To appear in the Journal of Physiology, 2003.

Cortical Connections and Early Visual Function: Intra- and Inter-Columnar Processing

Ohad Ben-Shahar Patrick S. Huggins Tomas Izo*
Steven W. Zucker†
Dept. of Computer Science and Interdisciplinary Neuroscience Progam
Yale University, New Haven, CT U.S.A.
{ben-shahar,huggins,izo,zucker@cs.yale.edu}

Abstract

While it is widely assumed that the long-range horizontal connections in V1 are present to support contour integration, there has been only limited consideration of other possible relationships between anatomy and physiology (the horizontal connections) and visual function beyond contour integration. We introduce the possibility of other relationships directly from the perspective of computation and differential geometry by identifying orientation columns in visual physiology with the (unit) tangent bundle in differential geometry. This suggests abstracting early vision in a space that incorporates both position and orientation, from which we show that the physiology is capable of supporting a number of functional computations beyond contour integration, including texture-flow and shading-flow integration, as well as certain relationships between them. The geometric abstraction emphasizes the role of curvature, which necessitates a coupled investigation into how it might be estimated. The result is an elaboration of layer-to-layer interactions within an orientation column, with nonlinearities possibly implemented by shunting inhibition. Finally, we show how the same computational framework naturally lends itself to solving stereo correspondence, with binocular tangents abstracting curves in space.

*Current address: Artificial Intelligence Lab., MIT, Cambridge MA
†Corresponding author. Research supported by AFOSR and DARPA. We thank S. Alibhai for the stereo computations.

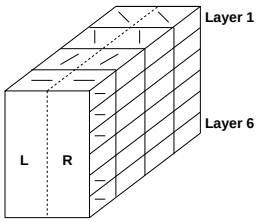

Figure 1: The standard Hubel-Wiesel "ice cube cartoon" shows how cortex is organized so that each local retinotopic area is covered by receptive fields spaning orientations(tangential penetration), size (normal penetration), and eye-of-origin.

1 Introduction

Orientation provides the basis for organizing much of visual cortex and provides the foundation for visual information processing. An examination of receptive field structure shows how the neurons in V1 form a position+orientation map of the visual field [HW77]. Roughly speaking, recordings along a short tangential (interblob) penetration reveal a collection of cells with about the same receptive field centers but with shifts in orientation preference, while normal penetrations reveal cells with similar orientation and position preferences but different receptive field sizes (or spatial frequency selectivity). Taken together these observations define an array of orientation columns which, combined with eye of origin, provide a basic representation for visual information processing; see Fig. 1.

The resulting hypercolumns are not independent. There exist long-range horizontal interactions between them, which greatly enlarges the domain of possibilities for information processing. The basic question is: which early visual information processing tasks take place within columns, and which between?

One such task is edge detection, and it is widely believed that the long-range horizontal connections support the sensory integration necessary for this. The nature of the evidence supporting this belief is reviewed in the next section. However, less analysis has been applied to the question of determining which tasks comprise sensory integration, and it is on this latter question that we concentrate in this paper. Assembling estimates of local orientation (within columns) into long curves (between columns) is clearly one such task, but this is not unique; there are several other tasks that also

arise naturally in early vision, all of which, we shall show, are consistent with the available data on long-range horizontal integration.

This paper is primarily computational. We first develop a computational theory of curve detection to illustrate that there is more to this task than is normally presupposed. We also use this opportunity to introduce the conceptual thread that runs through the paper: the use of differential-geometric ideas to articulate theories of sensory integration. We then sketch how these ideas can be extended for the analysis of texture, shading, and binocular information. Particular types of sensory integration arise for each of these tasks, although the general form remains invariant.

The computational ideas have a natural expression in physiological terms, and we will use the intra-columnar processing to develop local representations and inter-columnar processing to integrate them into coherent wholes. Some of the functional roles in which the long-range horizontal connections could be involved are quite unexpected, in that they do not appear related to curve integration without the analysis provided.

More detailed issues arise throughout the discussion, such as the role of certain nonlinearities in information processing. We also relate these to biology, at both a biophysical level (implementing the nonlinearities with shunting inhibition) and at a detailed anatomical level (elaborating intra-columnar processing across layers).

While evidence is provided to demonstrate that the above claims are well-founded computationally, in our current state of understanding any attempt to reduce such abstract models to physiological circuits must necessarily involve a degree of speculation. This said, we feel the time is right to start contributing such ideas to the neurophysiological community. Our hope is that they will broaden both the discussion around, and the experimental perspective on, sensory integration.

1.1 Background: Co-linear Facilitation and Curve Integration

Finding the boundaries of objects is one of the central problems of early visual information processing (Fig. 2), and cells with oriented receptive fields are often interpreted as local edge detectors, or at least as components of a system for local edge detection. That a system is required follows from the fact that many of the local responses are noisy or ambiguous: they may

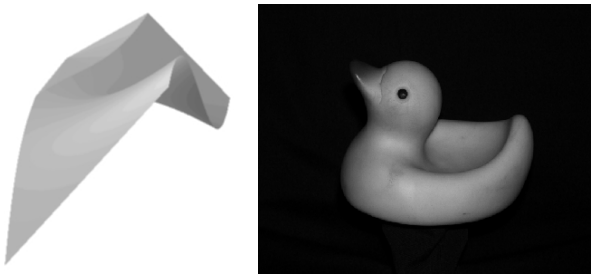

Figure 2: Two sample images suggest the range of problems for early vision, even when the task is limited to static, monocular, monochrome imagery. Both the sheet and the duck illustrate the importance of boundary and shading information for supporting 3-dimensional interpretations, but closer examination reveals how subtle the questions can become. Where, for example, should the boundaries be drawn on the sheet? Or how can edge and shading information combine to indicate that the duck's neck occludes its back rather than the other way around?

arise from an accidental alignment of viewing geometry and lighting, from a specularity, from noise in the sensory process, or from a myriad of other causes. The resolution of these noisy, ambiguous responses is sought from context, with those cells responding consistently along an edge facilitating one another to enhance the correct responses while eliminating the noisy, random ones.

This is perhaps the most basic type of contour integration, and it is classical. The Gestalt psychologists argued for a form of *orientation good continuation* nearly 75 years ago [Wer55], and it was suggested nearly 20 years ago that the anatomical substrate for such facilitation could be the long-range horizontal connections [MC82].

The long-range horizontal axons effectively connect cells in the superficial layers of nearby orientation columns, so to test the above hypothesis the orientation preferences of the cells involved must be known. Two experimental paradigms have been developed. First, to test for contextual effects, one can isolate a target unit with a particular orientation preference, and then plot how its activity is modulated by stimuli in the surround. This paradigm was first used to demonstrate influences from beyond the classical receptive field [MF76, AME85], and is still being applied in technically advanced ways [KIGW95b, KWG99]. Many such studies show that the target

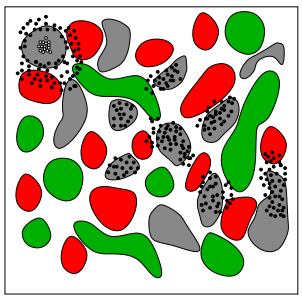

Figure 3: A cartoon of the relationship between orientation columns and long range horizontal connections as revealed in optical imaging combined with anatomical tracing. Different colours represents the different orientation columns. The white dots represent the region where an anatomical tracer is injected. The black dots represent axonal terminals coloured by this process. As is typically found, short range interactions span all possible orientations, while long range interactions tend to concentrate in cortical regions that share similar orientation specificity with the injection site.

cell's firing is facilitated by surround stimuli that have about the same orientation as the one preferred by the target cell. While this technique can provide detailed characterization of the influences on an individual cell, it is not suitable for gathering population distributions, it is stimulus dependent, and the system-wide functional links it reveals are more difficult to interpret in terms of the physical connections between neurons.

Alternatively, studies that are more suitable for population statistics have used optical imaging and anatomical tracing to reveal the entire connectivity structure of cells (Fig. 3). Here optical imaging techniques are used to "colour" (an image of) the cortex with the approximate orientation preference of the underlying cells. Tracers are then injected into a cell and mark its terminals. The distribution of these terminals can then be plotted against the "colour" (or rough orientation preference) of the domain in which they terminate to yield a complete characterization of cells connectivity structure in the orienation domain [MAHG93, BZBF97].

Physiological evidence of both types is accumulating, and a summary of the evidence is that the majority of such facilitory interactions are iso-orientational; i.e., between cells with similar orientation preferences [KIGW95a, MAHG93, TGW86]. Taken back to the edge detection problem, many researchers observe that this is essentially what the orientation good contin-

uation hypothesis would predict, and it suggests that a form of co-linear facilitation underlies edge detection.

This basic model for edge detection – filtering by operations analagous to linear, oriented receptive fields followed by co-aligned facilitation – can be implemented and tested on natural images. Thus far this is only an outline of an approach, however, because variations in filters, their interactions, and the detection process remain unspecified. Researchers in computer vision have considered these issues, and one of the most widely used edge detectors is that of Canny [Can86]. This system effectively implements the above outline, and consists of an initial filtering stage, with filters very similar in form to simple-cell receptive fields, followed by a facilitory (hysteresis) stage that implements a type of co-aligned facilitation. Allowing an analogy to driving a car along the edge, the output of the detector is essentially the strongest set of initial edge detector responses that continue the edge in the same orientation in which it was (recently) going. Of course, the continuation can be adjusted by several parameters, intuitively varying the "inertia" with which an edge continues. Both the "driving history" and the strength of responses affect the final result.

The Canny detector can be evaluated on an image to assess performance, and researchers unfamiliar with computer vision are often surprised at the results. However, edge detection is not an easy problem as it may initially appear. Different values of the parameters illustrate the variations that typically occur, and it is instructive to examine them (Fig. 4). Among the problems that emerge: boundaries can be broken apart or, what is perhaps worse, proper but physically disconnected boundary segments can be inappropriately connected. Note in particular how different values can bridge between nearby (but totally different) parts of edges by stressing how the inertia parameter jumps across gaps. Unfortunately, there is no agreement on how to select the parameters so that these problems do not arise.

The experience with edge detection in computer vision reveals some of the types of questions that can arise from an information processing perspective. First, for the initial operators, there is the question of how to obtain local estimates of orientation that provide a consistent bridge between the image and the scene. This must not induce incorrect linkages between disparate curves, even when they are close in the image, because this implies incorrect physical structure for objects in the world. Second, there is the question of orientation good continuation: What does this mean in terms of physical object properties and how should nearby oriented responses facilitate one

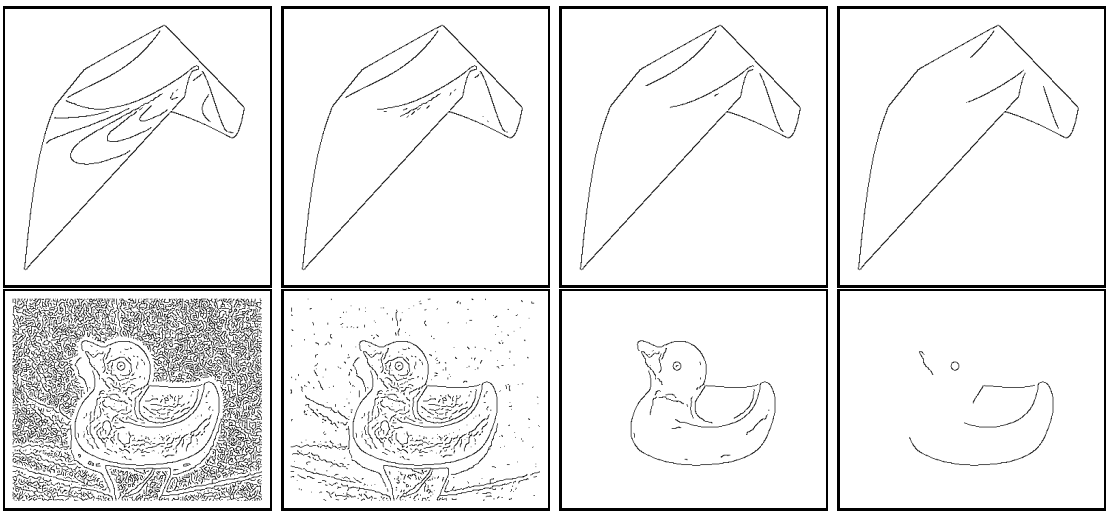

Figure 4: Geometric problems in edge detection using the Canny operator (Matlab implementation). The different edge maps depict the results of the edge detection process using different parameter values. Not only that at some levels non-boundaries are signaled as edges while at others levels genuine edges are missed or broken, but oftentimes proper but distinct boundaries are inappropriately connected (e.g., the boundary of the duck's back and the one of its neck).

another so that they induce precisely the right amount of facilitation, co-linear when appropriate and curved when not?

Thirdly, different types of questions arise from physiology. For example, if the the majority (as opposed to the entirety) of connections are between (approximately) iso-oriented cells, are the "outliers" to this simply noise? Is the correct abstraction for the majority co-linear facilitation, or is something else going on? And, how should the majority be defined; in particular, what is the proper spread in the distribution of connections in the orientation domain?

Finally, there are questions that arise from the interface between physiology and computer vision: are horizontal facilitations only participating in edge detection, or might they also be implementing other functional roles? If so, what might these be and are they consistent with the given data about connections?

We shall consider each of these groups of questions in this paper. To set the stage, we note that a closer examination of the physiological data suggests the story is more complex than co-aligned contour facilitation (Fig. 5). For example, while Kapadia et al. [KIGW95b] stress iso-orientation facilitation, they also provide examples of facilitation between cells with up to 50 deg orientation difference (their Fig. 10). The distributions provided by Bosking et al. [BZBF97] clearly show non negligible portion of anatomical connections even between cross orientations (their Fig. 6). And from the findings of Ts'o et. al. [TGW86] it is evident that there are facilitory functional interactions between iso-oriented cells whose receptive fields are parallel to one another, rather than co-aligned (see top pair in their Fig. 5). Interpreting such pairs as participating in contour facilitation is awkward. Instead, we shall argue that such exceptions are naturally explained from a series of computational tasks, including (curved) contour integration, texture, shading, and stereo processing.

To develop this argument, we shall have to consider how early visual information processing can be structured on orientation hypercolumns. We do this in two stages. First, we briefly review a model that captures enough of the structure of the columnar architecture that it can be related to neurophysiology and neuroanatomy, but is sufficently abstract that it can be analyzed mathematically and computationally. We then proceed to the analysis of curves, textures, shading, and stereo in it, and how these relate back to (certain aspects of) scene structure. As we show, facilitory interactions can be involved in all of them, but co-aligned facilitation by itself is insufficent.

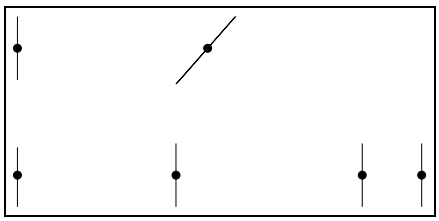

Figure 5: The fact that the majority of connections are between iso-orientation columns does not necessarily imply co-aligned facilitation; similarly, orientation good continuation does not necessarily imply only iso-orientation interactions. Three different pairs of cells are shown, with position corresponding to the relative placement of receptive fields and the short bar corresponding to the preferred orientation of the cell. (LEFT) Co-aligned facilitation is a natural implication of boundary detection and a common interpretation of iso-orientation interactions.. (CENTER) An "outlier" to co-aligned facilitation, with about a 50 degree rotation over a short (retinotopic) distance. This example is not iso-orientation facilitation although it reflects a possible receptive field arrangement along a coherent boundary curve. (RIGHT) An "outlier" to co-aligned facilitation that is iso-orientation facilitation. While this example may contribute to the majority of iso-orientation interactions, it cannot be interpret as serving curve integration. The question is whether this and other "outliers" can support useful visual information processing.

Differential geometry, and curvature in particular, is necessary to understand why co-aligned facilitation is dominant but not unique; there are an important (and predictable) number of non-co-aligned facilitory influences that play key processing roles.

2 The Columnar Machine

We begin by re-drawing the ice-cube model to illustrate the possibility of geometric information processing. We focus on columns with the same monocular specificity and drop deep layers. We depict the orientation hypercolumns as vertical fibres distributed over a retinotopic array (the tilted plane) and we display the orientation preference of cells within each hypercolumn as oriented segments (Fig. 6). When organized in this fashion, a geometric view of processing emerges, in which the fibre of orientations at each position in the

retinotopic array abstracts the orientation hypercolumn, and the arrangement of neighbouring fibres suggests an architecture that is specialized to support interaction between orientations. All orientations in nearby columns are clearly in a position to interact.

Mathematically we conceptualize this interaction by thinking of the different orientations at each position as "labels" and the measure of activity of each neuron as the probability that the neuron is signaling the correct orientation for that retinotopic position. Interactions are modeled as a compatibility function ($r_{ij}(\lambda_1, \lambda')$ in the top right of Fig. 6) between orientation λ' at position j with orientation λ_1 at nearby position i. In game-theoretic terms, one can think of this function as representing the *payoff* that player i gets by choosing strategy (label) λ_1 when player j plays strategy λ'. The goal is to seek an equilibium for this game. Diagramatically we illustrate this with the thickness of the orientated bars, with thicker bars denoting more activity or higher probability (bottom of Fig. 6). Note that this can be viewed as a selection process. The initial values are spread out along each fibre (bottom left, Fig. 6), but the final values concentrate on a single, or a small number, of distinct orientations (bottom right, Fig. 6).

Dynamics are not shown in this diagram, but are key to specifying the manner in which abstract computations map onto networks of real neurons. We have developed one way to do this, which is consistent with biophysics at least to low-order. Imagine neurons as players in the above game, with pure strategies being whether the neuron should depolarize (spike) or hyperpolarize [MZ92, MZ99]. Modeling neurons as piecewise-linear amplifiers, synapses (compatibilities) as conductances, etc., they can be placed in the above form. The key idea behind our model is to consider groups of tightly interconnected excitatory neurons capable of bringing themselves to saturation feedback response following a modest initial afferent bias current. We refer to each such group of neurons as a clique, and for this the connections can involve both short-range and long-range interactions.

The basic computation is in two phases, and builds upon the observed regular spiking behavior of pyramidal cells. Consider a patch of cortex, consisting of hundreds of thousands of cells. We view this as being further organized into many times that number of cliques of densely-interconnected neurons, with order-of-magnitude tens of cells per clique. (Miller and Zucker calculate 33 cells/clique for orientation hyperacuity computations.) Initially the cells in this patch are quiescent. A *computation* amounts to activating all the cells in a single clique to saturation levels, but no others. This begins,

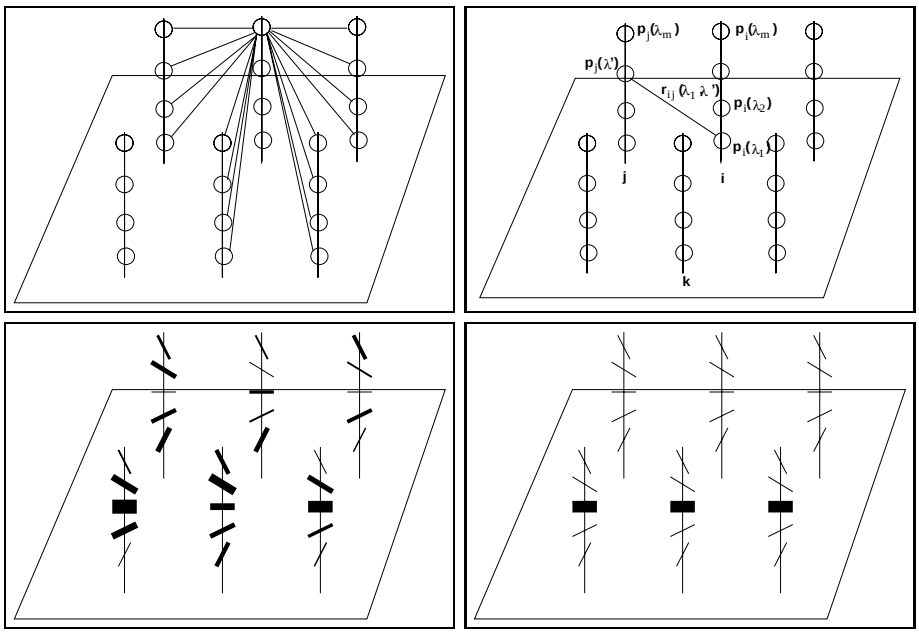

Figure 6: Two re-drawings of the ice-cube model, emphasizing (TOP) the functional anatomy and (BOTTOM) the processing. A retinotopic array (the tilted plane) and a sampling of superficial-layer cells drawn (scale not included) on fibres representing hypercolumns. (TOP LEFT) Inter-columnar connections are not only between iso-orientation cells. (TOP RIGHT) Notation for a game theoretic intepretation of inter-columnar processing; see text for details. (BOTTOM) Here cells are drawn as oriented segments to show their orientation preference. The thickness of the orientated bars denotes activity of cells; initial values in time (BOTTOM LEFT) and then later (BOTTOM RIGHT) denote processing in time, which may be viewed as a selection process along each fibre. Dynamics are not shown.

in Phase I, when afferent stimulation produces single spikes in a majority of the cells in one clique (those cells can be distributed among many different iso-orientation areas), as well as a certain number of other cells outside the clique (noise). Because the clique has a sufficient level of excitatory interconnections, all its cells drive themselves to saturation response levels of about 5 spikes in about 25 msec (end of Phase I), whereas the initially activated cells outside the clique return to their resting membrane potentials and do not spike further (end of Phase II). Thus the clique has been "retrieved", or the orientation specified, at each position, and it is these cliques that define the equilibrium of the game. See [MZ00] for further motivation, development, and analysis.

This model of an abstract machine provides a framework for specifying computations. Given that the labels at each position represent orientation, what remains is to specify the "payoffs" for having, say, a vertical label at one position, and another vertical label at the co-aligned position. Specifying such payoffs requires building an abstraction for the curve detection problem, starting with the geometry underlying connections between columns in the superficial layers. This is followed with discussions of inter-layer intra-columnar processing, stressing first some of the non-linearities that are motivated by our earlier computational demonstrations of Canny edge detection. We then build upon this foundation by developing the geometry and the corresponding connections for texture and shading analysis, and for stereo.

3 The Curve Inference Problem

Assume that orientation selectivity defines a substrate for representing those tangents that approximate the curves that bound objects in the scene and that define highlights and other surface markings. We now analyze how long-range horizontal interactions can reduce the errors inherent in locally estimating tangent orientation. Orientation change can be used to localize corners and junctions (as may occur at the point where one object occludes another in depth), and for the perceptual integration necessary to connect fragments of the same contour, e.g., as a result of occlusion or specularities [ZDI89b]. Orientation change is also essential for curve detection: it enables us to establish *good orientation continuation* and it is the connection between tangents and curves. The mathematical theory of how tangents interact with one another is provided by differential geometry.

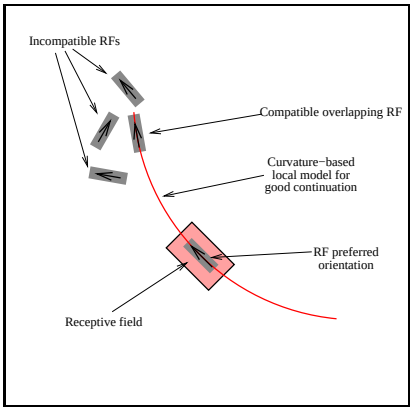

Figure 7: The problem of receptive field interaction. Which tangent (receptive field) continues the curve? Here strict co-aligned facilitation is clearly inappropriate.

The interaction between receptive fields for curve detection is illustrated in Fig. 7. Differential geometry dictates that the interactions must involve curvature [DC76]. The formal question is how to transport a tangent at one location to a nearby location. The analysis is not unlike driving a car, in the following sense. At each instant of time (that is, position along the curve) the axis of the car defines its (tangent) orientation, and the relationship between the orientation of the car at one instant with that at the next depends on how much the road bends; in operational terms, it depends on how much the steering wheel has to be turned during transport.

The endstopping property of visual cortical neurons is sufficent for representing curvature [DZC87]. (We shall return to the issue of how this might be developed by intracolumnar circuitry shortly.) The majority of superficial, interblob orientation selective cells in V1 are also endstopped to some extent, and these bi-selective dimensions of orientation and endstopping are precisely what is required to represent tangent and curvature.

We have used such notions of transport to derive the horizontal interactions [PZ89b, ZDI89a]; see Fig. 8 . Observe that they agree with available data that emphasize co-aligned facilitation [NF85, MAHG93, TGW86] as well as outliers [KIGW95a]. That is, most interactions are between cells with approximately the same orientation preference, but some (for the higher values of curvature) involve 50 - 60 degree shifts. A demonstration of the performance for this curvature-based system is in Figs. 9 and 10.

While it may seem counterintuitive at first, it is important to realize that building curvature directly into the system predicts a dominance of roughly co-aligned facilitation as well as the the spread of the distribution. Other models [YF97] adopt only the curvature = 0 case, and do not predict the non-co-linear data. Such results also question the interpretation of psychophysical data (e.g., [FHH93]).

Rather than make comparisons with "intuitive" notions of edge, as we did earlier, this example lets us appeal to the basic mathematics of the situation. Whitney has classified maps from smooth surfaces into smooth surfaces [GG73], and has shown that only two types of boundary points can occur generically (i.e., without changing under small changes in viewpoint): the fold and the cusp (the position where the fold disappears into the surface). Folds clearly indicate boundaries when viewed from a given position; in fact, the word implies that the tangent plane to the surface "folds" away from the viewer's line of sight. It is in this sense that they become singular (the two-dimensional tangent space collapses to a 1D tangent) and it is in this sense that boundaries are signalled. These ideas are illustrated in Fig. 11.

4 Intracolumnar Processes

The above example shows that boundaries need not be (in fact, are rarely) straight. Viewed locally, then, an approximation to a boundary over a very short distance is the tangent to the (boundary) curve. Viewed over a slightly larger neighbourhood, curvature begins to matter. Thus there are two problems that need to be addressed: (i) estimating tangents and (ii) estimating curvature.

Determining tangent directions with linear operators can be difficult, because different types of structure may fall within a single receptive field which would then be (inappropriately) averaged together. Thus we first consider the logical subunits of receptive fields and a mechanism to implement the non-linearities among them. The goal of these non-linearities is to prevent inappropriate types of structure from being combined into erroneous tangent estimates.

These local tangents are then combined to provide a curvature responses, based on the idea that curvature can be viewed geometrically as "deviation from straightness". Measurements over different spatial scales are utilized in this case. The result is two local circuits that illustrate the elaboration of

function across layers within a single orientation column. We now sketch each in turn, including new simulation results for the non-linearities in tangent estimation.

4.1 Logical Subunits and Shunting Inhibition

Boundaries that arise from surface folds, as discussed above, usually imply an intensity difference in the direction normal to the edge, with a generic dropoff in contrast in a neighbourhood around the edge (we shall consider this dropoff in detail in Sec. 6). Cracks appear as dark lines, with a bright-dark-bright contrast variation normally; while extended highlights can appear oppositely in contrast. Clearly, around "edges", there are different photometric events in in the direction normal to the edge [IZ95b].

In the tangential direction along an edge, continuity conditions arise. For the edge to be part of a boundary curve, locally its tangent must exist (almost everywhere) which implies that the limit as one point approaches another along it must exist. This can be realized as similar contrast signs along an edge element, for example.

Linear operators (or cells that combine receptive field sub-units linearly) would respond similarly to various contrast situations. They would respond to both cracks and edges, and low-contrast edges and complicated distributions of intensities would all give intermediate responses. In effect, all of these different physical situations would be mapped into a single response, making the inverse problem of labeling the type of physical situation from which the contrast difference arises untenable.

We have developed a class of non-linear local operators, called logical/linear operators [IZ90] to confirm such conditions operationally. Sub-unit combinators are defined with Boolean conditions testing for the above photometric and continuity conditions. If the structural conditions are met, they return the average; if not, they veto to zero. "Edge operators" are separated from "line operators", and edge patterns from multiple curves do not combine to give an artifactual response. These non-linearities are different from the compatibility fields above, because they are taking place at much smaller spatial scales.

We illustrate some of these cases below, but first set up the question of whether this type of non-linear response could be found in visual cortical circuits? We now illustrate one physiologically plausible possibility, beginning with the sub-unit interactions and then introducing the non-linearity.

The receptive fields of pyramidal neurons in layer VI are comprised of a number of zones, each of which may correspond, at least conceptually, to the receptive field of a layer V cell which synapses onto the layer VI cell. This view is supported by a Bolz and Gilbert experiment [BG86]: when one of the presynaptic layer V cells is pharmacologically inactivated, the layer VI cell shows no response to stimuli presented in the corresponding part of its receptive field. When the layer V neuron is re-activated, the receptive field of the layer VI cell returns to its original size.

This composite receptive field structure is a natural substrate for logical/linear operator construction. Both the tangential and the normal components consist of displaced linear operators joined together using logical/linear combinators. This is analogous to a number of layer V pyramidal cells sending axonal projections onto a layer VI neuron.

With the component construction now in place, logical/linear responses need a veto mechanism to assure no output response when at least one of the input conditions is not satisfied. Motivated by the demonstration of shunting inhibition in visual cortex [LBG98], and after considering several possibilities in simulation, we exploit this non-linear mechanism to dramatically decrease the efficacy of excitatory post-synaptic potentials, and thereby implement the veto. In simulations we approximated shunting inhibition by decreasing the passive conductance of the cell membrane [Koc99]. Following Borg-Graham et al [LBG98], we used a GABA type A synaptic model that was modified to increase the post-synaptic passive conductance by 200% for the duration of spiking and a short period of time (approximately 10ms) after each spike. This refractory period corresponded to the reuptake of the neurotransmitter and the closing of the ionic channels. Since we are working within the clique model for the columnar machine (Sec. 2), it is necessary for the post-synaptic conductance to remain increased between the spikes of the interneuron as it was firing at a high rate during the 3-5 spike burst. Since the soma of every cell in our model consisted of a single compartment [MS96], the post-synaptic conductance increase was a global change. Note that this does not make the model less realistic: the two excitatory synapses could easily be proximal to each other as well as proximal to the single inhibitory synapse.

We now illustrate the construction of a logical/linear edge operator over a receptive field with four subzones, Response to different input cases varies with the photometric configuration (Fig. 12). In cases B and C the operator should show no response due to the lack of an intensity edge. Case D is the optimal case, so the operator to respond maximally to this case. Case E

should generate a weak response because only the left side of the receptive field contains the optimal stimulus. In case F, a weak response should be observed because the edge has opposite contrast. Finally, case G should generate an intermediate response due to the weak edge in the right half of the receptive field.

While shunting inhibition provides us with a useful mechanism to implement the non-linear decision step of the logical/linear operator, the tricky design issue is how to activate the inhibitory interneuron when some of the conditions necessary for the operator to respond are not met. After experimenting with different configurations, a natural solution emerged by considering both the positive and negative contrast logical/linear edge operators at the same time and connecting them in a push-pull architecture that would enable each of the operators to inhibit its complement. This idea led to the circuit design shown in Fig 13. It is pleasing to observe this interaction between the photometric characterizations at the functional level and the neuronal characterizations at the implementation level.

The circuit consists of four layer V pyramidal cells, two layer VI pyramidal cells and two inhibitory interneurons. Each of the layer V cells corresponds to a subzone of the composite receptive field. The two layer VI cells represent the responses of the two opposite logical/linear edge operators. The performance of this circuit is perhaps best described using an example. Suppose that the input to the visual field is an ideal positive contrast edge, as in Fig. 13D. Given such input, the layer V cells corresponding to the top two subzones of the receptive field (subzones 1 and 2 in Fig. 12) will burst, exciting the top layer VI cell and causing it to burst as well. At the same time, the activity of these layer V cells will also excite the left inhibitory interneuron, which will in turn inhibit the lower layer VI cell. Since, however, the layer V pyramids corresponding to the bottom half of the receptive field are inactive, the lower layer VI cell shows minimal activity regardless of the inhibition. A much more interesting case to examine is the input pattern shown in Fig. 13G Since this is not an ideal edge, the cells monitoring the bottom half of the receptive field will burst, though with a lower frequency than the cells corresponding to the top half of the field.

Fig. 14 shows the results of our simulations. The top left of Fig. 14 corresponds to the optimal positive contrast edge input discussed above. As we can see, the top layer VI cell responds by bursting at a frequency much higher than baseline while the bottom layer VI pyramid shows only baseline activity.

The top right of Fig. 14 shows the results of a simulation, in which the input was a non-ideal edge such as in figure 13G. We simulated this by exciting the top two layer V cells with a strong current clamp (cca 1.2mA) and injecting the bottom right layer V cell with a weaker but still substantial current clamp (cca. 0.8mA). The high frequency activity of the top two layer V cells excited the left interneuron, which in turn inhibited the bottom layer VI cell. Had the inhibition not been in place, the bottom right layer V cell would have driven the bottom layer VI pyramid to fire at a higher than baseline rate. The top layer VI cell was excited by the top layer V cells and fired at a rate substantially higher than baseline.

In the bottom left of Fig. 14, we can see the results of a simulation with input similar to the pattern in figure 13E. The top left and the bottom right layer V cells were injected with a strong current clamp (1.2mA). The bottom right of Fig. 14 shows the results of stimulating all four layer V cells with a strong input (Fig. 13B).

The simulations illustrate how one cortical circuit can exhibit logical/linear-like response patterns. The compact push-pull circuit that gave the two contrasting logical/linear responses is natural and simple, and is intended as a first demonstration of concept. The effects of synaptic plasticity must also be considered in more realistic models, however, since facilitation and depression can profoundly affect the behavior of a cell. Many different types of such synapses have been found in the striate cortex, e.g., certain tufted layer V pyramidal cells that exhibit relatively fast frequency dependent depression [Mar97].

We now show how these operators can form the basis for the curvature response; again the differences between "edge" and "line" configurations will play a key role.

4.2 Building Up The Curvature Response

Our earlier scheme for integrating the information available from nearby tangent estimates is a function of curvature, but direct measurements of curvature are not common in the vocabulary of V1 functionality. Nevertheless, it is widely accepted that the (interblob) superficial layer neurons in V1 are often endstopped to some extent, and these bivariate selectivities for orientation and endstopping are precisely what we shall utilize in this sub-section.

Dobbins et al [DZC87] has developed a theory of endstopping which equates it to curvature. In its simplest form, an endstopped operator is

constructed by taking the difference between the responses of excitatory and inhibitory component cells. Matching these component cells for position, receptive-field profile, and orientation tuning, but not with size, results in a response that is sensitive to curvature. In particular, when the inhibitory component covers a larger spatial scale than the excitatory compent the combined response will peak for a specific, typically non-zero curvature.

While Dobbins' analysis assumes that the component operators are linear (combined via a "positive part" non-linearity to model the base firing rate and to prevent negatives from combining incorrectly), we now see the central role of the logical/linear non-linearities - to guarantee that receptive fields respond to tangents if and only if they properly exist within the component receptive fields. The large "surround" operator component in the curvature response must arise from a long curve for it to be meaningful; not from an unrelated series of fragments. Furthermore, Dobbins et al [DZC89] have found that even-symmetry (large) receptive fields provide an estimate of the magnitude of the curvature, while odd-symmetry surround operators provide information about the sign of curvature.

4.3 The Position-Orientation Representation

We now take a step back from the circuitry, to illustrate how this geometry is represented in the columnar machine sketched earlier (Fig. 6). As in that figure, the different orientation possibilities can be viewed as a structure "on top of" the image, with retinotopic (x,y) retinotopic coordinates extended into a third dimension ("height") corresponding to orientation. This orientation axis is different from the x and y axes, because orientation wraps around 2π with the circle S^1 being its domain. Thus this (position, orientation) space is not modeled as (x,y,z), but as (x,y,θ), where θ is the tangent angle. A point in this space is a point in $R^2 \times S^1$.

It is instructive to consider how different curves in the (x, y) plane lift into $R^2 \times S^1$, to understand some of its advantages for the cortical columnar machine. Not much happens for a straight line in the plane, which lifts to a "horizontal" straight line in $R^2 \times S^1$ at a "height" dependant only on the angle θ. A smooth, closed curve in the plane lifts into a smooth, closed curve in $R^2 \times S^1$. A real difference arises when we consider curves that are continuous but with corners. Such events are important because they could signal a monocular occlusion cue and as we mentioned earier, all these discontinuities in orientation lift into broken curves in $R^2 \times S^1$ (Fig. 15). In

the following we denote it as the (x, y, θ) space.

5 The Geometry of Texture Flows

With this understanding of the inference of tangent maps for individual curves, we move to patterns of multiple curves. Examples include pinstriped material and artist's etchings, animal coats and zebra's stripes. For such patterns orientation is distributed over two-dimensional regions, and as the requirement for perfect continuations is relaxed we obtain texture flows.

The importance of locally (almost) parallel structure has been observed psychologically [Kan79, Gla69, TR90]. In particular, the human visual system has the tendency to organize and group parallel structure into coherent units. Examples include Kanizsa's "social conformity of a line" [Kan79, GZ31], and Glass's random dot moire patterns [Gla69, Pra86]. In all cases, this organization drastically affect the interpretation of a scene bothin in 2D and in 3D (e.g., [TR90]).

Informally, texture flows are defined by their orientation content - a dense visual percept characterized by local parallelism and slowly varying dominant local orientation. However, texture flows are not simply a matter of completing a number of almost "parallel" curves, but rather involve dense interpolation of a field of orientations (Fig. 16). This should be compared with the gap-completion property for curves (e.g., [KFP99, WJ97, HvdHK94]).

Represented as a scalar orientation function, the local behavior of the flow is governed to first-order by its gradient. Unfortunately, this quantity does not capture all of the relevant geometry. As with curves, psychophysical evidence suggests that texture flow perception is linked to curvature [LZ85, HOZ90].

We model the differential geometry of texture flows by applying the frame field representation [BSZ01a, BSZ01b]. A frame is essentially a *local* coordinate system attached to each point in the flow in a natural way; one axis points in the direction tangential to the flow and the other in the normal direction. It is this frame that serves as a local "piece" of a texture flow in the sense that it generalizes the tangent (which we earlier used as a local "piece" of a curve). It remains to develop an analog to the osculating circle, or the local curvature, which was used to transport tangents, and to determine what needs to be measured locally. (Recall previously, for curves, local orientation and curvature were measured.)

The geometry of a texture flow is given by the rotation of the Frenet frame as it is displaced a small distance in an arbitrary direction within the flow. The resulting structure is modeled by the Cartan connection equation [O'N66]. Since the rotation may vary for different displacment directions, and since two coordinates are required to specify this displacement, it follows from the connection equation that the local behavior of a texture flow patch can be fully approximated by two scalars – two curvatures – which describe the rate of rotation along two independent directions. Since the tangential and normal directions are the natural local basis, we define both the *tangential curvature* and the *normal curvature* at each point; they represent the generalization of curvature (for curves) to texture flows and fully specify the behavior of the flow to first order approximation. To illustrate: for the case of strictly parallel straight lines there is no difference in the orientation for a displacement in either the tangential or the normal directions, while a texture flow composed of nested circles having the same center the orientation changes with displacement only in the tangential but not in the normal direction. On the other hand, a texture flow composed of straight line all converging and meet at the same point has zero tangential curvature but non-zero normal curvature everywhere.

Following the discussion in the introduction, it is now clear that the biological machinery of the primary visual cortex is capable of measuring the tangential direction, and thus the normal one as well. Hence the basic components for representing texture flow curvatures do exists in V1. While it was argued that endstopping relates to (tangential) curvature, the representation of normal curvature in V1 is still a matter of speculation. From a computational point of view, however, we end up with a natural representation for texture flow based on the tangent, the tangential curvature, and the normal curvature. Unfortunately, earlier computational efforts have omitted the normal curvature altogether (e.g., [VvVvdW98]).

If 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 at a point should induce a field of coherent measurements (i.e., an orientation function) in the neighbourhood of that point. This field, which we call the texture flow *osculating object*, generalizes the osculating circle for curves, and coherence of local measurements of texture flow can then be determined in a manner analogous to co-circularity for tangents to a curve [PZ89a] (Fig. 17).

For that purpose, we have developped a unique texture flow osculating object based on formal notions of good continuations. In (x, y, θ) space this object takes the form of a right helicoid [BSZ01b, BSZ01a] and it is the only geometrical object that does not prefer either of the two curvatures over the other. Using this object we then designed compatibility structures to be used in the columnar machine (Section 2), examples of which are illustrated in Fig. 18. The same "neural" machinery equipped with this new set of "long range horizontal connections" is now able to integrate and infer 2D coherent texture flows, as is exemplified in Fig. 19.

6 Shading Flows and Fold-type Edges

While texture flows are a generalization of perceptual organization beyond curves, they are still a somewhat special class of patterns. It is therefore difficult to understand why the cortex might have evolved specialized circuitry for them. We believe an important part of the answer to this question is that there are other, more universal perceptual features that share the basic structural properties of texture flows. Prominent among these features is the *shading flow field* - the vector field of tangents to the iso-brightness contours (or intensity level sets) of the gray level image of smooth surfaces [BZ96]. Measured from shading distribution as opposed to surface markings, the shading flow is directly analagous to the texture flow discussed earlier. Its geometry is a precursor to shape [HB89], and its interaction with edge geometry provides useful information for edge classification [HZ01].

When applied to shading flow fields instead of textures, a successful organization of coherent flow structure constrains the geometry of the smoothly curved surface which gives rise to the scene. The singularities of the shading flow field can arise from shape discontinuities and spectral highlights, and thus its reliable recovery, while preserving its discontinuities and singularities, is an important step toward the robust interpretation of shape from shading.

The shading distribution takes on a particular form in the neighbourhood around an edge. Recall that we earlier discussed how edges which correspond to object boundaries arise when the tangent plane to a surface "folds" out of sight. Such edges are distinct in appearance in that they enjoy a stable pattern of shading with respect to the edge. This shading pattern is illustrated in Fig. 20. As smooth surfaces fold away from the viewer, the shading

flow field becomes tangent to the edge. At the same time, the fold side of the edge "cuts" the background scene, which implies that the background cannot exhibit this regularity in general. Cuts, therefore, are those kinds of boundaries in which the shading distribution is transverse to the edge. This is fortunate as it naturally suggests a type of "figure" boundary, because it shows which is the occluding side and which is the occluded side of an edge [HZ01]. This complements earlier work which classified shadow edges on the basis of the shading flow field [BZ96].

Since shading information shares the same geometrical characteristics of texture flow, it opens up the possibility that shading flows may be signaled by cells that are orientation selective and tuned to low spatial frequencies (those tuned to higher spatial frequences would localize edges, as before). The geometrical analysis described in this section thus suggests two additional kind of computations and orientation-based interactions in the columnar machine. Firstly, the computation of coherent shading flow can be done through texture-flow-like interactions (see Section 5) between orientation-selective cells of low frequency, and the result can be used for the interpretation of shape from shading. Secondly, the interaction between shading and boundaries for edge classification may be achieved though interaction between low spatial frequency cells and high spatial frequency cells. Thus, following the ideas above, signaling a fold-type edge could follow from an iso-orientation interaction between low spatial frequency (shading) and high spatial frequency (edge) cells, while cut-type edges requires interactions between cells of different orientations. This is illustrated in Fig. 21.

7 The Geometry of Stereo Correspondence

As our final example of interactions between orientations we turn to stereo. Thus far we have been considering interactions between orientations within cells driven by one eye; we now consider both of the ocular dominance bands. Abstractly this implies an important construction for the columnar machine: the "product" of two machines, one for the left eye and the other for the right eye. Mathematically this suggests working in the *product space* $(R^2 \times S^1) \times (R^2 \times S^1)$ and designing compatability structures that also take a "product" form.

We have developed such a product structure and an algorithm for computing stereo correspondences [AZ00] that generalizes the tangent fields for

plane curves to those for general space curves. A curve in three space can be described by the relationships between its tangent, normal and binormal [DC76]. As the curve moves across depth planes, there exists a positional disparity between the projection of the curve in the left image and the projection in the right image. However, there also exist higher order disparities, in particular ones in orientation. It is these types of relationships that can be capitalized upon when solving the correspondence problem. Rather than correlating left/right image pairs, we require that there exists a curve in three space whose projection in the left and right image planes is commensurate with the locus of tangent pairs, and their orientation, in a neighbourhood of the proposed match.

For the stereo correspondence problem, we are given two edge maps (one for the left camera and one for the right); each of these will be consistent (in the sense that they satisfy the monocular transport constraint); our goal now is to make them consistent with a local approximation to the space curve from which they project. Again, an osculating object is required and the one that we have derived [AZ00] takes the form of a helix in (x, y, z). Based in this model it is possible to construct the compatibility fields that are required to facilitate responses of tangent pairs that are consistent with the same space curve. Examples of two compatibility fields are shown in Fig. 22 while the computation that results in using these structure in the columnar machine is illustrated in Fig. 23.

To conclude this section we observe that once again an early vision task, this time stereopsis, can be computed using the columnar machine when equipped with the appropriate orientation-based connectivity structure. Once again, these structures suggest a majority of connections between cells of similar orientations. Yet, connections between cells of significantly different orientations are intrinsic, not accidental. The only new aspect in connectivity that is required by the stereo compatibilities is links between monocular cells to binocular ones. Interestingly, some evidence showing such interactions already exist in the literatire [MAHG93].

8 Conclusions

We have argued for a differential geometric approach to vision, and have shown how this can be supported by the columnar architecture of visual cortex. Our computational model comprises non-linear local orientation mea-

surements and the refinement of these measurements using context, provided by curvature, through a cooperative network. We have demonstrated how the local orientation measurements could be computed by intra-columnar neural circuitry, via shunting inhibition. The observed physiology of inter-columnar connections suggests a more complex model for contour integration than co-aligned facilitation; we hypothesize that the connections implement curvature, which follows from the computational considerations of our model. Finally, we have shown how the differential geometric framework naturally extends to handle the analysis of texture, shading, and stereo correspondence, all using the same basic columnar architecture. We hope that these developments will provide a basis for future neurophysiological experiments.

References

[AME85] J. Allman, F. Miezin, and McGuinness E. Stimulus specific responses from byond the classical receptive field: Neurophysiological mechanisms for local-global comparisons in visual neurons. *Annu. Rev. Neurosci.*, 8:407–430, 1985.

[AZ00] S. Alibhai and S. W. Zucker. Contour-based correspondence for stereo. In *Computer Vision - ECCV 2000, Lecture Notes in Computer Science 1842*, June 2000.

[BG86] J. Bolz and C.D. Gilbert. Generation of end-inhibition in the visual cortex via interlaminar connections. *Nature*, 320:362–365, 1986.

[BSZ01a] O. Ben-Shahar and S.W. Zucker. Flowing towards coherence: On the geometry of texture and shading flow. In *IEEE Computer Society Workshop on Perceptual Organization in Computer Vision*, 2001.

[BSZ01b] O. Ben-Shahar and S.W. Zucker. On the perceptual organization of texture and shading flows: From a geometrical model to coherence computation. In *Proc. Computer Vision and Pattern Recognition*, pages 1048–1055, 2001.

[BZ96] P. Breton and S.W. Zucker. Shadows and shading flow fields. In *Proc. Computer Vision and Pattern Recognition*, 1996.

[BZBF97] W.H. Bosking, Y. Zhang, Schofield B., and D. Fitzpatrick. Orientation selectivity and the arrangement of horizontal connections in the tree shrew striate cortex. *J. Neurosci.*, 17(6):2112–2127, March 15 1997.

[Can86] J. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Machine Intell.*, 8(6):679–698, 1986.

[DC76] M Do Carmo. *Differential Geometry of Curves and Surfaces.* Prentice-Hall, Inc., 1976.

[DZC87] Allan Dobbins, Steven Warren Zucker, and Max S. Cynader. Endstopped neurons in the visual cortex as a substrate for calculating curvature. *Nature*, 329:438–441, 1987.

[DZC89] Allan Dobbins, Steven Warren Zucker, and Max S. Cynader. Endstopping and curvature. *Vision Research*, 29:1371–1387, 1989.

[FHH93] D. Field, A. Hayes, and R. Hess. Contour integration by the human visual system: evidence for a local association field. *Vision Research*, 33:173 – 193, 1993.

[GG73] M. Golubitsky and V. Guillemin. *Stable Mappings and their Singularities.* Springer-Verlag, 1973.

[Gla69] L. Glass. Moiré effect from random dots. *Nature*, 223(5206):578–580, 1969.

[GZ31] A. Galli and A. Zama. Untersuchungen über die wahrnehmung ebener geometrischen figuren die ganz oder teilweise von anderen geometrischen figuren verdecket sind. *Zeitschrift für Psychologie*, 123:308–348, 1931.

[HB89] B.K.P. Horn and M.J. Brooks, editors. *Shape from Shading.* MIT Press, Cambridge, MA, 1989.

[HOZ90] Y. Hel Or and S.W. Zucker. Texture fields and texture flows: Sensitivity to differences. *Spatial Vision*, 4(2/3):131–139, 1990.

[HvdHK94] F. Heitger, R. von der Heydt, and O. Kubler. A computational model of neural contour processing: Figure-ground segregation and illusory contours. In *Proc. From Perception to Action Conf.*, pages 181–192, 1994.

[HW77] 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.

[HZ83] Robert Hummel and Steven Warren Zucker. On the foundations of relaxation labeling processes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 6:267–287, 1983.

[HZ01] P.S. Huggins and S.W. Zucker. How folds cut a scene. In *Proc. of the* 4^{th} *Int. Workshop on Visual Form*, 2001.

[IZ90] Lee Iverson and Steven Warren Zucker. Logical/linear operators for measuring orientation and curvature. Technical Report TR-CIM-90-06, McGill Research Center for Intelligent Machines, Montreal, Canada, July 1990.

[IZ95a] L.A. Iverson and A.W. Zucker. Logical/linear operators for image curves. *IEEE Trans. Pattern Anal. Machine Intell.*, 17(10):982–996, 1995.

[IZ95b] Lee A. Iverson and Steven W. Zucker. Logical/linear operators for image curves. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 1995.

[Kan79] G. Kanizsa. *Organization in Vision: Essays on Gestalt Perception.* Praeger Publishers, 1979.

[KFP99] 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.

[KIGW95a] M. Kapadia, M. Ito, C. Gilbert, and G. Westheimer. Improvement in visual sensitivity by changes in local context: Parallel studies in human observers and in v1 of alert monkeys. *Neuron*, 15:843–856, 1995.

[KIGW95b] M.K. Kapadia, M. Ito, C.D. Gilbert, and G. Westheimer. Improvement in visual sensitivity by changes in local context: Parallel studies in human observers and in v1 of alert monkeys. *Neuron*, 15:843–856, October 1995.

[Koc99] C. Koch. *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press, New York, 1999.

[KWG99] M.K. Kapadia, G. Westheimer, , and C.D. Gilbert. Dynamics of spatial summation in primary visual cortex of alert monkey. *Proc. Natl. Acad. Sci. U.S.A.*, 96(21):12073–12078, October 12 1999.

[LBG98] Y. Fregnac L.J. Borg-Graham, C. Monier. Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature*, 393:369–373, 1998.

[LZ85] N.K. Link and S.W. Zucker. Sensitivity to corners in flow patterns. Technical Report TR-85-4R, Computer Vision and Robotics Laboratory, Department of Electrical Enginnring, McGill University, April 1985. Image texture flow.

[MAHG93] R. Malach, Y. Amir, M. Harel, and A. Grinvald. Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo targeted biocytin injections in primate striate cortex. *Proc. Natl. Acad. Sci. U.S.A.*, 90:10469–10473, November 1993.

[Mar97] H. Markram. A network of tufted layer 5 pyramidal neurons. *Cerebral Cortex*, 7(6):523–533, 1997.

[MC82] G. Mitchison and F. Crick. Long axons within the striate cortex: Their distribution, orientation, and patterns of connections. *Proc. Natl. Acad. Sci. U.S.A.*, 79:3661–3665, June 1982.

[MF76] M. Maffei and A. Fiorentini. The unresponsive regions of visual cortical receptive fields. *Vision Res.*, 16:1131–1139, 1976.

[MS96] Z.F. Mainen and T.J. Sejnowski. Influence of dendritic structure on firing pattern in model neocortical neurons. *Nature*, 382:363–366, 1996.

[MZ92] D. A. Miller and S. W. Zucker. Efficent simplex-like methods for equilibria of nonsymmetric analog networks. *Neural Computation*, 4:167–190, 1992.

[MZ99] D. A. Miller and S. W. Zucker. Computing with self-excitatory cliques: A model and an application to hyperacuity-scale computation in visual cortex. *Neural Computation*, 11:21–66, 1999.

[MZ00] D.A. Miller and S.W. Zucker. Cliques, computations, and computational tractability. *Pattern Recognition*, 22:535–542, 2000.

[NF85] J. Nelson and B. Frost. Intracortical facilitation among co-oriented, co-axially aligned simple cells in cat striate cortex. *Exp. Brain Res.*, 61:54–61, 1985.

[O'N66] B. O'Neill. *Elementary Differential Geometry.* Academic Press, 1966.

[Pra86] K. Prazdny. Some new phenomena in the perception of glass paterns. *Biol. Cybern.*, 53:153–158, 1986.

[PZ89a] P. Parent and S.W. Zucker. Trace inference, curvature consistency, and curve detection. *IEEE Trans. Pattern Anal. Machine Intell.*, 11(8):823–839, 1989.

[PZ89b] Pierre Parent and Steven Warren Zucker. Trace inference, curvature consistency and curve detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11(8):823–839, August 1989.

[TGW86] D.Y. Ts'o, C.D. Gilbert, and T.N. Wiesel. Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.*, 6(4):1160–1170, April 1986.

[TR90] 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.

[VvVvdW98] 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.

[Wer55] 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.

[WJ97] 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.

[YF97] S. C. Yen and L. Finkel. Salient contour extraction by temporal binding in a cortically-based network. *Advances in Neural Information Processing Systems*, 9, 1997.

[ZDI89a] Steven Warren Zucker, Allan Dobbins, and Lee Iverson. Two stages of curve detection suggest two styles of visual computation. *Neural Computation*, 1:68–81, 1989.

[ZDI89b] S.W. Zucker, A. Dobbins, and L. Iverson. Two stages of curve detection suggest two styles of visual computation. *Neural Comput.*, 1:68–81, 1989.

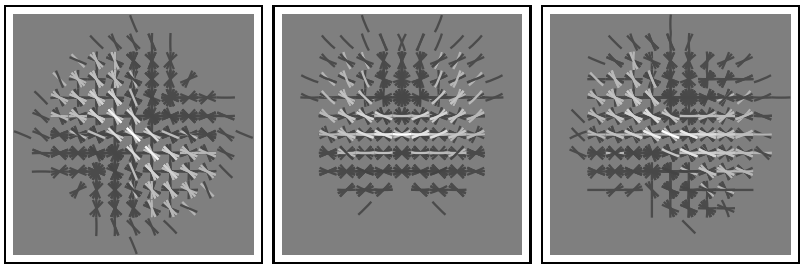

Figure 8: The geometry of inter-columnar interactions for curves. Co-circularity indicates how consistent a neighbouring tangent is with a given tangent at the origin (center of display). Since the analysis is local, smooth edge curves can be approximated by their osculating circle, and neighbouring tangents can be transported along it to fix a measure of mis-match. Our networks select those tangents that minimize such mismatches [HZ83, MZ99]. In network implementations, the transport results are precomputed and embedded in the connections. A very small mismatch results in an excitatory connection, and a larger mismatch in an inhibitory one. Three examples of the compatibilities derived from co-circularity are shown. The bars indicate orientation preference, and all compatibility fields are with respect to the central neuron (shown at maximal brightness). The brightness for each bar is the strength of the synapse with the central tangent, white indicating excitatory connections, black indicating inhibitory connections. Multiple bars at the same position indicate several cells in the same orientation hypercolumn. The connections are intended to model long-range horizontal interactions. Three cases are shown: co-aligned facilitation, curved a large amount in the negative sense and, curved a small amount. Notice in particular that most of the excitatory connections are between co-aligned cells, given the loose definition of alignment commonly used in the physiological literature (e.g., $\pm 15°$); however, in the high curvature example there are cases of excitatory connections with approx. 50 deg relative orientation.

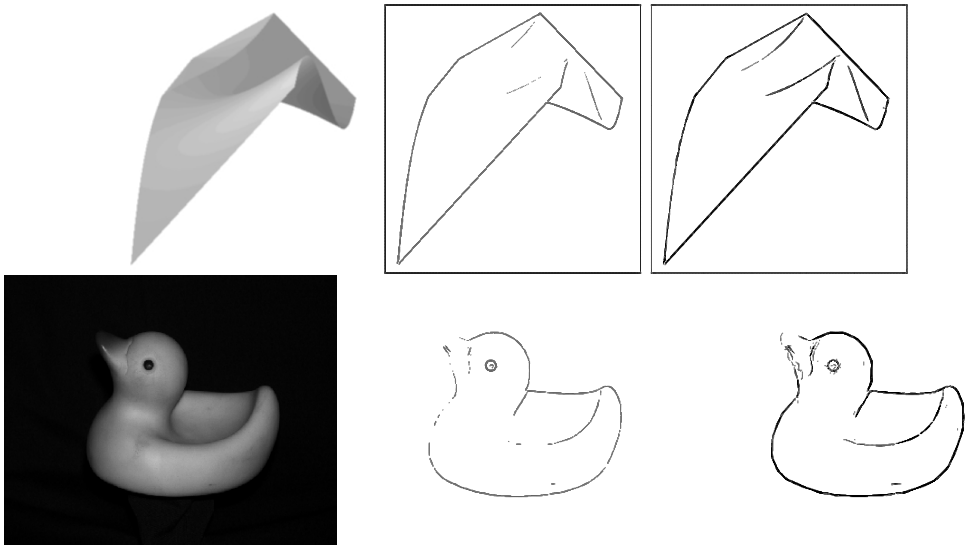

Figure 9: Results of curve detection using our network, with curvature-based connections (see Fig. 8). (LEFT) Original image, (CENTER) intial edge tangent measurements, (RIGHT) final edge tangents.

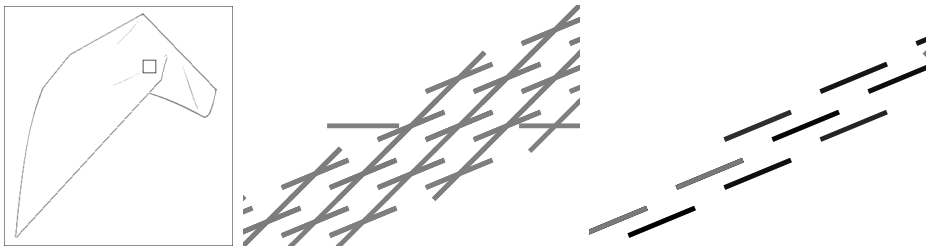

Figure 10: Performance of our cooperative model for boundary and edge detection and localization. (LEFT) the sheet edge map and a region of interest. (MIDDLE) Zoom on the ininitial edge measurements [IZ95b] in the region of interest. (RIGHT) Result of 5 iterations of relaxation labeling [HZ83] with Co-circularity compatibilities. Note the localization of the edge and the elimintion of sparious mesaurements.

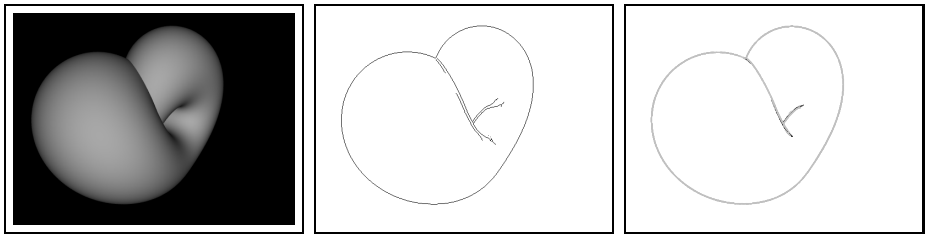

Figure 11: The motivation from differential topology for early vision. (LEFT) The image of the Klein bottle shows how "T"-junctions can arise from occlusion relationships (e.g., at the top and center of the figure), and how certain interior edges can end (e.g., where the fold smoothly joins the body). (MIDDLE) The Canny edge structure is inconsistent with both of these topological observations. Notice how the boundary "T"-junction is not connected, how it smooths the outline, and how the interior folds are blurred into the shading. (RIGHT) Output of the logical/linear operator [IZ95a]. Notice how the "T"-junctions are maintained, and how the contours end at cusps.

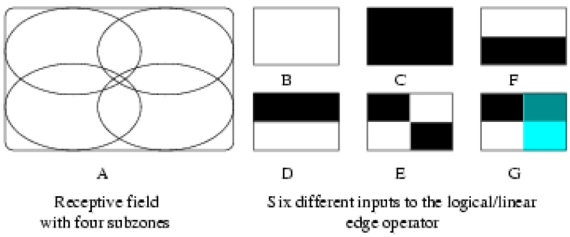

Figure 12: Logical/linear receptive field subunits and experimental stimuli used in our simulations.

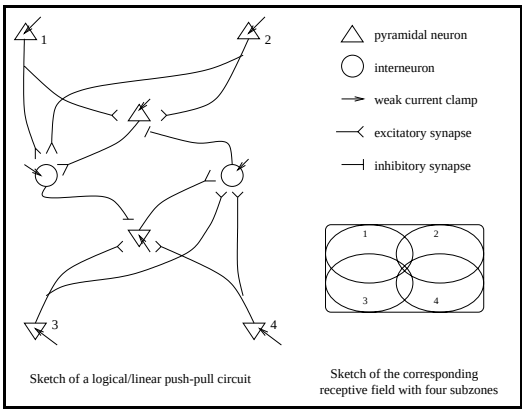

Figure 13: A neural circuit implementing a logical/linear operator.

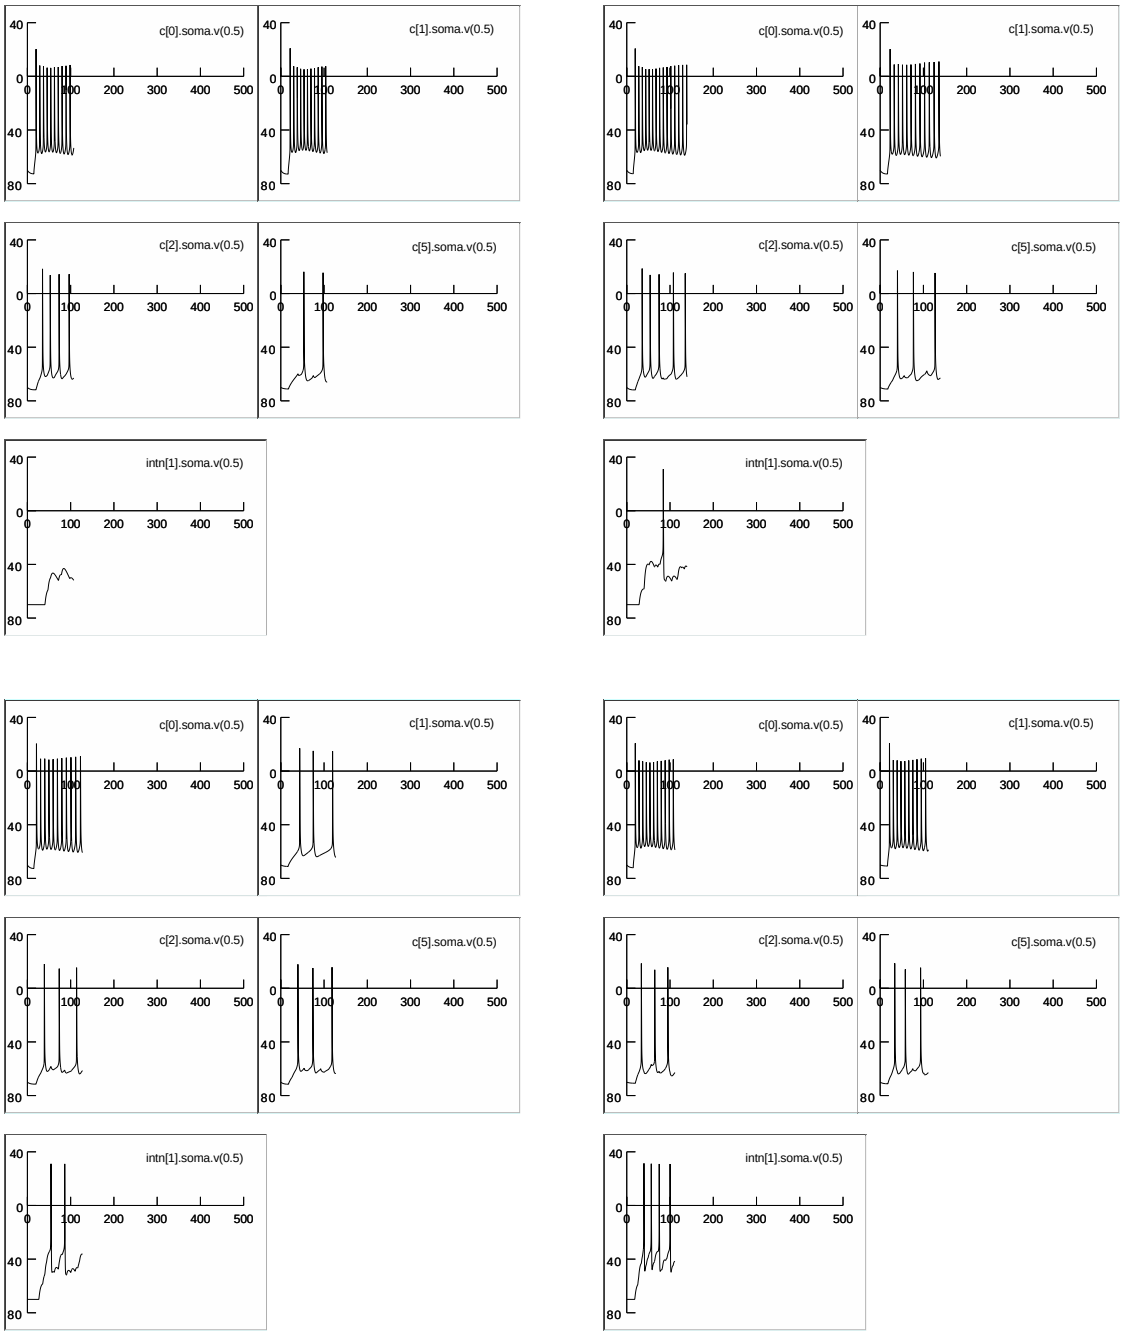

Figure 14: Simulation results on the logical/linear circuit shown in Fig. 12. See text for details.

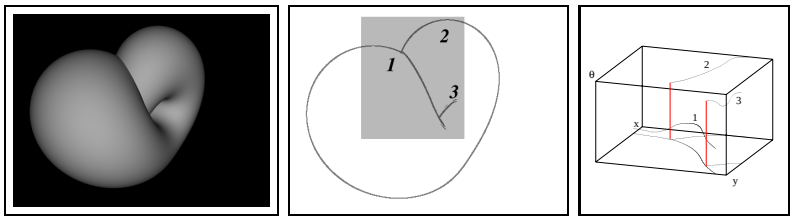

Figure 15: An illustration of the lift into position (x,y), orientation (θ) space. Three contour fragments from the tangent map in Fig. 5(right) are highlighted. Notice how the discontinuity in orientation at the 1-2 T-junction is separated, highlighting multiple orientations at the same position, the natural columnar representation for orientation discontinuities.

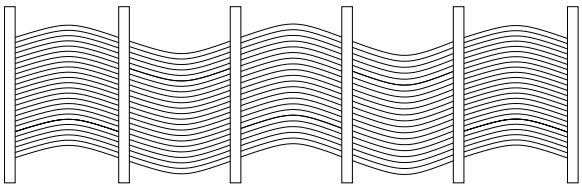

Figure 16: Example of a texture flow and it's perceptual completion quality. The wavy surface appears to form a single coherent unit behind the occluders, even though there are a different number of line segments in each visible region. This demonstrates that orientation is distributed densely over a region, and that computing texture flows is not simply the simultaneous completion of a number of "parallel" curves.

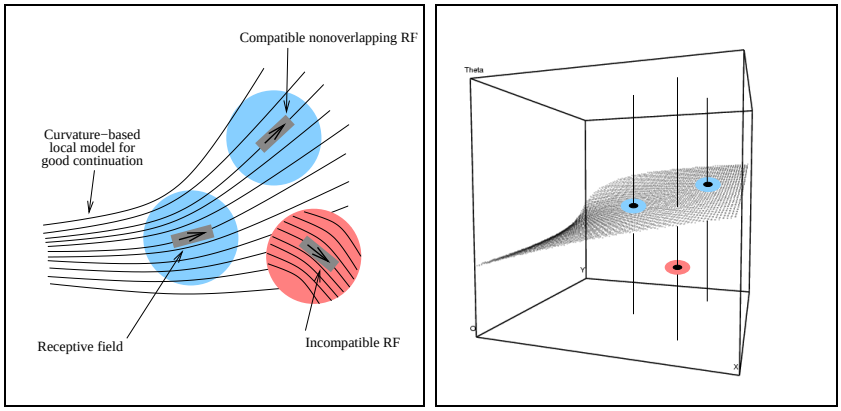

Figure 17: The problem of infering coherent texture flow both in the image plane (LEFT) and in the lifted space of (x, y, θ) (RIGHT). Given two nearby texture flow patches of given curvatures they may or may not be part of the same coherent structure. What affects this coherence is not only the local orientation of the patches, but their curvatures as well. In (x, y, θ) this task can be visualized as associating both texture patches, now represented as 3D points, to the same surface (whose height represents orientation). Here, the two blue patches are coherent and should facilitate each other while the red one is incompatible to both.

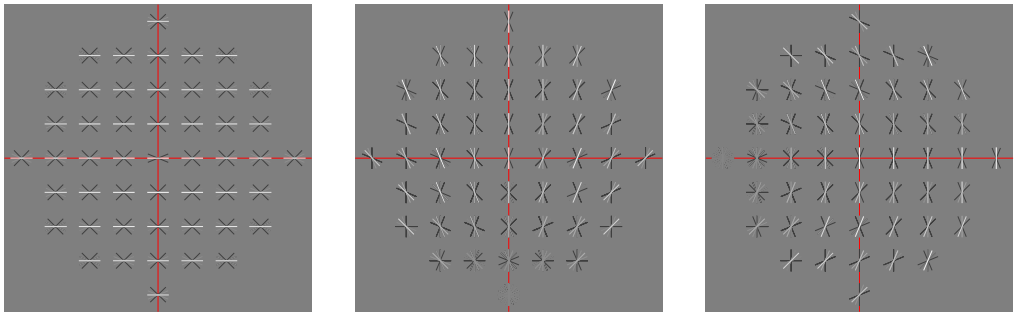

Figure 18: Three examples of texture flow compatibility structures for various tuning values for orientation and curvatures. Bright/dark indicates excitatory/inhibitory connections. (LEFT) Horizontal orientation tuning with both curvatures equal zero. (CENTER) Vertical orientation tuning positive normal curvature. (RIGHT) Vertical orientation tuning positive tangential curvature.

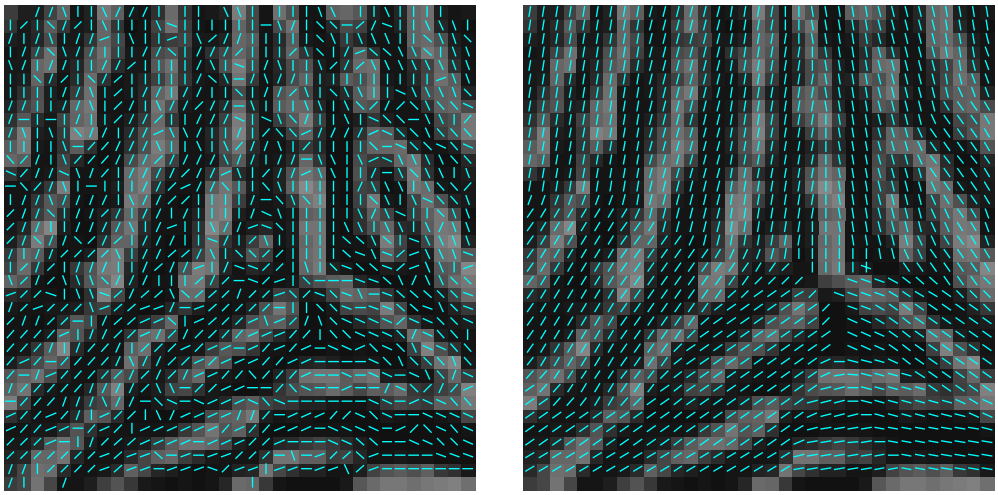

Figure 19: Example of texture flow computation with the columnar machine. In this piece of a fingerprint image the perceptual structure is overwhelmed by the noisy measurements of local orientation (LEFT). Based on the right helicoidal model for good continuation discussed in the text (see [BSZ01b] for details), the columnar machine can extract the perceptually coherent structure (RIGHT) in few iterations of relaxation labeling.

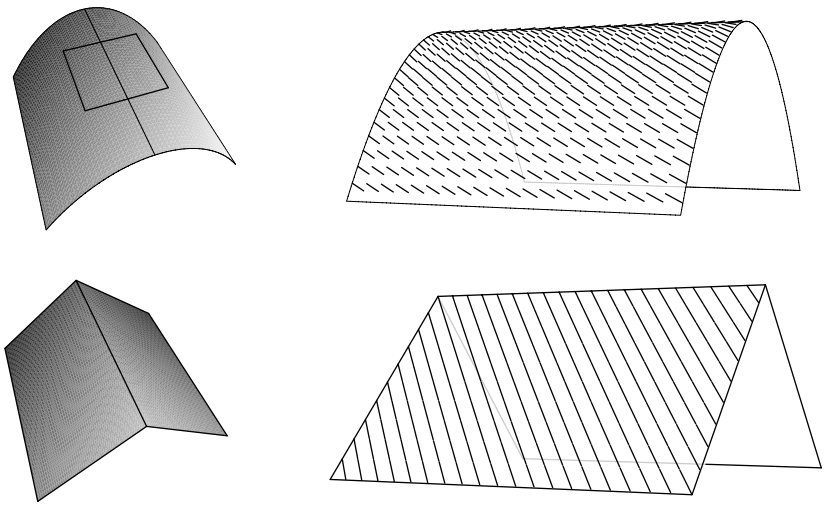

Figure 20: Illustration of shading flow in the neighbourhood of an edge. When a shaded surface (LEFT) is viewed such that an edge appears, the shading flow field (right) takes on different appearances depending on the nature of the edge. A fold occurs (TOP) when the surface bends smoothly away from the viewer (the typical occlusion case), and the shading flow field appears tangent to the edge. At a cut, the surface is discontinuous (or occluded), and shading flow is generally non-tangent to the edge.

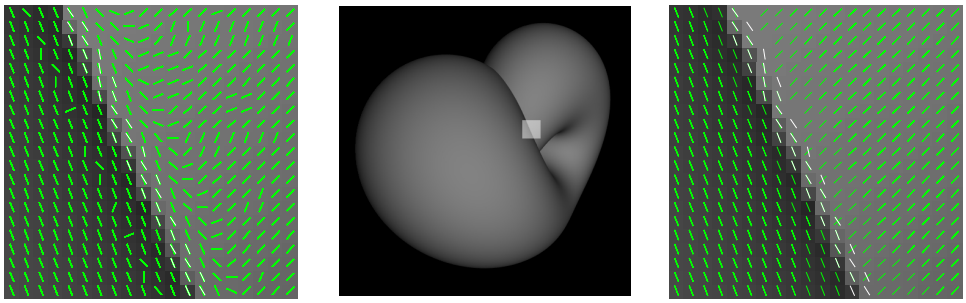

Figure 21: Example of shading flow relaxation and the possible interaction between the flow and the edges to signal either fold- or cut-type edges. Note how the initial shading flow (LEFT) of the marked region of interest (CENTER) becomes fully coherent after few iterations of relaxation labeling based on the texture flow compatibilities (RIGHT), and how the relationship of the relaxed shading to the active high frequency cells (in white) indicates a fold on the left and a cut to the right. Indeed, the surface to the left of the edge occludes the one to its right, as is clearly seen in the image itself.

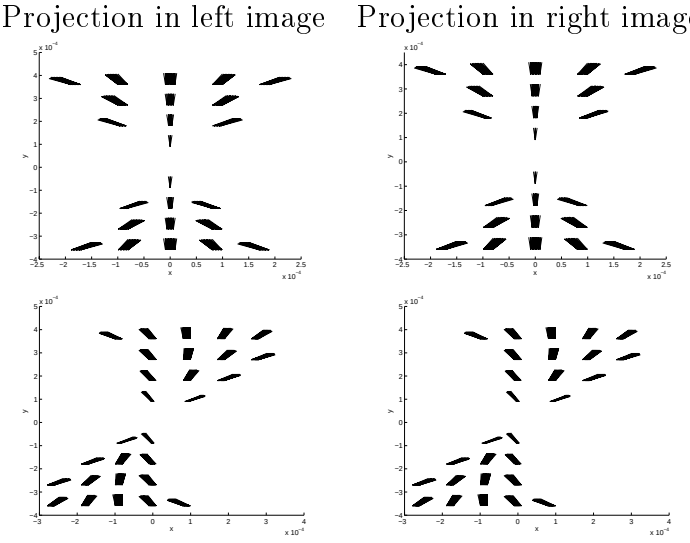

Figure 22: Two examples of facilitory compatibility structures for stereo are shown. Note how they incorporate both position and orientation disparity; this is especially evident in the lower pair of compatibility fields which are tuned for non zero values of curvature.

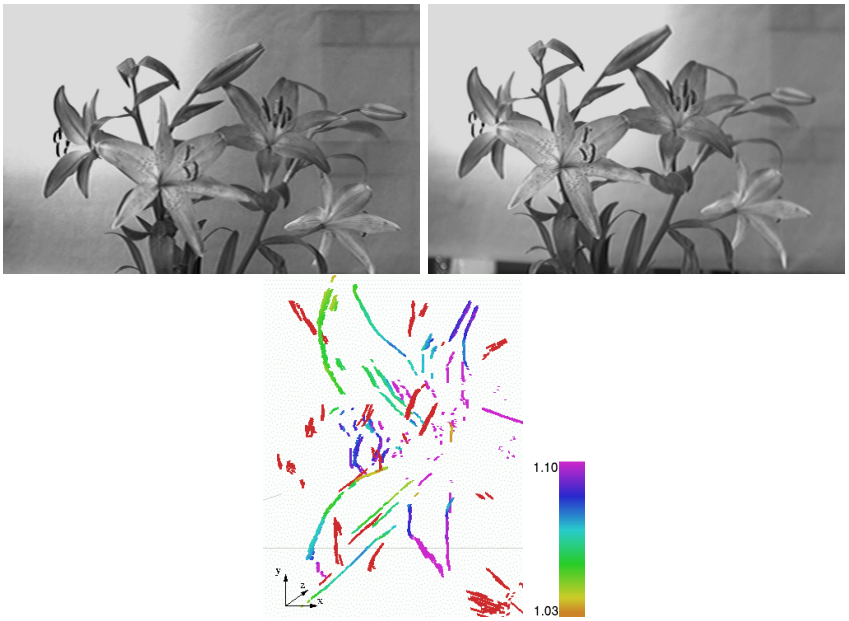

Figure 23: Example of a depth map computation using the stero compatibility structures in the columnar machine. The edge depth map is color coded (see bar on the right) and for clarity we show a magnified detail.