

The computational connection in vision: Early orientation selection

STEVEN W. ZUCKER

McGill University, Montreal, Quebec, Canada

It is widely accepted that computer implementations can play a role in verifying psychological theories. In this paper, I argue for a much broader and more abstract role for computation, in particular, one that includes formulation as well as verification. Consideration of issues of abstract computation—what should be computed and how—provides a level of analysis between ecological issues at the problem level and realization issues at the physiological level. This is the computational connection. The paper reflects my personal experience so that my argument can be made concretely. I concentrate on the evolution of one theory of orientation selection, and I show how we were led to differential geometry from “line detectors”; how parallel, distributed computational modeling led to novel proposals regarding curvature estimation; and how these proposals predicted psychophysical sensitivity to discontinuities.

The leap from neurophysiological observations through psychophysics to function has been a cornerstone of how theories are inferred in vision. This leap, however, has been clouded by the informality of many of the steps and the lack of feedback among them. Incompletely or heuristically specified theories are notoriously difficult to validate (or, for that matter, to invalidate). Computers and the theory of computation have begun to change this, however, and in this article I attempt to illustrate how computation has played a central role in my developing ideas about one aspect of early vision: orientation selection. I shall argue that computational theorizing, in addition to precision and validation, has led to several surprises regarding the subtlety of interpreting experimental data.

Constraints on visual theorizing arise from many different sources, and the general idea behind computational modeling is to utilize as many of these constraints as possible. Physical structure gives rise to constraints about the world, and projection gives rise to image constraints. Physiology provides constraints about (the details of) mechanisms, and psychophysics about their input/output properties. Viable theories are those that can account for a collection of these constraints at different levels of abstraction; successful theories are those that lead to new predictions.

This paper is a personal one and focuses on the evolution of my ideas about early orientation selection; for a wider view of the role of computation on the evolution of vision, see Marr (1982) and Zucker (1981; in press a).

S. W. Zucker is a senior fellow of the Canadian Institute for Advanced Research. This research was sponsored by NSERC Grant A4470. A. Dobbins, R. Hummel, L. Iverson, Y. Leclerc, N. Link, and P. Parent all participated in essential ways to the research project on orientation selection. The author's mailing address is: 3480 University St., McGill University, Montreal, Quebec, Canada H3A 2A7.

THE FORMATION AND INFERENCE OF IMAGE INFORMATION

Perhaps the most common definition of vision in computational circles is as *inverse optics*; that is, images are formed by a physical process in which light reflects off of objects and is then imaged by an optical system. Since the goal of vision is to produce descriptions of physical scenes, if the above process could be inverted—that is, if the light could somehow be run backward in time from the image to the world—then the physical properties of the objects could be inferred.

Inverse optics is clearly impossible without extra constraints, and much of the computational literature describes various attempts to find specific problems that can be solved under specific constraints (Ballard & Brown, 1982; Levine, 1985; Zucker, 1981, in press a). Many of these relate directly to the image formation process (e.g., various attempts at recovering *shape-from-shading* information), but further organization is necessary for the program resulting from these attempts to be successful.

Levels of Inference

The physical world is rich in structure that projects into images, and this richness implies a high degree of complexity among all attempts to deal with it. Of course, the antidote to complexity is organization. If the books in a library were arranged randomly, then on average half of them would have to be examined to find any particular one. But if the books were organized, for instance, according to hierarchical categories, such as those used in the Dewey decimal system, then the savings in search time could be enormous. Analogously, the knowledge in vision systems needs to be organized. More general structures—those that are likely to arise in wide classes of natural images—must be inferred first; such inferences then provide the foundation for subsequent inferences.

The Ubiquity of Orientation: Curves, Contours, and Surface Coverings

Perhaps the most ubiquitous structures in all scenes involve orientation. Contours arise from object boundaries, changes in material or lighting, and such surface markings as creases and pinstripes. Contours can be arranged into surface coverings as well; consider fur patterns or wheat fields. The local ingredient in all such patterns is orientation, and the prevalence of orientationally selective cells throughout the early visual system underlines its importance (Hubel & Wiesel, 1977). This raises the question for computational types: How can orientation information be computed? This question provides the launching point for our studies.

FEATURE DETECTORS: THE PHYSIOLOGICALLY MOTIVATED APPROACH

First, I will concentrate on the problem of inferring curves from images. The standard approach to the problem has its roots in the feature detection paradigm inspired by results from electrophysiology (Lettinger, Maturana, McCulloch, & Pitts, 1959). Stated in its simplest form, this paradigm maintains that templates for particular pattern features exist (perhaps as receptive fields) and that these templates can be matched against image information.

The matching mechanisms consist, in turn, of two distinct steps: measurement and selection. If the measurement step is linear, then it amounts to a convolution (or, roughly speaking, a computation indicating the correlation between a symmetric feature template and a piece of the image). Reichardt (1961) and Uttal (1975) have emphasized such procedures; however, the selection step that follows is inherently nonlinear, since it is designed to choose the templates that match the image best (i.e., those with the strongest convolution values).

Line Detectors and Orientationally Selective Neurons

The feature detection paradigm is particularly attractive for orientation selection, given the prevalence and arrangement of orientationally selective cells in the early visual cortex. In particular, leaving many of the details aside, the receptive fields of so-called *simple cells* (Hubel & Wiesel, 1977) seem ideally suited to function as line templates. They can be modeled by a Gaussian envelope in their preferred direction and a difference of Gaussians in the orthogonal direction (Wilson & Gelb, 1984; see Figure 1). Furthermore, such cells are organized (in cortical area V1 and elsewhere) as a system of *columns*, which suggests that each position in the image is covered by a cell "with" each orientation. Such an architecture seems to reveal the algorithm: try each orientation at every position, and select those orientations with significant matches. If the simple cells carry the convolutions, then

interactions between simple cells can be postulated to perform the significance test, for example, by maximum selection (Blakemore, Carpenter, & Georgeson, 1970). Unfortunately, the elegance of this scheme notwithstanding, in this form it cannot work in general.

The Computational Failure of the Line Detector Paradigm

Since the line detector paradigm is specified so well, it can be verified by computer. Computational experiments show that, if the curves are perfectly straight lines, and if the lines are of high contrast and far apart, the line detector scheme will work. But they also show that if the lines change direction (i.e., if they curve) and are close together (with respect to receptive field size), the scheme will fail (see Figure 2). Similarly, it will fail to distinguish between smooth curves and corners, and it will fail for textured patterns consisting primarily of oriented information (Zucker, 1982). Clearly, there is more to curve detection than convolutions and local maxima selection.

A more detailed examination of Figure 2 reveals the problem. Although the convolutions carry the information about orientation, they do not do so in a way that makes its subtleties explicit. The information is spread out around the contour and across different orientations at each position. Distinctions between different curves are blurred. The selection process must be sufficiently powerful to deal with these subtleties.

A COMPUTATIONAL ANALYSIS OF ORIENTATION SELECTION

A computational analysis of a problem implies more than simple verification; it consists in design as well. I have shown above how theories can be verified using computers; I will now show how ideas from computation can influence theory formation. In particular, before considering *how* something can be computed, an analysis of the problem can suggest *what* must be computed. I therefore begin with the question: What is a curve?

Traces, Tangents, Curvature, and Curves

A (parameterized, differentiable) curve is a differentiable map from an (open) interval to points on the image plane. The *trace* of a curve consists in the set of points through which the curve passes. In actual problems only a discrete sampling of the trace is given, but it is usually embedded in other image structures. Somehow the curve must be inferred from this noisy, incomplete, and embedded trace, while the trace itself is segmented from other image structure. Or, to state it differently, if the exact trace were known, then curve detection would simply be a matter of interpolating the curve between the trace points. The issues would then be ones of smoothness (e.g., order of interpolating polynomials and of localizing discontinuities [Pavlidis, 1982]). For the problem of orien-

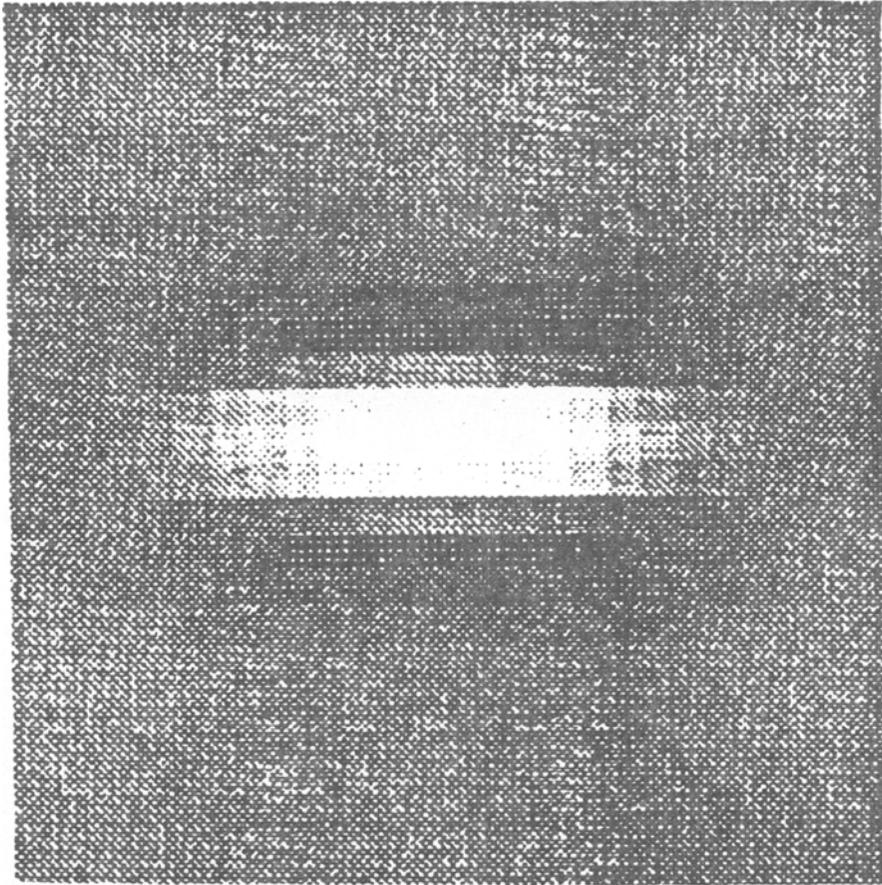


Figure 1. A model of a simple cell's receptive field. It consists of a Gaussian envelope in the preferred orientation and a difference of Gaussians in the orthogonal direction. Functionally, the Gaussian smooths over intensity variations along the contour, whereas the difference of Gaussians sharpens the spatial localization of the operator (Zucker & Hummel, 1986). Computationally, this receptive field can be thought of as an operator which, within its linear range, is convolved against image information (Rosenfeld & Kak, 1982). (Of course, the image information is first preprocessed, e.g., separated into positive and negative contrast channels.) The result of the convolution is then a linear average (L_2) estimate of the match between the local image pattern and the receptive field template; stronger convolutions indicate a better match.

tation selection, however, not only are the exact trace points unknown, but they are confused with other image contrast structure as well.

The mathematical properties of curves most relevant here come from differential geometry (do Carmo, 1976). If the mapping defining the curve were known, it could be differentiated at each point. The first derivative of the curve with respect to arc length, or a parameter that runs along the curve, is called the *tangent* to the curve. In an infinitesimal sense, then, the tangent is the best linear approximation to the curve at every point. The second derivative, appropriately normalized, is the *curvature*, or the rate of change of the tangent along the curve.

Unfortunately, neither the trace nor the mapping are given; these must be inferred during the process of curve detection. The above differential geometry, however, suggests a natural way to proceed when coupled with two lessons from the theory of computation.

The Two Stages of Curve Detection: Orientation Selection and Curve Synthesis

Recall the example of the Dewey decimal system discussed previously. As stated, the principal lesson was that if local and abstract decisions could be made, then the efficiency of searching a large data space would be improved substantially. For this principle to be applied to orientation selection, however, it needs to be modified to include a second point: in the face of uncertainty, each of the local decisions cannot be too fine grained (Zucker, in press b). Imagine, to continue the example, that someone had randomly erased some of the letters defining the categories and the books. As more noise is added, it becomes increasingly difficult to read the titles and categories, and at some critical level one would begin to use more effort in trying to use the classification system than would be incurred by a brute force search. Orientation selection, as we shall see, requires making coarse,

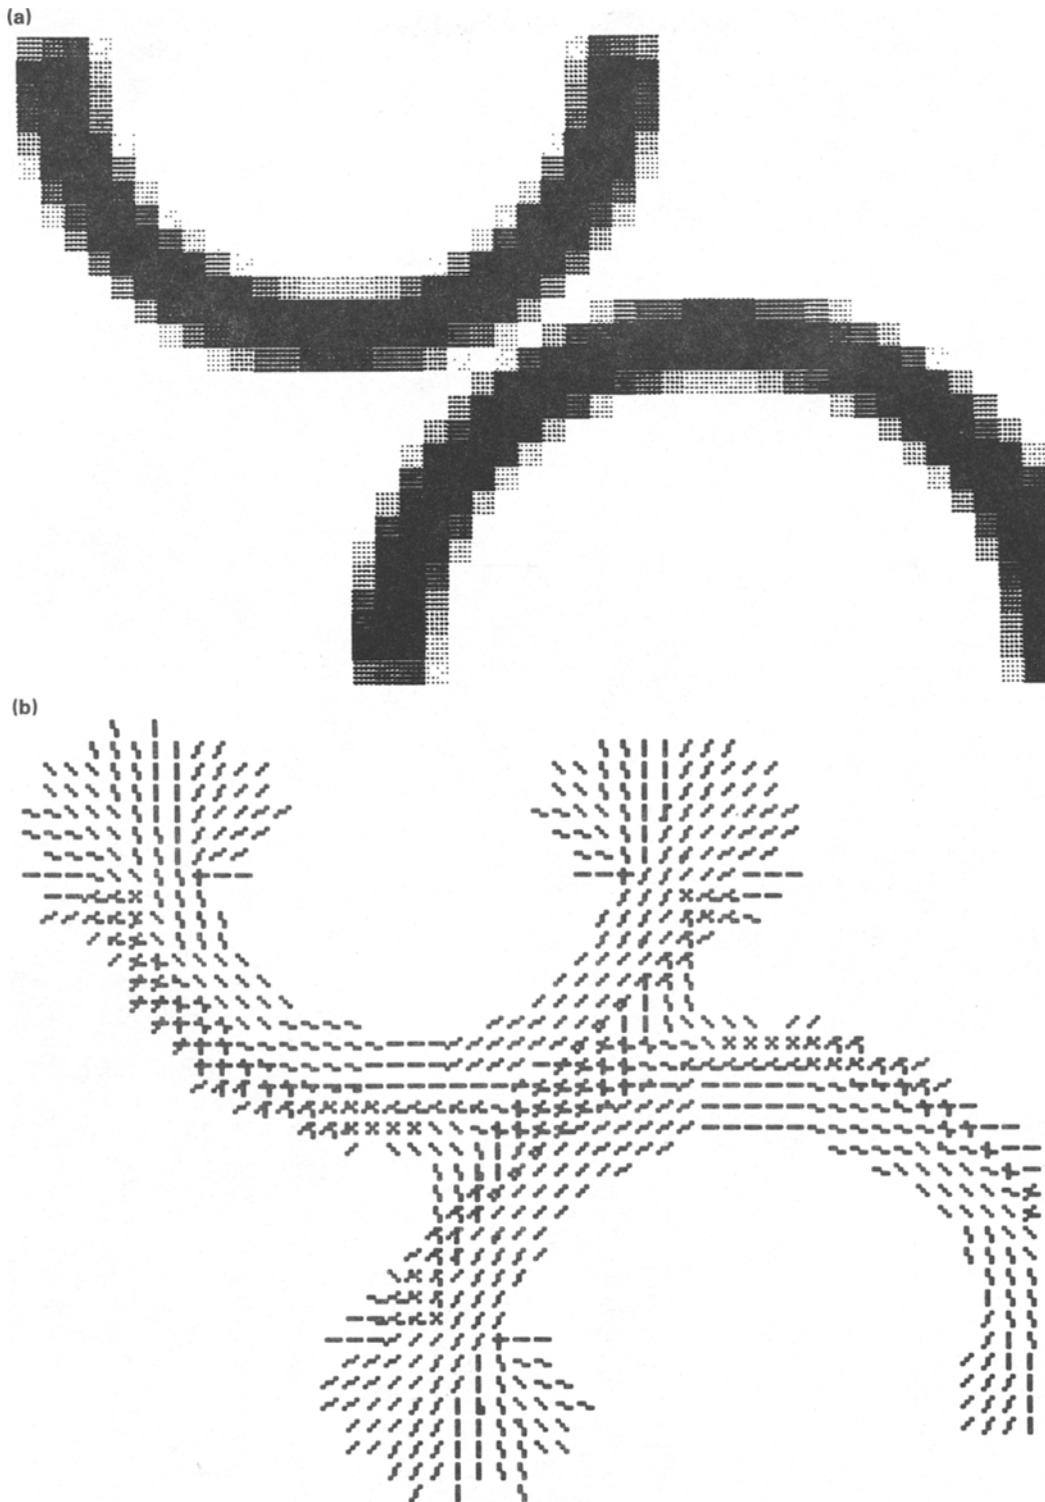


Figure 2. An example illustrating the failure of convolution followed by maximal response selection as an algorithm for orientation selection. (a) An image of two closely spaced curves. (b) The result of convolving operators, such as the one shown in Figure 1, against the image of curves above. Eight orientations were used at every position, and a short line segment indicates which of these convolutions were among the top 45% of the total responses. This thresholding algorithm is equivalent to selecting the strongest responses at each position, but eliminating those with only weak values. Observe the thick *tangent field* that is indicated, with responses flaring around the ends and between the curves. It is impossible to recover the original curves from this tangent field; thus, more sophisticated selection (or operator-interpretation) techniques are required.

local decisions first, and then following them with finer grained, more global ones.

Consider, for now, only the local/global distinction, because the differential geometry suggests such a breakdown. Recall that the tangent is the first derivative of the curve (with respect to arc length). Thus, if the tangents—the best local, linear approximations to the curve at each point of the trace—could be found, then they could readily be integrated into a global curve. Hence, we are naturally led to consider two conceptual stages for any process of curve inferencing, the first in which local information is recovered and the second in which global information is recovered. The first stage clearly amounts to orientation selection, for which the goal is the inference of a local description of the curve everywhere along it. We formalize this goal as a vector field of tangents. The second stage amounts to curve synthesis, or the formation of integrals through this vector field.

Stage 1 (orientation selection). Stage 1 estimates a discrete, quantized vector field of tangents. This is a spatial arrangement over quantized positions of unit vectors with discrete orientations. Each tangent touches the contour at exactly one position (on the quantized trace) and points approximately in the direction that the curve is going as it passes through that point. Thus, inherent in inferring the tangent field is the inference of the discrete trace of the curve.

Stage 2 (curve synthesis). Stage 2 finds integral curves through the vector field. The second stage of the model is based on the fact that the tangent is the first derivative of the contour, which suggests that contours can be recovered by a process of integration. Since the tangent field is actually quantized in space (into, say, retinotopic coordinates) and orientation (only n distinct tangent orientations are explicitly represented), the integration involves solving an approximation problem subject to inequality constraints about position, tangent, and curvature. This approximation problem will be treated formally in a separate paper.

This is precisely the breakdown that we have adopted (Parent & Zucker, 1985; Zucker, 1982). Since Stage 2 in the model is relatively unconstrained physiologically and psychophysically,¹ we shall now concentrate on Stage 1, or orientation selection. As in the full process of curve detection, orientation selection itself can be thought of as a two-step process.

Orientation Selection as a Two-Step Process: Measurement and Interpretation

The goal of Stage 1 in the model is to isolate the trace of the curve. But observe that this leads to something of a chicken-and-egg problem. If the curve were known, then its discrete trace (the set of pixels or, equivalently, the set of quantized retinotopic positions) could simply be calculated. But since it is the curve that we seek, how can the trace be calculated? The answer lies in simultaneously

inferring both the trace and local models for the curve. The tangent and the curvature define the local model at each point. The result is a structure that is referred to mathematically as a tangent field; the goal of orientation selection, I propose, is the computation of the tangent field. The simple-cell receptive fields discussed previously provide sufficient measurements; it is their interpretation that must be considered in more detail. Orientation selection, or the computation of a tangent field, can therefore be thought of as consisting of the following two steps:

Step 1. Perform measurements on a representation of the image patterns. The biological constraint to use simple-cell receptive fields as templates is very compelling. But the interpretation of these convolutions is not unique, so we must perform Step 2.

Step 2. Interpret the results of the measurements. This step is formulated both abstractly as a functional minimization problem and concretely as a cooperative network that computes solutions to the problem (Parent & Zucker, 1985). Physiologically, this amounts to interactions between the various receptive fields. Mathematically, it amounts to minimizing a functional through curvature variation [Parent & Zucker, 1985].

The first step is standard. It amounts to convolving the simple-cell receptive field operators at each explicit orientation against the image. The results are stored in what can be thought of as an *orientation column* (Hubel & Wiesel, 1977) at each position. With this given, I now concentrate on the interpretation step.

A fundamental constraint on both biological and artificial vision systems is the discrete sampling inherent in the retinotopic array. This spatial quantization is such that many different curves in the real world project into the same discrete trace. How then, given only an approximation to the discrete trace or a sampling of the points through which the curve passes, can the correct one be selected? In general it cannot; rather, the best that one can hope for is consistency among all the choices. The most natural choice is the smoothest curve from among the possible ones, and we use smoothness constraints to obtain estimates of how the curve is behaving in the neighborhood of each possible trace point. The estimates model the tangent and curvature locally, and smoothness dictates limits on their differences (Parent & Zucker, 1985).

It is important to note that the mathematics continue beyond tangents (first derivatives of the curve with respect to arc length), curvatures (second derivatives), change in curvature (third derivative), and so on. But it turns out that considerations through curvature are sufficient for situations such as the one in Figure 2, in which curvature information can be used to separate the two curves from one another by eliminating the (incorrect) responses between them.

The key to estimating the tangents and curvatures is to observe that they are not independent; for the proper parameterizations, the tangent is the first derivative of the

curve with respect to arc length, and curvature is the second derivative. The tangent must cover nearby trace points, and curvature estimates must cover the underlying tangents.² But these relationships are differential; on the discrete grid they can only hold approximately. This suggests the maximization of a functional: one would naturally like the above relationships to hold as much as possible. Fortunately, computational theory again suggests how this can be accomplished.

Distributed, parallel algorithms for functional maximization. Given the need to maximize certain functionals, which, in this case, can be specified in terms of the (estimated) trace, the (estimated) tangent, and the (estimated) curvatures, one could appeal to the existing literature to determine how this might be accomplished (Leuenberger, 1971); however, other constraints exist that are necessary for any such scheme to be biologically plausible. Neurons in the early visual system are not connected densely, but rather form local clusters of distributed networks (Gilbert & Wiesel, 1981). Parallelism is everywhere, over spatial distributions and across functions. Somehow, the algorithms for functional minimization must capture this distributed, parallel structure.

We have developed the mathematical theory for one class of such algorithms, which we call relaxation labeling processes (Hummel & Zucker, 1983; see also Ballard, Hinton, & Sejnowski, 1983, for related references). One set of specific interconnections that estimate the trace, tangent, and (coarse) curvatures so that they are mutually consistent (i.e., they maximize a certain natural functional) has been developed by Parent and Zucker (1985). The result is a cooperative network of interactions that runs between the initial convolution values, iteratively changing them so that they become more consistent with current curvature estimates. That it works successfully is exhibited in Figure 3.

The dynamics of the network intuitively run as follows. The initial convolutions are given, say as response values in orientation columns. Take these values as initial estimates of which tangents are actually indicated in the image, with large responses representing higher certainty. Now, initial estimates of curvature can be calculated from these initial tangent estimates by integrating information over a slightly larger neighborhood. Once curvature estimates are available, the underlying tangent estimates that they cover can be adjusted (recall that curvature varies with the change in tangent direction) to make them more consistent with each other. But now the tangent estimates have changed, so the curvatures must be recalculated. This process of re-estimation and adjustment for consistency is repeated until convergence.

There is a serious issue of how many iterations are required for convergence. My computational experience indicates that two iterations are sufficient for all applications that we have attempted. If the number of iterations is low enough so that they can be implemented by repetitions of the "wetware" in which each one is carried out, then two iterations are certainly within the limits of biological plausibility.



Figure 3. An example illustrating the success of our curvature-based orientation selection scheme. (a) A low-contrast image of a fingerprint, an image containing many closely spaced, thin curves. Observe that curvature varies from almost straight to very curved and that some of the curves stop. (b) The tangent field produced by our algorithm superimposed on the image in (a). Observe that the tangents coincide exactly with the bright curves and that none of the problems illustrated in Figure 2 are apparent here.

It is important to stress that results of computational experiments involving curves (rather than straight lines) led to the realization that curvature is so important. Other researchers (e.g., Grossberg & Mingolla, 1985) have thus far appeared to concentrate their computational experiments on (piecewise) straight contours, and therefore have not yet confronted curvature.

Endstopping and curvature estimation. It is clear from the above computational considerations that curvature plays an essential role in orientation selection. With this motivation, then, we sought ways in which curvature could be reliably calculated. Again, computational experiments—this time with our network (Parent & Zucker, 1985)—provided the key. The experiments revealed that only coarse curvature estimates are required at this step. That is, only a small number of possible curvature values is sufficient; it is not necessary to compute curvature to several decimal places (recall the argument above about noisy library searches). Current implementations use only seven possible curvature “classes”: straight; curved to the left in small, medium, and large amounts; and similarly curved to the right. Of course, the highest curvatures must respect the geometric limitations imposed by a discrete sampling grid.

The standard approach to computing curvature involves fitting splines or polynomials, and then differentiating them symbolically. But this approach is very susceptible to noise, so that small errors in positioning often lead to large errors in estimating curvature. We have developed another technique, however, that provides only coarse curvature estimates, but does so reliably (Dobbins & Zucker, 1986). It amounts to a certain rectified difference between “simple cells” at similar positions but with different receptive field characteristics. Interestingly, if one were to do the computational equivalent of electrophysiology on this part of the network, the curvature consistency computation would appear to have introduced *endstopping* onto the simple cells’ receptive fields. Endstopping is the property whereby a cell’s response increases with a line stimulus of increasing length, but only up to a point. Following this point, the response decreases. It is usually assumed that endstopping has something to do with end-point determination. Our computational perspective has suggested another alternative that is consistent with the known physiology (in particular, it could amount to layer 4 and layer 6 local circuits, mediated by an inhibitory interneuron) (Dobbins & Zucker, 1986).

PREDICTIONS FROM THE THEORY

The value of any abstract theory is, to a large extent, both in its explanatory power and in the predictions that arise from it. I have already shown that computational theories can suggest novel interpretations of physiological function. I now discuss two more predictions of the theory and some of their consequences.

Sensitivity to Discontinuities

The presence of discontinuities is one of the key features of contours. Discontinuities arise in the outlines of objects that stand in an occlusion relationship, and they often indicate abrupt changes in properties arising from different objects. They can certainly not be ignored in any theory of orientation selection, and since smoothness constraints play a key role in our computational theory, one might expect problems to arise in the neighborhood of

discontinuities. This is not the case, however, at least to the extent that the human visual system is capable of localizing discontinuities.

Imagine a pair of straight lines that meet at a corner, an inverted V shape. Now imagine rounding out this corner, say by smoothing it with the computational equivalent of sandpaper. What is the minimal amount of smoothing to which we are sensitive? Or stated in other terms, what is our sensitivity to distinguishing sharp corners from high curvatures? Clearly such a question is relevant ecologically, since corners, not high curvatures, arise from occlusions.

Computationally, one would expect the sampling quantization to determine this sensitivity, at least in part. We have examined it psychophysically, using a variant of the above idealized displays. Instead of solid lines, we have used dotted ones, and instead of rounding off the corner, we have rotated the dots in phase with respect to it. When a dot is placed right on the corner, of course, the display appears to be two (dotted) lines meeting at a corner (Figure 4). As the top dot is moved from the corner, it remains sharp until a certain displacement, when it suddenly looks rounded (Figure 4, middle); this smooth appearance persists for further displacements (until the next dot approaches the corner) (Link & Zucker, in press a). Our algorithm exhibits precisely this sensitivity to discontinuities (Figure 4).

Curves and Surface Coverings: One- and Two-Dimensional Oriented Patterns

The success of the theory thus far might lead one to believe that it will work on all patterns rich in oriented structure. In a preceding section, *The Ubiquity of Orientation: Curves, Contours, and Surface Coverings*, I outlined several ways in which patterns exhibiting orientation would arise: bounding contours of objects, surface creases, and surface coverings. Since curves of the sort that we have already discussed cover many examples of the first two possibilities, let us now consider surface coverings. Examples of such patterns might include grassy fields, fur, and hair coverings.

Observe that, by their very composition, surface coverings differ fundamentally from individual curves; they consist in very large numbers of curves (e.g., hairs), all of which are roughly parallel to each other, but which go in and out of occlusion relationships at close intervals and frequently. Thus, they are topologically two-dimensional patterns, as opposed to one-dimensional curves. Or to state it intuitively, one can only walk along a curve in the forward or backward directions; a sideways step takes one off the curve. But for two-dimensional surface coverings, one can walk both along the curves (hairs) and across them (from hair to hair). Surface coverings must reflect the dimensionality of the surfaces that they cover.

These topological differences are born out computationally. If we run our orientation selection network on an image of a fur pattern, the result is isolated short tangents associated primarily with the highlights (Figure 5a). Since the hairs interweave so much, the coherent image inten-

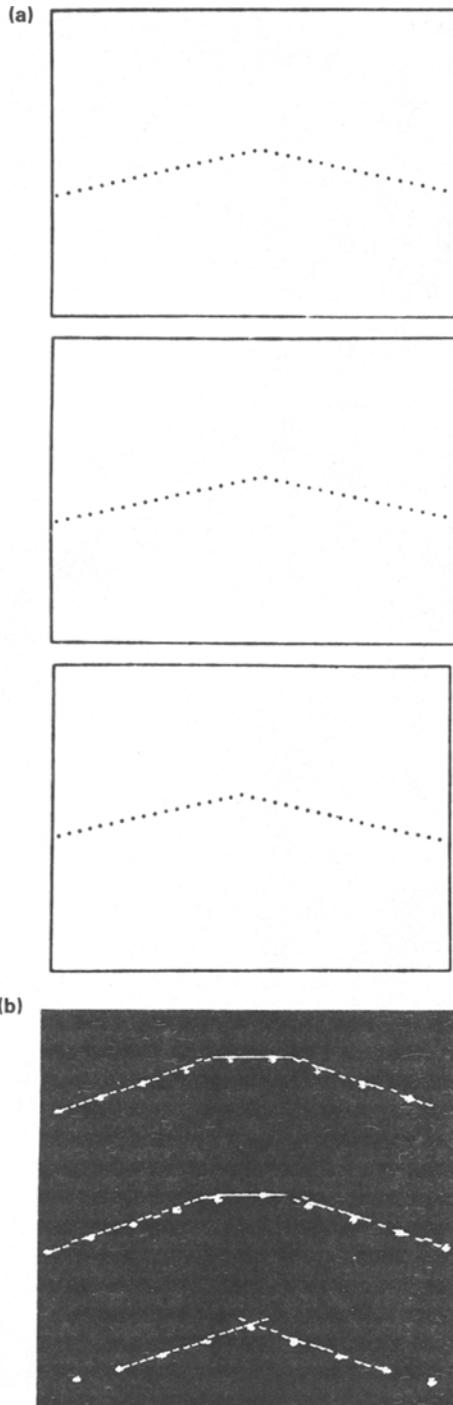


Figure 4. An illustration of the kind of psychophysical prediction that emerges from our computational modeling. (a) A collection of three dotted lines meeting abruptly at a corner. In the bottom display a dot is placed right on the corner, whereas in the middle and top illustrations the dot is displaced increasingly farther from the corner. The middle displacement corresponds to the position at which subjects stated the corner appeared smooth (rather than sharp) in Link and Zucker's (in press a) experiments. (b) The result of our orientation selection algorithm. Observe that the lower curve is discontinuous, representing a sharp corner (with multiple tangents represented there), whereas the other two are smooth.

sity structures necessary for localizing the hairs (contours) are missing. In fact, the hairs themselves cannot even be imaged, because they are under other hairs. What is required is an interpolation process that "spreads" the orientation information across contours—that is, in a direction perpendicular to the local orientation—as well as along them. Such a process is described in Zucker (1985); that it works is shown in Figure 5b.

The consequences of such a lateral spreading of orientation information would be expected to be most evident psychophysically in the neighborhood of orientation discontinuities, and this is in fact the case (Link & Zucker, in press b). One might also ask whether the physiological substrate could support such lateral spreading as well, and the evidence is beginning to emerge that certain complex cells exhibit processes with precisely this structure (T'so, Gilbert, & Wiesel, 1986). Like endstopping, however, the functional role for these lateral excitatory processes has been elusive, and again computational theorizing may have led to a viable explanation.

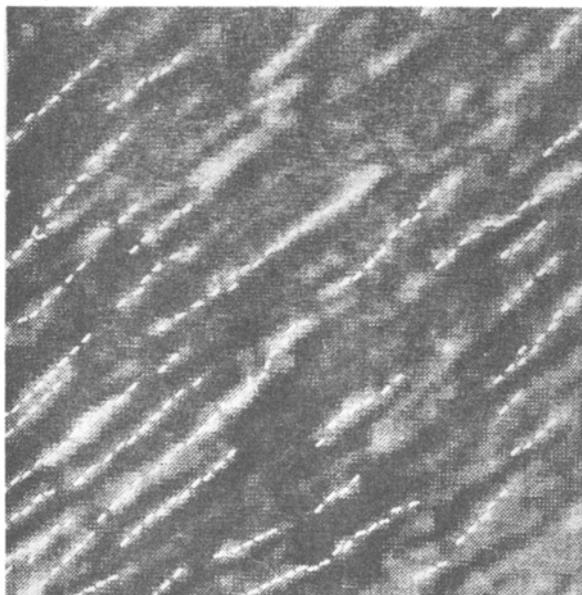
CONCLUSION

There is always a sense in which scientific disciplines are influenced by the current zeitgeist, and vision is no exception. Computers have greatly extended both the quantity and quality of experimental psychophysics and physiology, and computation has provided a theoretical perspective that has, up until now, largely been lacking. Perhaps if computation had been available 100 years ago, the great debates between Mach and von Helmholtz over neural networks and unconscious inference could have been far more productive. Computation provides the connection between problems and solutions, and between algorithms and implementations.

In this paper, I attempted to highlight the role of computers and computation in a much more focused problem domain, that of orientation selection. This is the process by which oriented entities are inferred from image structures. Computer implementations were used to demonstrate the weaknesses of standard approaches, and thereby to illuminate the more subtle aspects of the problem of orientation selection. The result forced us to analyze the problem more abstractly, and led to the introduction of techniques from differential geometry. The seemingly straightforward question of line detection evolved into the inference of the trace, tangent, and curvature of arbitrary curves. Basic constraints from physiology and psychophysics influenced the resulting theory, especially the shape of receptive fields and the regular interconnection arrangement suggested by neurons. The result led both to new computational ideas for maximizing functionals and to an algorithm for orientation selection. Finally, to complete the loop, the algorithm led to new psychophysical results (regarding sensitivity to discontinuities) and to new functional interpretations of physiology (endstopping and curvature estimation, and complex cells and oriented textures).

Orientation selection was chosen because of its prominence in early vision and because of the many different

(a)



(b)

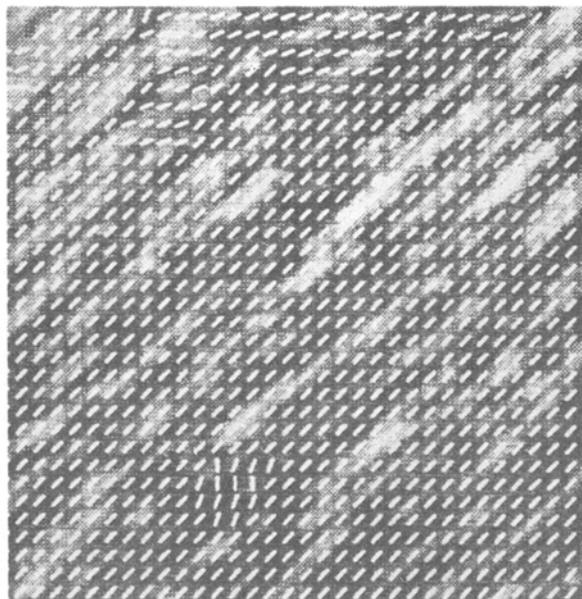


Figure 5. An illustration of the difference between one-dimensional curves and two-dimensional surface coverings. We show in particular that surface coverings, such as hair patterns, require lateral interpolation, or processing *across* curves, in addition to processing *along* curves. (a) The result of running the orientation selection algorithm for curves on an image of a hair pattern. Observe that the tangents are indicated essentially where the highlights are strong. But they are scattered sparsely and do not capture the flow of the hair. (b) The tangent field obtained when lateral interpolation is included in the orientation selection process. Observe now that the tangent field is dense in two directions (actually, only every sixth tangent is shown, so that the tangents' lengths can properly indicate orientation). It may be the case that some subset of complex cells perform this type of lateral interpolation via excitatory interconnections between neurons with similarly oriented receptive fields.

constraints from physiology and psychophysics that could be brought to bear on it. But what appeared to be a rather simple single process turned out to be a network of several rather different ones. These differences emerged between tangents and curvatures, and between curves and textures. As Uttal (1981) points out, the Gestalt laws have been notoriously difficult to formalize; the previously unrecognized differences mentioned above may be responsible, at least in part (Zucker, in press a). Moreover, indications are that a similar theory will underlie early optical flow computations (in which the two spatial dimensions are generalized to three spatiotemporal dimensions) (Zucker & Iverson, in press). Whether or not these ideas will stand the test of time, they certainly illustrate how computers and computation have begun to influence research in vision.

REFERENCES

- BALLARD, D., & BROWN, C. (1982). *Computer vision*. Englewood Cliffs, NJ: Prentice Hall.
- BALLARD, D., HINTON, G., & SEJNOWSKI, T. (1983). Parallel visual computation. *Nature*, **306**, 21-26.
- BLAKEMORE, C., CARPENTER, R., & GEORGESON, M. (1970). Lateral inhibition between orientation detectors in the human visual system. *Nature*, **228**, 37-39.
- DOBBINS, A., & ZUCKER, S. W. (1986). *The hypercomplex neuron: Has vision been thrown a curve?* (Report No. TR 86-15). Computer Vision and Robotics Laboratory, McGill University, Montreal.
- DO CARMO, M. (1976). *Differential geometry of curves and surfaces*. Englewood Cliffs, NJ: Prentice Hall.
- GILBERT, C., & WIESEL, T. (1981). Laminar specialization and intracortical connections in cat primary visual cortex. In F. Schmidt, F. Worden, G. Adelman, & S. Dennis (Eds.), *The organization of the cerebral cortex* (pp. 163-191). Cambridge: MIT Press.
- GROSSBERG, S., & MINGOLLA, E. (1985). Neural dynamics of perceptual grouping: Boundary completion, illusory figures, and neon color spreading. *Psychological Review*, **92**, 173-211.
- HUBEL, D., & WIESEL, T. (1977). Functional architecture of macaque monkey visual cortex. *Proceedings of the Royal Society (London)*, **B198**, 1-59.
- HUMMEL, R. A., & ZUCKER, S. W. (1983). On the foundations of relaxation labeling processes. *IEEE Transactions: Pattern Analysis & Machine Intelligence*, **PAMI-5**, 267-287.
- LETTVIN, J., MATURANA, H., MCCULLOCH, W., & PITTS, W. (1959). What the frog's eye tells the frog's brain. *Proceedings of the Institute of Radio Engineers*, **47**, 1940-1951.
- LEUENBERGER, D. (1971). *Optimization by vector space methods*. New York: Wiley.
- LEVINE, M. (1985). *Vision in man and machine*. Englewood Cliffs, NJ: Prentice Hall.
- LINK, N., & ZUCKER, S. W. (in press a). Corner detection in curvilinear dot grouping. *Biological Cybernetics*.
- LINK, N., & ZUCKER, S. W. (in press b). Sensitivity to corners in flow patterns. *Spatial Vision*.
- MARR, D. (1982). *Vision*. New York: Freeman.
- PARENT, P., & ZUCKER, S. (1985). *Trace inference, curvature consistency, and curve detection* (Report No. TR 85-12). Computer Vision and Robotics Laboratory, McGill University, Montreal.
- PAVLIDIS, T. (1982). *Algorithms for graphics and image processing*. Rockville, MD: Computer Science Press.
- REICHARDT, W. (1961). Autocorrelation, a principle for the evaluation of sensory information by the central nervous system. In W. Rosenblith (Ed.), *Sensory communication* (pp. 303-317). New York: Wiley.
- ROSENFELD, A., & KAK, A. (1982). *Digital picture processing*. New York: Academic Press.
- T'SO, D., GILBERT, C., & WIESEL, T. (1986). Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *Journal of Neuroscience*, **6**, 1160-1170.
- UTTAL, W. (1975). *An autocorrelation theory of form detection*. Hillsdale, NJ: Erlbaum.
- UTTAL, W. (1981). *A taxonomy of visual processes*. Hillsdale, NJ: Erlbaum.
- WILSON, H., & GELB, D. (1984). Modified line-element theory for spatial-frequency and width discrimination. *Journal of the Optical Society of America*, **A1**, 124-131.
- ZUCKER, S. W. (1981). Computer vision and human perception: An essay on the discovery of constraints. *Proceedings of the Seventh International Joint Conference on Artificial Intelligence* (pp. 1102-1116). Vancouver.
- ZUCKER, S. (1982). *Early orientation selection and grouping: Evidence for type I and type II processes* (Tech. Rep. No. 82-8). Department of Electrical Engineering, McGill University, Montreal.
- ZUCKER, S. W. (1985). Early orientation selection: Tangent fields and the dimensionality of their support. *Computer Vision, Graphics, & Image Processing*, **32**, 74-103.
- ZUCKER, S. W. (in press a). The diversity of perceptual grouping. In M. Arbib & A. Hanson (Eds.), *Vision, brain, and cooperative computation*. Cambridge: MIT Press.
- ZUCKER, S. W. (in press b). Early vision. In S. Shapiro (Ed.), *The encyclopedia of artificial intelligence*. New York: Wiley.
- ZUCKER, S. W., & HUMMEL, R. (1986). Receptive fields and the representation of visual information. *Human neurobiology*, **5**, 121-128.
- ZUCKER, S. W., & IVERSON, L. (in press). From orientation selection to optical flow. *Computer Vision, Graphics, & Image Processing*.

NOTES

1. Stage 2 can be readily formulated in mathematical terms at this point, but such a formulation is outside the scope of this paper.
2. Although curvature was defined in terms of derivatives, it can also be thought of in terms of "change in tangent" along the curve.