VOLUME 4 NUMBER 1 1989 ISSN 0169-1015

SPATIAL VISION

AN INTERNATIONAL JOURNAL OF PSYCHOPHYSICAL, PERCEPTUAL, AND COGNITIVE RESEARCH ON THE VISUAL PROCESSING OF SPATIAL INFORMATION

Editors-in-chief: D. H. Foster and P. C. Dodwell

UTRECHT, THE NETHERLANDS
TOKYO, JAPAN

///VSP///

Spatial Vision, Vol. 4, Nos 2/3, pp. 131–139 (1989)

Texture fields and texture flows: Sensitivity to differences

YAACOV HEL OR* and STEVEN W. ZUCKER†
Computer Vision and Robotics Laboratory, McGill Research Centre for Intelligent Machines, McGill University, Montréal, Quebec, Canada

Received for publication 1 July 1989

Abstract—There are two large classes of textures, those with an overall orientation structure (texture flows) and those without (texture fields). We investigate human sensitivity to detecting a patch of texture field within a texture flow psychophysically by using random not Moiré patterns. The resultant sensitivity, as a function of patch-size and path-length, is then related to a computational model of orientation selection, which reveals a connection between texture structure and the estimation of curvature. Finally, the connection back to curvature is confirmed by demonstrating a similarity between the patch sensitivity data and previous data on sensitivity to corners in flow patterns.

1. INTRODUCTION

Texture is a generic term for images of surface coverings, and fields of texture naturally fall into two classes: (i) those fields which consist of unit elements with no inherent orientation, such as sand or salt-and-pepper textures and (ii) those fields which consist of unit elements with a distinct orientation, such as hair, fur and grass textures. We shall refer to the former class (i) generically as (static) texture fields and the latter (ii) as texture flows. The rationale for these names is that a texture flow can be created by extending (or growing) the initial point elements in a texture field into the oriented elements comprising a texture flow. The question with which we shall be concerned is how much "growth" or extension is necessary before a texture field becomes a texture flow; or, stated differently, precisely how sensitive are we at detecting differences between texture fields and texture flows.

The experimental paradigm for studying these differences will be to create images of texture flows, but with holes in them filled with (unoriented) texture fields; we shall study the detectability of the holes as a function of the amount of structure in the texture flow. Sensitivity to such differences between textures is important to vision because the global orientation structure (or flow) of the patterns provides evidence about the structure of the underlying surface. In particular, an abrupt change in a flow pattern provides evidence for an abrupt change in surface property, as when one object occludes another or when two surfaces meet at an edge. In a previous study we examined sensitivity to changes in orientation ('corners') within flow patterns (Link and Zucker, 1987); this previous study quantified our sensitivity to the kinds of edges that arise between two surfaces covered with oriented texture flows. The current study extends these results to edges that arise between surfaces with different

---

*Department of Computer Science, The Hebrew University, Jerusalem, Israel.
†Senior Fellow, Canadian Institute for Advanced Research.

kinds of coverings, in particular, one with a texture flow, and the other with a (non-oriented) texture field.

The results of these experiments are also relevant to understanding the mechanisms by which descriptions of such textures could be inferrred biologically. Observe that, within a region of texture flow, the image cues derive from highlights, shadows and differential coloration cast by the individual texture elements (hairs, say). However, since these elements are so densely arranged, the photometric relationships between them are immensely complex: hairs pass into and out of occlusion relations extremely often, and they only occasionally align photometrically with the light sources and viewer. The resulting image structures—namely, the cast highlights and differential shading—are thus relatively sparse, and must be 'filled in' as part of whatever mechanism is responsible for inferring texture structure (Zucker, 1986). Such interpolation of structure implies averaging, which raises the question of when discontinuities in the flow are smoothed over. But smoothing over structure is just another way of stating that the texture process is insensitive to the change, and misses it. Our key finding here is that, analogously to detecting changes in orientation within a flow pattern (Link and Zucker, 1987), our sensitivity to changes between an oriented texture flow and an unoriented texture field takes a quantum jump when sufficient structure is available within the field to support estimates of curvature. Such is the prediction that comes directly out of a biologically-motivated approach to orientation selection being developed in our laboratory (Zucker, 1986; Dobbins *et al.*, 1987; Zucker *et al.*, 1988; Parent and Zucker, 1989), and it is elaborated more fully later.

The paper is organized as follows. In the next section, we introduce random dot Moiré patterns (RDMPs), or Glass patterns (Glass, 1969; Glass and Perez, 1973; Glass and Switkes, 1976), to model texture flow, and then describe the experiment in detail. Finally, the results are interpreted in the context of a model for orientation selection within flow fields. This mixture of psychophysics and computational modeling is very much in the tradition of K. Pradzny, who also worked extensively toward understanding the mechanisms underlying the perception of RDMPs (Pradzny, 1984; 1986a, b), and lends further support to the idea that cooperative mechanisms, rather than symbolic ones, provide the most feasible explanation.

2. TEXTURE FLOW: RDMPs

The key point about texture flows is that they correspond to patterns composed of dense arrangements of oriented segments distributed over a region. RDMPs are a class of patterns which provide a kind of mathematical abstraction of them. RDMPs are constructed according to the following algorithm:

OVERLAY 1 (Initialization):
Construct a field of dots distributed randomly:

OVERLAY n (Repeat $n-1$ times):
(i) Make a copy of overlay $n-1$.
(ii) Move each dot in the copy according to a chosen flow transformation (e.g., a rotation or a translation).
(iii) Superimpose this overlay on the other overlays.

The variable n in this algorithm, which we define as the *path length*, plays a key role in controlling structure: for larger values of n each little 'streak' in the pattern becomes longer and better defined. Observe that only two overlays ($n=2$) are required to produce the impression of flow, for now pairs of dots define the oriented elements, although larger values produce richer patterns. Part of the reason for this, we shall argue, is that at path length $n=3$ it becomes possible for the texture mechanism to estimate curvature (as well as orientation), and that this accounts for our most significant psychophysical result.

3. EXPERIMENT: DETECTION OF ANOMALOUS TEXTURE FIELDS IN RDMPs

The purpose of this experiment was to determine how sensitive human observers are to patches of (non-oriented) texture within larger regions of texture flow. The texture flow was modeled as translational RDMPs, and circular windows were placed within the flow field. The subject's task was to detect the presence of the dissimilar window as a function of the path length of the flow (the amount of structure within the flow field) and the size of the window. Sensitivity was then quantified as the size of the detectable hole as a function of path length.

3.1. Subjects

There were four subjects, three male and one female, all with normal or corrected vision. All subjects were associated with our laboratory, and were aware of the purpose of the experiment.

3.2. Apparatus

The stimulus images were generated on a Sun-3 workstation and displayed on its bit-mapped monochrome (black and white) monitor. The screen resolution is 1152 by 900 pixels, and the pixels are slightly elongated (see below). The subjects were seated comfortably about 55 cm in front of the monitor in a dimly lit room.

3.3. Stimuli

3.3.1. General. The stimuli consisted of RDMPs that were generated by translation-transformed overlays (see Section 2, step (ii)). A fixation mark was placed at the center of the flow field, and a circular test patch was generated at a random position. Within the circular test patch there was a (non-oriented) random dot texture field with the same density as in the RDMP, thus precluding the possibility of discrimination based on local contrast differences. Examples of the test stimuli are shown in Fig. 1.

To control for retinal eccentricity, the test patches were located on a circle surrounding the fixation mark, with the distance between the edge of the patch and the fixation mark fixed to 2°. The position of the patch was chosen randomly from 0° to 360°.

3.3.2. Sizes, domains and extents. The patterns consisted of 512 by 512 pixel arrays, in which each pixel represents either a dot (black) or a space (white) in the RDMP. One pixel width (x direction) was 1.9′ of visual angle and its height (y direction) was 2.0′. Δx and Δy, the displacement between corresponding dots in adjacent overlays, were chosen randomly from -4 to 4 pixel-units under the restriction that the

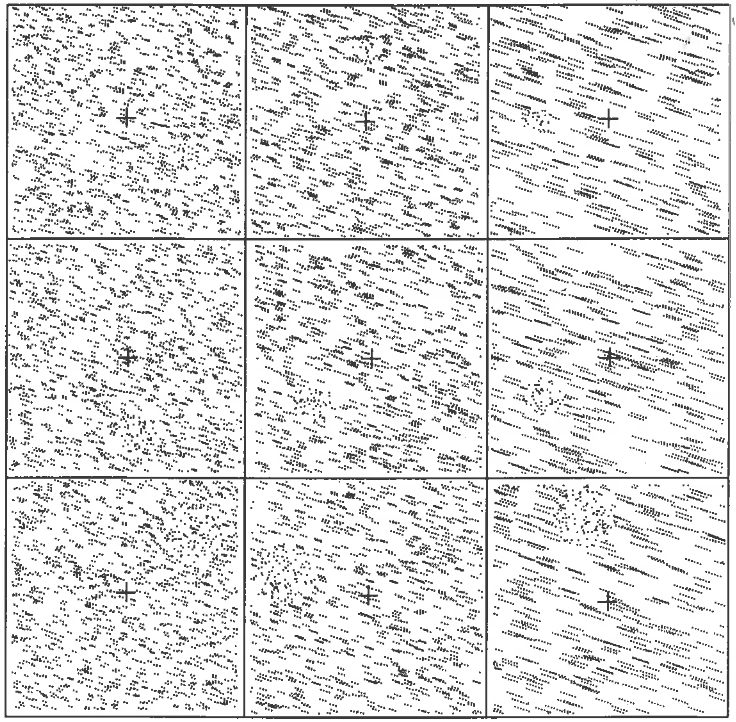

Figure 1. Examples of the test stimuli for the experiment. Nine patterns are shown, with the three columns (shown from left to right) indicating path-length 2, 3 and 6, respectively, and the three rows (from top to bottom) showing patch sizes of 14, 20 and 32 pixels, respectively. The patch is located randomly around a circle, as described in the text. Observe how the patch is most difficult to detect in the upper left image and most visible in the lower right.

dot:space ratio was in the range of 1:2 to 1:4 (Zucker and Davis, 1988). The densities of the dots used were 2, 4, 6, 8 and 10%, and the path lengths were 2, 3, 4, 6 and 8. The patch sizes used in the test pattern ranged over 8, 14, 20, 26 and 32 pixel-units (from 0.5° to 2.1°).

3.4. Procedure

The experiment was designed as a two-alternative, forced choice task. It consisted of three sessions, during which each subject viewed 750 patterns (250 for each session) obtained from: five values of path length, five quantizations of density, six sizes of path (including those patterns without patch, or where the patch size was 0) and five instances of each set of parameters. The patterns were displayed in a random order, constant across subjects, for a display duration of 0.5 s. The subject's task was to indicate the presence or absence of a patch inside the pattern.

During the inter-stimulus interval between patterns, the screen displayed a random dot pattern with a cross in the center on which the subject was asked to fixate before presenting the next stimulus.

Prior to the actual experiment, each subject was briefed on the purpose of the experiment and was put through an initial training set. The training set patterns were

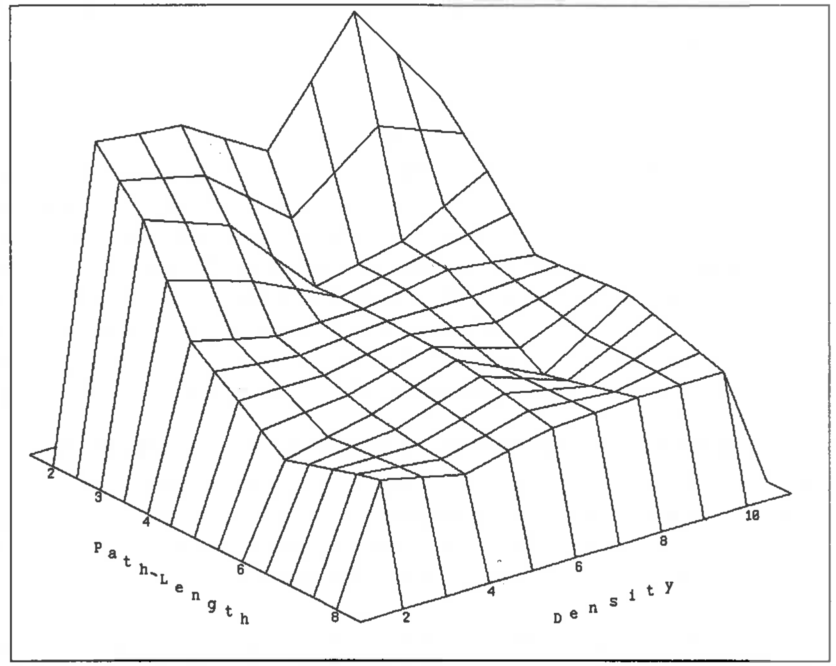

Figure 2. Plot of sensitivity (patch-size at the 50% correct level) as a function of path-length and density for a typical subject. Observe that there is little variation across the densities studied, and that the most rapid changes in sensitivity occur around path-length 3.

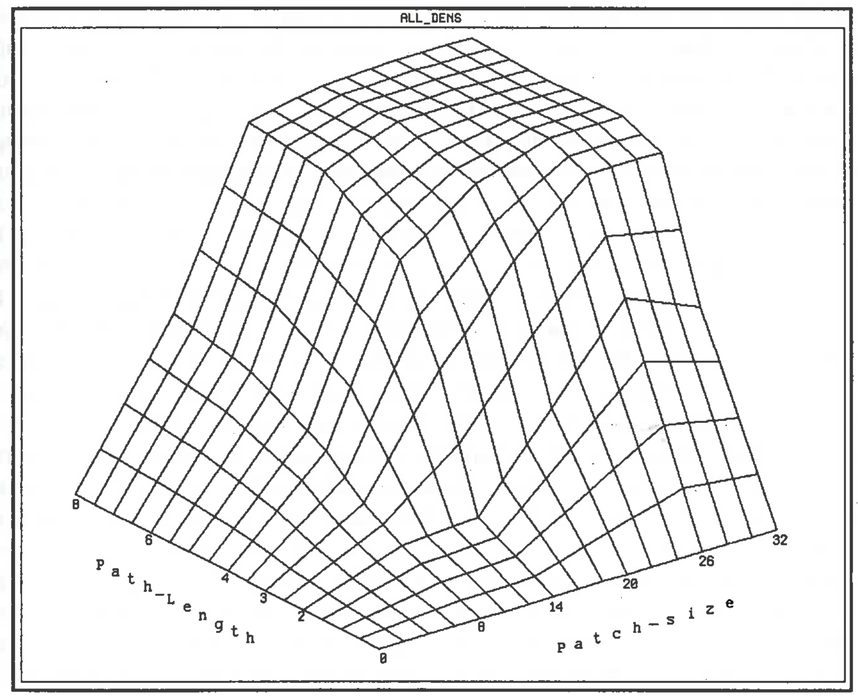

Figure 3. Plot of the percentage correct answers as a function of path-length and patch-size (in pixels). Now observe the increase in sensitivity for patches larger than about 10 pixels.

selected from the set of stimulus patterns, and were first flashed for 0.5 s. Then, following a brief interval, the training pattern was displayed again for the subject to analyse.

4. RESULTS

The results of the experiment for a typical subject are shown in Fig. 2, with sensitivity graphed as the patch size required for 50% correct responses. Clearly the most significant changes in sensitivity occurred at path-length 3, with little variation across density. This suggests that it is reasonable to average across density, and this is shown in Fig. 3 as the percentage of correct answers as a function of path-length and patch size (averaged across subjects). It can be observed that the sensitivity as a function of patch size increases after about 10 pixels, which corresponds to a patch of about 20 min at 2 deg eccentricity. In general, the standard deviation between subjects decreased as a function of path-length (standard error bars are shown in Fig. 5).

5. DISCUSSION: ORIENTATION SELECTION, TEXTURE FLOW AND CURVATURE

Since local orientation is an integral part of the description of a texture flow, a process of orientation selection should be an integral part of the mechanism for analyzing texture flows. We concentrate here on a computational view, which holds that orientations election consists of two (conceptually) distinct steps:

Step 1: Initial measurements performed locally by operators functionally analogous to receptive fields (that are orientationally-selective and endstopped): and

Step 2: a refinement of these initial measurements into sharper hypotheses about the orientation (and curvature, for reasons explained below) at each point in the texture flow.

The above stages are strongly influenced by visual physiology, with Step 1 consisting in the initial responses of neurons with orientationally-selective receptive fields arranged in 'hypercolumns'. Given the broad tuning of such cells, it follows that there will be a wide dispersion of responses along each column, a dispersion that will smooth over variations (and specifically the boundaries of texture patches) necessitating a second, refinement step. This refinement process can be conceptualized as building up a local model for the flow at each point based on the initial measurements at that point together with those measurements taken in a neighborhood surrounding the point. The natural form for such a model is differential, with the pointwise orientation-selective measurements carrying information about tangents, and the neighborhoods carrying information about curvature (or, equivalent, relations over tangents at nearby positions). In physiological terms, the initial estimates of curvature can be provided by endstopped cells (Dobbins *et al.*, 1987) and the refinement by inter-columnar interactions (Zucker, 1986; Parent and Zucker, 1989).

The refinement stage proceeds iteratively, conceptually increasing the firing rate of those neurons representing (orientation, curvature) hypotheses consistent (through curvature) with the responses in their neighborhood, and decreasing those that are inconsistent. Such a process amounts, in a sense, to the propagation of information along (virtual) curves, or along the local differential models referred to above, in the direction indicated by the orientation responses. However, for texture flows, information must be propagated in the orthogonal direction as well. This is because the image evidence for flows is intermittent and carried mainly by highlights (or, in the case of our laboratory examples, by the dotted arcs comprising the RDMPs); some process of interpolation or spreading of orientation information across (in addition to along) the flow must be included. While there are different ways to model this interpolation, they share the property that positional (or phase) sensitivity will decrease while orientation sensitivity will remain high, a property that is definitive of complex cells (Zucker, 1986). We thus conclude that the elements of our model have their physiological support in complex, endstopped cells. The end result, in terms of our texture patterns, would then be consistently high activity in the flow regions, and low activity in the patches.

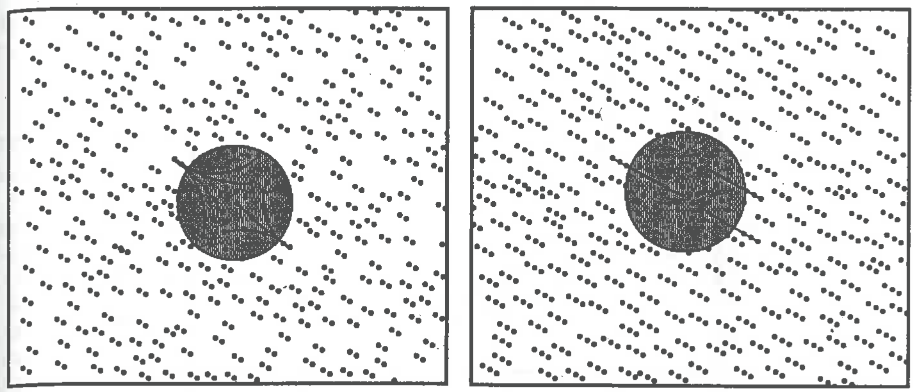

Figure 4. Illustration of the role of curvature hypotheses in constraining orientation estimates near the border of the noise patch. (left) A texture flow with a patch-length of 2, and little curvature constraint. Many different smooth continuations are possible within the patch. (right) When the path-length is increased to 3, it becomes possible to estimate curvature, an hypothesis under which many fewer consistent continuations are possible, more rapidly driving the initially scattered orientation responses to nil and increasing sensitivity to detecting the patch.

The problem is at the border of the patches. Since the initial responses are ambiguous here, additional information is required to separate local patches of consistent ones from the inconsistent 'noise' responses within the patch. The estimates of curvature provide precisely this sort of information, and without it the visual system has little choice other than to smooth over the discontinuity.

The role of curvature is illustrated schematically in Fig. 4. In the left panel, a flow field with path-length 2 is shown, for which little direct information about curvature is available; that is, the dot pairs indicate tangent orientation only. In fact, for our model in this situation the initial endstopped estimates of curvature will largely be random; it is only after a few iterations of the refinement step that any curvature constraint will arise. The result, then, is that there is a wide range of possible smooth continuations of each of the tangents near the border of the patch, some of which are bound to agree with the random responses over the patch. The result is that, for small sizes, the patch will be entirely smoothed over (recall Fig. 3). However, if a tighter curvature constraint is available, as in the right panel of Fig. 4, then few (if any) of the noise responses will be consistent, and the patch will be much more noticeable. In large patches, of course, the inconsistent responses cancel each other out, since there are no curvature responses consistent with them. This is exactly what was found experimentally.

Since we are claiming that patch detectability depends on curvature, or more

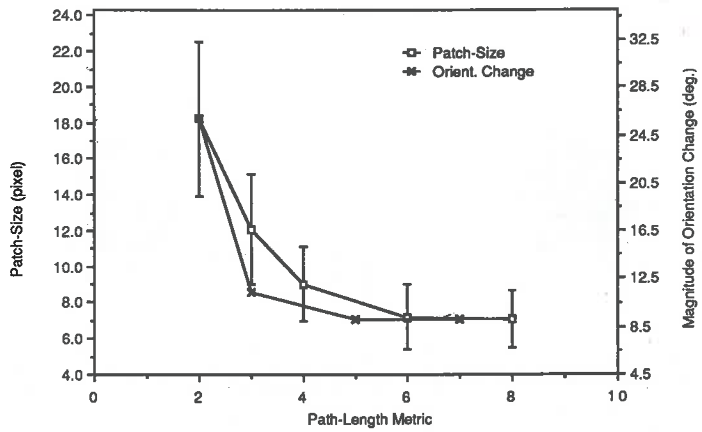

Figure 5. Comparison of data for sensitivity to detecting patches of texture fields (open squares; left scale) with sensitivity to change in orientation within a texture flow (crosses; right scale. Data from (Link and Zucker, 1987)). Observe the rapid change in sensitivity around a path-length of 3, or the point where initial measurement of curvature becomes feasible within the model. The standard error bars are across all subjects and across all densities for the patch detection task.

precisely on inconsistencies (discontinuities) in the orientation response, it follows that there should be a relationship between sensitivity to patches and sensitivity to orientation changes within texture flows. We have previously tested this latter point, by evaluating sensitivity to discontinuities in texture flow patterns (Link and Zucker, 1987). The results are shown in Fig. 5, plotted together with our current data. The similarities are obvious.

6. CONCLUSIONS

Textures come in many varieties, which somehow must be segmented from one another during the image understanding process. In this paper we highlighted two such varieties: those patterns which consist of random, salt-and-pepper fields, and those patterns which possess a clear orientation, or flow. We quantified human sensitivity to patches of random texture within a flow psychophysically, and showed that this sensitivity varied as a function of patch size, path-length (or the length of each stroke comprising the flow), but not as a function of density (over the range in which the streaks integrated into a homogeneous flow). This variation was then related to a computational model for orientation selection, and curvature was isolated as a key feature: sensitivity increased when there was sufficient image structure (path-length) to estimate curvature. For short path-lengths and small patches, the lack of a tight curvature constraint permits too many possible smooth continuations, and sensitivity goes down.

There is, of course, more to texture than its flow structure. For example, differences between texture fields can be detected by non-oriented operators (Zucker *et al.*, 1975; Bergen and Adelson, 1988; Beck *et al.*, 1983; Beck *et al.*, 1987). However, a biologically feasible mechanism for performing this discrimination has yet to be proposed; to date, all schemes involve statistics over local neighborhoods. Whether there is an analog to curvature in these schemes remains to be seen.

Acknowledgements
This research was sponsored by NSERC grant A4470.

REFERENCES

Beck, J., Prazdny, K. and Rosenfeld, A. (1983). A theory of textural segmentation. In: *Human and Machine Vision.* Beck, J., Hope, B. and Rosenfeld, A. (Eds). Academic Press, New York.

Beck, J., Sutter, A. and Ivery, R. (1987). *Comp. Vision, Graphics Image Process*, **37**, 199–325.

Bergen, J. and Adelson, E. (1988). Early vision and texture perception. *Nature*, **333**, 363–364.

Dobbins, A., Zucker, S. W. and Cynader, M. (1987). Endstopped neurons in the visual cortex as a substrate for calculating curvature. *Nature*, **329**, 438–441.

Glass, L. (1969). Moiré effect from random dots. *Nature*, **223**, 578–580.

Glass, L. and Pérez, R. (1973). Perception of random dot interference patterns. *Nature*, **246**, 360–362.

Glass, L. and Switkes, E. (1976). Pattern recognition in humans: correlations which cannot be perceived. *Perception*, **5**, 67–72.

Hubel, D. and Wiesel, T. (1977). Functional architecture of macaque monkey visual cortex. *Proc. R. Soc. London*, **B 198**, 1–59.

Link, N. and Zucker, S. W. (1987). Sensitivity to corners in flow patterns. *Spatial Vision*, **2**, 233–244.

Parent, P. and Zucker, S. W. (1989). Trace inference, curvature consistency, and curve detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, **11**, 823–839.

Prazdny, K. S. (1984). On the perception of Glass patterns. *Perception*, **13**, 469–478.

Prazdny, K. S. (1986). Some new phenomena in the perception of Glass patterns. *Biol. Cybern.* **53**, 153–158.

Prazdny, K. S. (1986). Psychophysical and computational studies of random-dot Moire patterns. *Spatial Vision*, **1**, 231–242.

Zucker, S. W. (1985). Early orientation selection: Tangent fields and the dimensionality of their support, *Comp. Vision, Graphics Image Process.* **32**, 74–103.

Zucker, S. W. (1986). The computational connection in vision: Early orientation selection. *Behaviour Res. Meth., Instrum. Comp.* **18**, 608–617.

Zucker, S. W., Rosenfeld, A. and Davis, L. S. (1975). Picture segmentation by texture discrimination. *IEEE Trans. Comp.* **C-24**, 1228–1233.

Zucker, S. W., David, C., Dobbins, A. and Iverson, L. (1988). The organizaton of curve detection: Coarse tangent fields and fine spline coverings. *Proc. 2nd Int. Conf. on Computer Vision*, Tarpon Springs, FL.