustrates a progressively more noisy texture flow along with the result of 20 iterations of the relaxation process. Failure to extract the original signal occur well beyond the perceptual threshold. (B) An illustration of the dynamics of $A(\bar{p})$ as a function of initial noise level and number of iterations. Note how for most noise levels the process converges to a particular average local consistency which represents the original signal. (C) Further experiments reveal three distinct zones; two stable ones, in which even large variations in initial data are monotonically suppressed till convergence to one predictable outcome; and one unstable, in which small variations are amplified and the result is unpredictable. Graphs represent mean $A(\bar{p})$ from 75 runs. Bars are 1 s.d.

for which the system is able to reconstruct the signal (panel A) which is best illustrated through the behavior of the average local consistency (panel B). In most cases, and below a certain level of noise ($MSE \leq 0.407$), the computation quickly converges to one particular average local consistency which represents the noiseless signal. Beyond a certain point ($MSE \geq 0.569$), it converges to a different value which represents the lack of flow everywhere. These two stable regions surround a narrow range of noise levels for which the behavior is unstable, i.e., the results are unpredictable and very sensitive to the initial conditions. This indeed was verified experimentally by repeating the test for a large number of trials, as is illustrated in panel C. Fig. 13 also indicates a stopping criterion for the relaxation iteration. As the average local consistency $A(\bar{p})$ approaches convergence, its time derivative (i.e., slope) quickly approaches zero. Thus, we stopped our relaxation when $\frac{d}{dt}A(\bar{p})$ was sufficiently small.

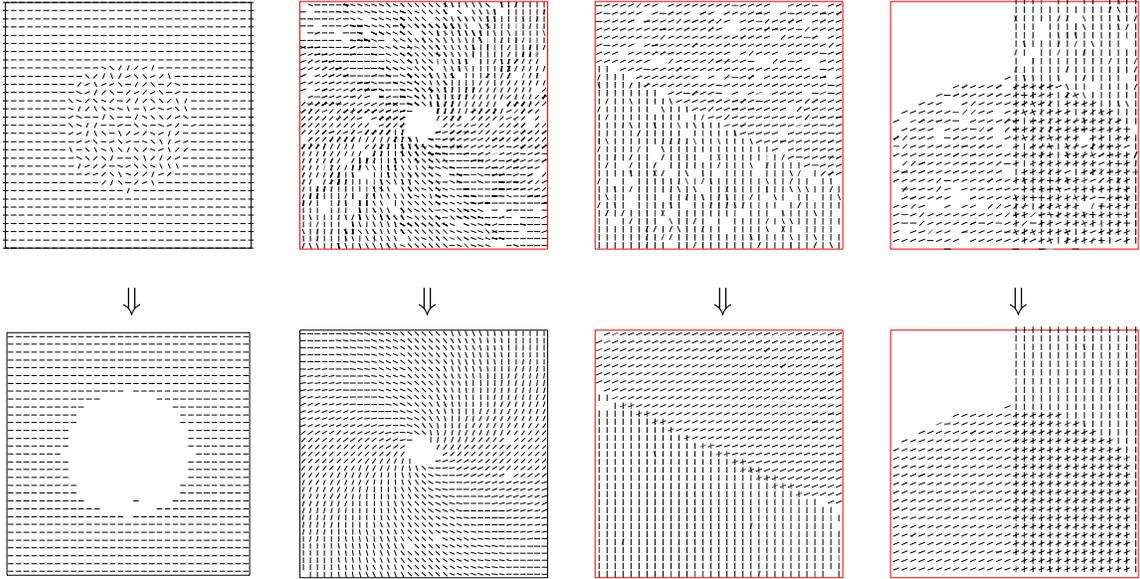


Figure 14: The organization of coherent, synthetic flow structures based on right helicoidal compatibilities.

Fig. 14 illustrates the result of the relaxation process (20 iterations, $\delta = 0.25$) on other synthetic inputs. These particular data sets were selected to demonstrate the different properties of our system, namely (1) the ability to handle missing measurements and “holes”, (2) the preservation of discontinuities and boundaries, (3) the rejection of non-flow structures, and (4) the capacity to process overlapping or multi-oriented flows. Most of these counter the inherent limitations of orientation diffusion, as we discussed in Section 1.5.

Figs. 15 and 16 show the results of applying our system to natural texture flows. Initial orientation measurements for these flows were computed directly from the intensity gradient while initial curvatures were approximated by the application of Eq. 3. Although this crude method provides very noisy measurements, the relaxation process succeeded to extract the coherent structure in the image⁴.

As we mentioned in Section 1.2, texture flow-like patterns are not limited to textured surfaces and may find applications in other aspects of computer vision too, most notably shading analysis. One example where a failure to preserve the geometrical information in the shading may result in a wrong interpretation of the visual data was illustrated in Fig. 6. Another example is shown in Fig. 12A, which may reflect the shading flow field of a bullet-like object, composed of a cylindrical (parabolic) part and a spherical (elliptical) part. The preservation of the (2^{nd} order) flow’s discontinuity in this data now becomes critically important since it indicates a second-order discontinuity in the observed shape, which ultimately may be used for part decomposition and object recognition.

⁴These and many other examples are available online at <http://www.cs.yale.edu/homes/shahar/demos.html>.

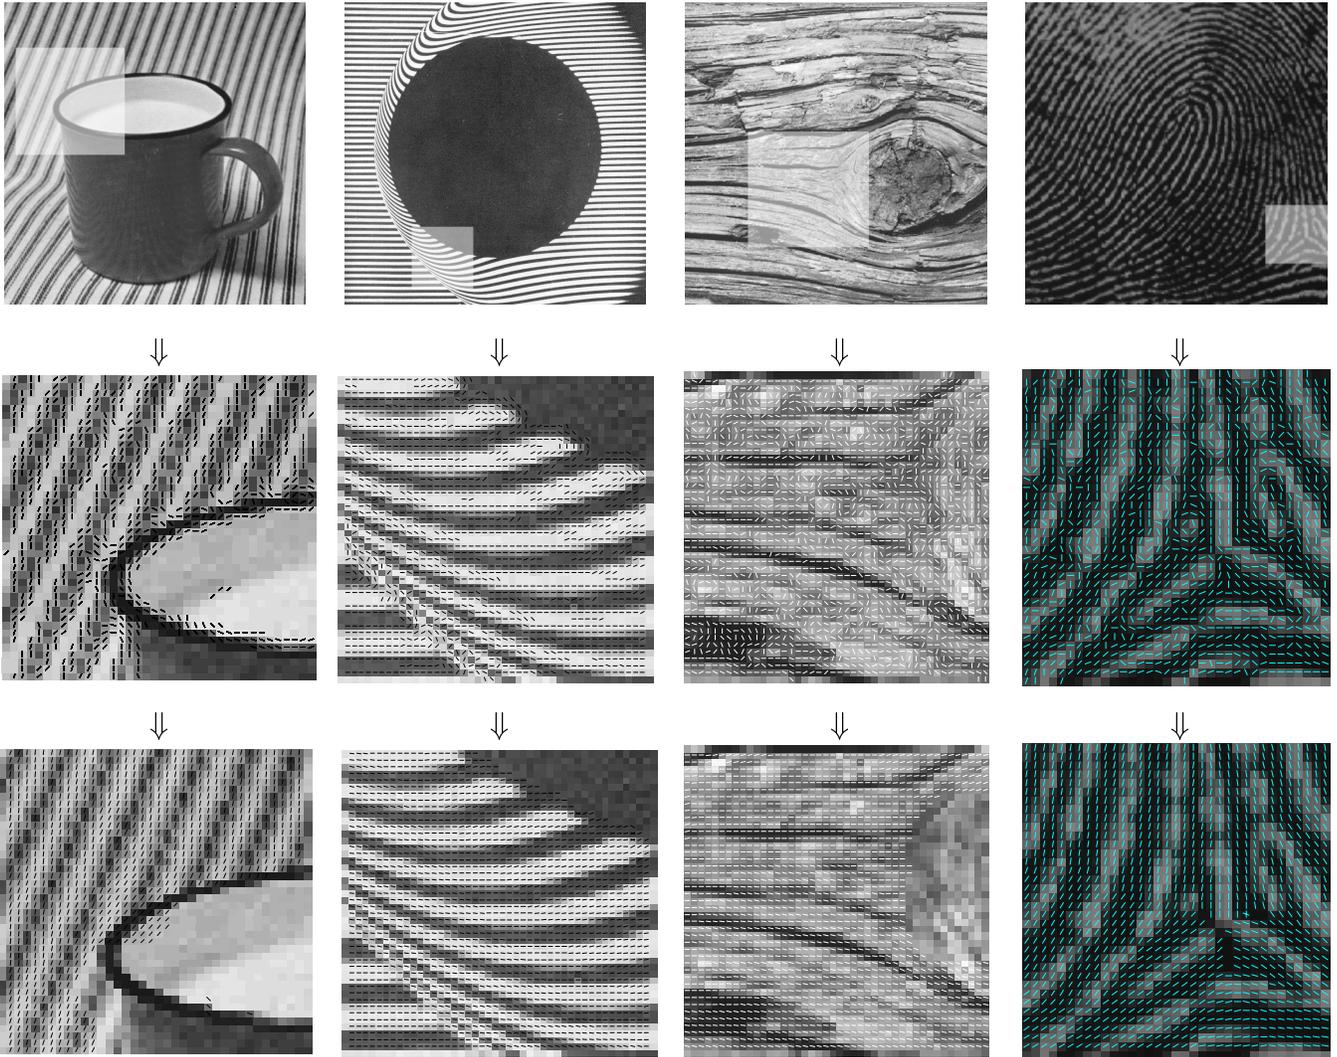
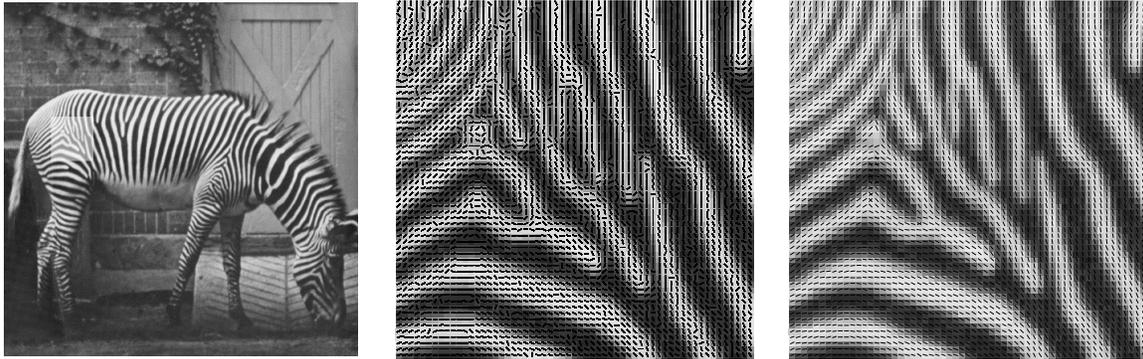


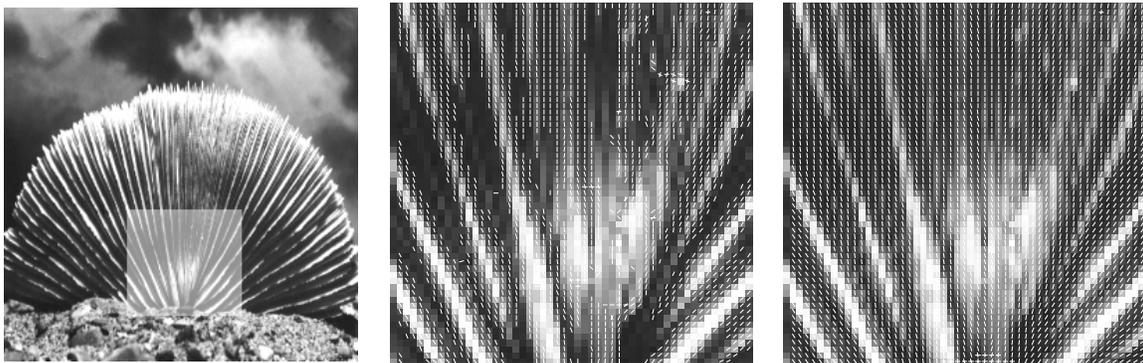
Figure 15: Natural texture flow organization based on right helicoidal compatibilities. Top row shows the images and a region of interest to which we applied the computation. Middle row shows the initial measurements superimposed on the intensity ROI. Bottom row shows the resultant flows. Compare, in particular, the recovery of the flow in the tree stump image to the desired perceptual structure in Fig. 5C and the rejection of measurements in the region of the knot.

We applied our relaxation labeling process to shading flow field of synthetic and natural images. (Fig. 17). Among other data, we analyzed the shading flow field of the object in Fig. 6 to get a globally coherent pattern with its vertical discontinuity well preserved. Also demonstrated in this figure is how non-coherent shading flow fields, typically originating from highly textured or non smooth surfaces, are rejected altogether. As we mentioned, this can be particularly useful to prevent confusion in edge classification [17] and shape interpretation tasks.

Motivated by the columnar architecture of the primary visual cortex, the computational system developed in this paper is distributed in nature and best suited to run on massively parallel hardware. The experimental results presented here were obtained on a serial computer, though, which affected the 5D



(A)



(B)

Figure 16: Natural texture flow organization. (A) A zebra image similar to that in Fig. 3 and a ROI around its scapular stripes. The initial gradient-based measurements (center) are dominated by the stripe terminations and the singularity region is corrupted by an intensity artifact. The relaxed flow (right) correctly restores the perceptual structure and signals the short discontinuity below the singularity. Contrast the almost everywhere smooth change in the dominant orientation to the segmentation in Fig. 3. (B) The Shell image from Fig. 3 with initial measurements based on the Logical/Linear edge detector [19]. The noisy and sparse data set is relaxed to a coherent flow all over the shell. Compare to the segmentation in Fig. 3.

relaxation run-times significantly. Experiments ran on a 860MHz Pentium III desktop PC equipped with 128 MBytes of memory and a Linux OS. While the typical (serial) computation time consumed by each node was less than 1 millisecond per iteration, the total iteration time on a typical image size of 128x128 pixels was approximately one hour. The much smaller synthetic flows were 35x35 pixels in size and their relaxation took approximately 5 minutes per iteration.

5 Summary

In this work we argued for the importance of texture flow from a perceptual organization point of view and how its organization into piecewise coherent parts can be used for segmentation, shape from shading, motion analysis, edge classification, and parts decomposition.

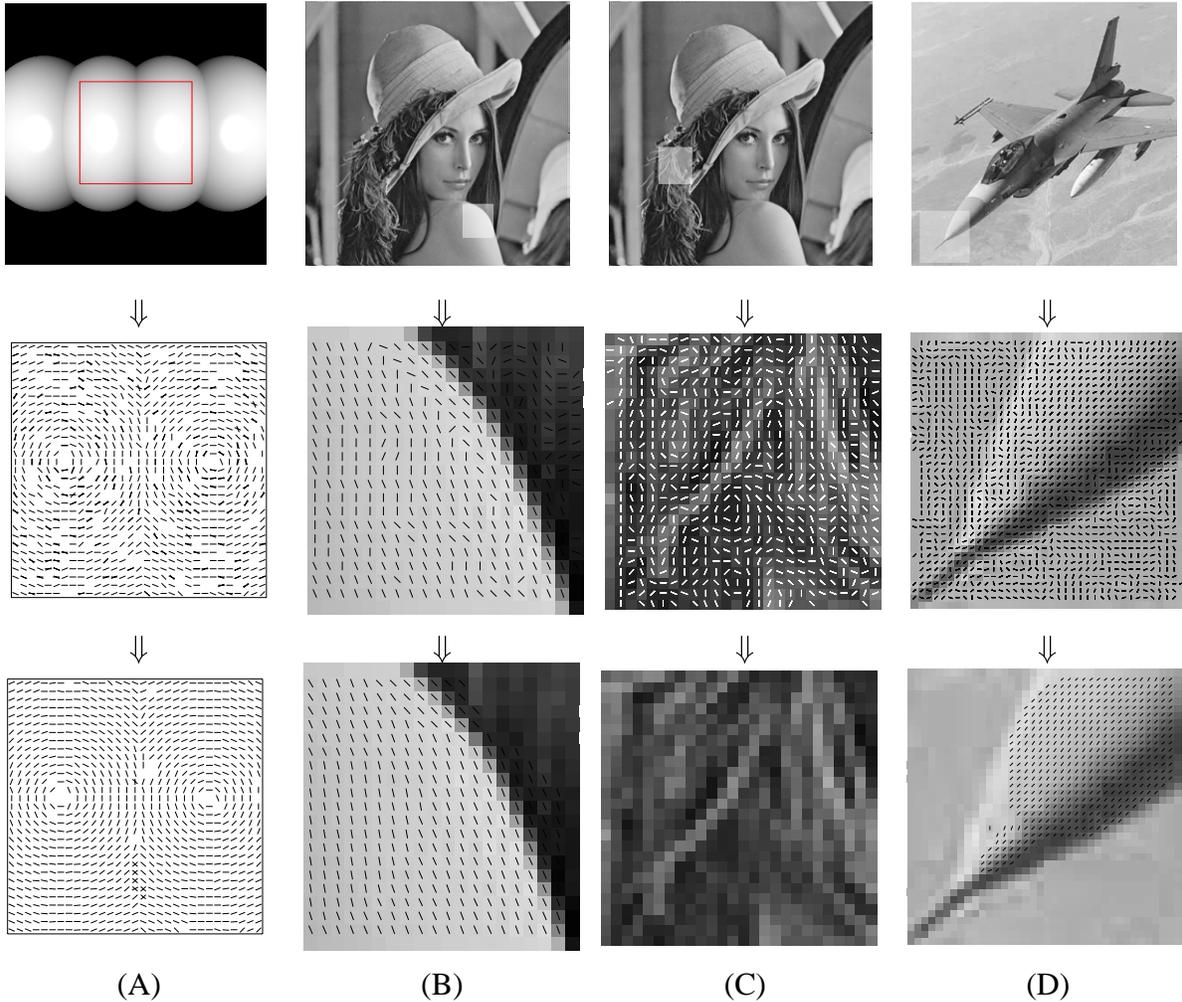


Figure 17: The organization of coherent shading flow field with right helicoidal compatibilities. (A) The worm-like object from Fig. 6. Unlike orientation diffusion (Fig. 6), the relaxation process preserves the vertical discontinuity. (B) Lena’s shoulder ROI from Fig. 4. Note how its shading flow field is cleaned and preserved, as opposed to the feathers ROI in (C) which was completely rejected. (D) An image of an F16. The computation extracted the coherent shading flow on the plane’s nose while suppressing the non coherent structure at the background.

We analyzed the geometry of texture flow and showed that in order to account for its varying nature, its representation must incorporate two curvatures. By applying minimization criteria and investigating the behavior of the curvatures we then derived a unique local model for the flow’s local behavior. Taking the form of a right helicoid in $XY\theta$, this model is the only one that has identical covariation of the two curvatures, and thus it is the only model that puts these variations on equal ground.

Using the right helicoid as a model for texture flow good continuation, we then proposed a relaxation labeling network for the organization of piecewise coherent texture flow. We embed the network in a 5D space in which every node represents a texture flow patch of a particular position, dominant orientation, and two curvatures. We then derived its compatibility function directly from the right helicoidal model and normalized it appropriately to provide a practical degree of boundary stability, under the assumption

that the locus of boundaries is locally a straight line in the image plane.

We demonstrated the behavior of our relaxation network on numerous examples, both synthetic and natural, and showed its robustness to noise. As opposed to existing methods to the analysis of oriented patterns, our approach handles non dense data, it is able to reject non flow structure, and to represent and process overlapping flows.

Finally, it should be mentioned that the advantages of an explicit modeling of the dynamic nature of texture flows, as proposed in this paper, may extend beyond pure computational efforts. In particular, it may open a path for a new psychophysical investigation of related perceptual issues such as orientation-based segmentation, pop out, and sensitivity to changes in orientation. We consider all these as part of our future research.

Appendix - Proofs

Proposition 2: Given any texture flow $\{E_T, E_N\}$, its curvature functions κ_T and κ_N must satisfy the relationship

$$\nabla \kappa_T \cdot E_N - \nabla \kappa_N \cdot E_T = \kappa_T^2 + \kappa_N^2 .$$

Proof: Rewrite the system in Eq. 3 as follows

$$\begin{aligned} \theta_x &= \kappa_T \cos \theta - \kappa_N \sin \theta \\ \theta_y &= \kappa_T \sin \theta + \kappa_N \cos \theta \end{aligned} \tag{14}$$

and observe that

$$\|\nabla \theta\|^2 = \kappa_T^2 + \kappa_N^2 . \tag{15}$$

Differentiate θ_x and θ_y in Eq. 14 by y and x , respectively, and substitute Eq. 14 to obtain:

$$\begin{aligned} \theta_{xy} &= \kappa_{T_y} \cos \theta - \kappa_{N_y} \sin \theta - \theta_y^2 \\ \theta_{yx} &= \kappa_{T_x} \sin \theta + \kappa_{N_x} \cos \theta + \theta_x^2 \end{aligned}$$

Finally, impose integrability, i.e. $\nabla \times \nabla \theta = \theta_{yx} - \theta_{xy} = 0$, to get

$$\nabla \kappa_T \cdot (-\sin \theta, \cos \theta) - \theta_y^2 = \nabla \kappa_N \cdot (\cos \theta, \sin \theta) + \theta_x^2$$

and substitute both Eq. 15 and the expressions for E_T and E_N .

Q.E.D.

Proposition 3: Assume (w.l.o.g.) that $q = (0, 0)$ and $\theta(0, 0) = 0$. A function $\theta(x, y)$ which satisfies $(\kappa_T(q), \kappa_N(q)) = (K_T, K_N)$ and is a critical point of both functionals (5) and (6) is either

$$\begin{aligned}
\text{a plane} & \quad \theta(x, y) = K_T x + K_N y, \\
\text{a right Helicoid} & \quad \theta(x, y) = \tan^{-1}\left(\frac{K_T x + K_N y}{1 + K_N x - K_T y}\right), \\
\text{or a left Helicoid} & \quad \theta(x, y) = \tan^{-1}\left(\frac{K_T x + K_N y}{1 - K_N x + K_T y}\right).
\end{aligned}$$

Proof: The critical points of functional (5), functions whose Laplacian vanishes, are known as *harmonic functions* [1]. The critical points of functional (6), surfaces whose mean curvature vanishes, are known as *minimal surfaces* [35, 8]. It was in the early part of this century that Hamel [13] (see also in [35]) proved that besides the plane, the only real harmonic minimal surfaces in Euclidean space of three dimensions are the helicoids, which can be represented in a suitable coordinate system by $\theta(x, y) = \tan^{-1}\left(\frac{y}{x}\right) + C$. This key result in the theory of minimal surfaces was later extended by Graustein [12] to include imaginary solutions as well.

With the application of Hamel's theorem it remains to include the fact that the surface must be a graph of a 2D function over the image plane up to 2π -periodicity in the θ direction. In addition, the surface solution should induce a given pair of curvatures $(\kappa_T(q), \kappa_N(q)) = (K_T, K_N)$ at q . Without loss of generality, set q as the origin, i.e., $q = (0, 0)$, and let

$$\theta(0, 0) = 0 \quad . \quad (16)$$

This aligns the frame $\{E_T, E_N\}$ at the origin in the direction of the global coordinate system and results, according to Eq. 3, in

$$\nabla\theta(0, 0) = (K_T, K_N) \quad . \quad (17)$$

Consider a helicoid in $XY\theta$ (the planar case is trivial). It is clear that in order for it to be a graph of a 2D function (up to 2π -periodicity) its *directrix* must be perpendicular to the XY plane and its pitch must be exactly $\pm 2\pi$. All these helicoids are equal up to translation $(\bar{x}, \bar{y}, \bar{\theta})$ in $XY\theta$ and thus can be expressed in Cartesian coordinates by

$$s(x, y) = (x, y, \theta(x, y)) = \left(x, y, \tan^{-1}\left(\pm \frac{y - \bar{y}}{x - \bar{x}}\right) + \bar{\theta}\right). \quad (18)$$

The given data (16) and (17) is now used to determine the vector $(\bar{x}, \bar{y}, \bar{\theta})$. From (18) and (17) we get

$$\begin{aligned}
\theta_x &= \pm \frac{\bar{y} - y}{(\bar{y} - y)^2 + (\bar{x} - x)^2} \quad \Rightarrow \quad \theta_x(0, 0) = \pm \frac{\bar{y}}{\bar{y}^2 + \bar{x}^2} = K_T \\
\theta_y &= \mp \frac{\bar{x} - x}{(\bar{y} - y)^2 + (\bar{x} - x)^2} \quad \Rightarrow \quad \theta_y(0, 0) = \mp \frac{\bar{x}}{\bar{y}^2 + \bar{x}^2} = K_N
\end{aligned}$$

which together with (16) becomes

$$(\bar{x}, \bar{y}, \bar{\theta}) = \left(\mp \frac{K_N}{K_T^2 + K_N^2}, \pm \frac{K_T}{K_T^2 + K_N^2}, \pm \tan^{-1}\left(\frac{K_T}{K_N}\right)\right).$$

Substituting these two translation vectors into (18) yields the following two solutions

$$\theta(x, y) = \pm \tan^{-1} \left(\frac{K_T}{K_N} \right) \pm \tan^{-1} \left(\frac{\mp K_T + (K_T^2 + K_N^2)y}{\pm K_N + (K_T^2 + K_N^2)x} \right).$$

Applying the following trigonometric identity

$$\tan^{-1}(A) + \tan^{-1}(B) = \tan^{-1} \left(\frac{A+B}{1-A \cdot B} \right)$$

results in the two helicoidal functions in the Proposition and completes the proof.

It is worth mentioning that a more familiar parametrization of the helicoidal functions in the proposition can be achieved by a proper substitution. In particular, substituting

$$x = v \cos(u) - \frac{K_N}{\xi^2} \quad y = v \sin(u) + \frac{K_T}{\xi^2} \quad \xi = \sqrt{K_T^2 + K_N^2}$$

into the surface representation of the *right* helicoid

$$s(x, y) = \left(x, y, \tan^{-1} \left(\frac{K_T x + K_N y}{1 + K_N x - K_T y} \right) \right)$$

yields the following parametrization

$$s(u, v) = \left(v \cos(u) - \frac{K_N}{\xi^2}, v \sin(u) + \frac{K_T}{\xi^2}, \tan^{-1} \left(\frac{K_T \cos(u) + K_N \sin(u)}{-K_T \sin(u) + K_N \cos(u)} \right) \right)$$

which then can be simplified to the familiar helicoidal form

$$s(u, v) = \left(v \cos(u) - \frac{K_N}{\xi^2}, v \sin(u) + \frac{K_T}{\xi^2}, u + \text{sign}(K_T) \cos^{-1} \left(\frac{K_N}{\xi} \right) \right). \quad Q.E.D.$$

Proposition 4: The harmonic minimal solutions are the only functions (that satisfy the initial data and) whose p-Laplacian $\Delta_p(\theta)$ vanishes simultaneously for all $p \geq 0$.

Proof: We need to show that the three harmonic solutions are the only ones to satisfy

$$\Delta_p \theta(x, y) \triangleq \nabla \cdot (|\nabla \theta|^{p-2} \nabla \theta) = 0. \quad (19)$$

First note that

$$p = 2 \Rightarrow \Delta_p \theta(x, y) = \Delta \theta(x, y) = \theta_{xx} + \theta_{yy}$$

$$p = 1 \Rightarrow \Delta_p \theta(x, y) = \nabla \cdot (|\nabla \theta|^{-1} \nabla \theta) = (\theta_x^2 + \theta_y^2)^{-\frac{3}{2}} (\theta_{xx} \theta_y^2 - 2\theta_x \theta_y \theta_{xy} + \theta_{yy} \theta_x^2).$$

Assuming non vanishing gradient, both operators are well defined, thus a function $\theta(x, y)$ that satisfies both $\Delta_1 \theta(x, y) = 0$ and $\Delta_2 \theta(x, y) = 0$ must satisfy the following two constraints simultaneously

$$\begin{aligned} \theta_{xx} + \theta_{yy} &= 0 \\ \theta_{xx} \theta_y^2 - 2\theta_x \theta_y \theta_{xy} + \theta_{yy} \theta_x^2 &= 0. \end{aligned}$$

Summing up both sides we get

$$\theta_{xx}(\theta_y^2 + 1) - 2\theta_x\theta_y\theta_{xy} + \theta_{yy}(\theta_x^2 + 1) \stackrel{\Delta}{=} H = 0$$

which requires the mean curvature of the surface $(x, y, \theta(x, y))$ in $XY\theta$ to vanish everywhere. In other words, a function whose p -Laplacian vanishes for all p must be a minimal surface. Naturally, it is time to apply Hamel's theorem [13] once again to conclude that the only candidates for $\Delta_p = 0 \quad \forall p$ can be either a plane or a helicoid, namely, one of our harmonic minimal solutions (Proposition 3). It remains to check which of these functions satisfies this constraint. Substituting each of them into Eq. 4 verifies that all of them satisfy it and thus completes the proof. *Q.E.D.*

References

- [1] S. Axler, P. Bourdon, and W. Ramey. *Harmonic Function Theory*. Springer-Verlag, 1992.
- [2] O. Ben-Shahar and S.W. Zucker. Beyond orientation diffusion: A curvature-based approach to the analysis of oriented patterns. (In preparation).
- [3] M.J. Black, G. Sapiro, D.H. Marimont, and D. Heeger. Robust anisotropic diffusion. *IEEE Trans. Image Processing*, 7(3):421–432, 1998.
- [4] P. Blomgren and T.F. Chan. Color TV: Total variation methods for restoration of vector-valued images. *IEEE Trans. Image Processing*, 7(3):304–309, 1998.
- [5] P. Breton and S.W. Zucker. Shadows and shading flow fields. In *Proc. Computer Vision and Pattern Recognition*, 1996.
- [6] P. Brodatz. *Textures: A Photographic Album for Artists and Designers*. Dover Publications, 1966.
- [7] T. Chan and J. Shen. Variational restoration of non-flat image features: Models and algorithms. CAM-TR 99-20, UCLA, 1999.
- [8] M Do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Inc., 1976.
- [9] A.R. Forsyth. *Lectures on the differential geometry of curves and surfaces*. Cambridge, University press, 1912.
- [10] L. Glass. Moiré effect from random dots. *Nature*, 223(5206):578–580, 1969.
- [11] G.H. Granlund and H. Knutsson. *Signal Processing for Computer Vision*. Kluwer Academic Publishers, 1995.
- [12] W.C. Graustein. Harmonic minimal surfaces. In *Trans. of the American Mathematical Society*, volume 47, pages 173–206, 1940.
- [13] G. Hamel. Zum gedachtnis an hermann amandus schwarz. *Jahresber, D.M.V.*, 32:6–13, 1923.

- [14] J.J. Hopfield and D.W. Tank. Neural computation of decisions in optimization problems. *Biological Cybernetics*, 52:141–152, 1985.
- [15] C.P. Hornung. *Background Patterns, Textures, and Tints*. Dover Publications, 1976.
- [16] D. Hubel and T. Wiesel. Functional architecture of macaque monkey visual cortex. In *Proc. R. Soc. London Ser. B*, volume 198, pages 1–59, 1977.
- [17] P.S. Huggins and S.W. Zucker. How folds cut a scene. In *Proc. of the 4th Int. Workshop on Visual Form*, 2001.
- [18] R.A. Hummel and S.W. Zucker. On the foundations of the relaxation labeling processes. *IEEE Trans. Pattern Anal. Machine Intell.*, 5:267–287, 1983.
- [19] L.A. Iverson and S.W. Zucker. Logical/linear operators for image curves. *IEEE Trans. Pattern Anal. Machine Intell.*, 17(10):982–996, 1995.
- [20] G. Kanizsa. *Organization in Vision: Essays on Gestalt Perception*. Praeger Publishers, 1979.
- [21] M. Kass and A. Witkin. Analyzing oriented patterns. *CVGIP*, 37:362–385, 1987.
- [22] D.R.T. Keeble, F.A.A. Kingdom, and M.J. Morgan. The orientational resolution of human texture perception. *Vision Res.*, 37(21):2993–3007, 1997.
- [23] B.K. Kimia, L. Frankel, and A.M. Popescu. Euler spiral for shape completion. In *IEEE Computer Society Workshop on Perceptual Organization in Computer Vision*, 1999.
- [24] R. Kimmel and N. Sochen. Orientation diffusion or how to comb a porcupine? CIS 2000-02, Technion - Israel Institute of Technology, 2000.
- [25] J. Kittler and J. Illingworth. Relaxation labeling algorithms - a review. *Image and Vision Computing*, pages 206–216, 1985.
- [26] J.J. Koenderink. The structure of images. *Biological Cybernetics*, 50:363–370, 1984.
- [27] J.J. Koenderink and A.J. van Doorn. Local structure of movement parallax of the plane. *J. Opt. Soc. Am.*, 66(7):717–723, 1976.
- [28] M.S. Landy and J.R. Bergen. Texture segregation and orientation gradient. *Vision Res.*, 31(4):679–691, 1991.
- [29] T.S. Lee. A bayesian framework for understanding texture segmentation in the primary visual cortex. *Vision Res.*, 35(18):2643–2657, 1993.
- [30] D.G. Lowe. *Perceptual Organization and Visual Recognition*. Kluwer Academic Publishers, 1985.
- [31] D. Marr. *Vision*. W.H. Freeman and Company, 1982.
- [32] D.A. Miller and S.W. Zucker. Computing with self-excitatory cliques: A model and application to hyperacuity-sclae computation in visual cortex. *Neural Comput.*, 11:21–66, 1999.

- [33] J.L. Mohammed, R.A. Hummel, and S.W. Zucker. A gradient projection algorithm for relaxation methods. *5(3):330–332*, 1983.
- [34] J.D. Murray. *Mathematical Biology*. Springer-Verlag, 1989.
- [35] J.C.C Nitsche. *Lectures on Minimal Surfaces*. Cambridge Univ. Press, 1989.
- [36] H.C. Nothdurft. Sensitivity for structure gradient in texture discrimination tasks. *Vision Res.*, *25(12):1957–1968*, 1985.
- [37] H.C. Nothdurft. Texture segmentation and pop-out from orientation contrast. *Vision Res.*, *31(6):1073–1078*, 1991.
- [38] B. O’Neill. *Elementary Differential Geometry*. Academic Press, 1966.
- [39] P. Parent and S.W. Zucker. Trace inference, curvature consistency, and curve detection. *IEEE Trans. Pattern Anal. Machine Intell.*, *11(8):823–839*, 1989.
- [40] P. Perona. Orientation diffusion. *IEEE Trans. Image Processing*, *7(3):457–467*, 1998.
- [41] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Machine Intell.*, *12(7):629–639*, 1990.
- [42] A.A. Rao and R.C. Jain. Computerized flow field analysis: Oriented texture fields. *IEEE Trans. Pattern Anal. Machine Intell.*, *17(7):693–709*, 1992.
- [43] A.R. Rao and B.G. Schunck. Computing oriented texture fields. *CVGIP: Graphical Models and Image Processing*, *53(2):157–185*, 1991.
- [44] D. Regan, L.V. Hajdur, and X.H. Hong. Two-dimensional aspect ratio idscrimination for shape defined by orientation texture. *Vision Res.*, *36(22):3695–3702*, 1996.
- [45] G. Sapiro and D.L. Ringach. Anisotropic diffusion of multivalued images with applications to color filtering. *IEEE Trans. Image Processing*, *5(11):1582–1586*, 1996.
- [46] J. Shi and J. Malik. Normalized cuts and image segmentation. In *Proc. Computer Vision and Pattern Recognition*, pages 731–737, 1997.
- [47] M. Sigman, G.A. Cecchi, C.D. Gilbert, and M.O. Magnasco. On a common circle: Natural scenes and gestalt rules. *Proc. Natl. Acad. Sci. U.S.A.*, *98(4):1935–1940*, February 2001.
- [48] N. Sochen, R. Kimmel, and A.M. Bruckstein. Diffusion and confusions in signal and image processing. *J. of Mathematical Imaging and Vision*, To appear 2001.
- [49] N. Sochen, R. Kimmel, and R. Malladi. A geometrical framework for low level vision. *IEEE Trans. Image Processing*, *7(3):310–318*, 1998.
- [50] K.A. Stevens. Computation of locally parallel structure. *Biol. Cybern.*, *29:19–28*, 1978.
- [51] K.A. Stevens. The line of curvature constraint and the interpretation of 3d shape from parallel surface contours. In *Proc. of the Int. Joint Conf. on Artif. Intell.*, pages 1057–1061, 1983.

- [52] B. Tang, G. Sapiro, and V. Caselles. Diffusion of general data on non-flat manifolds via harmonic maps theory: The direction diffusion case. *Int. J. Comput. Vision*, 36(2):149–161, 2000.
- [53] J.T. Todd and F.D. Reichel. Visual perception of smoothly curved surfaces from double-projected contour patterns. *J. Exp. Psych.: Human Perception and Performance*, 16(3):665–674, 1990.
- [54] P.W. Verbeek, L.J. van Vliet, and J. van de Weijer. Improved curvature and anisotropy estimation for curved line bundles. In *ICPR*, pages 528–533, 1998.
- [55] D.L. Waltz. Understanding line drawings of scenes with shadows. In P.H. Winston, editor, *The Psychology of Computer Vision*, pages 19–91. McGraw-Hill, New York, 1975.
- [56] T. Watanabe and P. Cavanagh. Texture laciness: The texture equivalent of transparency? *Perception*, 25:293–303, 1996.
- [57] J. Weickert. A review of nonlinear diffusion filtering. In B. ter Haar Romeny, L. Florack, J.J. Koenderink, and M. Viergever, editors, *Scale-Space Theory in Computer Vision*, volume 1252 of *Lecture Notes in Computer Science*, pages 3–28. Springer-Verlag, 1997.
- [58] M. Wertheimer. Gestalt theory. In W.D. Ellis, editor, *A source book of Gestalt Psych.*, pages 1–11. Routledge & Kegan Paul Ltd., 1955.
- [59] M. Wertheimer. Laws of organization in perceptual forms. In W.D. Ellis, editor, *A source book of Gestalt Psych.*, pages 71–88. Routledge & Kegan Paul Ltd., 1955.
- [60] L.R. Williams and D.W. Jacobs. Stochastic completion fields: A neural model of illusory contour shape and salience. *Neural Comput.*, 9(4):837–858, 1997.
- [61] A.P. Witkin. Scale-space filtering. In *Proc. of the Int. Joint Conf. on Artif. Intell.*, pages 1019–1022, 1983.
- [62] A.P. Witkin and J.M. Tenenbaum. On the role of structure in vision. In J. Beck, B. Hope, and A. Rosenfeld, editors, *Human and Machine Vision*, pages 481–542. Academic Press, 1983.
- [63] S.W. Zucker. Computational and psychophysical experiments in grouping: Early orientation selection. In J. Beck, B. Hope, and A. Rosenfeld, editors, *Human and Machine Vision*, pages 545–567. Academic Press, 1983.