Curvature Cues and Discontinuity Detection in Early Orientation Selection

Norah K. Link
B.A.Sc. (Systems Design)

Computer Vision and Robotics Laboratory
Department of Electrical Engineering
McGill University
Montréal, Québec, Canada

A thesis submitted to the Faculty of Graduate Studies
in partial fulfillment of the requirements for the degree of
Master of Engineering

December 1985

Curvature Cues and Discontinuity Detection in Early Orientation Selection

Norah K. Link
B.A.Sc. (Systems Design)

Abstract

Orientation selection is defined as the detection in an image of orientation structures, the components of curves or flows, and their representation in terms of tangent fields – the orientation at each point. Curves in images are one-dimensional orientation structures that often correspond to bounding contours of objects; flow patterns are two-dimensional, and they provide surface information within the contours. The local orientation of these structures, recovered through orientation selection, provides an initial description of the shape of objects and surfaces. Discontinuities in orientation often signal important events, such as surface creases or surface occlusions. This thesis demonstrates that human sensitivity to such discontinuities also reflects on the kinds of mechanisms by which we might reconstruct and represent curves and flow patterns. A computational theory of orientation selection is outlined, and predictions of discontinuity sensitivity that arise from this model are developed. Psychophysical experiments designed to test these predictions are presented and analysed. The experimental results indicate that the human visual system uses curvature information to reconstruct orientation structures and it uses at least change-in-curvature information to locate discontinuities in them.

December 1985

Curvature Cues and Discontinuity Detection in Early Orientation Selection

Norah K. Link
B.A.Sc. (Systems Design)

Résumé

La sélection d'orientation se définit comme la détection dans une image de structures orientées, c'est-à-dire des composantes des courbes ou des flux, et de leur représentation en termes de champs de vecteurs – l'orientation à chaque point. Les courbes dans une image sont des structures orientées à une dimension qui souvent correspondent aux contours délimitatifs d'objets; les flux ont deux dimensions, et ils donnent de l'information sur les surfaces incluses dans les contours. L'orientation locale de ces structures, obtenues par sélection d'orientation, permet une description initiale de la forme des objets et des surfaces. Les discontinuités d'orientation signalent souvent des événements importants, tels que les plis d'une surface ou l'occlusion d'une surface par une autre. Cette thèse démontre que la sensibilité chez l'humain à de telles discontinuités se reflète aussi sur les genres de mécanismes par lesquels nous pouvons reconstruire et représenter courbes et flux. Une théorie de la sélection d'orientation est esquissée, et des prédictions sur la sensibilité aux discontinuités sont élaborées à partir de ce modèle. Des expériences psychophysiques conçues pour tester ces prédictions sont présentées et analysées. Les résultats expérimentaux indiquent que le système de vision humain utilise l'information sur la courbure pour reconstruire les structures orientées, et qu'il utilise au moins l'information sur le changement de courbure pour localiser leurs discontinuités.

Décembre 1985

Curvature Cues and Discontinuity Detection in Early Orientation Selection

Norah K. Link
B.A.Sc. (Systems Design)

Acknowledgements

I would like to thank first of all my supervisor, Dr. Steven Zucker, for his unflagging enthusiasm for the work, his technical support, and his vote of confidence. For comments on this manuscript and help with the two technical papers preceding it, I would like to thank Yvan Leclerc, Peter Sander, Kamal Gupta, Pierre Parent, Frank Ferrie, Nevine Nassif, and Gracy Mimran. Pierre Lamoureux provided programming support at the outset of the project. A special note of gratitude goes to Pierre Parent for running his simulation of the model on my test data.

I would like to thank a few special friends for their emotional support: Yvan Leclerc; my office-mate and drinking buddy Frank Ferrie; my room-mate, friend and confidante, Gracy Mimran; Maria Mastromonaco and Mindle Levitt for coffee, cookies, good cheer, and encouragement; and the remaining members of CVaRL for good-natured fun. Last but never least, I would like to thank my family who always have been and always will be there to help me: my sisters Shelagh and Martha, and their husbands John and Chuck, my Aunt Elizabeth, and most of all, my mother Abbie and my father Murray. They have my everlasting love.

December 1985

Contents

List of Figures

Chapter 1 Introduction

A particularly striking feature of our visual systems is how quickly and spontaneously we impose three-dimensional interpretations on images. Even when we are presented with a completely homogeneous image that fills the entire visual field, we perceive a three-dimensional space. When inhomogeneities (such as dots, lines, or regions of different colours) are introduced into the image, this space is segmented into distinct surfaces at specific depths (Koffka 1935). How do we construct descriptions of these surfaces, or objects?

Consider a sphere and a cube. One salient feature which allows us to differentiate between them is the shape of their bounding contour (see Figure 1.1. The shape of a contour is described by the direction it takes across an image – its *orientation* – and by the way it changes direction – its *curvature*. The sphere's outline is smooth and has constant but non-zero curvature, while the cube's has orientation discontinuities and zero curvature everywhere else. Now, think of putting the cube in front of the sphere so that only part of the sphere can be seen. The outline of the two objects together contains orientation discontinuities both at the corners of the cube and at the points along the sphere's boundary where the cube occludes it. We can see, then, that while orientation is an essential descriptor of local contour segments, *orientation*, *curvature*, and their *discontinuities* are essential in the full description of contours. In this thesis, I shall be concerned with human sensitivity to orientation discontinuities in contours and with how this sensitivity can be explained within a computational model of curve reconstruction.

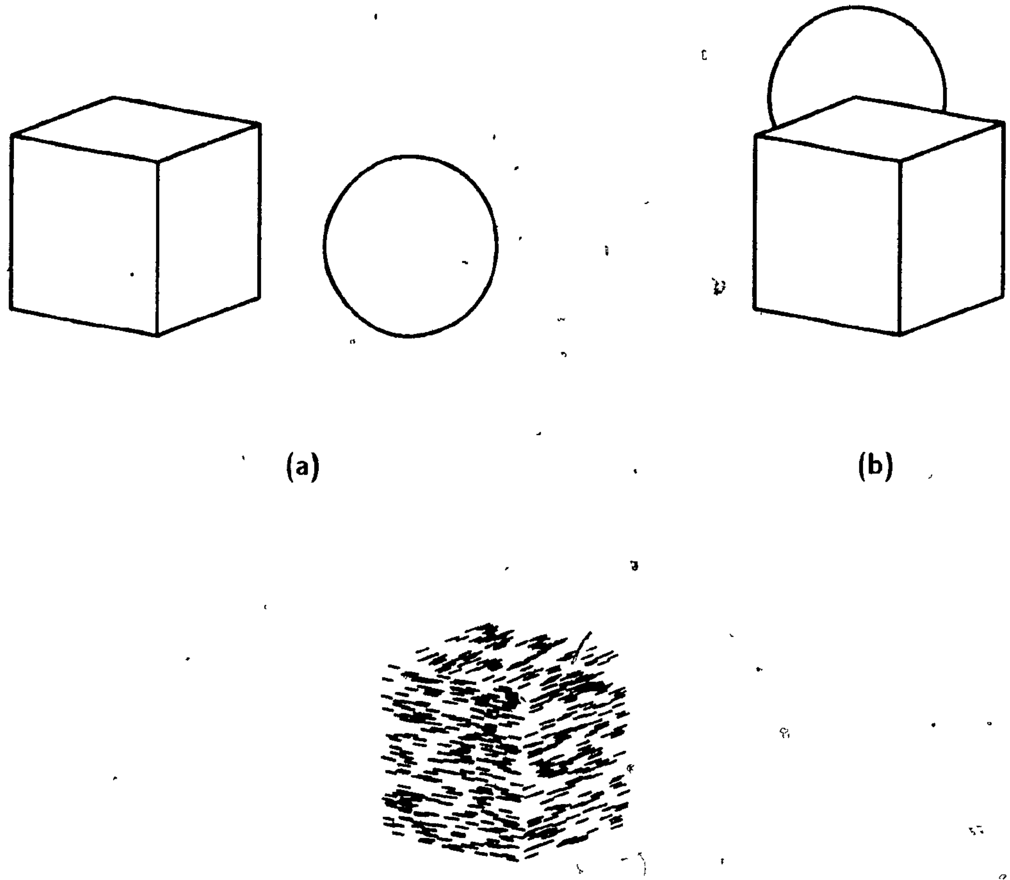

(a) (b)

(c)

Figure 1.1 An example of the importance of orientation, curvature, and orientation discontinuities in shape description. In (a), the outlines of a cube and a sphere are presented. They can be distinguished by the orientation discontinuities in the cube's outline, and the constant non-zero curvature of the sphere's. In (b), the cube is placed in front of the sphere, creating new orientation discontinuities in the outline at the points where the cube occludes the sphere. In (c), flow patterns are used to describe the cube (as if it were covered with fur, for example). Note how the shape is defined by the discontinuous changes in flow orientation at the edges between the two front faces and between the left and the top face.

Now suppose that the cube is covered with fur. Locally, the individual hairs all lie in the about same direction, and the overall impression is of a flow, or a two-dimensional orientation structure. While the flow orientation may or may not change within a single face, one would *expect* it to change abruptly along the edge between two faces; see Figure 1.1c. Therefore, it is also important to be able to detect orientation changes in flow patterns, and in particular to distinguish between smooth and abrupt changes in the flow. The second aspect of this thesis is to investigate human sensitivity to orientation discontinuities in flow patterns within the context of a model, similar to the one mentioned above, for flow reconstruction.

These examples demonstrate that both contour and flow reconstruction are important first steps in the process of shape description. Their orientation structure – including curvature and discontinuities – provides an initial description of shape. For this reason, the description of curves and flows in terms of their local orientation structure is referred to as *orientation selection*. The model of early orientation selection presented in Chapter 2 treats this local orientation structure as a *tangent field* – or set of unit-length line segments, appropriately oriented, at each point of the curve or flow. From this, the actual curves or flows could be reconstructed. Discontinuities in the orientation structure often signal important events, such as surface boundaries, and can be used to segment images into distinct regions in the same way as discontinuities in image intensity, surface depth, or surface orientation (Marr 1976, 1982; Witkin and Tenenbaum 1978). In addition, if curves and flows can be reconstructed from the image and represented accurately, they will provide a set of basis functions along which to integrate surface properties (Zucker 1984*b*, 1985). Contours provide accurate positional information, particularly about the boundaries of objects. Flow patterns provide surface information within the boundaries described by the contours. Since orientation discontinuities segment contours and flow patterns into piecewise smooth regions, locating them is intimately tied to orientation and curvature perception. Our sensitivity to them will therefore provide some insight into the mechanisms by which we reconstruct curves and flows.

The purpose of this thesis is to evaluate human sensitivity to orientation discontinuities and to investigate its implications with respect to a model of early orientation selection that was first proposed by Zucker (1982, 1985). I shall present the results of psychophysical experiments that were designed to test certain aspects this model, and by analyzing the results within the context of the model, I shall show that detection of orientation discontinuities requires curvature and change-in-curvature information. Hence, mechanisms for estimating or representing this information must exist in the cortex. The following paragraphs motivate the methodology used for developing computational vision models such as the one I shall be investigating, and introduce the particular technique I have used to evaluate human sensitivity to discontinuities.

1.1 The Influence of Neurophysiology and Psychophysics on Computer Vision

Computer vision is aimed at developing systems which can analyse and interpret grey-level or colour images of natural scenes. For better insight into the problem of vision – at the task specification level (what functions are needed) and at implementation levels (how these functions might be performed), computer visionists often draw on knowledge of biological vision systems. Different kinds of information are supplied by the fields of neurophysiology and psychology. Specific functions inspired by neurophysiological structures or psychophysical phenomena would include, for example, edge detection (Marr and Hildreth 1980), curve detection (Marr 1976), motion detection (Wallach and O'Connell 1953; Ullman 1979), stereo depth processing (Marr and Poggio 1979), and orientation selection (Glass 1982; Zucker 1985). Specific implementation schemes would include zero-crossings of convolutions with Laplacian of Gaussian operators (Marr and Hildreth 1980), networks of cells that perform Boolean algebra (Poggio and Torre 1978), and local excitatory and inhibitory networks of convolution operators (Zucker 1984*a*).

Once fully designed, a model of a specific task provides predictions about the relationships between stimuli and perceptions through the implemented mechanism. These predictions arise from assumptions made during the development of the model and the design of the mechanism by which the model is implemented. Psychophysical experiments can be used to evaluate the relationships between stimuli and percepts in human vision. Analysis of the results within the context of the model and its predictions leads to conclusions about the validity of the model (as being representative of human vision) and constraints on model parameters. This, then, closes the loop between computer vision and biological vision, as the psychophysical results are used to tune the model, causing it to describe more closely the performance of the human visual system.

The model of orientation selection for curve and flow reconstruction that I shall analyse in this thesis was developed with particular attention paid to the neurophysiology. The analysis that I shall perform involves a psychophysical evaluation of the relationship between different kinds of stimuli and our perception of discontinuities in orientation structures. It is intended to validate and tune the model, and it leads to predictions about the human visual system itself.

1.2 The Use of Dotted Stimuli to Study Orientation Structures

To investigate the processes of curve and flow reconstruction, I shall use dotted stimuli. Under certain conditions, dotted stimuli can be considered equivalent to continuous stimuli. Using them in the experiments allows the manipulation of particular orientation cues by changing the dot positions. In the case of curves, this allows the orientation structure in the neighbourhood of a discontinuity to be changed. In the case of flow patterns, it limits the number of curve derivatives that are made explicit within each curve segment that acts as a flow cue. In both cases, it enables the use of interpolation theory to explain the psychophysical results and to measure the differential properties of the mechanism described by the model. In Chapter 3, I shall expand on the relevance of using dotted stimuli to study each class of orientation structure, and I shall show how I intend to use them.

1.3 Thesis Overview

The model of early orientation selection that I shall investigate through psychophysical experiments is presented in Chapter 2. To provide the reader with some background, the treatment of orientation selection in other computational models of early vision is briefly outlined, and it is shown to be deficient due to its lack of correspondence to the neurophysiology and its computational properties. Following this, the neurophysiology and psychophysics relating to orientation selection are reviewed. A model of curve and flow reconstruction, first proposed by Zucker (1982, 1984*b*, 1985), is then described. This model is based in the neurophysiology, and it is not subject to the same computational problems as previous models. It is this model that my analysis of human sensitivity to discontinuities is based on.

Chapter 3 outlines again the importance of detecting orientation discontinuities and develops the predictions for discontinuity sensitivity that arise out of the model. These predictions are based on the fact that curves and flow patterns are discretely sampled by the imaging process (for example, by the retinal grid), and the model must therefore be able to reconstruct curves and flows from dotted as well as from "continuous" stimuli. Changes in discontinuity sensitivity are sure to arise, then, when certain aspects of the quantisation are changed. These issues are explored in this chapter and put to use in the development of dotted stimuli for the experiments.

Experiments to evaluate human sensitivity to orientation discontinuities in curves and in flow patterns are presented and analysed in Chapters 4 and 5 respectively. The results indicate that detection of orientation discontinuities is a non-local process and that it requires mechanisms for estimating curvature and higher-order curve derivatives (for example, change in curvature) at least over a local neighbourhood. In this way, experiments based on predictions arising from a computational model of vision are used to enhance our understanding of the psychology of human vision.

The content of this thesis is based on two papers written in conjunction with the author's thesis supervisor, Dr. S.W. Zucker (Link and Zucker 1985*a*, 1985*b*). As previously stated, the model of curve and flow reconstruction presented in Chapter 2 was developed by Dr. Zucker. The predictions arising from the model, the experimental design and execution, and the analysis of the psychophysical results, represent my contribution to this work.

Chapter 2 Orientation Selection and Its Role in Early Vision

Vision involves the inference of three-dimensional structures from two-dimensional images. As a first step in constructing this inference, it is particularly useful to extract from the image structures that are unlikely to have arisen randomly, but rather which stand in correspondence with real-world physical structures and can be used to describe them (Koffka 1935; Marr 1982; Witkin and Tenenbaum 1984; Zucker 1985). We have seen that two basic types of orientation structures are representative of such correspondence: one-dimensional contours which arise, for example, from projections of surface occlusions or surface creases; and two-dimensional flow patterns which arise from projections of surface coverings, such as fur. The mathematical properties of image projection guarantee that certain aspects of shape are preserved in the image,[1] so that the shape of these image orientation structures is representative of the shape of objects.

The shape of a contour is described not only by the direction it takes across an image – its *orientation* – but also by the way it changes direction – its *curvature*. Similarly, curvature is useful in the description of the "shape" of a flow pattern. *Discontinuities* – abrupt changes in the orientation structure – often signal important events, such as the edge between two objects. The discontinuity locations also constrain the particular reconstruction of a curve or flow pattern. In this chapter, a model for curve and flow reconstruction is

---

[1] For example, under orthographic projection and under perspective projection when the object is at a distance, the component of orientation that is parallel to the image plane is preserved.

presented. To begin, the treatment of orientation selection processes in early computer vision models is outlined and shown to be inadequate. The pertinent neurophysiology is then reviewed as background to the development of the model. In the remainder of the thesis, the model is analysed with respect to the detection of discontinuities.

2.1 The Reconstruction of Curves and Flows in Early Computer Vision Models

A major problem that models of curve and flow reconstruction must deal with is the trade-off between computational complexity and sensitivity to noise. Since many curves in images are formed by the bounding contours of objects, most early models of vision have concentrated on the latter. The techniques used to find these contours fall into two groups: those that find intensity edges, under the assumption that these often correspond to object boundaries; and those that segment an image directly into regions of interest, and trace the boundaries of these regions to find the contours. The most common methods of contour representation include chain-coding (a sequential representation of the contour points; see Freeman 1974), polygonal approximation (Ramer 1972), and point-for-point representation in a bit-map (a binary copy of the imaging grid with feature points recorded in it; see Marr 1981). While the detection of bounding contours is not exactly equivalent to curve detection, these techniques are typical of the way in which orientation structures have been detected, reconstructed, and represented. They therefore serve to illustrate the trade-off between complexity and noise sensitivity and to demonstrate that orientation selection must be treated as a complex problem.

Among the earliest examples of (intensity) edge detectors are the Robert's cross and the Sobel operators (Levine 1985). These operators yield a measure of the intensity gradient, including both the magnitude of the gradient and its orientation. Intensity step-edges are asserted at local maxima in the gradient magnitude. The problem with these operators is that they are extremely local (the Robert's cross is a 2×2 convolution operator, and the Sobel operator is 3×3), and hence they are highly susceptible to noise in the image. The two most common methods of dealing with this noise have been to threshold the output of the operator and to smooth either the orientation information provided by it or the image itself. Unfortunately, thresholding completely ignores potentially valuable orientation information. Smoothing blurs the operator output, so that the orientation is locally less accurate – thus defeating the purpose of using a local operator.

As an attempt to resolve the problems of noise sensitivity, later operators used several sizes of masks at varying orientations to detect edges (Rosenfeld and Thurston 1971; Marr 1976). However, these methods were computationally very expensive, and it was not clear either how to correlate the information from the different mask sizes or how to use the orientation information to advantage. Since the orientation information was the primary cause of complexity, Marr and Hildreth (1980) later developed a rotationally symmetric operator to detect points along step edges in intensity. The recovery of contours was then treated as a "grouping" process that uses principles of similarity and proximity to link together feature points detected by these edge operators (Marr 1981; Stevens (1978) also used such grouping processes to reconstruct flow patterns).

Other models segment the image into fairly homogeneous regions using statistical methods of region-merging (Meurle and Allen 1968; Pavlidis 1972; Gupta and Wintz 1975; Levine and Shaheen 1981) or region-segmentation (Levine 1973; Schacter, Davis, and Rosenfeld 1976; Tomita and Tsuji 1977; Ohlander, Price, and Reddy 1978; Weszka 1978). The contours fall out of such processes automatically by tracing the boundaries of the regions. These methods ignore orientation information altogether.

In early models of vision, then, the issues of local orientation and of smooth or abrupt changes in curve orientation generally have not been addressed at the level of curve reconstruction. Several researchers have, however, considered how to detect orientation discontinuities in curves at later levels of processing. In general, these methods find the "average" orientation of a curve over neighbourhoods, and they classify discontinuities as changes in the average orientation that exceed some threshold and that are isolated (localised to a point). The actual detection scheme has been implemented as an extension to the chain-coding method of curve representation (Freeman and Davis 1977; Feng and Pavlidis 1973; Rosenfeld and Johnston 1973; Rosenfeld and Weszka 1975), and also as a match-filtering problem acting directly on the image or on a bit-map of the curve (Kruse and Rao 1978).

The fallacy in these early approaches is that they implicitly assume a particular orientation structure. Rather than suppressing orientation information, they impose on the image an orientation structure which may be incorrect, because they never test for violations of the assumption. This occurs in three ways: first, the region segmentation or edge-finding operations impose particular structure (such as planar surfaces) on the regions; second, these same operations impose an orientation structure on the region boundaries; and third, the methods for curve representation, and especially for detecting discontinuities in the curves, assume a particular orientation structure around discontinuities and hence impose this structure over the entire curve. To illustrate the first two points, consider again the direction-independent edge detector mentioned above. The orientation structure of the intensity edge is assumed to be a step discontinuity *and* the edge (or bounding contour) is assumed to be straight within the spatial extent of the operator (Marr and Hildreth 1980). When the underlying image structure deviates from this assumption, the asserted edge (or contour) locations are displaced from their true locations. This usually results in an incorrect assignment of the contour's orientation and curvature. To illustrate the third point, consider the method of detecting corners in chain-coded curves introduced by Freeman and Davis (1977). This method first averages orientation over neighbourhoods and then imposes a (relatively low) threshold on orientation changes before considering them as candidate discontinuities, in order to reduce noise and quantisation effects. However, when the orientation changes locally because the underlying curve changes direction, the averaging will reduce the curvature effect, possibly smoothing out changes in curvature. In addition, discontinuities are defined as isolated above-threshold changes in orientation. Since the change in orientation is above-threshold wherever the curve bends, this definition assumes that orientation discontinuities occur only between straight segments of curves.

In fact, an orientation discontinuity *can* occur between curved segments. Therefore, this method results both in the detection of false discontinuities and in the failure to detect and accurately locate some true discontinuities – thus providing an inaccurate representation of the orientation structure.

To produce an accurate representation, we must reconstruct curves and flow patterns by *first* determining their differential properties, including their local orientation, curvature, and discontinuities (Leclerc and Zucker 1984). Ignoring or overly smoothing orientation information is not, then, an adequate solution to the conflict between complexity and noise sensitivity. It may be, in fact, that to achieve robustness while still deriving accurate information, we will be forced to use more complex models. Furthermore, as was earlier suggested, if curves and flow patterns were recovered first, they could provide both valuable early descriptions of shape and a set of basis functions to anchor such calculations as depth and surface orientation. Given their usefulness, orientation structures should be recovered and their properties computed as early as possible (Witkin and Tenenbaum 1984; Zucker 1984*b*, 1985); neurophysiological research suggests that we may reconstruct curves directly from image-like structures (Hubel and Wiesel 1962, 1977; Schiller, Finlay, and Volman 1976).

Some methods have been developed to extract line and edge orientation directly from images (Rosenfeld, Thomas, and Lee 1970; Hueckel 1973). However, these methods also generally assume a locally straight line structure. A relaxation labelling technique for line and curve enhancement (Zucker, Hummel, and Rosenfeld 1975) tries to use local orientation structure to perform the dual functions of filling in gaps and reducing noise in the image. However, this first attempt at curve enhancement by consideration of a slightly more global structure does not take into account curvature consistency. While it does not rule out or suppress curved lines, it gives preference to low curvatures. (No attempt is made to explicitly detect orientation discontinuities in curves. The algorithm simply asserts the presence of a curve and its orientation at each point by allowing neighbouring orientation assertions to either support or inhibit each other. The inhibition increases as the difference between neighbouring orientations increases, thereby favouring low curvatures.) As the experiments presented in Chapters 4 and 5 shall show, such curvature information is certainly an important factor in how we perceive curves. In particular, it is required to explain our perception of orientation discontinuities in curves. The model developed in Section 2.3 is based in differential geometry, and it determines the differential structure of curves and flow patterns before imposing any interpretations on the image. The relevant background in neurophysiology and psychology follow.

2.2 Orientation Selection in Neurophysiology and Psychophysics

How do biological vision systems reconstruct curves? Neurophysiologists have postulated that so-called simple cells might be involved. Their spatially localised response and orientation selectivity have led Hubel and Wiesel (1962) to suggest that, by selecting the strongest response at a position, simple cells can become *line detectors*. While this may be true for isolated straight lines, the receptive field structure is insufficient to explain curve perception. The optimal orientations for stimulating simple cells appear to be discretely distributed (Hubel and Wiesel 1962, 1977; Schiller, Finlay, and Volman 1976) in about 10° steps. In addition, each cell actually responds to a range of orientations that is 10° to 20° wide. Thus, each simple cell acting individually as a line detector could not explain either our sensitivity to all orientations or to small changes in orientation. To further complicate matters, since the average orientation response for each cell is broader than the step-size between optimal stimuli, one line possessing an orientation somewhere in between would stimulate two cells, although it would stimulate neither one optimally. It would seem that orientation, and curve, perception is not as simple as "detection" of oriented line or contour segments by individual simple cells.

This does not imply, however, that simple cells are not involved in curve perception or that the orientation information they provide is unimportant. But the above description oversimplifies the neurophysiology in attempting to ascribe to it a particular function. It does not consider, for example, other properties of simple cells – that they vary in size, and that some exhibit end-stopping inhibition – or how they are arranged in the cortex. In particular, the missing element is the *global interactions* – or computations – between these local operators that occur before a percept is constructed. Such interactions are necessary to sort out possibly conflicting responses and to fill in gaps, and they have been shown to be an important factor in generating response patterns of simple cells (Blakemore and Tobin 1972; Sillito et al. 1980). Interactions between neighbouring orientations have also been noted in human perception that are strikingly similar to the broad overlapping orientation tuning of simple cells (Movshon and Blakemore 1973; Carpenter and Blakemore 1973). More recently, the possible existence of "curvature detectors" has been investigated (Timney and Macdonald 1978; Riggs 1973; Heggelund and Hohmann 1975; Crassini and Over 1975; Foster 1983), but this research is fraught with the same pitfalls discussed above: namely, that it does not take into account global interactions in the cortex.

Simple cells and their interactions are perhaps the earliest level of orientation processing. At a higher level, Attneave (1957; also Attneave and Arnoult 1956) ~~has~~ asserted that "shape" can best be described by a set of line segments, each possessing a position, orientation, and length, and the set further described by their connections (or their relative positions). Consider Attneave's famous drawing of a cat (1954), in which he used straight line segments *appropriately placed* to construct a recognizable outline of a cat. However, to construct and place the line segments appropriately (i.e., to determine length and orientation, as well as connectivity), a mechanism is needed for calculating curvature, changes in curvature, and curvature maxima.

In this thesis, I shall concentrate on the lower-level inference of contours, rather than on the higher level representation of shape. A new model of orientation selection, described in the following section, shows that interaction between simple-cell-like operators can provide the necessary information about higher-order derivatives, and that curves can be fit in this way to stimuli in the discrete retinal array. Further consideration of how the model detects discontinuities will lead in Chapter 3 to predictions suitable for psychophysical experiment. The results presented in Chapters 4 and 5, analysed within the context of the model, also point to the estimation of higher-order derivatives by the visual system. Interestingly, for computational and geometric reasons, information about orientation and curvature appears to be compiled everywhere, while derivatives higher than curvature need be estimated only over local neighbourhoods. This neighbourhood information then permits the precise identification of discontinuities.

2.3 The Reconstruction of Curves and Flows from Orientation Cues

To build a computational model for the reconstruction of curves, we first need to consider the neurophysiology. But in order to avoid either oversimplifying its function or simply building an imitation of it, we must then abstract ourselves from the level of neurons and concentrate more on the mathematics of the problem. In this section, one such model is presented in sufficient detail to provide a context for making predictions about sensitivity to discontinuities and for explaining the psychophysical findings presented in Chapters 4 and 5. To simplify the presentation, a model for curve reconstruction is presented first, and it is then extended to encompass flow reconstruction. This model of early orientation selection was first developed by Zucker (1982) and was later refined by Dr. Zucker, P. Parent, and myself.

2.3.1 A Model for Curve Reconstruction

The computational scheme for curve reconstruction consists of a two-stage procedure (Zucker 1982, 1985):

Stage I. Construction of a tangent field corresponding to the curve orientation at each position. This is accomplished in two steps:

1. Convolution against simple-cell-like operators to produce initial orientation estimates;
2. Interactions between these convolutions to:

   (a) estimate curvature

(b) eliminate noise and parameterisation effects.

Stage II. Interpolation of the curve from the tangent field.

Stage I of this procedure reduces to a physiological model as layers of simple cells, or networks of excitatory and inhibitory connections between simple cells to produce the desired feedback effects. It is important to note that Stage I produces a tangent field, in contrast to earlier methods that simply detect lines. The tangent to a curve at a point is the best linear approximation to the curve in a local neighbourhood around that point, and it is also the first derivative (with respect to arc length) of the curve at that point. The tangent field is a set of unit-length line segments having the same orientation as the tangent to the curve at each point on the curve. Curvature (change in orientation) information is used in Stage I to recover the tangent field, while in Stage II information about change in curvature becomes relevant.

2.3.1.1 The Relevance of Discrete Input to the Model

Recall from the Introduction that the experiments will use dotted stimuli to evaluate human sensitivity to orientation discontinuities in curves and flow patterns. At this point in the discussion, it is instructive to note that discrete inputs to the model are highly relevant, and they can actually be used to help develop the model. The input to our visual system is a pointillist array given by discrete retinal receptors. In addition, simple cells are quantised both spatially and in orientation. Although the visual system apparently *infers* continuous curves and continuous objects from the image, it is never *presented* with continuous data. In addition, noise in the optical system may degrade the discretised image so that a continuous curve may not even stimulate adjacent receptors. Therefore, any mechanism for detecting and reconstructing curves must infer the curve out of discrete points.

The goal of Stage II, then, is to infer a continuous curve through discrete points using the orientation and curvature information provided by Stage I. Note that this is made possible initially by the fact that the convolution operators in Stage I are linear. They can therefore be stimulated equally well by pairs of dots having the appropriate orientation as by short line segments of the same orientation, the only constraint being that the dots must be close enough together to both fall within the excitatory range of the operator. The model was developed with exactly this property in mind, enabling the approach to Stage II as an interpolation process between sample points in a discretised curve. This in turn leads to an explicit definition of discontinuities and a method for detecting them.

The linearity of the convolution operators in Stage I, corresponding to the linear operating range of simple cells, also makes the use of dotted stimuli in the experiments both valid and highly effective. Dotted stimuli, as I shall show in Chapter 3, stretch the model to its limits and provide isolated and well-controlled orientation and curvature cues as input to Stage I and, therefore, also to the interpolation process. This method thereby allows us to infer which information is used by the orientation selection mechanism to reconstruct curves and flow patterns and to detect discontinuities in them. To clarify the way in which these predictions arise, the details of the model will be presented in the following sections as a reconstruction process acting on sampled curves. The reader should keep in mind that the same process applies whether the sample points are connected (as in continuous curves) or separated by small spaces (as in dotted curves). Of course, a significant amount of processing precedes the convolutions, and the specific structure of the convolution operators themselves will affect how the curve is reconstructed. However, the analysis of the model in this thesis does not rest on these details, and hence I shall also assume binary images.

2.3.1.2 Obtaining the Tangent Field from Imprecise Orientation Cues

Stage I of the model involves the construction of a tangent field to the curve, or a representation of the curve in terms of its orientation (the first derivative with respect to arc length) at each point. Assuming that simple-cell-like operators provide local orientation information everywhere along a curve, how can these localised responses be interpreted? In other words, how can they interact to detect, reconstruct, and represent *curves* which pass from one location to the next (out of one receptive field and into a neighbouring one) and which change orientation? How can higher-order derivatives be calculated? This requires a comparison of orientation estimates at several locations, together with assumptions of smoothness and of how quickly curves are expected to twist around.

Such comparisons can be performed, and the assumptions made explicit, within a *local excitatory and inhibitory network* of cells (or operators) that continue to signal new estimates until the network reaches equilibrium (Zucker 1984*a*). However, rather than operating on actual estimates of orientation, such as might be obtained by selecting the maximum response at a particular location, the network can compare the overall pattern of responses to an expected pattern of response for any particular curve structure (Zucker, 1984*b*). In this way, not only orientation but local curvature information can be derived from the response pattern. Computationally, representing curvature explicitly has several advantages (Parent and Zucker 1985). In particular, it constrains the orientation estimates within a local neighbourhood, permitting more accurate initial guesses and reducing the number of comparative iterations required to reach equilibrium; see Figure 2.1. It requires, however, that the size of the operators vary so that the larger ones can accommodate sections of curves that bend while the smaller ones provide more local information and help to define the spatial resolution of the curve. The larger operators, however, will respond to straight stimuli (thin enough to stimulate smaller cells) over a slightly broader range of orientations than smaller operators. Thus, the broad orientation tuning of some simple cells (Hubel and Wiesel 1977) may actually be an *advantage* when curved stimuli are considered, as long as the ranges of orientation responded to by individual cells overlap to cover all possibilities. The hypercomplex, or end-stopping, properties of some of these cells also result in response patterns that reflect the local curvature (Orban 1984).

Higher-order derivatives can be calculated implicitly with each comparison, and the more global we allow the comparisons to become, the more we can confine the curve to specific locations, orientations, and curvatures. The notion of arbitrarily many curve

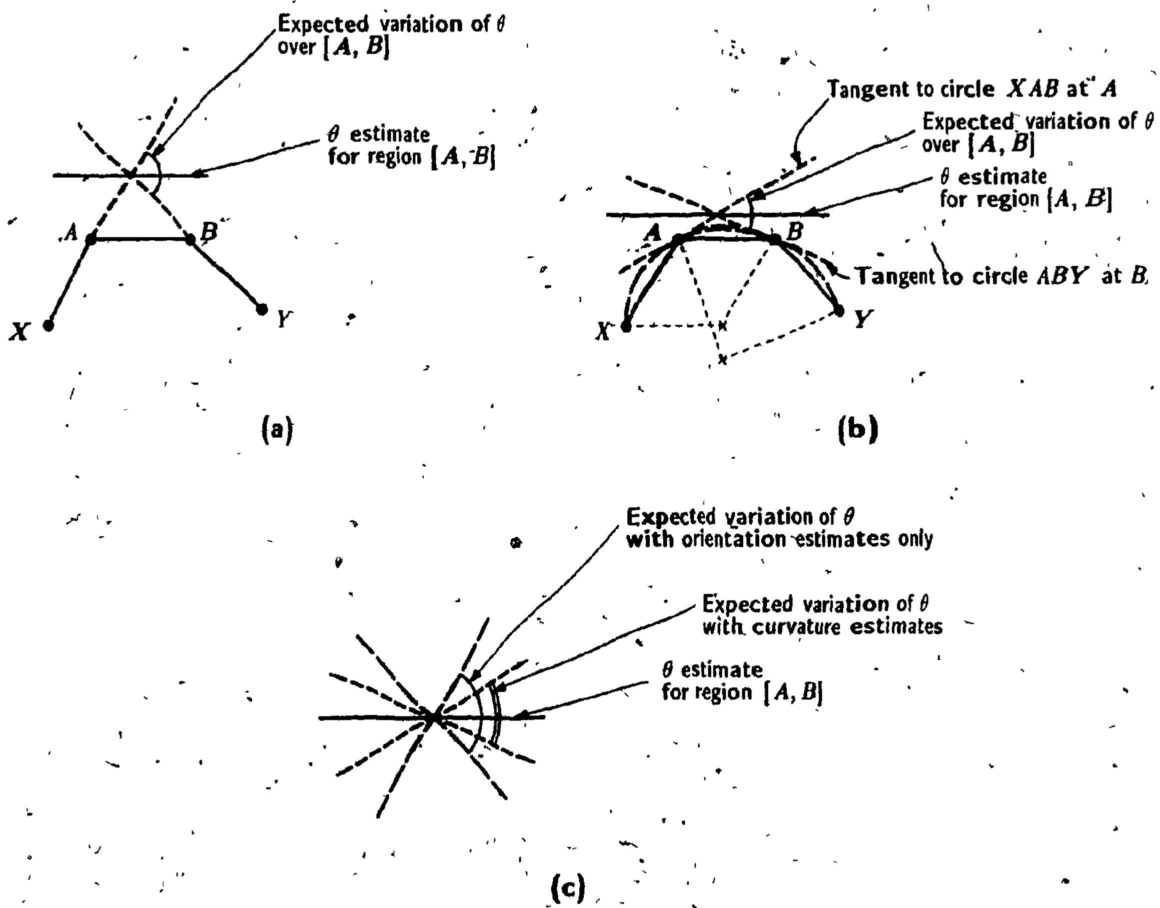

Figure 2.1 This example illustrates the variability of orientation estimates (a) when only orientation estimates are available and (b) when both orientation and curvature estimates are available. The difference is shown in (c). The smaller variation when curvatures are known indicates how the criteria for locating orientation discontinuities can be changed, allowing discontinuities corresponding to smaller changes in orientation to be identified.

derivatives being calculated and represented by the visual system has been advanced by Watt and Andrews (1982). However, Watt and Andrews postulate several mechanisms, perhaps working in parallel, to calculate the various derivatives. The model presented here calculates the first two derivatives using a single mechanism whose purpose is to construct the tangent field of the curve and to provide direct input to an interpolation process. It is in this second stage that higher-order derivatives need to be represented, but only over neighbourhoods rather than at every point.

2.3.1.3 Curve Interpolation and Discontinuity Detection

Let us turn now to Stage II of the model. The process of constructing a curve which is constrained to pass through certain given locations is known as *interpolation*. Note that an infinite number of curves can be made to pass through any discrete set of sample points. Interpolation processes, therefore, impose certain assumptions about the properties of the inferred curve to constrain it sufficiently so that it will always be unique. For example, the derivatives of the curve may be constrained to lie within a certain range. Usually, the number of times the curve can change direction (i.e., that the sign of the curvature can change) between adjacent sample points is limited. The previous section showed that networks of simple-cell-like operators can construct and represent the tangent fields (or first derivatives) of curves (Stage I of the model). The advantage of speaking of Stage II in terms of interpolation is that interpolation theory provides a framework for defining and detecting discontinuities in the curve.

Interpolation theory states that if n positions are represented, then the underlying curve can be approximated by a polynomial of degree $n-1$, and all derivatives of the curve of order n or greater must be assumed to be zero. In practice, however, the resolution of Stage I is limited – the system can only accommodate, say, m derivatives. Then even if the number of sample points n is larger than m, all derivatives of order higher than m must be assumed to be zero. *Discontinuities* must be asserted at points where this assumption is violated in order to cause the interpolated curve to pass through the dots. The order of the discontinuity refers to the lowest-order curve derivative which undergoes an instantaneous step change. The mathematical definition of a discontinuity states that the limit of this derivative, as predicted by the integral of higher-order derivatives, depends on the direction along the curve from which the limit was obtained. This thesis concentrates on *orientation discon' ties* – first-order discontinuities – that occur, for example, at corners, when the orientation changes suddenly and unpredictably at a point. At this point, curvature and all higher-order derivatives are infinite or undefined, and hence we assert an orientation discontinuity.

The assumption that higher-order derivatives must be zero can, however, be relaxed during Stage II to effectively increase the degree of the representation. Recall that the input to Stage II is a set of quantised orientation and position estimates, as well as coarse curvature estimates. During the interpolation stage, the estimates can be allowed to vary within a limited range that reflects both a maximum expected variability (exemplified by how coarse the curvature estimates are) and the characteristics of other neighbouring estimates. One way to achieve this is to represent higher-order derivatives of the curve, not necessarily explicitly at every point, but at least as changes in the Stage I estimates over some open neighbourhood. To illustrate, consider the first degree approximation in which each pair of points is joined by a straight line. Unless the curve is perfectly straight (that is, unless the change in orientation – or curvature – is zero) the interpolated curve will have orientation discontinuities introduced at each point during Stage I; see Figure 2.2. During Stage II, the relaxation of the assumption that the curve is locally straight might permit a smooth interpretation. Notice in particular that the orientation change at P is the same for both curves. If relaxation of the original assumptions were permitted by representing and comparing curvature over local open neighbourhoods, the angle θ at point P in part (a) would most likely be consistent with the other orientation changes in the neighbourhood – that is, the difference between them would be negligible, and a smooth interpretation would result. However, the same orientation change at P in part (b) would most likely differ significantly from the neighbourhood estimate for curvature, resulting in the assertion of a discontinuity at P. Unless a mechanism for comparing orientation changes over a neighbourhood (to determine consistency) exists, the orientation change at P must receive the same interpretation for both curves – the discontinuity in part (b) of the figure could not be distinguished from the smooth curve at P in part (a).

How, then, could higher-order derivatives actually be represented and used in Stage II? In principle, a curve can be assumed to be straight within some small neighbourhood.

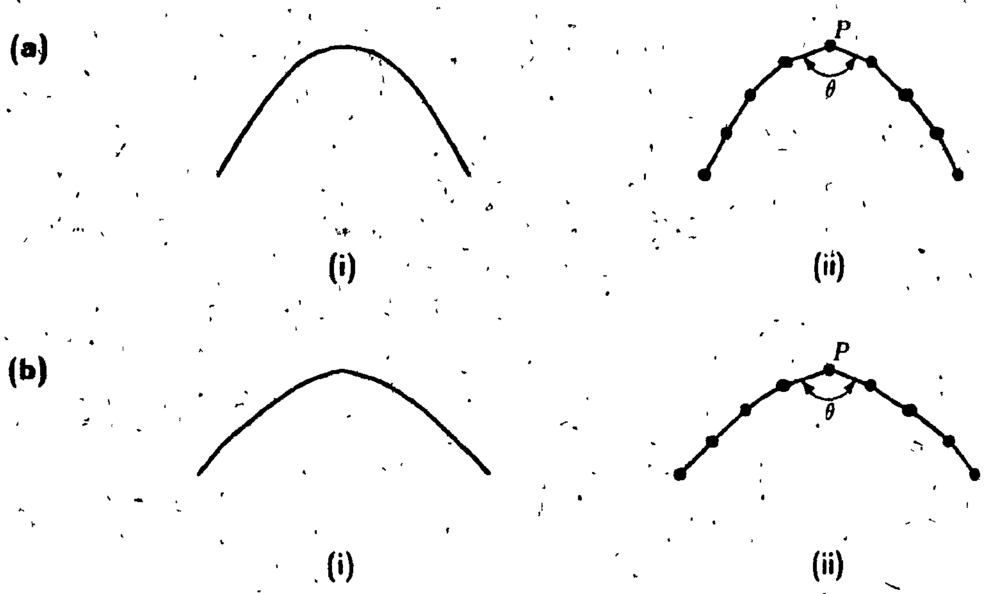

Figure 2.2 Straight-line interpolation of two piecewise smooth curves. In each part, (i) is the original curve, and (ii) is the sampled and interpolated curve. Note that the orientation discontinuity introduced at each sample point in the interpolated curve is ambiguous. The angle θ is the same in both part (a) and part (b), although in part (a) the curve is smooth at point P, while in part (b) P is the location of an orientation discontinuity. This demonstrates that we must have estimates of higher-order derivatives, or a knowledge of several neighbouring estimates and their relationships, in order to resolve such ambiguities and locate the discontinuities.

This permits the tangent – or the best straight-line approximation to the curve over the neighbourhood – to be estimated at the neighbourhood centre. Over a slightly larger neighbourhood, however, curvature cannot be assumed to be zero. If it is assumed to be constant, then it can be estimated by differencing two neighbouring orientation estimates and normalising by the distance between them. The curvature estimate obtained in this way provides an approximation to the osculating circle of the curve at the centre of the neighbourhood. As this neighbourhood is moved along the curve, the curvature can change – but only gradually. If too large a change in curvature occurs between adjacent neighbourhoods, and especially if this change is inconsistent with the local changes between neighbourhoods to either side along the curve, then an orientation discontinuity can be asserted.

In practice, however, both the spatial domain and the orientation domain are quantised.

Problems with the quantisation and with noise may even require several estimates to be obtained over neighbourhoods of varying sizes (Parent and Zucker 1985). Therefore, a curvature estimate obtained by differencing orientation estimates is not exact. Rather, the curvature is known to lie within some range of the estimate. The size of this range reflects an error tolerance that arises out of the particular quantisation. Consider again the example when only orientation estimates are obtained in Stage I, and follow the first line in Table 2.1. The estimates of the tangent are obtained over pairs of points by assuming that curvature is locally zero – the orientation of the tangent is then equal to the orientation of the dot pair. Over a slightly larger neighbourhood, the curvature can be approximated as the difference between adjacent (or nearby) orientations. This requires the assumption that curvature is constant within the neighbourhood. However, since these estimates are noisy (due to the spatial and orientation quantisation) we must relax this assumption so that curvature is only constant to within some error tolerance ϵ. As this neighbourhood is moved along the curve, then, the curvature would only be expected to change within this error tolerance (that is, it should not jump by more than one curvature range). Note that when these constraints are imposed, the orientation of the tangent at some point can be predicted from the orientation at a nearby point and the curvature estimates. However, if there is an orientation discontinuity at some point P, the predicted orientation at P will depend on which side of P (along the curve) the prediction is made from. This is similar to the mathematical definition of an orientation discontinuity. The discontinuity will equivalently affect the behaviour of the curvature representation – the change in curvature over the neighbourhood containing P will exceed the error tolerance ϵ. Thus, the *local* constraint is used to obtain the initial (orientation) estimates, but the *neighbourhood* constraint is imposed to estimate higher-order derivatives and to locate the discontinuities. The remaining rows in Table 2.1 generalise these constraints over several orders of Stage I estimation.

2.3.2 From Curves to Flows

Recovering flow in images is more complex than recovering curves because flows are

Stage I Estimates	Estimation Neighbourhood Size*	Local Assumptions	Constraints On An Open Neighbourhood (N)
θ	2	$\kappa = 0$	$\kappa \approx c_0 \; : \; \Delta_N(\kappa) \leq \epsilon_0$
θ, κ	3	$\partial\kappa = 0$	$\partial\kappa \approx c_1 \; : \; \Delta_N(\partial\kappa) \leq \epsilon_1$
$\theta, \kappa, \partial\kappa$	4	$\partial^2\kappa = 0$	$\partial^2\kappa \approx c_2 \; : \; \Delta_N(\partial^2\kappa) \leq \epsilon_2$
$\theta, \kappa, \partial\kappa, \partial^2\kappa$	5	$\partial^3\kappa = 0$	$\partial^3\kappa \approx c_3 \; : \; \Delta_N(\partial^3\kappa) \leq \epsilon_3$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\theta, \kappa, \partial\kappa, \ldots, \partial^{n-3}\kappa$	n	$\partial^{n-2}\kappa = 0$	$\partial^{n-2}\kappa \approx c_{n-2} : \Delta_N(\partial^{n-2}\kappa) \leq \epsilon_{n-2}$

* number of sample points used to produce initial estimates

θ = Tangent;
κ = Curvature;
$\partial\kappa$ = 1^{st} derivative of curvature;
$\partial^m\kappa$ = m^{th} derivative of curvature.

c_i = constant value of i^{th} derivative of curvature;
$\Delta_N(x)$ = magnitude of change in x over neighbourhood N;
ϵ_i = small bounded tolerance variable based on the quantisation.

Table 2.1 For each level of resolution – the estimates that are obtained directly from the image, the local neighbourhood size used to produce the estimate; the local assumptions imposed in obtaining the estimates; and the constraints used to relate neighbouring estimates. Discontinuities are asserted at locations where the neighbourhood constraints are violated: *where a change in the highest-level estimate for* neighbours is too high within a local neighbourhood (Δ_N locally exceeds ϵ). Note that since the neighbourhoods overlap, this is equivalent to saying that the change is unpredictable given the changes between other neighbouring estimates. Orientation discontinuities are asserted where this change percolates back up through the derivatives to also produce an unpredictable change in the curvature estimates over the neighbourhood N.

two-dimensional structures whose representation in images (as flow patterns) is more sparse than that of curves. Informally, a flow pattern is defined as a dense covering of a surface with a family of curves that are locally parallel almost everywhere.[2] think of this as arising from a limiting process: consider a surface covered by pin-stripes. Now, imagine adding more and more pin-stripes to the intermediate spaces while at the same time shrinking the width of each pin-stripe. A mathematical idealisation of a flow pattern is achieved

[2] For a more precise mathematical presentation, see Zucker 1984*b*.

when the stripes are infinitesimally thin and densely packed. In practice, however, only short segments of these stripes can be displayed, and each successive piece of contour that is displayed (in the direction of the flow) must be displaced from the last in the direction perpendicular to the flow; see Figure 2.3. Information along any one flow contour is therefore extremely sparse: noise in the image will degrade even further the initial orientation estimates. There is not enough information to recover the curves that locally represent the flow transformation using exactly the same process described in Section 2.3.1. Additional constraints must be imposed. The most prominent one arises from the fact that flows are locally parallel; in flow patterns, even though the representation of the flow is sparse, this is approximately true (Glass 1969; Stevens 1978). This can readily be incorporated into Stage I of the model by requiring that the orientation estimates be approximately equal in the perpendicular direction – or, equivalently, by averaging the tangents perpendicularly.

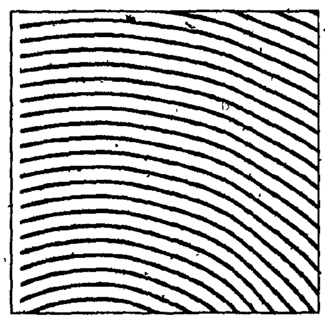

(a)

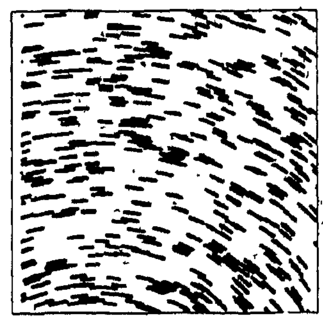

(b)

Figure 2.3 Examples of (a) a pin-stripe pattern and (b) a flow pattern. When the stripes in part (a) are made infinitesimally thin and close together, they define a continuous flow. However, this flow is impossible to display. Therefore, flow patterns are composed of short curve segments that are displaced from one another both in the direction of the flow and perpendicular to it, in order to represent the flow over both dimensions. Note that in a continuous flow, the pinstripes are, by definition, locally parallel. For (discrete) flow patterns, this can only be approximately true – the parallelism breaks down as the flow changes direction.

Imposing new constraints, however, will affect later processing and, hence, the final percept. In particular, when the constraints are violated locally, the perceived flow will differ from the true one in a predictable way. One effect of averaging the orientation estimates in a direction perpendicular to the flow is that the exact positional information associated with the cues is lost – a property of flow patterns noted by Zucker (1982). Technically, then, Stage II becomes an estimation process where flows are involved, rather than an interpolation process; the resulting percept is of the best-fitting flow transform, not an exact fit. This does not change the way in which discontinuities are defined or detected by Stage II. It does, however, also result in a smoothing of orientation information across the image, leading to a loss of sensitivity to changes in orientation, or equivalently a loss of curvature information. Note that the assumption of local parallelism is *by definition* violated in regions where the flow changes orientation. In these regions the model will attempt to equalise orientation estimates that are inherently different, causing the apparent change in orientation to decrease in magnitude. Therefore, changes in orientation are generally detected only when they also line up and thereby provide support for a region of flow change. The width of this region depends both on how local and how large the change in orientation is. Since discontinuities are asserted only where this region is (ideally) infinitesimally thin, it is to be expected that the orientation averaging that results from the assumption of local parallelism would decrease our sensitivity to discontinuities; see Figure 2.4.

It is desirable, therefore, to inhibit the assumption of local parallelism where we have evidence that it is violated. Obtaining higher-order estimates from individual flow cues would provide such evidence. For example, curvature estimates would tell us not only how orientation estimates should change in the direction of the flow, but also how they should change in the direction perpendicular to the flow. This would allow the region over which estimates are averaged perpendicularly to be reduced and the assumption of local parallelism to be relaxed. We might expect, then, that the effect on discontinuity sensitivity of obtaining, wherever possible, various orders of approximation from the flow cues – see Table 2.1 – would be magnified for flow patterns.

2.4 Discontinuity Sensitivity and Models of Orientation Selection

Since the detection of discontinuities is an important aspect of the model, we can

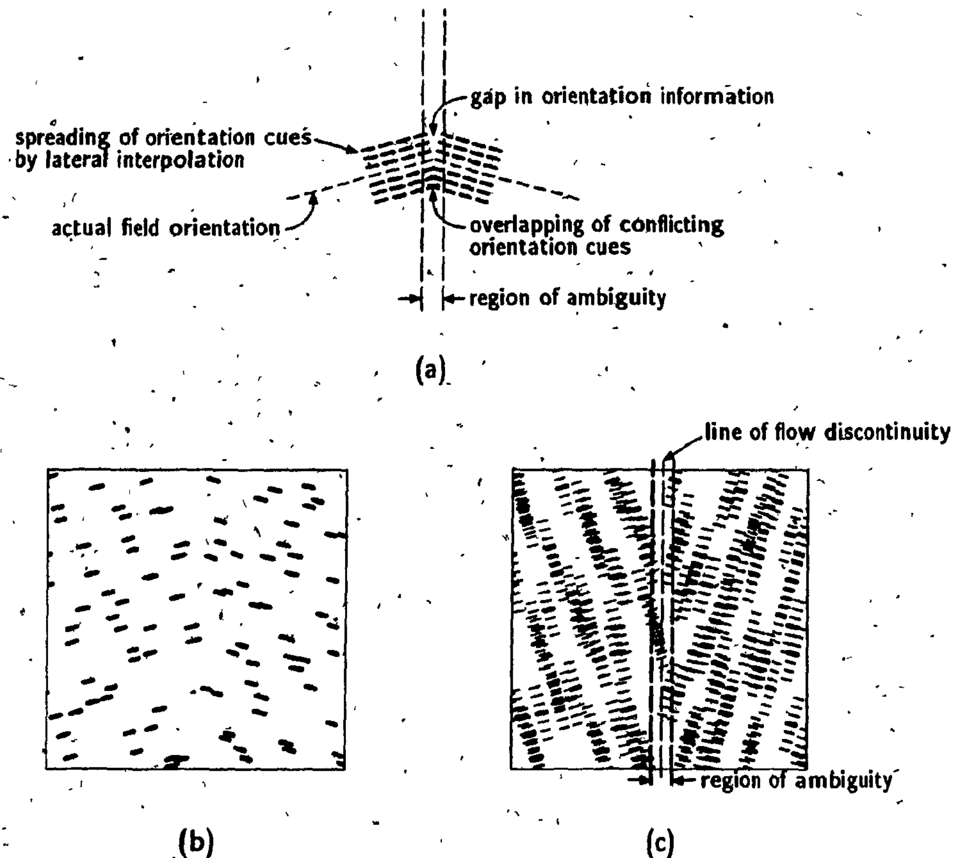

Figure 2.4 The ambiguity introduced by the assumption of local parallelism at changes in orientation. In part (a), the lateral propagation of support for the orientation cues results in conflicting orientation support inside the corner. A gap in orientation support outside the corner compounds these ambiguities. A flow pattern with this field orientation is shown in (b) and is copied in (c) with some lateral spreading of orientation cues marked to show the conflicting information in the central region.

study human sensitivity to them in order to constrain the mechanism. In particular, as was pointed out in relation to early models of vision, the complexity of the model is directly connected to how discontinuities are smoothed. This complexity is determined by the size of the neighbourhood over which information is integrated and by the number of curve derivatives that can be represented by the model. In the next chapter, I shall discuss in more detail how and why we can expect sensitivity to orientation discontinuities in curves and in flow patterns to vary, with particular reference to the order of the interpolation. This gives rise to specific predictions about how discontinuity sensitivity changes when dotted curves and flow patterns are used to manipulate the orientation and curvature content of the image. I shall use these predictions to set up paradigms for studying human sensitivity to this manipulation. The psychophysical experiments performed to test out the predictions are presented in Chapters 4 and 5.

Chapter 3 Discontinuity Sensitivity in Curve and Flow Reconstruction

Discontinuities in orientation structures mark possible boundaries between objects. This is important both for the initial segmentation of an image into coherent surface regions and for qualitative shape description. Our sensitivity to them also provides a means for assessing the model of early orientation selection presented in Chapter 2. Specifically, it will allow us to make inferences about the order of the estimation and interpolation in the orientation selection process.

We have seen that any model for orientation selection must be able to recover curves and flows from discrete stimuli. This does not mean, however, that the model is insensitive to variations in the quantisation. The model will always provide some reconstruction, but it is not clear that this reconstruction will always be correct – this depends on how well the assumptions (of smoothness and maximum curve degree) describe the stimulus. This provides a basic paradigm for studying discontinuity detection: we can use different quantisations of curves and flow patterns to control how well the assumptions describe the underlying functions (in the case of curves) and whether or not the quantisation itself limits the way the mechanism can function (in the case of flows). In this chapter, I shall discuss in more detail the effects of quantisation on the model with particular reference to the detection of discontinuities, first for curve and then for flow reconstruction. I shall show how the smoothness assumptions and the finite order of the system can be expected to result in the misperception of discontinuities under certain circumstances that are controlled by the actual quantisation. The specific predictions that arise from this discussion have been tested psychophysically, and the results are presented in Chapters 4 and 5.

3.1 Discontinuity Detection in Curves

The detection of discontinuities by Stage II of the orientation selection process (described in Section 2.3.1.3) depends on the structure of orientation cues in a neighbourhood around the discontinuity. Using dotted – or sparsely sampled – curves allows us to manipulate this structure in a controlled manner, and hence to study in some detail its effect on discontinuity sensitivity. The following paragraphs justify the use of dotted stimuli to study the psychophysics of early curve perception. Their use in studying discontinuity sensitivity is also motivated by a demonstration of the effect of differential sampling on the perception of orientation discontinuities. This leads to a paradigm for studying human sensitivity to orientation discontinuities in curves as a function of the local orientation structure.

3.1.1 From Dotted to Continuous Curves

Recall from the discussion in Section 2.3.1.1 that the mechanism for detecting and reconstructing curves must operate on a pointillist input array, such as that given by the discrete retinal receptors. Because noise in the optical system degrades the discretised image, a continuous curve may not even stimulate adjacent receptors, and so the mechanism must be able to infer curves out of discrete, non-connected points. From this, two related observations emerge. First, one can assume that it is legitimate to study curve perception by studying the perception of *sampled* or *dotted curves*, as long as the sample points are close together. Secondly, if the quantisation of the curve changes when it is imaged – perhaps due to a slight change in position relative to the sampling grid, or due to noise – the arrangement of data points will change. This will affect the *geometry* of the curve: if the quantised positions of the curve change relative to each other, then the orientation of neighbouring points will change, and so may the curvature. Since geometry varies with discrete input, we can take advantage of the first observation and use discrete stimuli to study curve detection. The following paragraphs expand on these observations and show how I will use them

3.1.1.1 Densely Dotted Curves Are Equivalent to Continuous Curves

Since continuous curves are presented to our visual systems in a quantised form, how close together do the sample points have to be for us to perceive the curve as a unit? Is it possible that some dotted curves are processed *equivalently* to continuous ones? Theoretically, the Nyquist criterion answers this question for us (Oppenheim and Schafer 1975), but this line of reasoning begs the issue of what mechanisms are used to reconstruct curves by the human visual system. From the point of view of perception, the relevant questions are:

1. How does the curve percept vary as a function of the density of sample points?
2. How does the curve percept vary as a function of the sampling phase (given a regularly sampled curve)?

In answer to the first question, consider Figure 3.1. In this figure, a sinusoid has been sampled with a uniform dot size but using three different sampling intervals.[3] In part (a), the curve is essentially continuous. In part (b), although the curve is dotted, the peaks are still smooth. In part (c), however, the dots are far enough apart that the peaks no longer appear to be smooth – rather, they appear triangulated. Somewhere in between (b) and (c) is a transition point where the equivalent curve percept actually breaks down (disregarding phase). This point is actually dependent on the size of the dots used to represent the sample points, and I shall refer to it as the size/separation constraint after Zucker and Davis (1985). Under the assumption that simple cells are involved in curve perception, the size/separation constraint would relate directly to the spatial distribution of a simple cell's response: one would expect that over the linear operating range for simple cells, two dots appropriately oriented are equivalent to a solid bar stimulus as long as they

[3] In all figures, the dots sizes and sampling intervals quoted were accurate before photocopying

are close enough together to both fall within the receptive field of the cell. (This is certainly true of the convolution operators used in Stage I of the model presented in Section 2.3.1.) The effect of dot size and separation has been fully demonstrated and measured by Zucker and Davis (1985) with respect to the property of well-placed endpoints of continuous and dotted lines. I shall be concerned only with curves that are sampled on the dense side of this constraint.

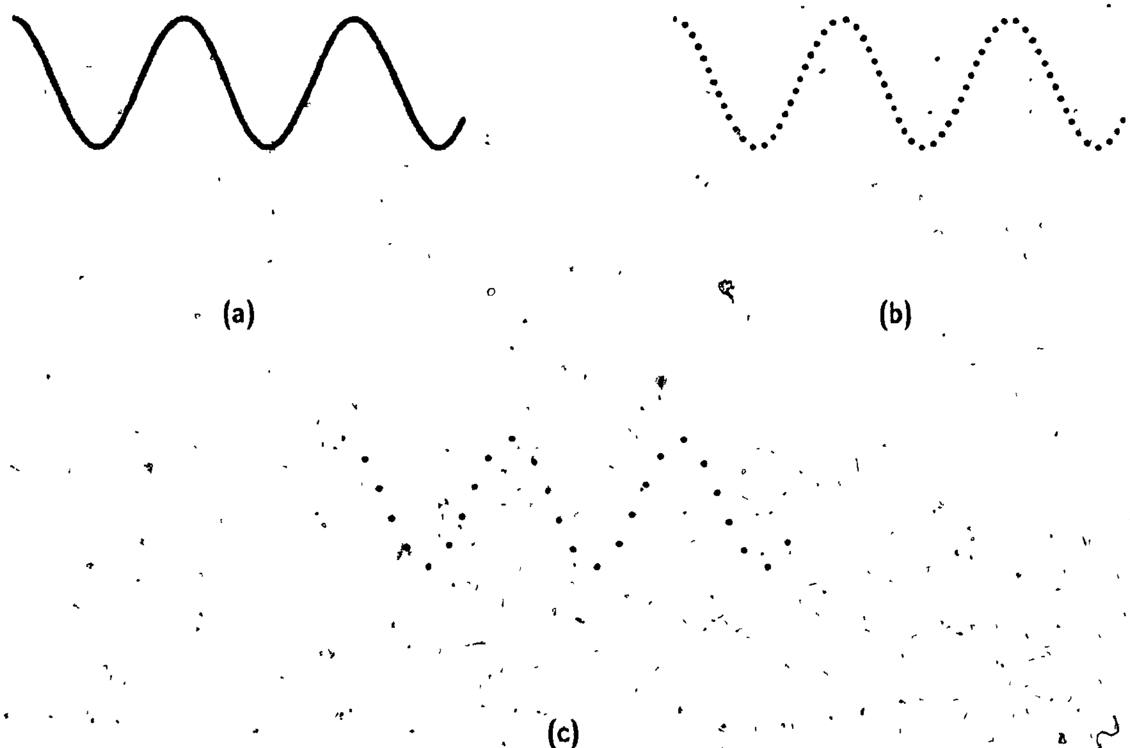

Figure 3.1 The effect of different sampling intervals on curve perception. In part (b), the sinusoid has been sampled on the dense side of the size/separation constraint (Figure 10 in Zucker and Davis 1985), and it looks the same as the continuous sinusoid in part (a). In part (c), however, the sinusoid has been sampled on the sparse side of the constraint, and the peaks appear to be triangulated. In all parts, the dot size is 2.9 minutes of visual angle (m.v.a.) when viewed from a distance of 1 meter (before printing). In part (b), the dot to space ratio is 1:1.5; in part (c) it is 1:4.

The second question, how the curve percept varies with the sampling phase, provides a means for studying the sensitivity of the curve reconstruction mechanism to local orientation structure. This sensitivity is the subject of this thesis. The following paragraphs demonstrate how changing the sampling phase, and thereby changing the geometry of the dot pattern, changes the percept. This leads to an experimental paradigm for studying human sensitivity to orientation discontinuities as a function of the local orientation structure.

3.1.1.2 Quantising a Curve Differently Changes the Geometry of the Pattern

To produce a discrete trace of a curve, we begin with some continuous curve, sample it at regular intervals, and then display only the sample points. When a curve has been sampled, we must distinguish between:

1. the underlying continuous curve, or virtual curve;
2. the discrete trace of the curve; and
3. the apparent curve.

Consider Figure 3.2. Notice that the curve in part (a) of the figure has an orientation discontinuity at P. In part (b) the curve appears to be smooth everywhere. Although one's attention may be drawn to the part of the curve near P, perceptually there is no break in the curve such as in part (a). In fact, parts (a) and (b) were constructed from the same underlying curve, but their quantisations – *within* the size/separation constraint – are phase shifted: the sampling interval is the same, but the sampling points are different. Changing the position of the dots along a virtual curve changed the geometry of the dot pattern. The portion of the model presented in Section 2.3.1.3 makes explicit how the detection of discontinuities is dependent on the local orientation structure. Specifically, the curvature (or higher-order derivatives), approximated by differences of orientation estimates, must be consistent across open neighbourhoods for a smooth interpretation to result; otherwise a discontinuity must be asserted. Changing the geometry of the dot pattern is equivalent to changing the structure (the local orientations and the spatial arrangement) of the orientation cues. As in the example, therefore, it further changes the apparent curve. Other sampling phase shifts would affect in varying degrees our ability to see the discontinuity in the curve.

Figure 3.2 A curve which has been quantised at slightly varying positions (but at the same rate) relative to the discontinuity at P (see part (a)). Note how the discontinuity appears sharp in (a) but smooth in (b). The dot size is 2.9 m.v.a. when viewed from a distance of 1 meter.
(a) no offset in quantisation from the discontinuity (seen at P).
(b) quantisation offset 0.7 dot diameters to the left of P.

3.1.2 A Corner and a Curve

By taking together the two observations discussed above, we know that we can use *differently sampled curves* to study how changes in the geometry affect the *apparent curve*. Specifically, we can use phase shifts in the sampling to affect the geometry of the pattern without changing the sampling interval (and therefore without changing any other possible characteristics of the inference). To obtain stimuli for studying this effect, consider two straight lines approaching a point, and allow for two cases. In the first case, let the lines continue along a straight path until they meet at the point, forming a sharp orientation discontinuity in the resulting curve: see Figure 3.3a. In the second case, let the lines curve (say, following a low-frequency sinusoid) to meet with the same orientation, forming a smooth curve; see Figure 3.3b. These two examples provide a means of assessing whether the *orientation discontinuity* – the corner – is detectable: it must not appear like the smooth, or control, curve. Notice that by sampling the test (discontinuous) curve with different phase shifts, we can change the orientation cues in the neighbourhood of the corner in a measurable way: see Figure 3.4. This then provides a paradigm for studying how neighbouring orientation cues interact in discontinuity detection. I shall use similar transformations and varying sampling phases in a psychophysical experiment, presented in Chapter 4, to measure human sensitivity to orientation discontinuities as a function of the local orientation structure.

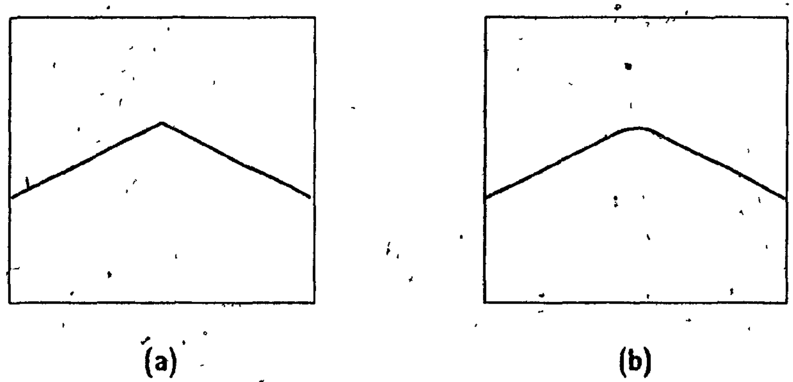

Figure 3.3 In (a), the curve is described by two straight lines joining at a point, where there is an angle created by the difference in orientation of the two lines. In (b), the curve is described by the two straight lines, with the same angle as in (a) between them, but this time the lines arc (with a low-frequency sinusoid) to join with no discontinuity in orientation.

In the following section, I shall show how we can use dotted stimuli to study the effect of the assumption of local parallelism on detecting orientation discontinuities in flow patterns. A new parameter of flows, the path-length of the flow cues, can be used to control the amount of local information about the flow derivatives that is present in the image. Our sensitivity to this parameter, therefore, will allow us to assess the order of the flow interpolation mechanism.

3.2 Discontinuity Detection in Flow Patterns

Recall that in flow reconstruction, a major obstacle to accurate discontinuity detection is the smoothing introduced by the assumption of local parallelism. Suppose, however, that the flow cues were long enough that the mechanism for detecting flows could extract not only orientation, but also higher-order flow information from each cue. Then in areas where

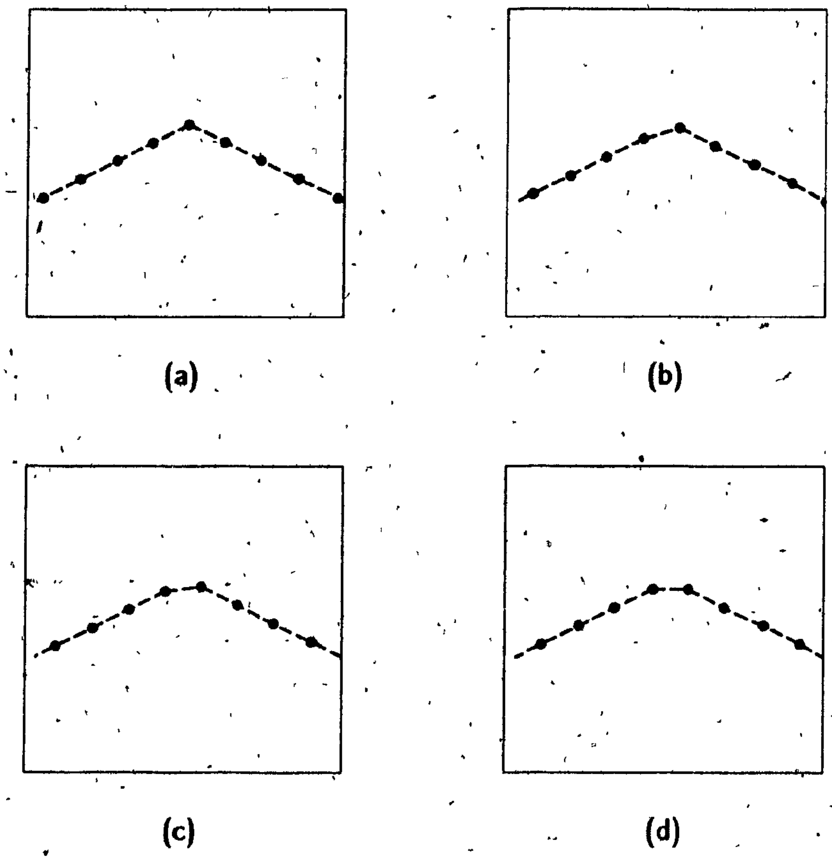

Figure 3.4 The effect of sampling Figure 3.3a with various phase shifts: (a) top dot offset = 0.0 sampling intervals; (b) offset = 0.2; (c) offset = 0.4; (d) offset = 0.5. Note the changes in the orientation cues, represented by the dashed line, brought about by the changing phase.

the flow is changing, we could predict how the assumption local parallelism will break down. Provided that this prediction was supported by both the orientation and the curvature cues in some neighbourhood, then orientation averaging could at least be partially inhibited. We can therefore see that sensitivity to discontinuities in flow patterns is dependent on two things:

1. the order of approximation that can be applied to each flow cue, as determined by its length.
2. the capabilities (or order) of the mechanism reconstructing the flow.

Dotted flow patterns allow us to explicitly restrict (under interpolation theory) the order of approximation that can be applied to each flow cue. Using dotted stimuli to study human sensitivity to orientation discontinuities in flow patterns as a function of this restriction will therefore allow us to infer the order of the mechanism for reconstructing flows in the human visual system.

The following paragraphs further motivate and elaborate on the effect of path-length on discontinuity detection. The use of random dot Moiré patterns, or dotted flow patterns, to study this effect is then described, resulting in a paradigm for the psychophysical measurement of human sensitivity to orientation discontinuities in flow patterns as a function of path-length.

3.2.1 Motivation: An Example from Art

To intuitively motivate the role of the path-length of individual flow cues in discontinuity detection, consider how artists convey three-dimensional shape using a two-dimensional medium, the canvas. They often use flow-like patterns, covering a region with several *short, roughly parallel* strokes to locally highlight what are in fact relatively global shape characteristics. Some of the clearest examples are the sketches by Leonardo da Vinci presented in Figure 3.5. He has used this technique to indicate both the roundness and the angle of the man's arm. The strokes fall along what might be folds in the cloth of the sleeve, which together take on the shape of the arm supporting the cloth. They are slightly curved with the roundness of the arm, and they either bend more sharply about the elbow or create a splayed pattern. Da Vinci has used the same technique to impart a sense of the flow of water around a post.

Note that the pen strokes in these figures are rarely exactly parallel, and even when they are, they rarely line up side-by-side or end-to-end. Rather, they are spaced to cover an entire two-dimensional region. The extent to which the strokes are parallel or do line up changes a qualitative impression of the turbulence or texture of the flow. Despite this, an impression of flow persists. The model for recovering flow described in Section 2.3.2

(a) (b)

Figure 3.5 Details from "Studies of an Old Man Seated and of Swirling Water," pen and ink sketch by Leonardo da Vinci. Housed at the Royal Windsor Library. No 12,579 *recto*. Reproduced from Popham (1945, no. 282). The full sketch contains notes by da Vinci which were translated by Richter (1939, no. 389) as follows: "Observe the motion of the surface of the water which resembles that of hair, which has two motions, of which one depends on the weight of the hair, the other on the direction of the curls: thus the water forms eddying whirlpools, one part of which is due to the impetus of the principal current and the other to the incidental motion and return flow."

takes into account this persistence of flow orientation over variable orientation cues and imposes an assumption of local parallelism to fill in the empty spaces using information from nearby orientation cues. The imposition of the assumption also results in a local averaging of orientation information so that small localised changes in orientation cues do not affect the recovery of flow. However, the assumption is by definition violated wherever the flow actually changes direction. Locations of high curvature change are only perceived, therefore, only when they also line up. Lines of orientation discontinuities that are detected in this way are assigned a special significance in the interpretation of the overall shape, such as boundaries between distinct regions of flow or distinct surfaces. But the changes in direction are often smoothed – they appear to occur over larger neighbourhoods (therefore with smaller curvature and changes in curvature) than is actually the case. In order to detect changes in orientation more accurately (more locally), more information is needed to allow the assumption of local parallelism to be relaxed.

Note that da Vinci has varied the length of the strokes to create different impressions. For the small folds in the cloth and the water turbulence, he has used short strokes.

While these short strokes give a well-defined impression of *constant overall shape*, there is a fluid impression that the *small details may change randomly*. The overall shape is perceived as being basically smooth. The small details that are lost are small crevices, or the places where the smoothness is disturbed. That is, high-curvature information, or more particularly high *changes* in curvature, are lost or made to appear random. While high curvatures can be perceived, such as the small circles in the turbulence, discontinuities or high changes in curvature are not always correctly interpreted. (Note that a discontinuity can be considered to represent both infinite curvature and an infinite change in curvature.) Around the bend of the man's elbow and in the smooth arcs of water, however, da Vinci has used mostly long, well-defined strokes. In these regions of longer strokes, the flow appears to be less variable, and changes in the flow orientation – particularly abrupt ones – are more readily perceived. It would seem that the longer strokes provide the extra information required to relax the local parallelism assumption.

In summary, two points emerge: First, da Vinci and others have observed that locally parallel structure in the world often maps onto locally parallel structure in images. But for natural flow patterns such as water, hair, or the folds of clothing, the structures are only roughly locally parallel. If made precisely so, then the character of the flow changes; see Figure 3.6. Second, within these roughly parallel flow patterns, the length of the pen strokes affects the perceived spatial (and orientation) precision. Longer pen strokes lead to more precision around bends and curves, while shorter strokes lead to less.

3.2.2 Path-Length as a Parameter of Orientation Information

The length of the pen strokes in da Vinci's drawings is analogous to the *path-length* of orientation cues in real images. Preliminary experiments indicate that the loss of curvature change information observed at some path-lengths takes two forms. First, corners which correspond to small changes in orientation may be blurred to give the impression of a smoothly bending flow, resulting in the possible loss of boundary information. Second, a smooth flow with a large and relatively local change in orientation that results in high

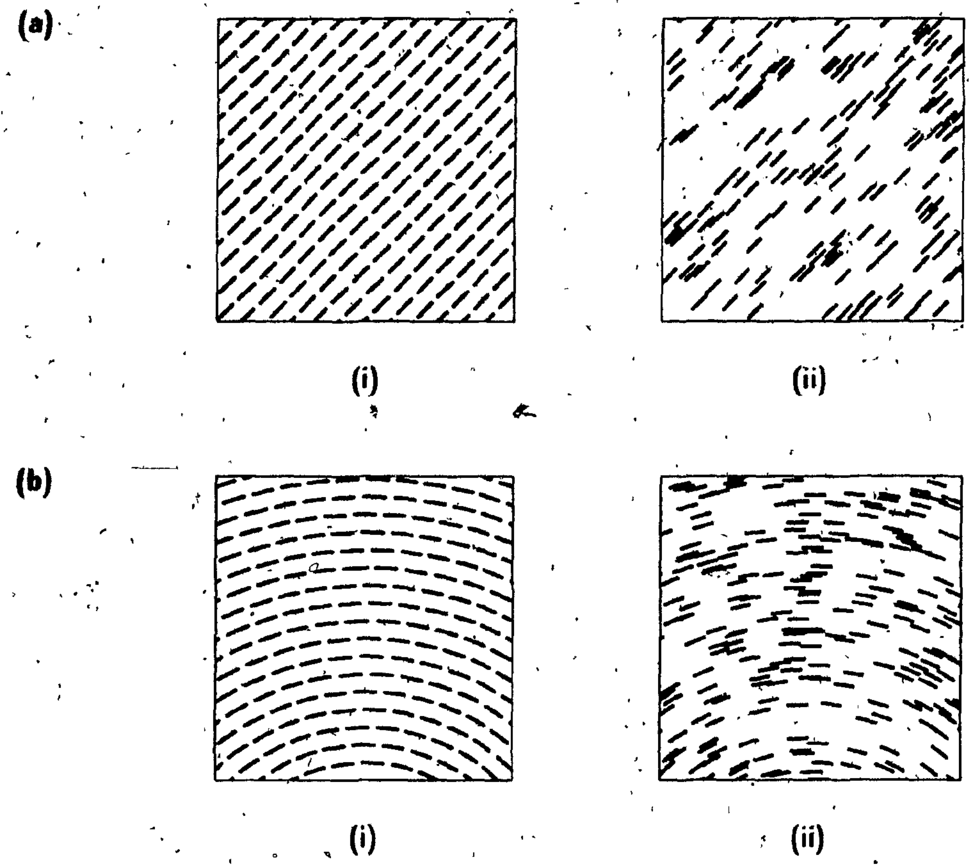

Figure 3.6 The difference between (i) pin-stripe flow patterns with orientation cues lining up exactly end-to-end and (ii) those with (the same number of) randomly positioned orientation cues. Part (a) is a linear flow (generated by pure translation), and part (b) is a section of a circular flow (generated by pure rotation). This thesis is concerned with "random" flow patterns such as those in (ii).

curvature change may be misinterpreted as having a discontinuity, resulting in the possible perception of a surface boundary where there is in fact none. These situations are illustrated in Figure 3.7, and the effect of path-length on these perceptions is illustrated in Figure 3.8.

Note that increasing the path-length has a greater effect on curvature perception than increasing the density of orientation cues; compare Figures 3.8 and 3.9. Attneave (1954; 1959) and Klemmer and Frick (1953) have shown that the amount of information transferred in an image is dependent both on the number of informational parameters and on the amount of information carried within each parameter. Specifically; they conjecture that information within a parameter is in some sense additive. In contrast, the information across parameters is multiplicative because it is combinatorial. It is my assertion that increasing the density of the orientation cues adds orientation information within the existing parameter set, while the path-length adds a new parameter – the correlation of orientation cues. The model described in Section 2.3.2 reflects this parameterisation.

As was explained in Section 2.3.2, curvature information – as long as it receives support over a neighbourhood – allows us to locally predict the way in which the flow will change and the assumption of local parallelism will break down. It thus allows this assumption to be relaxed, and the orientation information is no longer averaged over as large a neighbourhood. Therefore, we would expect that our sensitivity to discontinuities in the flow would increase when curvature information is available. Path-length is important because, as the experimental results presented in Chapter 5 shall show, short path-length cues only provide local information about orientation, while longer ones provide information about curvature as well. The following section demonstrates how the derivative information provided by different path-lengths can be controlled with dotted stimuli. This leads the way to a paradigm for investigating the loss of sensitivity to discontinuities as a function of path-length when the assumption of local parallelism is imposed. I shall concentrate on the smoothing of discontinuities corresponding to small changes in orientation.

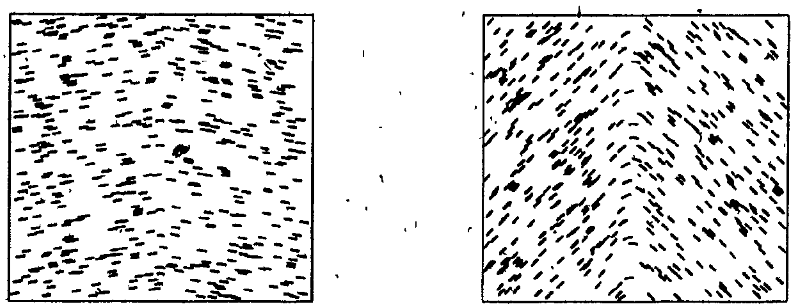

(a) (b)

Figure 3.7 In (a), the flow pattern has a discontinuity running down the centre that is not visible without careful inspection. In (b), the flow pattern has no discontinuity, although at first glance it appears to have one running down the centre.

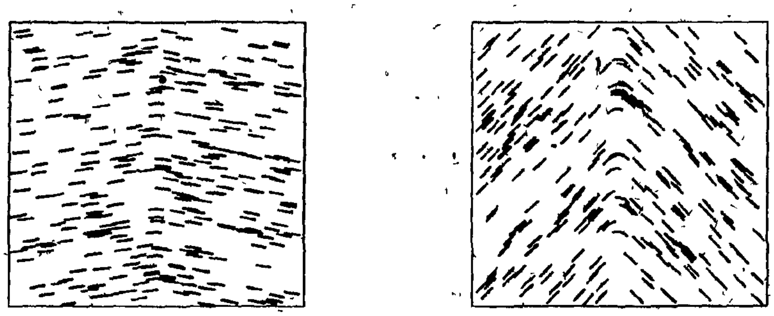

(a) (b)

Figure 3.8 . The same flow functions were used to generate these figures as those in Figure 3.7. but the path-length is twice as long. Note that the discontinuity in (a) is now visible at a glance, and there is no false discontinuity in (b).

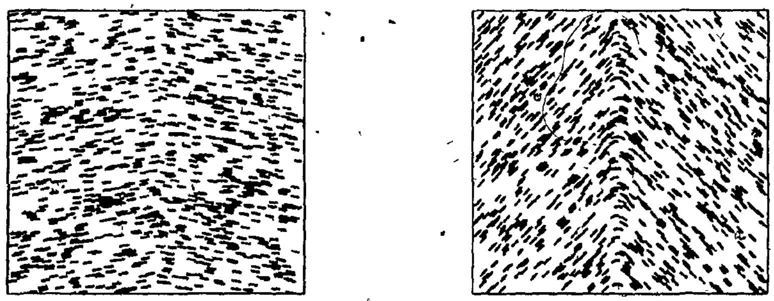

(a) (b)

Figure 3.9 The same flow functions were used to generate these figures as those in Figure 3.7. but the density of pen strokes is twice as high. Note that although the true flow is more easily reconstructed, the effect of doubling the density is not as pronounced as doubling the path-length; see Figure 3 8.

3.2.3 Random Dot Moiré Patterns Mimic Natural Flows

The flow patterns that I shall use to study path-length are called *Glass patterns* or *random dot Moiré patterns* (Glass 1969; Glass and Pérez 1973). They are illustrated in Figure 3.10 and are constructed as follows:

OVERLAY 1:

Construct a field of randomly distributed dots.

OVERLAY n:

i) Make a copy of overlay $n-1$.

ii) Move each dot in the copy according to a chosen flow transformation, for example a rotation or a translation.

iii) Superimpose this overlay on the other overlays.

Only two overlays ($n = 2$) are required to produce an impression of flow.

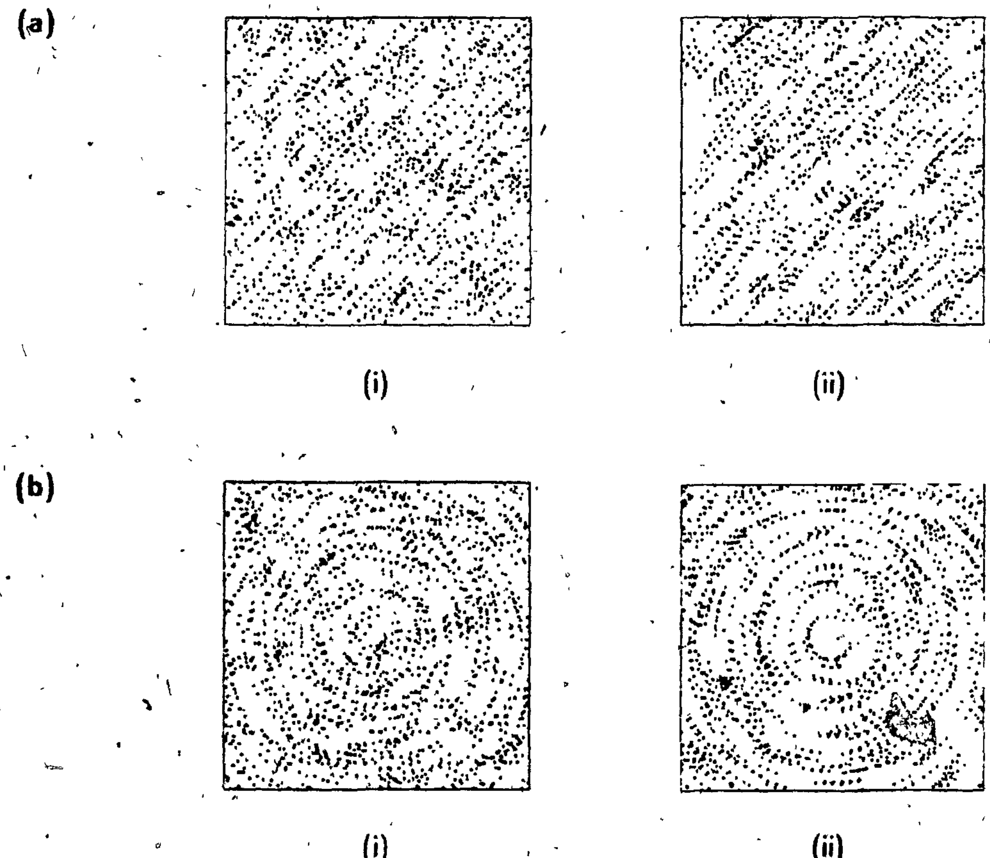

Figure 3.10 Sample random dot Moiré patterns: (a) was created by translating each dot in the previous overlay at an oblique angle; (b) was created by rotating each dot in the previous overlay about the centre of the figure. Part (i) in each has two overlays, part (ii) has three overlays. All parts have the same overall density, and in all parts the dots were moved a distance of three dot diameters (1 dot : 2 spaces).

Random dot Moiré patterns create flows rich in locally parallel structure as follows. Each dot in the original pattern can be traced through each overlay. The resulting set of corresponding dots – the original plus its transformed copies – can be considered as a randomly distributed pen stroke along one of the curves described by the transformation used in step (ii). Define the number of overlays used, n, as the *path-length metric* or *path metric* of each stroke. The path metric times the dot separation gives the actual path-length of a stroke. The patterns used in the experiment all have dot separations that are small enough to guarantee that there is no functional difference between the solid pen strokes and the dotted contours.[4] The path-length metric provides a direct and meaningful measure of what orientation information is available locally and therefore allows us to determine how changes in orientation could ideally be processed at different path-lengths. It is for this reason that random dot Moiré patterns are used in this study.

3.2.4 Corners and Curves in Flow Patterns

Consider again the paradigm set up to study discontinuity sensitivity in curves. Stimuli were constructed from two straight lines with some difference in orientation and joined with either an abrupt or a smooth transition between the two orientations; see Figure 3.3. This provided us with a means of assessing whether the "corner" is detectable: it must not appear like the smooth, or control, curve. By replicating these transformations along the vertical axis, we can produce flow patterns that provide a similar paradigm for studying sensitivity to discontinuities in flow patterns; see Figure 3.11. As was discussed previously, one would expect the angle between the lines to be significant, particularly when only orientation information is available for reconstructing the flow. I shall therefore use such

[4] This functional equivalence between dotted and continuous contours refers to the size/separation constraint measured by Zucker and Davis (1985); see Section 3.1.1 for elaboration. Since a similar mechanism has been proposed to extract orientation cues from random dot Moiré patterns (note that flow patterns with high path-length metrics are collections of one-dimensional contours), I shall assume that staying on the dense side of this constraint ensures that the consequences of changes in the path-length correspond to the consequences of similar changes in the path metric while the dot separation remains constant.

transformations, varying the angle and path metric, to measure (psychophysically) human sensitivity to orientation discontinuities in flow patterns as a function of the path-length of flow cues. The experiment and its results are presented in Chapter 5.

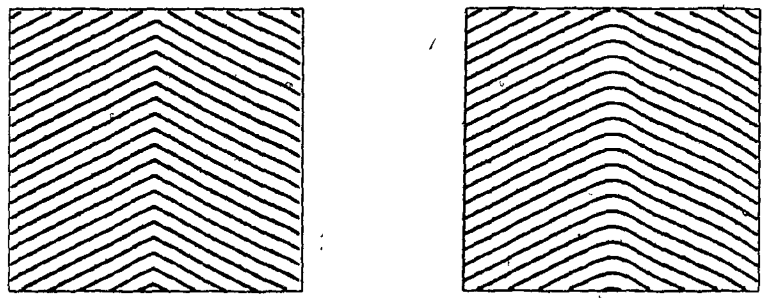

(a) (b)

Figure 3.11 The flow transformations described by these figures correspond to the same transformations that were used to construct the curves in Figure 3.3. For examples of the actual flow patterns used in the experiment, see Figure 5 1

The remainder of the thesis is devoted to the experiments that were set up in this chapter to study human sensitivity to orientation discontinuities in curves and in flows. Because the experimental paradigms were derived from predictions that came directly from the orientation selection model and its assumptions, the results will reflect properties of the model. Specifically, they will provide an indication of the order of approximation used by the visual system for detecting orientation discontinuities during curve and flow reconstruction. The experiment relating to curve reconstruction is presented first, followed by the experiment relating to flows.

Chapter 4 Experiment: Corner Sensitivity vs. Sampling Offset in Curvilinear Dot Grouping

The purpose of this experiment was to determine how the changes in orientation structure which result from changes in the sampling of a curve affect the sensitivity of human observers to orientation discontinuities in the curve. It was expected that this would shed light on the order of the mechanism used to detect and reconstruct these curves.

Subjects were shown several instances of dotted contours resembling those in Figure 3.3. These contours consist of two straight lines approaching a point. Some of the contours were constructed with the lines meeting discontinuously in orientation at the point (such as in Figure 3.3a), and some were constructed to meet with a smooth transition between the two orientations (such as in Figure 3.3b). These two transformations were used to assess the subjects' ability to detect the discontinuity: a discontinuous curve must not look like the smooth, or control, curve. All the contours were constructed using the same sampling interval but with different phase shifts. The phase was determined by the offset of the "top dot" (the one closest to the peak in curvature) from the discontinuity or the peak in curvature of the underlying curve and was measured in terms of the sampling interval. Note that for smooth underlying curves, the dot placement should not affect the percept as long as the curvature is within the sensitivity range of the process. Other factors which might affect the sensitivity to discontinuities, such as dot size, dot separation, and noise, were held constant.[5]

---

[5] The dot size and spacing were within the size/separation constraint measured by Zucker and Davis

4.1 Method

4.1.1 Subjects

There were five subjects, four male and one female, all with normal or corrected vision. Three subjects were aware of the goals of the experiment, two having participated in preliminary experiments. All subjects received training sets to become familiar with the task, and all subjects were presented with the full range of images.

4.1.2 Apparatus

The stimulus images were generated using a DEC VAX 11/780 computer and displayed on an AED-767 colour graphics monitor. The experiments were conducted in a dimly lit room. After the subjects had adapted to the illumination level for a few minutes, the luminance of the display screen was adjusted so that the dots did not appear to the subject to be self-luminous. The subjects were seated 3.6 meters from the monitor (looking directly at it), although they were permitted slight movement.

4.1.3 Procedure

The experiment consisted of four sessions, during which the subject viewed 120 images. These images were dotted contours as described above, with eight quantisations of orientation change between the straight-line portions of the contours. During each session, only four of these quantisations were used. A randomly mixed but equal number of smooth and discontinuous curves were shown. For each underlying curve, every sampling offset was used. (The "top dot" was offset from the peak curvature position by 0.0, 0.1, 0.2, 0.3, 0.4, and 0.5 sampling intervals.) The images were displayed for 1.5 seconds, and the screen was blank during the inter-stimulus interval, which was not timed. Subjects were

(1985).

allowed to view the image for one additional 1.5-second interval if they were unable to focus the first presentation. After this, the subject was required to indicate the nature of the underlying curve with one of three choices: (1) a smooth change in orientation; (2) a single abrupt change in orientation; or (3) ambiguous.

The subjects were shown sample images such as those in Figure 4.1 for each of the three possible choices at the beginning of each session. These defining images were followed by a short training set (eight images) selected from the images to be presented during the session. Each training set image was first flashed for 1.5 seconds and, following a brief interval, was displayed again for the subject to analyse. This was the only reinforcement used during either the training set or the formal experiment.

4.1.4 Stimuli

The stimuli consisted of dotted one-dimensional contours that were constructed from two straight lines with some difference in orientation and joined with either an abrupt or a smooth transition between the two orientations. The smooth transition was accomplished by "capping" the straight lines with a sinusoid of the same *average slope*.[6] (See the Appendix for the precise equations used to generate the curves.) Eight quantisations of orientation changes were used, ranging from 2.3° to 33.4°. Dot size and sampling interval were constant across all images, with dot size 2.4 minutes of visual angle (m.v.a.) and a sampling interval of 10.3 m.v.a. (a ratio of dot size to intervening space of 1:3.3). This size and spacing were chosen to minimise the quantisation error of the display while staying within the size/separation constraint for early contour inferencing processes (see Zucker and Davis 1985).[7]

---

6 This method of construction for the smooth contours resulted in only one pattern of curvature change for any one orientation change. However, the subjects were not instructed that this was the case. Rather, they were given, as much as possible, the impression that any pattern of smooth orientation changes might be present, as well as other types of discontinuities.

7 Although this combination of size and spacing appears to fall on the borderline of Zucker and Davis's

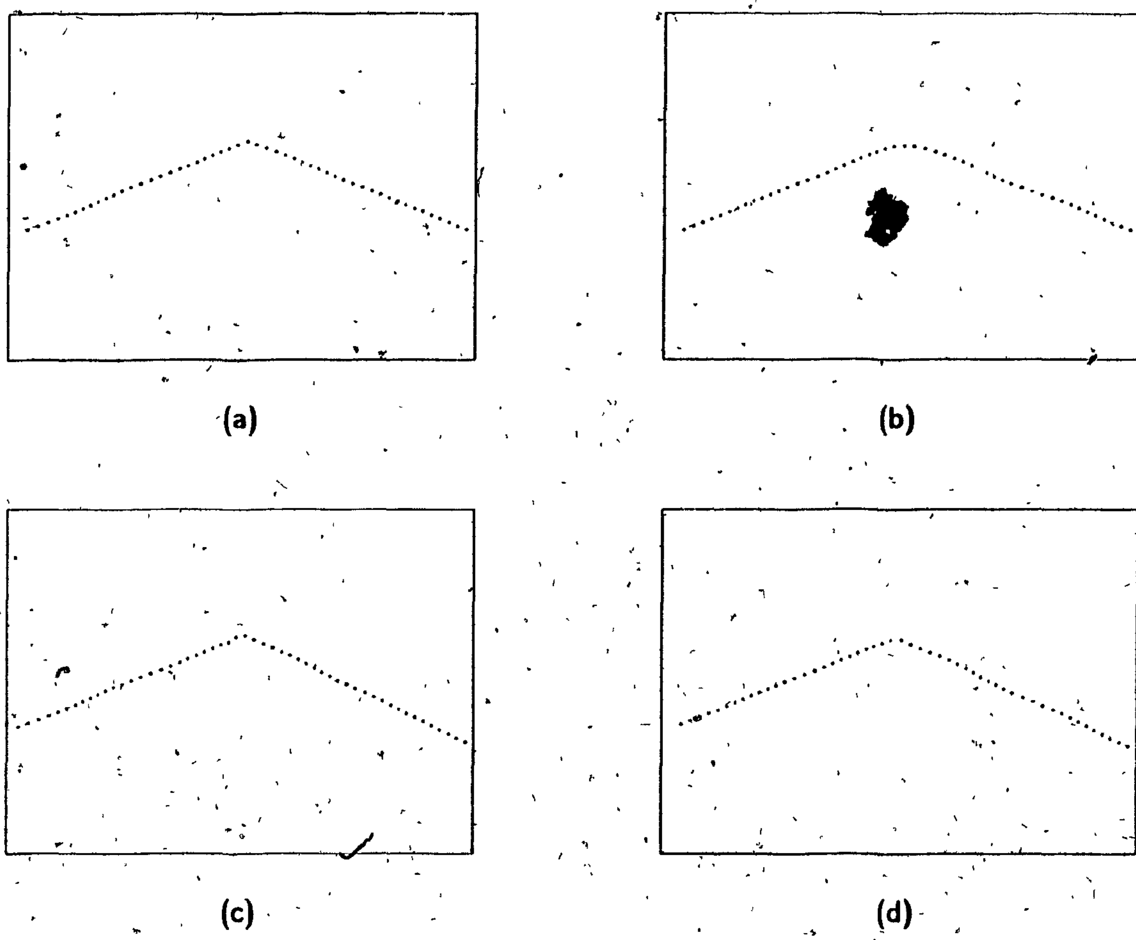

Figure 4.1 Images that were shown to subjects before the experiment, as examples expected to fit into each of the three possible response categories. Part (a) is an angular curve; part (b) is smooth. For part (c), if the subject noticed that the dot directly to the left of the "top dot" appears to be displaced from the rest of the curve, he was instructed to choose the ambiguous response, whether the curve otherwise appeared to be smooth or angular. The subject was also instructed to choose the ambiguous response for part (d) if he saw the curve as bending smoothly from the left, but then changing orientation discontinuously at the top dot. The angular response category was reserved, then, for curves made up of two straight lines with a single orientation discontinuity at their juncture. The above images have a dot size of 2.4 m v a and a sampling interval of 7.7 m.v.a. (a dot to space ratio of 1:2 2) when viewed from about 0.6 meters.

Each of the resulting 16 curves (8 orientation changes, 2 patterns each) was sampled six times, with the sample positions ranging from being symmetrically distributed about the peak curvature position with a dot centred on the peak (a "top dot" offset of 0.0 sampling

size/separation constraint (see their Figure 10), independent experiments using exactly this combination confirmed that it does lie within the constraint and is suitable for our purposes.

intervals), to being symmetrically distributed about the peak with no dot on the peak (a "top dot" offset of 0.5 sampling intervals). Figures 4.2 and 4.3 show sample images for the angular and smooth curves respectively, with zero, intermediate, and full offset of the top dot. Since the underlying curves were symmetric, offset to the left or to the right was expected to have no effect on the sensitivity to the discontinuities. However, left and right offsets were alternated in an attempt to average out any preferences on the part of the subjects that could have affected the results.

To prevent the possibility of subjects "memorising" the discontinuity position and using this as a cue for discontinuities, each of these 96 images (16 curves, 6 sampling patterns) was displayed with the peak randomly chosen to lie anywhere within 10 m.v.a. of the centre of the screen. Each image was then presented five times over two sessions (three times in one session and two in the other), totalling 480 images presented over four days.

4.2 Results

The results for angular curves are tabulated for each offset in Figure 4.4. Shown for each offset are the average percentage of discontinuities that were correctly reported, as a function of the orientation change. (As expected, the curves having underlying smooth orientation changes were correctly perceived as being smooth over all presentations.) For offsets of 0.0 and 0.1 sampling intervals, small orientation changes were seen as smooth, while for large orientation changes, the discontinuity was usually detected. The results for offset 0.2 follow the same pattern, but with higher variance in the responses and somewhat less overall sensitivity to discontinuities. At offset 0.3, this pattern of responses is broken, and the results are fairly ambiguous. (In fact, most of the "ambiguous" responses were given for discontinuous curves with offsets 0.2 and 0.3.) Finally, the curves with offsets 0.4 and 0.5 were predominantly seen as being smooth for all changes in orientation.

The results for offset 0.0 can be examined to determine an "absolute threshold" for accurately perceiving discontinuities; see Figure 4.4a. In this case, the subjects on average

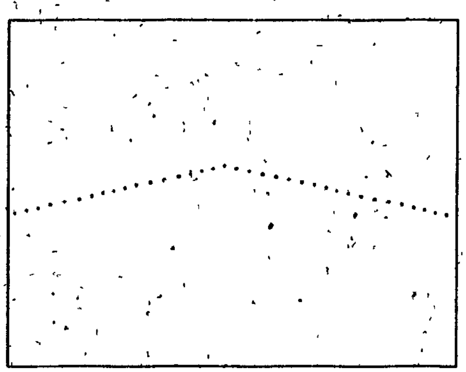

(a) Offset 0.0

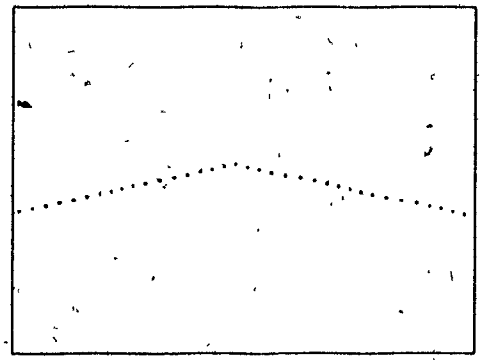

(b) Offset 0.1

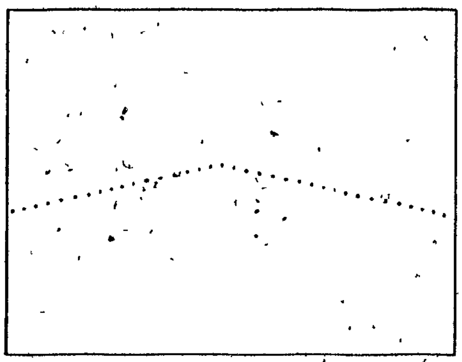

(c) Offset 0.2

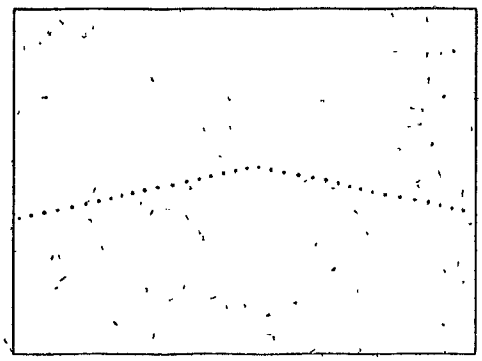

(d) Offset 0.3

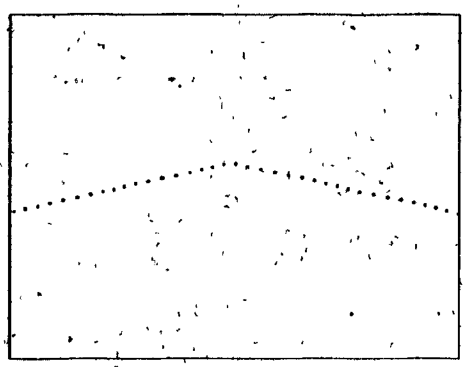

(e) Offset 0.4

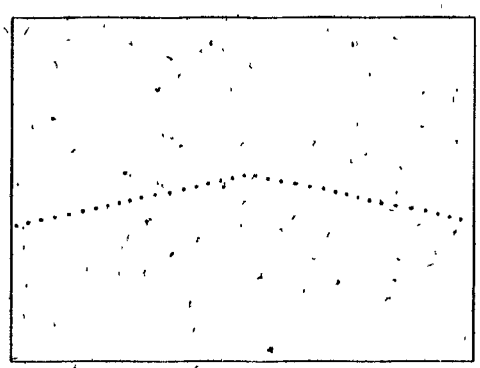

(f) Offset 0.5

Figure 4.2 Sample images from the experiment for angular curves Shown are all possible offsets for the curve with underlying orientation change of 24.8°. The dots in these images are 2 4 m.v.a. (the same size as in the experiment) at a viewing distance of approximately 0.6 meters.

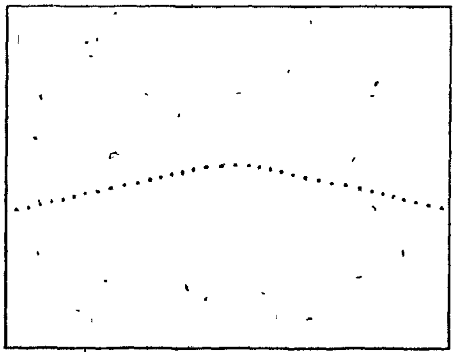

(a) Offset 0.0

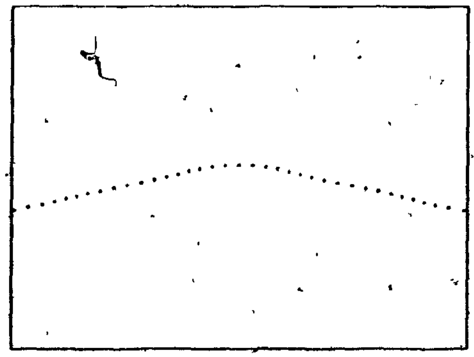

(b) Offset 0.1

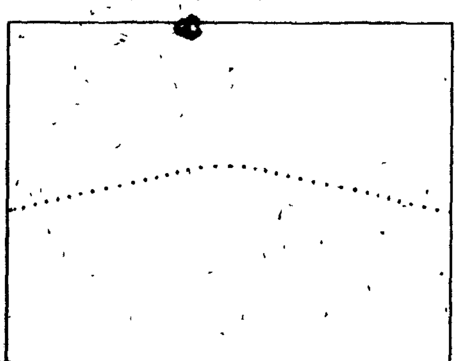

(c) Offset 0.2

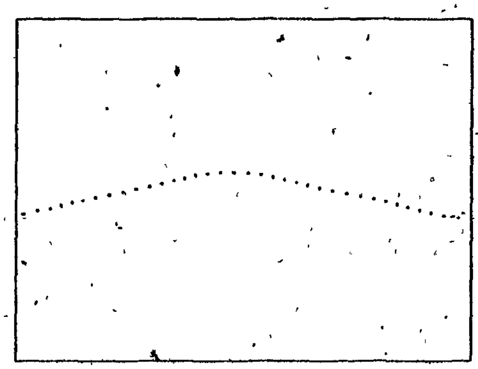

(d) Offset 0.3

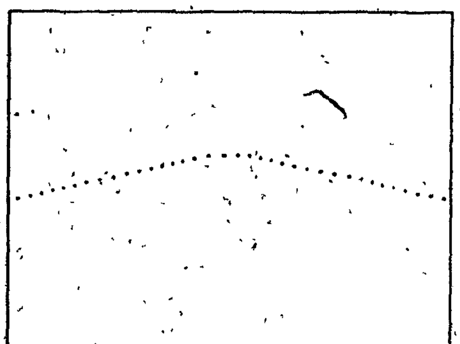

(e) Offset 0.4

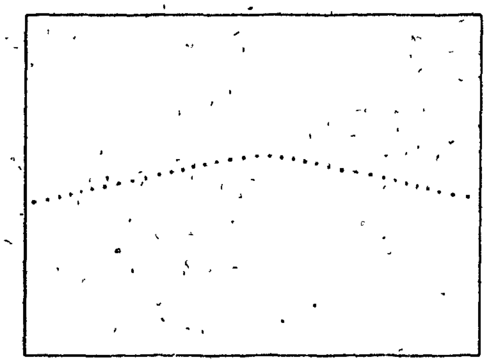

(f) Offset 0.5

Figure 4.3 Sample images from the experiment for smooth curves. Shown are all possible offsets for the curve with underlying orientation change of 24.8°, the corresponding curves to those shown in Figure 4.2. The dots in these images are 2.4 m.v.a.. (the same size as in the experiment) at a viewing distance of approximately 0.6 meters.

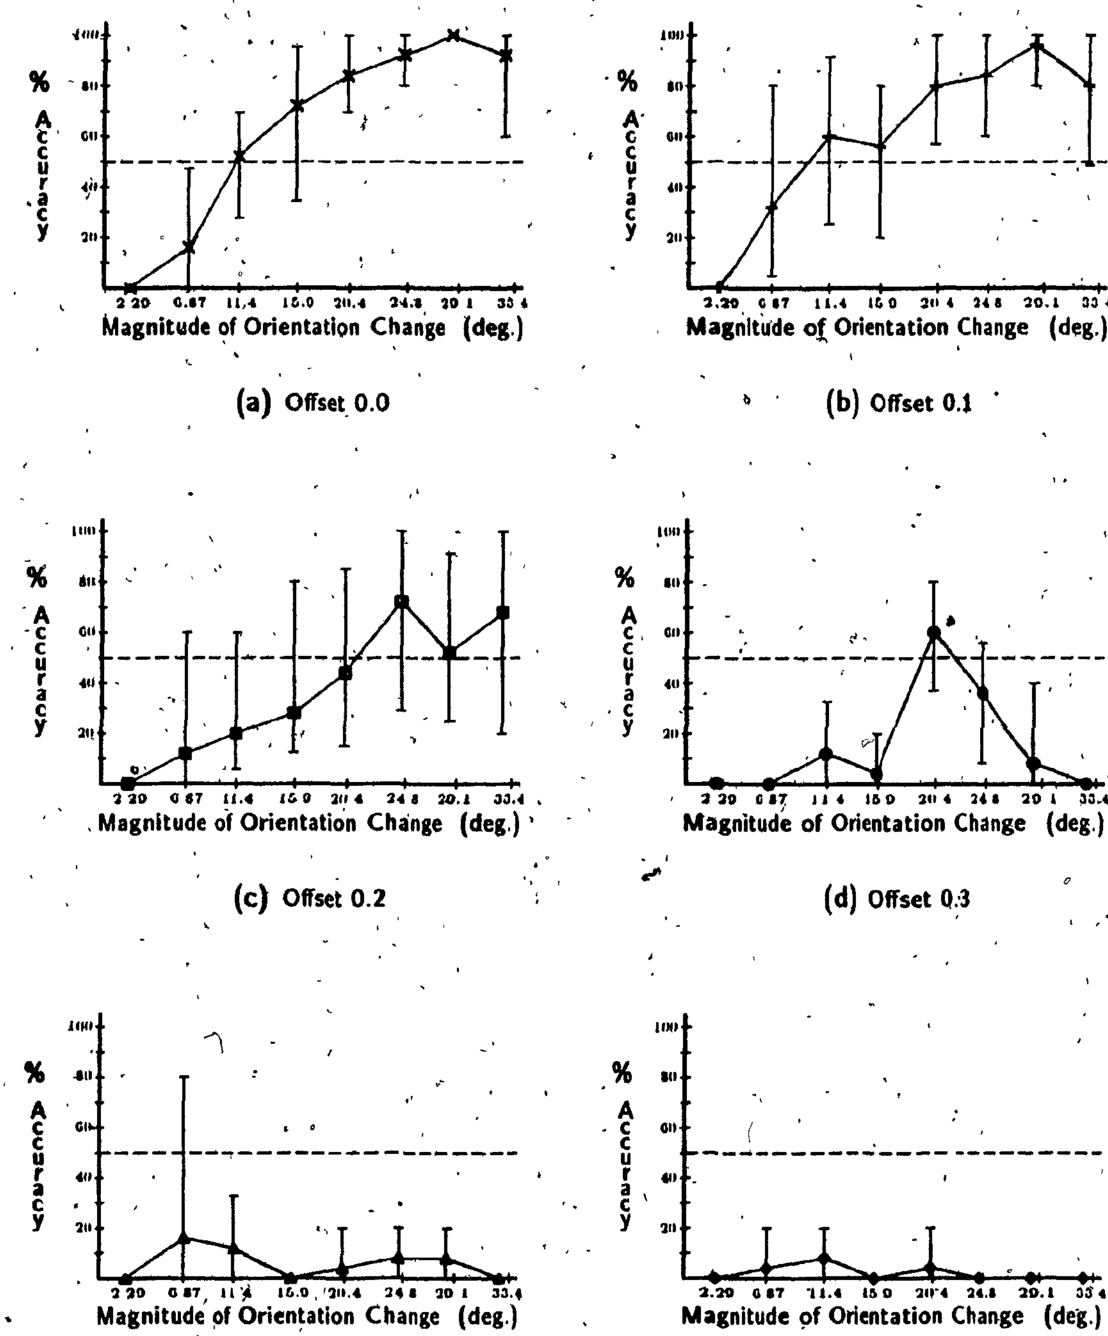

(a) Offset 0.0

(b) Offset 0.1

(c) Offset 0.2

(d) Offset 0.3

(e) Offset 0 4

(f) Offset 0 5

Figure 4.4 The results by offset phase shift (measured in units of the sampling interval) of the experiment. Illustrated are the mean percentages across all subjects (and one standard deviation) of sampled curves with discontinuities that were correctly identified, as a function of the magnitude of the orientation change at the corner. (All smooth patterns were correctly identified) Notice that there are two patterns of results. one for offsets 0 0. 0 1. and 0 2. and the other for offsets 0 3. 0.4. and 0.5

reported seeing the discontinuity with an accuracy greater than 50% for an orientation change of 11.4°. This is the same threshold that will be reported in Chapter 5 for sensitivity to discontinuities in flow patterns when curvature information is explicit in the image (Link and Zucker 1985*a*). It is also of the same order of magnitude as the quantisations of dominant orientations in simple cells found by Hubel and Wiesel (1962, 1977), and of orientation-selective channels found by Movshon and Blakemore (1973). However, for offset 0.5, when the underlying angle is 24.8° or higher, two adjacent angles each greater than 11.4° are formed by the dots surrounding the discontinuity. In this case (when two such angles are adjacent), no discontinuity was detected; see Figure 4.4f.

To demonstrate that these results are not tied to the underlying symmetry in the images, see Figure 4.5 for a parallel example using a non-symmetric discontinuity.

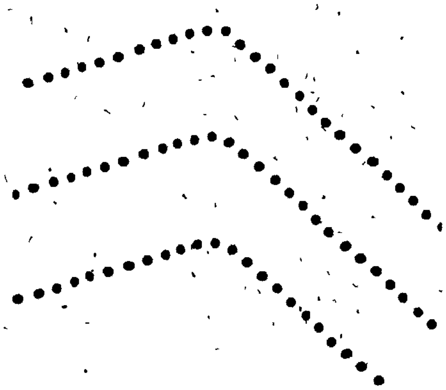

Figure 4.5 An example of a curve with an asymmetrical orientation discontinuity, that has been sampled with different phase shifts. The same change in percept, from discontinuous to smooth, can be seen when the offset varies from 0.0. In part (a), the offset is 0.0; in part (b) it is 0.5; in part (c) it is 0.2. The dots are approximately 3 m.v.a. when the figure is viewed from 2 meters, and the dot to space ratio is about 1:1.5

4.3 Theoretical Discussion: Discontinuity Sensitivity as a Function of Curvature

Recall that the model for curve reconstruction is comprised of two stages. The first stage is the construction of a tangent field – or the first derivative with respect to arc length of the curve – through:

1. convolution against simple-cell-like operators to obtain initial estimates of orientation;
2. interactions between the convolutions to estimate curvature, fill in gaps, and to eliminate noise and parameterisation effects.

The second stage is an interpolation, in which discontinuities are detected, and smooth curves are fit to the sample points in between the discontinuities. The discontinuities are defined explicitly in terms of the curve's derivatives given the order of the process, or the number of derivatives that can be represented. Discontinuities must be asserted at locations where the curve has non-zero derivatives of a higher order than can be represented by the system, or where the derivatives that are represented in the neighbourhood of a point cannot be used to predict the orientation estimate derived during the first stage; see Section 2.3.1.3. In this way, the segments in between the discontinuities can be interpolated by polynomials of a degree that can be represented by the system.

The model reduces to a physiological one as networks of simple cells (for accomplishing Stage I), and it provides the mechanistic background to interpret the experimental results. In particular, it permits the following predictions with respect to the reconstruction of a curve from a discrete, quantised trace:

change in geometry (by quantisation phase shift)	$\Rightarrow$	change in simple cell convolutions (Stage I, Step 1)
		$\Downarrow$
change in percept (Stage II and higher)	$\Leftarrow$	change in simple cell interactions (Stage I, Step 2)

How much change in geometry is necessary to change the resulting percept? Consider a single receptive field and a pair of dots in its excitatory centre. To a first approximation, one would expect that shifting one dot in the pair from the excitatory to the inhibitory region would produce a significant change in the overall pattern of responses: see Figure 4.6. Given the relative dot separation used in the experiment, a rectangular-shaped receptive field with an aspect ratio of 3:1 would require a change in orientation on the order of about 19°. For the same dot size and a covering receptive field with aspect ratio of 5:1, the required change in orientation would be on the order of 14°. Thus, depending on specifics, a shift of 10° to 20° is necessary to move a dot from the excitatory centre to the inhibitory surround. The results for offset 0.0 show that this is not an unreasonable prediction. However, to analyse the results for other offsets, we must replot them as a function of the angles formed by the dots near the discontinuity.

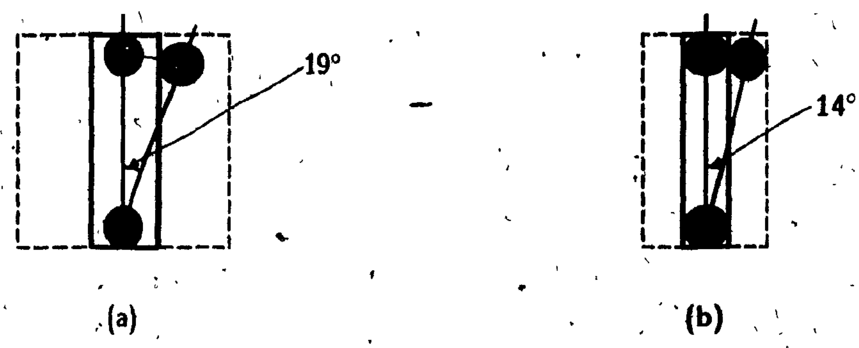

Figure 4.6 A first approximation to the shift in orientation that would be required to significantly change individual simple cell responses, thereby changing the percept. Shown are rectangular approximations to a simple cell's receptive field that will cover a dot pair with the dot:space ratio of 1:3.3. The thick solid line outlines the excitatory part of the receptive fields, and the dotted line outlines the inhibitory parts. Both fields have been constructed to cover a pair with the same size dots. In part (a), the aspect ratio of the excitatory field is 3:1, and the rotation required to shift one dot into the inhibitory part of the field is 19°. In (b), the aspect ratio is 5:1, and the angle of rotation is 14°.

Two angles of interest emerge near the discontinuity when the sampling of the angular curves is phase shifted. I shall call these angles α and β; see Figure 4.7. $\beta + \alpha$ is the primary angle between the lines, where α is formed as a result of the offset. Note $\alpha = 0°$, and β is the primary angle, when the offset is 0.0 for lines meeting at any angle. If curvature estimates are formed by differencing two neighbouring orientation estimates (see Section 2.3.1.3), then we might expect the sensitivity to discontinuities in the experiment to show some relation to these two angles. In Figure 4.8, the experimental results are replotted as a function of $\beta - \alpha$, for different values of β. Note in particular that for β between 15° and 25° (just over the threshold for discontinuity detection given 0.0 offset), $\beta - \alpha$ must be at least 15° to 20° for discontinuities to be detected. That is, α must be less than 10° (note the initial estimate for producing a significant change in the response pattern) *and* β must be above the absolute threshold (i.e., the threshold given $\alpha = 0°$) for discontinuity detection. In fact, this result is verified by Chi-square tests which show both β and $\beta - \alpha$ to have significant input to the responses (at the 0.995 and 0.999 levels respectively). For β in the range 10° to 25°, however, β no longer has a significant contribution (at the 0.062 level), but the contribution of $\beta - \alpha$ remains high. When the results are structured in this way, the response variance across subjects was insignificant (at the 0.0 level overall, and at the 0.198 level for β between 10° and 25°).

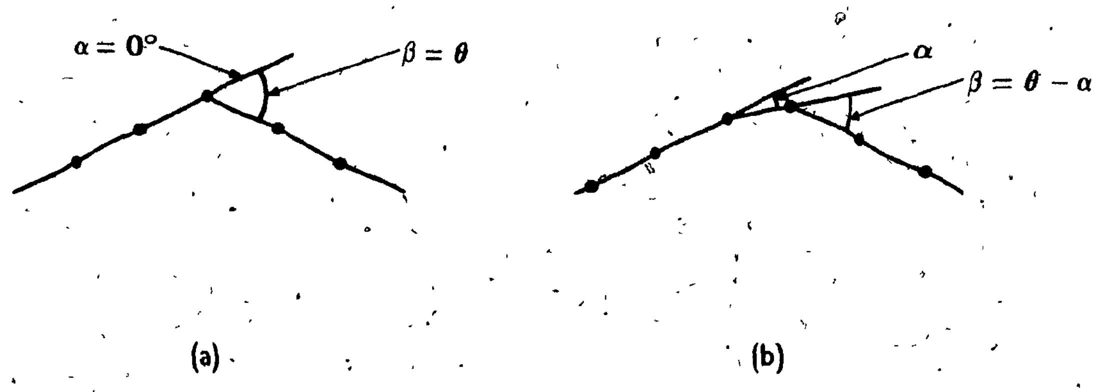

Figure 4.7 The two angles, α and β, that are created when an angular curve from the experiment is sampled with an offset other that 0.0. Illustrated in part (a) is the degenerate case when the offset is 0.0, $\alpha = 0°$, and β is equal to the primary orientation change θ. Illustrated in part (b) is an arbitrary sampling offset to the right of the discontinuity, causing α to increase and β to decrease in magnitude. For offset 0.5 (not illustrated here), $\alpha = \beta$.

As a practical illustration of how α and β would affect the detection of discontinuities, consider the change in the response patterns of the model's operators that would occur as these angles are varied. In Figure 4.9, a close-up is shown of the corner of an angular curve with $\beta = 24.8°$ at offset 0.0, and the resulting curves when the sampling is offset by 0.1 and 0.3. Overlaying the curves are operator receptive fields, quantised by 10° in

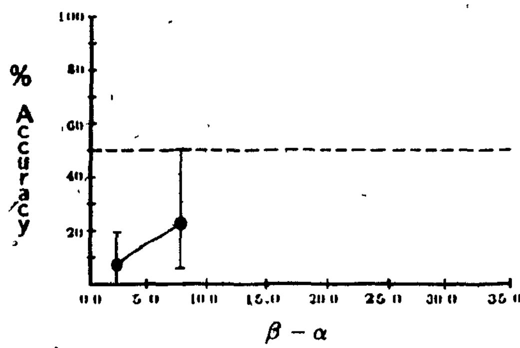

(a) $5^\circ < \beta < 10^\circ$

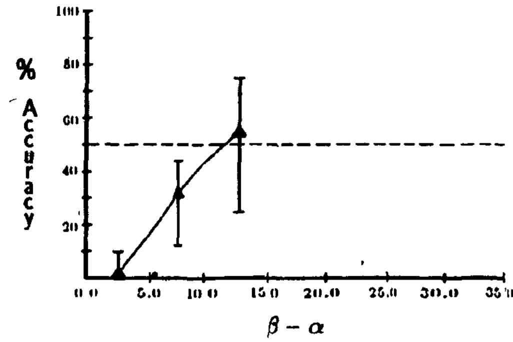

(b) $10^\circ < \beta < 15^\circ$

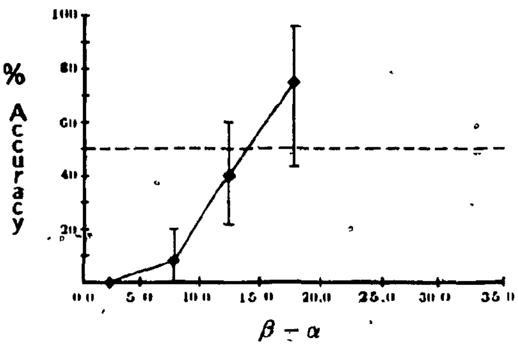

(c) $15^\circ < \beta < 20^\circ$

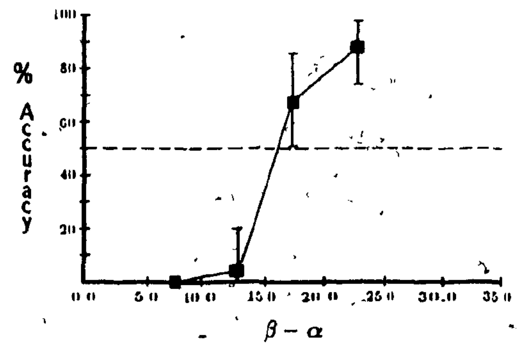

(d) $20^\circ < \beta < 25^\circ$

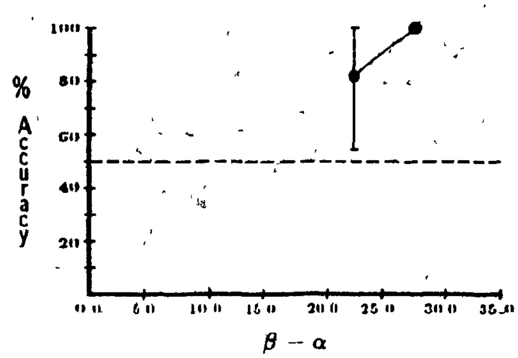

(e) $25^\circ < \beta < 30^\circ$

Figure 4.8 The results of the experiment by $\beta - \alpha$. Illustrated are the mean percentages across all subjects (and one standard deviation) of sampled curves with a discontinuity that were correctly identified, as a function of $\beta - \alpha$ for each value of β (each in 5° blocks) Notice that parts (c) and (d) in particular show a clear increase in sensitivity to the discontinuity with increasing $\beta - \alpha$

orientation, that would be among those responding optimally to these dot patterns. Note that in parts (a) and (b) of the figure, these operators maintain a constant orientation except at one location, where a large change in orientation occurs. By contrast, in part (c) of the figure, there is a small and constant change in orientation over a larger section of the curve. Referring to the experimental results for these stimuli, parts (a) and (b) were reported as being discontinuous, while part (c) was reported as being smooth.

In principle, since α and β represent change in orientation, or curvature estimates, $\beta - \alpha$ must reflect *change in curvature.* The experimental results therefore show that simply differencing orientation cues and then thresholding this difference is insufficient for detecting discontinuities. Rather, the curvature estimate resulting from the difference must be compared with other curvature estimates in the neighbourhood. In fact, the situation in this experiment is simplified because all other angles are 0°, so by considering α, β, and $\beta - \alpha$ only we have *implicitly* taken into account other neighbouring curvature estimates. The visual system must therefore be *at least* third order in its overall capabilities. The mechanisms employed at the lower levels can be composed of functions of orientation and curvature: note that comparing curvature and change in curvature over a neighbourhood is not necessarily the same thing as evaluating and representing the derivative of curvature at each point. But the detection of discontinuities, and curve reconstruction in general, must be a non-local process.

The arguments presented here have been necessarily vague. To actually compute theoretically what angles, or combination of angles, would be required to produce discontinuous or smooth percepts would be extremely difficult, since the interactions between the convolutions are complex and not entirely known. In addition, precise measurements of the orientation and curvature resolution of the process would be required – but these measurements in themselves are difficult to sort out from other effects because of the many levels of processing involved in producing the final percept. Finally, the actual computation of curvature and of change in curvature must in general be more complex than the first differences discussed above. However, a version of the model has been implemented and

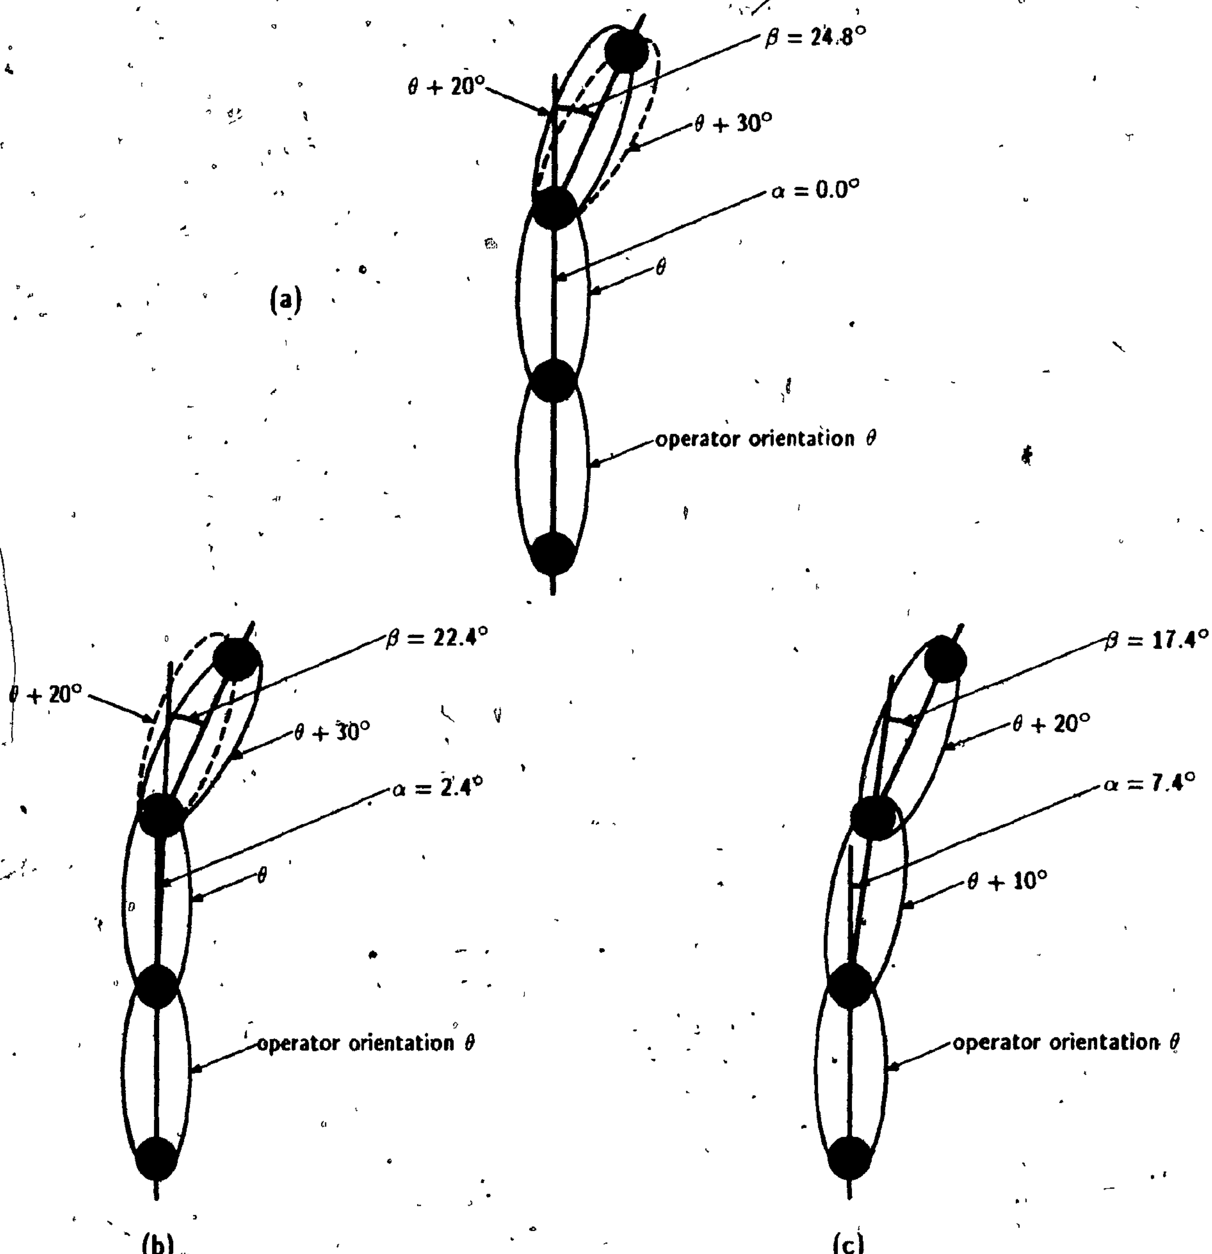

Figure 4.9 A close-up of three sampling phases for an angular curve with orientation change 24.8° to show the angles α and β and the change in response pattern. In part (a), the offset is 0.0, $\alpha = 0.0°$ and $\beta = 24.8°$, in part (b), the offset is 0.1 $\alpha = 2.4°$, and $\beta = 22.4°$, in part (c), the offset is 0.3, $\alpha = 7.4°$ and $\beta = 17.4°$ Subjects reported (a) and (b) as angular and (c) as smooth. The receptive fields shown, with orientations quantised by 10°, are among those that would be responding optimally to the patterns. Note the changes in response pattern that would result from the different offsets The excitatory regions of the receptive fields have been modelled as ellipses with aspect ratios of 3:1 to reflect the Gaussian shape of the implemented operators (Parent and Zucker 1985).

run on the test patterns for the experiment. The results, shown in Figure 4.10, indicate that the interpretation discussed here is in fact characteristic of the model.

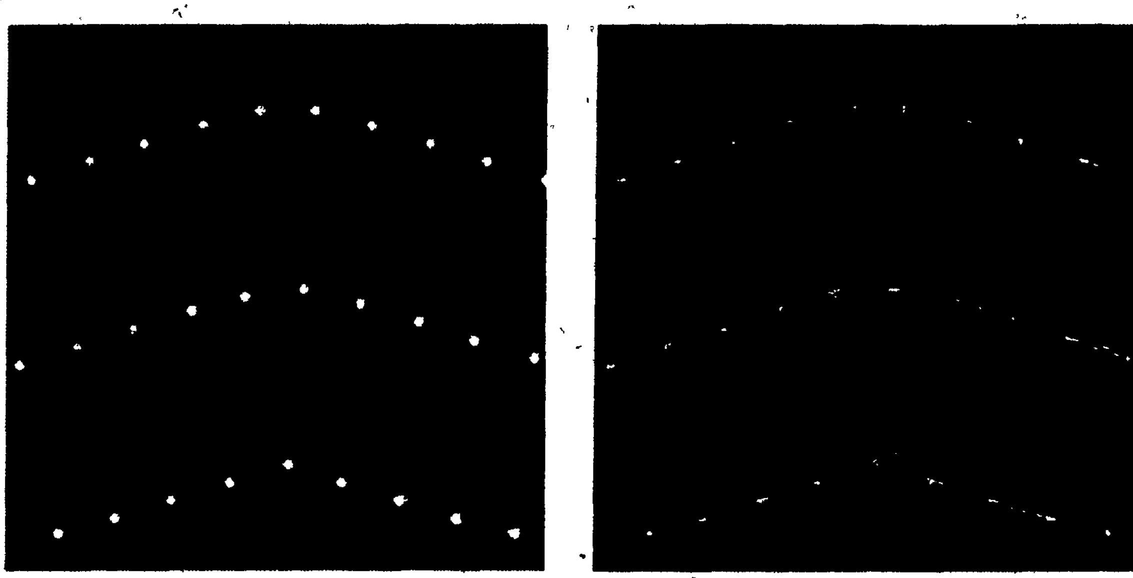

(a) (b)

Figure 4.10 The result of running an implementation of the model presented in Chapter 2 on the test images from the experiment. The implementation used eight orientation quantisations (at regular intervals of 22.5°), and seven rather coarse curvature quantisations. Shown in part (a) is the test image, having an underlying orientation change of 45°. The lower curve was constructed with a sampling offset of 0.0 sampling intervals, the middle with an offset of 0.3 sampling intervals, and the upper with an offset of 0.5. The resulting tangent field is displayed in part (b). The curvature values for the lower curve indicated two straight lines meeting discontinuously at a point. The upper two curves, however, both received smooth interpretations, with non-zero curvature over a broad range of the curve. This result was consistent with the psychophysical results reported earlier. The only problems with the simulation resulted from the regular sampling pattern of the image grid, which implicitly favours orientations that are multiples of 45°. P. Parent is gratefully acknowledged for implementing the model and running this simulation.

In the following chapter, an experiment to assess human sensitivity to corners in flow patterns as a function of the path-length of flow cues is presented. The results will reflect the order of the flow reconstruction mechanism.

Chapter 5 Experiment: Corner Sensitivity vs. Path Metric in Random Dot Moiré Patterns

The purpose of this experiment was to determine how the path-length of a flow pattern affects the sensitivity of human observers to changes in curvature of the underlying flow. Specifically, we are interested in how well discontinuities, or corners, representing small changes in the flow orientation are detected, as a function of the path metric in random dot Moiré patterns. The path metric makes explicit what information about the flow derivatives is available in the image. How the sensitivity to discontinuity sensitivity changes as a function of the path metric will, therefore, reflect on the order of the mechanism used to reconstruct flows.

Subjects were shown several instances of random dot Moiré patterns. These patterns consisted of two translational flow regions meeting in the centre of the image either at a discontinuity or with a smooth transitional region of flow; see again Figures 3.3 and 3.11. Several differences in orientation between the two translational flow regions were used. The independent variable in the images was the path metric, or the number of overlays used to produce the patterns. Subjects were given three possible responses: (1) a single abrupt change in orientation, (2) a smooth change in orientation, or (3) ambiguous. Other factors which could affect sensitivity to discontinuities, such as dot density, dot size, and dot spacing, were held constant at a value within the range that permits these patterns to be clearly visible.[8] Note also that by keeping the dot spacing constant, each increase

[8] The particular value of dot size and dot separation used were within the size/separation constraint for

orientation change, five images with smooth and ten images with discontinuous flows were presented. A different path metric was used for each session. (Path metrics used were 2, 3, 5, 7, and 15.) All subjects viewed the same set of images, although the presentation order was random. The images were displayed for 1.5 seconds, and the screen was blank during the inter-stimulus interval, which was not timed. Subjects were allowed to view the image for one additional 1.5-second interval if they were unable to focus the first presentation. After this, the subject was required to indicate the nature of the underlying curve with one of three choices: (1) a smooth change in orientation; (2) a single abrupt change in orientation; or (3) ambiguous. It was expected that images near the threshold for perceiving discontinuities would elicit the most ambiguous responses. Therefore, if the subject indicated an ambiguous response, he was presented with the *same* image one more time later in the session. At that time, an ambiguous response was accepted as such. (The instances of initial ambiguous responses were also recorded.)

The subjects were shown sample curves such as those in Figure 3.3 to describe the meaning of the different responses. These defining images were followed by a short training set selected from the images to be presented during the session. Each training set image was first flashed for 1.5 seconds and, following a brief interval, was displayed again for the subject to analyse. This was the only reinforcement used during either the training set or the formal experiment.

Images with a path metric of 3 dots were used for the first session, and path metric 2 for the second. Preliminary experiments showed that the patterns were in general easier to see in images with path metric 3 than with path metric 2, although there were still no individual curves visible. Therefore, testing images with path metric 3 first allowed the subject to become familiar with the task during the first session, and acted as an additional training set for the path metric 2 session. For this reason, any increase in curvature sensitivity from path metric 2 images to path metric 3 images could be attributed to properties of the images and the grouping process acting on them, and not to familiarity with the task or with the images. All other path metrics were tested in increasing order.

5.1.4 Stimuli

The stimuli consisted of random dot Moiré patterns that were constructed using an underlying function of two straight lines with some difference in orientation and joined with either a discontinuous or a smooth transition between the two orientations. The smooth transition was accomplished by "capping" the straight lines with a sinusoid of the same *average slope*, in the same way as for the experiment reported in Chapter 4.[9] (See the Appendix for the precise equations used to generate the curves.) Fifteen quantisations of orientation changes were used, ranging from 2.3° to 33.4°. Dot size and spacing were constant across all images, with dot size 2.4 minutes of visual angle (m.v.a.) and dot separation 7.9 m.v.a. (a ratio of dot size to intervening space of 1:3.3). This size and spacing were chosen to minimise the quantisation error of the display.[10] The number of dots used in the images was also constant at 600 dots, a dot density of approximately 3%.

The resulting 30 functions (15 orientation changes, 2 patterns each) were used to generate random dot Moiré patterns with 2, 3, 5, 7, and 15 overlays. For each path metric, ten images were generated using the discontinuous transformation, and five images were generated using the smooth flow function. Several images were generated and used for each instantiation of path metric and underlying function in an effort to overcome the random nature of the process used to generate the images. (The first overlay was always an instantiation from a uniform random number generator.) See Figure 5.1 for sample images.

Not all fifteen quantisations of orientation change were used for each path metric.

---

[9] This method of construction for the smooth flows resulted in only one pattern of curvature change for any one orientation change. However, the subjects were not instructed that this was the case. Rather, they were given, as much as possible, the impression that any pattern of smooth orientation changes might be present, as well as discontinuous flows

[10] This dot size and separation were also chosen to be within the size/separation constraint for one-dimensional contours; see note 7.

(a)

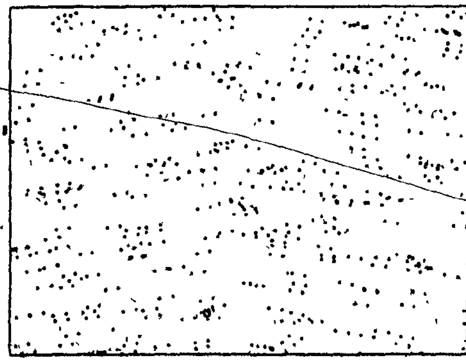

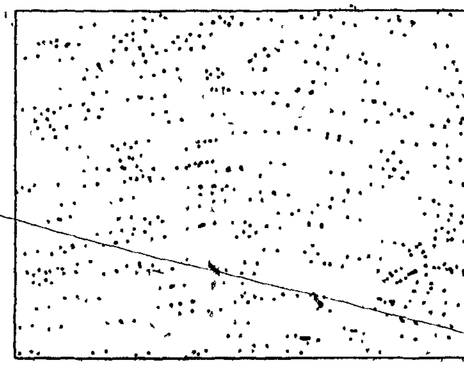

(i) (ii)

(b)

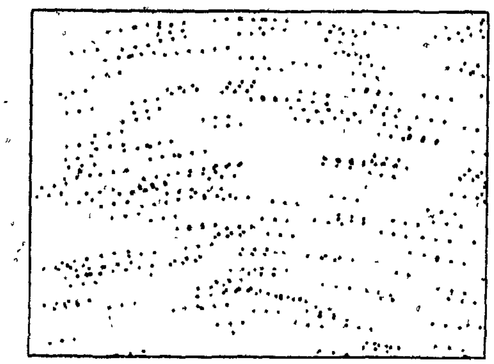

(i) (ii)

Figure 5.1 Sample test images for the experiment: a) path metric 2; b) path metric 3 Part (i) of each is the discontinuous transformation; part (ii) is the smooth transformation for the same overall angle. All parts show the transformation with an underlying angle of 13.69° The dots in these images are 2.4 m v.a (the same size as in the experiment) at a vieweing distance of approximately 0 6 meters. The underlying flow transformations are illustrated in Figure 3.11. (Path metrics 5, 7, and 15 shown on following page)

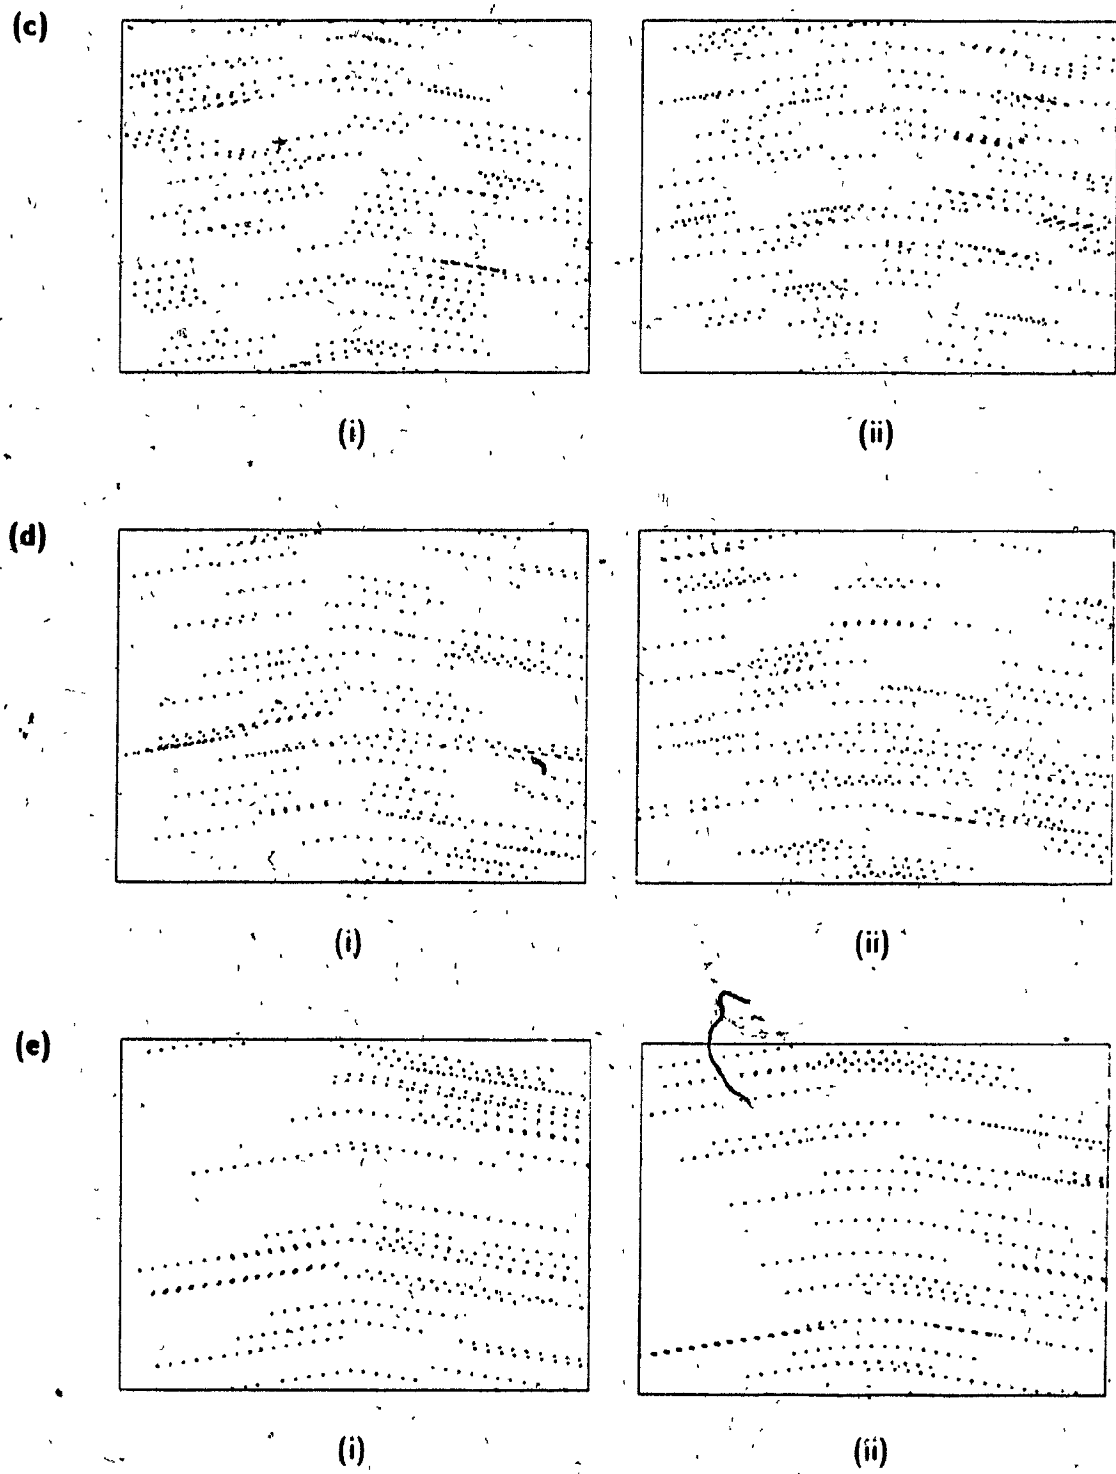

Figure 5.1 (cont'd) Sample test images for the experiment: c) path metric 5; d) path metric 7; e) path metric 15

Rather, preliminary experiments were conducted to determine a narrower range of orientation changes within which the threshold was expected to lie. This range usually encompassed eight quantisations of orientation change, or about 18° difference between the smallest and the largest orientation change. The average session therefore involved having the subject view 120 images. The shortest session consisted of 90 images (six orientation changes, fifteen images at each), and the longest of 150 images (ten orientation changes).

5.2 Results

The results for discontinuous flow fields are tabulated for each path metric in Figure 5.2. (All smooth flows were correctly reported.) A Chi-square test indicates that the magnitude of the orientation change provides input to the response significant to the 0.999 level for all path metrics and for all subjects. This indicates that we were successful in predicting the threshold and in structuring the stimuli to surround it for each path metric. Shown in the figure are the average percentage of discontinuities that were correctly perceived, as a function of the orientation change. The 50% threshold is plotted as a function of path metric in Figure 5.3. Note that the most significant change in threshold occurs between path metrics of 2 and 3 dots, where it decreases from 27.0° to 11.4°. Observe that the threshold sensitivity for dotted one-dimensional contours with the same dot size and dot spacing is also 11.4°; see Section 4.2 (Link and Zucker 1985*b*).

Of interest is the variation in sensitivity between individual subjects as the path metric increases beyond 3 dots, and most evidently beyond 5. For path metrics 2, 3, and 5, the change in threshold is verified by a Chi-square test indicating the path metric had a significant (beyond the 0.999 level) input to the response for those orientation changes tested for all path metrics (that is, between 13.69° and 18.18°; see Figure 5.3 and note that these orientation changes lie between the 50% thresholds for path metric 2 and all other path metrics). In this same region, the actual change in orientation had no effect on the response (significance less than 0.003). For all orientation changes tested, there was no effect on the responses by subject for these path metrics. However, the responses varied

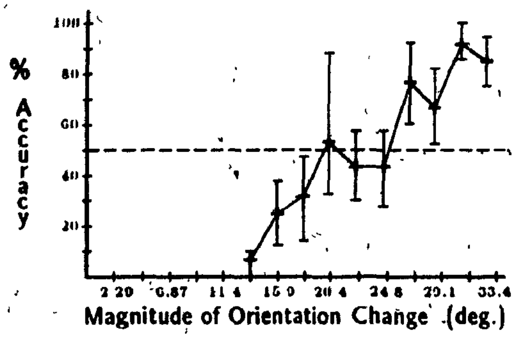

(a) Path Metric 2

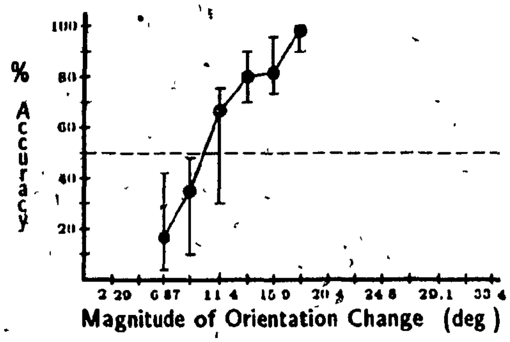

(b) Path Metric 3

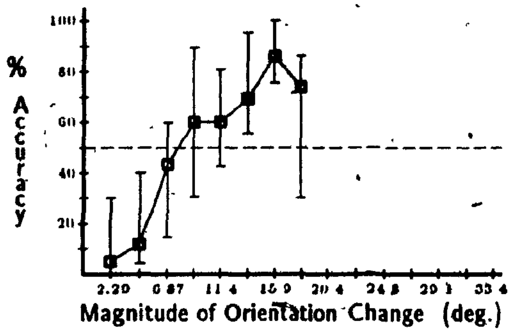

(c) Path Metric 5

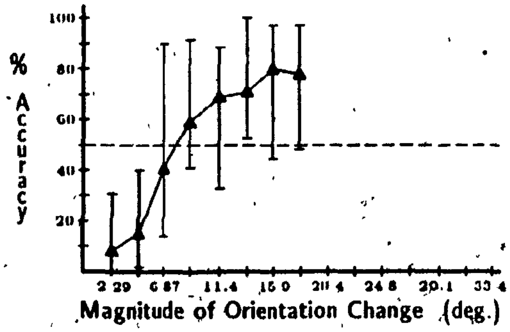

(d) Path Metric 7

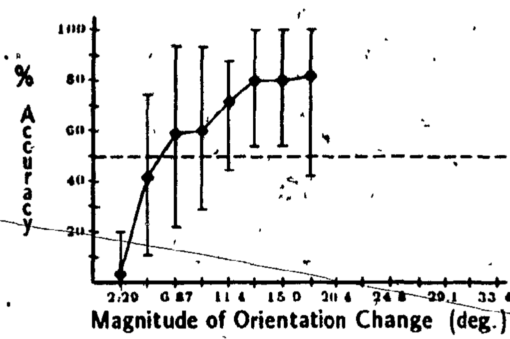

(e) Path Metric 15

Figure 5.2 The results by path metric of the experiment. Illustrated are the mean percentages across all subjects (and one standard deviation) of flow patterns with discontinuities that were correctly identified, as a function of the magnitude of the orientation change. (All smooth patterns were correctly identified.) Note the significant change in threshold between path metrics 2 and 3.

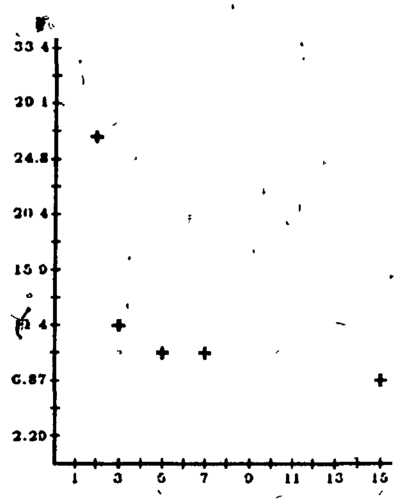

Figure 5.3 The mean threshold sensitivities (50% accuracy of detection) to discontinuities in flow patterns, as a function of path metric.

significantly across subjects (at the 0.999 level) for path metrics 5, 7, and 15. Thus, while the magnitude of the orientation change had a highly significant effect at these path metrics, and the path metric had some effect on the responses, the form of the contribution of the path metric was ambiguous. There *was* variation of the responses across path metric, but it was inconsistent across subjects. Scatter plots of the responses confirm this result.

Subjects were interviewed to determine their strategies for performing the task. Again, all subjects reported using the same strategy at the lower path metrics (2, 3, and 5 dots), but discrepancies appeared at the higher path metrics. The images with lower path metrics appeared to all subjects to have a significant random element and to require a kind of global focus of attention in order to see the pattern at all. At higher path metrics, some subjects maintained a global focus of attention, and these subjects showed a continued increase in sensitivity to corners in the flow. However, some subjects focused their attention on individual curves in the region of the discontinuity at the higher path metrics, and these subjects showed a decrease in sensitivity at these path metrics. In fact, one subject consistently reported the discontinuous flows to be smooth at path metric 15.

5.3 Theoretical Discussion: The Importance of Curvature and Curves Within Flow Patterns

A number of theoretical questions are raised by the psychophysical findings. Foremost among these is why the most dramatic increase in sensitivity occurs when the path metric is increased from 2 to 3 dots. And, second, why is it possible for subjects to have multiple strategies for assessing the flow at higher path metrics, but not at the lower ones? I shall deal with each of these questions in turn within the context of the model developed in Chapter 2.

Recall from Sections 2.3.2 and 3.2 that to reconstruct flows we were required to add to the model for curve reconstruction the additional contraint that the tangent estimates computed in Stage I be approximately equal (or should be averaged) in the direction perpendicular to their orientations. This averaging is required to fill in the gaps along any particular flow contour and to reduce the effects of noise, but it results in a loss of sensitivity to changes in orientation, and in particular to discontinuities. The introduction of higher-order estimates – perhaps by using longer curve segments to represent the flow – would permit this assumption of local parallelism to be relaxed since they would provide a prediction in both flow dimensions of the change in orientation. Under these circumstances, therefore, discontinuity sensitivity would be increased, but only within the order of the estimation mechanism.

The two questions which opened this discussion can now be answered. The increase in sensitivity from path metric 2 to path metric 3 was large because it is only with path metrics of 3 dots or more that curvature information could be reliably estimated: see Table 2.1. Therefore, the mechanism for reconstructing flows must encompass second-degree (curvature) estimation. Further increases in sensitivity with increasing path metric were much less, which suggests that higher-order information is not used by this process. Rather, the increases are more likely attributable to more accurate estimation of first- and second-degree information. At path metric 3 (that is, with curvature information), our sensitivity to discontinuities within flow patterns was the same as that found in contours (Link and Zucker 1985*b*).

One possible confounding of these results could arise from the fact that there is more orientation information available at path metric 3 than at path metric 2, since each triple of dots can be viewed as two pairs. (With 600 dots in each image, 300 orientation cues are present at path metric 2, and 400 are present at path metric 3). If this were the case, the empirical argument in support of curvature would be incorrect. But it is not the case, as we have shown in Figures 3.8 and 3.9. Even doubling the amount of orientation information (by doubling the density but leaving the path metric at 2) is not the same as adding curvature information, or the correlation between neighbouring orientation cues (using the same overall dot density, but constructing the pattern with path metric 3).

The answer to the second question involves a mixture of curves and flows. As path metric increases to 5 and beyond, it is possible to see the segments not only as part of a flow but as distinct contours as well. If a subject were actually to do this, then a further higher-level question would arise regarding how to mix cues about discontinuities from these two different kinds of processes. It is not surprising that this would cause some confusion, as the data showed for higher path metrics. In particular, it is highly probable that few if any of the contours overlapping the discontinuity in the images with higher path metrics met the criteria established in Chapter 4 for detecting discontinuities.

In conclusion, then, the anticipation and detection of changes in orientation structure within flow patterns is essential to accurate recovery of the flow. On the surface, the theory presented in Section 2.3.2 is similar to others that have been proposed for flow reconstruction (Glass 1969 and 1982; Glass and Switkes 1976; Stevens 1978). These theories also involve the two stages of extraction of orientation cues followed by an interpretation of these cues. But it is in the interpretation stage that the theory presented here differs.

Both Glass and Stevens fail to make provisions for changes in the flow. The assumption of local parallelism is rigidly applied, so that as the assumption breaks down the methods will ultimately fail. One way in which the assumption of local parallelism breaks down is with the introduction of noise. Another way is when *systematic* changes in orientation, or changes in the flow itself, occur. Since neither method makes provision for the structure of the flow to change locally, any curvatures or systematic changes in the flow resulting in loss of local parallelism are treated in the same way as noise, and hence are not detected.

The model presented here seeks to impose a structure *in the direction of the flow* which allows the assumption of local parallelism to be relaxed in regions of orientation change. The degree to which we can relax the assumption depends directly on the information we have about the structure of the flow. The model can take advantage of curvature information available at longer path-lengths only because it has a mechanism for representing the local structure of the flow in *two dimensions* – in the direction perpendicular to the flow (the notion of local parallelism) *and* in the same direction as the flow. The psychophysical results demonstrate that the human visual system also takes both dimensions of the flow structure into account when it reconstructs the flow, and that the path-length of orientation cues provides an important information parameter.

Chapter 6 Summary and Conclusions

Contours are important because they separate objects from each other and from the background, and their description provides meaningful information about the shape of objects. Flow patterns are oriented surface coverings that provide information about surfaces within the boundaries defined by contours. It is therefore essential to produce an accurate representation of them early in the visual process. In the past, orientation – or the first derivative with respect to arc length – has received the greatest attention in the research on the perception of contours, flow patterns, and shape in general. However, the change in orientation – or curvature – also plays an important role. In particular, orientation discontinuities are one of the most salient and useful properties of curves and flows. They often represent boundaries between objects or surfaces, and locating them is essential to constructing an accurate representation of the orientation structure. But our perception of curvature and of orientation discontinuities has rarely been addressed. In this thesis, I have attempted to confront discontinuities directly by measuring human sensitivity to them and by analysing these measurements within a theoretical context for computing them.

The experiments I performed introduced a new class of stimuli – dotted curves and dotted flow patterns (random dot Moiré patterns) – for quantitatively evaluating this sensitivity. The use of these stimuli is valid as long as the dots are small and close together. The discrete nature of such images made it possible to control the local differential structure of orientation cues and hence to evaluate the performance of the model in detecting discontinuities as a function of this structure. By concentrating on the detection of changes in orientation – curvature in the continuous case, and "corners" in the discontinuous one

– I was able to show that the estimation of higher-order derivatives is required for curve and flow perception and for the construction of higher-level shape representations. The psychophysical results indicate that human sensitivity to discontinuities is dependent on a comparison at least of curvature, and perhaps of change in curvature, over a neighbourhood. There has been something of a controversy concerning how curvature is used and whether curvature detectors exist in the cortex. The issue presented here, however, is not whether curvature detectors exist, but is rather how curvature, change in curvature, and discontinuities are computed.

The research on discontinuity perception in curves was motivated by the observation that changing the sampling phase of a curve changes the geometry of its discrete trace. In other words, if curves are represented by dots, the placement as well as the density of the dots will affect an interpolation of the curve. Psychophysical experiments were performed to evaluate human sensitivity to orientation discontinuities for pairs of straight lines meeting at a point as a function of dot placement. The results of these experiments were examined for a two-stage computational model of orientation selection. The first stage is the construction of a tangent field to the curve using convolutions against simple-cell-like operators. Within this model, and to a first approximation, changes in the percept arise when the sampling has been changed sufficiently to alter individual convolutions. In the case of straight lines, this amounts to requiring, perhaps, that a dot be moved sufficiently to leave the excitatory centre and enter the inhibitory surround. However, the simple cells (or their computational equivalent) must interact to eliminate conflicting responses and to fill in gaps. It is here, in this interaction step, that the need for curvature arises, and it is only when these interactions are taken into account that the amount of dot displacement (caused by changes in sampling) required to change the percept can be evaluated. In the second stage of the model, when the curve is actually interpolated, the change in curvature over a neighbourhood is also needed, particularly for the detection of discontinuities.

The representation of flows in images is considerably sparser than that of curves. The model for curve reconstruction can, however, be extended to reconstruct flows by imposing an assumption of locally parallel structure. This assumption permits averaging of orientation information in the direction perpendicular to the flow to fill in the large gaps along any one flow contour. It was noted, that this would result in a decreased sensitivity to changes in orientation (and hence to orientation discontinuities), particularly since the assumption of local parallelism is by definition violated in regions where the flow orientation is changing. This decrease in sensitivity could be counter-acted, however, by relaxing the assumption (and averaging over smaller neighbourhoods) when information about higher-order derivatives is available. The research on discontinuity perception in flows was motivated by the observation of a new parameter of flows – the path-length of the flow cues – which could be used to accomplish this relaxation. Random dot Moiré patterns – flow patterns in which the curve segments acting as individual flow cues are dotted – make explicit what order of information can ideally be extracted from the image. This information is embodied in the path-length metric, defined as the number of dots used to represent each curve segment. The psychophysical experiment was designed to assess human sensitivity to discontinuities in random dot Moiré patterns with two regions of straight flow meeting in a line at the centre of the image, as a function of this metric. The results, examined within the context of the model of orientation selection, showed that curvature information is also used for flow reconstruction when this information can be obtained directly from the image. Otherwise, the assumption of local parallelism is applied everywhere. Therefore, and contrary to previous methods, the anticipation and detection of changes in orientation structure within flow patterns is used to accurately recover flow and detect orientation discontinuities.

In summary, then, orientation structures such as curves and flow patterns can be recovered directly from image-like structures (with a minimum of preprocessing required). In this thesis, I have shown that the mechanisms which accomplish this task are not local – they must determine the differential structure of the curve or flow pattern through comparisons of orientation and curvature estimates over open neighbourhoods. I have demonstrated this fact by showing that human sensitivity to orientation discontinuities is a function of the curvature and change-in-curvature information in a neighbourhood surrounding the discontinuity.

References

Attneave, F. 1954. Some informational aspects of visual perception. *Psych. Rev.* 61:183-193.

Attneave, F. 1957. Physical determinants of the judged complexity of shapes. *J. Exper. Psych.* 53:221-227.

Attneave, F. 1959. *Applications of Information Theory to Psychology: A Summary of Basic Concepts, Methods, and Results.* New York: Holt, Rinehart & Winston.

Attneave, F., and Arnoult, M.D. 1956. The quantitative study of shape and pattern perception. *Psych. Bull.* 53:452-471.

Barrow, H.G., and Tenenbaum, J.M. 1978. Recovering intrinsic scene characteristics from images. In *Computer Vision Systems,* eds. A.R. Hanson and E.M. Riseman. New York: Academic Press.

Blakemore, C., and Tobin, E. 1972. Lateral inhibition between orientation detectors in the cat's visual cortex. *Exper. Brain Res.* 15:439-440.

Carpenter, R.H.S., and Blakemore, C. 1973. Interactions between orientations in human vision. *Exper. Brain Res.* 18:287-303.

Crassini, B., and Over, R. 1975. Masking, after-effect and illusion in the visual perception of curvature. *Perception and Psychophysics* 17:411-416.

Davis, P.J. 1965. *Interpolation and Approximation.* New York: Blaisdell Publishing Co.

Feng, F.Y., and Pavlidis, T. 1973. Finding "vertices" in a picture. *Comp. Graphics and Im. Proc.* 2:103-117.

Foster, D.H. 1983. Visual discrimination, categorical identification, and categorical rating in brief displays of curved lines: implications for discrete encoding processes. *J. Exper. Psych.* 9:785-806.

Freeman, H. 1974. Computer processing of line drawing images. *Computing Surveys* 5:57-97.

Freeman, H., and Davis, L.S. 1977. A corner-finding algorithm for chain-coded curves. *IEEE Trans. Comp.* C-26:297-303.

Freund, J.E. 1971. *Mathematical Statistics.* 2nd ed. New Jersey: Prentice-Hall.

Glass, L. 1969. Moiré effect from random dots. *Nature* 223:578-380.

Glass, L. 1982. Physiological mechanisms for the perception of random dot Moiré patterns. In *Chaos and Order in Nature: Proc. of the Int. Symp. on Synergetics*, at Schloss Elmau, Bavaria, April 27 – May 2, 1981, ed. H. Haken. (*Springer Series in Synergetics*, Vol. 11) Berlin: Springer-Verlag.

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.

Gupta, J.N., and Wintz, P.A. 1975. A boundary finding algorithm and its applications. *IEEE Trans. on Circuits and Systems* CS-22:351-362.

Heggelund, P., and Hohmann, A. 1975. Responses of striate cortical cells to moving edges of different curvature. *Exper. Brain Res.* 23:211-216.

Hubel, D., and Wiesel, T. 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol. (London)* 160:106-154.

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

Hueckel, M.H. 1973. A local visual operator which recognizes edges and lines. *J. Assoc. Comp. Mach.* 20:634-647.

Klemmer, E.T., and Frick, F.C. 1953. Assimilation of information from dot and matrix patterns. *J. Exper. Psych.* 45:15-19.

Koffka, K. 1935. *Principles of Gestalt Psychology*. New York: Harcourt, Brace & World.

Kruse, B., and Rao, C.V.K. 1978. A matched filtering technique for corner detection. *Proc. 4th Int. Joint Conf. Patt. Rec.*, Kyoto, Japan, Nov. 7-10, 1978: 642-644.

Leclerc, Y., and Zucker, S.W. 1984. The local structure of image discontinuities in one dimension. *Technical Report* 83-19R. Montréal: Dept. of Elect. Eng., McGill University.

Levine, M.D. 1973. Scene analysis for a breadboard Mars robot functioning in an indoor environment. *Technical Memorandum* 33-645. Pasadena: Jet Propulsion Laboratory, California Institute of Technology.

Levine, M.D. 1985. *Vision in Man and Machine*. New York: McGraw-Hill Book Co.

Levine, M.D., and Shaheen, S.I. 1981 A modular computer vision system for picture segmentation and interpretation. *IEEE Trans. on Patt. Anal. and Mach. Intell.* PAMI-3:540-556.

Link, N.K., and Zucker, S.W. 1985*a*. Sensitivity to corners in flow patterns. *Technical Report* 85-4R. Montréal: Dept. of Elect. Eng., McGill University.

Link, N.K., and Zucker, S.W. 1985*b*, in preparation. Corner detection in curvilinear dot grouping. *Technical Report* 85-5R. Montréal: Dept. of Elect. Eng., McGill University.

Marr, D. 1976. Early processing of visual information. *Phil. Trans. Roy. Soc., B* 275:483-524.

Marr, D. 1982. *Vision*. San Francisco: W.H. Freeman & Co.

Marr, D., and Hildreth, E. 1980. Theory of edge detection. *Proc. Roy. Soc. (London), B* 207:187-217.

Marr, D., and Poggio, T. 1979. A computational theory of human stereo vision. *Proc. Roy. Soc. (London), B* 204:301-328.

Meurle, J.L., and Allen, D.C. 1968. Experimental evaluation of techniques for automatic segmentation of objects in a complex scene. In *Pictorial Pattern Recognition*. eds. G.C. Chang, R.S. Ledley, D.K. Pollock, and A. Rosenfeld. Washington: Thompson.

Movshon, J.A., and Blakemore, C. 1973. Orientation specificity and spatial selectivity in human vision. *Perception* 2.53-60.

Ohlander, R.; Price, K.; and Reddy, D.R. 1978. Picture segmentation using a recursive region splitting method. *Comp. Graphics and Im. Proc.* 8:313-333.

Oppenheim, A.V., and Schafer, R.W. 1975. *Digital Signal Processing*. New Jersey: Prentice-Hall.

Orban, G. 1984. *Neuronal Operations in the Visual Cortex*. New York: Springer-Verlag.

Parent, P., and Zucker, S.W. 1985. Curvature consistency and curve detection. *Technical Report* 85-12R. Montréal: Dept. of Elect. Eng., McGill University.

Pavlidis, T. 1973. Segmentation of pictures and maps through functional approximation. *Comp. Graphics and Im. Proc.* 1:360-372.

Poggio, T., and Torre, V. 1978. A new approach to synaptic interactions. In *Approaches to Complex Systems*, eds R. Heim and G. Palm, pp 89-115. Berlin: Springer-Verlag.

Popham, A E. 1945. *The Drawings of Leonardo da Vinci*. New York: Reynal & Hitchcock.

Ramer, U. 1972 An iterative procedure for the polygonal approximation of plane curves. *Comp. Graphics and Im. Proc.* 1:244-256.

Richter, Jean Paul. 1939 *The Literary Works of Leonardo da Vinci*, 2nd ed., 2 vols. London: Oxford University Press.

Riggs, L.A. 1973. Curvature as a feature of pattern vision. *Science* 181:1070-1072.

Rosenfeld, A., and Johnston, E. 1973. Angle detection on digital curves. *IEEE Trans. Comp.* C-22:875-878.

Rosenfeld, A.; Thomas, R.B.; and Lee, Y.H. 1970. Edge and curve detection for texture discrimination. In *Picture Processing and Psychopictoris.* eds. B.S. Lipkin and A. Rosenfeld. New York: Academic Press.

Rosenfeld, A., and Thurston, M. 1971. Edge and curve detection for visual scene analysis. *IEEE Trans. Comp.* C-21:677-715.

Rosenfeld, A., and Weszka, J.S. 1975. An improved method an angle detection on digital curves. *IEEE Trans. Comp.* C-24:940-941.

Schacter, B.J.; Davis, L.S.; and Rosenfeld, A. 1976. Scene segmentation by cluster detection in color space. *Special Interest Group on Artificial Intelligence (of the ACM).* SIGART 58:16-17.

Schiller, P.H.; Finlay, B.L.; Volman, S.F. 1976. Quantitative studies of single-cell properties in monkey striate cortex: II. Orientation specificity and ocular dominance. *J. Neurophys.* 39:1320-1333.

Sillito, A.M.; Kemp, J.A.; Milson, J.A.; Berardi, N. 1980. A re-evaluation of the mechanisms underlying simple cell orientation selectivity. *Brain Res.* 194:417-520.

Stevens, K.A. 1978. Computation of locally parallel structure. *Biol. Cyber.* 29:19-28.

Timney, B.N., and Macdonald, C. 1978. Are curves detected by "curvature detectors"? *Perception* 7:51-64.

Tomita, F., and Tsuji, S. 1977. Extraction of multiple regions by smoothing in selected neighborhoods. *IEEE Trans. on Systems, Man, and Cyber.* SMC-7:107-109.

Ullman, S. 1979. The interpretation of structure from motion. *Proc. Roy. Soc. (London), B* 203:405-426.

Wallach, H., and O'Connell, D.N. 1953. The kinetic depth effect. *J. Exper. Psych.* 45:205-217.

Watt, R.J., and Andrews, D.P. 1982 Contour curvature analysis: hyperacuities in the discrimination of detailed shape. *Vision Res.* 22:449-460.

Weszka, J.S. 1978. Threshold selection techniques. *Comp. Graphics and Im. Proc.* 7:259-265.

Witkin, A.P., and Tenenbaum, J.M. 1984. On the role of structure in vision. In *Human and Machine Vision*, eds. A. Rosenfeld and J. Beck, pp. 481-543. New York: Academic Press.

Zucker, S.W. 1982. Early orientation selection and grouping: type I and type II processes. *Technical Report* 82-6R. Montréal: Dept. of Elect. Eng., McGill University.

Zucker, S.W. 1984*a*. Computational and psychophysical experiments in grouping. In *Human and Machine Vision*, eds. A. Rosenfeld and J. Beck. New York: Academic Press.

Zucker, S.W. 1984*b*. Type I and Type II processes in early orientation selection. In *Figural Synthesis*, eds. T. Caelli and P. Dodwell. Hillsdale, N.J.: Lawrence Erlbaum Assoc.

Zucker, S.W. 1985. Early orientation selection: tangent fields and the dimensionality of their support. *Technical Report* 85-13R Montréal: Dept. of Elect. Eng., McGill University. *Comp. Vision. Graphics, and Im. Proc.*, in press.

Zucker, S.W., and Davis, S. 1985. Points and endpoints: a size/spacing constraint for dot grouping. *Technical Report* 85-3R. Montréal: Dept. of Elect. Eng., McGill University.

Zucker, S.W.; Hummel, R.A.; and Rosenfeld, A. 1975. Applications of relaxation labelling, 1: line and curve enhancement. *Technical Report* TR-419. Maryland: Computer Science Center, University of Maryland.

Appendix

The equations used to generate stimuli for both experiments were as follows:

discontinuous:

$$
\begin{aligned}
y &= x \times \tan(\theta/2) && 0 \le x \le X/2 \\
&= (X - x) \times \tan(\theta/2) && X/2 \le x \le X
\end{aligned}
$$

smooth:

$$
\begin{aligned}
y &= x \times \tan(\theta/2) && 0 \le x \le b \\
&= X/4 \times \tan(\theta/2) \times \sin(\pi x/X) + d && b \le x \le X - b \\
&= (X - x) \times \tan(\theta/2) && X - b \le x \le X
\end{aligned}
$$

where $X =$ screen width

$\theta =$ orientation change across corner $(0° < \theta < 35°)$

$b = \arccos(\tan(\theta/2) \times 4/X) \times X/\pi.$

$=$ point of equal slope between line and sinusoid

$d = b\tan(\theta/2) - (X/4)\tan(\theta/2)\sin(\pi b/X)$

This method of construction for the smooth contours resulted in only one pattern of curvature change for any one orientation change. However, the subjects were not instructed that this was the case. Rather, they were given, as much as possible, the impression that any pattern of smooth orientation changes might be present, as well as other types of discontinuities.