

Evaluating Accounts of Human Contour Integration Using Psychophysical And Computational Methods.

Paul George Lovell, B.Sc.(Hons.), M.Sc.

Thesis submitted for the degree of doctor of philosophy

2002

Department of Psychology, University of Stirling, Scotland.

Acknowledgements

Firstly, I'd like to thank professors Roger Watt and Bill Phillips for their advice and help throughout my PhD.

I'd like to thank my partner, friend and willing psychophysical observer Dr Karen Spencer for her support and help throughout the past six years. Harry (the dog) and numerous fish and Canada (the country) for helpful distractions/diversions. My family also deserve my thanks and appreciation for their understanding and support of my continued presence within higher education. I'd especially like to thank my Grandfather, for encouraging an enquiring mind, despite the fact that my enquiries often involved the destruction of his valve-radios.

University of Stirling colleagues, Ben Craven, Darragh Smyth and Peter Hancock are owed a great debt of thanks for guiding me when I got lost down dark mathematical alleyways. Also, Tim Chapman for assisting in the reconstruction of my PC after Linux distractions. Emmanuel Stamatakis (Manoli) for being a willing observer and a helpful advisor. My thanks also go to Paul Miller for his helpful advice on undertaking a PhD, which I largely ignored – at my own cost.

I'd also like to thank a constant stream of observers: especially Liz Humphreys, Bill Phillips, Roisin Ash, Trish Carling, Charlie Frowd and numerous University of Stirling undergrads. Finally, I'd also like to thank the departmental technicians for equipment maintenance, particularly Bruce Sutherland and Peter Hucker, without whom much of my research would have been impossible.

No thanks go to Bill Gates and Microsoft, especially Windows 98/2000 and Microsoft Word, for always throwing a spanner in the works just when everything appeared to be working smoothly.

Table of Contents

Glossary.....	1-1
1 Introduction.....	1-1
1.1 Overview.....	1-1
1.2 Gestalt psychology.....	1-2
1.3 Experimental investigations of contour detection.....	1-3
1.4 Development of the path-paradigm.....	1-8
1.4.1 Field <i>et al's</i> experimental evidence for an association field.....	1-8
1.4.2 The extension of the classical receptive field.....	1-11
1.4.3 Evidence suggests that the mechanisms underlying contour integration are acquired through visual experience.....	1-14
1.4.4 Computational models that are deliberately sensitive to smooth contours.....	1-17
1.4.5 Alternative explanations of Field <i>et al's</i> findings.....	1-22
1.4.6 Effects of global contour characteristics upon contour integration.....	1-28
1.4.7 The association field is only a phenomena of the fovea.....	1-30
1.5 Summary.....	1-30
1.6 Directions for the current thesis.....	1-32
1.6.1 Testing the simple-filter model of contour integration.....	1-32
1.6.2 Do the global characteristics of contours influence detectability?.....	1-33
1.6.3 Does the presence of spurious contours in the background influence detection?.....	1-33
1.6.4 Contour integration in the periphery.....	1-34
2 General Methods.....	2-1
2.1 Overview.....	2-1
2.1.1 Stimulus Generation.....	2-1
2.1.2 Presentation sequencing.....	2-5
2.1.3 Displaying stimuli.....	2-5
2.1.4 Response storage and analysis.....	2-6
2.2 Experiment 1: Replication of Field, Hayes and Hess path-angle experiment.....	2-9
2.2.1 Introduction.....	2-9
2.2.2 Methodology.....	2-10
2.2.3 Results.....	2-11
2.2.4 Discussion.....	2-12
2.3 Experiment 2: Manipulating contour element spacing.....	2-12
2.3.1 Introduction.....	2-12
2.3.2 Methodology.....	2-13
2.3.3 Results.....	2-14
2.3.4 Discussion.....	2-15
2.4 Summary.....	2-15
3 Modelling Contour Integration with the Simple-Filter Model.....	3-1
3.1 Overview.....	3-1
3.1.1 Selection of filter scale in the simple-filter model.....	3-2
3.1.2 Differences from other models.....	3-4
3.2 Methodology.....	3-4
3.2.1 The simple-filter model.....	3-4
3.2.2 Presentation of experimental results.....	3-7
3.3 Simulation 1: Modelling the influence of path-angle (expt. 1).....	3-10
3.3.1 Stimulus image parameters.....	3-10
3.3.2 Procedure.....	3-11
3.3.3 Simulation results.....	3-11
3.3.4 Discussion.....	3-17
3.4 Why do the current simulation results differ from those of Hess and Dakin (1999)?.....	3-18
3.4.1 Comparing stimulus images.....	3-18
3.4.2 Comparing the filters currently employed and that used within the Hess and Dakin model.....	3-20
3.4.3 Comparing current performance with a filter that is comparable to that employed by Hess and Dakin.....	3-20
3.4.4 Discussion.....	3-21
3.5 Simulation 2: The influence of contour element spacing (Expt. 2).....	3-22
3.5.1 Overview.....	3-22
3.5.2 Stimuli image parameters.....	3-22
3.5.3 Procedure.....	3-23

3.5.4	Simulation results.....	3-23
3.5.5	Discussion.....	3-28
3.6	Simulation 3: Why is the simple-filter model impaired when detecting phase-alternated contours with low levels of curvature?.....	3-29
3.6.1	Overview.....	3-29
3.6.2	Methodology.....	3-31
3.6.3	Discussion.....	3-32
3.7	General Discussion.....	3-33
3.7.1	Performance of the simple-filter model with manipulations of contour path-angle.....	3-33
3.7.2	Performance of the simple-filter model with manipulations of contour element spacing.....	3-34
3.7.3	Type-1 error in estimates of simple-filter model performance?.....	3-35
3.7.4	Is integration occurring within the simple filter model?.....	3-35
3.7.5	A note on the coarse sampling of Gabor patches in the current simulations.....	3-36
3.7.6	On filter selection mechanisms.....	3-36
3.7.7	Summary.....	3-37
4	The Influence of Global Contour Structure upon Saliency.....	4-1
4.1	Overview.....	4-1
4.2	Experiment 3: Verification of the Kovacs and Julesz' closure effect.....	4-1
4.2.1	Introduction.....	4-1
4.2.2	Methodology.....	4-1
4.2.3	Results.....	4-4
4.2.4	Discussion.....	4-8
4.3	Experiment 4: Examining the influence of contour 'smoothness'.....	4-10
4.3.1	Overview.....	4-10
4.3.2	Methodology.....	4-11
4.3.3	Results.....	4-13
4.3.4	Discussion.....	4-16
4.4	Experiment 5: Is integration impaired at inflection points?.....	4-18
4.4.1	Introduction.....	4-18
4.4.2	Methodology.....	4-20
4.4.3	Results.....	4-22
4.4.4	Can probability summation account for the detectability of multiple 'C' contours?.....	4-25
4.4.5	Discussion.....	4-26
4.5	Experiment 6: Comparing the influence of contour breaks at inflection points and elsewhere.....	4-28
4.5.1	Introduction.....	4-28
4.5.2	Methodology.....	4-30
4.5.3	Results.....	4-33
4.5.4	Discussion.....	4-36
4.6	General Discussion.....	4-36
4.6.1	Summary.....	4-37
5	How does arc-length influence the performance of the simple-filter model?.....	5-1
5.1	Simulation 4: The performance of the simple-filter model with the arc-length stimuli.....	5-1
5.1.1	Introduction.....	5-1
5.1.2	Methodology.....	5-2
5.1.3	Results.....	5-4
5.1.4	Discussion.....	5-9
5.2	Simulation 5: Detection rates for the simple-filter model were influenced by the presence of inflections within contours.....	5-10
5.2.1	Introduction.....	5-10
5.2.2	Methodology.....	5-11
5.2.3	Results.....	5-11
5.2.4	Discussion.....	5-12
5.3	General discussion.....	5-13
6	Manipulating the level of co-linearity within the stimulus background.....	6-1
6.1	Overview.....	6-1
6.2	Manipulation of background element co-linearity.....	6-2
6.2.1	Measuring co-linearity using an association-field model.....	6-2
6.2.2	Procedure for the manipulation of contour membership of an individual line element.....	6-5
6.2.3	Manipulating all element orientations within a background.....	6-6
6.2.4	On the generation of stimulus backgrounds.....	6-10
6.3	Experiment 7: The influence of noise-contour prevalence upon the path-angle detection limits.....	6-12

6.3.1	Overview.....	6-12
6.3.2	Methodology.....	6-12
6.3.3	Results.....	6-14
6.3.4	Discussion.....	6-18
6.4	Experiment 8: The influence of CLE upon the detection of short contours.....	6-19
6.4.1	Overview.....	6-19
6.4.2	Methodology.....	6-20
6.4.3	Results.....	6-21
6.4.4	Discussion.....	6-25
6.5	General Discussion.....	6-26
7	Does the same process underlie contour detection in the fovea and in the para-fovea?.....	7-1
7.1	Overview.....	7-1
7.2	Experiment 9: Reducing ‘co-linearity energy’ in the background leads to improved contour detection in the periphery.....	7-2
7.2.1	Introduction.....	7-2
7.2.2	Methodology.....	7-3
7.2.3	Results.....	7-6
7.2.4	Discussion.....	7-9
7.3	Experiment 10: Do global properties influence contour detection in parafoveal areas?.....	7-10
7.3.1	Introduction.....	7-10
7.3.2	Methodology.....	7-11
7.3.3	Results.....	7-14
7.3.4	Discussion.....	7-21
7.4	General Discussion.....	7-23
8	General Discussion.....	8-1
8.1	The role of coarse-scale processes in the detection of path-paradigm stimuli.....	8-1
8.1.1	Detection of contours with a random global structure.....	8-1
8.1.2	Detection of smooth contours.....	8-2
8.1.3	Summary.....	8-2
8.2	Smoothness in path-paradigm stimuli.....	8-3
8.3	Contour detection in the periphery.....	8-4
8.4	Background element co-linearity energy.....	8-6
8.5	Acquisition of an interactive contour integration mechanism.....	8-9
8.5.1	Neurophysiological evidence for the roles of feed-forward, horizontal and feed-back connections.....	8-13
8.5.2	Could visual experience influence the development of contour integration processes?.....	8-15
8.5.3	Current experimental results.....	8-16
8.5.4	Questions raised by the interactive contour integration model.....	8-19
8.6	Conclusions.....	8-21
	References.....	
	Table of Contents for Appendices.....	
	Appendices.....	

Field, Hayes and Hess (1993) offered path-paradigm (*PP*) stimuli that enabled the investigation of the Gestalt processes of proximity and good continuation without the potential confounds present in earlier studies. They proposed that only an association field, a mechanism able to integrate dynamically the outputs of filters with different orientation preferences, could detect their stimuli. This thesis describes simulations which examined whether the *PP* task could be solved without recourse to an association field. A simple-filter model (*SFM*) was tested, where each response image was the result of the convolution of the stimulus image with an oriented filter. The lengths of zero-bounded regions (*ZBRs*) within each image were calculated (Watt, 1991). In a simulated 2AFC trial the status of *target* was assigned to the image containing the longest *ZBR*. Results confirmed the Hess and Dakin (1997) finding that Field *et al.*'s *PP* stimuli could be successfully detected by the *SFM*. Further simulations conflicted with Hess and Dakin; the *SFM* was also able to detect stimuli containing phase-alternated contours. Thus, it is not necessary to invoke an association field mechanism to explain contour integration, even for phase-alternated stimuli.

Psychophysical experiments indicated that the closure effect reported by Kovaks and Julesz (1993) may actually be caused by contour smoothness, rather than closure *per se*. Where local properties are held constant, manipulations of contour closure are inevitably confounded by smoothness. Further modelling experiments revealed that, for the *SFM*, detection rates were inversely correlated with global smoothness, a reversal of the pattern found for human observers. This phenomenon provides a useful means of investigating the *PP* task in the peripheral field. Hess and Dakin (1997) have argued that the *SFM* is sufficient to account for contour integration processes in the periphery. Thus, for the periphery there should be an inverse relationship between smoothness and detectability. Experimental results revealed that this is not the case, smooth contours were detected more frequently than jagged contours.

It is concluded that the Field *et al.* (1993) conception of the association field was largely correct, despite the fact that their stimuli were potentially flawed. However, the association field model may need modification in order to account to the effect of contour smoothness. It is suggested that the association field model should be considered a component in Roelfsema, Lamme and Spekreijse's (2000) model of curve-tracing. Within this framework the process of 'chunking' (Mahoney and Ullman, 1988) might explain the improved detection of frequently experienced contour shapes, i.e. straight lines and smooth arcs. Furthermore, it is suggested that this mechanism may be acquired through visual experience and that it could underlie *PP* performance in the whole visual field.

Declaration

The work described here was carried out in the Department of Psychology at the University of Stirling. All the work is my own, unless otherwise stated, and it has not been submitted previously for a degree at this or any other institution.

Signed..... Date

Glossary

<i>Arc-length</i>	The 'arc-length' of a contour, refers to the length of individual arcs that form a contour. Where an arc is an uninflected section of contour. The terms <i>arc-length</i> and <i>smoothness</i> are used interchangeably.
<i>Contour Integration</i>	This term represents the process of forming a representation that denotes which discrete items within a visual scene are members of a particular contour.
<i>Filter</i>	A filter is a function which is applied separately to each point in an image. When a filter is applied to any one point in the image, it takes a subset of values from about that point in the image, and combines them in a fixed manner to produce a single value which is the filtered output
<i>Filtered image</i>	A filtered image is the array of filtered outputs, one for and corresponding to every point in the input image.
<i>Jagged/Jaggedness</i>	The term 'jagged' refers to the local properties of the contour only, as it's these local properties that are most likely to influence performance of the contour detection models under investigation. So, a jagged contour is only jagged in a very local sense, i.e. the sign of successive path-angle changes varies with a high frequency. However, in a non-local sense the contours may actually appear quite smooth, i.e. where path-angle jitter is low the overall shape of the contour is relatively straight
<i>Jitter</i>	A small, usually random, value that is combined with another value. Values are usually <i>jittered</i> in order to reduce the likelihood of the occurrence of patterns that might ease contour detection.
<i>Line segregation</i>	A process which identifies which areas of the visual scene belong to a particular contour. In experimental settings line segregation is usually considered to have occurred when a subset of stimulus elements <i>pop-out</i> and become segregated from a stimulus background.
<i>Path-angle</i>	The difference in orientation of successive Gabor elements forming a contour. <i>See figure 1-5.</i>
<i>Path-paradigm</i>	An experimental paradigm developed by Field, Hayes and Hess (1993): Contours formed from Gabor patches are embedded in a background of randomly located Gabor patches. Two images are presented to the observer, one containing only the background and another also containing the embedded contour. Observers have to identify which image contains the

embedded contour.

Smoothness

A contour is said to be 'smooth' when the sign of path-angle changes between successive contour elements varies very little. Following this definition, the smoothest contour will feature no inflections, i.e. the sign of path-angle changes do not change throughout the contour. Ultimately the smoothest contour will be a circle (or a spiral) or perhaps a straight line, depending upon the values of path-angle and path-angle jitter.

Receptive field

A receptive field is the local mechanism which computes a filtered value at one point in an image.

Zero bounded region (ZBR)

A ZBR is a mutually-connected region of a filtered image, within which all values share the same sign. A description of a ZBR is a parameterisation of the values and locations encompassed within a ZBR.

1 Introduction

1.1 Overview

The Gestalt psychologists demonstrated that our experience of the visual world is often at odds with the actual content of the world. We often perceive structure in the world where none exists. Furthermore, we may apply structure in a consistent manner across different stimuli. For example, we may perceive a grid of dots as series of columns rather than rows. By examining the properties of images that give rise to particular perceptions, we can develop models of the processes underlying perception.

The phrase ‘contour integration’ describes a process through which contours that are formed from discrete elements can nevertheless be identified by the observer as a single object in the visual scene. The most obvious example of contour integration is our ability to recognise continuity in occluded lines. Imagine a discarded piece of rope lying within a patch of grass. We can easily identify areas of the visual scene that are *rope*. Furthermore, we can readily decide whether or not occluded parts of the rope are connected. We may even recognise whether two halves of a partially occluded blade of grass are continuous. In practice, experimental investigations of contour integration have examined the detection of contours composed of individual elements that are similar to each other. Such elements might be dots or rectangular line segments.

The underlying assumption in experiments investigating contour integration is that each individual dot or rectangle must be detected by a different processing unit. It follows that some additional mechanism must combine the information provided by each processing unit in order that the whole contour can be detected. The Gestalt psychologists (section 1.2) identified many stimulus properties that influenced contour integration. Their observations led to the creation of a number of rules. Each of these rules predicts how particular stimuli might be perceived by an observer. Unfortunately, these rules were largely descriptive. They did not explain the mechanisms underlying contour integration. Later experimental examinations of contour integration (section 1.3) attempted to explain how these rules might be implemented within our visual systems.

The development of the path-paradigm (Field, Hayes and Hess, 1993; section 1.4.1) is an important stage in the development of our understanding of contour integration. Field *et al.* argued that previous studies had not adequately ensured that integration was the only means of identifying the presence of the target contours. They designed their stimuli so

that other cues to the presence of a contour, such as element density and luminance, were eliminated. Furthermore, they used bandpass Gabor elements rather than rectangles. These stimuli were supposed to prevent coarse-scale filtering processes (section 1.4.5) from contributing towards the identification of embedded contours.

Recent models of contour integration are discussed in sections 1.4.2 and 1.4.3. Section 1.5 summarises current accounts of contour integration. Section 1.5 identifies unresolved issues within the literature that will be investigated within this thesis.

1.2 Gestalt psychology

The Gestalt school of psychology has identified many rules that predicted how particular arrangements of stimuli might be perceived by an observer. These rules tended to be descriptive rather than explanatory. Nevertheless, they provided a framework that has proven useful for the psychophysical research that followed. The main concern of Gestalt psychologists was that while the visual world is formed from many discrete perceptual phenomena, our *perception* of the world is of a coherent whole, literally a *Gestalt*. The passage below illustrates their emphasis:

“When we are presented with a number of stimuli we do not as a rule experience ‘a number’ of individual things, this one and that and that. Instead larger wholes separated from and related to one another are given in experience; their arrangement and division are concrete and definite.”, (Wertheimer, 1923, pg. 301).

The Gestalt school was concerned with identifying the rules that influenced our perception of structure in the world. Several of these rules are relevant to the current thesis. Wertheimer (1923) identified a number of principles that account for the grouping of individual stimuli. For example, Figure 1-1 illustrates how a *spontaneous arrangement* of dots might occur when certain spatial constraints are satisfied. In this instance, the factor of proximity has influenced the perceived arrangement. Dots within each vertical grouping (eg, *efghi*) are closer to one another than they are to other dots within the scene. Other factors, such as the *similarity* between stimulus items, were also found to influence grouping.




Figure 1-1. An example of spontaneous arrangement offered by Wertheimer (1923, pg. 302, figure iv). The dots are spontaneously arranged into groups of (a/bcd/efghi/jklmnop).

Wertheimer introduces the concept of *Prägnanzstufen* where the perceived arrangement of stimuli is in some way driven towards a particular representation. For example, the dots in Figure 1-1 are perceived as members of chevrons rather than as lines radiating from the 'a' dot. An example illustrating the *factor of direction* is provided in Figure 1-2. Here, we are more likely to perceive a semi-circle with a straight appendage than we are to perceive a corner with a semi-circular appendage. The perception of the whole semi-circle is dominant and the straight line is perceived as an appendage. Arrangements of stimulus dots that readily lent themselves to a particular co-linear perception were said to possess *good-continuation*.




Figure 1-2. This example illustrates the 'factor of direction', perception of a semi-circle with an attached straight line is dominant. (cf. Wertheimer, 1923. Figure 9 pg. 311).

Kohler (1929) introduced the concept of isomorphism, which suggests that perceptual phenomena correlate with physiological processes - Kohler himself cites Wertheimer as the founder of this concept. Consequently, one might suggest that the Holy Grail of the Gestalt school would be the discovery of the physiological substrate underlying the rules of perceptual grouping. Unfortunately, the science of cortical physiology was not sufficiently advanced to enable the achievement of such a goal. Ultimately the Gestalt school failed to explain the processes that underlay perceptual grouping phenomenon, nevertheless it did provide a useful basis for the analysis of perceptual phenomena.

1.3 Experimental investigations of contour detection

Uttal's (1975) experiments were concerned with the detection of target stimuli that differed from a background because of the geometric arrangement of constituent elements, in this

case a form defined by an arrangement of dots. He argued that such tasks must be accomplished in *pre-symbolic* occipital areas. Uttal's stimuli were formed from arrangements of dots because he argued that such stimuli avoided the potential confounds of photochemical and lateral neural interactions.

Measures of the detectability of curved lines revealed that straight lines were most easily detected. Detectability decreased with greater curvature. Uttal interpreted this pattern of results as an indication that a single cortical mechanism was responsible for the detection of both straight and curved lines. The increasing mismatch between the straight-line mechanism and the curved stimuli accounted for the poorer performance with increased curvature. Stimuli employed in this experiment were non-random. The relative magnitude and sign of angular offset of the path joining successive dots was constant for each target. Therefore, each of the contours featured within Uttal's experiments were continuous arcs.

Uttal reports that contour detection rates fell by as much as twenty percent when dots were irregularly spaced. As autocorrelation models are sensitive to periodicity within stimuli, this result was interpreted as support for Uttal's adoption of such a model. A further experiment revealed that lateral deviations in dot positions relative to the global contour also led to a reduction in contour detectability. Further experiments investigated the detection of various polygons such as triangles and diamonds. These findings are not discussed here, as they are not considered directly relevant to this literature review.

Uttal's autocorrelation model offered an example of how a model featuring an array of unspecialised processing elements might account for the stimulus feature sensitivities of observers in psychophysical experiments. Thus, his model is parsimonious, as it does not have to rely upon pattern matching operations. He offers a biologically plausible account of how such a model could be constructed from known neurophysiological elements. Though, Uttal accepts that models which rely upon spatial frequency analysis are related to his autocorrelational model. Furthermore, they would require a similar neural architecture. Thus, he suggests the important property of his model is the unspecialised nature of the individual processing elements, in this case the neurones.

Smit, Vos and Van Oeffelen (1985) offered a model (CODE-2) which attempted to represent the Gestalt phenomena of *proximity* and *good-continuity*. Their model achieved contour segregation by first estimating the relative proximity of neighbouring dots and then by calculating the extent of interaction between dots. The strength of dot interactions was determined by non-circular symmetric Gaussian distribution functions that were located

upon each dot. Each distribution function was then *stretched* along a major axis that was determined by calculating the proximity of neighbouring dots. A threshold was then applied to the overall function and perceptual grouping was predicted by the remaining supra-threshold contours.

The contour salience predictions of Smit et al's model were strongly correlated with human detection latencies. This implies some validity, at least at a functional level of description. Nevertheless, their model is only a more elaborate description of the interplay of Gestalt processes in contour segregation. The model does not explain how such processes might be implemented in a neurophysiological system. Smit et al. accept that the mechanisms employed within their model may not be biologically plausible, their only defence against this criticism is that other models, such as that offered by Caelli, Preston and Howell (1978), are even less plausible.

Beck, Rosenfeld and Ivry (1989) investigated the detectability of contours composed of line segments embedded in a background composed of similar line segments. In their first experiment, they investigated whether the detection of embedded contours was influenced by the alignment of line segment edges. They compared the detectability of straight contours that were formed from either 'bars' or 'blobs'. The 'bars' featured straight edges while the 'blobs' had irregular edges. The overall shape and aspect ratios of the bars and blobs were roughly equivalent. They found that, when line segments were aligned with the overall orientation of the contour, the detection of bars was faster and more reliable than for blobs. When the orientation of line segments was orthogonal to that of the contour, the blobs were more readily detected than the bars. Beck *et al.* interpreted this result as an indication that the straight edges of the bars contribute towards the detectability of the contour. Whereas these straight edges interfere with the detection of the lines composed of orthogonal bars, Figure 1-3 summarises their results.




Figure 1-3. (Left) Error rates for subjects detecting embedded contours composed of either bars or blobs. (cf. Beck *et al.*, 1989, figure 4b). (Right) Examples of the types of contours utilised in their experiment.

Their second experiment demonstrated that the lateral displacement of line segments further interfered with detectability. This phenomenon was confirmed at a number of spatial scales. A further experiment revealed that increased spacing of elements increased line detection latencies and errors. A third experiment revealed that there was no difference in the detectability of the solid and outline squares. They argued that this was because detectability must be based upon edges rather than whole squares. Their fourth experiment confirmed that contours in scaled-up images were detected more easily, even when contrast was controlled. Beck *et al.* accepted this as evidence that edge length was a key factor in contour detection. The fifth experiment featured Laplacian squares in addition to solid squares. They found no difference in observer performance between the solid and the checkerboard stimuli. Performance improved along with the edge length of both types of squares.

Beck *et al.* argued that if observer performance could be explained by coarse-scale filtering processes, then there should be a correlation between reaction times and the differences in the outputs of spatial-frequency channels for contour and background areas. In their final experiment, they found no such correlation and concluded that observer performance was not explained by coarse-scale filtering processes. This led them to conclude that line detection is the result of the integration of the outputs of discrete bar, edge or spot detectors. Rather than the result of a coarse-scale process that is simultaneously activated

by multiple line segments. However, their analysis assumes that the relationship between reaction time and spatial frequency channel output should be linear.

Moulden (1994) investigated the detection of contours formed from line segments that were hidden within a dense background of randomly positioned and oriented line segments. Moulden proposed a hierarchical model of contour integration. In this model local parts of the target contour were detected by individual filters. The outputs of these filters were integrated by *collator units* that are arranged hierarchically. Thus, collator units closer to the top of the hierarchy will integrate increasingly large portions of the target contour. Collator units were defined as second-order orientational filters. Moulden's psychophysical results indicated that collator units were insensitive to the both the contrast and colour of the individual line segments. This led him to suggest that simple cells might take the role of individual filters. The strongest criticism of this theory is that collator units must accommodate an almost infinite range of potential contours. Therefore, at the highest level a collator unit must exist that matches all potential contour shapes.

Each of the experimental studies described above has contributed towards an understanding of which stimulus parameters can effect contour integration. The important properties are detailed in Table 1-1 below.

Table 1-1. Summary of stimulus properties that can effect contour integration.

<i>Stimulus property</i>	<i>Correlation with contour detectability</i>	<i>Authors</i>
The number of line segments.	Positive	U B S M
Overall length of line	Positive	M
The regularity of the line segment spacing.	Positive	U
The spacing of line segments.	Negative	U S B
Density of the background elements.	Negative	U M
Curvature of the contour path	Negative	U
Lateral offset of contour elements.	Negative	U B

(Authors : B = Beck *et al.* 1990; U = Uttal 1975; S = Smit *et al.* 1985 ; M = Moulden 1994)

Principal axis orientation was also found to influence detectability by Smit *et al.* and Beck *et al.*, where vertical and horizontal orientations were favoured. However, Uttal did not record any effect of orientation. This was potentially due to the slightly different nature of the task undertaken by subjects. Line segment properties were also found to either support or interfere with integration (Beck *et al.*). Contrast and chromatic properties of segments appear to be unimportant (Beck *et al.*; Moulden), except where the detectability of elements is near threshold (Moulden).

1.4 Development of the path-paradigm

1.4.1 Field *et al.*'s experimental evidence for an association field

Field, Hayes and Hess (1993) devised an experimental paradigm that proved useful for identifying those stimuli properties that determined whether line-segments were integrated into a detectable contour. Field *et al.* embedded the contour elements into a background of elements that were matched for all properties except contour membership. Unlike the experimental paradigms described in the previous sub-section, the spacing density of contour elements and background elements was matched. Furthermore, rather than generating stimuli which consisted of simple line segments or dots they utilised Gabor patches. Oriented rectangles have a complex energy spectrum, especially at their corners. Hence, many cells in early visual areas would be activated following presentation of such stimuli. Consequently, Field *et al.* argued that utilising band-pass elements restricted the types of visual processing mechanisms that were likely to contribute towards detection of the contours. Furthermore, by employing band-pass Gabor patches they claimed that they were able to explore contour detection at a single spatial scale. Figure 1-4 shows an example of such stimuli.




Figure 1-4. A typical path-paradigm image. For illustrative purposes, the start and end elements of the embedded contour are indicated by arrows. The path-angle between successive contour elements is 20° and the mean spacing between elements is sixteen pixels, corresponding to 4λ .

Field *et al.* were interested in how contour detection might be achieved in a visual system composed of cells tuned for particular orientations and spatial frequencies. The underlying assumption of their study was that each embedded Gabor patch must be detected by an independent processing unit. Thus, an additional mechanism is required in order to detect a contour formed through the co-alignment of embedded Gabor patches. This mechanism

might form the substrate of good-continuation. In the context of the path-paradigm, the term contour integration refers to the process of identifying that a subset of Gabor patches within the stimulus are members of the embedded contour.

Field *et al* reported a series of five experiments that examined contour integration. Each one of these contributed towards the specification of a model that they termed the *association field*. They used a two alternative forced choice methodology. Two stimulus images were shown, but only one of these featured a deliberately embedded contour. The observer was required to indicate which of the two stimuli had contained the embedded contour. The results of these experiments are briefly summarised below and the most important findings are illustrated in Figure 1-5.

Field *et al.* placed their contour elements upon the centre of an invisible path. They found that detection rates fell as a function of the angle between successive path segments. This stimulus parameter will be termed '*path-angle*' for convenience. Their results indicated that detection rates approach a 100% level with a path-angle of 0°. Whilst detection rates fall to 75% for path angles of approximately 45°. When contour elements were oriented orthogonally to the contour-path the observers ability to detect contours was reduced. With orthogonally oriented Gabor patches detection rates fell to 75% with path-angles of approximately 15°. Detection rates approached chance levels when the orientation of elements relative to the path-segment was randomised by $\pm 30^\circ$. This manipulation is described as the addition of '*jitter*' to the path-relative element orientation.

In experiments 1-3 contour elements were separated by an average distance of 0.5° corresponding to 4λ (where λ is the period of the modulating sinusoid forming the Gabor element). In the fourth experiment, they manipulated the spacing of both background and contour elements. Performance improved by approximately 10% when element spacing was reduced to 1λ and fell by an equivalent amount when spacing was increased to 7.2λ . This was the maximum spacing achievable with their equipment.

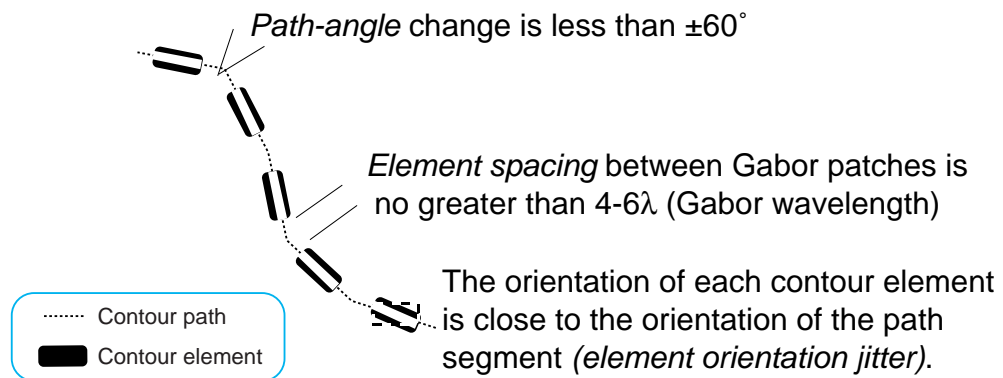


Figure 1-5. Some manipulations that influence the detectability of embedded contours. Italicised labels indicate terms used for commonly manipulated contour parameters.

Their final experiment repeated the stimuli utilised in experiment one, except that the phase of each Gabor patch sinusoid was randomised. This manipulation had no significant effect upon contour detection rates. Examples of phase manipulations are illustrated in Figure 1-6.




Figure 1-6. Examples of manipulations of the phase of contour Gabor elements.

The results of their experiments led Field *et al.* to conclude that facilitatory connections existed between neighbouring processing units within early visual areas. The hypothesised pattern of these connections led to the structure of the ‘association field’ illustrated in Figure 1-7.




Figure 1-7. Field, Hayes and Hess's (1993) conception of the 'association field'. The strength of facilitatory connections shared by the central receptive field (RF) and surrounding RFs is determined by the extent to which a smooth curve can be drawn between the central and surround major-axes. Thus, strong connections exist between the centre and the left-hand RFs whilst no such connections exist between the centre and right-hand RFs. (*cf. figure 16, pg. 190*).

The strength of these connections is determined by a number of constraints. The most important of these are the relative position and the orientation of the filters. The connection is stronger if the filters are near-neighbours and both are co-aligned with a smooth, uninflected, curve. Whilst the connection becomes weaker as the curvature or the length of the line that intersects the two receptive fields increases. Field *et al.* suggest that these connections may join complex cells, as observer performance on the path-paradigm task was unaffected by the randomisation of the phase of Gabor patch elements.

The association field model suggested by Field *et al.* is not equivalent to Moulden's collator unit model. This is because the latter is hierarchical in nature and hard-wired, whilst the former is not. The association field model posits that the integration of receptive field outputs occurs within a single cortical layer, possibly the long-range horizontal connections of the primary visual cortex. Whilst for Moulden's model contour integration is necessarily staged, larger segments of the contour are integrated by collator units that are higher within the hierarchy.

The proposed structure of the association field is informed by many developments within the field of neurophysiology. The following sub-section summarises recent neurophysiological evidence of connections within primary visual areas that may lend themselves to integrative functions.

1.4.2 The extension of the classical receptive field.

Within the striate cortex long-range horizontal collateral connections exist which are thought to be *intra-cortical* axonal links between columns with similar orientation preferences (Ts'o, Gilbert and Weisel, 1986). Studies that have stained single cells with

horseradish peroxidase have revealed that these connections can link neurones 4mm apart (Gilbert and Weisel, 1983; Rockland and Lund, 1993). Some studies have reported connections up to 8mm in length (Hirsch and Gilbert, 1991; Gilbert, Das, Ito, Kapadia and Westheimer, 1996). In area 17 of the cat, a 6mm lateral connection at 5° eccentricity corresponds to a retinotopic separation of approximately 6° (Tusa, Palmer and Rosenquist, 1978). Whilst there are slight discrepancies in estimates of the extent of long-range horizontals, it appears likely that their range will correspond to a visual angle of little more than a few degrees. The studies of Ts'o *et al.* have indicated that long-range horizontals tend to connect cells that differ in their orientation preference by no more than 30-40°. Gilbert (1992) has demonstrated that horizontal connections do not just link cells with a similar orientation preference. They also link cells with the same colour selectivity, ocular dominance channel etc. Bosking, Zhang, Schofield and Fitzpatrick (1997) examined the structure of long-range horizontal connections within the primary visual cortex of the tree shrew. They report that connections longer than 500µm preferentially link receptive fields that are both co-axial and co-aligned.

Recent physiological studies have suggested that the striate cortex is the neurological substrate for the task of contour integration (Knierim and Van Essen, 1992; Kapadia, Ito, Gilbert, Westheimer, 1995; Gilbert *et al.*, 1996). The Kapadia *et al.* study is summarised below, as it features similar stimuli to those used in the path-paradigm experiments of Field *et al.*

Kapadia *et al.* conducted the same psychophysical experiments with human observers and macaques (*Macaca Mulatta*). They measured the detection of a low contrast bar whilst the presence and position of neighbouring bars was manipulated. Responses of complex cells in area V1 of the macaques and the contrast thresholds of human observers were measured. The spiking rate of V1 cells showed a close correspondence to the performance of their human observers. For their human subjects the contrast detection was improved by forty percent when a single surrounding bar was placed in a collinear position. For the same stimuli forty two percent of the V1 cells that they recorded from showed an increased spiking rate. The pattern of results indicated that the second bar was not activating the receptive field directly. When the central bar was removed then the surrounding bar had no influence upon the response of the cell being recorded. The surrounding bar only had a modulatory influence upon the response of the cell.

Kapadia *et al.* found that facilitation was dependent upon the collinearity and proximity of the centre and surrounding bars. In a further experiment, analogous to the path-paradigm of Field *et al.*, they presented the central bar and surrounded it with a background of randomly oriented bars. When the background bars were present, performance was reduced. For macaques simple-cell firing rates fell and for human observers contrast threshold was increased. The surrounding bars suppressed the cells response to the central bar. This suppression was eliminated if a subset of the background elements were rotated so that they were collinear with the central bar. Their results indicate that the response of a receptive field can be influenced by stimuli presented beyond what has been classically defined as the receptive field. Kapadia *et al.* concluded that these modulatory influences occurred because of interactions between the outputs of neighbouring receptive fields. They suggested that the long-range horizontal connections of area V1 formed the substrate of these interactions.

Kapadia *et al.* used stimuli formed from rectangular bars whereas Field *et al.* employed Gabor elements. Field *et al.* argued that the use of rectangular bars as line segments does not adequately control for the potential contribution of coarse-scale mechanisms. However, the single unit responses in their macaque experiments suggest that such mechanisms were unlikely to have been operating during their experiment. Stimuli featuring surround stimuli without a central bar evoked little response in the monitored cell. If a coarse-scale mechanism had been operating, each bar element that influenced firing rates when the central bar was present should also have provoked a weaker response, even when the central bar was not present. Therefore, their results suggest a modulatory, as opposed to additive, influence of the surrounding bars. Kapadia *et al.* did not measure cell responses to curved lines embedded in a random background. That is, the bars within the background that they termed 'flanks' were always collinear with the central bar. If they had undertaken this manipulation then it would have been possible to make a direct comparison with the path-angle manipulations of Field *et al.*

The traditional characterisation of a striate cortex receptive field as sensitive to a small edge or contrast segment (Hubel and Wiesel, 1962) has gradually been modified. This modification accounts for the evidence that stimuli beyond the receptive field can influence firing rates within the receptive field. The area beyond the receptive field capable of influencing the neurones response has been variously termed the *integration field* (e.g., Li

and Li, 1994) and the *contextual field* (Phillips and Singer, 1999). These physiological concepts can be considered analogous to the association field suggested by Field *et al.*

There is considerable correspondence between the physical extent of long-range horizontal connections and the detection limits identified by the path-paradigm task. Long range horizontal connections link simple-cells that differ in orientation by as much as 45° (T'so, *et al.*, 1986). Contour detection rates fell to levels of around 75% when the path-angle between subsequent contour element is 45° (Field *et al.*, 1993). Field *et al.* report that contour integration occurred even when elements were separated by 0.9° visual angle, larger separations were prohibited by the limitations of their equipment. Long-range horizontal connections can link simple-cells that are as far apart as 5° visual angle (Hirsch and Gilbert, 1991; Gilbert *et al.*, 1996).

Citing these correspondences, many authors have claimed that the long-range horizontal connections may form the physiological basis of the contour integration (Field, Hayes and Hess, 1993; Kovaks and Julesz, 1993; and Polat and Sagi, 1992). Furthermore, when the interaction between stimulus elements occurs over visual angles greater than those spanned by the long-range horizontals authors have suggested that facilitation can 'cascade' through a chain of such local connections (Polat and Sagi, 1994). Giersch, Boucart, Danion, Vidailhet, Legrand (1995) found that observers administered lorazepam, performed worse during a contour matching task. Lorezepam interferes with the neuro-mediator GABA. They concluded that disruption of GABA channels interfered with the operation of the long-range horizontal connections and that these connections form the substrate of contour element binding.

Neurophysiological evidence suggests that receptive fields will fire synchronously if they are detecting line segments belonging to the same contour (see Singer and Gray, 1995, for a review). Singer *et al.* suggested that the long-range horizontals collaterals of area V1 may form the neural substrate for this phenomenon. Many of the contour integration models discussed in Section 1.4.4 incorporate this process.

1.4.3 Evidence suggests that the mechanisms underlying contour integration are acquired through visual experience

Converging evidence from neurophysiology, statistics from natural scenes, computer simulations and clinical research suggests that the mechanisms underlying contour integration may be developed as a consequence of early visual experience.

Kiorpes, Bassin & Movshon (2001) have reported that pigtailed macaques (*Macaca Nemistrina*) develop contour integration abilities later than they achieve adult levels of visual acuity. Kiorpes *et al.* examined the development of contour integration from 3 months to 2 years. Circular path-paradigm contours were embedded in a background formed from randomly oriented Gabor elements. Detection thresholds were measured whilst background density and contour element orientation jitter were manipulated. They found that infants (<5 months) were unable to detect contours without a density cue. Performance gradually improved towards adult levels over the course of two years. They speculated that contour integration mechanisms could only be established after early visual filters had reached adult acuity levels. As contour integration mechanisms took longer to develop than visual acuity measures, they argued that these mechanisms operated upon the outputs of early visual filters.

Evidence from studies which have examined co-occurrences of line segments in natural images suggest that sufficient information exists in natural scenes to direct the development of contour integration mechanisms (Giesler, Perry, Super and Gallogly, 2001; Sigman, Cecchi, Gilbert and Magnasco, 2001). Sigman *et al.* (2001) examined how correlations in the orientation and location of line segments and edges varied with their relative locations. They found that a co-circularity rule provided a parsimonious account of these correlations.

Geisler *et al.* analysed co-occurrences of segment locations in natural images. They used 2 methods, absolute co-occurrences and Bayesian edge co-occurrences. Both methods led to successful models. The responses of the model to path-paradigm stimuli correlated very closely to those of human observers (Pearson statistic = 0.87). They also reported that the best fitting parameters (to those of human observers) also gave rise to the best overall performance. Edge co-occurrence estimates were based upon the line tracing abilities of observers, these were required in order to indicate which segments belonged to the same contours. Contour membership was represented by a transitivity rule : *i.e.* “*if (edge element) a binds to b, and b, binds to c, then a binds to c*”. They offer little evidence to support the existence of such a rule, but speculate that this is likely to be a widespread heuristic within the cortex. The model successfully predicts the 'closure' effect described by Kovaks and Julesz (1993), though the strength of this effect has been questioned by Braun (1999). Giesler *et al.* also offer preliminary confirmation, in natural images, of Field *et al.*'s. (1993) assertion that edge co-occurrences were consistent across spatial-scales.

Finally, the model that they offer was capable of accommodating top-down influences by enabling local modulation of grouping strengths.

Feldman (2001) accepts that co-circularity rules can account for the grouping of pairs of elements but argues that these rules do not account for the subjective experience of curvature. Feldman presents a Bayesian model which successfully predicts whether a particular arrangement of dots would be perceived by an observer as either a smooth contour or a corner. Thus, Feldman's model extends existing models and attempts to provide a framework for the appropriate selection and application of alternative grouping algorithms. It is hypothesised that the selection of algorithms is based upon whether previous applications have correctly identified whether edges or line segments belong to the edge or surface.

Recent studies have found that strabismic amblyopes are relatively impaired in the detection of path-paradigm contours in comparison to normal individuals and to their own unimpaired eyes (Hess, McIllhagga and Field, 1997, *ibid* and Kovacs, Polat, Pennefather, Chandra and Norcia, 2000). Hess, McIllhagga and Field argue that the poor contour integration performance of strabismic amblyopes results from positional uncertainty (Hess, Campbell and Greenhaigh, 1978 *ibid*). The positional uncertainty hypothesis argues that normal interactions occur between receptive fields with disordered positions. Kovacs *et al.* (2000) recorded contour detection performance with closed-contours and found similar patterns to those described by Hess, McIllhagga and Field. Amblyopia generally results from abnormal visual experience during critical periods of visual development.

The evidence presented above is readily accommodated by a framework which proposes that association fields are developed through visual experience. Neurophysiological evidence has shown that lateral connections link the outputs of similarly oriented receptive fields in the primary visual areas (Gilbert, 1994). Following Hebbian learning principles (Hebb, 1949), it is reasonable to suppose that these lateral connections could be created and/or strengthened through correlations in receptive field activity. As a consequence, the strength of lateral connections would be determined by those properties in the visual world which are correlated. Both Sigman *et al.* (2001) and Giesler *et al.* (2001) present evidence for the existence of correlations between neighbouring oriented areas of the visual scene. The overall pattern of these correlations is readily accommodated by a co-circularity rule which is weighted by the difference in orientation preference and proximity of receptive fields (for example see Yen and Finkel, 1998). Finally, when visual experience is

disrupted, as is the case in of strabismic amblyopia, the development of association field mechanisms is impaired.

However, Hoyer and Hyvärinen (2002, *in press*) found little evidence for the development of co-circularity based rules in their model. Using a sparsely coded network trained according to information theoretic principles, they arrive at a hierarchical model. The higher-level cells within this model responded selectively to the presence of straight contours. They argue that such cells may be similar to the ‘collator units’ described by Moulden (1994). Hoyer and Hyvärinen demonstrate how these cells operate in a top down fashion, strengthening the outputs of receptive fields that were only weakly activated by incomplete contours within the incoming image. They speculate that curved contours may be detected by units responding to the outputs of the ‘collator units’.

1.4.4 Computational models that are deliberately sensitive to smooth contours

A number of models have attempted to explain the processes underlying contour integration. Each of these has claimed varying degrees of biological plausibility. Some even featuring discrete receptive fields located within hyper-columns that have a range of orientation preferences (Li 1998; Yen and Finkel, 1998). Most feature excitatory and inhibitory lateral connections between neighbouring receptive fields. Often these connections determine the models preference for contours. This section will briefly describe a number of current models whilst trying to highlight those elements of each model that are primarily responsible for the selection of embedded contours.

Gigus & Malik (1991) offered a model that was able to identify noisy curves within images. The first stage of their model involved rotating linear filters until the function of the convolved image was maximised. This identified local oriented areas within a response image, either for edges $E(x, y, \theta)$ or lines $O(x, y, \theta)$. The outputs of these response images were combined, leading to the identification of local dominant orientations. Dominant curves within a particular area are maintained whilst less pronounced curves are suppressed. This collator mechanism leads to a further response image $C(x, y, \theta, k)$ where k is the curvature. This is defined as the reciprocal of the radius of the circular arc.

Gigus and Malik compare the predictions of their model with the performance of human observers described by Beck, Rosenfeld and Ivry (1990), see Section 1.3 for a summary. The usefulness of this comparison is limited, as the Beck *et al.* study is concerned with the detection of embedded straight lines. They do not attempt to offer any comparison of curve detection for their model and human observers.

Yen and Finkel (1998) offered a model featuring temporal synchrony. Their model posited that modulatory interactions exist between cells and that these interactions occur via the long-range horizontal connections. They also claimed that temporal synchronisation of receptive field outputs was necessary to signal contour salience. Initially images were segmented using steerable filters (Freeman & Adelson, 1991). Within their model receptive field interactions are defined as modulatory because the cell will not fire if surround area alone is activated. This corresponds with the neurophysiological evidence offered by Kapadia *et al.* Receptive fields that only have a modulatory input can be ‘steered’ towards the dominant orientation of neighbours. Contours are not completed within their V1 model. However, the sub-threshold activation of receptive fields that only receive modulatory input offers the potential for the occurrence of illusory contours at a higher level. They suggested that this may correspond to the results of von der Heydt and Peterhans (1989) who have shown that illusory contours occur within V2.

Outputs of steerable filters then activate units that receive facilitatory modulatory connections from co-linear neighbours. The pattern of connectivity for these facilitatory interactions is defined by the Parent and Zucker’s (1989) co-circular connection scheme. Further, less extensive, facilitatory connections exist in a ‘trans-axial’ direction, i.e. for other, parallel, receptive fields flanking the target unit. The pattern of both groups of facilitatory connections is shown below in Figure 1-8. The presence of both trans-axial and co-axial connections suggests that the model is actually formed from two independent models. That is, both models are never tested simultaneously.




Figure 1-8. The pattern of facilitatory connections in the Yen and Finkel (1998) integration field. Here the central rectangle represents the preferred orientation and position of the target receptive field. Whilst the surrounding segments represent the orientation of a neighbouring element which will have the strongest facilitatory influence upon the target. The length of each surrounding line segment represents the relative strength of the facilitatory connection.

Inhibition takes a slower time-course than facilitation and serves to suppress activity for elements that have not previously received facilitation, i.e. those elements likely to belong to the stimulus background. Strong reciprocal-facilitatory connections link RF’s. These

connections feature coupled neural oscillators. Thus, elements that are members of different contours within the same stimulus image are likely to be desynchronised. They proposed that the salience of a contour is represented by the total of the activity across a chain of synchronised RF outputs. Yen and Finkel offer a number of simulation results that demonstrate congruency with both psychophysical and physiological experimental results.

Their model provided similar results to those of Kapadia et al. (1995) who reported a comparative study comparing single-unit recordings in alert macaques with human psychophysical data. The model exhibited facilitation of a low contrast target if high-contrast neighbours flanked the target. Kovaks & Julesz (1993,1994) demonstrated lowered contrast sensitivity for a Gabor target when it was situated within the bounds of an elliptical contour. The Yen and Finkel model showed a corresponding change in synchrony, largely due to the structure of the trans-axial facilitatory connections. The degree of synchronous firing within the model indicated the stimuli salience. They offered extensive comparisons of their models performance and Field *et al's* contour detection results. Stressing that their model offered similar results for manipulations of the following: contour path-angle, with co-linear and orthogonal element orientation; path-relative element rotation; and the spacing of contour elements. Their model exhibited an increased measure of contour-salience for contours that were closed, this matches the closure-effect reported by Kovaks & Julesz (1993). Finally, they offered an example of real-world image processing. They demonstrated that their model was capable of finding the outline of a camouflaged object. Despite the apparent success of their model, it is important to stress that the structure of their integration field largely determines the models preference for contours.

Pettet, McKee and Grzywacz (1998) described a three stage model which exhibits performance equivalent to human observers in contour detection tasks. In the first stage, receptive field units within their model were tailored to the stimuli employed in the experiment. In fact only those units optimally activated by a single Gabor patch received any further processing in later stages. In the second stage of their model receptive field co-facilitation was based on three factors, (i) the spatial separation of receptive fields, (ii) the length of splines fitted between pairs of co-activated receptive fields and (iii) the curvature of the splines. Model responses were based upon the presence of outliers in the response population of active RF outputs, i.e. images with embedded contours featuring a sub-set of RF's which were more active than their neighbours. Pettet et al. found that modifying the

relative importance of the three influences upon RF co-facilitation provided outputs that were equivalent to those of human observers. Specifically, the model showed similar levels of success in indicating the presence of embedded circle, bean and moon shaped contours.

Clearly, as with the other models described in this section, the three factors that they chose to represent co-facilitatory connections are a re-description of a smooth contour. By its very nature their model is tautological, manipulating their models parameters until co-facilitation occurs between elements which are members of a smooth extended contour will inevitably cause the model to favour such stimuli. Demonstrating that particular parameter sets make the model selective for the smoothest contour does not extend our understanding of contour integration. It only indicates that local interactions between elements can contribute to global salience. Their decision to allow only optimally active receptive fields to contribute to the processing of response outputs inevitably excludes any contribution from receptive fields that are activated by more than one contour Gabor element. This eliminates any potential for a coarse-scale solution to the problem of contour detection.

Li (1998) offered another biologically plausible model. It is similar to that offered by Yen and Finkel, except that it also has some properties derived from the physiology of V2 including modulatory feedback from V2 into V1. In this model the initial input layer codes locations and orientations of individual line segments within the stimulus pattern. The initial stage within this model corresponds loosely with the second stage within Yen & Finkel's model. Individual neuronal assemblies were simulated, modelling cellular membrane potentials and cellular interactions. However, the model's preference for smooth extended contours was largely due to the structure of the facilitatory 'horizontal' connections. The details of the neuronal interactions are ignored here, only the pattern of facilitatory connections will be considered. As it is the structure of the facilitatory connections that determined the models sensitivity for extended contours.

In Li's model the strength of excitatory connections between receptive fields was determined by three factors: (i) the similarity orientation preference; (ii) the proximate of each cell; and (iii) whether both cells fell upon a projected line determined by the orientation preference. Neighbouring cells that were located orthogonally to the preferred orientation of the receptive field (i.e. flanks) tended to be inhibitory regardless of their orientation. Thus, the spatial extent and pattern of facilitatory connectivity between neighbouring elements closely resembled that achieved when connectivity is defined using the co-circular rule of Parent and Zucker (1989). Consequently, receptive fields, which fell

upon smooth contours with a low level of circularity, were most likely to have strong excitatory links.

Li offered a number of simulation examples in which the magnitude of firing synchrony is affected by the properties of the stimulus input. The greatest level of firing synchrony occurred for a straight long contour and it was progressively weaker for a circular contour, a broken circle and finally for a short arc. Thus, Li argued that salience in the model, as represented by the degree of synchrony, is equivalent to the degree of contour salience for human observers. Further examples were offered in which the dual representations of a pair of embedded contours were disambiguated from one another by the feedback from higher areas. None of the example stimuli include a situation where the density of background elements and contour elements was matched. All of the stimuli described by Li feature a rather sparsely sampled background. Therefore, it's not clear whether contours would still be detected if the density of contour and background elements were matched.

The models offered by Yen and Finkel and Li, featured synchronous firing mechanisms. This mechanism leads to enhanced detection of closed contours. This occurs because synchronised firing patterns become stronger in closed loops of receptive fields. Interestingly the strength of the effect of closure has recently been questioned. Braun (1999) has demonstrated that when the overall extent of the contour was controlled the closure effect is actually very weak – and perhaps due to the inevitable remaining difference in eccentricity. Yen and Finkel emphasised the relevance of the closure effect, stressing the role of synchrony in its occurrence. However, many dispute whether synchrony is actually employed in perceptual grouping tasks (for example, Kiper, Gegenfurtner and Movshon, 1996).

The pattern of connectivity of the excitatory connections favoured smooth contours with a low curvature. Whilst the mechanisms featured within the model may be considered biologically plausible, the mechanisms responsible for the identification of potential contours are often a re-description of the techniques used to generate the contour. The pattern of connectivity between receptive field outputs embodies a description of the properties of a smooth contour. Hence, the model is circular. The generation and selection mechanisms both code the structure of the contour. Clearly, the existence of mechanisms within the visual system which match the properties of the incoming signal are probable and perhaps even desirable. However, it is important to determine whether it is necessary for this property to exist in order to solve the particular perceptual problem of contour

integration. Thus, the author suggests that these models may not be applying Occam's razor; they may feature more mechanisms than are required in order to solve the problem of contour integration.

It is important to stress that whilst particular mechanisms may be unnecessary when attempting to solve the problem of contour integration, there may be other tasks that require these mechanisms. For example, the Yen and Finkel (1998) model featured synchronised neuronal firing. This is not necessary for the achievement of contour integration. Co-facilitatory connections could simply influence the gain of the receptive fields. Consequently, contour and background elements could be segregated according to their contrast. Alternatively, if the perceptual task involved the detection and segmentation of two embedded contours. Then segmentation may be facilitated if elements belonging to the same contour shared synchronised firing patterns that differed between contours. The implementation of synchronised firing offers a potential solution to the binding problem.

1.4.5 Alternative explanations of Field *et al*'s findings

Field *et al.* offered alternative accounts of contour detection. One featuring pre-wired path detectors and another based upon a parallel filtering mechanism. They claimed that various properties of their stimuli precluded any contribution such mechanisms could make towards contour detection. Schematic representations of the association field mechanism, the path-detector and the parallel filter process are presented in Figure 1-9.

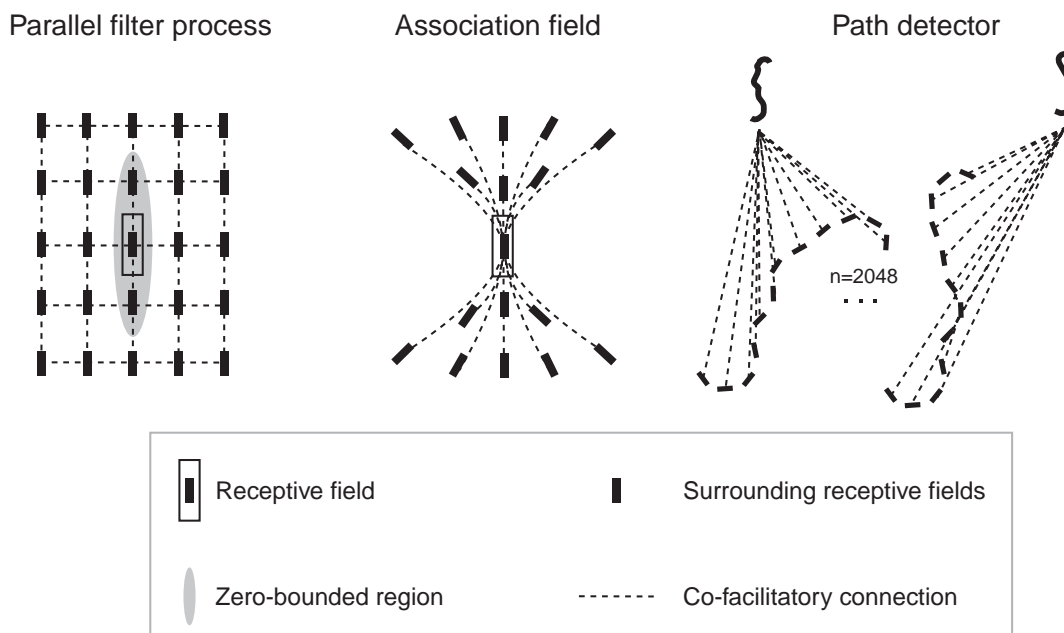


Figure 1-9. Schematic representations of the three solutions to contour detection offered by Field *et al.* (left) A parallel filter process which involves the convolution of the stimulus image and the subsequent extraction of a zero-bounded region. (middle) The association field. (right) Path-detectors, all possible path configurations are represented by hard-wired connections between filter outputs.

Path detectors were described as mechanisms that are hard-wired for particular contour configurations. Field *et al.* argued that it would be unlikely that path-detectors would exist that were matched to each of their randomly shaped and positioned contours. The contours employed in their experiments featured twelve Gabor elements with random changes in the direction of rotation between successive elements. For the path detector hypothesis to work, they claimed that 4098 different path detectors would be required at every spatial location and for every degree of path-angle. This is clearly an over estimation. If all of the hypothesised path detectors existed then subject performance would not fall as a function of the various contour parameters such as path-angle and element spacing. A path detector would exist that matched every one of their experimental stimuli. Regardless of this over estimation, the path detector mechanism appears the least plausible of the alternative hypotheses offered. The path detector solution is likely to be the most difficult hypothesis to disprove with psychophysical methods alone. However, it would appear that parsimony would offer sufficient justification for the rejection of this hypothesis.

Field *et al.* described a 'parallel filter process' (pg. 176) which involved the convolution of the stimulus image with an oriented filter. A filter combines a subset of values about a point in the stimulus image in a consistent manner and this produces a single value, known

as the filtered output. Convolution is the process of applying the filter at each point in the stimulus image. The result of convolution is an array of filtered outputs, described as a filtered image. Each point in this array corresponds to a point in the stimulus image. However, Field *et al.* did not explain how the subsequent filtered image might actually indicate the presence of an embedded contour. The filtered image was merely presented (Field *et al.* 1993, figure 4, pg. 178; Figure 1-10 below for an example). Their figure heading stated that ‘*The path cannot be segregated from the “noise” background by filtering along a single dimension...*’ (Field *et al.*, 1993, pg. 178). This implies that a segmentation process occurred following convolution.




Figure 1-10. A typical path-paradigm image [left]. Field *et al.* argued that an embedded contour could not be identified as a consequence of filtering with a vertically oriented bandpass filter. Thus, they claimed that identification of the embedded contour necessarily required the operation of a dynamic integration process. For illustrative purposes the start and end elements of the embedded contour are indicated by arrows. [right] The result of filtering the stimulus image with an oriented bandpass filter [middle]. The path-angle between successive contour elements is 20° and the mean spacing between elements is sixteen pixels, corresponding to 4λ . Filter dimensions were $1.5\sigma \times 3.0\sigma$.

This position was somewhat clarified by Hess and Dakin (1999). They described a process of ‘... *simple linear filtering without any subsequent orientation linking operations...*’ (Hess and Dakin, 1999, pg. 954). In practice, their simple filter model involves convolving the stimuli with a two-dimensional difference-of-Gaussian filter. Subsequently image values within one standard deviation were set to zero. This creates an image that contains zero-bounded regions (ZBR). These are mutually-connected regions of the filtered image within which all values share the same sign. Finally, Watt’s image description scheme (Watt, 1991) was used to parameterise the location and the values within each ZBR. The processes of convolution and the parameterisation of ZBRs is illustrated within Figure 1-11.

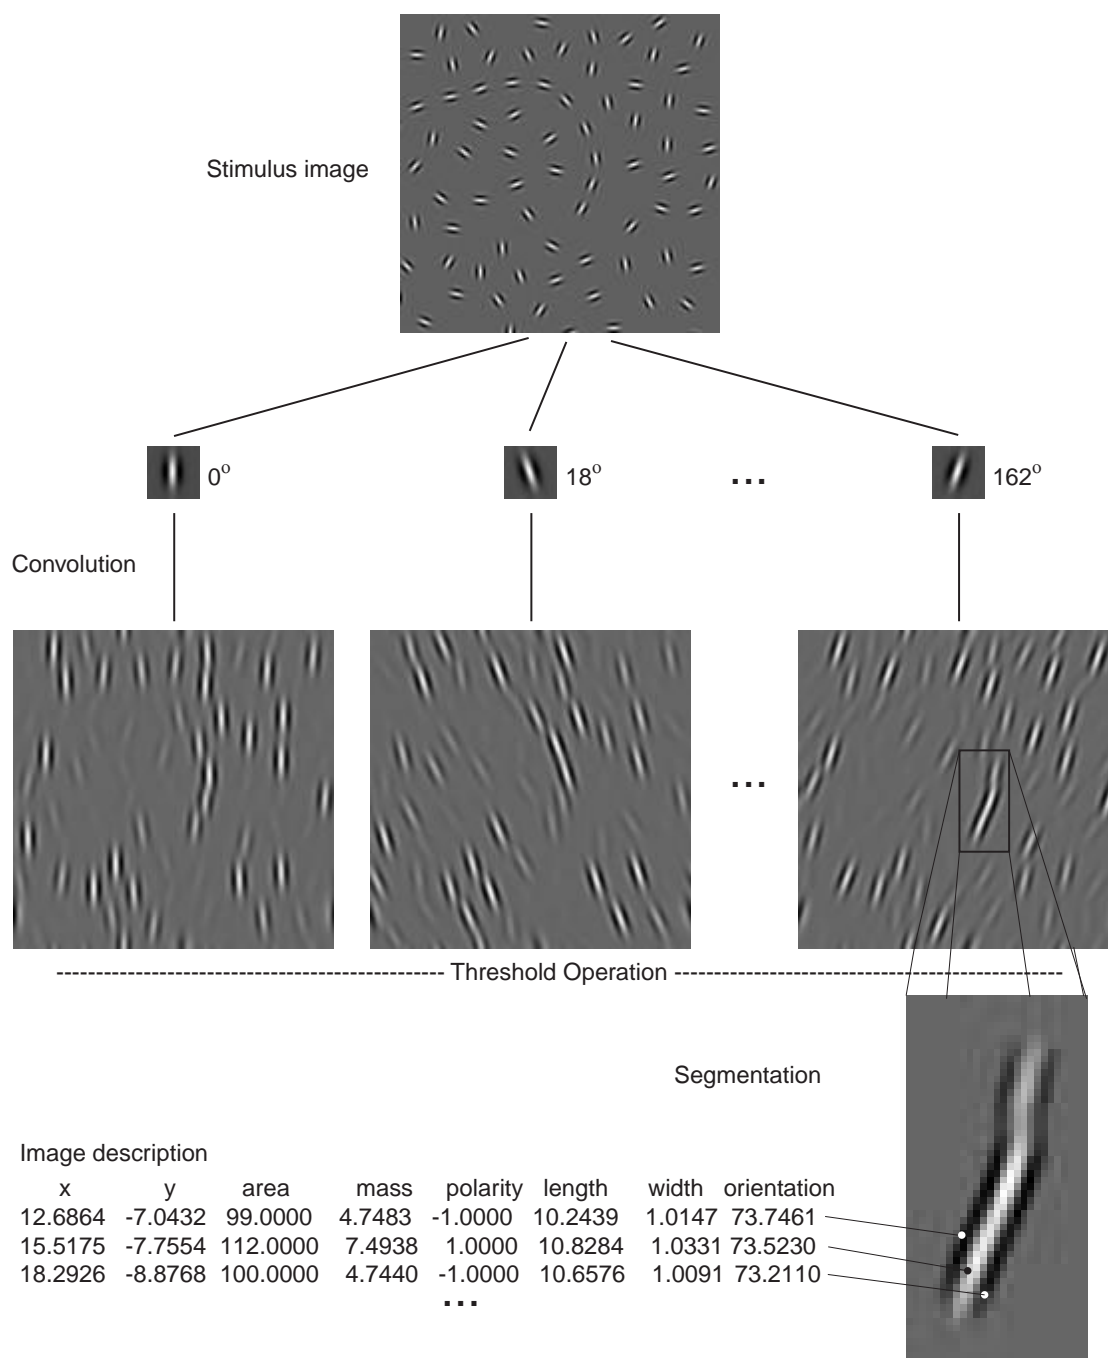


Figure 1-11. Outline of the image filtering and segmentation process. The first row shows the initial raw image. The second row shows the filters used within a particular spatial-frequency. The third row shows the result of convolving each filter with the raw image. Finally, for illustrative purposes the fourth row shows the filtered image section that contains the longest zero-bounded regions. The parameterisation of the zero bounded regions are shown to the left.

The path-paradigm task utilises a two-alternative forced-choice (2AFC) methodology. Images are presented in pairs. One image, the target image, within a pair contains a deliberately embedded contour. The remaining image, the noise image, does not contain a deliberately embedded contour. Therefore, the simple-filter model must identify one of

these images as the target image. This is achieved by comparing the maximum ZBR length within each 2AFC pair. The image containing the longest ZBR is selected as the target.

Field *et al.* argued that a parallel filtering mechanism was not sufficient to account for the contour detection performance of observers. They claimed that the bandpass nature of their stimuli and the fact that neighbouring elements along a contour differed in their orientations precluded such solutions to the contour integration task. They imply that as the parallel filter process only combines the outputs of similarly oriented filters, the ZBR that is formed cannot represent the whole contour. It is clear that a ZBR is unlikely to correspond to a whole contour. However, this should not prevent the parallel filter from performing successfully in the path-paradigm task. The task requires the selection of the image containing the embedded contour, not the identification of the whole contour.

Hess and Dakin (1997, 1999) compared the performance of the simple filter process with the performance of human subjects. For the human observers contours were presented to the fovea. Such contours were also presented to the peripheral field, this manipulation is described in Section 1.4.7. When contours were formed from symmetrical-phase Gabor elements, detection rates for human observers were approximately 20% better than the simple filter process. Detection rates for the simple filter process fell to chance when alternated-phase elements were used. Conversely, human observer performance was unaffected by the type of Gabor element used. This led Hess and Dakin to conclude that foveal contour detection could not be explained by a parallel-filter process.

The parallel filter process described by Field *et al.* is arguably equivalent to the simple filtering process described by Hess and Dakin. For the sake of clarity, the term ‘simple-filter model’ will be used henceforth throughout the current thesis. This term describes a process of convolution of the stimulus image with an oriented filter followed by a thresholding operation. A symbolic description of the zero bounded regions present within the filtered image is created using the Watt image description scheme (Watt, 1991). Finally, the identification of the maximal ZBR length within the target and noise images informs the 2AFC decision. The image containing the longest ZBR is selected as the target.

1.4.5.1 Cellular linking in the simple filter model?

The use of the phrase ‘...without any subsequent orientation linking operations...’ by Hess and Dakin (pg. 954) suggests that the outputs of differently oriented filters are not combined. This implies that they accept that the outputs of filters that share the same orientation may be indeed be linked. This interpretation seems to conflict with an earlier

statement within the same article, which suggests that no linking occurs between any filter outputs. They stated that ‘...*these results can be modelled by simple filtering without any subsequent cellular linking interactions.*’ (Hess and Dakin, 1999, pg. 947, emphasis added). The latter interpretation conflicts with the definition of a ZBR. A ZBR is a mutually-connected region of the filtered image, this region is likely to encompass multiple filter outputs. Therefore, the parameterisation of the ZBR necessarily implies the linking of filter outputs. If it is accepted that a single receptive field is the local mechanism computing a filtered value at one point within an image, then computing a ZBR parameterisation must involve the linking of receptive field outputs.




Figure 1-12. Illustration of a large zero bounded region that corresponds to the embedded contour within a stimulus image. (left) stimulus image featuring an embedded contour with a path-angle of 25° , (centre) filter [width = 3σ , length = 5σ , $\theta = 90^\circ$], (centre) longest zero bounded region within the image.

Figure 1-12 illustrates a large ZBR that corresponds to a significant proportion of a contour. To propose that this ZBR could be parameterised without any cellular linking operations would appear implausible. This would imply that the whole ZBR was detected by a single receptive field. Such a mechanism might be considered equivalent to the hard-wired path-detector proposed by Field, Hayes and Hess (1993) – see Figure 1-9 (left). Therefore, this conception of the parallel filter model ought to be considered equally implausible. However, it is presumed that any neural implementation of a simple-filter model must involve cellular linking between the outputs of similarly oriented receptive fields.

Another possible interpretation of the Hess and Dakin’s sentence, quoted in the previous paragraph, remains. Hess and Dakin may believe that there was no need for further linking

operations following the parameterisation of a ZBR. Such a mechanism might lead on to an additional stage of image processing which identifies relationships between the parameterised ZBRs. It is this definition of the simple-filter model that is accepted within the current thesis. In physiological terms the simple-filter model describes a mechanism that integrates the outputs of receptive fields that are co-oriented and spatially contiguous.

1.4.6 Effects of global contour characteristics upon contour integration

The association-field model of contour integration described by Field *et al.* suggests that integration is based solely upon the pair-wise interaction of neighbouring contour elements. Their experimental results and the presented replications (Sections 2.2 and 2.3) offer a pattern of evidence compatible with such a model. For this model, so long as local properties such as element spacing and path-angle are held constant, the influence of the overall shape of the contour upon its detectability is unclear.

The experimental results of Kovacs and Julesz (1993) demonstrated an increase in contour salience as the ends of the contour met, *i.e.* when contours were ‘closed’. They identified a ‘coherence distance’ between neighbouring contour elements where contour detection was 75% correct. This distance was 3.3λ for closed contours and 6λ for open contours. Where λ is the wavelength of the sine wave of the Gabor patches.

Their results suggested that for a given pair of contours (*A* & *B*) with matched local properties, but with differing global properties, the detectability of the contours will vary. For example, contour *A* is ‘open’ *i.e.* its beginning and end do not meet (Figure 1-13, A), and contour *B* is ‘closed’ *i.e.* its beginning and end are proximate (Figure 1-13, B). Kovacs and Julesz’s model predicts greater salience for contour B whilst Field *et al.*’s association field model makes no explicit distinction between the salience of the two contours.




Figure 1-13. Contours with varying levels of 'closure'. (A) a random contour, (B) a ‘closed’ contour. The local properties of the contour are matched, *i.e.* element spacing and path-angle.

Kovacs and Julesz (1993) suggested a model that contrasts with that of Field *et al.*, as their model predicts that the global properties of contours may also have an influence upon their salience. Specifically, Kovacs and Julesz suggest that a ‘synergistic process’ (pg., 7495) exists which increases the facilitation between contour elements when the contour forms a closed path.

Whilst there is an apparent conflict between the claims of Field *et al.* (1993) and Kovacs and Julesz (1993). It is not necessarily the case that the two claims are incompatible. It would appear likely that the local constraints suggested by Field *et al.* would take precedence over the global influences described by Kovacs and Julesz. If the claim that local integration is necessary to accomplish contour detection were accepted then it would seem unlikely that the contours utilised by Kovacs and Julesz would be detected by a completely different and independent mechanism. Consequently, their hypothetical ‘synergistic-process’ must occur within the same physiological substrate as the association field.

Pettet, Mckee and Grzywacz (1998) investigated the relative importance of the global cue offered by contour closure (Kovacs and Julesz, 1993) and the local constraints suggested by Field *et al.* (1993). They manipulated the presence or absence of ‘kinks’ within closed contours and demonstrated that near threshold closed contours became undetectable when sharp bends were introduced into the contour path. This result supports the claim that the local pair-wise constraints must be satisfied before global contour properties can have an influence. Kovacs and Julesz suggested that their *synergistic-process* occurs within a closed chain of co-facilitatory elements. Therefore, introducing a kink that breaks this chain will obviously suppress the closure-effect. Pettet *et al.* have merely confirmed what logic would suggest should be the case. That a collection of contour segments only form a contour when each member is co-aligned with each neighbour. Rotating a section of contour beyond a particular angle will break it into two independent contours. This finding suggests either that the closure-effect is caused by the association-field or that the mechanism responsible for the closure-effect occurs in areas later than those that form the substrate of the association field.

In their second experiment Kovacs and Julesz added elements to an incomplete contour and found that the closed contour became visible only with the addition of the last 1-2 elements. They suggested that this change in contour detectability was due to a synergistic process that is facilitated in closed contours. Braun (1999) argues that the second Kovacs

and Julesz experiment was flawed. He claims that their incomplete closed contours were too short to be detected at rates in excess of 75%. Thus, with the addition of 1-2 extra elements the contour attained a supra-threshold length. However, Braun's argument does not adequately explain why Kovacs and Julesz's open-contours fail to exhibit a similar psychometric function as additional elements are added

1.4.7 The association field is only a phenomena of the fovea

Hess and Dakin (1997,1999) tested the detection of contours in the periphery of human observers. They presented contours at eccentricities as large as 30°. Contour detection rates were measured for phase-aligned and phase-alternated Gabor patches (Figure 1-6). For contours presented to foveal areas, there was no significant difference in the detectability of either type of contour. Conversely, for contours presented at eccentricities beyond 10°, human observers were specifically unable to detect phase-alternated contours.

Hess and Dakin then tested the performance of the simple-filter model (described in section 1.4.5) with the same stimuli. They report that the simple filter model was only able to detect phase-aligned contours. Detection rates for phase-alternated contours barely exceeded chance levels. They conclude that the simple-filter model provides a sufficient account of human peripheral contour detection. As human observers can detect both phase-aligned and phase-alternated contours presented to the fovea, they argued that the simple filter model did not explain foveal contour detection. They conclude that the association field account of contour integration applies only to the fovea. For the periphery, contour detection was only achieved via a mechanism that is analogous to the simple filter model.

During their simple filter modelling, Hess and Dakin only reported filtering their stimulus images with filters of a particular scale ($4.5\sigma * 4.5\sigma$). They do not report why this particular filter was selected. This raises the question of whether or not the simple filter has been adequately assessed. It is certainly possible that other filters will exist which may be better suited to the detection of embedded contours. Furthermore, it possible that the most appropriate filter for detecting phase-aligned contours will not be suitable for the detection of phase-alternated contours.

1.5 Summary

The Gestalt school illustrated some of the conditions under which contour integration could occur. However, they failed to explain the processes underlying integration. Later studies (Uttal, 1975; Beck *et al.*,1989) developed methodologies that enabled a more

thorough investigation of the processes underlying contour detection. However, these studies often failed to successfully control for possible confounds in their stimuli, such as differences in the densities of background and contour elements. It is particularly important that earlier studies of contour integration often did not adequately control for the possibility that coarse-scale processes might account for detectability.

The CODE-2 model offered by Smit *et al.* (1985) successfully accounted for some contour integration phenomena. Unfortunately, the mechanisms employed within this model appear to have little biological plausibility. However, there are similarities between the CODE-2 model and the association field (Field, Hayes and Hess, 1993). Estimates of dot proximity might equate to the extent and strength of association field connections. Whilst the direction of elongation of the Gaussian distribution function might correspond to the major axis of these horizontal connections. A major difference remains though, the CODE-2 model is concerned only with dots which, individually, have no orientation component. Whereas the association field integrates the outputs of oriented filters. This reflects a difference in the stimuli employed in each study. Density and co-linearity alone define contours in Smit *et al.*'s experiments whilst an additional cue of line segment orientation included in the path-paradigm stimuli.

With the development of the path-paradigm, Field *et al.* argued that contour detection experiments could be conducted which precluded both coarse-scale (simple filter model) and hard-wired path-detector solutions. By carefully matching the spacing of contour and background elements, they claim that they controlled for any density cue present within their stimuli. Their results led them to propose a dynamic association-field, which integrates outputs from neighbouring oriented filters. Evidence from neurophysiological studies (for example, Ts'o *et al.*, 1986) revealed a possible substrate for the association field.

Studies have recorded the firing rates of simple cells in area V1, during the presentation of stimuli sharing some properties of those employed in path-paradigm experiments (Kapadia *et al.*, 1995). These have shown that the firing of cells in area V1 can be modulated by stimuli present in areas beyond the classical receptive field. Furthermore, parallel studies revealed that the firing of single-cells correlates with the responses of human observers. This study provides further support for the association-field hypothesis and suggests that long-range horizontal connections within area V1 may form the substrate of this process.

Converging evidence suggests that association fields may be acquired through visual experience.

Kovacs and Julesz (1993) have demonstrated that the overall structures of contours can influence their detectability – even when local properties are held equal. Furthermore, Pettet, Mckee and Grzywacz (1998) claim that local relationships, between individual contour elements, have priority over the overall structure of the contour.

Hess and Dakin (1997, 1999) have demonstrated that the simple filter model does not provide an adequate account of foveal contour integration. They report that the simple-filter model performs poorly when embedded contours are composed of phase-alternated Gabor elements, whilst human observers are equally capable of detecting both phase-aligned and phase-alternated contours, when the contours are presented to the fovea. However, they suggest that a simple-filter account does provide a sufficient explanation of the detection of contours presented to the periphery. Hess and Dakin concluded that the association field is only present within foveal areas.

1.6 Directions for the current thesis

1.6.1 Testing the simple-filter model of contour integration

Field, Hayes and Hess (1999) argued that contour detection occurs because of the operation of an association-field mechanism. An alternative explanation, involving a simple filter model, was rejected. Hess and Dakin (1997,1999) offered modelling evidence to support the rejection of the simple-filter account of human contour detection. The arguments against these alternative mechanisms are based upon the assumption that success in the path-paradigm task requires the detection of the whole contour. In the case of the simple-filter model, it is clear that the integration of whole contours is unlikely. This is largely due to the fact that distant contour elements are unlikely to share the same orientation. Therefore, it is unlikely that a single filter will match the orientation of distant elements belonging to the same contour. It follows that it is unlikely that a ZBR will correspond to the whole of an embedded contour.

The path-paradigm task itself simply requires the observer to indicate which of a pair of images contains a contour. If it were not necessary to detect the whole contour, then Field *et al's* arguments against the simple filter model are weakened. The degree of correspondence between the ZBR and the embedded contour does not need to be exact. It would be sufficient if there were a longer ZBR in the target image than there was in the

noise image. The overall length of the ZBR would not have to correspond to the length of the whole of the embedded contour.

Chapter 3 will investigate the possibility that the simple filter model can indeed account for human performance in path-paradigm experiments. Whereas Hess and Dakin selected a single filter scale for the convolution of their images (difference of Gaussians filter, symmetrical, 4.5σ pixels), the current modelling experiments will allow filter scale and elongation to vary independently in order to increase the power of the experiment.

1.6.2 Do the global characteristics of contours influence detectability?

The results of Kovacs and Julesz (1993) suggested that the overall structure of a contour can influence its salience. However, Braun (1999) suggested that their experimental methodology may have exaggerated the level of this influence. Neither Braun nor Kovacs and Julesz considered the influence that the frequency of inflection points within contours might have upon detectability. When local contour properties of spacing and path-angle are held constant, there will inevitably be a difference in the number of times that the path-angle between consecutive contour elements will change sign. Hence, for an open contour there must be more changes in the sign of the path-angle change than there would be for a closed contour. For terminological convenience, this contour property will be referred to as ‘smoothness’. Following this definition, the smoothest contour will feature no inflections, i.e. the sign of path-angle changes do not change throughout the contour. Ultimately the smoothest contour will be a circle (or a spiral) or perhaps a straight line, depending upon the values of path-angle and path-angle jitter. This influence of smoothness upon contour detection rates is explored in chapter 4.

1.6.3 Does the presence of spurious contours in the background influence detection?

The experiments of Field *et al.*, Kovacs and Julesz, Pettet *et al.* and Braun have all attempted to control for potential confounding cues that might indicate the presence of an embedded contour. An important example is the potential cue offered by the relative spacing of elements forming the contour and the background. Ideally, there should be no difference between the spacing of background elements and contour elements. However, Kovacs, Polat and Norcia (1997) have shown that subjects do not benefit from a density cue when the background elements are more closely packed than the contour elements.

Element spacing densities are controlled through various means. For example, Field *et al.* initially located elements upon the vertices of a square grid. The spacing of the rows and columns of the grid matched the spacing of neighbouring elements along the target

contour. Background element positions were then *jittered*. That is, their location was varied by applying a small degree of randomisation to each elements X and Y location. Finally, those background elements that overlapped with the target contour elements were removed. In each of the studies reported within section 1.4, the orientation of individual background elements was randomised.

As the locations and orientations of background elements are randomised, the stimulus backgrounds will inevitably feature contours that occur spontaneously because of randomly occurring co-alignments. These spurious contours will henceforth be referred to as *noise-contours*. Noise-contours are unlikely to be as long or as smooth as the target contours. It follows that there is a very low probability that noise-contours will be as salient as the deliberately embedded contours (*target contours*). Nevertheless, noise-contours may sometimes mislead subjects in 2AFC experiments. Therefore, the presence of noise-contours may have a negative effect upon subject performance. This is particularly likely when manipulations of target contours make them difficult to detect.

Chapter 6 describes a method of manipulating the occurrence of noise-contours. Experiments will illustrate how subject performance is influenced by the prevalence of these contours. It follows then that by eliminating noise-contours one is able to offer estimates of subject performance rates that are unaffected by distractions present in the background.

1.6.4 Contour integration in the periphery

Hess and Dakin (1997, 1999) have argued that contour detection in the periphery can be explained by the operation of a simple-filter model. Chapter 5 examines how the global structure of contours can influence contour detection in the simple-filter model. If the Hess and Dakin hypothesis is correct then these influences should have an equivalent influence upon human performance in the periphery. Chapter 7 examines how global contour properties can influence contour detection within the periphery.

2 General Methods

2.1 Overview

This section will briefly describe the details of experimental procedure that are common to all of the experiments described within this thesis. All aspects of experimental administration were conducted within the Matlab® software environment, i.e. stimuli generation; the ordering of image presentation sequences; the display of stimulus images and finally subject response storage and analysis. Each of these aspects will be described in the sub-sections that follow. Two experimental replications were conducted in order to validate the methodology and software described. These are described in sections 2.2 and 2.3.

2.1.1 Stimulus Generation

Prior to the administration of an experiment, stimulus images were generated offline and stored on disc. The appendices list some of the programs utilised during the administration of experiments. For example, appendix F shows the stimulus generation program used to generate the stimuli for the majority of experiments described in this thesis¹. There were three stages in the generation of a stimulus image: -

- 1) The generation of the contour path.
- 2) The generation of a random background and the process of embedding the contour path into the background.
- 3) The generation of an image with Gabor patches located at the x-y position of each background and contour element.

For clarity, each of these processes will be discussed in isolation within the following sub-sections.

2.1.1.1 Contour path

Target images featured embedded contours that were formed from a number of Gabor elements. The procedure for generating the positions of the contour segments was functionally equivalent to that described by Field et al. (1993) (Appendix D). Briefly, an invisible spine was generated, which consisted of a series of interconnected line segments. The angle and distance between successive segments was fixed. These variables are referred to as the *path-angle* and the *element separation* respectively. The contour

¹ Those excluded are the experiments described in chapters 6 and 7; these required additional specially tailored routines that are described within the relevant chapters.

elements were then positioned at the mid-point of each of these invisible lines. The difference between the orientation of each contour element and the section of spine upon which it was located is termed the *path-relative element orientation*; this parameter was 0° in all experiments.

In order to reduce the amount of regularity in the resultant contour, random *jitter* was applied to each of the three parameters defined in the previous paragraph. The length of each section of the path was modified by *element spacing jitter*; this value was randomly selected from within the range of $\pm 25\%$ of the element-spacing parameter. The orientation of each contour element was randomised by *element orientation jitter*, this was $\pm 5^\circ$. The path-angle between successive line segments was modified by *path-angle jitter*, this was also $\pm 5^\circ$.




Figure 2-1, schematic diagram illustrating the parameters which determine the overall structure of an embedded contour. Each of the three main parameters were modified by a small degree of random *jitter*, labels for the three types of jitter are provided within brackets.

2.1.1.2 *Stimulus background*

The underlying premise of the path-paradigm experiment is that all possible methods of detecting the embedded contour are prevented, except those involving the integration of the outputs of receptive fields that independently detect the presence of each line segment. Thus, the background performs an important role in relation to this premise. The density of elements belonging to the background and the contour must be matched in order that density does not cue the existence of the contour. Experiments have revealed that the presence of the contour is only cued when the density of the background is lower than the contour density (Kovacs, Polat and Norcia, 1997). This property is sometimes utilised in experiments described within this thesis. Contour element spacing may be increased whilst

background density is held constant. This enables the reduction of detection rates in order to avoid ceiling-effects.

The method developed by Field, Hayes and Hess (1993) for the creation of a stimulus background will be termed the *jittered-grid*. This background is achieved by placing elements within a lattice of rows and columns and then randomising the positions of these elements within the bounds of each row and column (Figure 2-2, dashed lines). The density of elements is determined by the number of rows and columns. When a contour element is placed into the grid, background elements are deleted where there is an overlap between contour and background elements. An overlap has occurred when two elements are placed into a single matrix square. Full details of the creation of a *jittered-grid* background are provided within appendix A.1. The orientation of each background element was randomly assigned a value between 0 and 180°.




Figure 2-2. Illustration of the *jittered-grid* background generation technique.

Experiments described in chapters 6 and 7 and section 4.5 utilise an alternative method of generating the stimulus background, this is termed the *spaced-fill* background. Brief details of this method are provided in sub-section 6.2.4 and the full software program is provided in appendix A.2.

2.1.1.3 *Gabor patch*

A major difference between the stimuli utilised by Uttal (1975) and those employed within the current thesis is the use of oriented line segments. Each dot element in Uttal's stimuli

did not have an orientation component. The orientation of Uttal's stimuli was defined by the overall arrangement of the dots. In contrast, line segments employed in the current thesis were oriented Gabor elements. Furthermore, arrays of dots can have a wide spatial frequency spectra, thus it is impossible to be certain that such stimuli are not being detected by a single, spatially extended, receptive field.

Gabor elements were used as they provided an oriented line segment that was bandpassed for both spatial scale and orientation; i.e. the position and orientation information is only available over a restricted range of spatial frequencies. Equation 2-1 defines a Gabor patch, which is the product of a circular Gaussian and an oriented sine wave. The location of the Gabor patch is defined by the Cartesian co-ordinates x and y . The slope of the circular Gaussian is defined by σ ($\sigma =$ two pixels). The orientation of the sine wave and thus the Gabor patch is determined by θ . The wavelength of the sine wave is defined by p ($p = 4$ pixels). The phase of the Gabor patch is defined by ϕ . Appendix E lists the computer program used to generate Gabor patches.

$$G(x, y) = e^{-\left(\frac{x^2 + y^2}{2\sigma^2}\right)} \cdot \cos\left(\frac{2\pi * (\cos\theta * x + \sin\theta * y)}{p + \phi}\right)$$

Equation 2-1

When the Gabor elements are described as *phase-aligned* the value of ϕ was 90° for all patches (i.e. cosinusoidal). This gives a patch with a positive (light) centre and two negative side-bars (dark). When the phase of Gabor patches is described as *phase-alternated*, then the phase of successive elements in a contour will alternate between 0° and 180° and background elements will be randomly assigned a phase of either 0° or 180° . Examples of phase manipulations are provided below in Figure 2-3.




Figure 2-3. Examples of contours which are formed from Gabor elements which are either phase-aligned or phase-alternated.

Where a stimuli featured phase-aligned contours, the phase (ϕ) of Gabor patches that formed the stimulus background was assigned a value of 90° . Where stimuli featured

phase-alternated contours the phase of Gabor patches was randomly assigned a value of either 0 or 180°.

2.1.2 Presentation sequencing

All experiments described within this thesis utilised a temporal two alternative forced choice paradigm. The order of presentation for images was randomised without replacement. For example, if five stimulus levels are to be presented during an experiment [1 2 3 4 5] then a randomised sequence might be : - [3 2 1 4 5] [4 3 1 5 2] ... [5 3 4 1 2]. During a typical experiment two hundred target images would be presented, each target image was paired with a noise image. The presentation order of each pair of target and noise images was randomised, with the constraint that a target would be shown first on 50% of presentations. As individual images were usually generated offline, prior to the beginning of an experiment, the selection of individual images was also randomised, again without replacement. Thus, in an experiment featuring ten discrete levels of stimulus strength and with twenty discrete images per level then the same image would not be presented more than once during a particular experimental session.

2.1.3 Displaying stimuli

The presentation of each image was cued by a fixation cross. The location of the fixation cross corresponds to the centre of the image. The details of presentation timing are provided in Figure 2-4, below. Each image was displayed for 250ms.

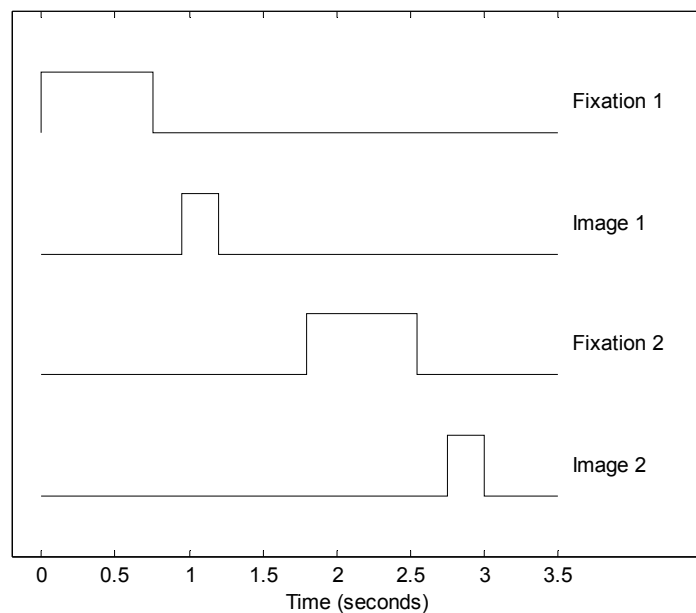


Figure 2-4. Stimulus timing details. Elevated sections of each line represent periods when the indicated items are visible to the observer.

The experiments described in the current chapter through to Section 4.4 were displayed upon a Microscan 4V ADI, 14" Colour Monitor with resolution set to 800x600 pixels. Images were viewed from a distance of 750mm.

Experiments described in sections 4.5 onwards were displayed with a Cambridge Research VSG system (VSG2/4-4MB) and displayed on a Panasonic S110 20" monitor. Screen resolution was set to 1024 x 768 pixels and the monitor frame rate was 100_{hz}. Gamma correction was achieved using the Cambridge research *OptiCal* hardware (OP200-E).

2.1.4 Response storage and analysis

The presentation sequence, individual image reference numbers, observer decisions and reaction times were retained following the completion of each experimental session. Sections 2.1.4.1 and 2.1.4.2 describe how observer responses were summarised and analysed.

2.1.4.1 Psychometric functions

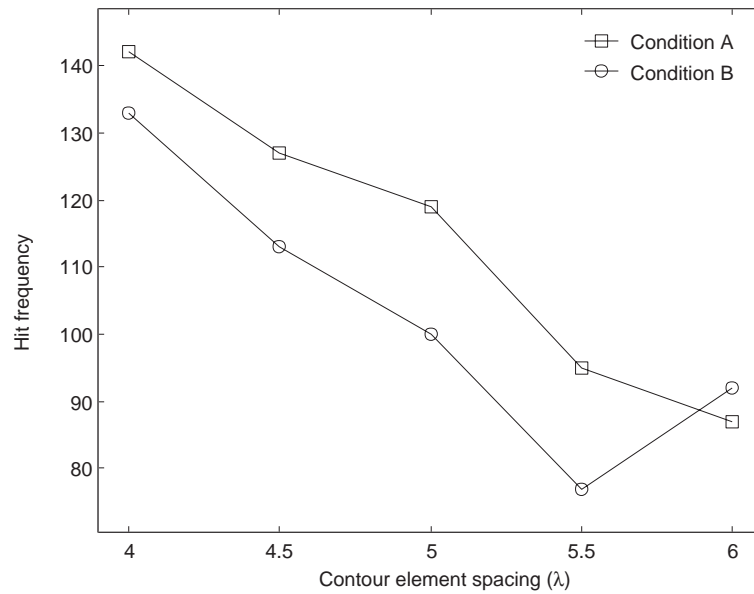
A psychometric function (Function 2-2) is used in order to summarise the performance of observers. This function provides an estimate of the probability of the detection of a target image as a function of the stimulus level (s). We start with a normal distribution $\exp(-x^2)$. The dimension of x is transformed by dividing by σ , it is scaled by raising to the power of pw and shifted by adding 50. The function is then integrated along the dummy variable x from $-\infty$ to the stimulus level. The parameters are σ and pw , the optimal values for these variables are found by the Matlab® *fmins* function. This function uses the Nelder-Mead simplex method.

$$P(s) = \int_{-\infty}^s \exp\left(-\left(\frac{x}{\sigma}\right)^{pw} + 50\right)^2$$

Function 2-2

2.1.4.2 χ^2 analysis of differences in the detectability of stimuli

Wherever necessary the difference in detectability between any two stimuli (A and B), was calculated using the χ^2 test, see Formula 2-3. An example analysis is shown in Figure 2-5. A χ^2 analysis was selected as it was suitable for non-parametric data and demands few unsupported assumptions about the nature of the data analysed.



λ	Ahits	An	Bhits	Bn
4.0	142	152	133	152
4.5	127	152	113	152
5.0	119	152	100	152
5.5	95	152	77	152
6.0	87	152	92	152
=====				
Total	570	760	515	760

	Correct (hits)	Wrong (miss)	Total
Stimulus A	570 <i>542.50</i>	190 <i>217.50</i>	760
Stimulus B	515 <i>542.50</i>	245 <i>217.50</i>	760
Total	1085	435	1520

(Expected values in *italics*)

Therefore the null-hypothesis can be rejected ($\chi^2(1) = 9.742, p = 0.0018$). The detection rates for stimuli A and B are significantly different.

Figure 2-5. An example of a χ^2 calculation. The plot represents the number of correct selections of the target stimuli in a 2AFC experiment as a function of contour element spacing and stimulus condition (A or B). (Bottom left) The total number of hits are calculated for each condition. (Bottom right) the χ^2 contingency table.

The application of the formulae defined in Formula 2-3 - Formula 2-6, enable the calculation of the probability that detection rates for the two stimuli are different.

$$\chi^2 = \sum \frac{(O - E_{hit})^2}{E_{miss}}$$

Formula 2-3

Where E represents the expected frequency and O the observed frequency for correct and incorrect A and B decisions. Estimated frequencies for the correct and incorrect selections of target stimuli (termed *hits* and *misses*) are calculated for A and B using the Formula 2-4 for *hits* and Formula 2-5 for *misses*.

$$E_{hit} = \frac{N_{hits} \times N_{Trials}}{N_{total}}$$

Formula 2-4

$$E_{miss} = \frac{N_{misses} \times N_{Trials}}{N_{total}}$$

Formula 2-5

Where N_{hits} is the observed total number of hits, N_{misses} is the total number of misses, and N_{trials} is the number of 2AFC decisions made for A or B stimuli type. N_{total} is the total number of target and noise image pairs observed for both A and B stimuli.

$$p = 1 - cdf(\chi^2 | df)$$

Formula 2-6

The probability p is calculated using Formula 2-6, where cdf is the χ^2 cumulative distribution function and df is the degrees of freedom, as comparisons are made between pairs of stimuli df is always one.

2.1.4.2.1 Comparisons of more than two psychometric functions

The example that is provided in Figure 2-5 shows an analysis that examines detection rate differences across all levels of element spacing. Differences in the detectability of stimuli may only be significant over a specific range of stimulus values. Ceiling and/or floor effects may cause detection rates to converge beyond this range. As many analyses involved comparisons of more than two psychometric functions, the analysis procedure attempted to find an optimal range, i.e. one that led to the largest number of significant differences between all comparisons.

In order to identify an optimal range two additional criteria were employed during statistical analysis. Firstly, the χ^2 statistic was calculated for all possible ranges of the stimulus level. The stimulus range that leads to the greatest number of significant p -values was then selected. If more than one range led to the same number of significant p -values then the larger range was selected. Figure 2-6 presents an example of a χ^2 analysis that has followed this process. Visual inspection of the plot indicates that observer responses approached chance levels when elements spaced further apart than 6λ . Conversely, when elements were separated by distances of between $4-6\lambda$, it appears that there are consistent differences in the detection rates for stimuli A, B and C. The χ^2 analysis confirms this, there are significant differences between all pairs of stimuli except for that between A and D.

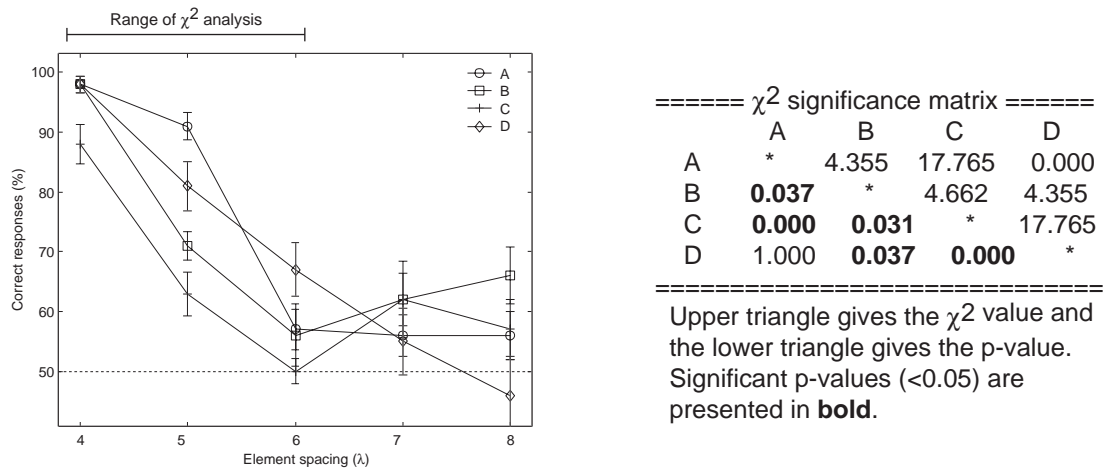


Figure 2-6. Examining the statistical significance of differences in contour detection rates over a range of contour spacing values. (Left) plots illustrating subject performance with stimuli (A-D) when presented with a range of contour element spacings (λ). (Right) The χ^2 significance matrix indicates that all differences are significant for spacings of 4-6 λ , except for those between stimuli A and D.

Experimental stimuli used in the current thesis were generally created offline. Therefore, it was necessary to select a range of stimulus levels that would be tested during a particular experiment. Consequently, this range did not always span all observer performance levels from detection rates between 100% correct and chance levels. Therefore, one should only consider the χ^2 analysis as an heuristic measure of the differences in detectability of any pair of contours.

2.2 Experiment 1: Replication of Field, Hayes and Hess path-angle experiment.

2.2.1 Introduction

As much of the software was created by the author, it was deemed necessary to replicate existing experiments in order to verify that each aspect of the experimental software was operating correctly. Field *et al.* (1993) report that contour detection rates fall as a function of path-angle, the difference in the orientation of successive contour elements. They suggest that this result reflects the pattern of facilitatory connectivity between receptive fields, where strong connections exist between co-linear receptive fields with similar orientation preferences. The following experiment aimed to replicate this experimental result.

2.2.2 Methodology

2.2.2.1 Stimulus parameters

The path-angle of embedded contours was varied between 0-80° in steps of 10°. Each element was oriented such that it was aligned with the orientation of the contour path. The spacing of contour elements was fixed at 16 pixels. This level of element spacing corresponded to a separation of 4λ when expressed in terms of the Gabor element wavelength. Further details of image scaling are provided below in Table 2-1. In order to reduce the regularity of contour structures, a degree of random ‘jitter’ was added to each of the parameters described above. Path-angle values were modified by $\pm 5^\circ$ jitter, path-relative element orientation was modified by $\pm 5^\circ$ jitter, the spacing of elements was varied by adding a random jitter of up to 4 pixels to the gap between successive elements.

Contours were randomly embedded within the stimulus background. The background was generated using the jittered-grid technique and the mean spacing of background elements was 16 pixels. Noise images featured a stimulus background generated with the same parameters as used for the generation of the corresponding target images. Noise images did not feature an embedded contour. All Gabor elements in the stimuli were phase-aligned.

2.2.2.2 Observers

Six observers participated in this experiment. Each participated for a single experimental session during which they observed two hundred target/noise image pairs. The observers were post-graduate students or academic staff, all had normal or corrected to normal vision.

2.2.2.3 Procedure

Stimuli were displayed upon a Microscan 4V ADI, 14” Colour Monitor with resolution set to 800x600 pixels. Images were viewed from a distance of 750mm. Stimulus scale properties corresponding to this distance are illustrated in Table 2-1. Observers were informed that they were to view pairs of images and that they should indicate whether the first or second image of each pair contained an organised area that drew their attention. A practice trial was administered which consisted of 10-15 target and noise image pairs. Further practice was available if it became apparent that subjects had misunderstood their instructions.

Table 2-1. Stimulus scale properties for the current experiment.

	Pixels	Lambda	Visual-angle
Gabor patch full-wave frequency	4	1λ	0.115°
Background and contour minimal nearest-neighbour spacing	16	4λ	0.458°
Contour element spacing	16	4λ	0.458°
Overall image size	256×256	$64 \times 64\lambda$	$7.37 \times 7.37^\circ$

2.2.3 Results

The overall pattern of each subjects responses was similar. They all showed lower detection rates as path-angle was increased. Therefore, only the mean performance of all subjects is shown in the figures below. Figure 2-7 shows the mean performance of all subjects, clearly detection rates fall as a function of contour path-angle.

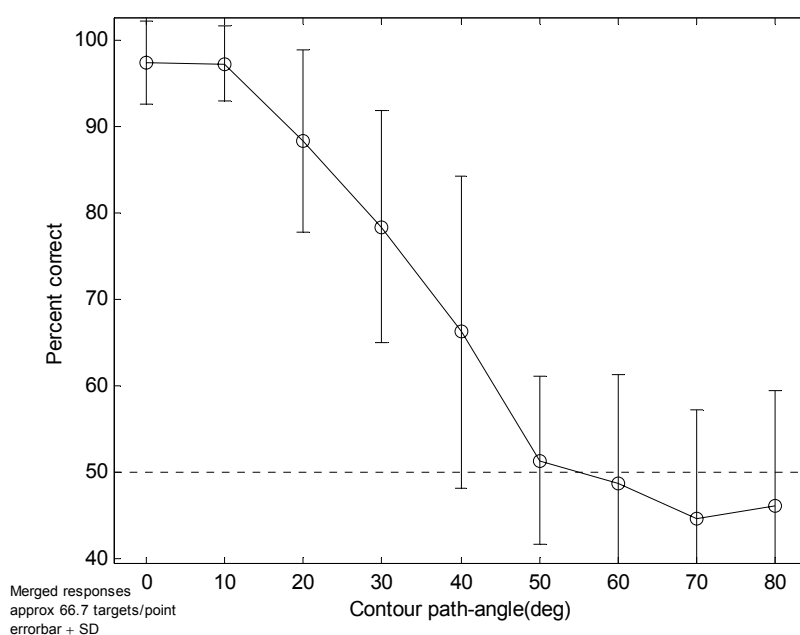


Figure 2-7. Contour detection as a function of contour path-angle.

Figure 2-8 shows subject performance fitted by a psychometric function of contour path-angle. The fitted function predicts that 75% correct performance would be achieved when contour path-angle was approximately 33° .




Figure 2-8. Contour detection rates as a psychometric function of contour path-angle. The error bars represent the standard deviation of observer performance.

2.2.4 Discussion

The overall pattern of observers' responses was similar to that reported by Field *et al.* (1993). However, for the current experiment the fitted psychometric function predicts performance levels of 60% correct when path-angle is 45°, whilst for Field *et al.* the rate was closer to 75% with equivalent stimuli. This slight discrepancy may have occurred because subjects in the current experiment were naïve to the path-paradigm and each only received minimal practice. The higher detection rates reported in the Field *et al.* study may be due to increased practice as two of the authors also acted as observers.

2.3 **Experiment 2: Manipulating contour element spacing.**

2.3.1 Introduction

Field *et al.* reported that contour detection rates fall as a function of contour element spacing. They interpreted this finding as an illustration of the spatial extent of the association field. Experiment 2.3 aims to partially replicate their experiment. There was an important difference between the stimuli utilised in this experiment and those employed by Field *et al.* (experiment IV). In their experiment, the spacing of contour elements was matched to the spacing of background elements, therefore as contour element spacing was increased, the density of background elements varied too. In the current experiment, the density of background elements remains constant. Kovacs, Polat and Norcia, (1997) have

shown that the relative density of contour and background elements only cues the presence of the contour when the density of the contour elements was higher. When the density of the contour elements is lowered relative to the background, there should be a reduction in the detectability of the contour. Any decrease in contour detectability with increases in element spacing may be due to a combination of background element crowding and increasing contour element spacing.

2.3.2 Methodology

2.3.2.1 Stimulus parameters

Embedded contours consisted of 15 Gabor patches. The path-angle was $\pm 20^\circ$ for all contours. The spacing between contour elements was manipulated between 16-25.6 pixels² in steps of 2.4 pixels. Path-angle jitter was $\pm 5^\circ$, contour element spacing jitter was 25% of the contour element spacing and element orientation jitter was $\pm 5^\circ$. The mean spacing for elements within the stimulus background was 16 pixels and the orientation of these elements was randomised. All Gabor elements were phase-aligned.

2.3.2.2 Observers

Thirty-six observers participated in this experiment³. The observers were naïve to the purpose of the experiment and all had normal or corrected to normal vision. Observers were undergraduate volunteers recruited from the University of Stirling subject pool. They received course-credit for participation. Each observer participated for a single experimental session, lasting approximately thirty minutes.

2.3.2.3 Procedure

Stimuli were displayed upon a Microscan 4V ADI, 14" Colour Monitor with resolution set to 800x600 pixels. Images were viewed from a distance of 750mm. Stimulus scale properties that correspond to this distance are illustrated below in Table 2-2. Observers were informed that they were to view pairs of images and that they should indicate whether the first or second image that they saw contained a single, continuous, contour. They were shown an example of a random (non-closed) embedded contour at this point. A practice trial was administered which consisted of 10-15 target/noise image pairs. Further practice was available for those that had clearly misunderstood the instructions.

² Clearly, a pixel value is necessarily an integer, the decimal values assigned to pixel spacing occur because element spacing is initially calculated in terms of λ . Mean distances between neighbouring contour elements will correspond to the value reported. As successive elements are unlikely to be either exactly vertical or exactly horizontal the distance between neighbouring elements is rarely an integer.

³ The number of observers was high as they composed the control group in an experiment investigating perceptual learning with path-paradigm stimuli.

Table 2-2. Stimulus scale properties for the current experiment.

	Pixels	Lambda	Visual-angle
Gabor patch full-wave frequency	4	1λ	0.265°
Background and contour minimal nearest-neighbour spacing	16	4λ	1.06°
Contour element spacing	16.0	4.00λ	0.458°
	18.4	4.60λ	0.527°
	20.8	5.20λ	0.595°
	23.2	5.80λ	0.665°
	25.6	6.40λ	0.733°
Overall image size	256 x 256	64 x 64 λ	7.37x7.37 $^\circ$

2.3.3 Results.

The patterns of overall responses for individual subjects were largely similar, therefore only the mean performance is presented. Figure 2-9 shows the mean detection rates for all subjects as a function of contour element spacing. The contour detection rates clearly fell as a function of contour element spacing. Performance drops to chance levels when element spacing approaches 6.4λ . Figure 2-10 shows the fitted psychometric function for element spacing. A visual inspection of this figure confirms that there was a good degree of fit between the raw observer performance and the fitted psychometric function.




Figure 2-9. Contour detection rates as a function of contour element spacing.




Figure 2-10. Contour detection rates as a psychometric function of contour element spacing.

2.3.4 Discussion

The contour detection rates of observers were much lower in this experiment than they were for Field *et al.* (expt. IV). For contours with a path-angle of 20° they report detection rates of near 100% over a range of element spacing values (0.25° - 0.9° visual angle). In the current experiment the spacing values of $4 - 6.4\lambda$ correspond to visual angles of $0.46 - 0.73^\circ$. However, it is important to note that the density of background elements was matched to that of contours in their experiment, whilst in the current experiment it remains constant. The current results may simply reflect the additional influence of background element crowding, i.e. the increased density of background elements relative to the density of the contour elements reduces contour detectability. Crowding effects have been previously utilised by Kovacs, Polat and Norcia (1997). The effect of background crowding was utilised in the experiments described in chapters 4 and 5 in order to examine the relative detectability of different stimuli across a range of performance levels.

2.4 Summary

Both of the replications were clearly successful, although the effect of contour element spacing may have been confounded by background element crowding. The results confirm that the methodology employed and the software created to implement the methodology were acceptable and that they provide useful measures of observer performance levels. The

stored observer responses from each of the experiments described within this chapter were retained and used in the comparisons of human observer and simple-filter model responses described in the following chapter.

3 Modelling Contour Integration with the Simple-Filter Model.

3.1 Overview

Section 1.4.5 summarised the stance taken by Field, Hayes and Hess (1993) and Hess and Dakin (1999) with regard to alternative accounts of human performance in the path-paradigm task. Two reasons were offered for the rejection of the simple-filter model. Firstly, Field *et al.* argued that the properties of the stimuli used precluded any contribution this mechanism might have made towards contour detection. They did not test this assertion however. Secondly, Hess and Dakin (1999) demonstrated that the simple filter model was specifically impaired when contours were composed of phase-alternated Gabor elements. In contrast to the assertions of Field *et al.*, the simple filter model was able to detect stimuli featuring phase-aligned contours; though detection rates remained approximately 20% lower than those reported for human observers. Hess and Dakin found that the foveal performance of human observers was unaffected by the phase of contour Gabor elements.

The majority of experiments that were reported by Field *et al.* (1993) utilised phase-aligned stimuli. Therefore, their experiments did not adequately control for contributions of coarse-scale mechanisms. Their final experiment featured Gabor elements with randomised-phase. With such stimuli, the occurrence of phase co-alignments between neighbouring elements will be reduced. Therefore, such stimuli should make the path-paradigm task more difficult for a simple-filter mechanism. However, phase-alignments between successive contour elements can still occur. So such stimuli would not entirely control for the contributions of coarse-scale mechanisms.

The strongest evidence against the simple-filter model of human contour integration is the fact that this model was specifically impaired when contours were formed from phase-alternated Gabor patches. Hess and Dakin only report the results of testing their model with a difference of Gaussians (DoG) filter of dimensions 4.5σ (pixels). They state that varying filter size and the threshold level did not improve the performance of the model. It is not clear if the reported scale of 4.5σ (pixels) for the DoG filter reflects the optimum scale for the detection of both phase-aligned and phase-alternated stimuli. Alternatively, this scale may actually be a compromise between the optimum filter scale for each type of stimuli. The experiments reported in this chapter were designed to test whether the

performance of the simple-filter model will improve when appropriately scaled filters are used.

The modelling experiments described within this chapter utilised the stimuli featured in the experiments already described in section 2.2 (manipulating path-angle) and section 2.3 (manipulation of contour element spacing). This enables a comparison with human-observer performance with the same stimuli.

3.1.1 Selection of filter scale in the simple-filter model

In contrast to Hess and Dakin (1999), it is not presumed that a filter with a particular bandwidth and degree of elongation will be best suited to the detection of all path-paradigm stimuli. Currently filters are selected in a *post-hoc* manner by the experimenter, generally with the aim of maximising the rate of contour detection. Early vision features filters that are tuned to many different spatial frequencies. Therefore, a complete implementation of a simple-filter model would require a mechanism that could select the filter output most appropriate to the particular path-paradigm task. It follows then that the simple-filter model only remains a parsimonious account of contour detection, while the mechanism responsible for selecting an appropriate filter remains relatively simple. If the relationship between the stimulus and the corresponding best filter⁴ is consistent and uncomplicated then the simple-filter model is parsimonious. If the process of selecting an appropriate filter is complex, then the simple-filter model may be no more parsimonious than alternative accounts of contour detection.

The relative plausibility of different implementations of the simple-filter model can be compared using the criterion defined below. Three broad categories of simple-filter model are described, each with an increasingly complex relationship between filter scale and stimuli. The plausibility of the model is inversely correlated with the number of differently scaled filters required to solve the path-paradigm task.

- i) A model that operates with a filter of only one scale. This presumes that a single optimal filter is appropriate for all experimental stimuli.
- ii) A model that deploys different filters for path-paradigm stimuli composed of Gabor patches that are phase-aligned than it does for stimuli that are phase-alternated.

⁴ The term ‘best filter’ refers to the filter that is best suited to the detection of contours embedded within particular stimuli; no other filter of the same type but with a different scale will detect contours as well.

- iii) Filters of different scales are deployed with every manipulation of a stimulus property. This implementation of the model would require a mechanism that was able to select the correct filter on an ad-hoc basis.

Clearly, the criteria described are an oversimplification. Some manipulations of stimuli will necessarily favour different filters. However, the relationship between the manipulation and the filter may be obvious. For example, if the viewing distance were increased, then the best filter would simply be a finer-scale version of the previous best filter. It follows that same heuristic should be able to select an appropriate filter-scale when viewing distance varies.

Hess and Dakin (1999) apply the most stringent criterion (i, above); they allow only a single filter scale for all their stimuli manipulations. The experiments in the current chapter apply the second, less stringent, criterion. It is hypothesised that different filters may be suited to the detection of embedded contours formed from phase-aligned and phase-alternated stimuli. Whilst the third criterion is the most lax, this model is also most likely to be the one which is the most realistic summary of human performance. It is likely that the selection of filter-scales is based entirely upon the contents of the visual scene at any instance.

The simulations reported in the current chapter employ a 2nd derivative of a Gaussian filter. The elongation and overall scale of this filter will be varied systematically in order to identify the optimal scale for each category of experimental stimuli. It is not presupposed that the optimal scale will be equivalent for both phase-aligned and phase-alternated stimuli. However, it is expected that the relationship between filter scale and stimuli will be broadly consistent. If this criterion is satisfied and the simple-filter model is able to identify stimuli containing embedded phase-alternated contours then it remains tenable as an explanation of human performance in path-paradigm experiments.

This thesis does not offer an account of how filter-scales appropriate to particular stimuli are selected. It is merely presumed that such mechanisms must exist. Clearly, a simple-filter account of performance in path-paradigm experiments remains incomplete until the details of such mechanisms are outlined. The aim of this section of the thesis was not to provide an alternative to the association-field account; the intention was to demonstrate that manipulations of Gabor patch phase do not control-for potential contributions of parallel-filter based mechanisms.

3.1.2 Differences from other models

The simple-filter model simulations reported in this chapter do not avoid the use of filters that have a coarser scale than the Gabor elements forming the stimuli. This is an important distinction from those models reviewed in section 1.4.3. In these models, processing units in the primary stage only detected individual Gabor elements. Therefore, in order to detect an embedded contour, these models integrate those filtered outputs that correspond to each contour Gabor element. The likelihood of integration was determined by the relative position and orientation of neighbouring filters. These models conform to the association field theory of Field, Hayes and Hess (1993), i.e. filtered outputs are combined across a range of different filter orientations.

For the simple filter model, the values within each zero-bounded region (ZBR) may correspond to the positions of multiple Gabor elements. However, all the filtered outputs that eventually form a ZBR come from the same filter shifted in space. Values within a particular ZBR are likely to reflect contiguous areas of the visual scene that have roughly comparable spatial-scale and orientation characteristics. As only neighbouring elements along a contour are likely to share a similar orientation, then it is likely that only sub-sections of an embedded contour will be identified by any single ZBR.

3.2 Methodology

3.2.1 The simple-filter model

For descriptive convenience, the operation of the simple filter model will be broken down and described in sub-sections. Section 3.2.1.1 will describe the processing of a single image; section 3.2.1.2 will explain how all images within a particular experiment are processed. Section 3.2.1.3 explains the administration of a simulated 2AFC experimental paradigm. The strategy for the presentation of simulation results is explained in section 3.2.2.

3.2.1.1 Processing individual images

Stimulus images are initially convolved with an oriented filter. The current implementation of the simple-filter model utilises a filter formed from a 2nd derivative of a Gaussian multiplied by a Gaussian.

$$\int (x, y) \mapsto -\left(1 - \frac{x^2}{\sigma_x^2}\right) \exp\left(-\frac{1}{2}\left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2}\right)\right) \quad \text{Function 3-1}$$

Following convolution, the values within the filtered-image that lie within one standard deviation of the mean are set to zero. This leaves an image containing discrete zero-bounded regions (ZBR). All values within a ZBR will be either positive or negative. The location and spatial structure of each ZBR is parameterised according to Watt's image description scheme (Watt, 1991). The process of convolution, thresholding and ZBR parameterisation is repeated whilst filters are rotated through ten discrete orientations⁵ spanning 0-180° in steps of 18°. The greatest ZBR length is retained. This value informs the model's selection of the *target* image in subsequent 2AFC decisions. The key stages of convolution and segmentation are illustrated in Figure 3-1, below.

⁵ 0, 18, 36, 54, 72, 90, 108, 126, 144 & 162°.

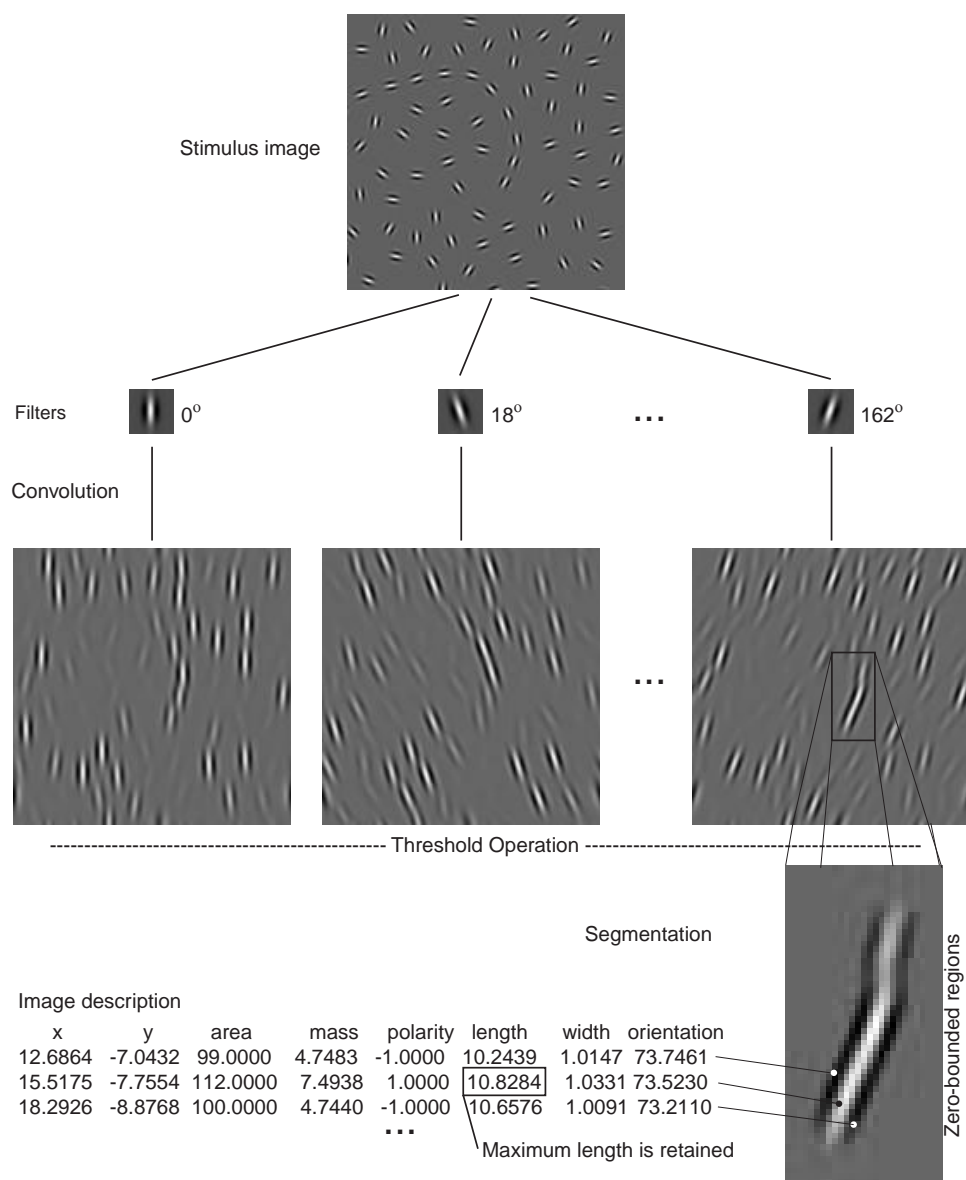


Figure 3-1. Outline of the simple filter model's filtering and segmentation process. The first row shows the initial raw image. The second row shows the filters used within a particular spatial-frequency. The third row shows the result of convolving each filter with the raw image. Finally, for illustrative purposes the fourth row shows the filtered image section that contains the longest zero-bounded regions. The verbal description of these is shown to the left. Finally, the length of the longest ZBR is retained – this value informs 2AFC decisions.

3.2.1.2 Convolution and segmentation of all images used within an experiment

The process described in section 3.2.1.1 was applied separately with varying filter scales. The width and length of filters was varied from 0.5 to 10σ in steps of 0.5σ , where σ corresponds to the standard deviation of the filters extent in pixels. Thus, 400 separate filters are employed in total. The longest ZBR was retained for each individual image and each particular filter scale.

3.2.1.3 Running a simulated 2AFC experiment

A simulated experiment was run for each filter scale. Where the original stimulus images used with human observers were available, the target/noise image pairings were exactly matched for the model and the human observers. Thus, the model was exposed to exactly the same target/noise image pairs as *all* the subjects. For every 2AFC decision, the maximum ZBR length for each image was retrieved. These values were then compared, if the target image has the longest zero-bounded region then the response was classified as a hit otherwise it was a miss. The maximum ZBR lengths of target and noise images were never equal. Subsequently a percent correct value was calculated for each stimulus level.

For some experiments (i.e. path-angle, below) the stimuli used with human observers were generated online. Thus, the original images were unavailable. In these circumstances, a ‘random’ experiment was generated as described in section 2.1.2. This random experiment consists of 2000 2AFC decisions.

3.2.2 Presentation of experimental results

Each simulation involved filters with 400 different scales (20 widths x 20 lengths) and ten different orientations. Furthermore, for each set of experimental stimuli there were a number of different levels of the independent variable and 10 images for each of these levels. It was deemed necessary to summarise the performance of the simple-filter model. Section 3.2.1.1 has already explained that only the maximum ZBR length across all filter orientations is retained. This is a reasonable simplification as the overall orientation of embedded contours is randomised. Sections 3.2.2.1 - 3.2.2.4 describe how the performance of the simple-filter model will be illustrated. The aim of these presentations is to summarise the performance of the model with filters of different scales, whilst enabling a comparison with the performance of human observers.

3.2.2.1 Presentation of the difference in maximum ZBR lengths for target and noise images.

The performance of the simple filter model is based upon the relative lengths of the ZBRs in target and noise images. More precisely, the longest ZBR in each image is retrieved and the assignment of ‘target’ status in 2AFC decisions is based upon which image has the longest ZBR. For each level of the independent variable, there were ten target images. There were fifty noise images which did not feature an embedded contour. The mean of the maximum ZBR lengths for all the images at each stimulus level was calculated. In order to

maintain clarity this value will be referred to as the maximum ZBR length, despite the fact that it is actually the mean of all the longest ZBRs for the images at each stimulus level.

The performance of the simple-filter model is based upon its ability to discover longer ZBRs for target images than it does for noise images. Therefore, the maximum ZBR length for noise images is then subtracted from the maximum ZBR length for each level of target image. Where the result is a positive value, the simple filter model is likely to successfully select target images. This process is repeated for each filter size and the results are plotted upon a contour plot. The peaks within these plots indicate those filters that are better suited to the detection of images containing embedded contours.

Separate contour plots are generated for each level of the stimulus variable under investigation. For example, where there are nine levels of path-angle in the experiment described in Section 3.3.1, then there will be nine different contour plots. Each individual contour plot has a heading that indicates the stimulus level that is represented. The heading “target @ 80° - noise” indicates that the path-angle between contour elements was 80°.

3.2.2.2 Counting the number of elements falling within the longest ZBR

In the thesis introduction, contour integration was defined as “... *a process through which contours that are formed from discrete elements can nevertheless be identified by the observer as a single object in the visual scene.*”. The extent to which the simple-filter model conforms to this definition is determined by the number of contour elements which are located within the longest ZBR.

It was possible to quantify the number of Gabor elements that contributed towards a particular ZBR. This was achieved by counting the number of elements that lie within the area defined by the ZBR. This value was calculated by taking the filtered and thresholded image that contained the longest ZBR and then setting values within this ZBR to one, all other areas are set to zero. Another image was created, in this image the locations that correspond to the centres of Gabor elements are set to one. The product of corresponding points in these images was then calculated and the sum of this value indicates the number of Gabor elements that contributed towards the formation of the longest ZBR. This process is illustrated in Figure 3-2. This could be considered a measure of the degree of integration that has occurred within the simple filter model. If all contour elements were represented within the longest ZBR then the whole contour was integrated. With fewer elements represented in the ZBR then integration could be considered to have been less successful.

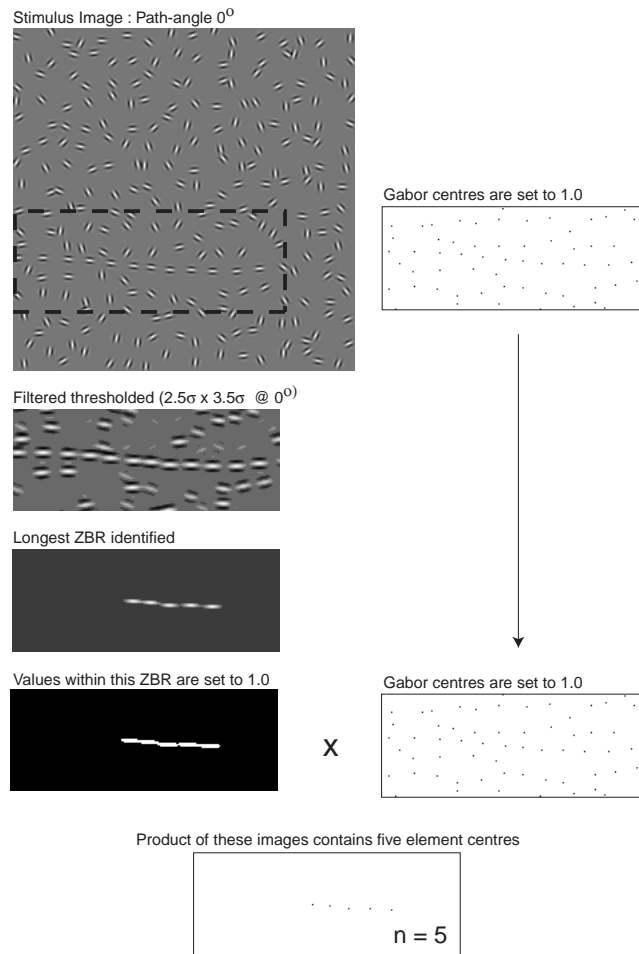


Figure 3-2. Illustration of the process of estimating the number of Gabor elements that fall within the area defined by the longest ZBR.

As this measure of contour integration requires prior knowledge of the position of contour elements, information that the simple filter model would not possess, it is not chosen as the output of the simple-filter model that determines 2AFC decisions. However, the measure is retained as a means of illustrating the extent of contour integration that has occurred within the simple filter model.

3.2.2.3 *Estimation of the biological plausibility of particular filter scales*

Parker and Hawken (1988) measured receptive field elongation in macaques. They found that the ratio was 3.2:1 (length:width) for the central excitatory region. Filter scales reported in the current thesis refer to the whole filter. There is a linear relationship between the extent of the central excitatory region and that of the whole filter and a conversion factor can be derived by plotting each measure for a range of filters. Translating Parker and Hawken's ratio using this method gives a ratio of 1.3:1 for the whole filter (pers. comm. Watt, 2000). This ratio is presented as a dashed line upon contour plots. This facilitates a visual estimate of the biological plausibility of particular filter scales. Those filter scales

that lie close to the line are more likely to be biologically plausible, or at least the occurrence of such filters in primary vision is higher.

3.2.2.4 *Presentation of simple filter model performance with particular filters*

The filter that gives rise to the performance levels that were closest to those of human observers was identified. This was achieved by calculating the sum of squared differences ($\sum D^2$) between the percent correct scores for human observers and the simple-filter model (Figure 3-3). Filters which gave rise to performance levels which were least different from the performance of human observers were selected for further examination.

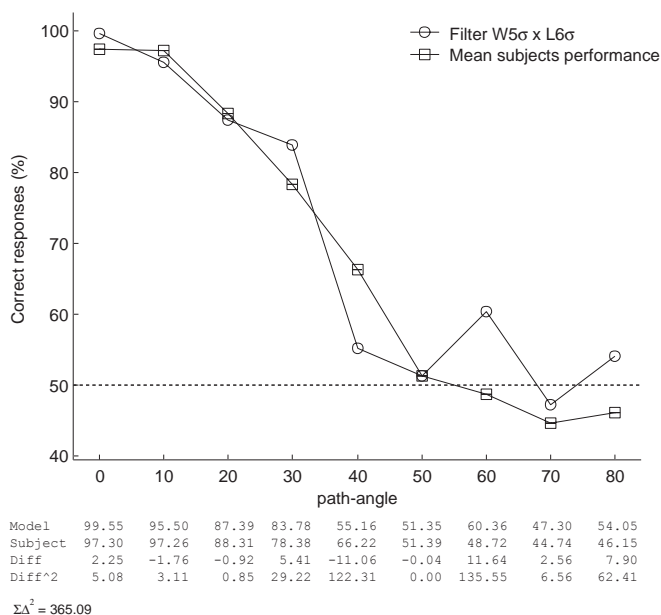


Figure 3-3. Illustration of the calculation of the difference between human subject and simple-filter model scores.

The filter scale that gives rise to the highest overall rate of detection for a set of target images was also identified. The performance of both these filters will be shown as a function of the manipulated stimulus parameter, i.e. path-angle or element spacing. Fitted psychometric functions are also presented, for the closest fitting filters, in order to aid comparison with human observer performance. Psychometric functions were generated according to the procedure described in Section 2.1.4.1.

3.3 **Simulation 1: Modelling the influence of path-angle (expt. 1).**

3.3.1 Stimulus image parameters.

A new set of images was constructed with identical parameters to those employed during experiment 1. A detailed description of the image properties is provided in Section 2.2.2.

Briefly, the overall size of image was 256x256 pixels. The contour path was generated using the conventional Field, Hess and Hayes (1993) method. Target contours consisted of 12 Gabor elements. The orientation change across consecutive path elements was varied from 0-80° in steps of 10° (path-angle) and the path-relative orientation of individual elements was jittered by $\pm 5^\circ$. A distance of 16 pixels separated consecutive path elements. The position of background elements was determined using the jittered-grid method described in Section 2.1.1.2. Background elements were separated by a mean distance of 16 pixels. The orientation of each background element was randomised between 0-180°. Ten images were generated for each of the nine levels of path-angle.

Fifty noise images were also generated. These were identical to target images, except that they did not feature an embedded contour. In order to compare the model's performance when Gabor patches were phase-symmetrical or phase-alternated two sets of stimulus images were constructed. In each set, the positions of all Gabor elements were identical, only the phase of Gabor patches differed. The parameters used to generate the Gabor patches were the same for those used in conjunction with human observers (details are available in Section 2.1.1.3).

3.3.2 Procedure

The simple-filter model was presented with each target and noise image following the procedures described in sections 3.2.1.1 and 3.2.1.2. A simulated experiment, featuring 2000 pairs of target and noise images, was constructed using the procedure outlined in section 3.2.1.3. Note that the selection of 2AFC image-pairs, that formed a simulated experiment, was randomised for each filter-scale.

3.3.3 Simulation results

The contour plots presented in Figure 3-4 illustrate the difference in the ZBR lengths for target and noise images at each level of stimulus path-angle and for each individual filter scale. The upper array of plots (*labelled A*) shows the results for stimuli featuring phase-aligned Gabor elements, whilst the lower array of plots (*labelled B*) shows performance with stimuli featuring phase-alternated Gabor elements.

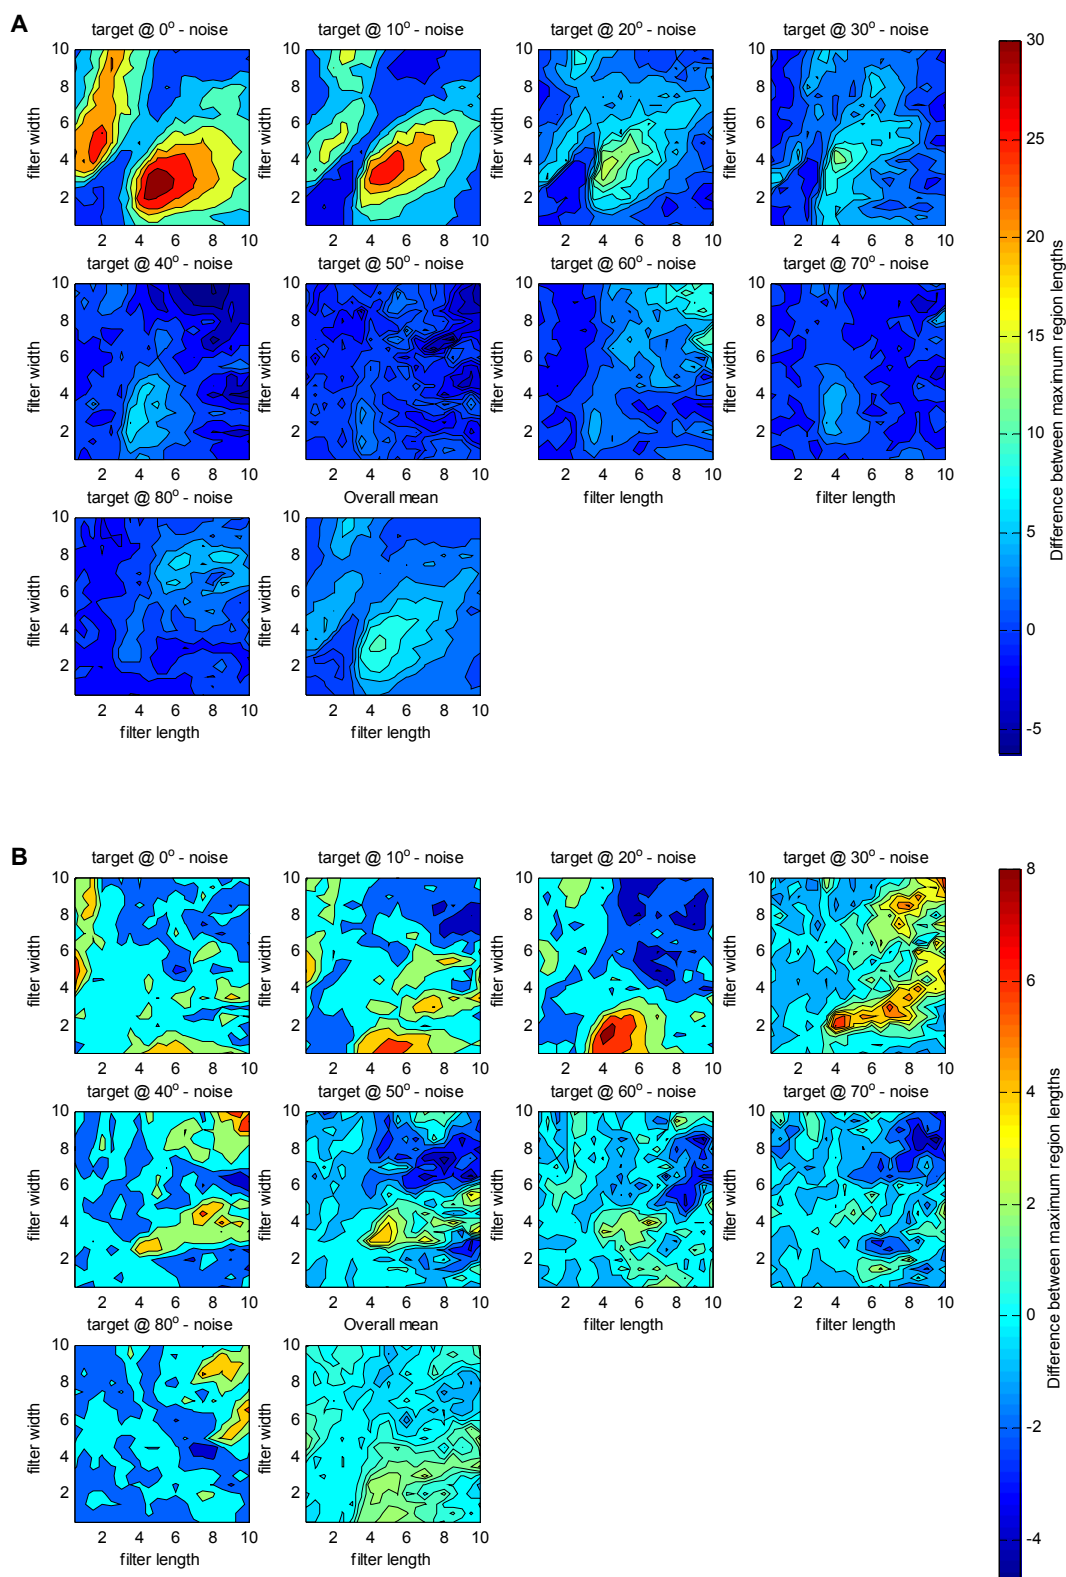


Figure 3-4. Contour plots indicating the difference in maximum ZBR length for noise and target images, each individual plot corresponds to a different level of the embedded contours path-angle. A larger difference indicates a filter that may be helpful in distinguishing between target and noise images. The data in these plots corresponds to images with phase-aligned elements (A) and phase-alternated elements (B).

For illustrative purposes, the elongated filter with percent correct scores that most closely matched those of human observers is indicated within each contour-plot (○). The filter with the highest overall performance across all levels of path-angle is also shown (□). The actual performance levels for the closest-fitting and the highest performing filters are plotted at each level of path-angle (Figure 3-5).

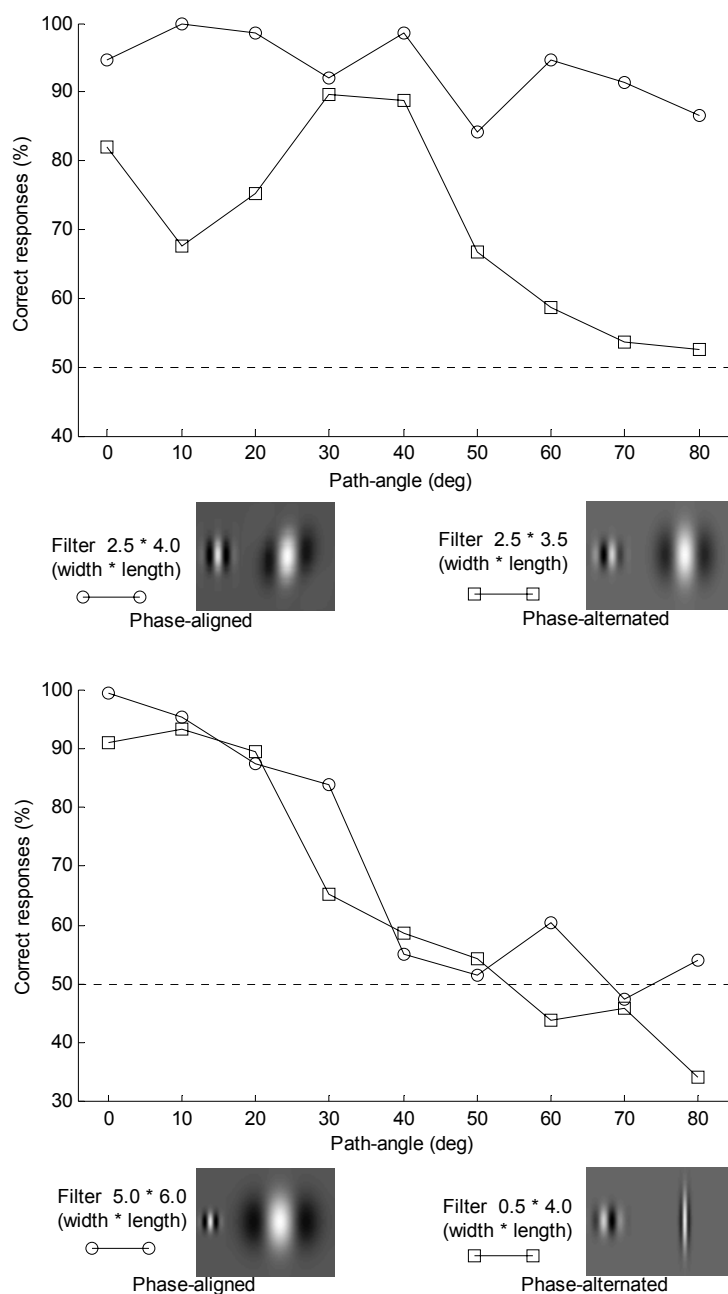


Figure 3-5. Contour detection as a function of path-angle for selected filters. The upper plot shows performance levels for those filters with the highest detection rates with phase-symmetrical and phase-alternating elements. The lower plot shows the performance levels for those filters with performance levels closest to the mean for human observers. The inset images represent the respective filters (not to scale). A Gabor patch is inset to the left of each image in order to illustrate the relative size of each filter. Symbols : (○) phase-aligned stimuli, (□) phase-alternated stimuli.

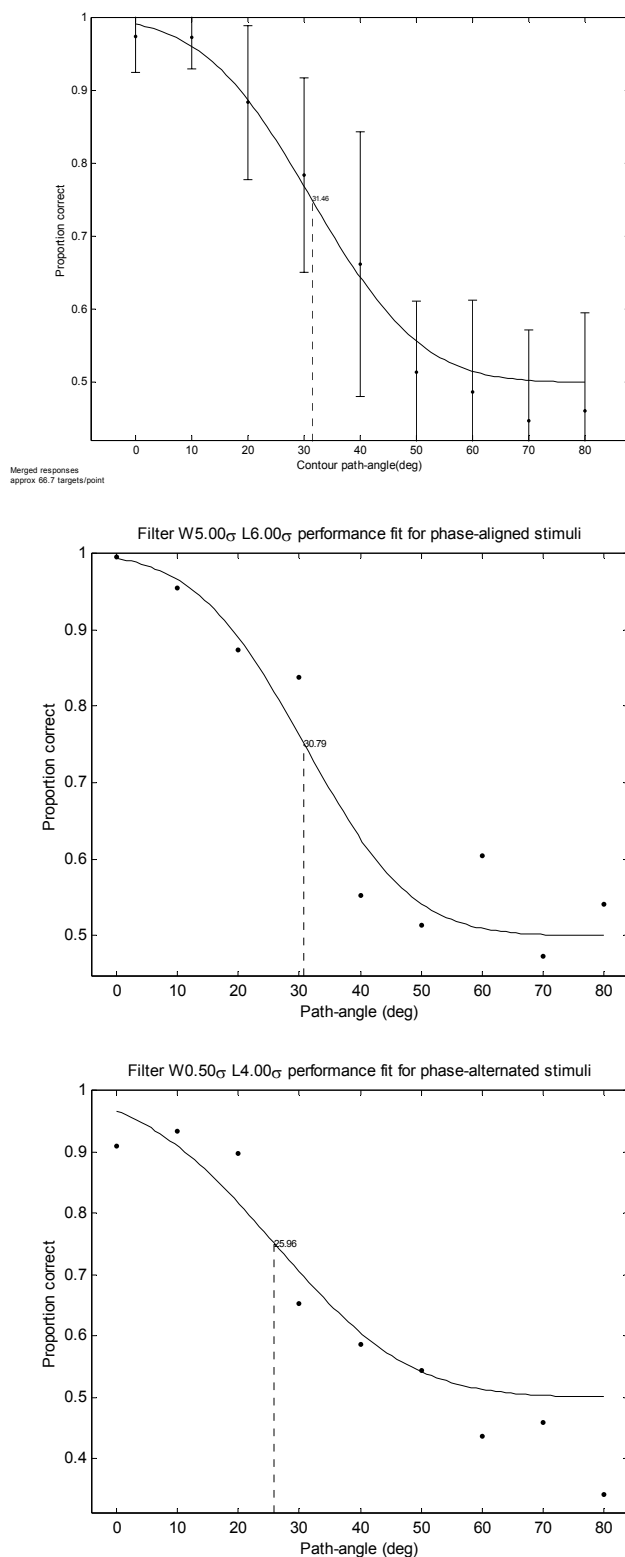


Figure 3-6. Fitted psychometric curves as a function of path-angle. Plots illustrate performance for mean performance of human observers ($n=6$) (top) and simple-filter model performance with phase-aligned elements (middle) and phase-alternated elements (bottom). Hess and Dakin (1999) have offered evidence that within foveal regions there is little difference in the detection rates for stimuli with different phase properties. Therefore, the human data illustrates performance with phase-aligned elements.

Three further plots are presented, Figure 3-6, these show the contour detection rates as a psychometric function of path-angle. These plots show performance of the simple-filter model with phase-aligned and phase-alternated stimuli. For comparison Figure 3-6(top) shows the psychometric fit for the mean performance of human observers.

3.3.3.1 *Assessing the biological plausibility of the best-fitting and best-performing filters.*

The filter elongation ratio provided by Parker and Hawken (1988) is shown as a solid diagonal line upon the contour plots below: for phase-aligned stimuli (Figure 3-7) and for phase-alternated stimuli (Figure 3-8).

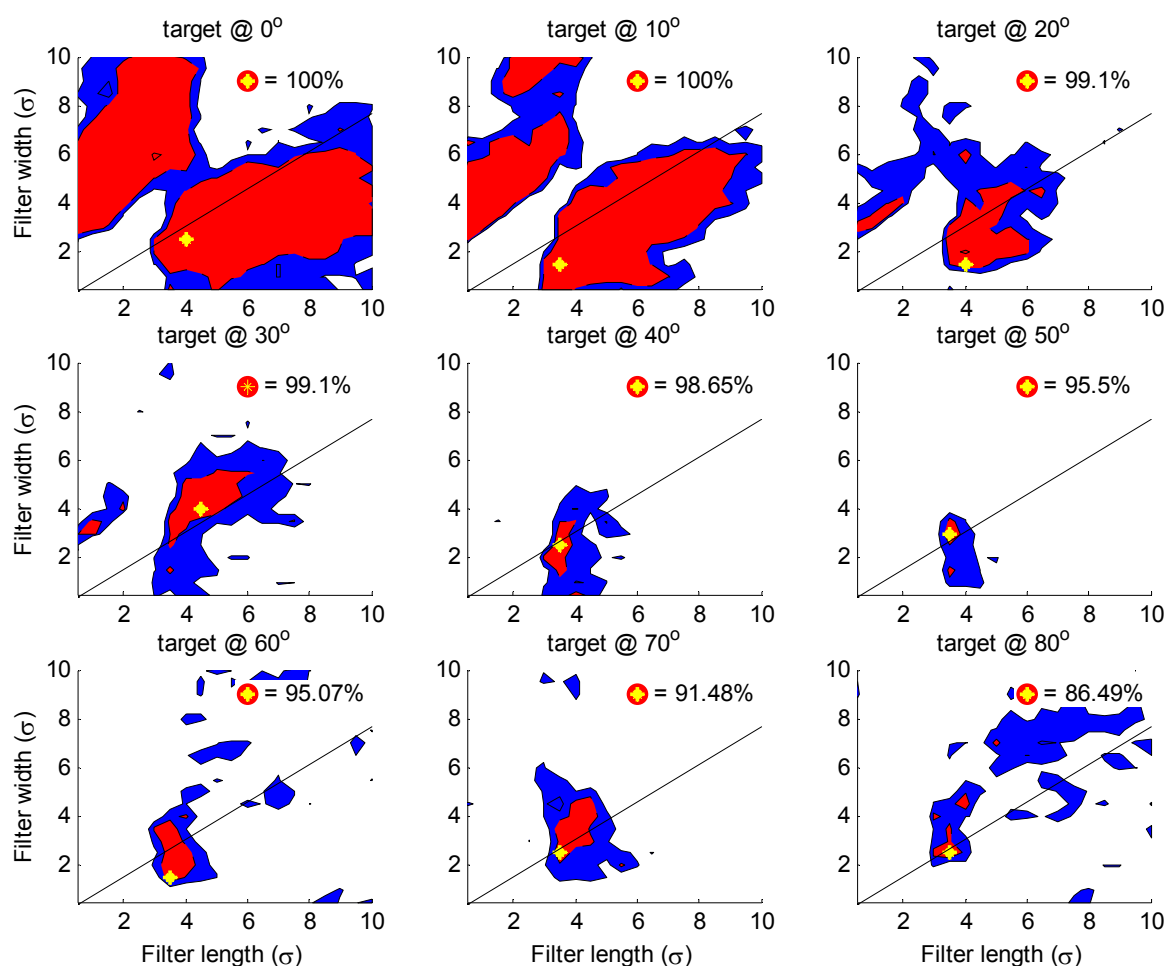



Figure 3-7. Contour plot illustrating those filter-scales which gave rise to the greatest contour detection rates for phase-aligned stimuli. The maximum level of performance achieved is represented by the  symbol. The red regions within each plot show those filter scales where the performance level was within 10% of the maximum. The blue regions represent those areas where performance was within 20% of the maximum. The diagonal line represents the simple-cell elongation ratio offered by Parker and Hawken.

For phase-aligned stimuli there is a close correspondence between the Parker and Hawken elongation ratio and those filters which gave rise to the best performance in the path-paradigm task.

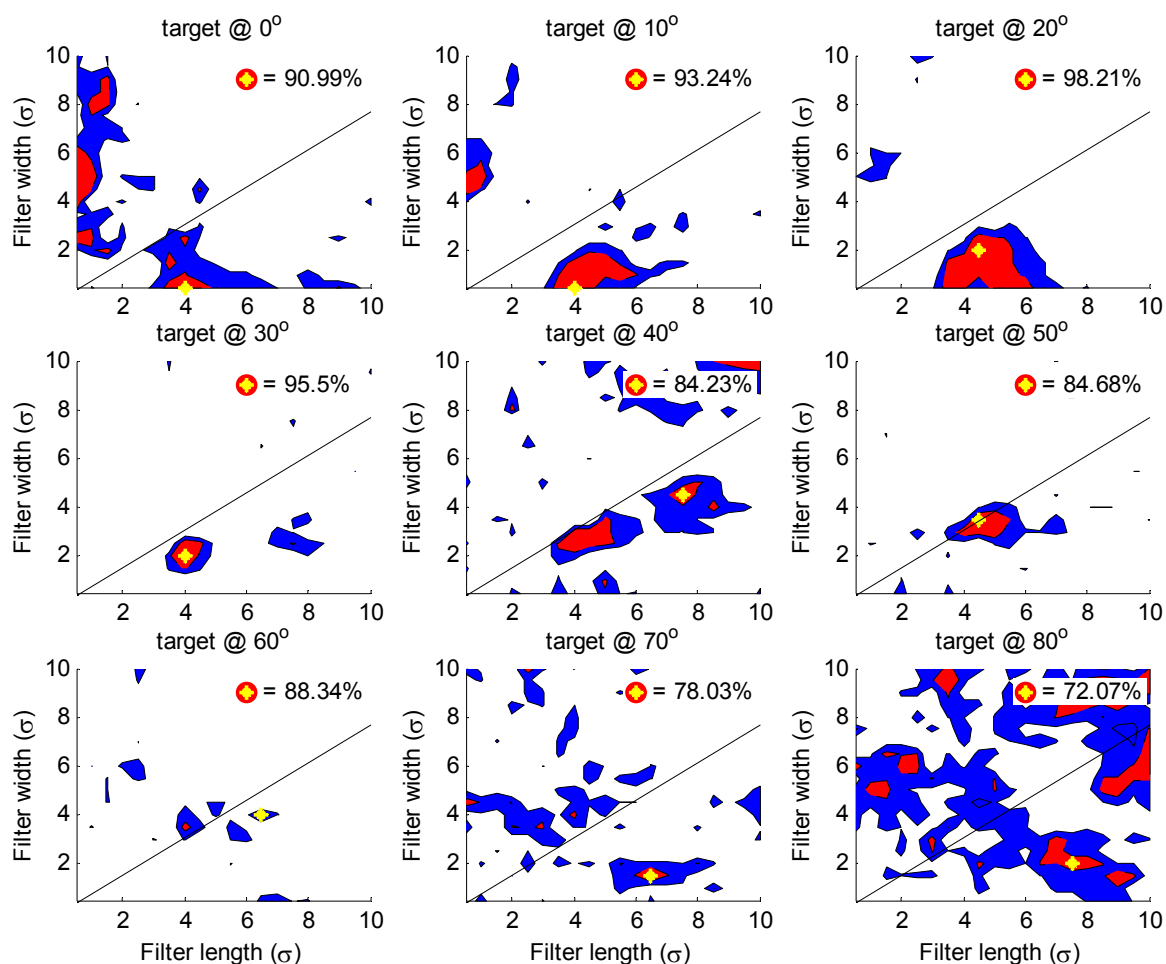


Figure 3-8. Contour plot illustrating those filter-scales which gave rise to the greatest contour detection rates for phase-alternated stimuli. The maximum level of performance achieved is represented by the symbol. The red regions within each plot show those filter scales where the performance level was within 10% of the maximum. The blue regions represent those areas where performance was within 20% of the maximum. The diagonal line represents the simple-cell elongation ratio offered by Parker and Hawken.

For the phase-alternated stimuli the correspondence between the filter elongation ratio and those filters which gave rise to the highest level of performance was only close for path-angle levels between 10° and 70°. Beyond this range the best performing filters were either relatively narrow or quite wide, this issue is discussed within 3.6.

3.3.3.2 *How many elements are represented within the longest ZBRs?*

Using the technique described in 3.2.2.2 the number of individual elements that contributed towards the formation of the longest ZBR are measured.

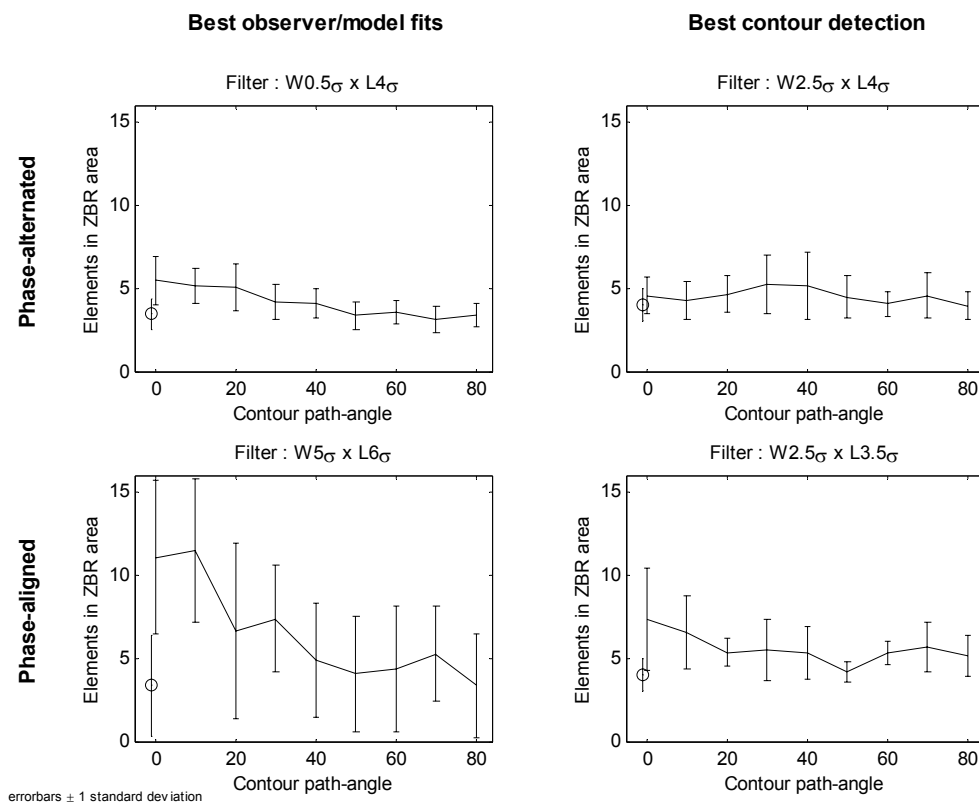


Figure 3-9. Plots illustrating the mean number of elements that are represented by the longest ZBRs found within the filtered images. Plots represent values for phase-aligned and phase-alternated stimuli and for the best fitting filters and the best performing filters. The plots show mean element counts for each level of path-angle, the mean element counts for noise images is also included (o symbols).

The data presented above confirm that for the simple-filter model fewer Gabor elements are represented by the longest ZBR when stimuli are phase-alternated. Only around five elements contribute towards the longest ZBRs in phase-alternated elements, while for phase-aligned stimuli as many as twelve elements are counted in the longest ZBRs.

3.3.4 Discussion

The simple-filter model is able to match, indeed with some filters it can even surpass, the performance of human subjects in the path-angle experiment. This is particularly true for stimuli composed of phase-aligned Gabor patches. The performance of the simple filter model is generally lower when stimuli feature phase-alternated Gabor elements, though at some filter scales (W0.5 σ x L4.0 σ) performance is comparable to that of human observers.

The psychometric function of human observers is closely matched to the psychometric functions for the best fitting filters, regardless of phase. However, the correspondence between the mean detection rates and the psychometric function is closer for the phase-aligned stimuli. For the phase-alternated stimuli the correspondence is reduced as there is a slight fall in detection rates with straighter contours (0-20° path-angle). This is apparent when Figure 3-4(B) is examined; for path-angles of 0-10° there are few filter-scales which give rise to longer target ZBRs for target images than noise images. Therefore, it is inevitable that performance in 2AFC decisions will be lower for these particular stimuli. The cause of the impaired detection of straighter phase-aligned contours is investigated in Section 3.6.

There is little correspondence between the results of the current modelling exercise and those reported by Hess and Dakin (1999). The following section examines possible causes for these differences in the reported performance of the simple-filter model.

3.4 Why do the current simulation results differ from those of Hess and Dakin (1999)?

The stimulus images used in the current simulations differ, on a number of measures, from those used by Hess and Dakin. This is because these stimuli were used in psychophysical experiments that were conducted prior to the publication of the Hess and Dakin report. However, it is possible to identify correspondences between the stimuli used. A direct comparison between the performance of the Hess and Dakin simple filter model and that currently reported is difficult as images were scaled differently and different types of filters were employed. The quoted value of σ for the Hess and Dakin simple filter model is 4.5 pixels. However, they report that their images and the filter that they employed were both scaled-down prior to processing. Images were reduced from 624² pixels to 128² pixels. This corresponds to a scaling ratio of 1:0.205. In order to compare the results of these simulations it is necessary to assess the equivalence of the stimulus images and of the filters employed. The following sections will attempt this assessment.

3.4.1 Comparing stimulus images

Estimates of the characteristics of Hess and Dakin's downscaled images are provided below, properties of the path-angle images used in this section are included for comparison.

Table 3-1. Stimulus properties for images used in the Hess and Dakin (1999) psychophysical and simple-filter modelling experiments. In the current thesis images were identical for psychophysical experiments and in simulations.

<i>Stimuli Property</i>	<i>Image used in psychophysical experiment (Hess and Dakin,1999)</i>	<i>Image used in simple-filter simulation (Hess and Dakin,1999)</i>	<i>Stimuli employed in the current simulation</i>
<i>Image size</i>	624 ² pixels	128 ² pixels	256 ² pixels
<i>Gabor wavelength</i>	20 pixels/cycle	4.1 pixels/cycle	4 pixels/cycle
<i>Contour spacing</i>	67 pixels	13.74 pixels	16 pixels
<i>Contour length</i>	8 elements	8 elements	12 elements
<i>Background spacing</i>	67 pixels	13.74 pixels	16 pixels
<i>Number of background elements</i>	13*13 = 169 169-8 = 161	13*13 = 169 169-8 = 161	16 x 16 = 256 256 -12 = 244
<i>Filter orientation resolution</i>		12 orientations in 15° steps	10 orientations in 18° steps

The wavelength of Gabor elements and their spacing are most likely to influence the performance of the simple filter model. A comparison of these properties reveals that images used in the current experiment and the scaled images of Hess and Dakin only differ in a few respects. Contours are longer within the current simulations, this will have increased contour detection rates. The overall size of the stimulus background is greater in the current experiment, hence it contains a larger number of elements (244 versus 161 elements). Raising the number of background elements will increase the occurrence of spurious co-alignments of background elements. However, this increase occurs in target and noise images, so the overall patterns of performance should be unaffected by this difference, though the variability of performance may increase as a function of the extent of the stimulus background.

The most important correspondences between the Hess and Dakin stimuli and those employed in the current chapter are the Gabor patch wavelength and element spacing. There is a close correspondence between the sinusoid cycles (4.0 Vs 4.1 pixels), whilst the element spacing for Hess and Dakin's stimuli is only slightly reduced. Therefore, a filter that performs well with the Hess and Dakin stimuli should perform as well with the stimuli used in the current simulations.

3.4.2 Comparing the filters currently employed and that used within the Hess and Dakin model

In order to make a direct comparison between the performance of the current simple-filter model and Hess and Dakin's model, it is necessary to identify the closest matching filter in the current experiment to that used by Hess and Dakin.

When the filter used by Hess and Dakin was scaled-down to the same extent as their images, the quoted filter scale (σ) of 4.5 pixels becomes 0.923 pixels. The current model utilises a 2nd derivative of a Gaussian filter, whilst Hess and Dakin use a difference of Gaussian filter (DoG). Unfortunately, the same value of σ within each function does not give rise to filters with an equivalent spatial bandwidth. Therefore, it was necessary to find closest matching 2nd derivative of a Gaussian filter to the 0.923pixel DoG filter.

This was achieved by minimising the sum of squared differences between the kernel of each filter while varying the length and width of the 2nd derivative of a Gaussian filter. Details of this procedure are included in Appendix H. The closest matching filter had a width and length of $1.586\sigma \times 2.758\sigma$. Thus, the closest matching filter employed in the current chapter had a scale of $1.5\sigma \times 3.0\sigma$ (width x length).

3.4.3 Comparing current performance with a filter that is comparable to that employed by Hess and Dakin

The performance of the 0.923σ (pixel) DoG filter utilised with the Hess and Dakin should be comparable to the performance of the $1.5\sigma \times 3\sigma$ filter in conjunction with the stimuli described in the current section. The detection rates of the current simple-filter model shown below (Figure 3-10), illustrate that performance was reasonably good for symmetrical-phase stimuli but lower for alternate-phase stimuli.




Figure 3-10. Target stimuli detection rates as a function of contour path-angle. The plots illustrate the performance for the current implementation of the simple filter model with phase-symmetrical and phase-alternated stimuli when filtering is conducted with a $W1.5\sigma \times L3.0\sigma$ filter, the nearest filter to that used by Hess and Dakin. The dashed lines show the performance levels reported by Hess and Dakin. (*Hess & Dakin performance levels are taken from Hess and Dakin (1999) figures 8 & 9.*)

The current implementation of the simple-filter model performs moderately well with the filter that most closely matches the filter employed by Hess and Dakin, the performance of each implementation of the simple filter model is shown in Figure 3-10. The overall pattern of results are loosely equivalent in both simulations, inasmuch as detection rates for phase-aligned stimuli are superior to the detection rates for phase-alternated stimuli in both cases. The remaining differences in detection rates for each simulation may be due to the differences between stimuli identified in Table 3-1.

3.4.4 Discussion

Hess and Dakin report that a DoG filter where σ is 4.5 pixels will give a moderate level of performance for phase-aligned stimuli while the performance is much lower for phase-alternated stimuli. The closest matching filter used in the current simulations led to a roughly equivalent pattern of performance ($W1.5\sigma \times L3\sigma$). This filter is slightly smaller than those found to be optimal in the current simulations. This filter may have been

favoured in the simulations of Hess and Dakin because elements in their stimuli were slightly more closely spaced.

Hess and Dakin report that a 4.5σ DoG filter led to the best overall performance in their implementation of the simple-filter model, regardless of element phase. In the current simulations, this result was not found, differently scaled filters are better suited to the detection of stimuli with different phase properties. Generally, much narrower filters were better suited to the detection of contours formed from phase-alternated stimuli. This discrepancy may occur because the stimuli used in the current simulations are not comparable to those that were used by Hess and Dakin. Alternatively, the criterion adopted for the selection of the filter scale may be the most important difference. Hess and Dakin, in their implementation of the simple-filter model, sought a filter-scale that performed optimally for both phase-aligned and phase-alternated stimuli. The results of the current simulations reveal that target images drawn from each of these broad categories of stimuli are best detected by a simple-filter model that utilises differently scaled filters.

3.5 Simulation 2: The influence of contour element spacing (Expt. 2).

3.5.1 Overview

This section describes simulation experiments where the spacing between consecutive path elements is manipulated. As with experiment 2, only the density of contour element was manipulated, thus the experiment is similar to that reported by Kovacs, Polat and Norcia, (1997), rather than Field, Hayes and Hess (1993, experiment 4). Only phase-aligned stimuli were utilised during this simulation, due to the fact that human observers were only tested with such stimuli.

3.5.2 Stimuli image parameters

The stimuli were identical to those used in experiment 2 (section 2.3). Briefly, the embedded contours consisted of 15 Gabor patches with a path-angle of $\pm 20^\circ$. The spacing between contour elements was varied between 16-25.6 pixels in steps of 2.4 pixels. Path-angle *jitter* was $\pm 5^\circ$, contour element spacing jitter was 25% of the contour element spacing and element orientation jitter was $\pm 5^\circ$. The mean spacing for elements within the stimulus background was 16 pixels and the orientation of these elements was randomised. All Gabor elements were phase-aligned.

3.5.3 Procedure

The simple-filter model was presented with each target and noise image following the procedures described in sections 3.2.1.1 and 3.2.1.2. In experiment 2 (section 2.3) each observer ($n = 36$) viewed 200 randomly selected pairs of target and noise images⁶. For the current simulation, the sequence of image-pairs viewed by these subjects was concatenated. Thus, the simple-filter model was required to make 7200 2AFC decisions. This sequence of image-pairs was presented to the model in conjunction with each separate filter-scale.

3.5.4 Simulation results

Figure 3-11 shows an array of contour plots, each illustrates the differences between the maximum ZBR length for each level of target images and the noise images. For each level of element-spacing there is distinct peak in the contour plot. The location of the peak varies with contour element spacing. There is a positive correlation between the spacing of contour elements and the length of the best-performing filter. The optimal filter length for contours spaced at 4λ (16 pixels) is 4.5σ and for elements at 6.4λ (25 pixels) the optimal filter length is approximately 7σ . The relationship between filter width and element spacing is less consistent. There appears to be a negative relationship between element spacing and filter width, at 4λ the best performing filter has a width of 4σ , whilst at 6.4λ the width is 1.5σ .

⁶ Details of the randomisation procedure are summarised in Section 2.1.2.

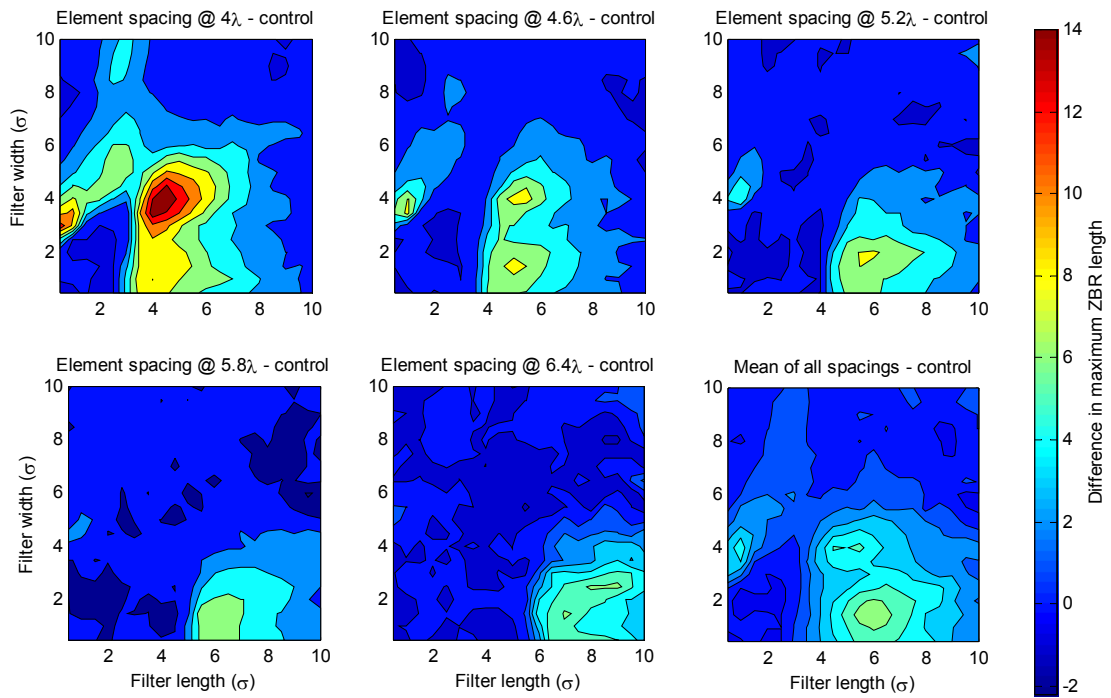


Figure 3-11. Contour plots indicating the difference in the maximum ZBR lengths of target and noise images. Each individual contour plot corresponds to a different level of the embedded contours element spacing. A larger difference indicates a filter that may be helpful in distinguishing between images which do and which do not feature an embedded contour.

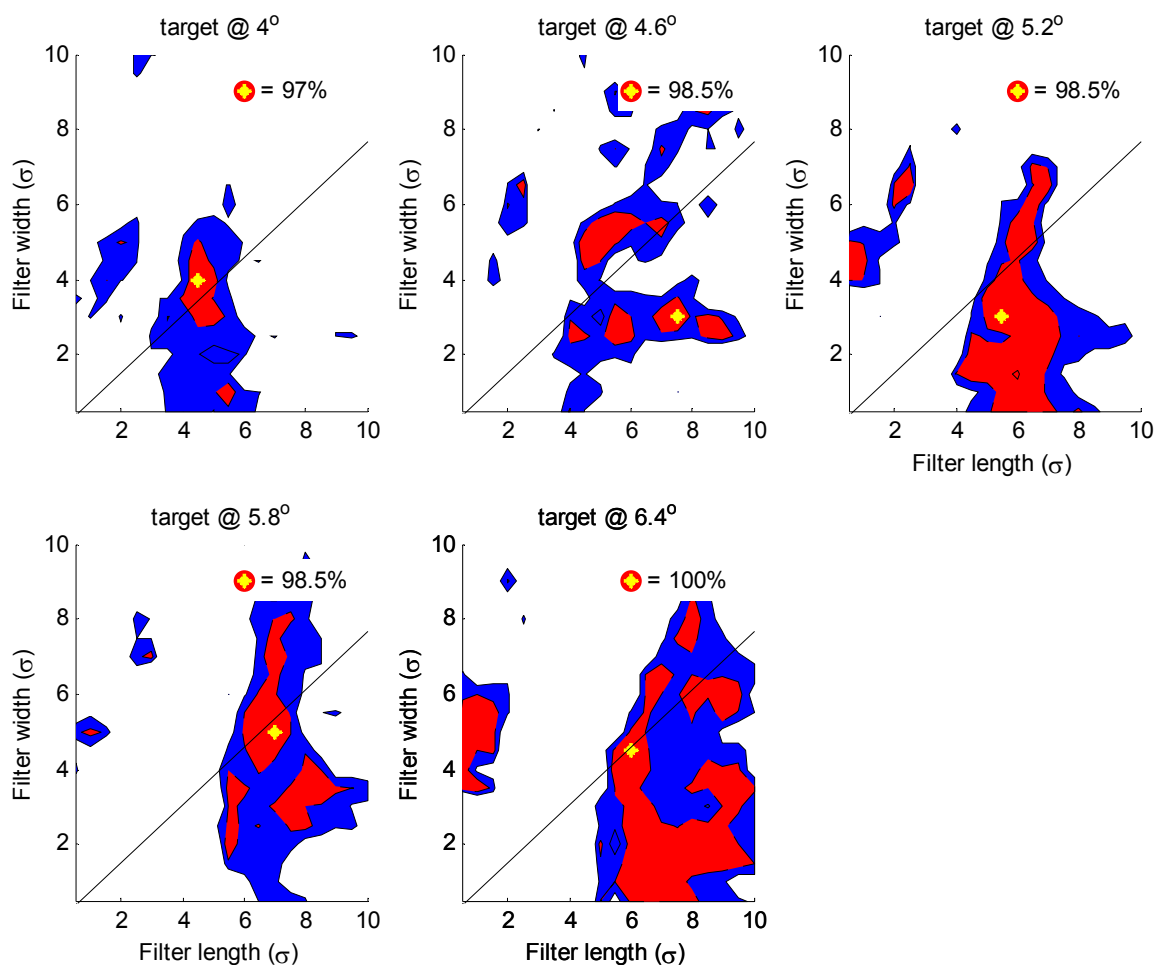



Figure 3-12. Contour plot illustrating those filter-scales which gave rise to the greatest contour detection rates (% correct) for each level of element spacing. The maximum level of performance achieved is represented by the  symbol. The red regions within each plot show those filter scales where the performance level was within 10% of the maximum. The blue regions represent those areas where performance was within 20% of the maximum. The diagonal line represents the simple-cell elongation ratio offered by Parker and Hawken.

Those filters-scales which gave rise to the best performance for each stimulus level are highlighted in red within Figure 3-12. There is a distinct region, centered about $L6\sigma \times W4\sigma$, where filters are particularly well suited to detecting the embedded contours. The closest match between the mean performance of the human observers and the performance of the simple-filter model was identified using the method described in Section 3.2.2.4. This filter has an elongation ratio ($W4.5\sigma \times L6.0\sigma$) that is very close to the estimate offered by Parker and Hawken (1988) for simple-cells in macaques.

Figure 3-13 shows the percent correct detection rates as a function of element spacing for the filter with the highest overall detection rate and for the filter with a detection rate that is closest to that of human observers.




Figure 3-13. Model performance for the filter which exhibits the best overall 2AFC success (\square) and for the filter with performance level most like that of human observers (\circ). The mean performance levels for human subjects (n=36) is also included for comparison (grey line with error bars representing 1 standard deviation). The inset images represent the respective filters (not to scale). A Gabor patch is inset to the left of each image in order to illustrate the relative size of each filter.

Figure 3-14 (bottom) shows the simple-filter models performance with the $W4.5\sigma \times L6.0\sigma$ filter as a psychometric function of contour element spacing. For comparison Figure 3-14 (top) shows the equivalent plot for human observers.

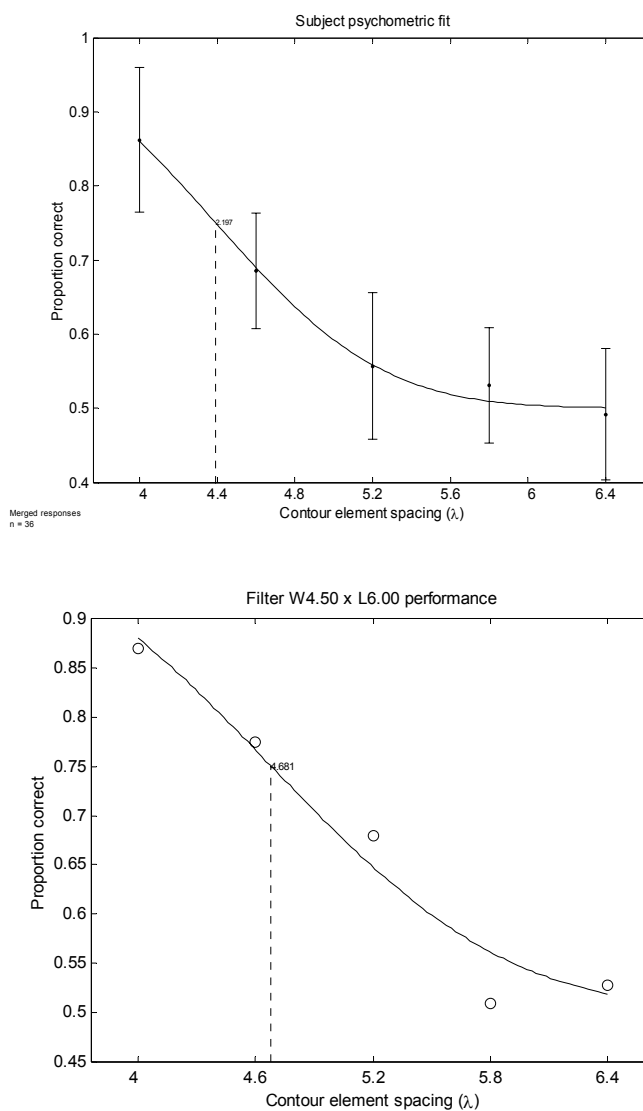


Figure 3-14. Fitted psychometric curves as a function of contour element spacing. Plots illustrate performance for mean subject performance (top) and model performance with filter $W4.5\sigma \times L6.0\sigma$.

Following the procedure described in section 3.2.2.2 the number of stimulus element that lie within the area of the longest ZBRs was calculated. These counts are presented as a function of element spacing within the plots below (Figure 3-15).




Figure 3-15. Number of elements within longest ZBR region as a function of element spacing. The (o) symbol represents the number of elements within the longest ZBRs of noise images.

3.5.5 Discussion

Clearly, the *simple-filter* model is able to detect contours with varying degrees of element spacing. Furthermore, the filter that produced detection rates that were the closest fit to observer detection rates had a length/width ratio that conforms to physiological estimates. There are filter-scales which lead to performance levels that exceed those found in human observers. However, these filter-scales (W1.5 σ x L6.0 σ) have an elongation ratio that is quite distant from the ratio suggested by Parker and Hawken (1998).

The contour plot illustrating overall detection rates Figure 3-12 reveals that the filter which gave rise to the closest match in performance to that of human observers was scaled quite differently from the filter which gave rise to the best overall performance. The implication is that, if humans do utilise a simple-filter mechanism, they are not using filters which might lead to better overall performance. This suggests that such filters are unavailable to the human observers, perhaps because they are quite distant from the scale of elongation offered by Parker and Hawken (1988). Alternatively, such filters may be available,

however their output is disregarded by the mechanism responsible for the selection of filter scales.

3.6 Simulation 3: Why is the simple-filter model impaired when detecting phase-alternated contours with low levels of curvature?

3.6.1 Overview

Section 3.3 reported that performance of the simple-filter model fell slightly for straighter contours (0-10°), where stimuli were composed of phase-alternated elements. A close examination of Figure 3-4 (B) reveals that for contours with a path-angle of 0-10° there are few filter scales which result in longer ZBRs for target images than for noise images. The filter that led to the highest overall level of performance for phase-alternated stimuli was actually close to the ratio of simple-cell elongation offered by Parker and Hawken (1988). However, while performance levels with this filter ($4.0\sigma_L \times 2.5\sigma_W$) fell for the straightest contours (see Figure 3-5), for human observers the reverse is true, i.e. there is a monotonic increase in contour detection rates as path-angle is reduced. This weakens the case for the simple-filter account of contour detection, at least in conjunction with phase-alternated stimuli.

Figure 3-16 illustrates how slight misalignments of contour elements lead to longer ZBR lengths in phase-alternated contours. In contours A-C the path-angle of contours was varied between 0-20°, whereas in contours D-F the orientation jitter of elements was varied by the same degree. For phase-aligned stimuli (Figure 3-16, lower image) the length of ZBRs was negatively correlated with path-angle and was not affected by the relatively small manipulations of orientation jitter. However, for phase-alternated stimuli (Figure 3-16, upper image), the relationship between ZBR length and element misalignment is more complex, ZBR lengths peaked when path-angle was 10° (B) and where orientation jitter was 20° (F).

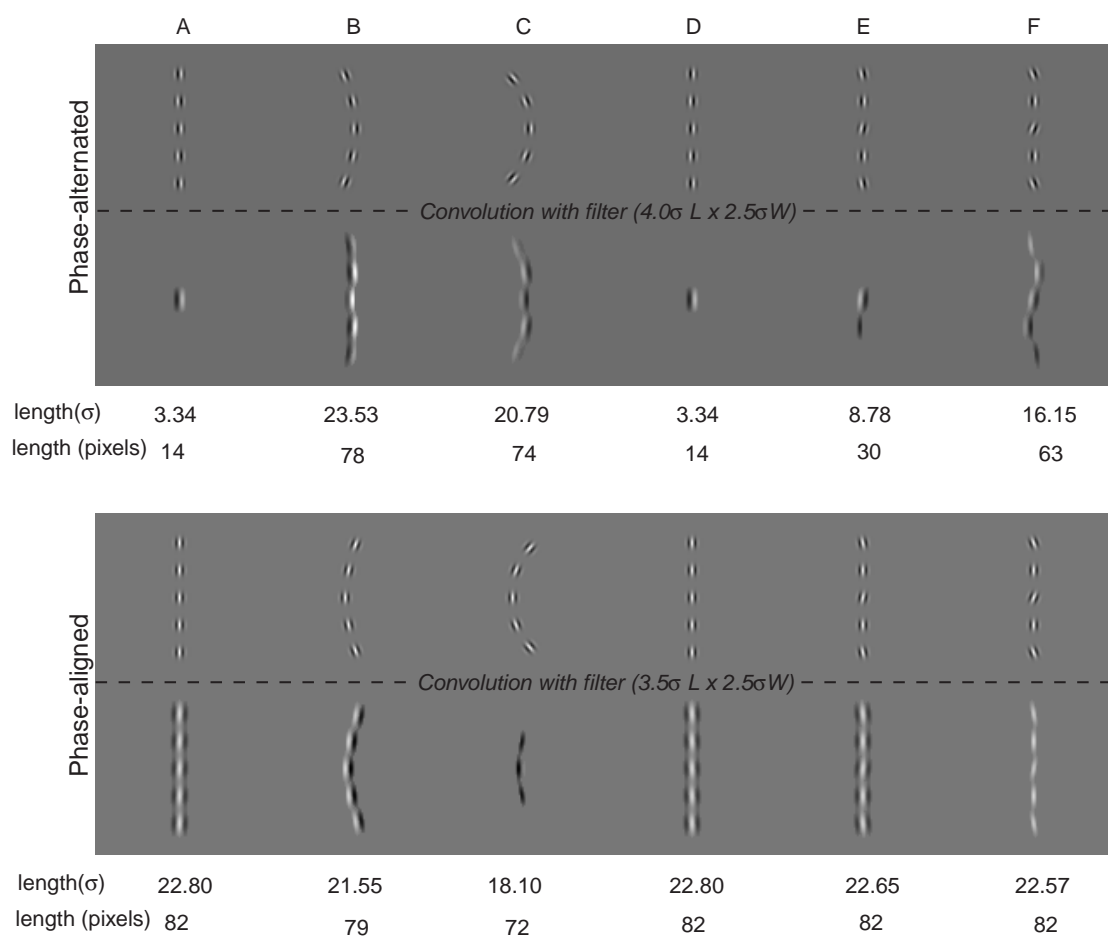


Figure 3-16. Diagram illustrating that misalignment of contour elements aids contour detection when contours are formed from phase-alternated elements. (A-C) contours with gradually increasing path-angle ($\pm 0^\circ$, $\pm 10^\circ$ & $\pm 20^\circ$). (D-F) all contours have path-angle fixed at 0° whilst the orientation jitter of elements is increased ($\pm 0^\circ$, $\pm 10^\circ$ & $\pm 20^\circ$). A comparison between the ZBR lengths of the convolved phase-alternated and phase-aligned images reveal that ZBRs are longer where there is misalignment between neighbouring elements. For both stimuli, the optimal filter indicated within section 3.3.3 is utilised. The orientation of each filter is determined by maximising the length of the ZBR whilst varying filter orientation.

Stimuli presented in Figure 3-16 omit background elements and feature uninflected contours. The presence of inflection points within contours also influences the performance of the simple-filter model (see Chapter 5). Consequently, extrapolations based upon such stimuli may be unfounded. This experiment examines the interaction between contour phase and element misalignment more closely. Stimuli featured contours with smaller increments in path-angle and element orientation jitter. Contour element locations were recorded, in order that lengths of ZBRs which corresponded to the contour could be examined. Stimuli featured background elements, consequently the performance of the model is based upon the longest ZBR length for the whole image, as it has been for previous simulations.

3.6.2 Methodology

Stimuli

Images featuring contours consisting of five elements were generated. The overall structure of contours was randomised, i.e. the sign of path-angle changes were randomised. The spacing between elements was fixed at sixteen pixels. Both path-angle and element orientation jitter were varied from 0-90° in steps of 5°. Fifty target images were constructed for each level of path-angle and orientation jitter. Contours were rotated so that their principal axis was at 90° and were located centrally within a 100² pixel image. Contours were embedded within a jittered-grid background, background element spacing was 16 pixels. Noise images were identical to target images, except that contours were omitted.

3.6.2.1 Simple-filter model

The operation of the simple-filter model was exactly as described in section 3.2.1.1. As the key stimulus parameters of element spacing and Gabor patch wavelength were the same as those used in simulation 1 (Section 3.3), only the optimal filters reported in that section were used, i.e. phase-alternated ($W2.5\sigma \times L4\sigma$) and phase-aligned ($W2.5\sigma \times L3.5\sigma$). The orientation of filters was determined individually for each image, by maximising the longest ZBR length as a function of filter orientation.

Following the convolution and thresholding operations, parallel analyses of image properties were performed : -

- i) Only those ZBRs that coincided with the location of contour elements were retained, the length of these ZBRs were recorded.
- ii) The maximum ZBR length for each individual target and noise image was retained, this value informed the 2AFC decision.

3.6.2.2 Simulation results

The average length of the longest ZBR for each target-contour was calculated for each of the combinations of path-angle and orientation jitter. Figure 3-17 presents this value as a function of path-angle and element orientation jitter for phase-aligned and phase-alternated contours. As illustrated in the Figure 3-16 there is an inverse correlation between ZBR length and both path-angle and orientation jitter in phase-aligned contours, whilst for phase-alternated contours ZBRs are longest when the level of path-angle is between 20-30° and the element orientation jitter is between 5-20°.




Figure 3-17. Contour plots illustrating mean ZBR length as a function of contour path-angle and element orientation jitter. (left) For phase-aligned contours, there is a peak in ZBR lengths at a path-angle of 20° and an orientation jitter of 10° . (right) for phase-alternated contours ZBR length is maximised when path-angle and orientation jitter are minimised.

The target contour detection rate for the simple-filter model is illustrated Figure 3-18. As anticipated, performance for phase-aligned stimuli fell as a function of path-angle and element orientation jitter. Whereas, for phase-alternated stimuli performance peaked where path-angle was approximately 30° and orientation jitter was between $0-5^\circ$.




Figure 3-18. Simple-filter model performance in detecting embedded contours as a function of contour path-angle and contour element orientation jitter.

3.6.3 Discussion

The simulation results confirm that for phase-alternated contours a small amount of misalignment between neighbouring contour elements leads to longer ZBRs. Whilst this result confirms that a degree of inter-element misalignment aids performance in phase-alternated stimuli, it does not explain why this should be the case.

In order for the simple-filter model to create extended ZBRs, the filtering operation must link together the same-sign regions of consecutive contour elements. Convolution of a

stimulus image with an elongated filter results in a filtered image with elongated positive and negative regions. Following the thresholding operation, these regions become ZBRs. For phase-aligned stimuli ZBRs are most likely to correspond to the location of multiple Gabor elements when they are collinear. However, when consecutive elements are phase-alternated, the same-sign regions are only linked when they are slightly misaligned. When phase-alternated elements lie upon a straight line only opposite-sign regions are co-aligned, this results in shorter ZBRs.

Examination of the influence of orientation jitter and path-angle upon ZBR length for phase-aligned stimuli reveals a pattern of results that is divergent from the performance of human observers. Field, Hayes and Hess (1993, expt. iii) found that element orientation jitter had a strong negative influence upon contour detection rates. When element orientations were randomised by $\pm 30^\circ$ relative to the orientation of the path, the contour was barely detectable. For human observers then, the contour integration process is more tolerant to path-angle manipulations than it is to manipulations of path-relative orientation. The reverse appears to be the case for the simple filter model. ZBR lengths are reduced at a faster rate for deviations in path-angle than they are for deviations in path-relative orientation.

3.7 General Discussion

3.7.1 Performance of the simple-filter model with manipulations of contour path-angle

For phase-aligned contours, the performance of the model, with specific filters, is very similar to that found in psychophysical experiments. Target images were most frequently detected when they featured contours that have a low level path-angle ($< 40^\circ$). The filters that lead to comparable performance have a biologically plausible ratio of elongation. Phase-alternated contours are also detected reliably. However, performance falls slightly when path-angle is lower than 20° . This occurs because the same-sign regions of phase-alternated Gabor patches are not combined during convolution if the Gabor patches are perfectly co-aligned. The filter that led to the closest match in contour detection rates for phase-alternated stimuli was narrower than the physiological estimates suggested. This filter had a closer match than its more plausible neighbours because these filters were poorer at detecting straighter contours. Thus, the plausibility of the simple-filter model hinges upon whether it is possible to accommodate the model's impaired detection of straight, phase-alternated, contours.

In the current implementation of the simple filter model, the stimulus image is convolved with a filter of a particular orientation. No interactions occur between the outputs of differently oriented filters. This restriction partly explains the reduced detection of phase-alternated contours when path-angle is minimised. It is likely that the interactions between filters would not be so precisely delimited in a physiological implementation of the simple-filter model. If a degree of noise were added to the wiring of filter outputs, this might be sufficient to counter the failure of the current model to detect the straightest phase-alternated contours.

Hess and Dakin (1997, 1999) revise the stimuli employed by Field *et al.* and demonstrate that the *single-filter* model can perform almost as well as a human subject when contours feature elements that are phase-aligned. They found that the *simple-filter* model fails when the phase of successive Gabor elements is alternated. The second assertion is clearly not supported by the simulation results reported herein. The current implementation of the *simple-filter* model is also able to detect phase-alternated contours, at a rate similar to that of human subjects. This raises the question of why the results of the current modelling experiments differ from those of Hess and Dakin. This discrepancy might occur because (i) the density of elements in Hess and Dakin's experiment was slightly higher and (ii) their embedded contours were 2/3 shorter than those featured in the equivalent experiment. However, manipulating element density and contour length should not selectively impair the detection of phase-alternated contours. The criterion for the selection of filter-scale appears to be the essential difference.

Evidence presented in section 3.4 suggests that the filter scale used by Hess and Dakin (1999) may not have been suitable for the detection of phase-alternated targets. Equivalent filters in the current experiment produced moderate detection rates for phase-aligned stimuli and poor detection rates for phase-alternated stimuli. The results of the current path-angle simulations indicate that the optimal filter for detecting phase-aligned contours is not necessarily the most suitable for detection of phase-alternated contours. Contours formed from phase-alternated elements tended to be detected more successfully when narrower filters are employed.

3.7.2 Performance of the simple-filter model with manipulations of contour element spacing

The simple filter model's performance matched that of human subjects on identical stimuli. Indeed with a subset of filters (centred upon $W1.5\sigma \times L1.5\sigma$) performance exceeded that of

human observers. The filter that had the closest match in performance to that of human observers had an elongation ratio that coincided exactly with the ratio suggested by physiological evidence.

3.7.3 Type-1 error in estimates of simple-filter model performance?

One potential criticism of the current methodology is that even if there were no consistent relationship between filter-scale, ZBR length and the presence of a contour, so many combinations of filter-scale and orientation were tested then some of these may perform at better-than chance levels. However, the performance of the simple-filter model co-varied in a predictable manner with stimulus level, i.e. path-angle and element spacing. The distribution of successfully performing filters was not random, filters that have a similar spatial-scale generally had similar levels of performance as a function of the stimulus level. Contour maps, representing the performance of the simple filter model as a function of the spatial-scale of the filter, have extended regions where performance was either better or worse than chance.

3.7.4 Is integration occurring within the simple filter model?

In the thesis introduction, contour integration is defined as a process that enables the recognition of a contour, although the contour is actually formed from discrete line elements. The results presented in Figure 3-9 and Figure 3-15 suggest that the ZBRs selected by the *simple-filter* model may reflect the presence of many contour elements. For phase-aligned stimuli, as many as ten elements are represented by the selected ZBR. Whereas for phase-alternated stimuli, approximately five elements may have contributed towards the formation of the longest ZBR. These results seem to support the claim that integration is indeed occurring, albeit partially. The most distant elements upon a contour are unlikely to be represented by the same ZBR, because they are least likely to share the same orientation.

Section 1.4.5 suggested that the path-paradigm task does not directly test integration of whole contours in humans. The task only requires that the observer identify which image contains the largest embedded contour. The observer may indeed have integrated all of the contour elements. Nevertheless, the task itself only requires the identification of the image most likely to contain the longest contour. It follows then, that whilst the simple-filter model may be unable to integrate all contour elements, it may still provide an adequate account of observer performance in the path-paradigm task.

3.7.5 A note on the coarse sampling of Gabor patches in the current simulations

Stimuli employed in simulations described within this chapter featured Gabor patches with a wavelength of 4 pixels. This is close to that utilised in Hess and Dakin's (1997, 1999) simulations, however there is a danger that the simple-filter model is favoured by coarse-sampling of Gabor patches. In order to verify that this was not the case, the key simulations (1, 2 and 4) were repeated with higher resolution images. Appendix A reports simulations in which the scale of the stimulus images was increased by a factor of five, giving a Gabor wavelength of 20 pixels. The results of these additional simulations were largely consistent with those reported in the current chapter and those reported in Chapter 5 (simulations 4 and 5) and Chapter 7 (simulations 6). There was a slight effect of image-scale, inasmuch that the specific filter-scales that gave rise to the best performance in the lower resolutions did not scale-up exactly with the higher-resolution images. However, the overall finding that manipulations of Gabor patch phase do not control-for contributions of coarse-scale mechanisms was supported.

3.7.6 On filter selection mechanisms

The results reported in Simulations 1-6 are based upon the performance of the model featuring filter-scales which specially suited the stimuli presented. The selection of these filter-scales was based upon a post-hoc examination of the performance of the model. Consequently, the strength of the simple-filter model is questionable until a mechanism is proposed that can automatically select the appropriate filter-scale. At least two filter scales are required in order for the simple-filter model to account for the ability of human observers to detect both phase-aligned and phase-alternated contours. Section 3.6 explains that narrower filters are required to detect phase-alternated contours, these filters were beneficial because they enable the blurring-together of the same-sign regions of neighbouring Gabor patches.

A key property of phase-alternated path-paradigm stimuli is that coarse-scale mechanisms, such as the simple-filter model, tend to lose the orientation information related to individual Gabor patches when they attempt to recover the location of contours as a whole. As a consequence, extended ZBRs are just as likely to be formed between background elements as they are between contour elements. This is because the opposite-sign areas of neighbouring Gabor patches tend to cancel one another out when filtered, with the result that co-alignment of patches does not increase the extent of ZBRs in filtered images. An alternative filter-based account of contour detection is explored in Appendix B. This model

deploys an edge-detecting filter prior to a rectification stage. The result of these processes is that the orientation of Gabor patches is retained, while their phase is discarded. As a consequence, further filtering with the existing SFM leads to extended ZBRs for both phase-aligned and phase-alternated stimuli without any need to vary the scale or elongation ratio of the filters deployed at any stage. As with the simple-filter model each stage of image filtering only involves filters of a single orientation, so no inter-orientation linking is required in order to detect path-paradigm stimuli.

3.7.7 Summary

For the simple-filter model, the detection of the longest ZBR within a filtered-image was sufficient to indicate the presence of an embedded contour. This finding does not prove that human observers solve the path-paradigm contour task using the same means. However, it does suggest that the path-paradigm experiments described by Field *et al.* (1993) were not sufficiently rigorous to discount coarse-scale accounts of contour integration.

As the simple-filter model does not feature any processing operations that are not available within primary vision it remains a tenable account of contour integration. The only weakness in this assertion is the finding that detection of straight phase-aligned contours is slightly impaired relative to human observer performance. It may be possible to resolve this weakness by enabling greater flexibility in the combination of filter outputs during the convolution stage. This might be considered an invocation of inter-cellular linking, however in a neural implementation of the simple-filter model, this could be a by-product of *noise* in the specification of connections between receptive field outputs.

If the path-paradigm is to be used as a methodology for investigating contour integration then an additional contour manipulation is required, one that differentiates between the operation of the association-field and coarse-scale mechanisms. Current experiments, such as the manipulation of path-angle and element spacing, are not sufficient to establish whether human performance is due to either of these models. The following chapter discusses variations on the path-paradigm methodology that may successfully differentiate between these models.

4 The Influence of Global Contour Structure upon Saliency

4.1 Overview

Section 1.4.6 described experimental evidence that suggested the global characteristics of contours could influence detection rates. Kovacs and Julesz (1993) reported that a closed contour was more easily detected than an open contour. They argued that this was evidence of a ‘synergistic process’ (pg., 7495) which favoured a closed path of contours. If their hypothesis were accepted then the association-field theory may need extension in order to account for their findings.

The current chapter describes a series of experiments that investigated the closure effect. Initially a partial replication of the Kovacs and Julesz experiment was undertaken in order to establish the validity of their claim that closed contours are detected more frequently. Later experiments investigated the hypothesis that the lengths of the arcs that form a contour may influence detection rates, rather than closure *per se*. Thus, the closure effect may be explained by the fact that closed contours generally feature longer arcs.

4.2 Experiment 3: Verification of the Kovacs and Julesz’ closure effect

4.2.1 Introduction

Prior to the implementation of newer experimental designs, it was deemed necessary to first replicate the effect of closure described by Kovacs and Julesz. They report that a closed contour was more easily detected than an open contour when local properties such as path-angle and element spacing were matched. Furthermore, they demonstrated that the spacing between neighbouring elements could be increased by a greater extent for closed contours than it could for open contours, before detection rates fell below 75% correct. The current experiment was created in order to compare the detectability of contours with varying degrees of ‘closure’. The level of closure was manipulated by varying the distance between the first and last elements of a randomly generated contour. The level of spacing between successive contour elements was also varied. This enabled an examination of the effect of spacing upon the detectability of contours with different levels of closure.

4.2.2 Methodology

4.2.2.1 Stimulus parameters

For any given contour, where the local properties of contour path-angle and element spacing are determined, there will be a maximum possible gap between the start and end

elements, this is defined in Formula 4.1. Three levels of contour closure were defined. An ‘open’ contour was defined as a contour that had start and end elements that were separated by the maximum possible gap. A semi-closed contour was defined as one in which the contour elements were separated by a distance that is 50% of the maximum possible gap. Finally, a closed contour in which start and end elements were separated by a distance that matched the spacing of neighbouring contour elements. Note that for closed contours the start-end gap was not zero, otherwise there would be an overlap in the contour elements that may have provided the observer with an additional density cue. In order to generate a contour with a particular level of closure, random contours were repeatedly generated until a particular start-end gap was achieved. Closure gaps that fell within $\pm 5\%$ of the target were accepted, except for closed-contours where only gaps that fell between the target level and $+5\%$ were allowed.

$$\cos(\text{path-angle}) \times n \times \text{spacing}$$

Formula 4-1

The actual gap achieved between start and end elements is shown in Figure 4-1 (left) below. When local properties such as path-angle and nearest-neighbour spacing are fixed, the overall eccentricity of contours will inevitably co-vary. The extent of the longest axis of each contour was calculated and the overall mean for each level of spacing and contour closure is shown in Figure 4-1 (right).

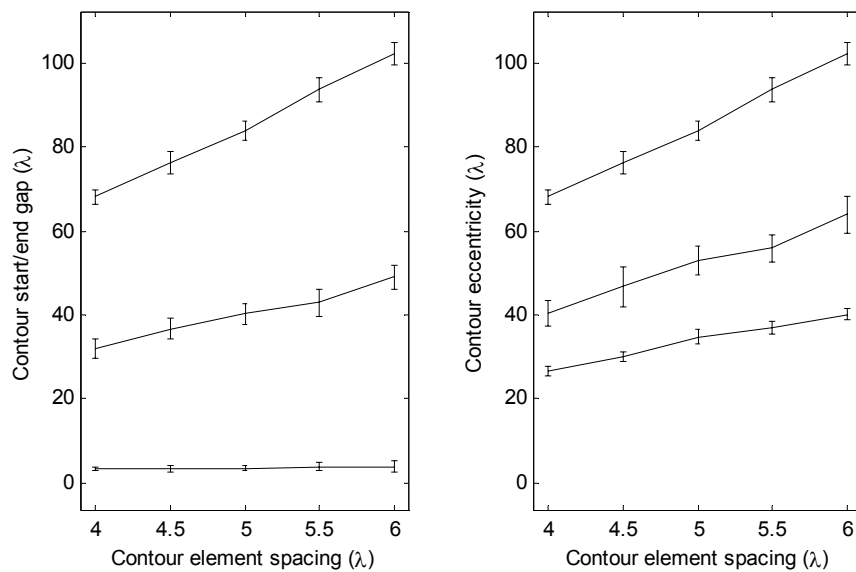


Figure 4-1. Contour start/end element gap (left), overall contour size (right). In both plots the bottom line corresponds to ‘closed’ contours, whilst the middle and top lines correspond to ‘semi-open’, and ‘open’ contours respectively. The error bars correspond to ± 1 standard deviation.

The spacing of neighbouring contour elements was varied between 4λ - 6λ in steps of 0.5λ , where λ is the wavelength of the embedded Gabor elements. Varying the relative density of contour and background elements affords the opportunity to examine performance across a range of salience levels. Previous experiments have demonstrated that this provides a useful method of manipulating contour salience (see Kovacs, Polat and Norcia, 1997). Hence, ceiling and floor effects might be avoided. Combinations of the three levels of contour closure and the five levels of contour spacing led to fifteen different sets of target images.

The contour was randomly located within the bounds of the stimulus image. Contours were embedded within a background of randomly oriented Gabor patches. The positions of background elements were determined using the jittered-grid generation technique. The mean spacing of background elements was 16 pixels, which corresponds to a spacing of 4λ . Noise images featured a randomly generated background without a deliberately embedded contour. The overall size of target and noise images was 320x320 pixels. All Gabor patches were phase-aligned. Images were generated online, therefore each target and noise image was novel.

4.2.2.2 Observer details

As the structure of embedded contours was largely prescribed within this experiment, there was some concern that observers might learn the properties of the stimuli that they were viewing. In order to minimise the possibility that such knowledge might have influenced their responses, it was decided that observers should not participate in the experiment for an extended period. Therefore, multiple observers were recruited and each of these only participated for one $\frac{1}{2}$ hour period. Observers consisted of eleven undergraduate volunteers who received course credit in exchange for participation. They were naïve to the experimental hypothesis and to the path detection paradigm in general. Each individual observed 200 target and noise image pairs, which involved approximately thirty minutes of observation. In order to enable a comparison between naïve and experienced observers, another observer (KAS) was recruited who had previously participated in path-detection experiments. KAS undertook five observation sessions featuring a total of 1000 image pairs.

4.2.2.3 Procedure

Observers were informed that they would be shown pairs of images and that they should indicate whether the first or second image in each pair contained a single continuous

contour. At this point, they were shown an example of a random (non-closed) embedded contour. A practice trial was administered which consisted of 10-15 target/noise image pairs. Further practice was available for those observers who had clearly misunderstood the instructions.

Stimuli were displayed upon a Microscan 4V ADI, 14" Colour Monitor with resolution set to 800x600 pixels. Images were viewed from a distance of 750mm. Stimulus scale properties that correspond to this distance are illustrated below in Table 4-1. Timing details were the same as those provided in Section 2.1.3.

Table 4-1. Stimulus scale properties for the current experiment.

	Pixels	Lambda	Visual-angle
Gabor patch full-wave frequency	4	1λ	0.103°
Background element nearest-neighbour spacing	16	4λ	0.412°
Contour element spacing	16	4.00λ	0.412°
	18	4.50λ	0.464°
	20	5.00λ	0.515°
	22	5.50λ	0.567°
	24	6.00λ	0.618°
Overall image size	320 x 320	$80 \times 80\lambda$	$8.3 \times 8.3^\circ$

4.2.3 Results

Whilst the overall performance levels of individual subjects varied, the relative pattern of detection rates for contours with different degrees of closure was consistent for all observers. For each observer the closed-contours were most frequently detected, followed by the semi-closed contours and finally the open-contours. As the data for the 11 naïve observers was similar, these were combined and are presented below as a single plot (Figure 4-2, top). The standard deviations of the merged detection rates are high, reflecting differing degrees of overall contour detection performance. Fitted psychometric functions illustrate the greater detectability of closed contours (Figure 4-2, bottom).

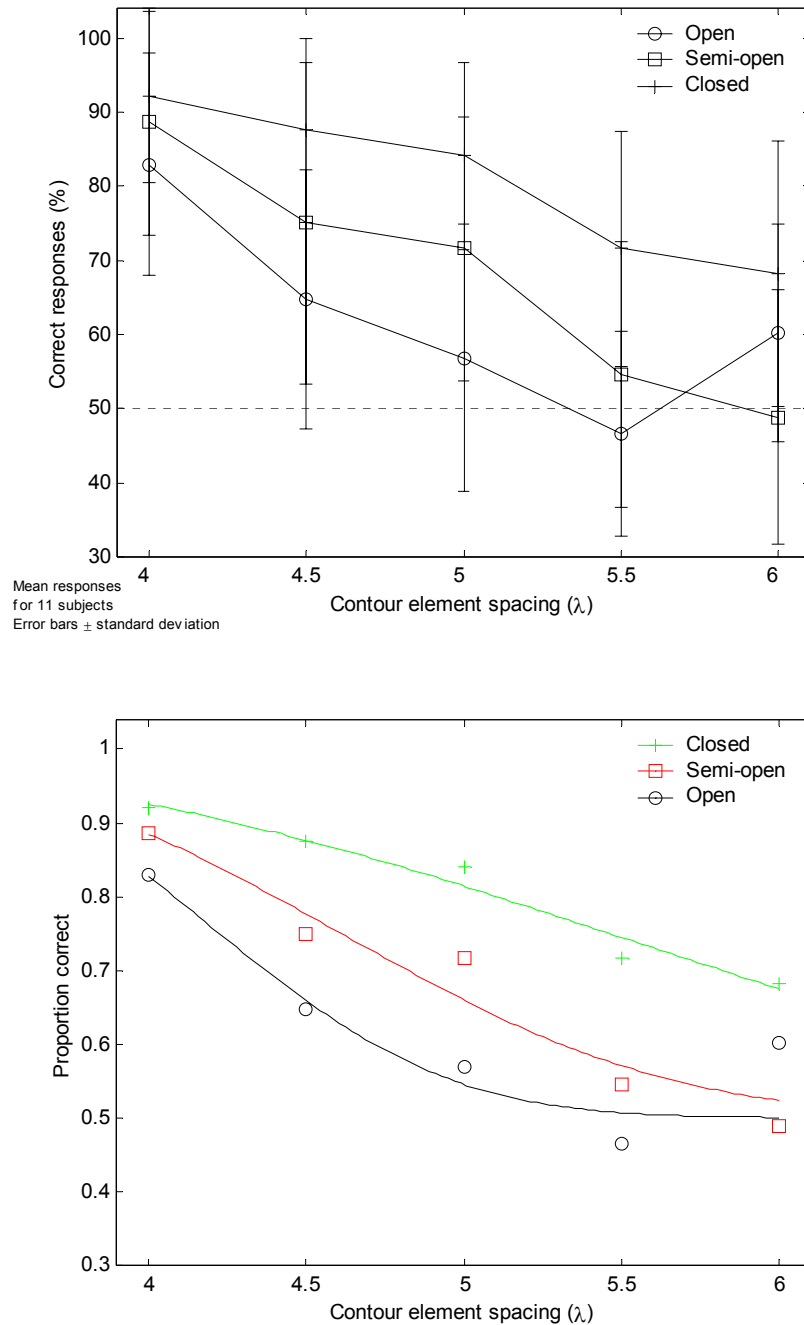


Figure 4-2. Contour detection performance as a function of contour closure and contour element spacing. The data-points reflect the mean contour detection rate for the 11 naïve observers. The lower plot shows the fitted psychometric functions for each of the three levels of contour closure. Error-bars are omitted from this plot in order to avoid clutter, however the bars presented in the upper plot provides a measure of the variance at each stimulus level.

A χ^2 test confirms that there was a significant difference in the detection rates for each of the three levels of contour closure, Table 4-2. The significant difference spans contour element spacing levels of 4-5.5 λ . When contour element spacing exceeds 5.5 λ detection rates for open and semi-open contours approach chance levels.

Table 4-2, χ^2 analysis of differences in detection rates for contour closure stimuli. Analysis of merged response data for all naïve observers (n=11).

χ^2 significance matrix			
Contour Closure	Open	Semi-open	Closed
Open	*	7.499	39.740
Semi-open	0.0062	*	13.299
Closed	< 0.0000	0.0003	*

The upper triangle gives the χ^2 values and the lower triangle the p-values. Statistically significant differences are highlighted in **bold**

The results for observer KAS are presented separately as KAS had participated in five times as many sessions as each naïve observer. Figure 4-3 shows raw results and Figure 4-4 shows the fitted psychometric function for observer KAS.

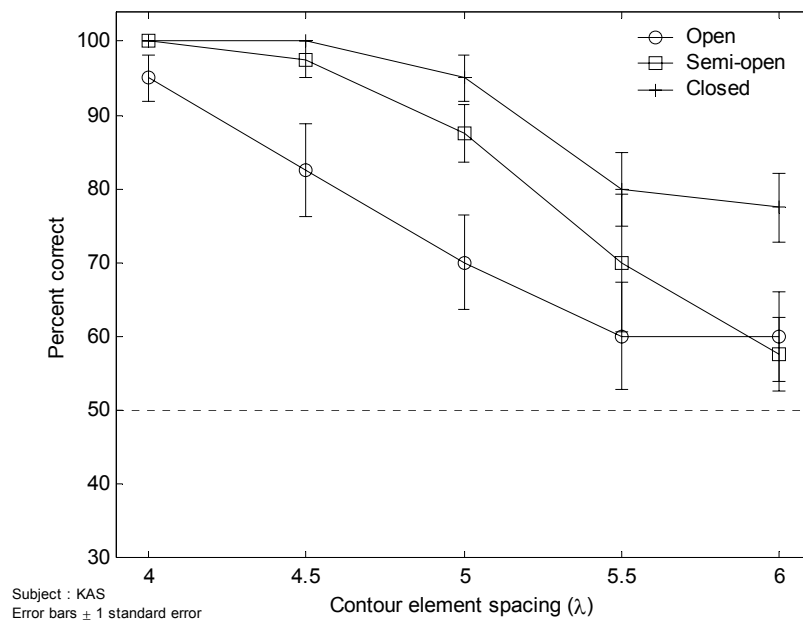


Figure 4-3, contour detection rates for observer KAS, as function of contour closure and contour element spacing.




Figure 4-4. Performance for observer KAS as a function of contour element spacing and contour type. Whilst there is evidence of a ceiling effect for closed contours, there is still a clear distinction in the detection-rates of each type of contour.

For observer KAS there was a significant difference in contour detection rates across all levels of element spacing (see Table 4-3).

Table 4-3, χ^2 analysis of differences in detection rates for contour closure stimuli with observer KAS.

===== χ^2 significance matrix =====

Contour Closure	Open	Semi-open	Closed
Open	*	4.720	19.580
Semi-open	0.030	*	5.481
Closed	< 0.000	0.019	*

=====

The upper triangle gives the χ^2 values and the lower triangle the p-values. Statistically significant differences are highlighted in bold

Whilst overall detection rates differed between the naïve observers and KAS, the relative differences in the detectability of contours with varying degrees of closure was maintained. The mean performance levels for KAS were approximately 15% higher than they were for the naïve observers. However, the influence that contour closure had upon detection rates was the same. This is illustrated in Figure 4-5, which shows the degree of element spacing

that the fitted psychometric functions predict would lead to performance levels of 75% correct.




Figure 4-5. Fitted 75% detection limits as a function of contour element spacing and contour-type. Practice appears to have increased tolerance to manipulations of the element spacing gap but has not had any influence upon the relative detectability of the three contour types.

4.2.4 Discussion

The results confirm those reported by Kovacs and Julesz (1993). It appears that there is a clear effect of contour closure. Closed contours are detectable at a level of 75% correct when the spacing of contour elements is $5.4-6\lambda$, whereas for open contours performance levels are similar for element spacings of between $4-5\lambda$. Kovacs and Julesz reported values of 6λ and 3.3λ . This slight difference in element spacings is likely to be due to the choice of background generation technique or differences in the location of individual contours within images.

In the current experiment the gaps in the contour stimuli varied from 6λ for closed contours to 45λ and 100λ for semi-open and open contours respectively. Despite these large gaps, there were significant differences in detection rates between all levels of closure. This suggests that there was a linear relationship between contour closure and detection rates. This contrasts with the qualitative difference in the processing of closed and open contours hypothesised by Kovacs and Julesz.

Kovacs and Julesz (1993) reported that salience increased suddenly with the addition of the last 1-2 elements, i.e. when an open contour was closed. They interpret this as evidence for the existence of a ‘synergistic process’, where synchronous firing patterns are facilitated by the presence of a closed-loop of co-facilitatory connections between each of the receptive fields that have detected the contour elements. This account of their results is dependent upon the fact that contours were detectable when closed and then become undetectable when 1-2 elements were removed. The additional elements that Kovacs and Julesz added to a contour would correspond to a gap of approximately $12-18\lambda$. Clearly, the semi-closed contours featured in the current experiment had gaps that significantly exceeded this amount ($30\lambda-45\lambda$). This difference suggests either that the synergistic process is no longer a tenable account, or that it must be much more tolerant to gaps than previously suggested.

As already mentioned in the introduction (Section 1.4.6), Braun (1999) has argued that the results reported by Kovacs and Julesz may have been confounded by contour length. The length of their contours was very close to the detection threshold that he identifies. Consequently, regardless of the overall shape of the contour, contour detectability would inevitably have improved with the addition of further Gabor elements. In the current experiment contour lengths were held constant, yet there was still an effect of contour closure.

Whilst the experiments conducted by Kovacs and Julesz may have been confounded by contour length the results of the current experiment may have been confounded by the overall size of contours. Contour size is inevitably correlated with global structure (Figure 4-1, right). Hence, the influence of this variable must be explored before any firm conclusions can be made. In an attempt to minimise the influence of size, an additional experiment was conducted which is only briefly described here. In this experiment, the location of embedded contours was also manipulated. Closed contours were located furthest from the centre of the stimuli whilst the open contours were positioned such that they crossed the position of the fixation point. The overall pattern of detectability remained the same as that reported within this section. A further experiment introduced additional degrees of contour closure. Again, the detectability of contours varied as a function of contour closure in the same manner as reported within this section.

The results of the current experiment appear divergent from what would be predicted by Kovacs and Julesz’ ‘synergistic process’ hypothesis. There was a linear relationship between contour closure and the detection rates of observers. This differs from the

qualitative difference in processing described by Kovacs and Julesz. Consequently, an alternative explanation of the current results is required. The following section investigates the influence of contour ‘smoothness’. When path-angle and element spacing are held constant then the frequency of inflection points within a contour will inevitably co-vary with closure. As contours are closed, they will feature fewer inflections. Contours with fewer inflections are inevitably smoother; i.e. they are composed of longer arcs. The following section describes an experiment that investigated the influence of this variable.

4.3 Experiment 4: Examining the influence of contour ‘smoothness’

4.3.1 Overview

The association field theory argues that contour integration is based upon pair-wise interactions between neighbouring receptive fields that lie upon a smooth curve. Consequently, experiments have generally concerned themselves with local variables, such as element spacing, path-angle, Gabor patch phase and spatial frequency (i.e. Field, Hayes and Hess, 1993; Hess and Field, 1995; Field, Hayes and Hess, 1997; Hess, Dakin and Field, 1998; Dakin and Hess 1999). Those studies that have examined the influence of global contour structure, for example closure (Kovacs and Julesz, 1993; Current thesis, Experiment 3), must hold local properties constant. Otherwise, differences in contour detectability may be confounded by differences in local properties such as path-angle or element spacing. However, if closure is manipulated whilst path-angle, contour length and element spacing are held constant then an additional difference exists between open and closed contours, this is termed *smoothness*. A closed contour is likely to feature fewer inflections, thus the contour as a whole will feature longer, but fewer, continuous arcs. In contrast, an open contour will feature more inflections, thus the contour as a whole will feature many, but shorter, arcs. The ‘arc-length’ of contours, i.e. the length of arcs that form contours, is directly correlated with a contours *smoothness*. Throughout this chapter the terms *arc-length* and *smoothness* are used interchangeably⁷.

The relationship between contour-closure and contour-smoothness described above occurs where the contour parameters of path-angle and element spacing are held constant. When such a contour is closed, the sign of path-angle changes between consecutive contour elements is likely to remain constant. Whereas for open contours the sign of path-angle changes is more likely to switch with each consecutive element. Figure 4-6 illustrates the correlation between contour closure and the length of component arcs within a contour.

⁷ Both terms are defined in the glossary.

Kovacs and Julesz tried to avoid this confound by basing all their contours upon closed contours, open contours were then created by deleting elements, thereby inadvertently confounding their results with contour length (Braun, 1999).




Figure 4-6. Where local contour properties are held constant open-ended contours inevitably have a greater number of path-angle sign changes than closed contours. The lengths indicated correspond to the number of elements encountered until a change in path-angle, this will generally underestimate the actual arc-length by 1 element.

The hypothesis that contour smoothness, rather than closure, is influencing contour detectability implies a process of contour integration which favours continuous arcs. The influence that the smoothness of contours has upon contour detectability was investigated by manipulating the frequency of inflection points within embedded contours.

4.3.2 Methodology

4.3.2.1 Stimuli

Contours were generated with the technique described in Section 2.1.1.1. However, an additional constraint was introduced, the length of arcs within contours were pre-determined. This was achieved by specifying the sign of consecutive changes in path-angle. For an arc-length of one, the changes in path-angle signs were “+ - + - + -...”, whilst for an arc-length of 4, the path-angle signs were “+ + + + - - - - + + + +...”. Examples of the five contour types employed in this experiment are shown below (Figure 4-7). The number of elements that formed each contour was fixed at 16 elements except for the closed-contour that featured only 15 elements. Fewer elements were included in the closed-contour so that the number of pair-wise interactions in each contour was held constant. The path-angle between subsequent contour elements was 24°. Therefore, with an arc-length of 15 elements a closed contour was achieved.




Figure 4-7. Manipulating the length of component arcs within a contour. Units represent distance spanned along the contour before an inflection point is encountered.

In order to manipulate contour detectability levels, contour element spacing was varied between 4λ and 6λ in steps of 0.2λ . The contour was randomly located within the bounds of the stimulus image. The positions of background elements was determined using the jittered-grid generation technique (Appendix A). The mean spacing of background elements was 16 pixels (4λ). The overall size of target and noise images was 320x320 pixels. Noise images consisted of random backgrounds that did not feature a deliberately embedded contour. All Gabor elements were phase-aligned.

4.3.2.2 *Observer details*

Observers consisted of seven undergraduate volunteers who received course credit in exchange for participation. They were naïve to the experimental hypothesis and to the path detection paradigm in general. Each individual observed 200 target/noise image pairs which corresponds to approximately thirty minutes of observation.

4.3.2.3 *Procedure*

Subjects were informed that they were to view pairs of images and that they should indicate which image contained a single continuous contour. They were then shown an example of a random embedded contour. A practice trial was administered which consisted of 10-15 target/noise image pairs. Further practice was available if it became apparent that a subject had misunderstood the instructions.

Target/noise image pairs were selected randomly (without replacement) from all available stimulus levels following the procedure described in Section 2.1.2. Each image was visible for 250ms following presentation of a fixation point. Full details of presentation timing have been given in Section 2.1.3. Stimuli were displayed upon a Microscan 4V ADI, 14" Colour Monitor with resolution set to 800x600 pixels. Stimulus scale properties that correspond to the viewing distance of 75cm are given below in Table 4-4.

Table 4-4. Stimulus scale properties for the current experiment.

	Pixels	Lambda	Visual-angle
Gabor patch full-wave frequency	4	1λ	0.103°
Background element mean nearest-neighbour spacing	16	4λ	0.412°
Contour element spacing	16	4.00λ	0.412°
	18.67	4.66λ	0.481°
	21.33	5.33λ	0.550°
	24	6.00λ	0.619°
Overall image size	320×320	$80 \times 80\lambda$	$8.3 \times 8.3^\circ$

4.3.3 Results

Whilst the overall performance levels of individual subjects varied, the pattern of detection rates for contours with different arc-lengths was consistent for all subjects. For contours with arc-lengths of 2-15 elements, there was generally a monotonic increase in detection rates with increased arc-length. Detection rates are less consistent for contours with an arc-length of one element. For all contours detection rates fell as a function of element spacing. The least detectable contours were those featuring an arc-length of two elements. As individual subject performance patterns were largely homogenous, a single plot is provided below illustrating the mean performance levels as a function of contour arc-length and contour element spacing (Figure 4-9).




Figure 4-8. Contour detectability as a function of contour element spacing and arc-length. Data represents the means for the performance of seven observers.




Figure 4-9. Contour detectability as a function of arc-length and contour element spacing. Data represents the means for the performance of seven observers.

When contour detection rates are presented as a function of arc-length, the pattern of results is non-monotonic; i.e. there are peaks at both the smallest and largest arc-length. However, contour detection rates fall predictably as a function of element spacing. Figure 4-10 shows contour detection rates as a fitted psychometric function of contour element spacing for each level of arc-length. A visual inspection of these plots confirms that the correspondence between the raw detection rates and the fitted psychometric function is satisfactory for arc-lengths of 4, 8 and 15 elements. However, for arc-lengths of 1 and 2 elements the fit is quite poor, possible explanations for these poor fits are considered within the discussion.




Figure 4-10. Fitted psychometric functions of contour arc-length and element spacing. In order to avoid clutter, error-bars are not shown on this plot. Error-bars shown in Figure 4-8 are appropriate to this plot.

A useful summary of the influence of contour arc-length upon detection rates is achieved when the contour element separation that is necessary to achieve a 75% detection rate is calculated for each level of arc-length (Figure 4-11). These estimates are based upon the fitted psychometric functions presented in Figure 4-10.




Figure 4-11. 75% detection limits as a function of contour element spacing and contour-arc length.

Table 4-5, below, examines the statistical significance of the differences in the detection rates of contours with varying arc-lengths. The χ^2 analysis reveals that detection rates were significant different between contours of all arc-lengths except for 1 Vs 4 ($\chi^2 = 2.666$, $p = 0.103$) and 1 Vs 8 ($\chi^2 = 2.462$, $p = 0.117$). These differences fail to achieve statistical significance because the function for an arc-length of one element intersects the functions for arc-lengths of 4 and 8 elements.

Table 4-5, χ^2 analysis of differences in detection rates for manipulations of arc-length.

Arc-length	1	2	4	8	15
1	*	19.715	2.666	2.462	28.087
2	0.000	*	8.038	35.319	86.682
4	0.103	0.005	*	10.150	46.079
8	0.117	0.000	0.001	*	14.694
15	0.000	0.000	0.000	0.000	*

The upper triangle gives the χ^2 values and the lower triangle the p-values. Statistically significant differences are highlighted in **bold**.

4.3.4 Discussion

The results appear to support the argument that contours become more detectable as their smoothness is increased. Unfortunately, overall size inevitably co-varies with arc-length. Figure 4-12 (below) illustrates how the size of a contour and detection rates were inversely related. For arc-lengths greater than two elements, there is a strong inverse correlation between detection-rates and contour size. When examined in isolation the results of the current experiment cannot be considered conclusive, as the influence of arc-length may have been confounded by contour size. In an experiment that featured similar stimuli to those reported here, Pettet (1998) attempted to minimise the effect of contour size by rotating each contour so that its principle axis was vertical. Pettet also reports a positive correlation between arc-length and contour detectability, even when the influence of size was minimised.




Figure 4-12, bar graph illustrating the extent of individual contours in visual angle units. The size of the contours with arc lengths of 1, 2 and 4 do not differ to a great extent, whilst for the ‘s’ and circle stimuli the size varies much more. The overlaid line plot shows the 75% detection limits as a function of arc-length and element spacing. Clearly, for arc-lengths of 2 elements and above there appears to be an inverse correlation between contour size and detection rates.

The negative correlation between contour size and the detection rates of observers breaks down for contours with an arc-length of one element. These contours are actually detected more often than those with an arc-length of two elements despite having a slightly larger eccentricity. A possible explanation of this discrepancy is that the detection of smooth contours and the detection of the jagged⁸ contours (arc-length < 2) may be achieved via different mechanisms. Apart from overall size and smoothness, another variable differentiates smooth and jagged contours. For jagged contours, non-neighbouring elements may still share a similar orientation, whilst for smoother contours, the orientations of elements become increasingly dissimilar as they are separated by a greater number of contour elements. Thus, detection of the jagged contour might be achieved by a coarse-scale mechanism such as that described in Chapter 3, whilst detection of smoother (curved) contours would be less likely with such a mechanism. Therefore it may be that in order to detect smooth contours it is necessary to utilise an association field mechanism, whilst for jagged contours a simple-filter mechanism might be sufficient to indicate the presence of a contour. This possibility will be explored in Chapter 5.

⁸ See the glossary for a definition of this term.

During debriefing the majority of observers reported that they had only seen contours that were ‘C’ shaped or circular. None of the observers reported viewing ‘S’ shaped stimuli and often expressed surprise when they were told that such stimuli had been included. A possible explanation for the experience of observers is that the integration of contour elements occurs within component arcs and that subsequently the arcs themselves become integrated with neighbouring arcs. This might be achieved in two ways; firstly, integration might occur for all elements but the influence of global context may speed integration within arcs. Secondly, the integration of individual contour elements might only occur within arcs and then the resultant arcs are further integrated with one another at a later stage of processing. Both possibilities might account for the reported experience of observers. As images were displayed with a fairly short duration (250ms) then the integration of elements at the point of inflection might not have occurred. A third possibility is that we have ‘grandmother’ cells that enable the detection of arcs, though this explanation may fall foul of the usual arguments of computational inefficiency.

4.4 Experiment 5: Is integration impaired at inflection points?

4.4.1 Introduction

The previous section offered the smoothness hypothesis as an alternative explanation of the closure-effect described by Kovacs and Julesz (1993). It was argued that the frequency of inflections within a contour could be inversely related to that contour’s detectability. As a consequence of the need to hold constant local properties of contours, overall size tends to co-vary with smoothness. Therefore, it is difficult to construct stimuli that provide a reasonable test of the smoothness hypothesis without confounding the results on a local level (path-angle, spacing etc) or on a global level (size). An alternative approach was adopted; this involved examining some specific implications of the smoothness hypothesis.

The smoothness hypothesis implies a mechanism that integrates contour elements more strongly for members of the same arc than it does for elements that are members of different arcs. Neighbouring elements within an arc might be integrated sooner, or more strongly, than elements that are neighbours but which lie upon different sides of an inflection point. If this argument were correct then there would be no difference in the detectability of an ‘S’ shaped contour and two proximate ‘C’ shaped contours. Alternatively, if the ‘S’ shaped contour were detected more frequently, then this would imply that that integration had occurred at the point of inflection. Thus, the smoothness

hypothesis could be rejected. With such stimuli, variables such as path-angle, overall size and density could be held constant between contours.

An assumption underlying this experiment is that longer contours are detected more frequently. A great deal of experimental evidence supports this claim (Beck *et al.* 1990; Uttal, 1975; Smit *et al.*, 1985; Moulden, 1994; Braun, 1999). Therefore, a target image featuring a contour composed of sixteen elements should be detected more frequently than an image featuring two contours that are each eight elements long. If the sixteen element 'S' contour were no more detectable than two eight element 'C' contours then this implies that no integration occurred at the inflection point of the 'S' contour.

An experiment was designed to assess the relative detectability of such contours. Stimuli featured four types of contour, these are described below: -

- (i) An 'S' shaped contour composed from sixteen Gabor elements.
- (ii) A 'C C' stimuli in which the two eight element C's were randomly positioned within the stimulus image. A check was made, during the embedding of the contours, that no overlap occurred between the 'C's.
- (iii) A single 'C' contour, composed of eight elements, was also embedded in order to establish the detectability of this contour in isolation.
- (iv) A 'C-C' 'contour' which was identical to the 'S' stimuli, except that one half of the contour was rotated by 45°.

Typical stimuli are shown within Figure 4-13, below. Following the reasoning outlined above, there should have been little difference in the detectability of the 'S' and 'C-C' stimuli. Furthermore, if there were some benefit for the 'C' contours being proximate then the 'C-C' stimuli might be more easily detected than the 'C C' stimuli. Single 'C' stimuli were included in order to establish the level of detectability of each individual 'C' segment. This enabled a post-hoc assessment of a probability summation account for the detectability of the other contours (i, ii & iv).

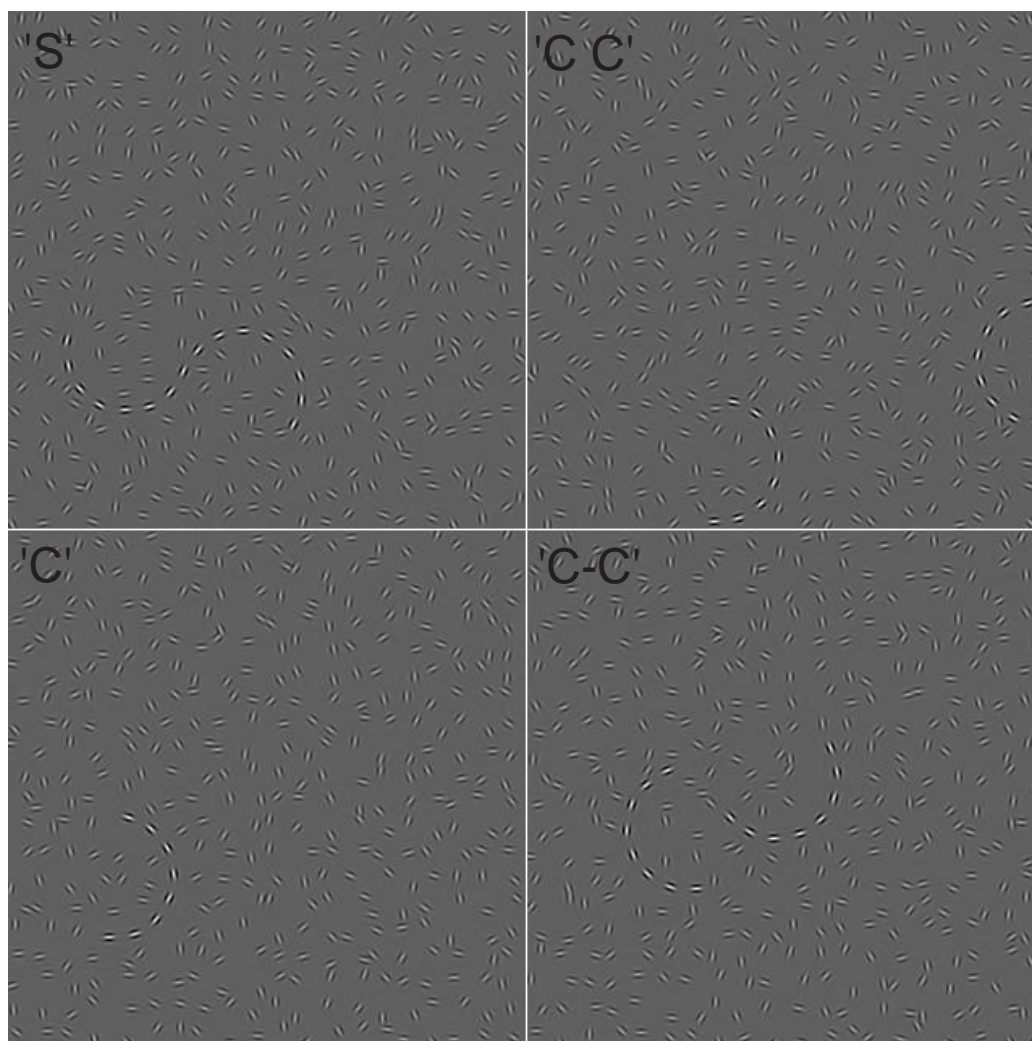


Figure 4-13. Examples of stimuli utilised in the current experiment. For illustrative purposes the contrast of background Gabor elements has been reduced in order that the embedded contours are more easily recognised.

The stimuli employed in the current experiment differ from those utilised by Petted *et al.* (1996); their stimuli were designed to investigate ‘kinks’, i.e. they deliberately manipulated path-angle. The current hypothesis claims that inflection points *per se* may influence salience even when path-angle is held constant.

4.4.2 Methodology

4.4.2.1 Stimuli

Contours were generated with the previously described technique discussed in Section 4.3.2.1. Each contour⁹ was composed of arcs that were eight elements long. Four types of target stimuli were generated, each featuring either one or two arcs. (1) ‘S’ shaped contours composed of two, 8 element long, arcs that were contiguous with one another. (2) Other

⁹ Note that for terminological convenience even contours that are in fact broken to form two smaller contours will be referred to as single contours.

stimuli featured a pair of embedded ‘C’ shaped arcs, which were randomly located relative to one-another – with the single constraint that they did not overlap. This is termed the ‘C C’ contour. (3) A ‘C’ shaped arc composed of eight elements was included in order to ascertain the detectability of a single arc. (4) A further contour featured two ‘C’ shaped arcs that were proximate to one another. This contour was created by generating an ‘S’ shaped contour and then rotating one of the arcs by 45°. Again, there was an additional constraint that the two ‘C’ shaped arcs were not allowed to overlap. This contour is termed the ‘C-C’ contour – the hyphen represents the fixed proximity of the two arcs.

For each of the contours (1-4) the path-angle between consecutive elements was $\pm 24^\circ$. The arcs were all composed of eight Gabor elements thus the ‘S’, ‘C C’ and ‘C-C’ contours were composed of sixteen elements, whilst the ‘C’ contour was composed of only eight. An additional ‘element orientation jitter’ of $\pm 5^\circ$ was applied to each contour element. Contour element spacing was manipulated between 4λ and 8λ in steps of 1λ in order to measure observer performance across a range of contour detectability levels.

As with the previous experiment contours were embedded within a background of randomly oriented Gabor patches. The contour was randomly located within the bounds of the stimulus image. The positions of background elements were determined using the jittered-grid generation technique (Appendix A). The mean spacing of background elements was 16 pixels (4λ). The overall size of target and noise images was 320x320 pixels. Noise images were identical to target images except that they did not feature embedded contours. All Gabor patches were phase-aligned.

4.4.2.2 Observer details

As with the previous experiment, the stimuli featured were non-random. Hence, there was a possibility that observers might form an idea of the nature of the contours they were selecting. Consequently, many observers were tested for relatively short periods, as opposed to the more usual psychophysical procedure of testing a few individuals many times. Eleven undergraduate volunteers were recruited, each received course credit in exchange for participation. They were naïve to the experimental hypothesis and to the path detection paradigm in general. Each individual observed 200 target/noise image pairs which corresponded to approximately thirty minutes of observation.

4.4.2.3 Procedure

Observers were informed that they were to view pairs of images and that they should indicate whether the first or second image contained a single continuous contour.

Observers were shown an example of a random embedded contour at this point. A practice trial was administered which consisted of 10-15 target/noise image pairs. Further practice was available if it became apparent that an observer had misunderstood the instructions.

Image presentation was determined in the same manner as described in Section 2.1.2. Details of presentation timing were identical to that of the previous experiment (Section 4.2.2.3). Images were viewed from a distance of 750 mm. Stimulus scale properties that correspond to this viewing distance are provided in Table 4-6.

Table 4-6. Stimulus scale properties for the current experiment.

	Pixels	Lambda	Visual-angle
Gabor patch full-wave frequency	4	1λ	0.103°
Background element nearest-neighbour spacing	16	4λ	0.412°
Contour element spacing	16	4λ	0.412°
	20	5λ	0.515°
	24	6λ	0.619°
	28	7λ	0.722°
	32	8λ	0.825°
Overall image size	320 x 320	80 x 80λ	$8.3 \times 8.3^\circ$

4.4.3 Results

The pattern of relative detectability for each contour was similar for each of the eleven observers. Hence, only the mean performance levels of all observers are presented below. As with the previous experiment, the detection rates for individual observers varied. Consequently, the standard deviation of the mean detection rates was quite large (Figure 4-14). A cursory examination of the raw detection rates for contours indicates that there was little difference in the detectability of the 'S' and 'C-C' contours, whilst lower detection rates were recorded for the 'C C' contours and the single 'C' contours were lowest of all.




Figure 4-14. Contour detection rates as a function of contour-type and contour element spacing. The results illustrate the overall performance levels for all of the eleven observers.

Psychometric functions of contour detection rates were calculated as a function of element spacing and contour type, these are illustrated in the plot below (Figure 4-15). The shapes of the four psychometric functions reveal that the detection rates for the 'S' and 'C-C' stimuli are almost indistinguishable, whilst the rates for the single 'C' stimuli and the non-proximate 'C C' stimuli are reduced.




Figure 4-15. Contour detection rates as a fitted psychometric function of contour element spacing and contour type. Error bars are omitted in order to ease interpretation, the error bars in Figure 4-14 represent the variance about the points shown.

A χ^2 test was conducted to test the statistical significance of these differences, the results are presented within Table 4-7. Except for 'C-C' and 'S' contours, there was a significant difference in the detection rates for all types of contours, when contour elements are separated by 4-6 λ . For element separations of greater than 6 λ detection rates are not consistent. This suggests that observers were poor at detecting contours when element spacing was greater than 6 λ , regardless of the type of contour.

Table 4-7, χ^2 analysis of differences in detection rates for the stimuli employed in the current experiment. Analysis examines observer responses for element spacing sizes of 4-6 λ .

===== χ^2 significance matrix =====

Contour type	S	C C	C	C-C
S	*	4.355	17.765	0.000
C C	0.037	*	4.662	4.355
C	< 0.000	0.031	*	17.765
C-C	1.000	0.037	< 0.000	*

The upper triangle gives the χ^2 values and the lower triangle the p-values. Statistically significant differences are highlighted in **bold**.

The level of contour element spacing when the psychometric fits for detection rates were at 75% correct are shown in Figure 4-16. Clearly, the detection rates of the 'S' and 'C-C' stimuli are the same whilst for the 'C C' and 'C' stimuli, detectability falls off more quickly as a function of increasing element spacing.




Figure 4-16. 75% detection limits as a function of contour element spacing and contour-type.

4.4.4 Can probability summation account for the detectability of multiple ‘C’ contours?

An examination of the relative detectability of the non-proximate arc stimuli ‘C C’ and the individual ‘C’ arc stimuli suggests that probability summation alone accounts for the observer detection rates for the ‘C C’ stimuli. This hypothesis was tested by generating a predicted psychometric function of the detectability of the ‘C C’ stimuli based upon the measured detection rates of the single embedded arc stimuli (‘C’). The probability that a single ‘C’ shaped arc would be detected was calculated (Equation 4-1) and then the predicted detectability of two embedded arcs was calculated (CC_{fit} , Equation 4-2). ‘C’ represents the fitted psychometric function for detection of the C shaped contour. Figure 4-17 shows the psychometric functions of the ‘C’ and ‘C C’ stimuli and the predicted detectability of the ‘C C’ stimuli (CC_{fit}).

$$p(C) = 1 - \left(\frac{1 - C}{0.5} \right)$$

Equation 4-1

$$CC_{fit} = \frac{2p(C) - (p(C))^2}{2} + 0.5$$

Equation 4-2




Figure 4-17. The detection rates for the 'C C' contours are very similar to the summed probability of detecting two embedded 'C' stimuli. This differs from the psychometric function for 'C-C' stimuli where higher detection rates reflect the increased proximity of the arcs.

Clearly, there is a close correspondence between the fitted psychometric function of the 'C C' stimuli and the 'C C_{fit}' estimate. This indicates that the increased detectability of the 'C C' stimuli over the 'C' stimuli is likely to be due to probability summation. Furthermore, it is reasonable to infer that the higher detection rate of the 'C-C' stimuli contours cannot be explained by probability summation alone.

4.4.5 Discussion

The current experiment has not recorded any increase in the detectability of the 'S' contour over that of the 'C-C' contour(s). This suggests that there has been no advantage conferred by the increased length of the embedded 'S' contour over the proximate arcs of the 'C C' stimuli. This might lead one to conclude that the 'S' and 'C-C' stimuli were equivalent to the observers. This suggests that integration at the point of inflection is either weaker or has not yet occurred. Otherwise, one would predict an increase in detectability as contour length increases from eight to sixteen elements within the 'S' stimuli. However, no firm conclusions should be drawn from a failure to find a significant difference.

For Field et al's association field model, contour detectability is largely a function of the number of elements that form the contour and of the distance and path-angle between these elements. If the contour was broken into two halves through a manipulation of these parameters then detectability ought to be based upon the summation of the probability of

detecting each of the contour halves. The current experimental results suggest that this is not the case. When a contour was broken into two, at the point of inflection between two arcs, the detectability of the resultant target stimulus ('C-C') was not predicted by the combined detectability of each of the individual contours. The detectability of the 'S' and 'C-C' contours is clearly not a simple function of the summed detectability of each of the arcs that form those stimuli. Conversely, when the embedded arcs are not proximate then detectability is indeed predicted by probability summation. Thus, contour detectability appears to increase if the two arcs are located nearby one another.

Another possible explanation for the failure to find an 'S' / 'C-C' detection difference might be that detection rates as a function of contour length had already reached asymptote. Thus, increasing the contour beyond the length of the 8-element arc did not lead to any significant increase in detection rates. This conflicts with the performance of observers in Kovaks and Julesz' (1993) 'jagged-contour' condition (pg. 7496, Figure. 4); they report a gradual increase in performance as contour lengths are increased from 5 to 12 elements. However, this difference may simply reflect differing element densities or viewing conditions. Section 6.3 of the current thesis describes experiment 7, featuring similar stimuli, which reports observer performance approaching a ceiling level for contours composed of only eight elements, though these contours had a larger path-angle value.

Whilst increasing contour length across the inflection point did not increase detection rates, reducing the spacing between individual elements did improve detectability. This suggests that there was no interaction between the influence of these two contour properties, length and spacing. This raises the question, why was the performance in the 'S' and 'C-C' stimuli conditions higher than it was for the non-proximate 'C C' stimuli? One possible explanation is that a representation is formed that is spatially weighted, i.e. where two arcs are nearby one another they somehow become more salient to the observer. Alternatively, observers may be less confident when making 2AFC decisions if they perceive two contours, within the same stimuli, which are spatially removed.

The current experiment has shown that proximate arcs are more readily detected than non-proximate arcs. Probability summation provides a sufficient account of the detectability of two individual arcs. However, it is not a sufficient explanation of the detectability of the 'C-C' and the 'S' stimuli, relative to the 'C' stimuli. The experiment failed to find any difference in the detectability of 'C-C' and 'S' stimuli, despite the fact that the latter

stimulus features a longer contour. This might lead to a conclusion that integration does not occur across contour inflection points. The proximity of the two arcs was sufficient to explain detection rates. However, conclusions based upon a failure to find a significant difference in the detection rates for two similar stimuli are likely to be equivocal. The experiment introduced in the following section was created in order to counter this criticism.

4.5 Experiment 6: Comparing the influence of contour breaks at inflection points and elsewhere

4.5.1 Introduction

The results of the previous experiment suggested that, when local properties were held constant, contour detectability might be initially determined by the properties of continuous arcs. That is, inflections may essentially ‘break’ a contour, giving rise to what will be termed a *virtual* break. One explanation of this result is that the strength of ‘integration’ may be weaker where the sign of the path-angle changes between subsequent contour elements. Before the hypothesised influence of inflections can be accepted a number of alternative explanations must first be ruled out: -

- i) *Contour-length*: The arcs featured in the previous experiment consisted of eight Gabor elements and detection rates as a function of contour length may already have achieved asymptote. Consequently, there may have been no benefit for adding additional elements to a contour by conjoining the arcs.
- ii) *Coarse-scale explanations*: Contours were formed from phase-aligned Gabor patches. Results presented in Chapter 3 have shown that coarse-scale mechanisms are well suited to the detection of such stimuli. It may be that the ‘S’ and ‘C-C’ stimuli featured in the previous experiment would be indistinguishable to such a mechanism. However, if results were due to a simple-filter mechanism the non-proximate ‘C C’ stimuli ought to have been detected at a similar rate. Furthermore, the following chapter demonstrates that the simple-filter model is not well suited to the task of detecting smooth contours. Therefore, a coarse-scale explanation is discounted for the moment.
- iii) *Type II error*: No difference was found in detection rates for ‘S’ and ‘C-C’ contours. It is possible that the previous experiment was not sensitive enough to reveal a difference, when in fact a difference did exist.

In order to address the criticisms raised above, a number of enhancements were made to the stimuli featured in the current experiment. Inflection points occurred twice as frequently, addressing the first criticism. When contours were broken, the remaining segments were either 4 or 8 elements long. The arc length of 4 elements ought to have been well below the 8 element asymptote level previously suggested. In order to counter the second criticism, the Gabor patches were phase-alternated. Therefore, the embedded contours were less susceptible to detection using a coarse-scale mechanism.

In order to address the third criticism additional breaks were introduced into contours. Whilst in the previous experiment all breaks were introduced at the points of inflection, the current experiment also introduces breaks between inflection points, i.e. in the centre of arcs. Therefore, it should be possible to make a comparison between the relative detectability of contours that were broken at their inflection points and those that were broken elsewhere. It was hypothesised in the previous section that inflection points *per se* do not contribute towards contour detectability, implying that contours were already broken at these points. Therefore, there should have been a greater penalty for breaking contours in the centre of arcs than there should be for breaking contours between arcs.

4.5.1.1 Summary of methodology and working hypotheses

The stimuli used in the current experiment are illustrated below, Figure 4-18, the dashed lines represent the *virtual* breaks predicted by the smoothness hypothesis.




Figure 4-18. Typical stimuli utilised in the current experiment. Dashed lines represent points at which the sign of path-angle change changes from positive to negative or vice-versa.

Two patterns of relative detectability can be predicted for the stimuli pictured above. The first pattern conforms to the hypothesis that contours are not integrated at their inflection

points. Whilst the second pattern of results is what one would expect if the quality of integration was unaffected by the presence or absence of inflection points.

Hypothesis If one accepts that integration at contour inflection points is weaker, or that it does not occur, then breaking contours at their inflection points should have little influence upon the detectability of the stimuli. The detectability of stimuli A, C and D should be equal and greater than the detectability of stimuli B. Stimuli B features an additional break in the middle of the central arc. Consequently, this stimuli consists of 2x four element arcs and two arcs with only two elements. The B stimuli should be least detectable as it features fewer extended contours. *Prediction: (A=C=D)>B*

Null hypothesis If the degree of integration of elements within arcs and between arcs is the same then there should be a penalty for breaking a contour at the points of inflection. Consequently, the lengths of individual contours within stimuli A-D will be the sole determinant of the detection rates of observers. If this were the pattern of experimental results, then one would have to reject the hypothesis that integration only involves elements that are members of the same arc. *Prediction: A>C>B>D*

4.5.2 Methodology

4.5.2.1 Stimuli

Two independent variables were manipulated during this experiment. The most important variable was the shape of each of the contours (A-D, Figure 4-18). Contour path-angle was also manipulated in order that the relative detection rates for each contour type could be examined across a range of salience levels, thereby reducing the likelihood that a ceiling or floor effect could mask any difference in salience that exists between the contours. Additionally, this manipulation enabled an examination of whether the depth of the contour arcs had any influence upon integration.

Each of the broken contours pictured within Figure 4-18 (B-D) was created by deforming contour A. The deformations were achieved by rotating sub-sections of contour A about their centres. By creating stimuli in this manner, the overall size of each contour is approximately maintained. In order to maintain clarity, a description of each deformation will follow a full description of contour (A).

Contour (A) consisted of twelve Gabor elements. The sign of path-angle change was manipulated such that the contour featured three arcs, each four elements long. Thus, there

was a change in the sign of path-angles after the fourth and the eighth Gabor elements. The magnitude of the path-angle was manipulated between images; the path-angle value for each contour was selected from values between 30 and 50° in steps of 5°. The spacing between successive contour elements was matched to that of the stimulus background, i.e. it was 16 pixels. The path-relative orientation of Gabor elements was modified by a *jitter* of 5°. Contour (B) was achieved by breaking contour (A) into two halves. One of the half-contours was rotated by 45°. For contour (C) the (A) contour was broken into two contours one 4 elements long and the other 8 elements long. The 8-element contour was rotated by 45°. Finally, contour (D) was created by rotating the central 4-element arc by 45°.

As the relative differences in detectability of the embedded stimuli were likely to be small, a number of additional influences upon detectability were rigorously controlled. Contours were positioned such that their centres were placed upon the circumference of a circle centred upon the fixation point. The radius of this circle was 3° (visual angle). The stimulus background was generated using the *spaced-fill* technique (see Section 6.2.4 for a summary). This background was utilised instead of the *jittered-fill* technique, as it was felt that the relative density of the elements within target contours and the background should be matched as closely as possible. The minimum level of separation between background elements was 16 pixels.

The extent to which background elements were co-linear with one another was also controlled, minimising the variability of an influence that is normally random. This was achieved by rotating individual background elements, whilst measuring the degree to which they were co-linear with their neighbours. The aim of this manipulation was to produce a background that featured fewer of the spurious contours that occur when background elements are randomly oriented. The target level of background *co-linearity energy* was 1.0 (this term is defined in Section 6.1). Details of this manipulation are provided in Section 6-2.

The overall size of each stimulus image was 360 x 360 pixels. Gabor patches for both the embedded contours and the stimulus backgrounds were phase-alternated. Twenty images were created for each combination of path-angle (5 levels) and contour-type (4 levels). A set of noise images was also generated. The noise images were created using exactly the same technique, described above, as the target images. They were identical to target images, except for the fact that the orientations of Gabor elements belonging to the

embedded contours were randomised. Thus, there were no differences in the spacing density of elements within target and noise images.

4.5.2.2 *Observer details*

Two observers participated in the current experiment, KAS an experienced observer in path-paradigm experiments and PGL (the author); both observers have corrected to normal vision. Each observer participated in a total of six observation sessions. Each session lasted approximately half an hour. During a session 200 pairs of target and noise images were viewed.

4.5.2.3 *Procedure*

Images were presented upon a Cambridge Research VSG system, full details of this system are provided in Section 2.1.3. Briefly, the screen resolution was 1024 x 768 pixels and the screen frame-rate was 100_{hz}. Images were viewed from a distance of 500mm, stimulus scale properties that correspond to this viewing distance are illustrated below in Table 4-8.

Table 4-8. Stimulus scale properties for the current experiment.

	Pixels	Lambda	Visual-angle
Gabor patch full-wave frequency	4	1λ	0.173°
Background and contour minimal nearest-neighbour spacing	16	4λ	0.69°
Contour offset's, relative to fixation point	69	17.25λ	3°
Overall image size	320 x 320	$80 \times 80\lambda$	$14.11 \times 14.11^\circ$

Presentation timing was unchanged from that used in the previous experiment (Section 4-21).

4.5.3 Results

Contour detection rates for each observer are presented as a function of contour-type and contour path-angle in the Figure 4-19. Fitted psychometric functions for contour-type and path-angle are presented in Figure 4-20.

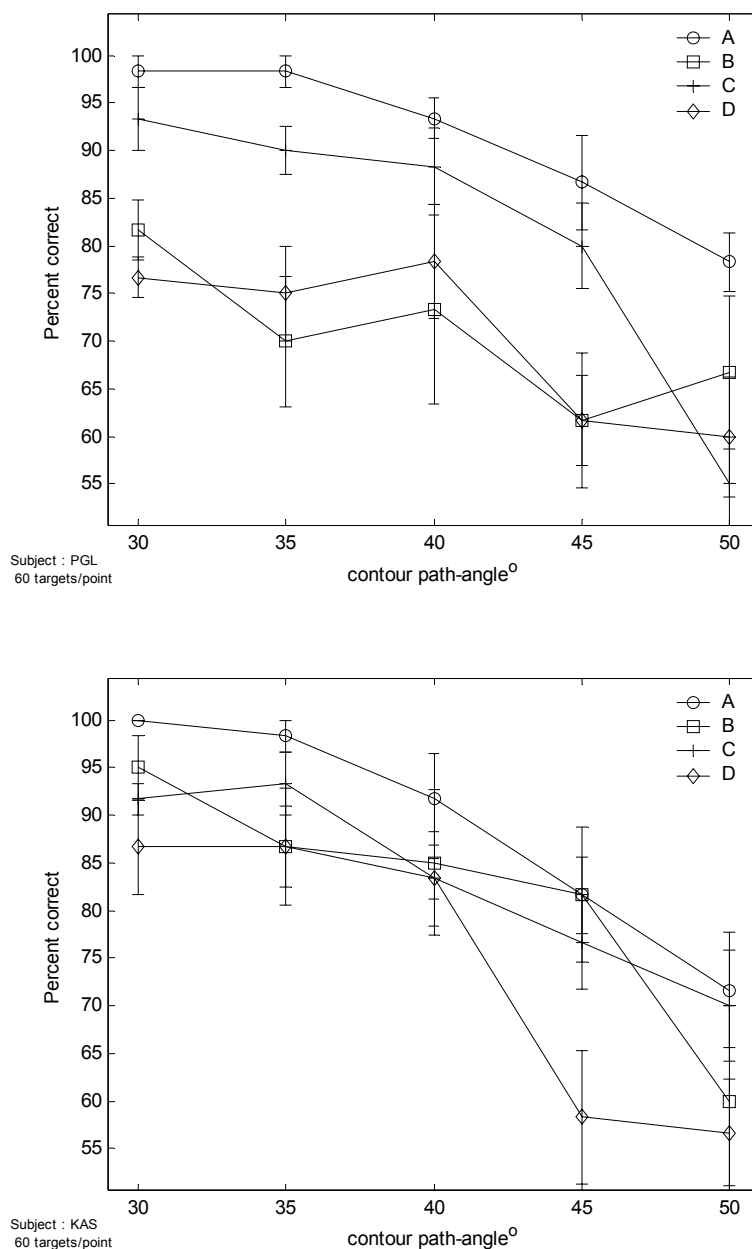


Figure 4-19. Contour detection rates as a function of contour-type and contour element spacing. [top] contour detection rates for observer PGL. [bottom] contour detection rates for observer KAS. In both plots the error bars represent ± 1 standard error.

The quality of the fit between the raw observer performance rates and the fitted psychometric functions is variable. For condition A there is a close correspondence

between the actual performance rates and the fitted psychometric functions. However, the correspondence is increasingly poor for conditions B, C and D respectively.

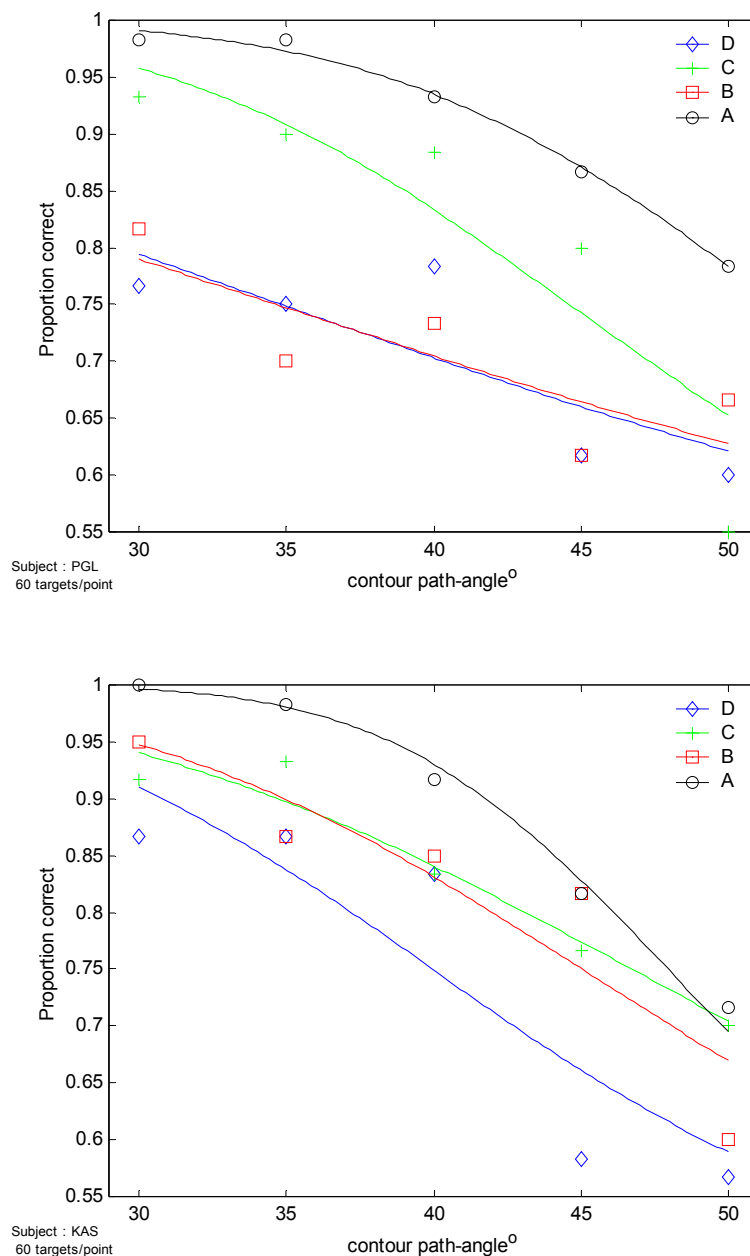


Figure 4-20. Fitted psychometric curves as a function of contour-type and contour path-angle. Refer to Figure 4-19 for standard-error estimates of variance.

In order to compare the relative detectability of the four types of contour (A-D) a scalar estimate was required. The path-angle value at which performance rates are 75% correct was plotted for each type of contour, Figure 4-21.




Figure 4-21. Contour path-angle when fitted psychometric functions predict 75% correct detection rates.

The statistical significance of the differences in contour detection rates for the four types of contour (A-D) was examined using a χ^2 analysis (Table 4-9). For observer PGL all differences were significant except that between contours B and D. For observer KAS all differences except that between B & C were significant.

Table 4-9, χ^2 analysis of differences in detection rates for observers KAS <top> and PGL <bottom> for the stimuli employed in the current experiment. The analysis examines observer responses across all path-angle values.

===== χ^2 significance matrix =====				
Contour type	A	B	C	D
A	*	5.818	3.961	20.439
B	0.016	*	0.183	4.701
C	0.047	0.669	*	6.713
D	< 0.000	0.030	0.010	*
Contour type	A	B	C	D
A	*	40.029	11.759	41.080
B	< 0.000	*	9.357	0.008
C	0.001	0.002	*	9.904
D	< 0.000	0.929	0.002	*

The upper triangle gives the χ^2 values and the lower triangle the p -values. Statistically significant differences are highlighted in **bold**.

Contours were ranked according to their relative levels of detectability (Figure 4-21). This was achieved by comparing the degree of path-angle deviation that leads to a predicted detection rate of 75%. The most detectable contour was that which had the largest path-angle value when detected at a rate of 75%. For observer KAS the pattern was $A > C > B > D$ and for observer PGL the pattern was $A > C > (B = D)$. If only those differences which were statistically significant are accepted (Table 4-9), then these patterns become - $A > (C = B) > D$ (observer KAS) and $A > C > (B = D)$ (observer PGL). Whilst these patterns are not identical and the quality of the fitted psychometric functions was quite poor, it is clear that they do not correspond to the pattern suggested by the experimental hypothesis $(A = C = D) > B$. The patterns for both observers most closely matches that predicted by the null-hypothesis $A > C > B > D$. Thus, breaking contours at their inflection points does reduce the detectability of the contour.

4.5.4 Discussion

The levels of detectability of the contours A-D most closely resembled that predicted by the null-hypothesis. This suggests that integration does occur at contour inflection points. Consequently, one must conclude that the hypothesis that inflection points *break* contours is not supported by the results of the current experiment.

4.6 General Discussion

Experiment 3, a partial replication of Kovacs and Julesz's (1993) experiment, confirmed that closed contours are detected more frequently than open contours. However, the pattern of results did not suggest that there was a qualitative change in the integration process for closed contours. Kovacs and Julesz suggested that the addition of 1-2 elements, that finally closed a contour, enabled an oscillatory process that raised a contour's salience. In the current closure experiment, there was a linear relationship between the degree of closure and the contour's detectability. This pattern of results does not correspond that that which would be predicted by the 'synergistic process' explanation offered by Kovacs and Julesz.

Whilst the stimuli employed by Kovacs and Julesz may have been confounded by contour length (Braun, 1999), it appears that the stimuli used in the current closure experiment may have been confounded by contour arc-length. When local variables such as path-angle and element spacing are held constant, the frequency of inflection points within a contour has a negative correlation with closure. Consequently, closed contours featured much longer arcs than open contours; in other words, they are smoother. Experiment 4 (Section 4.3), indicated that the *arc-length* within embedded contours has a significant influence upon

detectability. Where arc-length and consequently smoothness is increased, the contour is more readily detected. Consequently, predictions of contour salience should take the frequency of changes in the sign of the path-angle into account.

The discussion of experiment 4 (Section 4.3.4) offered the hypothesis that the *smoothness* effect occurred because integration initially only occurs within individual arcs. Experiment 5 attempted an examination of this hypothesis. No difference was found in the detectability of embedded 'S' shaped contours and similar contours in which one half of the contour, a 'C' was rotated. Thus, disrupting the continuity of the contour across the inflection point did not reduce contour detectability. However, the detection rates for the 'S' and 'C-C' contours were higher than would be have been expected if rates were merely based upon the salience of individual arcs. This suggests that the proximity of the arcs led to an increase in contour detection rates.

Experiment 6 (Section 4.5) was designed with the intention of clarifying the results of the third experiment. Here, a comparison was made between the influence of contour breaks at inflection points and elsewhere. It was hypothesised that if integration occurred solely within arcs then detection rates should not fall when contours were broken between arcs. Conversely, it was predicted that breaks at points between inflections should reduce contour detection rates. This hypothesis was not supported by the experimental results. Recently, Lamote and Wagemans (1999) have used the deletion detection paradigm to investigate the relative strength of integration at different points within a contour. They found that integration was actually strongest at the inflection points. This finding certainly seems to undermine an explanation of the smoothness effect that is based upon greater inter-arc integration.

4.6.1 Summary

The experimental results presented within this chapter have indicated that contours featuring fewer inflection points are detected more frequently than contours that have a larger number of inflection points. However, it is not clear what mechanism underlies these detection rate differences. Experimental evidence did not support the hypothesis that contour integration initially only occurs for elements within continuous arcs. There are many reasons why it may be beneficial for a visual system to be sensitive to the presence of smooth contours. The ability to detect smooth uninflected lines would certainly be beneficial when attempting to detect structure in noisy images, for example when viewing camouflaged objects.

The results of the second experiment prompted the suggestion that multiple mechanisms may be involved in contour detection. The results show increased detectability for smoother contours and an increased level of detectability for contours in which the sign of path-angle change was alternated after every element. The latter contour might be most readily detected by a coarse-scale mechanism, such as the simple-filter model. Hence, of the contours represented within Figure 4-22, the *jagged* contour (right) should be most compatible with a *simple-filter* mechanism. This possibility is investigated within the following chapter.




Figure 4-22. Two short contours featuring either ‘smooth’ (left) or ‘jagged’ (right) global structures. Locally all properties are identical. Only the frequency of path-angle sign changes differs between the two contours.

-0-0-0-

5 How does arc-length influence the performance of the simple-filter model?

The results of the modelling experiments presented in Chapter 3 have shown that the simple-filter model remains a tenable account for the performance of human observers in path-paradigm experiments. During these simulations only local stimulus properties such as contour path-angle, element spacing and Gabor patch phase were manipulated. Experimental evidence presented in Chapter 4 suggests that the global structure of a contour can influence the detection rates of human observers. Contours with fewer inflections were detected most frequently. The experiments reported within this chapter examined how the presence of inflections within target-contours influenced the detection rates of the simple-filter model. In Section 4.6.1 it was suggested that coarse-scale mechanisms, such as the simple-filter model, might only be suitable for the detection of contours that are frequently inflected, i.e. jagged contours. If the simple-filter model were relatively impaired when detecting smooth contours, then such stimuli would provide a useful means of evaluating the extent to which coarse-scale mechanisms underlie human performance in path-paradigm experiments.

5.1 *Simulation 4: The performance of the simple-filter model with the arc-length stimuli*

5.1.1 Introduction

Experiment 4 demonstrated that contours composed of longer arcs were most readily detected by human observers. However, contour detection rates were not simply a linear function of arc-length. If the relationship had been purely linear then contours that were inflected between each successive element should have been the least detectable. In fact, detection rates for such contours were higher than they were for contours featuring inflections after every second and even after every fourth element. Figure 5-1 shows the relative levels of detectability for contours composed of arcs of various lengths. The relative levels of detectability were estimated by calculating the degree of element separation that was necessary to achieve a 75% detection rate these are based upon the fitted psychometric functions.




Figure 5-1. 75% detection limits as a function of contour element spacing and contour-arc length. For contours with arc-lengths of two or more elements, there is a near linear relationship between detectability and arc-length. This linear relationship fails for contours with an arc-length of one element.

It has been suggested that different mechanisms might underlie the detection of smooth and jagged contours (Section 4.3.4). Where the term jagged describes those contours which feature frequent inflections, whilst smooth contours feature fewer inflections. It was proposed that a coarse-scale mechanism, such as the simple-filter model, might underlie the detection of jagged contours. In these contours, the orientations of neighbouring contour elements do not diverge significantly. Therefore, it was expected that such stimuli would lend themselves to the creation of extended ZBRs, which would correspond to large segments of the embedded contour. Conversely, coarse-scale mechanisms would be poorly suited to the detection of smooth contours. In smooth path-paradigm contours, the orientations of neighbouring elements diverge as a function of their separation along the contour. Consequently, it was expected that those ZBRs within a filtered image which represent smooth contours would be shorter. The current simulation presents the arc-length stimuli utilised within experiment 4 (Section 4.3) to the simple-filter model. Thereby, enabling a direct comparison of the influence that arc-length has upon the simple-filter model with that already reported for human observers.

5.1.2 Methodology

5.1.2.1 *The simple-filter model*

All the details of the simple-filter model are exactly as described in Section 3.2.1.

5.1.2.2 *Stimuli*

Images are the same as those utilised in experiment 4 (Section 4.3). Briefly, the size of images was 320 x 320 pixels. Target images featured contours with varying arc-lengths. Arc-lengths were varied by manipulating the sign of path-angle changes between elements. Contours featured arc-lengths of 1, 2, 4, 8 and 15 elements. The spacing of contour elements was also manipulated between $4-6\lambda$ in steps of 0.66λ .

5.1.2.3 *Comparing performance measures for human observers and the simple-filter model*

In order to compare the simple-filter model's performance with that recorded for human observers a summary of the relative degree of detectability of the various contour stimuli was required. The predicted 75% detection limits of contours as a function of contour spacing and arc-length gave a useful measure of the relative detectability of contours with different arc-lengths (Figure 5-1). Clearly, for arc-lengths of two and above there was a gradual increase in contour detectability, whilst detection rates for contours with an arc-length of one element fell between those for arc-lengths of four and eight elements.

The relative levels of detectability of contours for human observers are reported below in (i) and the predicted levels of relative detectability for the simple-filter model are offered alongside in (ii).

		Most detectable	→	Least detectable
i)	Human observers	15 > 8 > 1 > 4 > 2		
ii)	Simple-filter model prediction	1 > 2 > 4 > 8 > 15		

Thus, performance for human observers largely correlates with the global smoothness of the contours, whilst it was hypothesised that performance for simple-filter model would be inversely related to contour arc-length.

5.1.2.4 *Maximising the likelihood of hypothesis falsification*

Whilst it was accepted that early vision involved the parallel operation of filters with widely varying scales, the assessment of the simple-filter model presented in Chapter 3 was based upon the operation of a restricted set of filters. It was suggested that different filters would be suited to the detection of contours composed of phase-aligned and phase-alternated Gabor patches. This restriction to two filter-scales reflects a compromise between the position of Hess and Dakin (1999) who allowed only a single-filter scale, and the accepted reality of early vision, where filters of many spatial-scales are available.

Despite this constraint, the performance levels for the simple-filter model approximated those of human observers.

For the current experiment, the hypothesis states that detection rates for the simple-filter model will differ from those recorded for human observers. In order to maximise the likelihood that this hypothesis could be falsified the previous constraint was removed. For each level of stimuli, the ideal-filter was identified, i.e. the filter most successful at detecting the target stimuli. Subsequently, the performance of the model at each stimulus level was based upon the model's performance in conjunction with the ideal-filter. A number of restrictions were placed upon the range of filters that could be deployed by the simple-filter model : -

- i) As with previous simulations the lengths and widths of filters were restricted to values of 0.5 to 10σ .
- ii) Filters were categorised as either elongated (length > width) or broad (width > length). This enabled an examination of the performance of the simple-filter model when the ideal-filters were elongated, broad or drawn from either category.

Whilst the performance of the simple filter model will be assessed using both broad and elongated filters, those estimates that rely solely upon the contribution of elongated filters are likely to be considered most biologically plausible (eg. Parker and Hawken, 1988). When the ideal-filters are used the reported performance should be the best that could be achieved by a coarse-scale process, which was able to select the most appropriate filter-scale for any given stimuli. If the performance of the model most closely matched that predicted then it must be concluded that the simple-filter model does not account for human contour detection where arc-length is manipulated. Conversely, if it were found that the performance of the simple-filter model corresponded to that reported for human observers, then this model would remain a tenable account of human performance in path-paradigm experiments.

5.1.3 Results

5.1.3.1 Performance levels for all filters

The differences in ZBR lengths for target and noise images are presented upon an array of contour plots in Figure 5-2. Each contour plot corresponds to a particular combination of contour arc-length and element spacing. For contours with an arc-length of one element there is a clear range of filter-scales which lead to ZBR lengths which are longer for target

images than they are for noise images. This suggests that such filters will give lead to high performance levels in simulated 2AFC experiments. For all other levels of arc-length, there are few clearly defined peaks of equivalent extent and amplitude.




Figure 5-2. The array of contour plots above illustrates the differences between the maximum ZBR lengths for target and noise images. Each column corresponds to a particular level of contour element spacing and each row corresponds to different levels of arc-length. Clearly, for images featuring embedded contours with an arc-length of 1 element there was discrete range of filters which led to differences in the ZBR lengths of target and noise images (centred at around a scale of $W3.5\sigma * L8\sigma$). For other arc-lengths, there are far fewer filters suitable which gave rise to differences in ZBR lengths for target and noise images.

It appears that the maximum lengths of ZBRs for target and noise images differ over a larger range of filters when embedded contours are jagged. This was confirmed by

counting the proportion of filters for which there was a significant difference in the means of the longest ZBRs for target and noise images. This was calculated using a two-tailed unrelated t-test. Differences were considered significant when the p -values were less than or equal to 0.05. Figure 5-3 confirms that the number of filters for which there was significant difference in ZBR lengths for target and noise images fell as a function of arc-length.




Figure 5-3. The proportion of filters that led to a significant difference in maximum region lengths between the target and noise images decreased as a function of arc-length and of element spacing.

The performance of the simple-filter model was evaluated using a simulated 2AFC paradigm (procedural details have already been given in Section 3.2.1.3). The contour detection rates of the model are plotted upon an array of contour plots in Figure 5-4. The model's performance was generally better for contours which featured many inflections. For contours with the smallest element spacing distance ($\lambda = 4$) and an arc-length of one element, the simple-filter model achieved a detection rate of 90% correct.




Figure 5-4. Contour plots illustrating percent correct scores for the simple filter model as a function of filter-scale and contour arc-length and element spacing. Each column corresponds to a particular level of contour element spacing and each row corresponds to different levels of arc-length.

5.1.3.2 *Examining the performance of the simple-filter model in conjunction with the ideal-filters*

Figure 5-5 presents the highest contour detection rates achieved by the simple-filter model using elongated ideal-filters; detection rates are shown as a function of contour arc-length and element spacing. The equivalent plot for broad filters is presented in Figure 5-6, whilst Figure 5-7 shows the models performance when the ideal filters were selected from all filters, regardless of the elongation ratio. With very few exceptions, the relative levels of contour detectability correspond to those predicted by the experimental hypothesis. For the simple-filter model contour detection rates fell as a function of contour arc-length.




Figure 5-5. Contour detection rates for the simple-filter model in conjunction with the elongated ideal-filters. Contour detectability was inversely related to arc-length and element spacing. Details of the scale of the ideal-filters used in conjunction with each stimulus level are provided within Appendix B.




Figure 5-6. Contour detection rates for the simple-filter model in conjunction with the broad ideal-filters. As with the elongated filters, contour detectability was inversely related to arc-length and element spacing. The relationship between arc-length, element spacing and the detection rates is less consistent with broad filters.




Figure 5-7. Contour detection rates for the simple-filter model in conjunction with the elongated ideal-filters. Contour detectability was inversely related to arc-length and element spacing. Where inconsistencies with this relationship occur they are due to the contribution of broad filters.

A post-hoc examination of the contour detection rates of the simple-filter model with each filter-scale was carried out. This revealed that there were no elongated filters-scales which led to detection rates that were at least 10% greater for smooth contours (arc-length ≥ 2) than they were for jagged contours. Where performance was better for smoother contours, albeit with a difference of less than 10%, the detection rate was well below the level reported for human observers.

5.1.4 Discussion

Regardless of the assessment criterion selected, the performance of the simple-filter model worsens as a function of arc-length. As arc-length was increased, fewer filter-scales gave rise to a significant difference between the maximum lengths of ZBRs of filtered target and noise images. Detection rates of the simple-filter model, in conjunction with ideal-filters, fell as a function of arc-length. This was found to be the case whether filters were elongated or broad. There was no instance where performance at a particular filter scale was at least ten percent higher for smooth contours than it was for jagged contours. At some filter-scales, the simple-filter model detected smooth contours more frequently than jagged contours. However, the performance levels of the model were far below those reported for human observers. If coarse-scale mechanisms were responsible for the detection of smooth contours, it follows that the filters that were best suited to the detection

of smooth contours would also be available when detecting jagged contours. Consequently, performance would still be higher for the detection of jagged contours, as the ideal-filters for smooth contours still led to higher detection rates for jagged contours.

One potential explanation of the performance of the simple-filter model, is that target-contour ZBRs were longer where the contour was frequently inflected. In path-paradigm contours, the product of arc-length and path-angle gives the maximum difference in the orientation of any pair of contour elements. Therefore, the maximum difference between the orientations of contour elements will be smaller where there are more inflections. As arc-length is increased, there will be a reduction in the number of contour elements that share the same orientation. Consequently, fewer contour elements are likely to share the same orientation as a particular filter. Simulation 4 (Section 3.6) has demonstrated that it is unlikely that a pair of Gabor elements will be represented by the same ZBR if their orientations differ significantly. Consequently, the maximum ZBR length in an image will decrease in smoother contours.

In Section 4.5.1(ii) it was suggested that the detection of the eight element arcs might have been achieved by a coarse-scale mechanism. The evidence presented in the current chapter suggests that this was unlikely to have been the case. Detection rates for the simple filter model were low where contours featured no inflections.

5.2 Simulation 5: Detection rates for the simple-filter model were influenced by the presence of inflections within contours

5.2.1 Introduction

The results of the previous simulation were consistent with the experimental hypothesis; the simple-filter model was better suited to the detection of jagged contours than it was to the detection of smooth contours. With the current target stimuli, the smoothness of embedded contours was manipulated. This manipulation led to three consistent differences in contours that had different levels of smoothness: -

- i) The frequency of inflections within contours.
- ii) The overall size of contours.
- iii) Arc-length (synonymous with smoothness in the current experiment).

The first of these differences, inflection frequency, is most likely to be the reason why the simple-filter model is better suited to the detection of jagged contours. In jagged contours, which featured an inflection after every element, the orientations of any pair of elements

within the contour would not have differed by more than 24° . Consequently, the simple-filter model is best suited to detection of jagged stimuli.

The jagged stimuli featured in the previous experiment featured many inflections. So there is little point in investigating the correspondences between these inflections and the longest ZBRs. After all, nearly all elements within the most jagged contours were the centre of inflections. Instead, the following sub-section examines the spatial correspondence of inflections and longest ZBRs for images containing contours with a randomised global structure.

5.2.2 Methodology

If the arguments detailed above are correct, then the location of inflections within contours should correspond with the longest ZBRs found by the simple-filter model. In order to test this claim, stimulus images were examined and the locations of inflection points within target-contours were logged. Subsequently, it was noted whether any contour inflections fell within the area of the longest ZBR. It was also necessary to indicate whether the target-image was assigned target or noise status in the 2AFC decisions made by the model. The assignment of target or noise status to a particular image is variable as it depends upon the relative extent of ZBRs in both 2AFC images. Therefore, the mean length of the longest ZBRs in noise images was calculated for each filter-scale. This value was subtracted from the maximum ZBR lengths of target images at the corresponding filter-scales. If the result was positive, then the target image was assigned ‘hit’ status, ‘miss’ status was assigned for negative or zero values.

The locations of inflections in the target contours of the path-angle stimuli (described in Section 3.3.1) were logged. These were compared with the longest ZBRs found in images filtered with the filter found to be most successful at detecting images containing embedded contours (length 2.5σ x width 3.5σ).

5.2.3 Results

5.2.3.1 *Contour inflection and longest ZBR overlaps in path-angle stimuli*

Figure 5-8 shows the number of overlaps between the longest ZBRs and the inflections within the path-angle stimuli. For 63 hits with the $W2.5\sigma$ x $L3.5\sigma$ filter, the longest ZBR overlapped an inflection within the contour, a remaining 25 hits did not overlap a contour inflection. Only two misses were recorded, for both of these images, the ZBRs did not overlap the inflections.




Figure 5-8. Overlaps between ZBRs and inflection points in contour which have a random global structure.

For images featuring contours with a path-angle of 60° , all longest ZBRs overlapped inflections. A visual examination of the images revealed that these contours featured many inflections. As the path-angle is relatively large, the likelihood that elements will overlap in randomly generated contours was increased. Those contours that did not feature overlapping Gabor elements would feature many inflections, thereby preventing overlaps between contour elements. Contours with even larger path-angle values also featured many inflections, however the relative difference in orientation between Gabor elements was large enough to prevent the creation of extended ZBRs between neighbouring elements. Examination of Figure 3-18 (left) reveals that the performance of the simple-filter model approaches chance levels as path-angle exceeds 60° .

5.2.4 Discussion

The analysis of the overlap between contour inflections and the longest ZBRs was confounded in the arc-length stimuli. A post-hoc examination of these image properties in conjunction with the path-angle stimuli revealed that the vast majority of longest ZBRs corresponded to the location of inflections within target contours. This confirms that for the simple-filter model, inflections within target contours lead to longer ZBRs in the model's output. Consequently, the model's performance is significantly improved when it is exposed to contours featuring many inflections.

5.3 General discussion

The simple-filter model was specifically impaired when detecting stimuli containing smooth contours. Contour detection rates fell as a function of arc-length. Even when the filter-scale was selected in order to maximise the detection of smooth contours, detection rates were consistently lower for smoother contours than they were for jagged contours. This pattern of results is radically different from that reported for human observers. For human observers performance increases as a function of arc-length.

The simple-filter model is best suited to the detection of jagged contours because these contours feature many inflections. Inflections within a contour serve to reduce the magnitude of orientation differences between contour elements. Consequently, more of the Gabor elements in a jagged contour will have similar orientations, this leads to longer ZBRs in the output of the simple-filter model.

The modelling evidence presented in the Chapter 3 has shown that previous path-paradigm experiments have not adequately controlled for the potential contribution of coarse-scale mechanisms. These experiments featured contours with a random overall structure, such contours may feature many inflections. Stimuli featuring contours that manipulate the frequency of inflections may provide a means of reducing the potential contribution that coarse-scale mechanisms may make towards the detection of path-paradigm stimuli. This would enable a more accurate estimate of the relative contributions of coarse-scale and associative-field mechanisms during contour detection. Hess and Dakin (1997, 1999) have argued that contour detection in para-foveal areas is achieved using only coarse-scale mechanisms. If this assertion is correct then human observer performance should be impaired when contours presented to the para-fovea feature few inflections. This possibility is investigated in Chapter 7.

6 Manipulating the level of co-linearity within the stimulus background

6.1 Overview

The current literature on contour integration largely ignores the presence of spurious contours that occur within the stimulus background when element orientations are assigned randomly. In a randomly generated background, it is possible that many such contours will be present. Therefore, the subject's task becomes one of selecting the deliberately embedded contour (henceforth termed the *target-contour*) from an image that includes many contours that have occurred through chance alignments of background elements (*noise-contours*). Estimates of the structure of the association field and of its spatial extent are based upon the results of experiments that may not have sufficiently taken into account the influence that noise-contour visibility has upon subject performance. This omission is understandable as such contours should occur with equal frequency in both target and noise images. However, estimates of the limits of contour detectability may be influenced by the presence of noise-contours. Performance may approach chance levels when the salience of target-contours is similar to that of noise-contours.

Braun (1999) has suggested that the stimulus backgrounds created by researchers such as Field, Hess and Hayes (1993) and Kovacs and Julesz (1993) do not feature an element spacing density that adequately matches the separation of successive contour elements. Braun assessed the distribution of contour and background element separations within stimulus images. This was achieved by calculating the pair-distribution functions for element separations. Pair distributions are achieved by measuring the distances between each element within a stimulus and all other elements. This gives a frequency distribution of element spacing. Peaks within these distributions reveal patterns in the spacing of stimulus elements (see Figure 6-7, section 6.2.4 for an example). The pair distribution function illustrated differences in spacing density for contour and background elements. The pair-distribution for embedded contours has two peaks reflecting the separations of nearest-neighbour and next nearest-neighbour elements. Whereas backgrounds have either a flat distribution (Kovacs and Julesz, 1993) or an increasing element density up to separations of 4λ (Field *et al*). Braun argued that because the pair-distribution function of background elements differed from that of target-contours, any noise-contours that were present would have a different spacing distribution than target-contours.

In order to create a closer match between the pair-distribution functions of contour and background elements, Braun (1999) offered an alternative method of generating backgrounds. By generating backgrounds with a spacing density that is similar to that of target-contours, Braun claimed to have maximised the likelihood that noise-contours will be present. Thus, contour detection thresholds are more likely to reflect the average visibility of noise-contours. However, the occurrence of such contours is still largely uncontrolled, as the orientation of background elements is still randomised. If Braun's claim that his detection thresholds reflect the average salience of noise-contours is correct, then manipulating the salience of noise-contours should affect the detection limits for various contour parameters. Detection limits should be affected if the visibility of noise-contours is manipulated whilst the spacing pair-distribution function is held constant. Conversely, if manipulation of the visibility of noise-contours does not affect detection limits, then Braun's experimental results may simply reflect the change in background generation technique.

Manipulations of the visibility of noise-contours were achieved, in this chapter, by iteratively applying a model association field to each element within the stimulus background. The model association field was based upon that described by Yen and Finkel (1998). Their association field model gave a measure of the strength of co-facilitatory connections between neighbouring receptive fields. In the current implementation, this estimate of connection strength was used as a measure of the degree of co-linearity between a central element and surrounding background elements. The orientation of each background element was then manipulated in order to vary the degree of co-linearity between the central element and surrounding elements. By applying this process iteratively to all background elements, it was possible to manipulate the overall visibility of noise-contours within stimuli.

The experiments described here investigated whether the detection limits for contour path-angle and contour-length would vary as the visibility of noise-contours was manipulated. Specifically, reducing the visibility of noise-contours should result in an extension of the path-angle limit and in a reduction of the length of a detectable target-contour.

6.2 Manipulation of background element co-linearity

6.2.1 Measuring co-linearity using an association-field model.

Before the presence of noise-contours can be manipulated, some means of measuring their prevalence is required. The association field theory described by Field *et al.* (1993)

suggests that a linking process determines whether a specific element will be integrated with those nearby. The strength of integration between any given pair of elements is determined by the relative orientation and position of these elements. In simple terms, contour integration is strongest when a short line of low curvature can be drawn through the orientation axes of the neighbouring elements. Therefore, the strength of integration between a specific element (henceforth referred to as the target) and those elements that fall within its association field (henceforth referred to as the surround elements) will indicate the likelihood of contour membership.

A simple model of the association-field is adopted, one that has already been employed in the published literature (Yen and Finkel, 1998; Parent and Zucker, 1989). The model's prediction of the strength of the co-facilitatory connections between a target receptive field and surrounding receptive fields was calculated. These values were then summed and used as a measure of the likelihood that the target element was a member of a contour. In its current guise the association-field model was employed as a measure of the likelihood of an element's contour membership, therefore the phrase *co-linearity energy* (CLE) will be substituted for the phrase co-facilitation. An element that has a very high level of CLE is likely to form part of a contour, whereas an element with low CLE is unlikely to be a member of a contour.

Within this model, the CLE of an element is defined by the product of three functions. The 'fan-out' of the association field is defined by function 6-1, where θ is the orientation of the 'target' element, k_t is the trans-axial fan-out (10°) and k_c is the co-axial fan-out (30°). Function 6-2 defines the optimal orientation of surround elements (ϕ) using the co-circularity constraint, where θ is the orientation of the target element. Function 6-3 defines activation as a Gaussian function of the acute angle between the neighbour's orientation and ϕ , where ψ is the actual orientation of the surround element and σ_ψ^c defines the range of angles over which facilitation occurs (20°). Function 6-4 defines the fall-off in co-linearity energy with increasing distance between the target and surround line segment, where σ_d^c defines the range of distances over which facilitation occurs (40 pixels). The product of functions 6-1, 6-3 and 6-4 determines the co-linearity energy (Function 6-5).

$$\Gamma(\theta, i, j) = \begin{cases} 1, \text{ if } \tan^{-1}\left(\frac{j}{i}\right) - \theta < k_c, \\ \text{and } \tan^{-1}\left(\frac{-j}{-i}\right) - \left(\theta + \frac{\pi}{2}\right) < k_c. \\ 1, \text{ if } \tan^{-1}\left(\frac{j}{i}\right) - \theta = \frac{\pi}{2} \pm k_t, \\ \text{and } \tan^{-1}\left(\frac{-j}{-i}\right) - \left(\theta + \frac{\pi}{2}\right) = \frac{\pi}{2} \pm k_t. \\ 0, \text{ otherwise.} \end{cases} \quad \text{Function 6-1}$$

$$\phi(\theta, i, j) = 2 \tan^{-1}\left(\frac{j}{i}\right) - \theta \quad \text{Function 6-2}$$

$$B(\theta, i, j, \Psi) = G\left(\left|\psi - \phi(\theta, i, j)\right|_{\text{acute}}, \sigma_{\psi}^c\right) \quad \text{Function 6-3}$$

$$D(i, j) = G\left(\sqrt{i^2 + j^2}, \sigma_d^c\right) \quad \text{Function 6-4}$$

$$F(x, y, \theta) = \iint\limits_{(i,j) \in N} \Gamma(\theta, i, j) D(i, j) B(\theta, i, j, \Psi) \quad \text{Function 6-5}$$

Figure 6-1 shows the extent and structure of the associative field that is defined by functions 1-5.




Figure 6-1, the model associative field defined by functions 1-5. The central bold line represents the 'target' element whilst the surrounding line segments represent the optimal orientation of neighbouring line segments. The areas within the dashed and dotted lines represent the co and trans-axial fan-out areas respectively, defined by function 6-1. The orientation of surround line segments is determined by function 6-2. The length of the surrounding line elements represents the strength of facilitation, this is determined by the product of functions 6-1, 6-3 and 6-4.

If function 5 is applied to all the positions of Gabor patches within a stimulus image where patch centres and orientations (i, j, θ) are known, then a measure of the CLE of each patch is generated. In order to illustrate the utility of this measure, an image is generated within

which the contrast of each Gabor patch was determined by its CLE (Figure 6-2, right). Clearly, those Gabor patches that had the greatest CLE are the ones that were members of an embedded contour, whilst the background elements have a lower level of co-linearity energy.




Figure 6-2. (Left) An image containing an embedded contour. (Right) this is the same image but here the summed product of function 5 has determined the contrast of each Gabor element.

6.2.2 Procedure for the manipulation of contour membership of an individual line element.

In order to manipulate the CLE of a particular line element the difference between the actual CLE and the desired CLE was minimised. This was achieved by identifying the level of CLE of the target element (indicated by arrows in Figure 6-4) as it was rotated through a range of orientations (circles within Figure 6-3). For the sake of computational efficiency, energy was only calculated at ten discrete orientations and energy between these points was interpolated using a fitted cubic spline within Matlab®, this is shown as a solid line within Figure 6-3. Subsequently, the point along the fitted curve that is nearest to the desired energy level was selected, determining the orientation of the target element. If multiple points exist due to the curve crossing the desired energy level a number of times then a random selection is made from each of the potential orientations. For example, if the goal was to increase the prevalence of contours then the orientation that corresponds to the maximal fitted energy level was identified (point A within Figure 6-3). Alternatively, if a decrease in the prevalence of noise-contours is required then the target element orientation that corresponds to the minimal energy was selected (point B within Figure 6-3).




Figure 6-3. Measures of co-linearity energy strength as the target element is rotated.




Figure 6-4. The left figure indicates the orientation of the target Gabor patch when the target patches co-linearity energy is maximised. The right-hand figure shows the orientation when it is minimised. In both images an arrow indicates the target element.

Rather than simply maximising or minimising the co-linearity energy, a particular energy level could be specified. Subsequently the orientation of the target element will be set to that which gave rise to the desired energy level. By manipulating the orientation of each background element within an image the prevalence of noise-contours could be varied.

6.2.3 Manipulating all element orientations within a background.

The process of manipulating the orientation of patch was repeated for all background elements. The processing of all element orientations once is termed a *sweep*. The order in which each element was manipulated during a sweep was randomised. As the CLE of an

element will be influenced by the orientation of surrounding elements a number of sweeps were made. Pilot manipulations indicated that element orientations tended to settle within five sweeps. Hence, manipulations were terminated after five sweeps.

6.2.3.1 Prevention of 'recruitment' by target-contours

When the CLE of background elements was manipulated, it was found that the orientation of the target-contour elements could influence the orientation of background elements. When the aim of CLE manipulation was to reduce the prevalence of noise-contours, those background elements that surrounded the target-contour tended to adopt the prevailing orientation of the contour (Figure 6-5, middle) i.e. they were recruited by the target-contour. Conversely, when co-linearity energy was reduced the background element orientations diverged from those of the target-contour (Figure 6-5, left). However, the divergence effect was reduced because the orientations of target-contour elements themselves soon diverge from one-another with successive sweeps. In order to prevent the occurrence of recruitment in experimental stimuli, the orientations of the target-contour elements were fixed at randomised orientations during CLE manipulation sweeps. Thus, any recruitment should be reduced (Figure 6-5, right).




Figure 6-5. When the co-linearity energy of an unmodified stimulus image (left) is raised, elements that surround the contour patches are recruited (middle). When the orientations of the contour elements are randomised and then fixed during manipulation sweeps there is a lower incidence of recruitment (right).

6.2.3.2 Measures of element co-linearity energy with proximity to the image edge

The CLE of a target element is based on the summed strength of measured co-linearity between a target and surrounding elements. The number of neighbours that a target element has was not taken into account. This leads to two consequences. (i) In an image featuring randomised background orientations, CLE measures will be lower near to the edges of an image. (ii) When the CLE manipulation algorithm tries to raise the CLE level of background elements, the frequency of noise-contours may be slightly higher towards the edges of an image.

The contour plots within Figure 6-6 illustrate these changes in energy levels across images where element orientations are random (left column) and where they have been manipulated with CLE targets of 0 (middle column) and 1.5 (right column). The data presented in each column represents the mean values of 100 320x320 pixel images with a background generated using the *spaced-fill* technique (appendix 1 and Section 6.2.4) with a minimal spacing distance of 16 pixels. Clearly, for images featuring patches with random orientations (Figure 6-6, left column) there is a decrease in co-linearity energy for elements close to the image boundary. Conversely, for manipulated images the mean-energy levels do not decrease to the same degree.

The majority of relative target/surround positions are likely to have a low CLE energy. Consequently, when CLE levels are lowered the standard deviation is likely to be relatively small (Figure 6-6, middle). Background elements can easily be oriented so that they are co-linear with few, if any, of their neighbours. Conversely, the standard deviation of energy levels is large for images featuring randomly oriented elements (Figure 6-6, middle left) and for images where the CLE levels have been raised (Figure 6-6, middle right). In unmodified backgrounds this is to be expected, as the relative orientations of background elements are randomised. Whereas, in raised CLE images, the large distribution of CLE values reflects the fact that it is generally not possible to orient an element so that it is co-linear with all its neighbours.

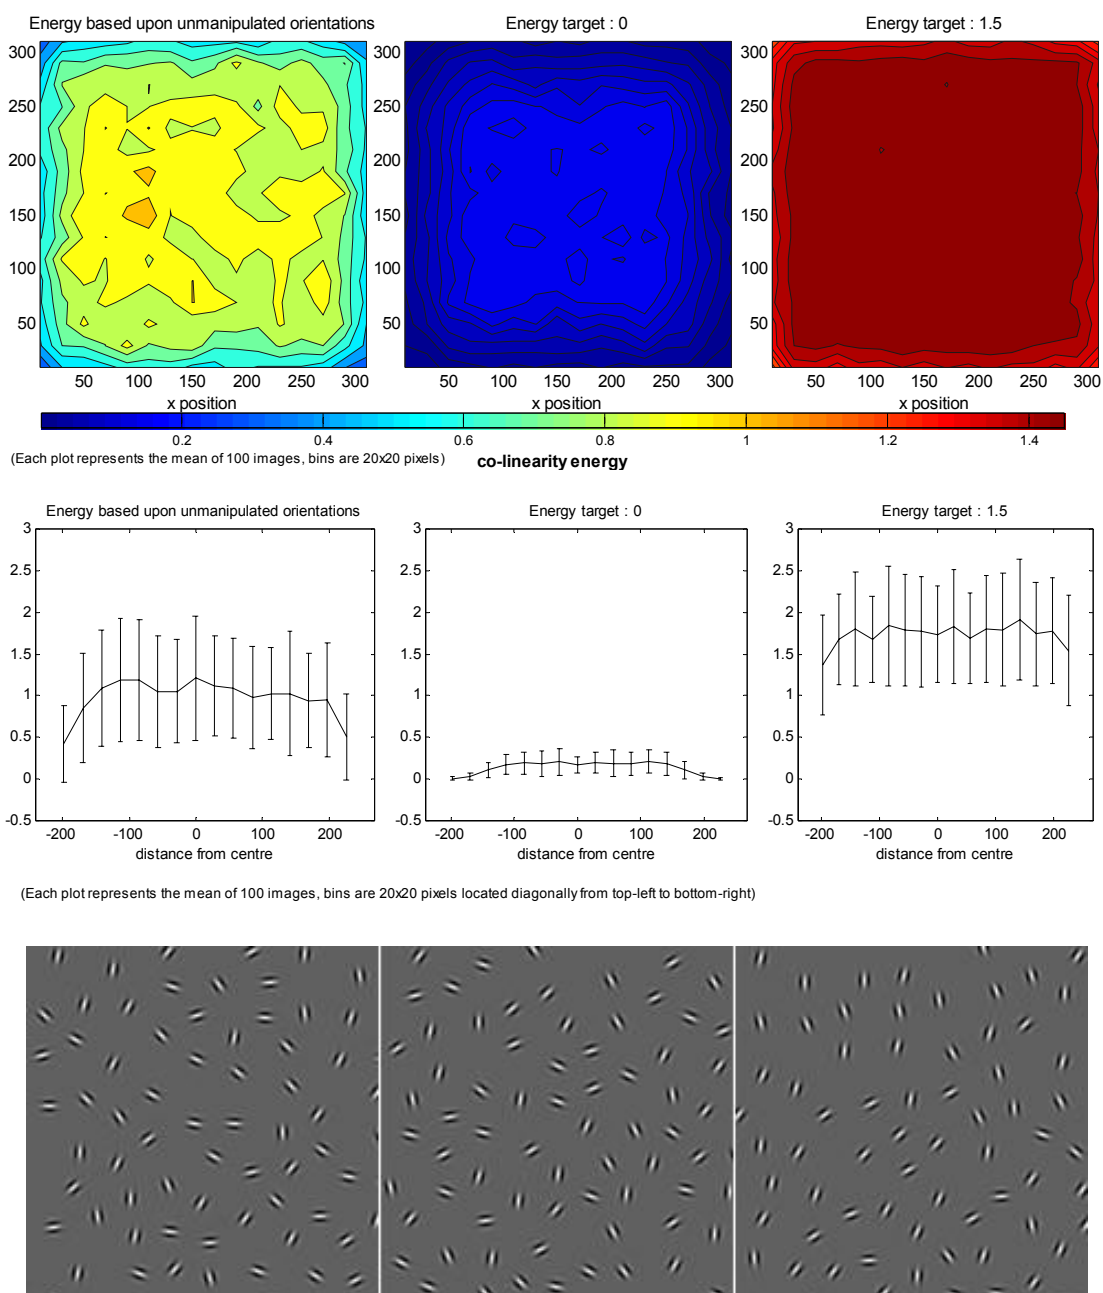


Figure 6-6, Examining mean CLE levels at varying locations across a stimulus image. The three columns correspond to images with different co-linearity energy manipulations. (Column 1) Images were orientations are random. (Column 2) where energy levels are lowered to zero. (Column 3) Energy levels are increased towards 1.5. The upper row illustrates the change in mean co-linearity energy with varying x-y positions within an image. The middle row illustrates the distribution of energy values from the top-left to bottom-right of the images. The bottom row shows the upper left quadrants of example images with matching co-linearity energy targets.

As the CLE measures across a particular image are influenced by the relative location, orientation and the quantity of surrounding elements, measures of this energy should only be considered ordinal in nature. Furthermore, whilst images featuring the same degree of co-linearity energy within a specific experiment will have a similar distribution of noise-contours, images drawn from different experiments may not. The relative size of images

and spacing of background and contour elements will determine the equivalence of the co-linearity energy values between such images.

6.2.4 On the generation of stimulus backgrounds

Unless otherwise stated experiments in earlier chapters utilised a background generation algorithm that was similar in its operation to the method described by Field, Hess and Hayes (1993). This has already been defined as the *jittered-grid* method in Section 2.1.1.2. The overview of this chapter discussed differences between the nearest-neighbour spacing distributions of contour and background elements generated using the jittered-grid method. Braun (1999) suggested that these differences might be a source of bias when estimating detection limits for various contour parameters. In order to minimise this potential confound a new method of generating background element locations was developed. This is henceforth termed the *spaced-fill* method. This method differs from that offered by Braun. The latter method requires that the experimenter should decide when the iterative randomisation of element locations should cease, whereas the *spaced-fill* method does not move elements once their random location has been selected, therefore it is computationally more efficient.

The *spaced-fill* method achieves a filled background by placing Gabor elements within a background until there is no further space for additional elements. The centre of each successive element is randomly selected from remaining empty space within the background. Empty space is defined as the area of the image that remains beyond a specified distance (B_{gap}) from all existing Gabor centres. The resultant spacing for nearest-neighbours has a non-symmetrical distribution, the majority of elements have nearest-neighbours within a distance of $B_{gap}+2$ pixels. No background elements have a separation that is less than B_{gap} . If a target-contour is to be placed within the background then the centres of the contour elements can be added to the background prior to the execution of the iterative fill process. A small constant is subtracted from the B_{gap} value for target-elements (1.219 pixels¹⁰). This prevents the occurrence of a noticeable step in the spacing between the target-contour elements and surrounding background elements. Full details of both the *jittered-grid* and *spaced-fill* background generation methods are provided within appendix A.

Figure 6-7 shows pair-distribution functions for *jittered-grid* and *spaced-fill* images (n=200), the leftmost peaks on the dotted lines represent the nearest-neighbour spacing of

¹⁰ The calculation of this constant is explained within appendix A

contour elements (at 16 pixels), the following peak represents the ‘next nearest-neighbour’ and so on. For the *spaced-fill* background there is a corresponding nearest-neighbour peak, this is not present for the *jittered-grid* background. Hence, any noise-contours within the *jittered-grid* images are likely to have a spacing distribution that differs from that of target-contours.

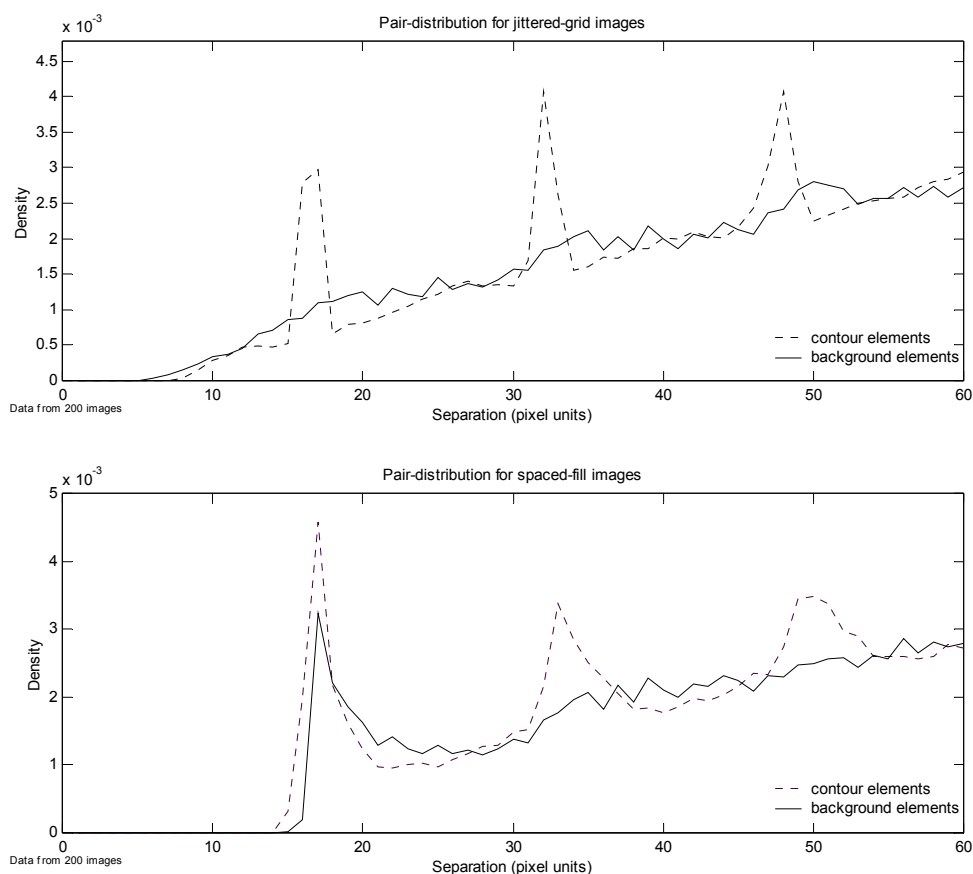


Figure 6-7. Pair-distribution plots for *jittered-grid* (upper) and *spaced-fill* (lower) images. The regular peaks on the dotted-lines correspond to the mean spacing separation for contour elements.

For each set of experimental images, the spacing distribution of nearest-neighbours for contour and background elements was matched. This was achieved by generating a set of 50 backgrounds with the desired spacing parameters. The distance between each element and its nearest neighbour was then stored. Subsequently, the spacing of consecutive contour elements was randomly selected from the stored distribution of background nearest-neighbour separations. It follows then that the distribution of target-contour separations matched the distribution of separations between neighbouring elements within the background (see dashed-line in Figure 6-7, lower plot).

In all the experiments described, a two-alternative forced-choice paradigm was employed. Each target image was paired with a noise image. In order to control for any remaining

density cue, contours were also embedded within the noise images. The single distinguishing feature of noise images is that the orientations of the embedded contour elements were randomised. Therefore, any remaining density cue would have been present within both the target and noise images.

6.3 Experiment 7: The influence of noise-contour prevalence upon the path-angle detection limits.

6.3.1 Overview

The findings of Field, Hayes and Hess (1993) suggested a range of path-angles over which contour integration will occur. Observer performance in a 2AFC paradigm was 75% correct for contours with path-angles of around 45°. As the stimulus backgrounds employed by Field *et al.* featured randomised element orientations, it is not clear to what extent the occurrence of noise-contours has influenced the range of path-angle values that the integration process can tolerate. Field *et al.* claim that range extends to differences of about 60°. However, as already discussed, the spacing distributions of the background and target-contour elements were not appropriately matched in the stimuli created by Field *et al.* In the current experiment a *spaced-fill* background is adopted in order to create a closer match between the spacing distributions of target-contour and background elements.

The CLE of background elements was also manipulated in the current experiment. It was expected that the detection of target-contours would improve, as the CLE was reduced and vice-versa. Such a result would indicate that a significant component of the claimed path-angle range might be determined by the presence or absence of noise-contours. Manipulating background CLE may reveal how much the path-angle detection limit is contaminated by the visibility of noise-contours.

6.3.2 Methodology

6.3.2.1 Stimulus image parameters

The overall size of stimulus images was 320x320 pixels. There were five discrete levels of co-linearity energy and nine levels of contour path-angle, thus there were 45 combinations of target image parameters and 45 corresponding levels of noise stimuli. Ten images were generated for each of these levels. In all images, Gabor patches were phase-alternated, thus ensuring, according to criteria suggested by Hess and Dakin (1999), that the task was not being solved by a coarse-scale mechanism. However, as indicated by the modelling results that were presented in Chapter 3, it is not certain that this manipulation of Gabor element phase actually prevents the operation of coarse-scale mechanisms. Prior to the

manipulation of CLE, the mean level of CLE for stimulus images was measured. The CLE of images with random backgrounds was approximately 1.0.

Backgrounds were generated using the *spaced-fill* technique (Section 6.2.4 and Appendix A). The minimal background element spacing constant (b_{gap}) was fixed at 16 pixels, which corresponds to an element gap of 4λ . The target co-linearity energy of backgrounds was varied from 0-2 in steps of 0.5. The overall structure of noise images was identical to that of target images, except that the orientations of target-contour elements were randomised.

Pilot studies revealed a pronounced ceiling effect when contours featured more than eight elements. Therefore, each target-contour was composed of eight elements. The path-angle between consecutive contour elements was varied from 10-90° in discrete steps of 10°. An additional path-angle *jitter* of $\pm 5^\circ$ was applied to the path-angle between each contour element. A rotational *jitter* of $\pm 5^\circ$ was applied to the orientation of each element. The spacing of contour elements was matched to the spacing of background elements using the technique described within sub-section 6.2.4. The position of each contour was random except for the constraint that its centre was 69 pixels from the centre of the image. This distance corresponds to an eccentricity of 3° visual-angle with the viewing distance of 500mm.

6.3.2.2 *Observers*

Two observers participated. One observer was the experimenter (PGL) and another (KAS) was familiar with contour integration experiments but naïve to the hypothesis underlying the current experiment. Both have corrected to normal vision.

6.3.2.3 *Procedure*

Participants undertook observation sessions that consisted of 200 2AFC decisions. Typically, a session lasted 30 minutes. The presentation sequencing of target and noise image pairs was determined using the method described within Section 2.1.2. The target and noise images in a 2AFC pair were always matched in their stimulus level. Images were presented for 250ms. Full details of fixation point and stimuli and presentation timing are illustrated within Figure 6-8.




Figure 6-8. Stimulus timing details. Elevated sections of each line represent periods when the indicated items are visible to the observer.

Stimulus scale properties are given below in pixel, λ and visual-angle units (Table 6-1).

Table 6-1. Stimulus scale properties for the current experiment. Visual angle units correspond to the viewing distance of 500mm.

	Pixels	Lambda	Visual-angle
Gabor patch wavelength	4	1λ	0.173°
Background and contour minimal nearest-neighbour spacing	16	4λ	0.69°
Contour centre offset, relative to fixation point	69	17.25λ	3°
Overall image size	320 x 320	$80 \times 80\lambda$	$14.11 \times 14.11^\circ$

Stimuli were displayed on a Panasonic S100 monitor using the Cambridge VSG display system. Details of the monitor parameters have been provided in section 2.1.3.

6.3.3 Results

Observer performance is presented as a psychometric function of contour path-angle and CLE (Figure 6-9).





Figure 6-9. Psychophysical plots of observer performance. (Upper plot) detection rates for observer KAS. (Lower plot) detection rates for observer PGL.

Figure 6-10 shows the fitted psychometric functions for subjects KAS (top) and PGL (bottom), the creation of these is explained in Section 2.1.4.1. The actual detection rates of observers are also shown as symbols within the same plots.

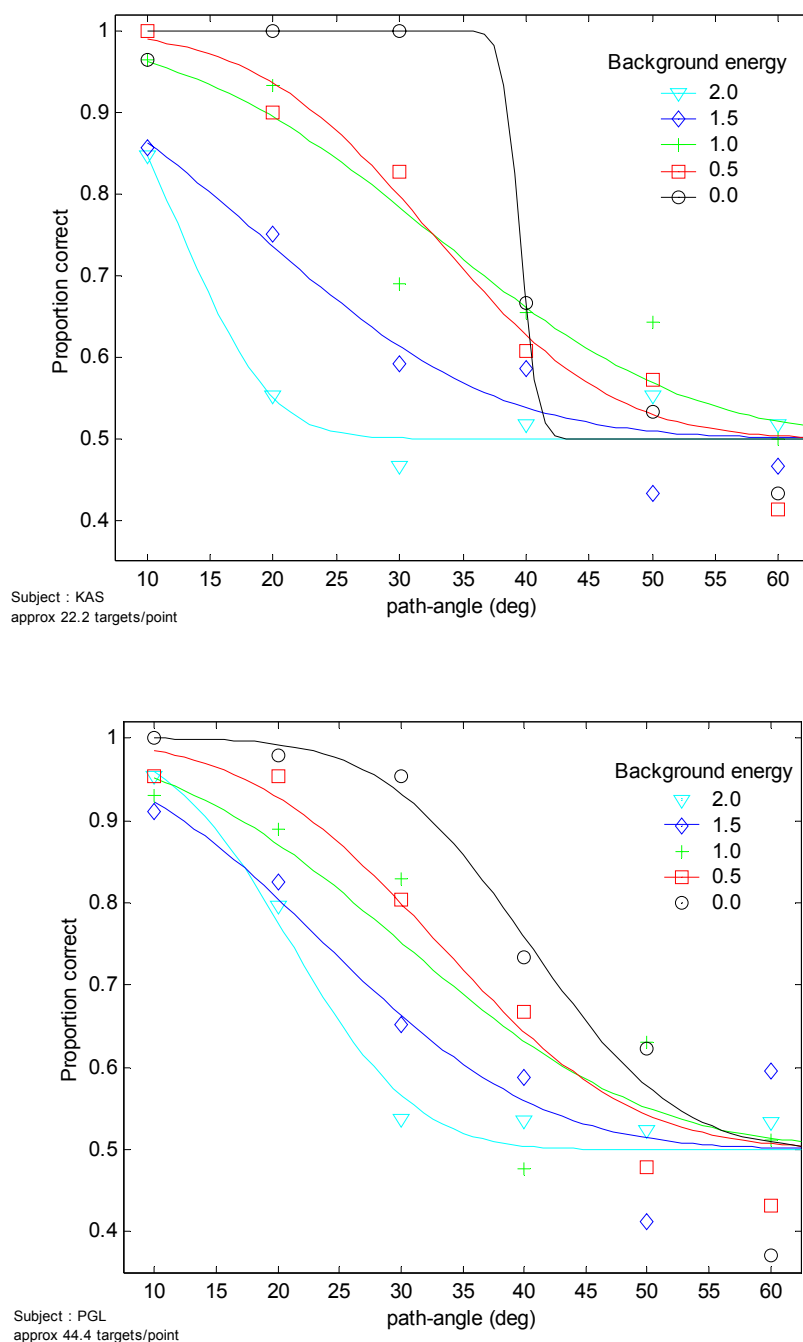


Figure 6-10. Fitted psychometric functions for performance in the path-angle experiment with varying levels of background co-linearity energy. Error-bars are omitted in order to maintain clarity. The variation around each point is shown in Figure 6-9. (Upper plot) performance for observer KAS. (Lower plot) performance for observer PGL.

The influence of CLE upon the detectability of target-contours is clearly illustrated when the value of path-angle that leads to a detection rate of 75% in the psychometric fits is plotted as a function of CLE (Figure 6-11). The level of path-angle that led to a performance level of 75% correct varied as the CLE of the background elements was manipulated. There was an inverse correlation between CLE and the path-angle value that led to a 75% detection rate. Increasing CLE reduces the detectability of the target-contour.





Figure 6-11. Plots illustrating the values of path-angle and CLE at the 75% points for the fitted psychometric functions. The 75% detection limit for contour path-angle is clearly lowered when the visibility of noise-contours is increased, and *vice-versa*.

6.3.3.1 Contour detection rates are influenced by changes in CLE

The χ^2 statistic was calculated for the detection rates of target-images with different levels of CLE. Using the criteria defined in Section 2.1.4.2 it was found that there were significant differences between the majority of CLE levels over the path-angle range 10-30° for observer KAS, whilst for PGL the maximum number of significant differences were present at path-angles of 20° and 30°. For both observers ceiling and floor effects tended to mask differences beyond the ranges identified.

Table 6-2. Results of χ^2 tests examining differences in the contour detection rates with manipulations of CLE. The upper table presents results for subject KAS and the lower for subject PGL. For KAS values represent the path-angles of 10, 20 and 30° and for PGL values represent the path-angles of 20 and 30°.

<i>CLE</i>	<i>0</i>	<i>0.5</i>	<i>1</i>	<i>1.5</i>	<i>2</i>
<i>0</i>	*	5.410	9.659	22.802	37.647
<i>0.5</i>	0.020	*	0.954	9.152	21.487
<i>1</i>	0.002	0.329	*	4.441	14.252
<i>1.5</i>	< 0.000	0.002	0.035	*	2.893
<i>2</i>	< 0.000	< 0.000	< 0.000	0.089	*

<i>CLE</i>	<i>0</i>	<i>0.5</i>	<i>1</i>	<i>1.5</i>	<i>2</i>
<i>0</i>	*	4.957	6.362	18.643	26.290
<i>0.5</i>	0.026	*	0.116	5.630	10.835
<i>1</i>	0.012	0.733	*	4.056	8.601
<i>1.5</i>	< 0.000	0.018	0.044	*	1.000
<i>2</i>	< 0.000	0.001	0.003	0.317	*

(Upper triangles represent the χ^2 values and the lower triangles p-values. Statistically significant values are highlighted in **bold**).

6.3.4 Discussion

The fitted 75% path-angle detection limit was inversely correlated with CLE. Target-contours with larger path-angles are detected with greater frequency when CLE was reduced.

As the CLE energy of unmodified backgrounds was close to 1.0 (see Appendix I), it follows then that experiments testing contour detectability as a function of path-angle may have underestimated the angular range within which contour elements are integrated. In the current experiment, with a CLE of 1.0 target-contours were detected at a rate of 75% when path-angle is 30°. When CLE was minimised path-angle was 40° for contours detected at 75%. This suggests an increase in detection rates of 10% relative to the detectability of images featuring unmodified backgrounds. However, the distribution of CLE energy in unmodified backgrounds is less homogenous so this performance prediction may be unfounded. Observers that participated in pilot studies that compared target-contour detection rates for modified and unmodified backgrounds detected contours at higher rates when CLE was minimised¹¹, supporting this prediction.

If it is accepted that some form of association field is necessary for target-contour detection then the differences in orientation of receptive fields over which such integration might occur may be larger than previously reported.

6.4 Experiment 8: The influence of CLE upon the detection of short contours

6.4.1 Overview

Braun (1999) reports that a contour length of greater than six elements is required in order for observers to achieve a detection rate of 75% correct when the element spacing of the background and target-contour were matched. This result is based upon measures of contour detection performance when the orientation of background elements was randomised. The aim of the current experiment was to establish whether shorter contours might have been detected more frequently if the visibility of noise-contours was decreased. Contours composed of four to seven elements were embedded in a background with a matched spacing distribution. The co-linearity energy of the background was varied between 0.0 and 1.5. It was expected that detection rates for target-contours would be increased when CLE was reduced.

There are some differences between the target-contour properties in the current experiment and those reported by Braun. It is difficult to indicate exactly the path-angle for the contours generated by Braun. This is because his contours are generated by placing elements along the circumference of a deformed polygon with 21-24 sides. The average path-angle value must have been between 15-17°. The shorter contours utilised by Braun

¹¹ Section 7.3.3.1 presents results that support this assertion.

were achieved by deleting surplus elements from the longer contours. In the current experiment a path-angle of 20° was employed, thus it was greater than the average used by Braun and at a level that ought to preclude detection by a coarse-scale mechanism¹² (Dakin, personal communication).

6.4.2 Methodology

6.4.2.1 Stimulus images

The overall size of all images was 220 (height) x 800 (width) pixels. Images had a large width because they were paired with stimuli utilised in a further experiment, which featured contours that were placed a large distance from the centre of the image¹³. As with the previous experiment (Section 6.3), both the target and noise images featured an embedded contour, which controlled for any residual density cue.

Backgrounds were generated using the *spaced-fill* technique. The minimal background element spacing constant (b_{gap}) was fixed at 16 pixels. The CLE of the background elements was manipulated through five discrete levels from 0 – 1.5 in steps of 0.375.

All embedded contours were composed of 4-8 Gabor elements. The path-angle between consecutive contour elements was fixed at $20^\circ \pm 5^\circ$. A rotational *jitter* of $\pm 5^\circ$ was applied to the orientation of each element. The position of each contour was random except for the constraint that its centre was located within 69 pixels from the centre of the image, this corresponds to an eccentricity of 3° visual-angle when viewed from a distance of 500mm. Twenty images were constructed for each combination of CLE and contour length level and noise/target status. Gabor elements within all images were phase-alternated.

6.4.2.2 Observers

The same observers participated as described previously (Experiment 7). PGL undertook at least two observation sessions for each target-contour length. KAS undertook one session for each of the four lengths of target-contours.

6.4.2.3 Procedure

Participants undertook a number of observation sessions, which consisted of 200 paired target/noise image pairs that were presented within a 2AFC paradigm. Typically, a session lasted 30 minutes. Within a particular 30-minute session, all embedded contours featured the same contour length. Prior to the onset of any one session, the target-contour length

¹² Experiments described in the current chapter were actually conducted prior to the modelling experiments described in Chapter 3. The modelling experiments reveal that this path-angle value is actually close to the optimal for the detection of phase-alternated contours by a coarse-scale mechanism.

was randomly selected. As contour lengths were randomly selected, it was not expected that the detection rates for contour of different lengths would be confounded by practice. Target images to be presented during an individual observation session are selected randomly - without replacement - from all relevant stimulus levels. Subsequently, matching noise images were selected. Stimulus images were presented for 250ms following presentation of a fixation point – full details of presentation timing are presented in Section 2.1.2.

Images were viewed from a distance of 500mm, scale properties for the stimuli displayed in the current experiment match those for the previous experiment (Experiment 7) thus scale values in Table 6-1 apply. Stimuli were displayed on a Panasonic S100 monitor using the Cambridge VSG display system. Details of the monitor parameters have been provided in section 2.1.3.

6.4.3 Results

Figure 6-12 plots contour detection rates as a function of contour length and CLE. Predictably, contour detection rates were correlated with length; contours that were composed of eight elements were detected more frequently than those with four elements. The CLE of images also had an influence upon the detectability of target-contours. When co-linearity energy is reduced even short contours became detectable at rates above 75%. For observer PGL this was the case even with four element contours, whilst for observer KAS the best performance achieved with contours featuring four elements was 69%.

¹³ This experiment is described in chapter 7.

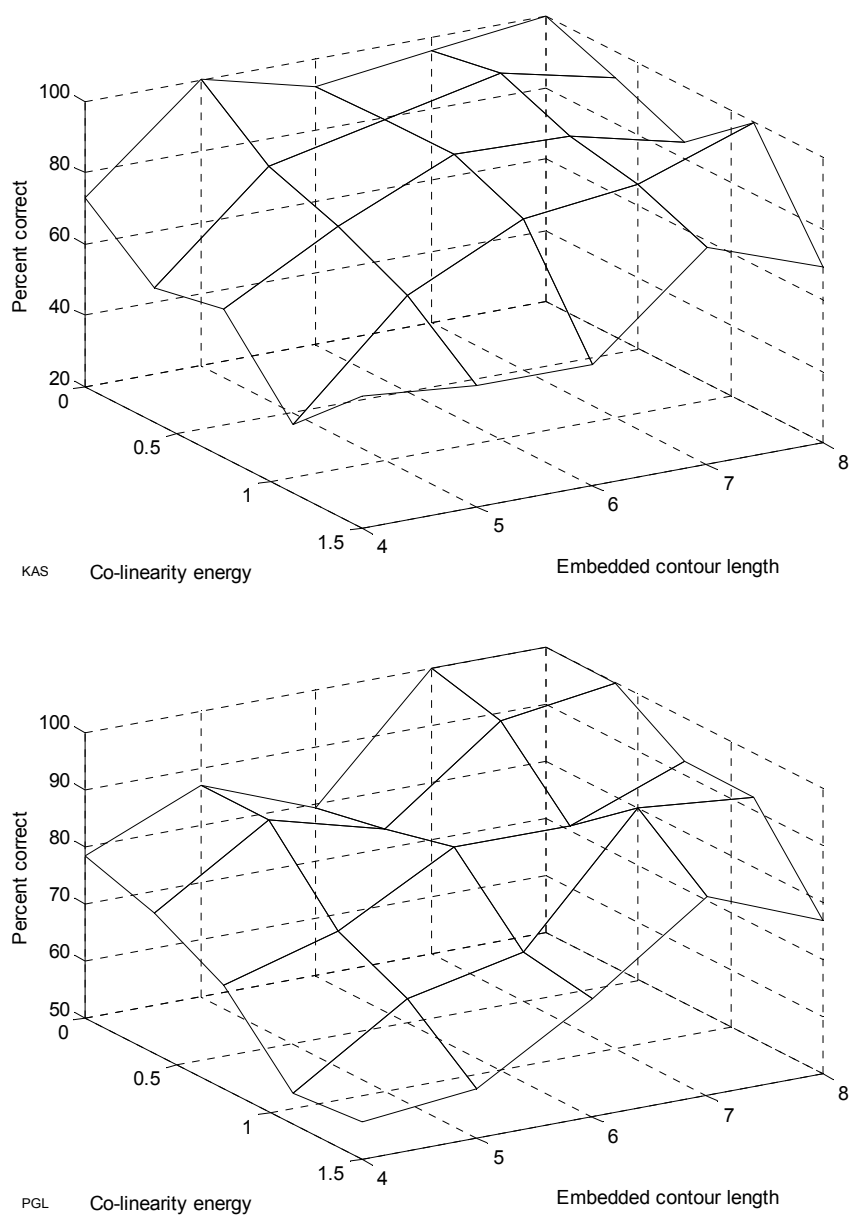


Figure 6-12. Lowering the CLE increases the detection rates for short target-contours. The upper plot presents data for observer KAS. The lower plot presents data for observer PGL.

The point at which subject performance reached 75% correct is plotted as a function of the psychometric fits of contour length and CLE (Figure 6-13). This plot illustrates that observers were able to detect shorter contours when CLE was reduced.




Figure 6-13. Fitted 75% detection limits as a function of contour length and CLE.

6.4.3.1 *Examining the differences in the detection rates for contours as CLE is varied*

The main aim of this experiment was to test the hypothesis that manipulating CLE would influence contour detection rates. A χ^2 analysis was used to examine whether there was a significant difference in the detection rates for contours embedded in backgrounds with different CLE levels. The χ^2 values and p -values are presented in Table 6-3. The χ^2 analyses presented represent comparisons of contour detection rates for all contour lengths (4-8 elements).

Table 6-3. Results of χ^2 tests examining differences in the contour detection rates with manipulations of CLE. The upper table presents results for subject KAS and the lower for subject PGL.

CLE	0	0.375	0.75	1.125	1.5
0	*	1.197	2.323	10.569	13.805
0.375	0.274	*	0.196	4.991	7.396
0.75	0.127	0.658	*	3.228	5.225
1.125	0.001	0.025	0.072	*	0.249
1.5	< 0.000	0.007	0.022	0.618	*

CLE	0	0.375	0.75	1.125	1.5
0	*	0.071	4.026	10.152	16.468
0.375	0.790	*	3.070	8.658	14.621
0.75	0.045	0.080	*	1.449	4.449
1.125	0.001	0.003	0.229	*	0.843
1.5	< 0.000	< 0.000	0.035	0.359	*

(Upper triangles represent the χ^2 values and the lower triangles p -values. Statistically significant values are highlighted in **bold**).

With only a single exception, it was found that where the difference in co-linearity energy was 0.75 or more, then there was a significant difference in the detection rates of target-contours. For subject KAS there was no significant difference in contour detectability for CLE values of 0 & 0.75.

6.4.3.2 Examining the differences in the detection rates for contours of different lengths

The detection rates for target-contours of different lengths were compared using the χ^2 analysis detailed in Section 2.1.4.2. This analysis examined detection rates across all levels of CLE.

Table 6-4. Results of χ^2 tests examining differences in the contour detection rates with manipulations of contour length. The upper table presents results for observer KAS and the lower for observer PGL.

Contour length	4	5	6	7
4	*	7.888	13.613	41.290
5	0.005	*	0.764	7.851
6	< 0.000	0.382	*	3.021
7	< 0.000	0.005	0.082	*

Contour length	4	5	6	7
4	*	4.036	6.086	36.537
5	0.045	*	0.433	19.674
6	0.014	0.510	*	13.548
7	< 0.000	< 0.000	< 0.000	*

*(Upper triangles represent the χ^2 values and the lower triangles p-values. Statistically significant differences are highlighted in **bold**).*

There is clearly a strong effect of contour length. For both observers contours were detected at significantly different rates when the lengths differed by two or more elements. For differences in length of only a single element, the results were less reliable, with only 1 of 3 differences significant for KAS and 2 of 3 for PGL.

6.4.4 Discussion

Detection-rates for the short target-contours were indeed increased when the CLE of background elements was reduced. This suggests the estimates of contour length detection limits are influenced by properties of the stimulus background, even when the densities of contour and background elements were carefully matched.

The length of contours was positively correlated with their rates of detection. This finding is consistent with previous studies (Braun, 1999). In the current study, significant differences in contour detection rates were only found when lengths differed by at least two elements. Had more trials been undertaken by observers then it appears likely that significant differences would have been found for length differences of only a single element.

One might enquire why detection rates were not higher when CLE levels were minimised. It could be argued that the co-alignment of only 2-3 elements ought to be sufficient to signal the presence of a contour. However, a number of additional stimulus properties might have led to reduced performance levels. Firstly, contours were randomly presented at locations between $\pm 13^\circ$ (visual angle). Consequently, attention was necessarily distributed over a relatively large area. Secondly, contours in this experiment featured a path-angle of 20° with an additional jitter of $\pm 5^\circ$. Therefore, contour properties and the increased attentional demands might have led to reductions in performance.

6.5 General Discussion

Manipulations of contour path-angle (Experiment 7) and contour length (Experiment 8) in conjunction with CLE demonstrated that there is an inverse relationship between contour detection rates and CLE. The presence of noise-contours reduces the detection rate for target-contours. Where the prevalence of noise-contours was increased, it is unsurprising that detection rates fall. Observers are required to detect embedded target-contours. Clearly, their task is made harder when additional contours are placed within both target and noise stimuli. When the prevalence of noise-contours was reduced contour detection rates improved. This result suggests that the geometrical limits of contour integration may have been underestimated in previous studies. Manipulations of contour path-angle in conjunction with CLE have demonstrated that contours with greater path-angle values can be detected when CLE is reduced.

The salience of short contours is also influenced by the existence of noise-contours within the stimulus background. Reducing the prevalence of noise-contours increases the detection rate for contours of 4-5 elements in length. This suggests that the Braun (1999) estimate of six elements was an overestimate. Uttal (1975) has reported that straight contours formed from just three dots could still be detected at rates of 65-75% even with very brief exposures of 3ms. However, for these stimuli the relative densities of target and noise elements were not matched. In the current experiment contours as short as four elements are detected at similar rates even when density cues are carefully controlled. With element density controlled, only the relative position and orientations of elements indicate their presence. Therefore, the mechanism responsible for their detection must be sensitive to element co-alignment.

Regardless of which model of contour detection is correct, it appears likely that two processes are involved in the detection of embedded contours. The first process involves

the integration of those elements that could potentially form parts of contours. The second process involves the selection of the longest contour from amongst all those that have been identified. When the CLE of background elements is minimised, the task of selection is made much easier, effectively isolating the process of integration from the process of selection. Therefore, by reducing the CLE of background elements to a minimum, one can achieve a more accurate estimate of the range of spacing and orientation differences that limit contour integration.

-0-0-0-

7 Does the same process underlie contour detection in the fovea and in the para-fovea?

7.1 Overview

Hess and Dakin (1997, 1999) investigated contour detection rates for para-foveal areas. They found that for eccentricities beyond 10°, contour detection was specifically impaired for contours that were formed from phase-alternated Gabor patches. They also reported a corresponding pattern of impairment for their simple-filter model. This led them to conclude that whilst both association-field and coarse-scale mechanisms were available for foveal contour detection, only coarse-scale mechanisms underlie contour detection in the periphery. Table 7-1 summarises the results of Hess and Dakin (1997, 1999) and includes relevant findings reported in Chapters 3, 4 and 5 of the current thesis.

Table 7-1. Experimental results from various sources indicate performance for fovea, periphery and simple-filter models. The ticks (✓) indicate successful detection and crosses (✗) indicate poor detection of contours. The numbers/letters in brackets relate to the source of the experimental result.

Contour properties	Fovea	Para-fovea	Simple-filter model (Hess and Dakin)	Simple-filter model (This thesis)
Elements phase-aligned	✓ (A,B)	✓ (B)	✓ (B)	✓ (3)
Elements phase-alternated	✓ (B)	✗ (B)	✗ (B)	✓ (3)
Smooth contours	✓ (4)	Untested	Untested	✗ (5)

(A. Field, Hayes and Hess (1993); B. Hess and Dakin (1997, 1999); Numbers 3, 4 and 5 refer to chapters in the current thesis)

Chapter 3 has shown that when constraints upon the selection of filter-scale were relaxed, the simple-filter model was able to detect phase-alternated contours. Whilst detection rates for these contours were lower than they were for phase-aligned contours, they were close to the detection rates achieved by human observers for foveally presented phase-aligned contours (see Figure 3-6, this thesis). However, this evidence does not weaken the case for a coarse-scale account of contour detection in the periphery. Nevertheless, it does undermine one of the reasons offered by Hess and Dakin for suggesting that coarse-scale mechanisms might underlie para-foveal contour detection, i.e. that both the human and model were impaired when presented with phase-alternated contours.

In Chapter 5, it was shown that the simple-filter model performed badly when detecting smooth contours. This was because ZBRs were generally shorter when embedded contours featured fewer inflections. If a simple-filter mechanism underlies peripheral-field contour detection, then performance should be correlated with the frequency of inflections within contours. This predicted pattern of results is the reverse of that reported for the fovea (Section 4.3). For contours presented to the fovea, detection rates increased for contours that featured fewer inflections.

Two experiments are described in the current chapter. The first experiment investigated whether the reduction of co-linearity energy (*CLE* has been defined in Section 6.2.1) led to improved contour detection rates in para-foveal areas. Chapter 6 reported that reductions of *CLE* led to improvements in the detection rates for short contours presented to the fovea. It was hypothesised that reducing *CLE* may facilitate the measurement of para-foveal contour detection abilities that were previously masked by the presence of noise-contours. A second experiment investigates the influence of global contour structure upon contour detection rates in para-foveal areas. According to Hess and Dakin (1999) peripheral performance should conform to that predicted by the simple-filter model. Consequently, smoother contours should be least detectable. A reversal of this relationship would undermine arguments that peripheral contour detection was achieved via mechanisms analogous to the simple-filter model.

7.2 Experiment 9: Reducing 'co-linearity energy' in the background leads to improved contour detection in the periphery

7.2.1 Introduction

This experiment investigated whether the detection of phase-alternated contours would improve when the occurrence of noise-contours within the stimulus background was manipulated. Experiment 8, described in Chapter 6, has shown that the detection rates for short contours presented to the fovea were improved when *CLE* was lowered. That is, shorter target-contours were detected more frequently when the prevalence of noise-contours was reduced. Hess and Dakin (1997, 1999) have shown that the detection of phase-alternated contours was extremely poor when contours were located beyond 10°. The current experiment was designed to test whether phase-alternated contours might only be detected in the periphery when *CLE* levels were reduced. It was hypothesised that, where background element orientations were randomised, these processes may have been

too weak to have had a measurable influence upon performance in the path-paradigm experiments of Hess and Dakin (1997, 1999).

Contours were located at various eccentricities ($\pm 3, 8$ and 13°) facilitating a comparison of contour detection rates within and beyond the fovea. Contour lengths were varied between 4 and 8 elements. The spatial-frequency and overall scale of stimuli was not varied with greater degrees of eccentricity. Hence, there is a possibility that any performance fall-off with increasing contour eccentricity might have been due to a mismatch between the spatial-frequency of stimuli and para-foveal acuity. Therefore, whilst comparisons at each level of eccentricity are valid, comparisons that are made between detection rates at different levels of eccentricity should only be accepted with caution. In order to enable a comparison in the detectability of phase-aligned and phase-alternated contours, stimuli featuring contours that were five elements in length were also generated with phase-aligned Gabor patches.

7.2.2 Methodology

7.2.2.1 Stimulus parameters

Images have overall dimensions of 220x800 pixels (height x width), this ‘letterbox’ format was adopted as contours were to be presented across a relatively large range of eccentricities, whilst the fixation point was located within the centre of all images.

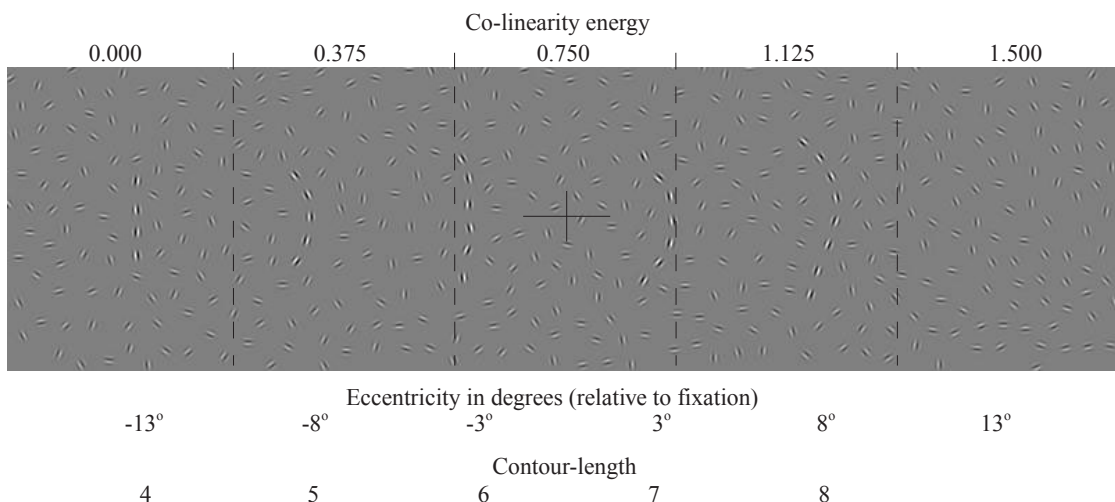


Figure 7-1. An illustration of the three stimulus parameters manipulated in the current experiment. These are *eccentricity* the distance (in visual angle) of the median contour x value from the fixation point; *Co-linearity energy*, the measured occurrence of co-linearity between background elements; *Contour-length*, the number of Gabor elements that form the embedded contour. For illustrative purposes, the contrast of the background elements has been reduced in order to ease the detection of the embedded contours.

Three independent variables were manipulated within this experiment: Contour length, contour location and background element CLE. Figure 7-1 illustrates each of these manipulations. The lengths of target contours were varied between 4 and 8 elements. The path-angle between consecutive contour elements was 20° with an additional path-angle jitter of $\pm 5^\circ$. Each contour element was co-aligned with the contour path and an additional random orientation jitter of $\pm 5^\circ$ was applied. The spacing of consecutive contour elements was matched to the spacing between nearest neighbours within the stimulus background (Section 6.2.4). The horizontal positions of contours were manipulated such that the mean x-offset relative to the fixation point was 3, 8 or 13° . For each combination of contour-length, CLE and eccentricity levels, half of the contours were embedded to the left of the fixation point and half were embedded to the right of the fixation point. Contours were vertically oriented in order to minimise the possibility that manipulations of eccentricity might have been confounded by manipulations of contour-length.

The *spaced-fill* technique (Section 6.2.4) was used to determine the positions of background elements. The minimal background element spacing constant (b_{gap}) of 16 pixels, was matched to the spacing of consecutive contour elements. The level of co-linearity energy (CLE) of background elements was manipulated using the technique described in Section 6.2. CLE was varied between 0 and 1.5 in steps of 0.375.

Twenty target images were created for each combination of contour-length, CLE level and eccentricity. Twenty corresponding noise images were also created using the same parameters, except that the orientations of the target-contour elements were randomised. Consequently, any residual density cue that may have indicated the location of the target-contour was also present within noise images.

Contour elements were phase-alternated. The phase of gabor patches in the background were randomised – randomly selected values of 0 or 180° were assigned to ϕ (see Section 2.1.1.3 for details). In order to examine the influence of Gabor patch phase upon contour salience, an additional set of images was created which were identical to those that featured 5 element target-contours, except that Gabor patches within these stimuli were phase-aligned. For these images the background elements were cosinusoidal, ie. ϕ values were fixed at 90° .

7.2.2.2 *Procedure*

Each observer undertook a number of observation sessions, which consisted of 200 paired target/noise image pairs, presented within a 2AFC paradigm. During an observation session CLE and eccentricity were manipulated while contour length was held constant. Contour length was selected randomly prior to the onset of each session. A post-hoc examination of observation data revealed that observers initially viewed stimuli featuring shorter contours, whilst longer contours were viewed towards the end of the experiment. Each observer undertook two observation sessions for each level of contour-length, with the exception that observer KAS only participated in one observation session each for contours of 6 and 8 elements. Observer PGL participated in two observation sessions for each contour length, except for length 8, only one session was undertaken for this level. The sequencing of images presented during an individual trial conformed to the procedure described in Section 2.1.2.

As with all the proceeding experiments, images were visible for 250ms. Observers were required to fixate upon a cross that was displayed prior to the stimulus image onset. The location of the cross corresponded to the centre of the image. Observers were advised not to move their eyes whilst stimuli were visible.

Images were viewed from a distance of 500mm, Table 7-2 below indicates some of the key stimulus parameters in pixel, lambda and visual-angle units. Stimuli were presented on a Cambridge Research VSG system (VSG2/4-4MB) and displayed with a Panasonic S110 20" monitor, further details of the display system have been given in Section 2.1.3.

Table 7-2. Spatial-scale properties of the stimuli utilised in the current experiment.

	Pixels	Lambda	Visual-angle
Gabor patch full-wave frequency	4	1	0.173°
Background and contour minimal nearest-neighbour spacing	16	4	0.69°
Contour offset's, relative to fixation point	69, 186, 306	17.25, 46, 74	3°, 8°, 13°
Overall image size	220 x 800	55 x 200λ	9.6 x 39.5°

7.2.2.3 *Observer details*

Two observers participated. One observer was the experimenter (PGL) and another (KAS) was familiar with contour integration experiments but naïve to the hypothesis underlying the current experiment. Both have corrected to normal vision.

7.2.3 Results

Figure 7-2 shows the contour detection rates of observers as a function of contour-length, CLE and distance from fixation. Regardless of the distance between the fixation point and the location of the contour, detection rates varied as a function of contour length and as an inverse function of CLE. Detection rates fell as a function of the distance from fixation. At the largest offset (13°) detection rates only approached 75% correct when CLE was reduced to a minimum and contours were at least eight elements in length.

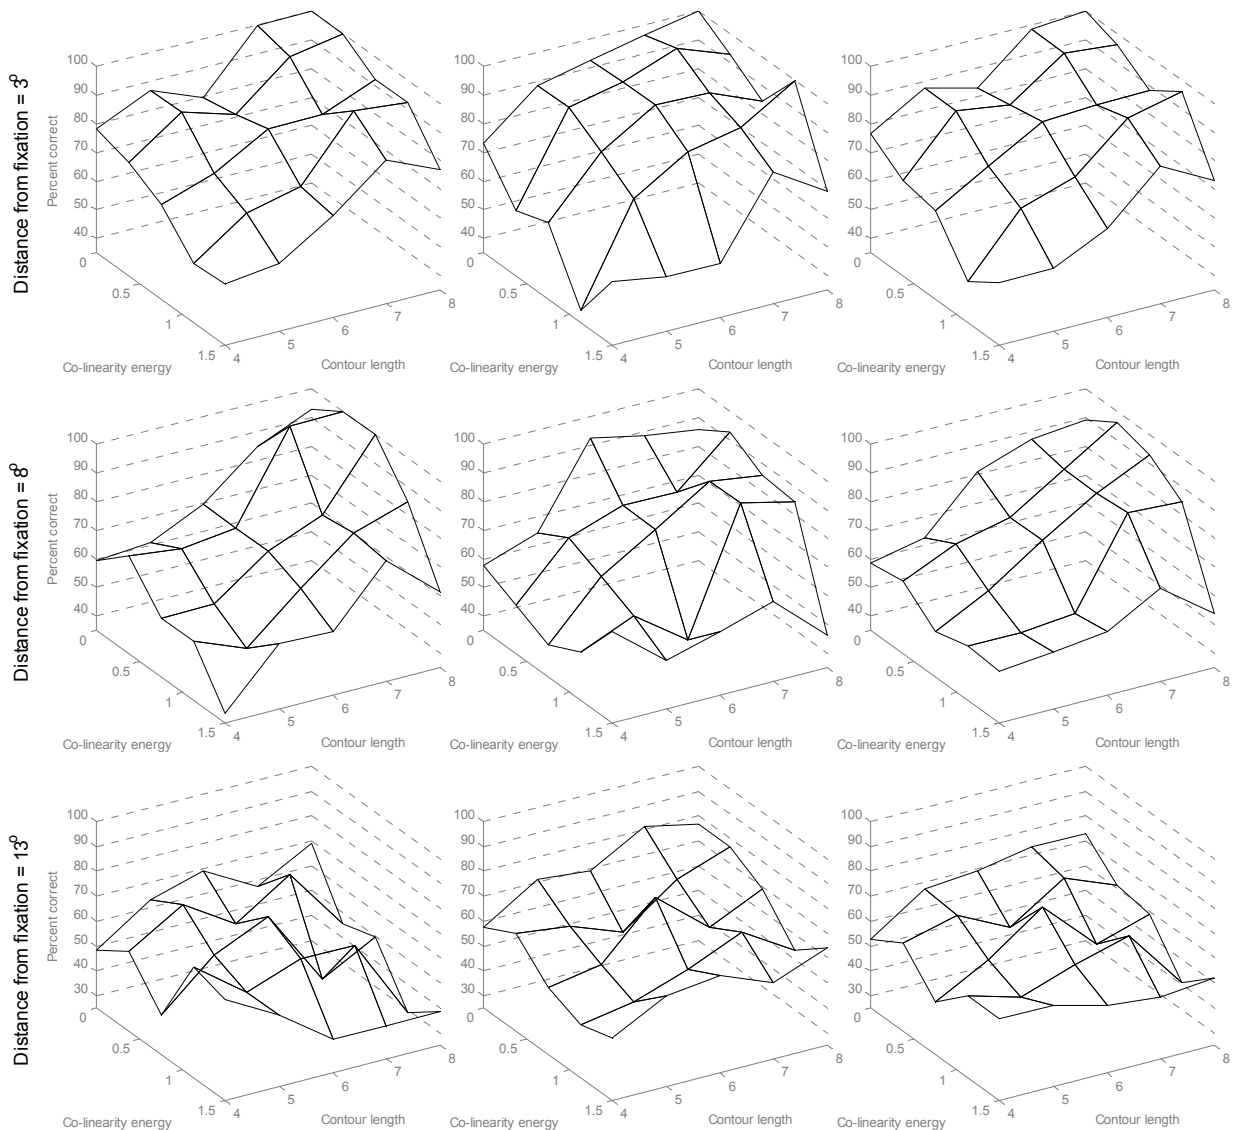


Figure 7-2. Contour detection rates as a function of length, distance from fixation and CLE. Three rows of surface plots show contour detection rates at ± 3 , 8 and 13° distance from fixation. The columns represent the detection rates for observers PGL (left) and KAS (middle) and the rightmost column shows the mean detection rates for both observers.

Detection rates for contours located at a distance of 3° and 8° from the fixation point fell as a function of co-linearity energy and as an inverse function of contour length. Detection

rates for contours located at eccentricities of 13° were less consistent. Detection rates for observer KAS are largely consistent with those recorded for the fovea, though performance was 10-20% lower. For observer PGL however, contour detection rates did not vary consistently as a function of contour length. In order to more clearly illustrate how contour length influenced detection rates at 13° , the mean detection rate for each contour length over all levels of CLE was calculated. These means are plotted in Figure 7-3 (top row), the corresponding plot of the influence of CLE, over all contour lengths, is presented in Figure 7-3 (bottom row).

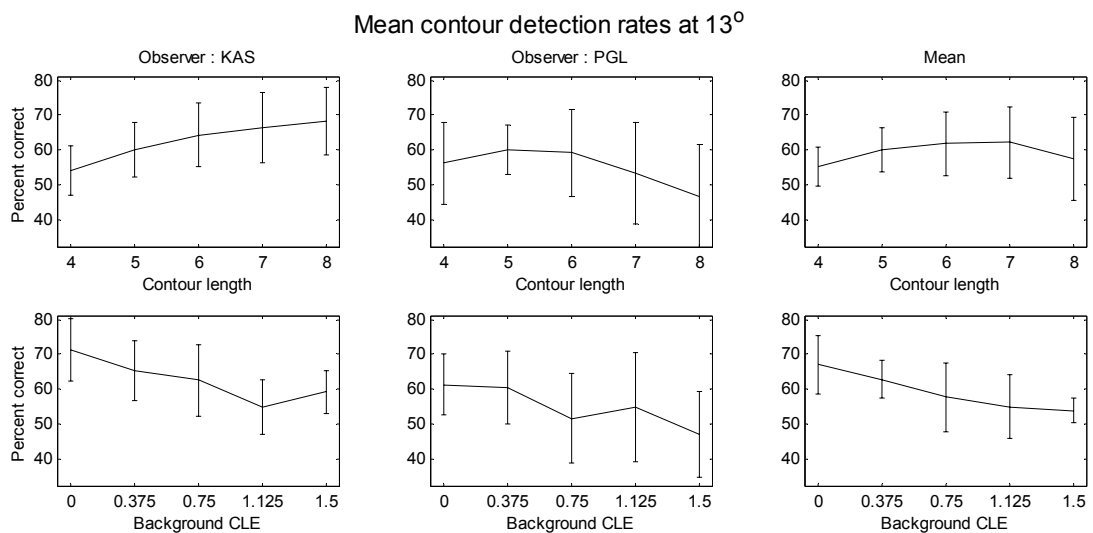


Figure 7-3. (*Top row*) Mean contour detection rates as a function of contour-length, each point represents the mean over all levels of CLE. (*Bottom row*) contour detection rates as a function of CLE, each point represents the mean over all levels of contour-length. Error bars show ± 1 standard deviation.

A visual inspection of Figure 7-3 confirms that detection rates for both observers fell as a function of CLE. However, contour detection rates increased as a function of contour length for observers KAS, whilst they fell as a function of contour length for observer PGL – possible reasons for this discrepancy are discussed in Section 7.2.4.

7.2.3.1 *Analysis of variance of observer performance*

The minimal adequate model was achieved by iteratively deleting non-significant terms. Stepwise deletion of non-significant terms from the minimal adequate model was performed. The order in which terms were deleted from the model took account of the p value ($p > 0.05$) and the size of the F ratio; the smallest F ratio being excluded first. Table 7-3 presents the minimal adequate model and each of the excluded terms, in the order that they were excluded. A post-hoc Anderson-Darling normality test confirmed that the

residuals were normally distributed. Residuals were also inspected for homogeneity of variance.

Table 7-3. Minimal adequate model for observer performance in the current experiment. (Upper section) Minimal adequate model. (lower section) excluded terms.

Minimal adequate model				
Source	<i>d.f.</i>	<i>F-ratio</i>	<i>p-value</i>	<i>R</i> ²
Eccentricity	2	81.1	< 0.001	27.23%
Length	4	31	< 0.001	20.82%
CLE	4	23.4	< 0.001	15.71%
Eccentricity * length	8	6.88	< 0.001	9.25%
CLE * length	16	1.79	0.043	4.80%
Observer * Eccentricity * length	8	2.11	0.041	2.84%
Observer * Eccentricity	2	4.09	0.020	1.37%
Observer * length	4	1.24	0.300	0.83%
Observer	1	2.08	0.152	0.35%
Error	100			
Total	149			83.21%

Deleted terms				
	<i>d.f.</i>	<i>F-ratio</i>	<i>p-value</i>	
Observer * CLE	4	0.87	0.49	
Observer * Eccentricity * CLE	8	0.58	0.78	
Eccentricity * CLE	8	1.23	0.30	
Eccentricity * CLE * length	32	1.46	0.10	

7.2.3.2 *Assessing the influence of Gabor patch phase upon contour detection*

The influence that Gabor element phase had upon contour detection rates was investigated for contours that featured five elements. Figure 7-4 shows the detection rates for these contours as a function of distance from fixation and CLE. There was no consistent effect of Gabor element phase upon contour detection rates, at any eccentricity. This was confirmed by a χ^2 analysis which indicated that there were no significant differences between detection rates for contours with different Gabor phase properties at any level of eccentricity or CLE.

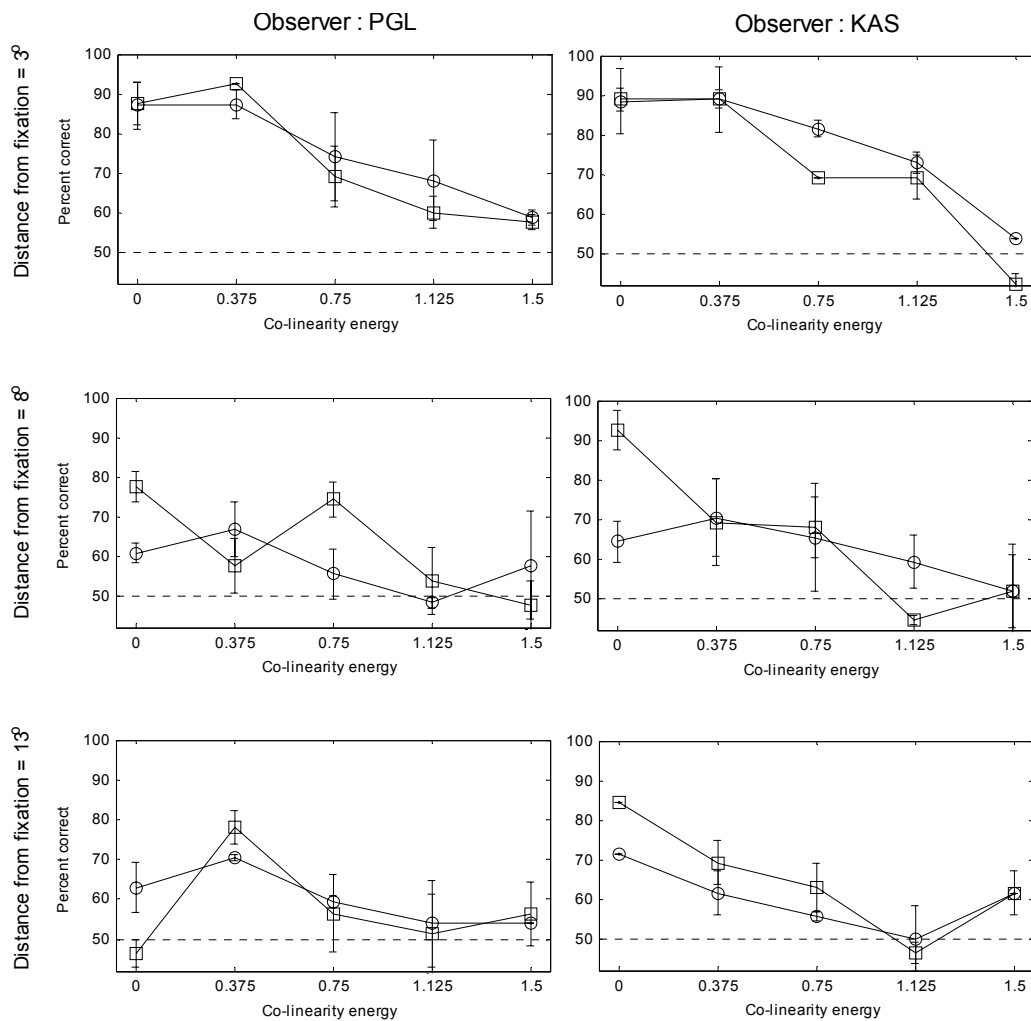


Figure 7-4. Plots illustrating the effect of Gabor patch phase. The left column presents the results for observer PGL, whilst the right column presents the results for observer KAS. Each row presents detection rates for contours that were located 3, 8 and 13° from the fixation point. *Symbols: (○) phase-alternated Gabor patches, (□) phase-aligned Gabor patches.* (Error bars ± 1 std.err.)

7.2.4 Discussion

Contour detection rates increased as a function of contour length, this corresponds to previous findings (for example, Uttal, 1975; Beck *et al.*, 1990; Smit *et al.*, 1985 and Moulden, 1994). However, for contours located 13° from the fixation point, observer PGL detected fewer 7 and 8 element contours than shorter contours (4, 5 and 6 elements). This fall-off in performance accounts for the significant interaction between observer, contour-length and eccentricity reported within Section 7.2.3.1. This may reflect an unanticipated order-effect. A post-hoc examination revealed that observation sessions featuring longer contours were undertaken later than those featuring shorter contours. Therefore, it is suggested that the poorer para-foveal performance in conjunction with longer contours may simply reflect the effect of fatigue. The experimental design required observers to

detect contours that might be located anywhere within a range of $\pm 13^\circ$, without moving their eyes, both observers found this task to be demanding.

Contour detection rates improved as the prevalence of noise-contours within the stimulus background was decreased. This effect was found at all contour eccentricities. However, even when CLE was minimised detection rates at 13° eccentricities only approached a level of 75% correct where contours were eight elements in length. A pilot study (Appendix I) has shown a close correspondence between the detection rates for contours embedded in randomly oriented background elements and where background element orientations are manipulated in order to impose a CLE of 1.0. The current psychometric functions indicate that contour detection rates for contours located at a distance of 13° from fixation were near to chance when the CLE was greater than 0.75. This suggests that the embedded contours would not have been detected at rates greater than 75% correct if background element orientations had been randomised. The significant interaction between CLE and length simply reflects the fact that when contours were short detection rates did not improve if CLE levels were reduced and vice-versa.

The fact that reductions in CLE led to improvements in peripheral contour detection may explain why detection rates in the current experiment for phase-alternated contours presented at 13° were higher than those reported by Hess and Dakin (1999). However, the manipulation of CLE does not explain why there was no consistent difference in the detectability of phase-alternated and phase-aligned contours presented to the peripheral-field, as was previously reported by Hess and Dakin. Therefore, it must be concluded that the results of the current experiment do not support the findings of Hess and Dakin, *i.e.* that the detection of phase-alternated contours is specifically impaired within the peripheral field. However, the distribution of observer responses was quite large at this eccentricity - an examination of the error-bars within Figure 7-4 confirms this. It may be that if additional observation sessions had been administered, then a slight difference in the detection rates for these stimuli might have been found.

7.3 Experiment 10: Do global properties influence contour detection in parafoveal areas?

7.3.1 Introduction

Chapter 5 has shown that the detection of 'smooth' contours is very poor for the simple-filter model. Hess and Dakin (1997, 1999) have argued that the simple-filter model is a sufficient account of contour detection in the peripheral field ($>10^\circ$). It follows then that if

smooth contours were presented at eccentricities beyond 10° then they should be undetectable by human observers. The relative detection rates for contours of various arc-lengths should correspond to that which was found for the simple-filter model (ii, below); rather than that found for foveal presentations of contours (i, below).

Table 7-4 . Relative contour detection rates for the simple-filter model (chapter 5) and the human fovea (chapter 4).

i) Human observers	:	15	>	8	>	1	>	4	>	2
ii) Simple-filter model	:	1	>	2	>	4	>	8	>	15

For contours presented at eccentricities of 3° and 8° , it was expected that detection rates would increase as a function of contour-arc length, confirming the findings of experiment 4. For contours presented at eccentricities of 13° two contrasting hypotheses were available. (i) If contour integration within the peripheral field were achieved via a coarse-scale mechanism then it was expected that contour detection rates should fall as a function of arc-length – because element orientations inevitably diverge as a function of distance. (ii) If the mechanism underlying peripheral-field contour detection were the same as those present in the central field, then detection rates should increase as a function of arc-length – because this is what we found for subjects.

7.3.2 Methodology

7.3.2.1 *Stimulus parameters*

The size of stimulus images in this experiment was the same as that used in the previous experiment (section 7.2) i.e. 220 x 800 pixels. Pilot studies revealed that contours featuring 15-16 elements were detected too easily by observers where the density of background and contour elements was matched¹⁴. Therefore the stimulus parameters selected in the current experiment reflect a compromise between stimuli that produced a ceiling effect in the fovea and stimuli that were undetectable when presented to the peripheral field.

In order to eliminate the foveal ceiling effect, the lengths of target contours featured in the current experiment were reduced to 10 elements and the path-angle value was increased from 24° to 40° . Consequently, a closed contour was achieved in a contour featuring nine elements. The higher path-angle value also ensured that coarse-scale mechanisms would be less likely to contribute towards contour detection. Each of these manipulations led to reductions in contour detection rates, thereby avoiding ceiling effects. When the simple-

¹⁴ In previous arc-length experiments the density of background elements was increased relative to the density of contour elements in order to manipulate the salience of target contours.

filter model was tested in conjunction with the current stimuli (Section 7.3.3.1), contour detection rates were largely consistent with those reported in Chapter 5.

Figure 7-5 illustrates the various contour-types and background manipulations utilised in the current experiment. For each contour the arc-length was fixed at one of five levels, these were 9, 5, 3, 2 and 1. The arc-length value refers to the number of contiguous elements that share the same sign of path-angle change, with the count commencing with the first contour element. Hence, arc-lengths within the contour may actually be one element longer than the arc-length value. This is because for any contour the line segment at the inflection point could be considered a member of both the preceding and of the following arc.

All contours consisted of 10 Gabor elements with the exception of the closed-contour which consisted of 9 elements. This ensured that the number of pair-wise co-alignments with contour neighbours was held constant. Consequently, any increase in detection rates for the closed-contours would not have been due to there being more pair-wise interactions for the closed contours.

In order to reduce the effect of any potential interaction between contour size and presentation eccentricity, all embedded contours were rotated so that their principal axis was vertical. The spacing of neighbouring contour elements was matched to the spacing of nearest neighbours in the background, hence the pair-distribution function for the spacing of contour and background elements was equivalent to that presented in Figure 6-7 (lower). The location of contours within each image was varied such that the contour's centroid was located at one of three offsets relative to the fixation point ($\pm 3^\circ$, $\pm 8^\circ$ and $\pm 13^\circ$ visual angle). In order to ensure that the observer would not benefit from fixating anywhere other than the fixation point, half of the contours were located to the left of the fixation point and the other half were located to the right.




Figure 7-5. A stylised version of the stimuli reflecting each of the independent variables manipulated in the current experiment. The six embedded contours illustrate the contour arc-length parameters. The background elements to the left of the fixation point have been manipulated in order to minimise the co-linearity energy. To the right of fixation, the orientation of background elements is randomised. For illustrative purposes, the contrast of the background elements has been reduced in order facilitate the detection of the embedded contours.

The location of background elements was determined using the *spaced-fill* technique (for details see section 6.2.4). The mean spacing distance between nearest-neighbours was 16 pixels, this corresponds to a separation of 4λ units. The co-linearity energy of the background elements was also manipulated. It was either minimised (*CLE minimised*), using the technique described in section 6.2, or the orientation of background elements was randomised (*CLE random*). The phase of all Gabor elements was alternated between 0° and 180° , thereby ensuring that the task could not be solved by coarse-scale mechanisms, according to Hess and Dakin (1999).

Each noise image was constructed using exactly the same procedure as its paired target image, except that the orientation of contour Gabor elements was randomised prior to display. Hence, any density difference that might exist between contour and background elements was also present within noise images.

7.3.2.2 *Procedure*

Observers undertook a number of observation sessions. Each session involved viewing 200 pairs of target and noise images presented within a 2AFC paradigm. Typically, a session lasted thirty minutes. Target images to be presented during an individual observation session were selected randomly - without replacement - from all available stimulus levels. Subsequently, matching noise images were selected. As with all the proceeding experiments, images were visible for 250ms. The schedule for fixation point and image presentation was the same as described in Section 7.2. The viewing distance and image-

scale properties were also unchanged. Therefore, values identified in Table 7-2 are also applicable to the current experiment. Stimuli were presented on a Cambridge Research VSG system (VSG2/4-4MB) and displayed with a Panasonic S110 20" monitor, further details of the display system have been provided in Section 2.1.3.

7.3.2.3 Observer details

Three observers participated in the current experiment. All observers were post-graduate research students. PGL (the experimenter) and KAS were experienced observers in path-paradigm experiments. EMH was naïve to path-paradigm experiments, however EMH received additional practice prior to onset of the experiment. KAS and EMH were both unaware of the current hypothesis.

7.3.3 Results

7.3.3.1 Simulation 6: Examining simple-filter model performance

The simulations presented in Chapter 5 demonstrated that the performance of the simple filter model was inversely related to arc-length. Because many other aspects of the current stimulus images differed from those described in chapters 5, it was necessary to confirm that the performance of the simple-filter model remained inversely related to arc-length. The performance of the simple filter model was evaluated in conjunction with the current stimuli following the procedures described in Section 3.2. Figure 7-6 shows contour detection rates as a function of stimulus level and filter-scale.

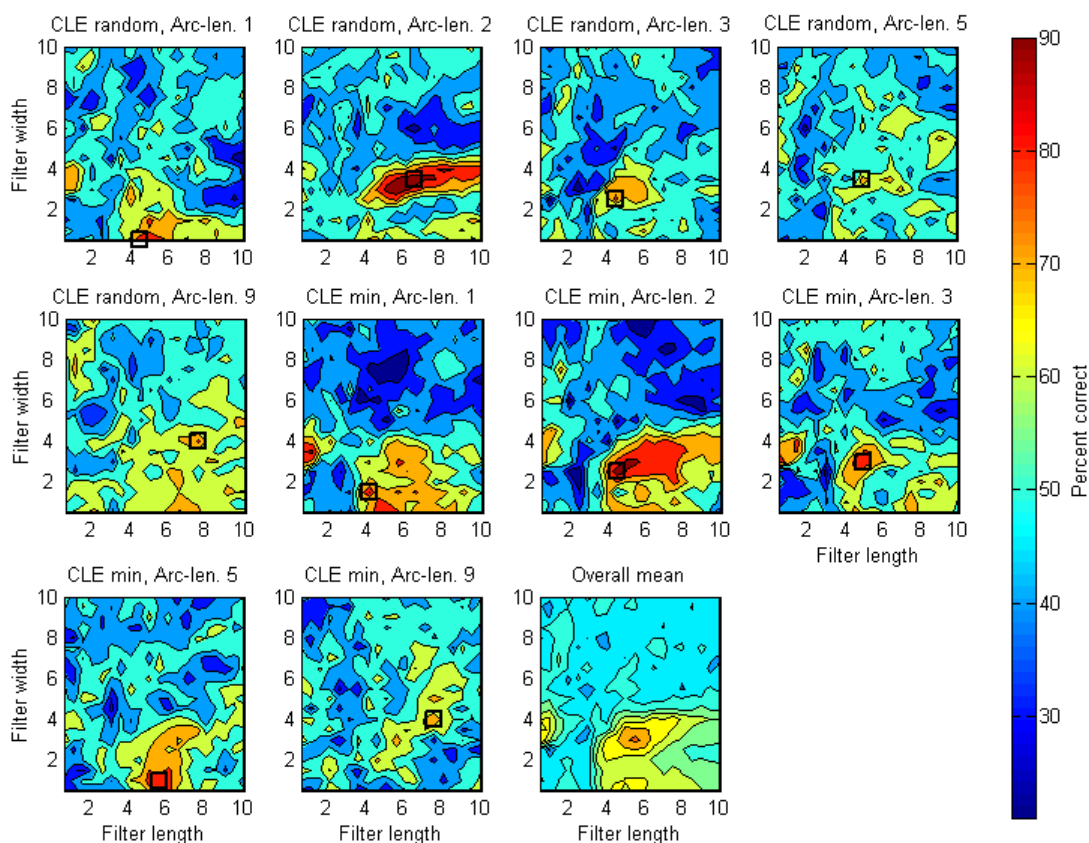


Figure 7-6. Contour detection rates for the simple-filter model as a function of filter-scale, CLE and arc-length. Symbols (□) indicate those filter-scales which gave rise to the highest target-contour detection rate.

An unrelated t-test was used to test whether there was a significant difference between the maximum ZBR length of target and noise images. This test was repeated for each filter-scale and each stimulus level. The proportion of filters for which there was a significant difference in ZBR lengths is presented in Figure 7-7 (left) as a function of CLE and arc-length. Following the criterion specified in Section 5.1.2.4 the ideal-filters, i.e. those elongated filters which gave rise to the highest contour detection rate for each combination of stimulus levels were identified. These filters are identified as square symbols upon the contour plots within Figure 7-6. The performance of the simple-filter model where the ideal-filters were utilised is shown in Figure 7-7 (right).




Figure 7-7. Plots showing the performance of the simple-filter model as a function of arc-length. (*left*) the number of filter-scales at which there was a larger maximum ZBR length in target images than in noise images (unrelated t-test, $p \leq 0.05$). (*right*) Performance of the simple-filter model, each point reflects the maximum detection rate achieved across all elongated filters. The filter-scale of each point on the plot corresponds to that identified by the symbol (\square) on the contour plots in Figure 7-6.

For both evaluation criteria, detection rates generally fell as a function of arc-length. However, detection rates for contours featuring an arc-length of one element were slightly lower than those for contours with an arc-length of two elements (approximately 90% Vs 96%), this unanticipated result may reflect the degree of co-alignment of same-sign areas of Gabor patches of the phase-alternated contours (Simulation 3, see Section 3.6). Furthermore, where the orientations of background elements were randomised (*CLE Random*), the detection rates for contours with an arc-length of nine elements exceeded those achieved for contours with shorter arc-lengths. The detection rates of the ideal-filters were adopted as an indicator of relative detection rates for the arc-length stimuli; these are summarised within Table 7-5.

Table 7-5. Relative contour detection rates for the simple-filter model

CLE Minimised	2 > 1 > 3 > 5 > 9
CLE Random	2 > 1 > 3 \geq 9 > 5

If similar patterns of relative performance were found in the peripheral field then this would support Hess and Dakin's (1999) hypothesis that a simple-filter mechanism formed the basis of contour detection in the periphery. If peripheral field performance diverged from that presented above, then this would suggest that a simple-filter mechanism did not underlie contour detection in the peripheral field.

A number of those filter-scales identified as ideal-filters, actually had elongation ratios that were quite divergent from Parker and Hawken's biological estimates, this issue is discussed in the following sub-section.

Do the ‘best-performing’ filters have a biologically plausible elongation ratio?

Figure 7-8 illustrates how filters with different scales and elongation ratios were best suited to the detection of contours with differing arc-lengths. The majority of the ideal-filters had an elongation ratio that is less than 1:0.5 (length : width). Whilst three of the ideal-filters (a, i and f) were perhaps elongated beyond limits of biological plausibility, alternative filters were available (x, y) which led to only slight reductions in performance.




Figure 7-8. Contour line-plot illustrating those regions of ‘filter-space’ which give rise to the highest detection rates for the current stimuli. The letters identify those filter-scales giving rise to the highest overall detection rate for each arc-length and the polygons identify those contiguous filter-scales that give rise to performance rates within 5% of the maximum. Regions a, i and f correspond to filter-scales which are relatively distant from physiological estimates of simple-cell elongation, alternate regions giving only slightly reduced detection rates are indicated (x & y). Parker and Hawken’s estimate of simple-cell elongation is represented on the plot as a solid diagonal line.

Table 7-6 presents those filter-scales that gave rise to the highest contour detection rate for each stimulus level.

Table 7-6. (Upper table) Filter-scales which gave rise to the highest contour detection rates for the simple-filter model. (Lower table) alternative filter-scales to those in the upper table that have a relatively large elongation ratio compared to estimates of elongation offered by Parker and Hawken.

CLE mode	Arc-length	Filter width (σ)	Filter length (σ)	Percent correct	Symbol
Random	1	0.5	4.5	89.0%	a
Random	2	3.5	6.5	97.0%	b
Random	3	2.5	4.5	82.0%	c
Random	5	3.5	5.0	74.0%	d
Random	9	4.0	7.5	81.0%	e
Minimised	1	0.5	5.0	87.5%	f
Minimised	2	2.5	4.5	96.0%	g
Minimised	3	3.0	5.0	87.5%	h
Minimised	5	1.0	5.0	87.5%	i
Minimised	9	5.0	7.5	78.5%	J

Random	1	2.5	4.0	77.0%	X
Minimised	1	3.0	5.5	83.5%	y
Minimised	5	3.0	5.5	78.0%	y

If one were to accept only those filter-scales which lie close to estimates of simple-cell elongation, then the performance of the simple-filter model may be characterised by Figure 7-9, below.




Figure 7-9. Simple-filter model performance, where only filter-scales falling within an elongation ratio of 1:0.5 are accepted.

7.3.3.2 *Human observer performance*

Figure 7-10 shows contour detection rates as a function of eccentricity, arc-length and co-linearity energy. For all eccentricities, there was a clear effect of arc-length. With few

exceptions, contour detection rates increased as a function of arc-length. Consistent with the results of experiments 7 and 8, minimising the co-linearity energy of background elements led to improved detection rates for embedded contours. Detection rates for all contours fell as a function of eccentricity.

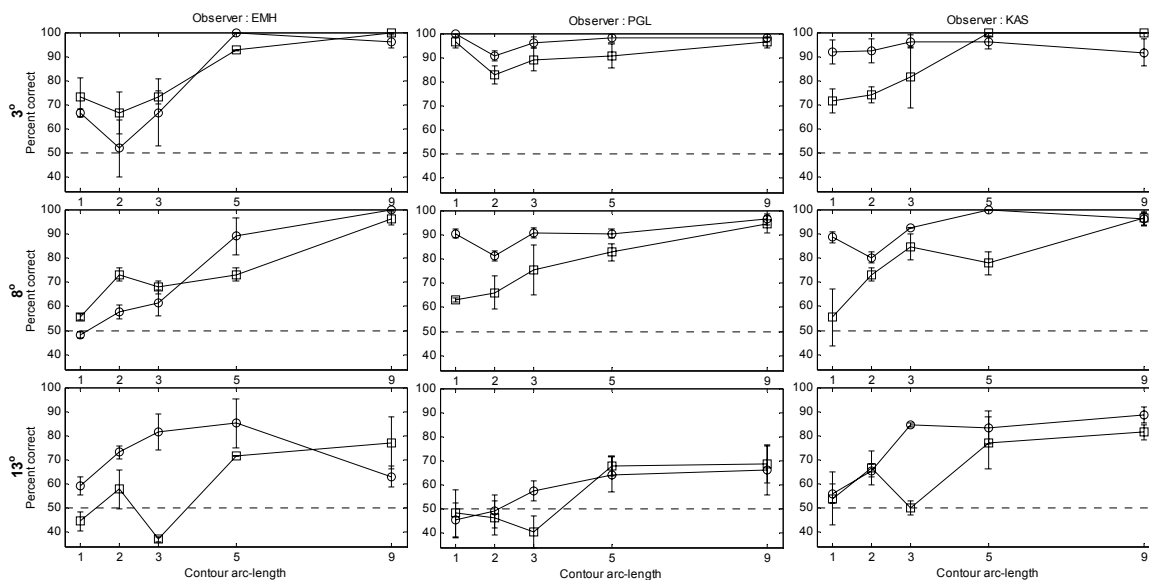


Figure 7-10. Contour detection rates as a function of arc-length, CLE and distance from fixation. Each column presents the psychometric plots for EMH (left), PGL (centre) and KAS (right). Each row presents detection rates for contours presented at eccentricities of 3° (top), 8° (middle) and 13° (bottom). Symbols: (o) = CLE minimised, (\square) = CLE random. Error bars = ± 1 std. err.

The mean contour detection rates for all observers are presented in Figure 7-11. Regardless of eccentricity there is a clear effect of contour arc-length, smoother contours were detected most frequently. Contours which featured an arc-length of three elements were poorly detected by all observers at 13° eccentricity where the orientation of background elements was randomised; this unexpected result is discussed in Section 7.3.4.




Figure 7-11. Mean contour detection rates for all subjects. Symbols: (o) = minimised CLE, (\square) = random CLE.

A post-hoc examination of the performance of each observer was conducted. This examined whether detection rates for each type of contour were higher when they were presented either to the left or to the right of the fixation point. Mean contour detection rates

across all arc-length levels did not differ significantly for left or right presentations. This finding suggests that observers did not fixate to the left or to right of the fixation point. Furthermore, as detection rates were highest for centrally presented contours ($\pm 3^\circ$) this implies that observers had followed their instructions and fixated centrally.

7.3.3.3 *Analysis of variance for observer performance*

A minimal adequate model was created using the procedure outlined in Section 7.2.3.1. The final model accounted for 75% of the variance in observer detection rates. Table 7-7 shows the terms within the model and lists those that were excluded. A post-hoc Anderson-Darling normality test confirmed that the residuals were normally distributed. Residuals were also inspected for homogeneity of variance.

Table 7-7. ANOVA analysis of observer contour detection rates.

Minimal Adequate Model				
Source	<i>d.f.</i>	<i>F-ratio</i>	<i>p-value</i>	<i>R</i> ²
Eccentricity	2	51.56	< 0.001	33.01%
Arc-length	4	20.54	< 0.001	26.30%
Observer * Eccentricity	4	5.65	< 0.001	7.24%
Observer	2	8.13	< 0.001	5.20%
CLEMode	1	12.25	< 0.001	3.92%
Error	76			
Total	89			75.67%

Excluded terms				
	<i>d.f.</i>	<i>F-ratio</i>	<i>p-value</i>	<i>R</i> ²
CLEMode * Eccentricity	2	1.41	0.273	
Observer * CLEMode * Eccentricity	4	2.65	0.072	
CLEMode * Eccentricity * arc-length	8	1.84	0.143	
Observer * CLEMode	2	1.38	0.268	
Observer * CLEMode * arc-length	8	0.47	0.866	
CLEMode * arc-length	4	2.06	0.104	
Observer * arc-length	8	1.41	0.218	
Observer * Eccentricity * arc-length	16	0.73	0.750	
Eccentricity * arc-length	8	1.14	0.348	

7.3.3.4 *Comparing simple-filter model and peripheral field performance*

Contour detection rates for the human peripheral field (13°) and for the simple-filter model are presented in Figure 7-12. The detection rates for the simple-filter model are clearly divergent from those of human observers. For the human periphery, the two smoothest contours (arc-length = 5 & 9 elements) were most frequently detected. Conversely, for the simple-filter model those contours which were most jagged (i.e. arc-length = 1 & 2 elements) were detected most frequently.

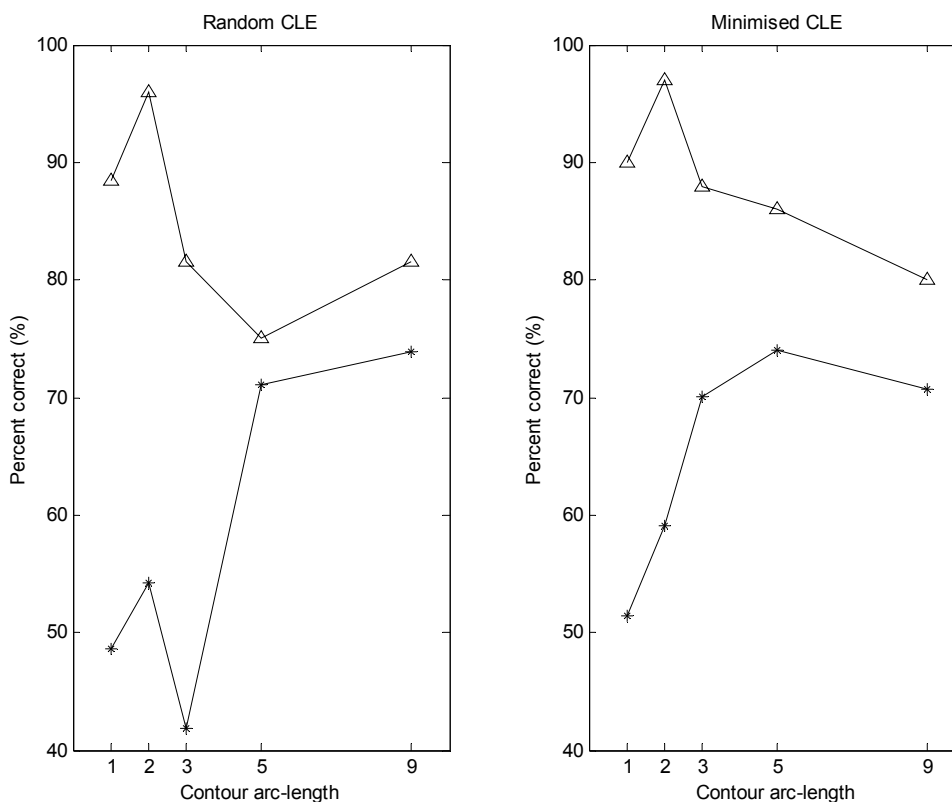


Figure 7-12. Comparing para-foveal and simple-filter model performance. Plots present contour detection rates as function of contour arc-length and co-linearity energy. Symbols: (*) mean ($n=3$) detection rates for contours presented to human para-foveal vision, (Δ) simple-filter model.

7.3.4 Discussion

The findings show a clear effect of arc-length in both the fovea and the periphery. Contours that featured longer arcs were detected most frequently, regardless of their location. This pattern was especially apparent in the peripheral field where jagged contours were barely detected at all ($\approx 55\%$). Conversely, for the simple-filter model, detection rates were generally inversely related to arc-length, jagged contours were most readily detected ($\approx 95\%$). This suggests that the detection of smooth contours that occurs in the periphery must be due to something other than a simple-filter mechanism. Any fall-off in performance with greater displacement appears to reflect a change in overall performance levels, rather than a qualitative change in the mechanism underlying contour integration. However, there remain some potential flaws in this assertion, these are addressed below.

As detection rates for the simple-filter model were higher than those reported for the peripheral field, one could argue that a simple-filter mechanism may utilise sub-optimal filters, that favoured smooth contours, within the periphery. A close examination of Figure 7-6 reveals that one could assemble a collection of filter-scales that would give rise to

performance levels which correspond to those reported for the peripheral-field. For example, it could be argued that as detection rates for the peripheral field were generally better for contours with an arc-length of nine elements and as the simple-filter model detected these stimuli most successfully when relatively coarse filters were utilised, then it might follow that the peripheral-field may favour smooth contours because of a greater reliance upon coarse-scale filters.

However, detection rates for contours with an arc-length of five elements were equally high in the peripheral field. The simple-filter model requires relatively fine-scale filters in order to detect these stimuli, suggesting that a larger range of filter-scales is required in order to accommodate a simple-filter account of peripheral field performance. Finally, close inspection of Table 7-6 and Figure 7-8 reveals that those filter-scales which were best suited to the detection of contours with an arc-length of two elements lie exactly between these extremes, consequently one would expect higher detection rates for these contours. In fact, peripheral field detection rates were poor for these contours, undermining this argument. It could be argued that coarse and fine scale filters were available, whilst those with an intermediate scale were not, however this stance appears counter-intuitive.

The arguments raised above are based upon the assumption that impaired peripheral field performance, regardless of arc-length, could be due to a poor match between the filters available to a simple-filter mechanism and the stimuli which were presented. This assumption has a number of weaknesses. Firstly, manipulations of arc-length have a similar influence upon detection rates in both the central and peripheral fields, despite the fact that appropriate filter-scales would have been available within the central-field. Secondly, it is likely that eccentricity *per se* is responsible for the decrease in detection rates within the peripheral field, rather than a specific impairment in the contour integration mechanism. Observers were required to fixate centrally while contours could be presented anywhere within a range of $\pm 13^\circ$ (visual-angle), therefore attention was inevitably distributed over a relatively wide area. Evidence from the visual search paradigm suggests that performance is likely to be poorer when targets are presented further from fixation (Wolfe, O'Neill & Bennett, 1998).

Arc-length and contour-size are inevitably inversely correlated when element spacing and path-angle are held constant. Therefore, the possibility remains that detection rates were determined by the overall size of contours, rather than arc-length. However, if this account were accepted then the simple-filter model would require extension in order to

accommodate this phenomenon, as ZBRs rarely correspond to more than a small subsection of the target-contour (see Simulation 1, Section 3.3.3.2).

As with experiment 9 (Section 7.2) the level of co-linearity energy appears to have had a small, though consistent, effect upon contour detection rates. The reduction of the prevalence of noise-contours within both target and noise stimuli leads to improved observer performance. Where the orientation of background elements was randomised, observers were specifically impaired when detecting contours with an arc length of 3 elements which were located at 13° . This may reflect the possibility that there are many randomly occurring 3-element arcs present within randomised backgrounds. Consequently, performance would be impaired when target contours share these properties. For foveal presentations, this effect was less pronounced. This may be because additional attentive mechanisms may have been available during foveal detection that enabled the rejection of noise-contours.

7.4 General Discussion

The current results do not support those of Hess and Dakin (1997, 1999), phase-alternated contours were detected when presented to the peripheral field. Detection rates for both phase-aligned and phase-alternated contours fell as a function of eccentricity. This decrement reflects the normal fall in performance as stimuli are presented further away from the fovea. Peripheral detection rates were quite low ($\leq 75\%$), this reflects the fact that few elements were featured within contours. Pilot studies confirmed that contour detection rates were higher when longer contours were utilised. Only one observer (KAS) demonstrated a consistent improvement in detection rates where Gabor patches were phase-aligned. However, this improvement was much smaller ($\approx 15\%$) than that reported by Hess and Dakin.

There was a consistent effect of arc-length in both the central and peripheral field. For human observers detection rates rose as a function of contour arc-length whereas for the simple-filter model detection rates tended to fall as a function of arc-length, though detection rates for the most jagged contours (*arc length* = 1) did not fit this pattern. The arc-length effect occurred in the peripheral field despite the fact that path-angle values were relatively large (40°) and Gabor patches were phase-alternated. It was concluded a coarse-scale mechanism that is equivalent to the simple-filter model is unlikely to underlie contour detection in the peripheral field.

In both experiments reported within the current chapter there was a small, though consistent, effect of CLE manipulation. Contour detection rates improved where the collinearity energy of background elements was minimised. This is likely to be due to a reduction in the occurrence of distractions and false alarms that occur when the orientations of background elements are merely randomised. Minimising CLE improved the detection of both phase-alternated and phase-aligned contours. Consequently, the influence of CLE manipulations cannot be invoked as an explanation of the divergent findings of the current experiments and those of Hess and Dakin (1997, 1999).

-0-0-0-

8 General Discussion

8.1 *The role of coarse-scale processes in the detection of path-paradigm stimuli*

8.1.1 Detection of contours with a random global structure

Field, Hayes and Hess (1993) suggested that path-paradigm stimuli enabled the investigation of grouping mechanisms that occurred after the initial edge-detection processes of early vision. They argued that coarse-scale processes of early vision would be unable to detect path-paradigm contours as they were formed from Gabor elements which were band-pass in nature. The presence of path-paradigm contours is indicated by the co-alignment and co-orientation of multiple Gabor patches. In essence, they claim that mechanisms which only employ fine spatial scales are insensitive to co-alignment information, whereas mechanisms that only employ coarse-spatial scales lose orientation information. In their study they found that observers were proficient in the detection of contours embedded in path-paradigm stimuli. This prompted their proposal of the association field, a mechanism which integrated the outputs of those filters that detected the presence of individual contour Gabor elements.

The majority of the path-paradigm experiments presented by Field *et al.* utilised stimuli which were phase-aligned. Hess and Dakin (1997, 1999) raised concerns about whether such stimuli succeeded in precluding the contribution of coarse-scale mechanisms to the detection of path-paradigm contours. They demonstrated that a simple-filter model, based upon Watt (1991), was able to indicate the presence of contours in phase-aligned stimuli. The performance of the simple-filter model was based upon the maximum length of zero-bounded regions (ZBRs) in images that had been filtered and then subjected to a thresholding operation. However, they also demonstrated that the simple-filter model was unable to detect phase-alternated contours. Hess and Dakin argued that whilst coarse-scale processes might be involved in the detection of phase-aligned stimuli, such mechanisms could not be involved in the detection of contours within path-paradigm stimuli that are composed of phase-alternated Gabor patches.

Simulations 1-3 (Chapter 3) investigated the contour detection performance of the simple-filter model over a range of independent variables, these included specific properties of the stimulus: path-angle; element spacing and Gabor patch phase and properties of the model: filter width; filter length and filter orientation. These experiments demonstrated that the simple filter model was able to detect both phase-aligned and phase-alternated stimuli,

often at rates surpassing those reported for human observers. However, whilst many different filters were employed, the *output* of the model, i.e. the longest ZBR, was based upon convolution with only a single filter. That is, the model requires only “*intra-cellular filtering*” (Hess and Dakin, 1999, *pg.* 947), as opposed to inter-cellular integration, i.e. an association field mechanism.

The simulations presented in Chapter 3 have demonstrated that the performance of the simple-filter model is comparable to that of human observers when an appropriately scaled filter is utilised. Where the model was successful, the filter-scale reflects a compromise between those filters ideally tuned to individual Gabor patches and those sensitive to co-alignments of multiple Gabor patches. Where the filter scale is too fine then ZBRs only correspond to the light and dark areas of individual Gabor patches. Conversely, local orientation information is lost if filters are too coarse. As a consequence, the lengths of ZBRs increase but these ZBRs span the locations of Gabor patches that were not co-aligned, i.e. ZBR lengths no longer correlate with the presence of target contours.

8.1.2 Detection of smooth contours

Simulations 4-6 (Chapter 5) demonstrated that the simple-filter model was specifically impaired in the detection of contours that were composed of extended arcs, i.e. smooth contours. Contour detection rates fell as a function of contour arc-length. Whilst detection rates for smooth contours sometimes exceeded 75% for particular filter-scales, detection rates were always higher for jagged contours. For the simple-filter model ZBR lengths covary with the degree of co-alignment and co-orientation of target contours. Since co-orientation decreases as a function of arc-length, it follows that the detection rates of the model will also fall as a function of arc-length.

The consistent relationship between arc-length and detection rate provides a useful means of dissociating the performance of the simple-filter model from the performance of human observers. For the latter, performance does not appear to fall as a function of arc-length. In fact, experiments 4 and 10 demonstrate that detection rates have a positive relationship with arc-length.

8.1.3 Summary

Coarse-scale mechanisms can account for the detection of targets in path-paradigm experiments, where the global structure of the contour is randomised. However, this mechanism does not form a complete representation of the whole contour, especially

where the contour features smooth continuous arcs. This distinction illustrates that successful performance in path-paradigm experiments does not necessarily require that all contour elements are integrated.

Where the global structure of embedded contours was randomised, *i.e.* changes in the sign of the path-angle between successive elements were randomised, the simple filter model performed as well as human observers. Thus, it is concluded that path paradigm experiments (for example: Field *et al.* 1993; Hess and Dakin 1997 and 1999) may not have adequately controlled for the operation of coarse-scale mechanisms. It is necessary to make certain that target contours embedded in path-paradigm stimuli cannot be detected by coarse-scale mechanisms. One potential solution is to manipulate arc-length. However, this methodology involves additional potential confounds, as arc-length is negatively correlated with the overall extent of contours.

8.2 Smoothness in path-paradigm stimuli

Experiment 3 (Section 4.2) confirmed Kovacs and Julesz' (1993) finding that closed contours are detected more frequently than open contours. However, whilst Kovacs and Julesz described a step-change in contour detection rates as closure was achieved, Section 4.2 did not support this claim. Further experiments (numbered 4-6; Sections 4.3-4.5) suggested that a more parsimonious explanation for the closure effect was available. These results demonstrated that contour detection rates increase as a function of smoothness, where smoother contours feature fewer inflections. This explanation would predict that closed contours would be most salient as they are most likely to feature arcs that are continuous and smooth. This account of the closure-effect undermines the Kovacs and Julesz explanation, *i.e.* that neuro-physiological mechanisms underlying contour integration might signal contour membership through the synchronised firing of receptive field outputs (*e.g.* Singer and Gray 1995). However, recent evidence suggests that closed contours may indeed have privileged status in perceptual processing (Hess, Beaudot and Mullen, 2001, see Section 8.5.3.4).

Experiment 4 (Section 4.3) proposed that the smoothness-effect might occur because the mechanism underlying contour detection favoured arcs. If this were correct then contours which were inflected would only be as salient as the arcs which composed the whole contour. Whilst initial experimental results supported this hypothesis, more rigorous experiments did not. An alternative explanation of the smoothness effect is that the

association field could respond more strongly to smoother contours. Appendix J describes modifications to the association field model that can increase the relative detection rates for smooth contours over jagged contours without recourse to firing synchrony mechanisms.

Contour size remains a potential confound for experiments which manipulate contour smoothness. Smoother contours tend to be smaller overall than their jagged counterparts. Pettet, McKee and Grzywacz (1998) rotated contours, so that their principle axes were vertical, in order to minimise this potential confound. A similar methodology was followed in experiment 10 (Section 7.3), the smoothness effect was still recorded, suggesting that this result is not due to an artefact of the experimental design.

8.3 Contour detection in the periphery

Hess and Dakin (1997, 1999) claimed that the mechanisms underlying contour detection in the peripheral field differed from those responsible for contour detection in the central field. They argued that in the periphery only coarse-scale processes underlie contour integration, whereas for the fovea, both coarse-scale and association field mechanisms are available.

Experiment 9 (Section 7.2) examined contour detection rates for phase-aligned and phase-alternated contours in the central (3°) and in the peripheral field (13°). These contours had a randomised (*jagged*) global structure. Both phase-aligned and phase-alternated contours were successfully detected in the periphery. Sub-section 7.2.3.2 showed that the detection rates for phase-aligned and phase-alternated contours did not differ significantly. The finding that phase-alternated contours can be detected in the periphery does not undermine Hess and Dakin's hypothesis that coarse-scale processes alone might underlie peripheral-field contour detection. Simulations 1-3 have shown that the simple filter model can detect phase-alternated contours when constraints upon the selection of filter-scale are relaxed. Both the model and observers were able to detect phase-aligned and phase-alternated contours that had a random global structure. If this were the only pertinent evidence, then the simple-filter model would remain tenable as an account of peripheral field contour detection. The following paragraph introduces evidence that undermines the SFM account of para-foveal contour detection.

In order to establish whether the same mechanisms might underlie both central and peripheral field contour detection, a further experiment was conducted which manipulated

contour smoothness and the eccentricity of presentation. Simulations 4 and 5 have shown that the simple filter model is specifically impaired in the detection of smooth contours. Therefore, if coarse-scale mechanisms underlie peripheral field contour detection then performance should fall as a function of smoothness. In fact, the inverse relationship was found for human observers, contour detection rates generally increased as a function of contour smoothness, regardless of the eccentricity of presentation, whereas performance for the simple filter model generally worsens as contour smoothness was increased.

This finding informed the conclusion that the same mechanism might underlie the detection of contours that are presented to both the central and to the peripheral field. Furthermore, as smooth contours were favoured at all presentation eccentricities, it can be concluded that the simple-filter model does not provide an adequate account of contour detection at any eccentricity. If a mechanism analogous to the simple-filter model were responsible for contour detection in the periphery, then the proposed mechanism would have to be modified in order to account for the increase in detection rates for smooth contours. It is difficult to envisage a mechanism that could favour smooth contours without invoking inter-cellular linking between the outputs of differently oriented receptive fields. This, hypothetical, mechanism would require receptive fields that were specially tailored for the detection of contours with particular shapes, spatial frequencies and locations, i.e. a grandmother cell for contours.

Unlike the simple-filter model, the association-field mechanism can account for the detection of both smooth and jagged contours. Field, Hayes and Hess (1993) do not explicitly state that the association field model should favour smooth contours, however this model readily lends itself to extensions which would increase the salience of such contours (Appendix J).

Two potential weaknesses remain in the assertion that a coarse-scale mechanism does not underlie peripheral field contour detection:

- (1) The presented smoothness effects might have been confounded by the overall size of contours, as this decreases as a function of arc-length. Whatever the nature of the mechanism that underlies contour detection in the periphery, it might favour compact contours.
- (2) Smoothness effects might occur because processes that occur after initial linear filtering are influenced by global structure. For example, mechanisms responsible for

the identification and selection of potential target-contours might favour regularly shaped regions, i.e. arcs and circles.

Both arguments are undermined by the fact that the ZBRs generated by the simple-filter model rarely correspond to whole target contours. Where path-angle is greater than 20° the longest ZBRs in filtered images represent less than half of the target contour elements, regardless of Gabor element phase (see Figure 3-9). If fewer than half of the elements which form the target contour are detected then much of the global structure of a contour is lost. Consequently, an additional mechanism would have to be invoked that enabled the integration of the ZBRs representing disparate contour segments. It is precisely such a mechanism that is rejected within the simple-filter model.

Recently, Nugent, Keswani, Woods and Peli (2001) reported successful detection of phase-alternated contours at eccentricities of up to 22° . Furthermore, the detection rates for their phase-alternated contours were only approximately 10% lower than they were for phase-aligned contours. One might argue that their findings support the current conclusions. However, Nugent *et al.* presented phase-aligned and phase-alternated contours that had a path-angle of only 0° , simulations 1-3 have shown that such contours can be detected with relative ease by the simple-filter model. Consequently, whilst they claim that their results undermine the findings of Hess and Dakin (1997, 1999), their experiments actually failed to control for coarse-scale processes.

8.4 Background element co-linearity energy

Experiment 7 demonstrated that the occurrence of spurious contours amongst the background elements of path-paradigm stimuli (noise-contours) can influence detection rates for target-contours. The association-field function described by Yen and Finkel (1998) was utilised as a measure of co-linearity energy (CLE) within stimulus backgrounds. Experiment 8 (Section 6.2) described how CLE levels could be manipulated in order to increase or decrease the occurrence of noise-contours within path-paradigm stimuli. Consequently, this enabled an examination of the influence that CLE has upon the detection rates for embedded target-contours. The CLE of background elements was manipulated in conjunction with manipulations of contour path-angle, contour length and contour smoothness. In each case reductions of CLE led to slight improvements in contour detection rates. Experiments 9 and 10 (Chapter 7) demonstrated that reductions in CLE had

a similar influence upon detection rates in the peripheral field to that found in the central field.

Experiment 7 (Section 6.3) reported that detection rates for contours with path-angles of 0-50° increased as CLE levels were decreased. However, regardless of CLE levels, contour detection rates fell to chance levels where path-angles exceeded 50°. The pattern of performance would be consistent with an association field mechanism where co-facilitatory connections only occur between co-circular receptive fields that differ in orientation by less than 50°.

However, the contour path-angle range identified may not reflect the absolute limits of contour integration. The relative densities of contour and background elements were carefully matched. Therefore, the only cue towards the presence of a target contour should be the co-alignments between consecutive contour elements. When manipulating CLE, there are geometrical limitations to the amount of misalignment that can be achieved in a densely packed array of background elements. As a consequence, following the minimisation of CLE, the mean CLE levels of background elements may remain higher than they are for embedded contours with path-angles of more than 50°. For example, Figure 8-1 (top row) shows three path-paradigm stimuli. In each of these stimuli the CLE of background elements has been minimised and the path-angle is varied. The lower row of this figure shows the same stimuli with the exception that the contrast of Gabor patches is determined by the actual CLE level for each element (see 6.2.1 for details of CLE measures).

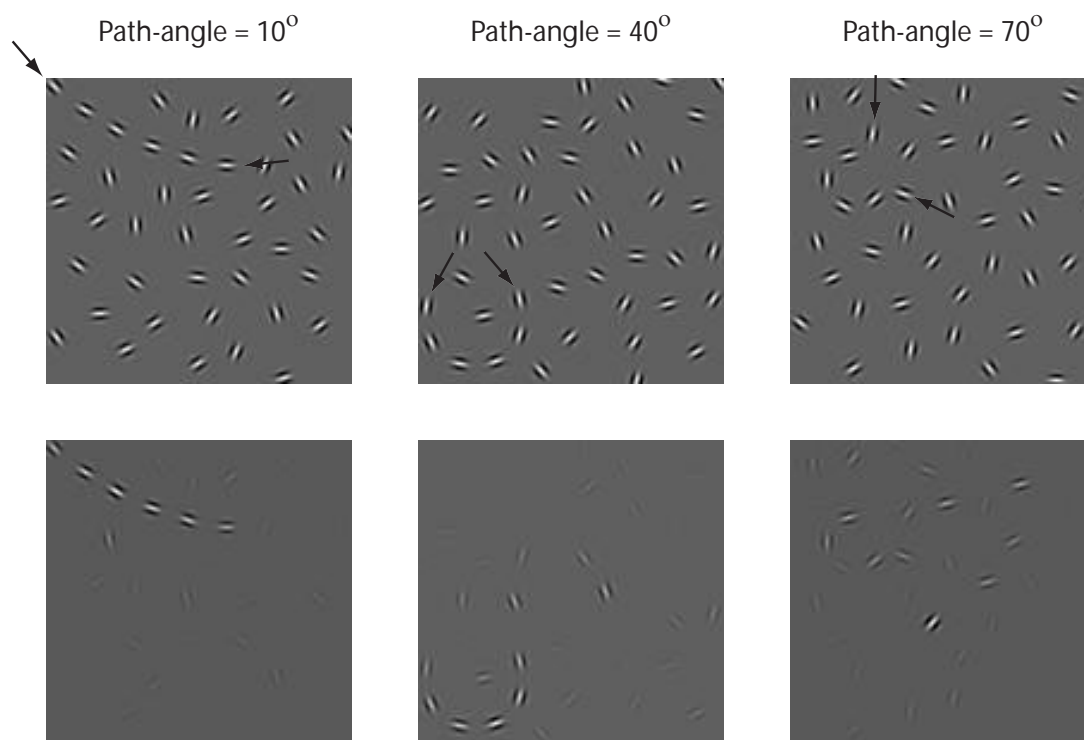


Figure 8-1. (*top row*) path-paradigm stimuli featuring embedded contours with varying contour path-angle these are indicated by arrows. (*bottom row*) The contrast of each Gabor patch is determined by its co-linearity energy.

It is apparent that where contour path-angle is 70° , the CLE level of the contour Gabor patches is approximately the same as some of the background elements. This may explain why, even when CLE is minimised, contours with a large path-angle are not detectable. Where contour path-angle levels are high, the cue provided by element co-alignment is equally strong for the background elements (within both target and noise images) as it is for the target-contour.

Braun (1999) estimated that a contour length of six elements was required in order to achieve a detection rate of 75% correct. Section 6.4 reported that contours consisting of only four elements could be detected at rates approaching 75% correct. This difference is likely to reflect the influence of CLE minimisation in background elements, in Braun's experiment background elements were randomly oriented.

Whilst manipulations of CLE provide a useful means of improving contour detection thresholds, they do not serve to distinguish between coarse-scale and association-field based detection mechanisms. In Section 7.3.3.1 it was reported that reductions in CLE led to an improvement in contour detection rates for the simple-filter model, the results of experiments 7-10 suggest comparable effects for human observers. Regardless of the

mechanism underlying contour detection, it is unsurprising that lowering CLE levels will lead to improvements in contour detection thresholds. The presence of contours in path-paradigm experiments is cued by co-alignments of neighbouring contour elements. Consequently, performance should inevitably improve where these co-alignments between background elements are minimised.

8.5 Acquisition of an interactive contour integration mechanism

The main emphasis of this thesis has been to establish whether or not the simple-filter model of early-vision provides an adequate account of performance in path-paradigm experiments, in either the central or the peripheral fields. The experimental and modelling results presented here suggest that a simple-filter account is not sufficient. That is, additional inter-cellular linking mechanisms are required in order to provide a full account for human observer performance in path-paradigm experiments. This section considers the nature of these inter-cellular interactions. A model is proposed that places the association-field (Field *et al.*, 1993) into a framework provided by Roelfsema, Lamme and Spekreijse (2000).

This model prescribes roles for feed-forward, horizontal and feed-back connections within the primary visual cortex. This account of observer performance in path-paradigm experiments is based upon the earlier stages of a model of curve-tracing offered by Roelfsema, Lamme and Spekreijse (2000).

In curve-tracing the first processing stage involves the formation of a base representation (Ullman, 1984). This is established by the feed-forward activation of neurones throughout the visual hierarchy, in response to the presentation of the visual stimuli. This representation is established quickly, prior to the operation of feed-back and horizontal connections. However, Roelfsema *et al.* accept that some modulatory interactions occur between feed-forward channels, these interactions are termed “feed-forward counterbalancing”. Next, an incremental representation is generated according to Gestalt grouping principles. Those principles that are most relevant to the current thesis are represented within the horizontal connections of the primary visual cortex (see Section 1.4.2). The result is a network of enabled connections called the *interaction skeleton* (Roelfsema and Singer, 1998). Finally, attentional influences lead to the enhancement of activation at a point upon the contour. Directed attention enables the spread of this

activation along the interaction skeleton until all parts of the curve receive enhanced activation relative to unconnected curves.

A key feature of their model is that the interaction skeleton is simultaneously activated at many levels of the visual hierarchy. For curve-tracing, this enables faster tracing of straight or regularly shaped contours, as tracing can be pursued in visual higher areas which have receptive fields spanning larger spatial distances. Roelfsema *et al.* propose that representations of familiar contour shapes can be “chunked” thereby speeding tracing (Mahoney and Ullman, 1988). Chunks may represent relatively complex shapes which can be detected as a whole during the establishment of the base-representation. Roelfsema *et al.* accept that chunks may be acquired through visual experience. For example, Vecera and Farah (1997) found that tracing speeds were relatively enhanced for letters when compared to inverted letters, suggesting that these chunks can be formed through visual experience.

For Roelfsema *et al.* the grouping of contour segments is only implicit when the interaction skeleton is formed. Further attentional processing is required before the connectedness of the contour is established (for example, Jolicoeur, Ullman and MacKay, 1986 & 1991). Physiological (Roelfsema, Scholte and Spekreijse, 1999) and psychophysical (Roelfsema, Lamme and Spekreijse, 1998) evidence has demonstrated that there is a temporal cost to this process that is related to the spread of activity through successive enabled connections. As path-paradigm tasks seem to involve no temporal costs, it is generally accepted that pre-attentive processes must underlie performance in path-paradigm experiments (Field, Hayes and Hess, 1993; Kovacs and Julesz, 1993). Roelfsema *et al.* propose that the base representation alone may support this form of contour grouping. They suggest that the model offered by Gigus and Malik (1991) provides an adequate account of this process. However, their suggestion gives rise to a contradiction: they propose that for path-paradigm stimuli, the base representation accommodates a feed-forward contour detection mechanism; whereas for curve-tracing they suggest that contour grouping is achieved within the incremental representation. This seems an unparsimonious stance, as it implies that two independent contour grouping processes, both of which are based upon Gestalt principles, exist within early vision.

The basis for the distinction between processes suggested by Roelfsema *et al.* is the fact that contour grouping is achieved pre-attentively in the path-paradigm whereas attention is required in curve-tracing tasks. In the path-paradigm there is an implicit acceptance that

the observer is aware of the connectedness of all contour elements at the point of the psychophysical decision. If this assertion were rejected then the basis for Roelfsema *et al.*'s proposal of independent grouping substrates for curve-tracing and contour integration is removed.

Section 3.3.3.2 has demonstrated that the performance of the simple-filter model can surpass that of human observers, even when only a fraction of the whole contour is encompassed by the longest zero-bounded region. This illustrates the fact that observers in path paradigm experiments do not have to establish the connectedness of all contour elements, they only have to identify which, from a pair of images, is the one most likely to contain a contour. Integration of all contour elements is only inferred because the observer has successfully identified a target image. This distinction is important, the co-alignment of elements belonging to a small part of the contour may provide sufficient evidence to inform the decision of the observer.

To reiterate, success in path-paradigm experiments does not require that the observer is aware of the connectedness of all contour elements, an assertion which may conflict with the introspections of observers. In the terms of the current model, the interaction skeleton is established following the application of Gestalt grouping principles. These principles involve pair-wise co-facilitatory interactions between co-aligned and co-activated receptive fields (i.e. the association field, Field *et al.*, 1993). Whilst the interaction skeleton may represent the whole of the contour, the attention of the observer may only be drawn towards a fraction of the skeleton. The detection of this fraction may be sufficient to inform the assignment of *target* or *noise* status in a 2AFC trial.

Figure 8-2 shows a schematic outline of the processes proposed to underlie path-paradigm contour detection and curve-tracing. The presence or absence of 'chunks' is established in the base-representation (Figure 8-2a). In the incremental representation the colinearity of line segments is established according to Gestalt principles (Figure 8-2b). Accounts of curve-tracing and contour detection in path-paradigm experiments diverge after the formation of the incremental-representation. For path-paradigm experiments attention is drawn towards the presence of an interaction skeleton, though attention may already have been drawn towards chunks within the base representation, if any were activated. Whereas, for curve-tracing, the connectedness of elements is only established following further, attentive, processing (Figure 8-2c). The following sub-section discusses

neurophysiological evidence for the roles that feed-forward, horizontal and feed-back connections play in these processes.

A key point is that the observers' awareness of the presence of an embedded contour is cued by the presence of an interaction skeleton, or a 'chunk'. In order for this assertion to be accepted, the role of attention in path-paradigm experiments must be established. In Roelfsema *et al's* experiments contour tracing is instigated by cueing attention to the origin of the target contour with a shape or a letter. Clearly, contour tracing can be instigated without such cues (for example Uhlhaas, Silverstein, Phillips and Lovell, *submitted*). Otherwise, an observer might be aware that a contour was present within an image, but would not be able to instigate contour tracing if required to do so. The implication is that attention is drawn to the presence of an embedded contour in both curve-tracing and path-paradigm experiments. In the case of path-paradigm stimuli this is the very basis of perceptual "pop-out", i.e. that some anomalous property of the stimulus image catches the attention of the observer. For path-paradigm stimuli this is likely to be the presence of an extended interaction skeleton.

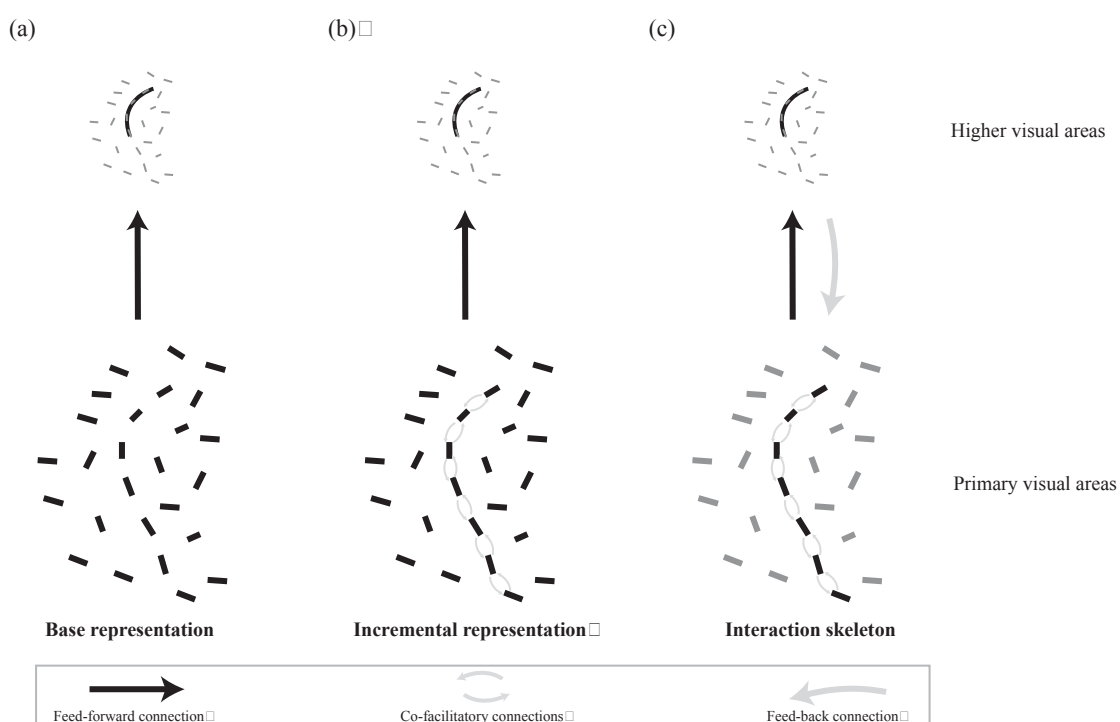


Figure 8-2. Processes underlying contour detection. (a) Feed-forward connections establish the base-representation of the stimulus viewed. Representations in higher-visual areas may also be activated, where matches to "chunked" contour portions are found. (b) An incremental representation is formed following the application of Gestalt grouping processes. (c) An interaction skeleton spans activated areas of the visual hierarchy. Attention is drawn to this interaction skeleton, however the 'connectedness' of the contour is only established through further attentive "curve-tracing".

Even though attention has been drawn towards the presence of an embedded contour it does not follow that tracing of the whole curve is mandatory. This process is only required to establish connectedness, in path-paradigm experiments it is sufficient to establish that a contour is present. Jolicoeur, Ullman and MacKay (1986, 1991) found that latencies in curve tracing tasks reflected the distance between markers on their curves, which suggests that curve tracing can halt upon achievement of a goal determined by the experimenter (or the observer). In other words the tracing of the whole curve is not mandatory. For path-paradigm stimuli, detection may occur once the interaction skeleton is established. Alternatively, tracing may be instigated but could halt as soon as the presence of a contour was established. As presentation times are generally short in path-paradigm experiments, this curtailment is generally enforced.

Whilst it is proposed that the detection of path-paradigm contours is achieved following the formation of an interaction skeleton, the process of chunking may enable the detection of contours within the base representation. In Roelfsema *et al's* conception of the base representation, many levels of the visual hierarchy are simultaneously activated. Consequently, some path-paradigm contours may be detected by the base-representation if they have been learned as a consequence of frequent observation (i.e. 'chunking'). The process of 'chunking' conflicts with the prevalent idea of a solely 'bottom-up' process of contour integration (Section 1.4.4). However, the process of chunking may be advantageous to accounts of performance in path-paradigm experiments. These advantages are discussed in Section 8.5.3.

8.5.1 Neurophysiological evidence for the roles of feed-forward, horizontal and feed-back connections

Feed-forward connections quickly carry activity upwards throughout the visual hierarchy. Gilbert (1994) and Lamme, Super and Spekreijse (1998) provide useful reviews of the role of these connections. Briefly, Lamme *et al.* propose that feed-forward connections are predominantly involved in definition of receptive field properties. Furthermore, they speculate that these connections may form the basis for fast behavioural responses, some occurring without awareness. This role would be consistent with the establishment of the base representation.

Many articles within the physiological and psychological literature speculate about the role that the horizontal collaterals of the primary visual cortex play in perceptual integration and figure-ground segregation (Gilbert and Wiesel, 1990; Field *et al.* 1993; Fregnac,

Bringuier and Chavane, 1996; Gilbert, 1998; Das and Gilbert, 1999; See Lamme, Super and Spekreijse, 1998 for a review). Section 1.4.2 has already described how correspondences in the spatial range of these connections and in the pattern of connectivity suggests that they are a suitable candidate for the neural substrate of the association field. In the current model of contour integration, the interaction skeleton is formed in primary visual areas when a network of co-facilitatory horizontal connections are enabled following the principles outlined by Field *et al.* (1993). However, Roelfsema *et al.* suggest that the interaction skeleton can be established at many layers of the visual hierarchy. Consequently, the horizontal connections of primary vision may only form part of the interaction skeleton.

There is a great deal of speculation about the role of connections that feed-back to the primary visual cortex (see Gilbert, 1998 or Lamme, Super and Spekreijse, 1998 for a review). From a cellular perspective it is suggested that feed-back connections have a role in the regulation of receptive field tuning (for example, Gilbert and Wiesel, 1990; Sillito and Jones, 1996; Kapadia, Westheimer and Gilbert, 1999). For example, Ito and Gilbert (1999) suggest that the contextual activity of horizontal connections can be gated by connections that feedback to area V1. From a systems perspective it is suggested that feed-back connections play a role in perceptual integration and figure-ground segregation (Fregnac, Bringuier, Chavane, Glaeser and Lorenceau, 1996).

Few studies have offered direct evidence for the role feed-back connections, though there are notable exceptions. For example, Lamme, Zipser and Spekreijse (1998) found that anaesthetisation stops figure-ground segregation whilst leaving receptive field properties unaffected. They conclude that that horizontal and feed-back connections must play a role in figure-ground segregation. Also, Hupe, James, Girard, Lomber, Payne and Bullier (2001) inactivated area MT of macaques by cooling. They found that inactivating MT had an influence upon firing rates in areas V1, V2 and V3 within the first 10ms recording bin. This led them to conclude that not only do feed-back connections have significant role in vision, but also that this role can influence receptive field responses within 10ms of the onset of feed-forward activation. Therefore monitoring response modulations over time may not be a reliable means of establishing the role of feed-back connectivity. Knierim and Van Essen (1992) have also found evidence of modulatory feed-back activity within 10ms of stimulus onset.

VanRullen, Delorme and Thorpe, (2001) suggest that connectivity patterns prior to primary vision may also contribute towards contour integration. They offer a model which makes use of differences in the temporal resolution of separate, pre-cortical, feed-forward channels. In their model contour integration is achieved with only feed-forward mechanisms. This suggests that pre-cortical processes could play an active role in the formation of the base representation.

8.5.2 Could visual experience influence the development of contour integration processes?

Section 1.4.3 has already discussed the possibility that the pattern of long-range horizontal connections, the hypothesised neural substrate of the association field, might be established through visual experience. As it is hypothesised that the interaction skeleton underlying contour integration is based upon the same neural substrate then it follows that the evidence presented in Section 1.4.3 applies equally to the interaction skeleton.

There is also evidence that the connections which feed-back to areas V1 and V2 are modified by visual experience. Burkhalter (1993) conducted a post-mortem examination of the development of feed-forward and feed-back connections in the primary visual cortex of infant brains. This investigation revealed that both types of connection develop in V1 and V2 shortly before birth, though they do not achieve maturity until sometime later. Whilst the feed-forward connections were apparently mature at four months of age, the feed-back connections remained immature at this time.

The phenomenon of chunking (Mahoney and Ullman, 1988; Vecera and Farah, 1997) suggests that frequently experienced shapes are learned so that they are activated earlier, during the feed-forward sweep of activity which establishes the base representation. This suggests that visual experience may not only establish the wiring of the contour detection mechanism, it may also play a role in the continued extension of the range of shapes which may be represented within the base representation.

What criteria would determine whether or not a particular contour shape would be chunked? Feldman (2001) hypothesised that grouping algorithms are selected on the basis of their previous success in identifying whether particular line segments belong to an edge. This hypothesis could be extended to account for the development of chunking processes. The acquisition of a chunk may be dependent upon the frequency with which a particular shape has been encountered and upon its previous success in establishing the

connectedness of contour segments. Studies of correlations within natural scenes have reported that the co-circularity rule provides a useful summary of co-dependencies in the orientations of line segments across space (Giesler, Perry, Super and Gallogly, 2001; Sigman, Cecchi, Gilbert and Magnasco, 2001, see 1.4.3). As both straight lines and smooth arcs are accommodated by the co-circularity rule and they are frequently experienced, it appears likely that these shapes would be heavily represented within the repertoire of chunked contour-shapes.

There are some similarities between what Field *et al.* (1993) described as path-detectors (Section 1.4.5) and the chunking mechanism currently under discussion. For Field *et al.*, these detectors were hard-wired to enable the detection of any conceivable contour configuration. Clearly, as they suggest, such mechanisms appear unlikely. Firstly, there are too many potential arrangements of contour segments that could form contours. Secondly, if tailored detectors existed for every contour shape then there would not be a fall off in detection rates with increases in path-angle and element spacing. However, a chunking mechanism need not fall foul of these criticisms. Firstly, the mechanism would only form representations of the most frequently experienced contour shapes. Secondly, the mechanism would not necessarily represent whole contours. It is likely that the chunking mechanism will only favour the smoother sections of contours.

8.5.3 Current experimental results

This section will attempt to evaluate how well the proposed Interactive contour integration mechanism accounts for contour integration evidence from this thesis and from the recent psychophysical literature.

8.5.3.1 Contour smoothness

Experiments 4 and 10 demonstrated that contours which featured longer arcs were detected at a higher rate than those featuring shorter arcs, Pettet, McKee and Grzywacz (1998) report similar findings. The process of chunking may provide an account of this effect. Smooth contours are more likely to have been experienced than randomly structured contours. Consequently it is more likely that these structures will be represented in higher visual areas, i.e. they will be added to the repertoire of chunked shapes. However, chunks representing smaller arcs may also be featured within the repertoire. Therefore, it is necessary to invoke some other property to explain higher detection rates for longer arcs. Two alternatives are apparent: Firstly, longer arcs may be more strongly represented within the chunked repertoire, or secondly, longer chunks may draw attention more reliably.

Whilst the chunking of regularly experienced contour shapes may account for a smoothness bias, it is also possible that the association field alone could favour smooth contours. Field *et al* (1993) do not state whether this would be the case, an entirely reasonable stance as their experiments only featured contours that had a random global structure. Appendix J demonstrates how the association field mechanism might be modified in order to account for smoothness effects by extending the range of lateral interactions.

8.5.3.2 CLE manipulation

Experiments 7 and 8 (Chapter 6) have shown that an increase in the degree of colinearity of background elements leads to poorer performance in path-paradigm experiments. Where the degree of colinearity is raised in background elements an extended interaction skeleton would be generated for both the target and the noise stimuli. If it is the presence of an interaction skeleton that draws the observers' attention towards the target-contour then creating a more diffuse skeleton may confound this process. However, it is likely that any other contour integration account would be likely to provide an equally plausible explanation of this finding: any mechanism that was specially tuned towards colinearity of contour elements would be interrupted by increased colinearity throughout the stimulus image.

8.5.3.3 Contour integration in the periphery

If the mechanisms underlying contour integration can be acquired, it may follow that the same processes could be acquired in both the central and peripheral fields. However, as visual acuity within the peripheral field is lower then it follows that the self-organising mechanisms within these fields would receive lower quality stimulation than that received by the central fields. As a consequence of this, the acquired mechanisms underlying contour integration in these regions would be relatively impaired when compared to those of the central field.

8.5.3.4 Temporal properties of contour integration

Psychophysical evidence suggests that the period stimuli must be visible prior to detection (the critical duration) increases as a function of path-angle (Hess, Beaudot and Mullen, 2001). No difference was found in the critical duration of jagged and smooth contours. They concluded that the same mechanism was likely to underlie the detection of both smooth and jagged contours and that this mechanism is initially tuned towards contours with a low path-angle. As critical duration was positively correlated with path-angle they

suggest this reflects a population-coded representation of contour membership, that might be guided by both intra- and extra-cortical feedback. A comparable effect of curvature has been reported in the curve-tracing literature, Jolicoeur *et al.* (1991) report that latencies in curve-tracing tasks are correlated with contour curvature. They argue that this reflects shifts in focal attention along the contour. In contrast Roelfsema *et al.* (2000) propose that the effect is due a correlation between curvature and the speed of the spread of *object-based* attention among the segments of the curve. If, as Roelfsema *et al.* propose, this temporal cost increases with curvature then it may be that connection strengths within the interaction skeleton are weaker or that they take longer to establish where curvature is greater. For path-paradigm stimuli, the establishment of the interaction skeleton is based upon colinearity. Increasing contour curvature may weaken the enabled connections within the interaction skeleton; with the consequence that it no longer clearly distinguishes between contour and background elements. This effect is illustrated in Figure 8-1, where increased contour curvature leads to lower levels of colinearity energy for the target-contour relative to the background. It is conceivable that responses based upon a weaker interaction skeleton would be delayed.

Experiment 4 and Pettet *et al.* (1998) have demonstrated that smoother contours are detected more reliably. However, Hess *et al.* (2001) report that there was no difference in the critical durations for smooth and jagged contours, suggesting that chunking may have no role in the detection of smooth contours. If chunking enabled the detection of some contour shapes within the base-representation, then it follows that these contour shapes ought to be detected earlier than those that are detected following the establishment of an interaction skeleton. However, where path-angles exceeded 20°, smooth contours were slightly more detectable (Hess, *personal communication*). This finding provides support for the chunking hypothesis, but it also lends itself to an association field model favouring smooth contours (this possibility is explored in Appendix J). Furthermore, Hess *et al.* (2001) report that the critical frequency for circles was twice that recorded for other stimuli. They concluded that circles have special status as objects, unlike contours. This finding also suggests a role for chunking.

8.5.3.5 Top-down influences upon contour detection

The Roelfsema *et al.* curve-tracing model allows attentional processes to influence activation at relatively early stages within the visual hierarchy. Scholte, Spekreijse and Roelfsema (2001) demonstrated that attention is required if an observer is required to trace

a contour that overlaps another contour. For example, they present evidence that the nature of the overlap between the curves determines the extent to which attention ‘leaks’ from the target curve to the distracter curve. If the association field mechanism were to be placed within the conceptual framework that they offer, then it is conceivable that there should also be evidence of attentional influences upon contour grouping processes.

However, it is generally accepted that performance in path-paradigm experiments is based upon pre-attentive mechanisms. Experiments tend to adopt brief exposure times and carefully control the fixation of observers (Field *et al.*, 1993; Kovacs and Julesz, 1993 and this thesis). Consequently, the author is unaware of any published articles that report on attentional influences upon the detection of path-paradigm contours. Indeed Polat and Bonneh (2000) argue that the mechanisms underlying contour integration and lateral masking paradigms share the same neural substrate. Evidence to support this claim comes from Wehrhahn, Li and Westheimer (1996) who found, in the lateral masking paradigm, that detection thresholds were increased as a function of length and decreased as a function of curvature. This equates to the effects of contour length and path-angle in path-paradigm stimuli.

Working within the lateral masking paradigm Ishai and Sagi (1995) have presented evidence that thresholds for low contrast visual targets are improved when observers were asked to imagine the presence of previously perceived masks. This facilitation occurred as much as 5 minutes after presentation of the original flanking mask. This suggests that low-level (monocular) visual processes can be influenced by high-level attentional mechanisms. Dufour (1996) has demonstrated that the influence that noise segments have upon decisions about line segment orientations can be changed by the nature of the task, this also suggests a top-down influence upon low level perceptual processes.

8.5.4 Questions raised by the interactive contour integration model

Burkhalter (1993) has demonstrated that feedback connections mature over a number of postnatal months. This may enable examination of the relative contributions of feedforward and feedback connections. However, visual acuity is relatively poor at this age. Consequently, it may be difficult to evaluate contour integration in young infants. An alternative approach could examine contour integration in observers that have had deprived early visual experience. For example, recent research has examined contour integration in amblyopes (Hess, McIlhagga and Field, 1997; Hess and Demanins, 1998; Pennefather, Chandna, Kovacs, Polat and Norcia, 1999; Kovacs, Polat, Pennefather, Chandna and

Norcia, 2000 and Simmers and Bex, 2001, See 1.4.3 for a summary). Future research could also examine contour integration in individuals that were born with cataracts which were later corrected. Contour integration in such individuals may be impaired immediately after corrective surgery, although improvement may be seen if neural pathways remain malleable, i.e. if the specification of relevant neural circuitry is not subject to a critical developmental period. If the bias favouring smooth contours was due to the process of chunking, then it follows that individuals with deprived visual experience and possibly infants should not exhibit this bias.

In section 8.5 it was stated that observers in path-paradigm experiments were not necessarily aware of the connectedness of contour elements, they were merely aware of the presence of a contour. This claim could be examined by requiring that observers state whether contour segments are connected whilst further ‘distracter’ contours are placed within the stimuli. Where overlaps between unconnected contours conform to those that led to attentional ‘leaks’ in the curve-tracing experiments (Scholte, Spekreijse and Roelfsema, 2001), one might expect errors in the decisions of observers. Where stimuli are presented for relatively brief times one might expect that correct responses should only follow those stimuli where target-contours are spatially discrete from neighbouring distracter contours. The hypothesis that the same interaction skeleton underlies contour integration and curve tracing would be supported by establishing whether the same stimulus properties influenced the success of an observer in both paradigms.

If chunking plays a role in observer performance in path-paradigm experiments then it may be possible to induce improvements in the detection rates for particular contour shapes. Conceivably, these improvements could be invoked for contours which have element spacing or path-angles which lie at the limits of normal path-paradigm performance levels. Observers could be exposed to contour shapes during training sessions which were given under the guise of an apparently unconnected experiment. Normal contour integration abilities could be assessed during a test phase that included contours that were random and those which had been previously viewed during training. If chunking had occurred then one might expect detection rates for the chunked contours that exceeded those expected for contours with an equivalent level of path-angle and element spacing. Alternatively, rather than training observers with novel contour configurations, detection rates for contours which form familiar shapes could also be investigated. For example, detection rates for letter or number-shaped contours could be examined. Such experiments would have to

ensure that element spacing and path-angle were carefully controlled. If it were confirmed that chunking aided contour detection, then further variables such as speed of acquisition, extinction and detection latencies could also be investigated using similar procedures.

Finally, animal models of contour integration offer a great deal of potential with regard to distinguishing between the different underlying processes. Multiple unit recordings and the anaesthetisation of target regions may enable the isolation of each processing component.

8.6 Conclusions

Contour integration experiments do not necessarily exclude the contributions of coarse-scale processes. However, by utilising contours with a smoother structure it is possible to minimise such contributions. It seems likely that the same processes underlie contour integration within the central and peripheral fields, at least to 13° eccentricity. It is speculated that these mechanisms are acquired through and continually tuned by visual experience. However, the lower quality of peripheral experience may account for the poorer contour integration performance often reported within the periphery.

References

- Beck, J., Rosenfeld, A. & Ivry, R. 1989 Line segregation. *Spatial Vision* **4**, 75-101.
- Bosking, W. H., Zhang, Y., Schofield, B. & Fitzpatrick, D. 1997 Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *Journal Of Neuroscience* **17**, 2112-2127.
- Braun, J. 1999 On the detection of salient contours. *Spatial Vision* **12**, 211-225.
- Burkhalter, A. 1993 Development of Forward and Feedback Connections Between Areas V1 and V2 of Human Visual-Cortex. *Cerebral Cortex* **3**, 476-487.
- Caelli, T. M., Preston, G.A.N. and Howell, E.R. 1978 Implications of spiral summation models for processes of contour perception: a geometric perspective. *Vision Research* **18**, 723-734.
- Dufour, A. 1996 Interactions between noise and signal segments at the decision level are task contingent. *Spatial Vision* **10**, 273-283.
- Feldman, J. 2001 Bayesian contour integration. *Perception & Psychophysics* **63**, 1171-1182.
- Field, D. J., Hayes, A. & Hess, R. F. 1993 Contour Integration by the Human Visual-system: Evidence for a Local "Association Field". *Vision Research* **33**, 173-193.
- Field, D. J., Hayes, A. & Hess, R. F. 1997 The role of phase and contrast polarity in contour integration. *Investigative Ophthalmology & Visual Science* **38**, 4643-4643.
- Freeman, W. T. & Adelson, E. H. 1991 The Design and Use Of Steerable Filters. *Ieee Transactions On Pattern Analysis and Machine Intelligence* **13**, 891-906.
- Fregnac, Y., Bringuier, V., Chavane, F., Glaeser, L. & Lorenceau, J. 1996 An intracellular study of space and time representation in primary visual cortical receptive fields. *Journal of Physiology-Paris* **90**, 189-197.
- Geisler, W. S., Perry, J. S., Super, B. J. & Gallogly, D. P. 2001 Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research* **41**, 711-724.
- Giersch, A., Boucart, M., Danion, J. M., Vidailhet, P. & Legrand, F. 1995 Effects Of Lorazepam On Perceptual Integration Of Visual Forms In Healthy-Volunteers. *Psychopharmacology* **119**, 105-114.

- Gigus, Z., and Malik, J. 1991 Detecting curvilinear structure in images. In *Computer Science Division (EECS)*: University of California at Berkeley.
- Gilbert, C. D. 1992 Horizontal Integration and Cortical Dynamics. *Neuron* **9**, 1-13.
- Gilbert, C. D. 1994 Circuitry, Architecture and Functional Dynamics of Visual-Cortex. In *Higher-Order Processing in the Visual System*, vol. 184, pp. 35-56. Chichester: John Wiley & Sons Ltd.
- Gilbert, C. D., Das, A., Ito, M., Kapadia, M. & Westheimer, G. 1996 Spatial Integration and Cortical Dynamics. *Proceedings Of the National Academy Of Sciences Of the United States Of America* **93**, 615-622.
- Gilbert, C. D. & Wiesel, T. N. 1983 Clustered intrinsic connections in the cat visual cortex. *Journal of Neuroscience* **3**, 1116-1133.
- Hebb, D. O. 1949 *The Organisation of Behavior*. New York: Wiley.
- Heitger, F., von der Heydt, R., Peterhans, E., Rosenthaler, L. & Kubler, O. 1998 Simulation of neural contour mechanisms: representing anomalous contours. *Image and Vision Computing* **16**, 407-421.
- Hess, R. & Field, D. 1999 Integration of contours: new insights. *Trends in Cognitive Sciences* **3**, 480-486.
- Hess, R. F., Beaudot, W. H. A. & Mullen, K. T. 2001 Dynamics of contour integration. *Vision Research*, **41**, 1023-1037.
- Hess, R. F., Campbell, R. F. & Greenhaigh, T. 1978 On the nature of the neural abnormality in human amblyopia: neural aberration and neural sensitivity loss. *Pflugers Archiv, fur die gesamte Physiologic*. **377**.
- Hess, R. F. & Dakin, S. C. 1997 Absence of contour linking in peripheral vision. *Nature* **390**, 602-604.
- Hess, R. F. & Dakin, S. C. 1999 Contour integration in the peripheral field. *Vision Research* **39**, 947-959.
- Hess, R. F., Dakin, S. C. & Field, D. J. 1998 The role of "contrast enhancement" in the detection and appearance of visual contours. *Vision Research* **38**, 783-787.
- Hess, R. F. & Demanins, R. 1998 Contour integration in anisometric amblyopia. *Vision Research* **38**, 889-894.

Hess, R. F. & Field, D. J. 1995 Contour Integration Across Depth. *Vision Research* **35**, 1699-1711.

Hess, R. F., Hayes, A. & Kingdom, F. A. A. 1997 Integrating contours within and through depth. *Vision Research* **37**, 691-696.

Hess, R. F., McIlhagga, W. & Field, D. J. 1997 Contour integration in strabismic amblyopia: the sufficiency of an explanation based on positional uncertainty. *Vision Research* **37**, 3145-3161.

Hirsch, J. A. & Gilbert, C. D. 1991 Synaptic physiology of horizontal connections in the cat's visual cortex. *Journal of Neuroscience* **11**, 1180-1809.

Hoyer, P. O. & Hyvärinen, A. 2002 A multi-layer sparse network learns contour coding from natural images. *Vision Research* *in press*.

Hubel, D. & Weisel, T. N. 1962 Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology* **195**, 215-243.

Hupe, J. M., James, A. C., Girard, P. & Bullier, J. 2001 Response modulations by static texture surround in area V1 of the macaque monkey do not depend on feedback connections from V2. *Journal of Neurophysiology* **85**, 146-163.

Hupe, J. M., James, A. C., Girard, P., Lomber, S. G., Payne, B. R. & Bullier, J. 2001 Feedback connections act on the early part of the responses in monkey visual cortex. *Journal of Neurophysiology* **85**, 134-145.

Ishai, A. & Sagi, D. 1995 Common Mechanisms Of Visual-Imagery and Perception. *Science* **268**, 1772-1774.

Ito, M. & Gilbert, C. D. 1999 Attention modulates contextual influences in the primary visual cortex of alert monkeys. *Neuron* **22**, 593-604.

Kapadia, M. K., Ito, M., Gilbert, C. D. & Westheimer, G. 1995 Improvement In Visual Sensitivity By Changes In Local Context - Parallel Studies In Human Observers and In V1 Of Alert Monkeys. *Neuron* **15**, 843-856.

Kiorpes, L., Bassin, S. A. & Movshon, J. A. 2001 Development of contour integration. *Investigative Ophthalmology & Visual Science* **42**, 654.

Kiper, D. C., Gegenfurtner, K. R. & Movshon, J. A. 1996 Cortical oscillatory responses do not affect visual segmentation. *Vision Research* **36**, 539-544.

- Knierem, J. J. & Van Essen, D. C. 1992 Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *Journal of Neurophysiology* **67**, 961-980.
- Kohler, W. 1929 *Gestalt Psychology*. New York: Horace Liverwright.
- Kovacs, I. & Julesz, B. 1993 A Closed Curve Is Much More Than an Incomplete One - Effect Of Closure In Figure Ground Segmentation. *Proceedings Of the National Academy Of Sciences Of the United States Of America* **90**, 7495-7497.
- Kovacs, I. & Julesz, B. 1994 Perceptual Sensitivity Maps Within Globally Defined Visual Shapes. *Nature* **370**, 644-646.
- Kovacs, I., Polat, U. & Norcia, A. M. 1996 Breakdown Of Binding Mechanisms In Amblyopia. *Investigative Ophthalmology & Visual Science* **37**, 3078-3078.
- Kovacs, I., Polat, U., Pennefather, P. M., Chandna, A. & Norcia, A. M. 2000 A new test of contour integration deficits in patients with a history of disrupted binocular experience during visual development. *Vision Research* **40**, 1775-1783.
- Lamote, C. & Wagemans, J. 1999 Rapid integration of contour fragments: From Simple Filling-in to Parts-based shape Description. *Visual Cognition* **6**, 345-361.
- Li, C. Y. & Li, W. L. 1994 Extensive Integration Field Beyond The Classical Receptive Field Of Cat's Striate. *Vision Research* **34**, 2337-2355.
- Li, Z. P. 1998 A neural model of contour integration in the primary visual cortex. *Neural Computation* **10**, 903-940.
- Mahoney, J. V. & Ullman, S. 1988 Image chunking defining spatial building blocks for scene analysis. In *Computational processes in human vision: an interdisciplinary perspective* (ed. Z. Pylyshyn), pp. 169-209. Norwood NJ: Ablex.
- Moulden, B. 1994 Collator Units - 2nd-Stage Orientational Filters. *Ciba Foundation Symposia* **184**, 170-184.
- Nugent, A. K., Keswani, R. N., Woods, R. L. & Peli, E. 2001 Contour integration in the peripheral field of normal and low vision observers. *Investigative Ophthalmology & Visual Science* **42**, 3287.
- Parent, P. & Zucker, S. 1989 Trace inference, curvature consistency and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**, 823-839.

- Parker, A. J. & Hawken, M. J. 1988 Two-dimensional spatial structure of receptive fields in monkey striate cortex. *Journal of the Optical Society of America A* **5**, 589-605.
- Pennefather, P. M., Chandna, A., Kovacs, I., Polat, U. & Norcia, A. M. 1999 Contour detection threshold: repeatability and learning with 'contour cards'. *Spatial Vision* **12**, 257-266.
- Pettet, M. W., McKee, S. P. & Grzywacz, N. M. 1996 Smoothness Constraints Long-Range Interactions Mediating Contour- Detection. *Investigative Ophthalmology & Visual Science* **37**, 4368-4368.
- Pettet, M. W., McKee, S. P. & Grzywacz, N. M. 1998 Constraints on long range interactions mediating contour detection. *Vision Research* **38**, 865-879.
- Phillips, W. A. & Singer, W. 1997 In search of common foundations for cortical computation. *Behavioral and Brain Sciences* **20**, 657-722.
- Polat, U. & Bonneh, Y. 2000 Collinear interactions and contour integration. *Spatial Vision* **13**, 393-401.
- Polat, U. & Sagi, D. 1992 Lateral Interactions Between Spatial Channels - Suppression and Facilitation Revealed By Lateral Masking Experiments. *Investigative Ophthalmology & Visual Science* **33**, 1345-1345.
- Polat, U. & Sagi, D. 1994 Spatial Interactions In Human Vision - From Near to Far Via Experience-Dependent Cascades Of Connections. *Proceedings Of the National Academy Of Sciences Of the United States Of America* **91**, 1206-1209.
- Rockland, K. S. & Lund, J. S. 1983 Intrinsic Laminar Lattice Connections in Primate Visual-Cortex. *Journal of Comparative Neurology* **216**, 303-318.
- Roelfsema, P. R., Engel, A. K., Konig, P. & Singer, W. 1996 The role of neuronal synchronization in response selection: A biologically plausible theory of structured representations in the visual cortex. *Journal Of Cognitive Neuroscience* **8**, 603-625.
- Roelfsema, P. R., Lamme, V. A. F. & Spekreijse, H. 2000 The implementation of visual routines. *Vision Research* **40**, 1385-1411.

- Roelfsema, P. R., Scholte, H. S. & Spekreijse, H. 1999 Temporal constraints on the grouping of contour segments into spatially extended objects. *Vision Research* **39**, 1509-1529.
- Scholte, H. S., Spekreijse, H. & Roelfsema, P. R. 2001 The spatial profile of visual attention in mental curve tracing. *Vision Research* **41**, 2569-2580.
- Sigman, M., Cecchi, G. A., Gilbert, C. D. & Magnasco, M. O. 2001 On a common circle: Natural scenes and Gestalt rules. *Proceedings of the National Academy of Sciences of the United States of America* **98**, 1935-1940.
- Sillito, A. M. & Jones, H. E. 1996 Context-dependent interactions and visual processing in V1. *Journal of Physiology-Paris* **90**, 205-209.
- Simmers, A. J. & Bex, P. J. 2001 Deficit of visual contour integration in dyslexia. *Investigative Ophthalmology & Visual Science* **42**, 2737-2742.
- Simoncelli, E. P. & Olshausen, B. A. 2001 Natural image statistics and neural representation. *Annual Review of Neuroscience* **24**, 1193-1216.
- Singer, W. & Gray, C. M. 1995 Visual Feature Integration and the Temporal Correlation Hypothesis. *Annual Review of Neuroscience* **18**, 555-586.
- Smit, J. T. S., Vos, P.G. and Van Oeffelen, M.P. 1985 The perception of a dotted line. *Spatial Vision* **12**, 163-177.
- The MathWorks, Inc. 1984-2000 *Matlab*. Natick, MA. 01760-2098, USA.
- Ts'o, D., Gilbert, C. D. & Weisel, T. N. 1986 Relationships between horizontal connections and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *Journal of Neuroscience* **6**, 1160-1170.
- Tusa, R. J., Palmer, L. A. and Rosenquist, A. C. 1978 The retinotopic organization of Area 17 (striate cortex) in the cat. *Journal of Comparative Neurology* **177**, 213-236.
- Uhlhaas, P. J., Silverstein, S. M., Phillips W. A. and Lovell, P. G. 2002 Visual Context, Schizotypy and Thought Disorder. *Journal of Abnormal Psychology* .
- Uttal, W. R. 1975 *An autocorrelation theory of form detection*. Hillsdale, N.J.: Lawrence Erlbaum Associates.

- VanRullen, R., Delorme, A. & Thorpe, S. J. 2001 Feed-forward contour integration in primary visual cortex based on asynchronous spike propagation. *Neurocomputing* **38**, 1003-1009.
- Vecera, S. P. & Farah, M. J. 1997 Is visual image segmentation a bottom-up or an interactive process? *Perception & Psychophysics* **59**, 1280-1296.
- Visser, T. A. W. & Enns, J. T. 2001 The role of attention in temporal integration. *Perception* **30**, 135-145.
- Vonderheydt, R. & Peterhans, E. 1989 Mechanisms of contour perception in monkey visual cortex 1. Lines of pattern discontinuity. *Journal of Neuroscience* **9**, 1731-1748.
- Watt, R. J. 1991 *Understanding Vision*. London: Academic Press Limited.
- Wehrhahn, C., Li, W. & Westheimer, G. 1996 Patterns that impair discrimination of line orientation in human vision. *Perception* **25**, 1053-1064.
- Wertheimer, M. (ed.) 1923 *Untersuchungen zur Lehre von der Gestalt II (translated)*. A source book of Gestalt psychology. London: Routledge & Kegan Paul.
- Wolfe, J. M., O'Neill, P. & Bennett, S. C. 1998 Why are there eccentricity effects in visual search? Visual and attentional hypotheses. *Perception & Psychophysics* **60**, 140-156.
- Yen, S. C. & Finkel, L. H. 1998 Extraction of perceptually salient contours by striate cortical networks. *Vision Research* **38**, 719-741.

Table of Contents

APPENDIX A	SIMPLE-FILTER MODEL PERFORMANCE WITH HIGHER RESOLUTION IMAGES	A-1
A.1.1	Introduction	A-1
A.2	SFM PERFORMANCE WITH MANIPULATIONS OF PATH-ANGLE AND PHASE.	A-1
A.2.1	Methodology.....	A-1
A.2.2	Results	A-2
A.2.3	Discussion.....	A-4
A.3	SFM PERFORMANCE FOR MANIPULATIONS OF PATH-ANGLE IN PHASE-ALTERNATED STIMULI OVER A RANGE OF FILTER-SCALES.	A-4
A.3.1	Introduction	A-4
A.3.2	Methods	A-4
A.3.3	Results	A-5
A.3.4	Discussion.....	A-7
A.4	SFM PERFORMANCE WITH NARROW FILTERS	A-7
A.4.1	Introduction	A-7
A.4.2	Methodology.....	A-7
A.4.3	Results	A-7
A.4.4	Discussion.....	A-8
A.5	MANIPULATIONS OF CONTOUR ARC-LENGTH AND SFM PERFORMANCE	A-9
A.5.1	Introduction	A-9
A.5.2	Methodology.....	A-9
A.5.3	Results	A-9
A.5.4	Discussion.....	A-10
A.6	GENERAL DISCUSSION	A-10
APPENDIX B	DETECTING PHASE-ALIGNED AND PHASE-ALTERNATED CONTOURS IN AN EXTENDED SIMPLE-FILTER MODEL.....	B-12
B.1	INTRODUCTION.....	B-12
B.2	METHODOLOGY.....	B-12
B.3	RESULTS.....	B-14
B.4	DISCUSSION.....	B-14
APPENDIX C	<i>CHERRY-PICKED</i> FILTER DETAILS.....	C-1
APPENDIX D	POST-HOC COMPARISONS OF THE FILTER USED IN THE CURRENT THESIS AND THAT EMPLOYED BY HESS AND DAKIN (1999).....	D-1
APPENDIX E	COMPARING CO-LINEARITY ENERGY IN RANDOM AND CLE-MANIPULATED IMAGES.....	E-1
E.1	INTRODUCTION.....	E-1
E.2	STIMULI.....	E-1
E.3	RESULTS.....	E-1
APPENDIX F	ENHANCED DETECTION OF SMOOTH CONTOURS WITH THE YEN AND FINKEL ASSOCIATION FIELD MODEL	F-1
F.1	INTRODUCTION.....	F-1
F.2	STIMULI.....	F-1
F.3	RESULTS.....	F-2

Appendix A Simple-filter model performance with higher resolution images

A.1.1 Introduction

All simulations reported in the current thesis used stimuli with a relatively low resolution, the wavelength of Gabor patches in all stimuli was only four pixels. It is possible that such stimuli may confer an additional advantage to the simple-filter model (SFM). In order to verify that this was not the case, a selection of simulations were repeated with higher resolution images. Images were scaled-up by a factor of five, giving a wavelength of twenty pixels for the sinewave forming the Gabor patch. If the results were broadly similar to those reported in simulations 1-6, then it can be concluded that the results were not confounded by the coarsely sampled nature of the stimuli. Conversely, if the performance for the SFM model were found to be poor then it could be concluded that the coarsely sampled nature of the stimuli featured in simulations 1-6 had indeed confounded the simulations 1-6.

Because of the increased processing overhead involved in filtering higher-resolution images, performance will only be examined at specific filter-scales. Initially, those filters which had given rise to the best overall performance and the best fit between the SFM and human observers was examined.

A.2 SFM performance with manipulations of path-angle and phase.

A.2.1 Methodology

Simulation 1 reported that the performance of the SFM was largely indistinguishable¹ from that of human observers over a range of path-angle values and where the phase of Gabor patches was manipulated. The current simulation replicates these simulations with stimulus images scaled by a factor of five. Consequently, the 256^2 images used in the original simulation become 1280^2 . All other scale dependent properties, such as contour and background element spacing were scaled accordingly.

Table A-1. Stimulus scale properties for the current simulation.

	Pixels	Lambda
Gabor patch full-wave frequency	20	1λ
Background and contour minimal nearest-neighbour spacing	80	4λ
Contour element spacing	80	4λ
Overall image size	1280 x 1280	$64 \times 64\lambda$

In order to make the simulations more time-efficient only those filter-scales that had performed well in simulation 1 were tested here (Table A-2).

¹ Excepting a drop in performance for phase-alternated stimuli where the path-angle value was low, this finding was explored in simulation 3.

Table A-2. The four filter-scales which had previously given rise to the best overall performance and the closest fit between human observers and the SFM were tested again with the higher-resolution images. The scale of the original filters is shown in the 'low-resolution' columns, whilst the corresponding scaled-up filters are shown in the 'high resolution' columns.

	Low-resolution filters		High-resolution filters	
	width	length	width	length
Best performance phase-aligned	2.50 σ	4.00 σ	12.50 σ	20.00 σ
Best performance phase-alternated	2.50 σ	3.50 σ	12.50 σ	17.50 σ
Closest observer/sfm fit phase-aligned	5.00 σ	6.00 σ	25.00 σ	30.00 σ
Closest observer/sfm fit phase-alternated	0.50 σ	4.00 σ	2.50 σ	20.00 σ

A.2.2 Results

Figure A-1 presents to SFM performance as a function of Gabor element phase and contour path-angle. The upper plot (A) compares the performance of those filters that gave rise to the best overall performance in the low-resolution images with the same filters in scaled-up images. The lower plot (B) shows performance with those filters which gave rise to the closest fit between the human observer performance and that of the SFM in the low-resolution and scaled-up images.

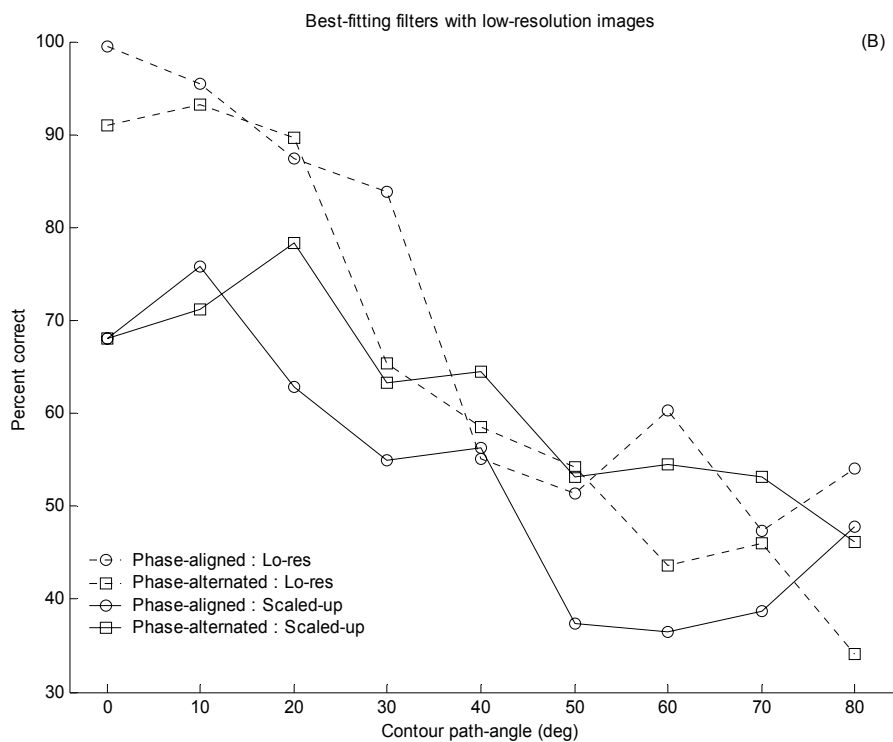



Figure A-1. A comparison of the SFM's performance for the original (low-resolution) and scaled-up stimuli. (A) Performance with the best-performing filters. (B) Performance with the best-fitting filters. See Table A-2 for the exact filter scales.

A.2.3 Discussion

Clearly the scaling-up of the filters and stimuli had an adverse effect on the performance levels at the specific filter scales examined. For both the phase-aligned and phase-alternated stimuli there is a drop of around 20% in performance. This result suggests that the coarsely-sampled nature of the original stimuli may have inadvertently aided the performance of the simple-filter model. However, the performance levels of the SFM still exceed 75% for both phase-aligned and phase-alternated stimuli. Consequently, even in high-resolution images, the manipulation of Gabor patch phase does not successfully control-for the contributions of coarse-scale mechanisms. Sections A.3 and A.4 examine whether the performance of the SFM is consistently lower than that of human observers over a range of filter-scales and elongation ratios.

A.3 SFM performance for manipulations of path-angle in phase-alternated stimuli over a range of filter-scales.

A.3.1 Introduction

The results reported in the previous sub-section suggest that detection rates for target stimuli decreased by around 20% when the resolution of stimulus images was increased. One may conclude that the reported performance levels of the simple-filter model were exaggerated because of the use of low-resolution images. However, the results of simulation's 1-3 revealed that the performance of the simple-filter model varied by a large amount as the size and elongation ratio of filters was manipulated. Consequently, the possibility remains that at particular filter scales the detection rates for high-resolution target stimuli may still be close to that reported for human observers. This is clearly the case for phase-aligned stimuli; the performance of the $(12.50\sigma \times 20.00\sigma)$ filter shown in Figure A-1 (upper-plot) exceeds that of human subjects.

A.3.2 Methods

The results of simulations 1-3 demonstrated that performance levels for phase-aligned stimuli were relatively high over a large range of filter-scales and elongations (see Figure 3-7). In contrast, for phase-alternated stimuli peaks of performance were relatively small and restricted to particular areas of 'filter-space' (see Figure 3-8). As a consequence, in order to reduce the processing overhead, the performance of the simple-filter model was initially only examined with high-resolution versions of the phase-alternated stimuli. Filter-scales were examined which corresponded to the highest-peaks of performance achieved for low-resolution phase-alternated stimuli². Filter widths were varied between 2.5σ to 17.5σ in steps of (2.5σ) and filter lengths were varied between 17.5, 20.0 and 22.5σ . Performance levels were examined at all combinations of these filter lengths and widths. The stimuli and other aspects of the methodology were unchanged from the previous sub-section. Once the optimal range of filter-scales was found for phase-alternated stimuli, performance levels were examined with the same filters in conjunction with phase-aligned stimuli.

² This range of filter-scales encapsulates the largest peak found in the surface plots shown in Figure 3-8 of the thesis.

A.3.3 Results

Figure A-2 shows the overall performance levels of the simple-filter model as a function of filter-scale. There is a clear peak at a filter length of 17.5σ spanning widths of around $7.5\text{-}12.5\sigma$. The performance levels of the SFM in conjunction with filters $W7.5\sigma \times L17.5$, 20 and 22.5σ were selected for closer examination; with both phase-aligned and phase-alternated stimuli.




Figure A-2. Detection rates for phase-alternated stimuli shown as a function of filter-scale. The values in brackets present the corresponding values in the low-resolution filters.

Figure A-3 shows the performance of the SFM with the selected filter-scales. For phase-aligned stimuli the detection rates peak at the lowest path-angle level and fall as a function of path-angle. However, while detection rates have similar overall shape for phase-alternated stimuli, detection rates are specifically lower for the lowest path-angle values, i.e. zero and ten degrees.





Figure A-3. Contour detection rates as a function of filter-scale, path-angle and Gabor patch phase. The upper-plot shows results for phase-aligned stimuli, the lower plot shows results for phase-alternated stimuli.

A.3.4 Discussion

Performance levels tend to fall as a function of path-angle for both phase-aligned and phase-alternated stimuli. However, the SFM model is specifically impaired in the detection of phase-alternated stimuli with low levels of path-angle. The reason for this dip in the straightest contours has already been explored in Simulation 3. Simulation 3 demonstrated that low-path-angle phase-alternated contours are best detected by narrow filters. These filters enable the linking of the same-sign regions of odd-symmetric Gabor patches. The following section describes simulations utilising such narrow filters in conjunction with higher-resolution stimulus images.

A.4 *SFM performance with narrow filters*

A.4.1 Introduction

The previous simulation reported that the simple-filter model was specifically impaired in the detection of phase-alternated stimuli when the path-angle of contours was low. Simulation 3 reported that such stimuli were successfully detected when very narrow filters were deployed. The following simulation examines whether this remains the case with higher-resolution images.

A.4.2 Methodology

The methodology remains unchanged. Filter scales were fixed at 2.50σ width and 17.50σ length. Performance was examined with phase-aligned and phase-alternated stimuli. The stimuli employed were the same as those described in Section A.2.1.

A.4.3 Results

Detection rates for target stimuli are shown in Figure A-4, the upper plot represents rates for phase-aligned stimuli and the lower phase-alternated. Detection rates for phase-aligned stimuli correspond between the SFM and human observers where path-angles are low ($0-30^\circ$), however this correspondence fails with higher path-angles. For phase-alternated stimuli the detection rates for the SFM peak at the lowest path-angle and fall as a function of path-angle. However, the slope of the function is much shallower for the SFM than it is for the human observers.

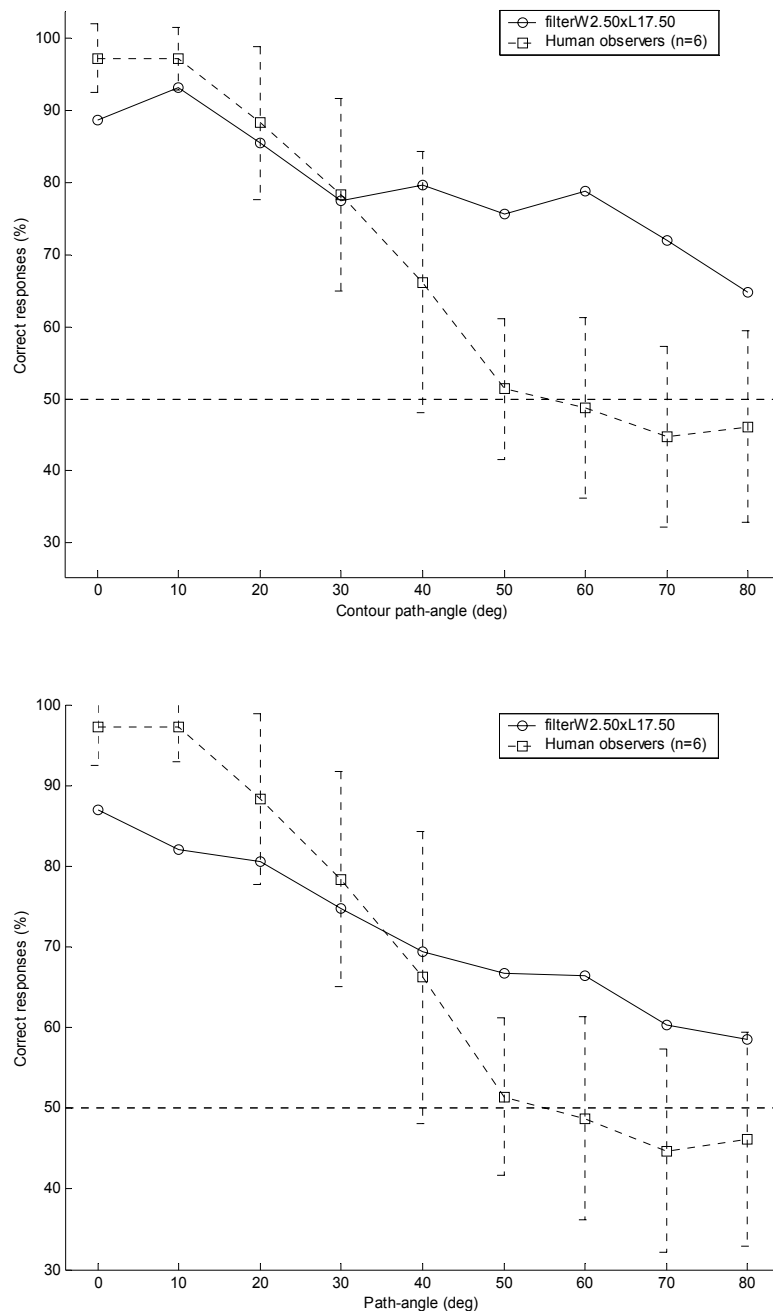


Figure A-4. Contour detection rates as a function of Gabor patch phase and contour path-angle. The upper plot presents the detection rates for phase-aligned stimuli. The lower plot presents the results for phase-alternated stimuli. The dashed-lines represent the detection rates for human observers with similar, though not identical, stimuli.

A.4.4 Discussion

As reported in Simulation 1 and 3, for narrow filters the dip in detection rates of the SFM for phase-alternated stimuli with low-path-angles disappears. The detection rates fall smoothly as a function of path-angle. However the slope of this function is much shallower than it is for human observers with similar stimuli. With phase-aligned stimuli detection rates correspond to those of human observers where path-angles are low but do not fall towards chance as path-angles increase.

A.5 Manipulations of contour arc-length and SFM performance

A.5.1 Introduction

A key claim of the simulation chapters in the current thesis is that manipulations of Gabor patch phase do not provide an adequate control of coarse-scale processes. The SFM is able to detect target stimuli regardless of Gabor patch phase. Simulation 4 demonstrated that manipulations of contour arc-length provide a more effective means of ensuring that coarse-scale processes do not contribute significantly to the detection of embedded contours; detection rates for human observers rise as a function of arc-length whereas they fall for the SFM. This section examines whether the effect of arc-length upon the performance of the SFM remain constant when the resolution of images is increased.

A.5.2 Methodology

Stimulus-scale parameters from Simulation 4 were increased by a factor of 5. Current stimulus properties are shown in the Table A-3.

Table A-3. Stimulus scale properties for the high resolution arc-length stimuli.

	Pixels	Lambda
Gabor patch full-wave frequency	20	1λ
Background and contour minimal nearest-neighbour spacing	80	4λ
Contour element spacing	80	4λ
Overall image size	1280 x 1280	$64 \times 64\lambda$

In Simulation 4 stimulus scale was varied from $0.5-10\sigma$ in steps of 0.5σ for both length and width. A corresponding range of filter scales was utilised in the current resolution ($2.50-50.00\sigma$). However, scale was varied in steps of 11.87σ (corresponding to 2.37σ in the low-resolution images); consequently lengths and widths had all combinations of the values: - 2.50, 14.38, 26.25, 38.13 and 50.00. This coarser-scale sampling of filter-scales was adopted due to the increased processing time required for higher-resolution images.

A.5.3 Results

The best performing elongated filters were identified for each level of arc-length and CLE manipulation. The methodology for identifying these filters is summarised in Section 5.1.2.4 of Chapter 5.




Figure A-5. Performance of the simple-filter model as a function of contour arc-length. Each point reflects the maximum detection rate achieved across all elongated filters. The scale of the filter giving rise to the performance level identified by each point is indicated in Table A-4. This plot should be compared to Figure 7-7 (left) of the thesis.

Table A-4. Lengths and widths of best-performing elongated filters.

CLE Mode	Arc-length	Percent Correct	Filter Width (σ)	Filter Length(σ)
Random	1	81.50	2.50	26.25
Random	2	94.50	14.38	26.25
Random	3	81.50	14.38	26.25
Random	5	68.00	14.38	26.25
Random	9	72.50	2.50	26.25
Minimised	1	91.50	2.50	38.13
Minimised	2	92.00	14.38	26.25
Minimised	3	76.50	14.38	26.25
Minimised	5	77.50	2.50	26.25
Minimised	9	68.50	2.50	50.00

A.5.4 Discussion

The performance of the SFM is consistent with the levels of performance reported in simulations 4 and 6; the detection of target stimuli varies as an inverse function of contour arc-length. Consequently, it is accepted that the coarse-scale of the original stimulus images did not inadvertently confound the reported experimental results.

A.6 General discussion

The resolution of stimulus images has had an influence upon the correspondence between specific filter-scales and the SFMs detection rates for target images. Where images and filters were scaled accordingly the detection rates for low and high resolution images varied. However, while performance at specific filter-scales varied, performance levels still exceeded chance by a significant degree at many filter-scales. Consequently, regardless of the resolution of stimulus

images, it is possible that coarse-scale processes may contribute towards the target-image detection rates of human observers in path-paradigm experiments.

Appendix B Detecting phase-aligned and phase-alternated contours in an extended simple-filter model

B.1 Introduction

An inherent problem with the simple-filter model (SFM) as an alternative account of human path-paradigm performance is the fact that the SFM is fundamentally a line-detecting mechanism whereas phase-alternated contours are composed of edges of alternating sign. The detection of phase-alternated contours within stimuli is achieved through a slight misalignment of narrow filters with the orientation of the contour (Simulation 3; Section 3.6). The current section examines an extended version of the SFM which initially deploys an edge-detecting filter; as a consequence the orientation and position properties of a contour can be retained, while the phase properties are eliminated. This means the disruption caused by phase-alternated stimuli to the detection of contours by the SFM is prevented. A key point is that this extension of the SFM does not integrate the outputs of differently oriented filters. Consequently, whilst the process is made more sophisticated, the processing still differs from an association-field based process which relies upon intra-orientation linking.

B.2 Methodology

The first stage of the current model involves convolving stimuli with an odd-symmetric filter. This filter was created by convolving phase-aligned and phase-alternated Gabor patches (see the sub-diagram in the right-hand column of Figure B-1). Subsequently, luminance values in the filtered-images are rectified. These two processes serve to retain orientation and location information while discarding Gabor patch phase information. Subsequent stages exactly match those described for the SFM (see Chapter 3). The orientation of the Gabor patches and all filters were matched. The scale of the Gabor patches that formed the filter were the same as that employed in the stimuli. The scale of the filter that formed the basis of the conventional SFM processing had a width of 12.5σ and a length of 20σ . As with simulations described throughout this thesis, all stimulus images were processed over a range of 10 filter-orientations.

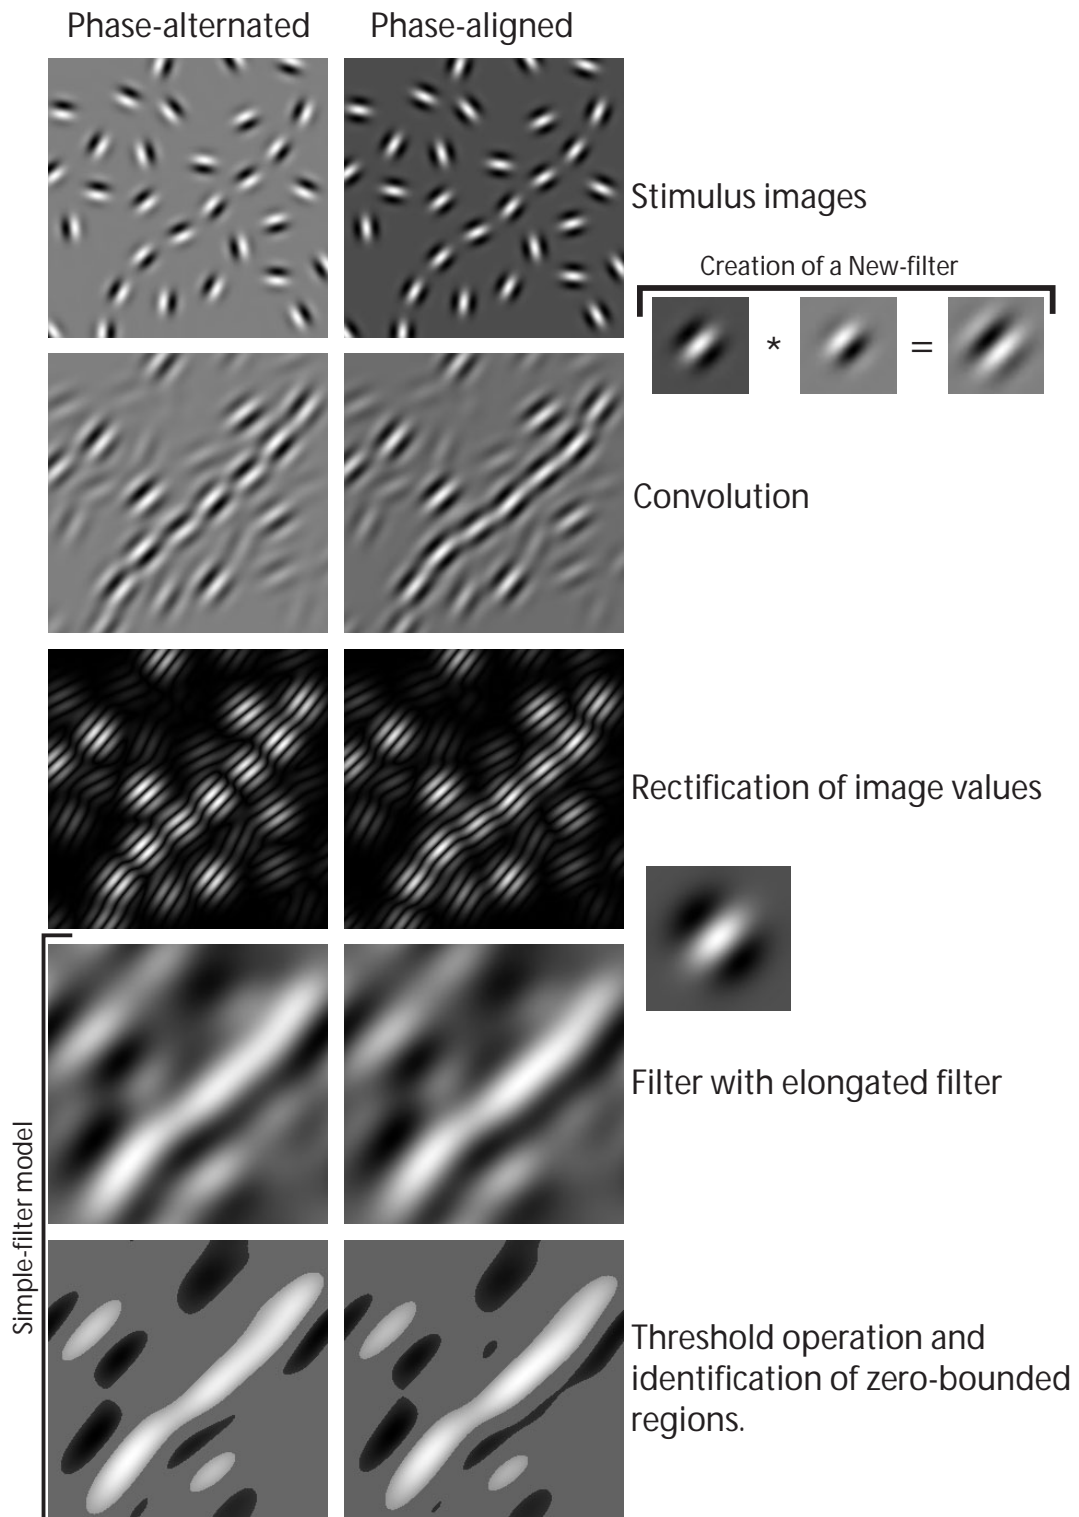


Figure B-1. Extending the simple-filter model: The first stage involves filtering stimuli with an edge-detecting filter. This is followed by a rectification stage. Subsequent stages are unchanged from the SFM described in Chapter 3.

B.3 Results

Figure B-2 shows detection rates for the extended simple-filter model as a function of contour path-angle and Gabor patch phase. With the exception of an anomalous peak at 70°, for both phase-aligned and phase-alternated stimuli detection rates fall as a function of contour path-angle.




Figure B-2. Fitted target stimulus detection rates as a function of Gabor patch phase and contour path-angle. For comparison, the dashed-line presents the human observer detection.

B.4 Discussion

The extended simple-filter model presented here successfully simulates observer performance with phase-alternated and phase-aligned stimuli. Furthermore, the model does not need to deploy different filters as the phase of Gabor patches within stimuli is manipulated – this is a problem of the simple-filter model discussed in Section 3.7.6.

Appendix C *Cherry-picked filter details*

Data for 'cherry-picked' elongated filters

Arc-Length	Spacing	Percent correct	Filter width	Filter length
1	4.00	100	2.50	3.50
1	4.66	100	2.50	4.50
1	5.32	100	3.50	5.50
1	6.00	100	4.00	8.50
2	4.00	100	4.00	4.50
2	4.66	88	4.50	5.00
2	5.32	89	4.00	7.50
2	6.00	81	1.00	5.50
4	4.00	88	3.00	3.50
4	4.66	86	5.00	8.00
4	5.32	81	0.50	3.50
4	6.00	83	1.50	8.00
8	4.00	88	1.50	2.00
8	4.66	85	4.00	5.50
8	5.32	81	1.50	2.50
8	6.00	81	2.00	3.00
15	4.00	83	1.50	4.00
15	4.66	75	8.50	9.00
15	5.32	77	1.00	6.00
15	6.00	79	1.50	9.00

Data for 'cherry-picked' broad filters

Arc-Length	Spacing	Percent correct	Filter width	Filter length
1	4.00	100	4.00	1.50
1	4.66	100	5.50	2.00
1	5.32	88	8.50	3.00
1	6.00	84	8.50	3.00
2	4.00	100	5.00	4.50
2	4.66	93	3.50	1.00
2	5.32	81	5.00	2.00
2	6.00	73	2.50	0.50
4	4.00	96	3.00	1.00
4	4.66	85	3.50	1.00
4	5.32	90	6.00	2.00
4	6.00	82	7.00	2.50
8	4.00	91	3.00	1.00
8	4.66	79	4.00	1.50
8	5.32	79	4.00	0.50
8	6.00	82	5.00	1.00
15	4.00	100	3.00	1.00
15	4.66	77	3.50	1.00
15	5.32	79	4.00	1.50
15	6.00	85	5.50	1.50

Data for 'cherry-picked' filters drawn from both wide and elongated filter populations

Arc-Length	Spacing	Percent correct	Filter width	Filter length
1	4.00	100	4.00	1.50
1	4.66	100	5.50	2.00
1	5.32	100	3.50	5.50
1	6.00	100	4.00	8.50
2	4.00	100	4.00	4.50
2	4.66	93	3.50	1.00
2	5.32	89	4.00	7.50
2	6.00	81	1.00	5.50
4	4.00	96	3.00	1.00
4	4.66	86	5.00	8.00
4	5.32	90	6.00	2.00
4	6.00	83	1.50	8.00
8	4.00	91	3.00	1.00
8	4.66	85	4.00	5.50
8	5.32	81	1.50	2.50
8	6.00	84	3.00	3.00
15	4.00	100	3.00	1.00
15	4.66	77	3.50	1.00
15	5.32	83	3.00	3.00
15	6.00	85	5.50	1.50

Appendix D Post-hoc comparisons of the filter used in the current thesis and that employed by Hess and Dakin (1999)

The Matlab M.files presented below enable an ad-hoc identification of a 2nd derivative of a Gaussian filter that is the closest match to the DoG filter employed by Hess and Dakin. For example, if a 4.5σ DoG filter is specified then the “matchDoGtoVgauss” program returns a filter of dimensions width = 7.74σ x length = 13.44σ . The figure below presents both of these filters.




Figure D-1. [left] 10σ DoG filter. [right] 2nd derivative of a Gaussian filter, width = 7.74σ , length = 13.44σ . The bars above each filter represent estimates of each filters wavelength.

```
function matchDoGtoVgauss (DoGsigma)
% Generate a Hess and Dakin DoG filter and then try to find a match from Vgauss...
%
if ~exist('DoGsigma')
    allsigmas=[1:.5:10];
else
    allsigmas=DoGsigma;
end

allsigmas=allsigmas(:)';

allscales = [];

figure;
ratplot = subplot(1,3,1);
elplot = subplot(1,3,2);
wdplot = subplot(1,3,3);

for i = 1:length(allsigmas)
    sigma = allsigmas(i);

    imsz = sigma*20;
    imhlf= ceil(imsz/2);

    [xx,yy] = meshgrid(-imhlf:(imhlf+1));

    th = 0*torad;
```

```

xt          = (xx.*cos(th)+yy.*sin(th));
yt          = (yy.*cos(th)-xx.*sin(th));

DoG         = [];

G1          = exp(-xt.^2 / (2*sigma)^2) - 1/2.23;
G2          = exp(-xt.^2 / (2*(2.23*sigma)^2) );
G3          = exp(-yt.^2 / (2*(3*sigma)^2));

DoG         = G1.*G2.*G3;

% Now try find the closest matching vgauss filter...

fops = optimset('MaxFunEvals',500,'TolX',0.001);

startscale = [4 4];
[scale,fval,exitflag] = fminsearch('f_matchFilters',startscale,fops,DoG);

allscales(i,1:2)=scale;

% Vgauss length/width ratio
axes(ratplot);
plot(allscales(:,1),allscales(:,2),'o-');
xlabel('width');ylabel('length');
title('vgauss filter sizes');

% filter width Vs sigma
axes(wdplot);
plot(allsigmas(1:i),allscales(:,1),'o-');
xlabel('sigma');
ylabel('width');
title('DoG sigma Vs Vgauss width');

% filter length Vs sigma
axes(elplot);
plot(allsigmas(1:i),allscales(:,2),'o-');
xlabel('sigma');
ylabel('length');
title('DoG sigma Vs Vgauss length');

drawnow;

end

if length(allsigmas)>1
    disp('sigma, vGauss width, vGauss length');
    disp([allsigmas(:) allscales]);

    % Vgauss length x width ratio
    disp('calculating Vgauss ratio');
    rats=(allscales(:,1)./allscales(:,2))
    meanrat = mean(rats)

    % Vgauss length to DoG sigma ratio
    disp('estimating ratio of DoG sigma to vGauss length');
    r=allsigmas(:)./allscales(:,2)
    meanr=mean(r)

    % Vgauss length to DoG sigma ratio
    disp('estimating ratio of DoG sigma to vGauss width');
    r=allsigmas(:)./allscales(:,1)
    meanr=mean(r)

    return
else

% Get the filter and display it alongside the DoG filter.

    [diff,fGaussFilter,diff2] = f_matchFilters(scale,DoG);

    figure;
    dsp([expimg(DoG,64) expimg(fGaussFilter,64)]);

    figure;
    surf([fGaussFilter DoG diff2]);

```

```

    figure;
    pcolor([expimg(DoG,1),expimg(fGaussFilter,1)])
end
function varargout = f_matchFilters(scale,targetFilter)
%Take the supplied filter and attempt to match a VGauss filter to it.
%Free parameters are the width and length scales. targetFilter orientation
%should be zero degrees.
%
% Round off scales to two decimal places.
fwid = round(scale(1)*100)/100;
flen = round(scale(2)*100)/100;

disp([fwid flen]);

if nargin == 1
    global fwidths flengths targetnoisediffs imgStats;
    fwidths(end+1) = fwid;
    flengths(end+1) = flen;
end

% Determine the size of the target filter
sz = size(targetFilter);
emptyimg = zeros(sz);
emptyimg(round(sz(1)/2),round(sz(2)/2)) = 1;

% Create a filter at the current scale
vGaussFilter = vfilter(emptyimg,[fwid flen],0,[2 0]);

% Make sure that the overall amplitude of each filter is the same
ftarget = expimg(targetFilter,1);
fvGauss = expimg(vGaussFilter,1);

% Make sure that the background level is approximately the same
backdiff = ftarget(1)-fvGauss(1);
ftarget=ftarget-backdiff;

% Now get the summed squared differences...
d = ftarget - fvGauss;
d2 = d.^2;
d2sum = sum(d2(:));

%=====
% Processing complete, now return output values...

varargout{1} = d2sum; % Return the current difference estimate...

if nargin > 1 % Return the generated filter too...
    varargout{2} = vGaussFilter;
end

if nargin > 2 % Return the squared difference between the two filters
    varargout{3} = d2;
end
end

```

Appendix E Comparing co-linearity energy in random and CLE-manipulated images

E.1 Introduction

This section describes briefly a pilot experiment which compared the detection rates for contours embedded within randomly oriented background elements and those which are CLE-manipulated.

E.2 Stimuli

Target stimuli featured contours which had lengths of 3, 4 and 5 elements. For the CLE manipulated stimuli the CLE was varied from 0-1.5 in steps of 0.375. For the CLE random stimuli the orientations of background elements were simply randomised. The locations of background elements were determined using the spaced-fill technique (see appendix A). The spacing of both contour and background elements was 16 pixels. All Gabor elements had symmetrical-phase. The overall size of stimulus images was 320x320 pixels.

E.3 Results

Detection rates are presented within Figure E-1. The detection rates for the contours embedded within randomly oriented backgrounds correspond quite closely to the detection rates for contours embedded in CLE manipulated backgrounds, where the CLE target was 1.1.




Figure E-1. Target image detection rates as a function of CLE and contour length manipulations. The solid lines present detection rates as a function of CLE manipulations. Whilst the isolated symbols (asterisk, diamond and square – corresponding to lengths of 5, 4 and 3 elements respectively) represent the detection rates for contours embedded in backgrounds where elements were randomly oriented.

Appendix F Enhanced detection of smooth contours with the Yen and Finkel association field model

F.1 Introduction

The parameters of the Yen and Finkel version of the association field can easily be modified in order that performance for smooth contours is greater than for jagged contours. This is achieved by increasing the spatial range over which interactions occur and by increasing the range of orientation differences that are tolerated between target and surround elements. The model is outlined in Section 6.2.1 (Functions 6-1 to 6-5). In the following example, the values of two parameters of the association field are manipulated: the orientation range of the co-axial fanout (k_c) and the spatial distance of the fall off in co-facilitation between the target and surround elements (σ_d^c).

F.2 Stimuli

Target images featured one of two levels of arc-length, these were one element or ten elements (Figure J-1). Ten images were generated for each of these arc-lengths. Contours were embedded in a 256^2 (pixel) background of randomly oriented elements. The location of the background elements was determined with the spaced-fill technique (appendix A). The mean spacing of background and contour elements was 16 pixels. Noise images were identical to the target images, except that the orientations of elements within the embedded contours were randomised.




Figure F-1. Stimuli featuring (left) jagged contour, (right) smooth contour.

Association field model

The model was exactly as described in section 6.2.1 (Functions 6-1 to 6-5), except that the values of co-axial fanout (k_c) and spatial extent (σ_d^c) were manipulated. The values of these parameters are shown below in Table J-1.

<i>Parameter</i>	<i>Default (Yen and Finkel)</i>	<i>Modified to favour smoothness</i>
k_c	30	53.7
σ_d^c	40	64.2

Table F-1. Parameters of the Yen and Finkel association field.

The performance of the association-field model was based upon the maximum co-linearity energy (CLE) value for the whole image. That is the maximum CLE value for each image was retained and this value was compared for the paired images in a 2AFC decision. The image with the largest CLE value was assigned ‘target’ status.

F.3 Results

Figure J-2, below, shows how the relative detection rates for smooth and jagged contours can vary where the overall extent of the association field is increased.

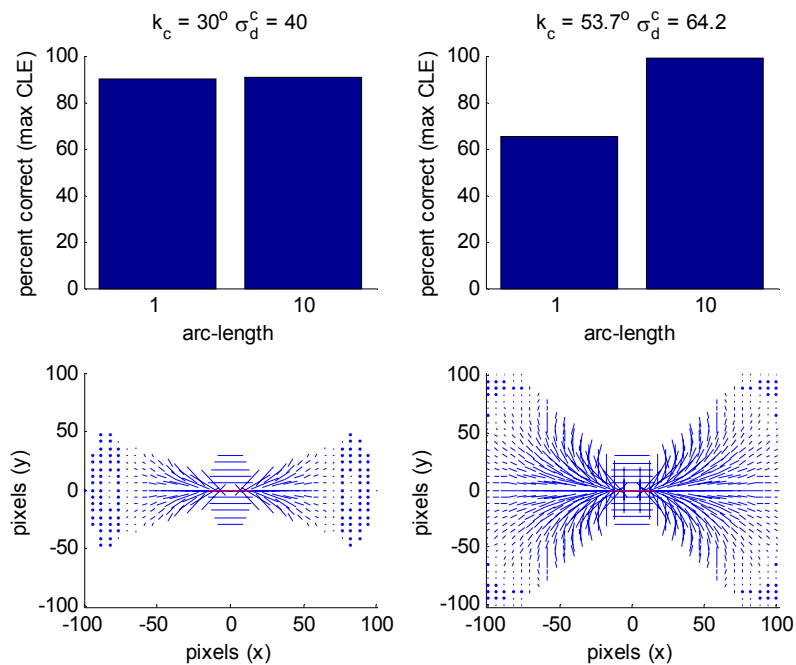


Figure F-2. (lower) integration field structure. (upper) model performance with default parameters (left) and modified parameters (right).