DOCTOR OF PHILOSOPHY

Visual texture integration processes and the role of selective attention

O'Donnell, Helen Louise

Award date:
2000

Awarding institution:
University of Wales, Bangor

UNIVERSITY OF WALES

PRIFYSGOL CYMRU

Visual Texture Integration Processes and the Role of Selective Attention

by

Helen Louise O'Donnell

A Thesis submitted to The School of Psychology, University of Wales, Bangor, in partial fulfillment of the requirements for the Degree of Doctor of Philosophy.

September 29, 2000

Summary

Considering the heterogeneous nature of textures in real-world environments, the ability to perceptually distill heterogeneous texture information into a single feature is an essential aspect of our visual experience. For example, despite wide variation in orientation among blades within a patch of grass, we can still perceptually extract the predominant orientation. The challenge of objectively measuring this ability was addressed by developing the Texture Coherence Paradigm. Observers were presented with an array of line segments in which a percentage of lines (signal) were oriented identically and the remaining lines oriented randomly (noise). The minimum percentage of signal required to correctly determine signal orientation was taken as an observer's coherence threshold. Naïve observers required only 16 to 20% signal for correct orientation discrimination. While it seemed that coherence thresholds varied as a function of signal orientation, results showed that Global Precedence of outer texture patch contours mediated the orientation anisotropies. This suggests that larger scales of analysis, i.e., at figural levels, have a significant effect on perception of inner local features and also that texture integration is isotropic with neutral patch shapes. Observers required a relatively long exposure (200 ms) to reach maximal sensitivity. Taken together, this evidence suggests that a spatially integrative cooperative network of orientation analyzers mediates orientation perception in heterogeneous textures. Since integrative processes are thought to be higher level in nature, i.e., extrastriate, I presented varied textures to examine construction of featural representations. Features were inefficiently combined but focused selective attention assisted significantly. I also induced a texture orientation contrast effect, analogous to the motion aftereffect, in which sensitivity to textures was significantly changed depending on the texture previously viewed. This effect took time to develop and was relatively long lasting suggesting an attentional component. Further, the effect could be "switched off" by redirecting selective attention away from the inducer.

Acknowledgements & Dedication

My mother and father always used to tell me as I was growing up that I could do anything I wanted to if I put my mind to it. This doctorate is the embodiment of putting my mind to an idea for the last four years. I thank them, and my brother Nick, for their love and support through it all and for giving me such a happy and loving home from which to grow and learn.

I also thank my advisor, Professor Jane Raymond, for her guidance and patience as well as for the fun we have had together. I'd like to thank Dr. Jesse Martin, Dr. Heather Jordan, and Mr. Gary Hill for numerous useful discussions that helped to shape the ideas presented in this thesis.

The School of Psychology was very supportive of my Ph.D. program and I would like to thank the department for giving me the opportunity to teach while I pursued my studies in Wales.

Finally, I would like to dedicate this thesis to my fiancé Sean. He has supported me from the penning of the first word to the penning of the last word of this thesis and has spent many an hour in my darkened lab looking at thousands of little tiny lines. He is my life, my soul mate, my best friend, my soon to be husband, and the best psychophysical observer any Ph.D. student could ever ask for. Sean, I look forward to spending all my time with you from now on, thanks for being so patient with me.

Table of Contents

Chapter 1 - Introduction

Abstract

I review here experimental evidence and resulting models regarding perceptual segregation and integration processes in texture perception. Texture perception research has focused mainly on texture segregation processes though integration is also an important function in the efficient construction of an accurate perceptual representation of a visual scene. Related paradigms such as visual search, orientation gradient, contour completion, pattern integration and motion integration are also discussed. The aims of the present thesis are discussed in the last section.

Perception of Complex Textures

Considering the heterogeneous nature of textures in real-world environments, the ability to perceptually distill heterogeneous texture information into a single "feature" is an essential aspect of our visual experience. Picture a patch of grass (Figure 1A). We have a "sense" of approximately how long most of the blades are and what orientation they have without looking at each blade individually. Despite wide variation in orientation among blades, we can still extract a sense of the predominant orientation of the patch. Similarly, a rocky shoreline composed of stones of differing sizes and shapes evokes a global texture quality representative of its surface (Figure 1B). We perceive an overall impression of the predominant size of the stones, for example, or the predominant shape of the stones. In Figure 1C, we can judge the typical shape of a leaf as well as its size. Without this ability to distill complex information into "visual summaries", the world would truly be an overwhelming place. The perceptual process by which we derive this sense of the most common element within locally varying textures has not been measured extensively. By using heterogeneous textures, I investigate here how "noisy" textures are integrated perceptually to form such a sense of the predominant element.

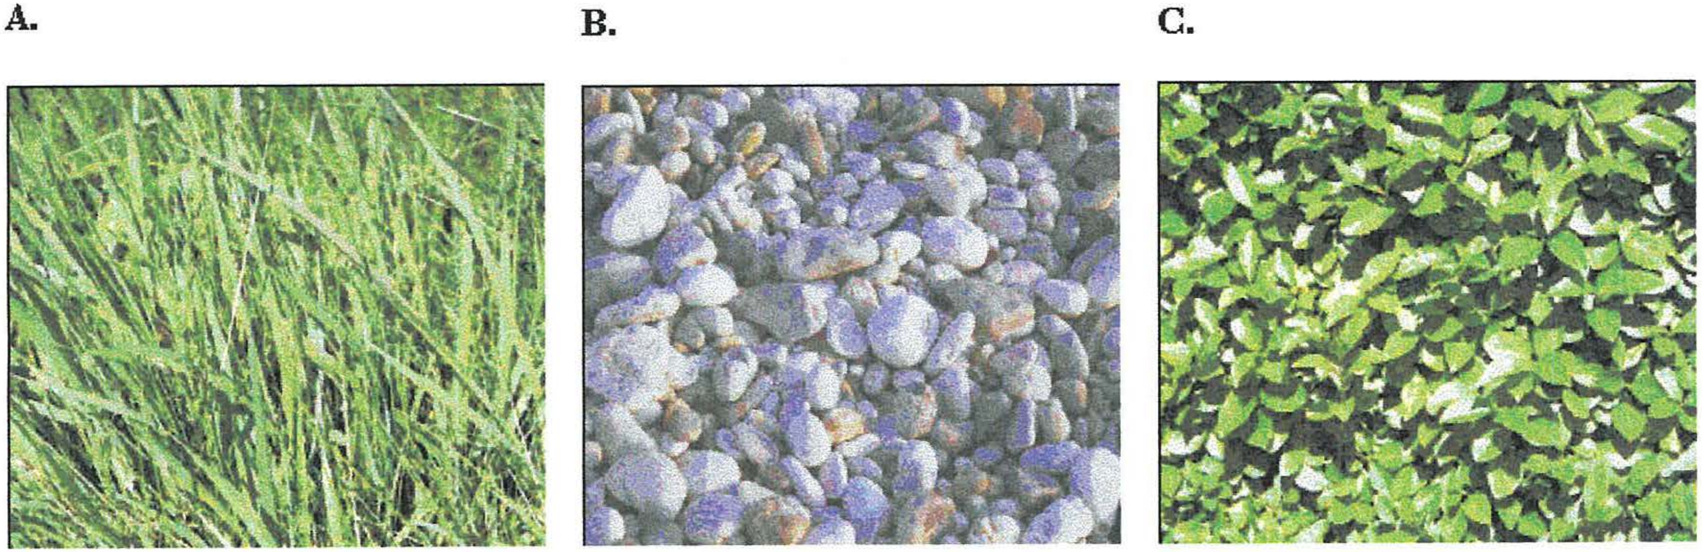

Figure 1. Real world textures vary locally but we can integrate information over space to extract a global percept of that texture's features. For example, we can determine the (a) orientation, (b) size and (c) shape typical of the elements making up a texture despite local variation.

Arguably, our visual experience of the world is defined by the perception of separable objects and coherent surfaces. Seeing objects in complex visual scenes depends on successful segregation of the visual array into objects and boundaries and on the integration of visual information occurring within those spatial, and temporal, boundaries. There are many related experimental paradigms that address segregation and integration-like issues but ask different questions regarding visual processing (Figure 2). Texture segregation studies (illustrated in Figure 2, Panel A) are mainly concerned with how observers perceptually segregate homogeneous texture areas solely on the basis of featural differences within the textures, in this case orientation. By texture segmentation, I refer to the perceptual impression of a border between two differing textures that is not accompanied by a luminance or colour border but rather that is defined merely by differences in the spatial structures or features of the textures in the two regions. Through examining which features result in "pre-attentive" or effortless perceptual segregation between two textures (as opposed to perceptual segregation that requires focused scrutiny), the nature of pre-attentive vision mechanisms has been examined (Julesz, 1981; Beck, 1982; Bergen & Adelson, 1988; Caelli, 1982; Nothdurft, 1985a; 1985b; Nothdurft & Li, 1985). A related task is the detection of a single featurally distinct target within an array of distractors that are different, as is commonly addressed in visual search tasks (Duncan & Humphreys, 1989; Humphreys & Muller, 1993; Sagi, 1990; Treisman & Gelade, 1980; Treisman & Gormican, 1988; Wolfe, 1992; 1994). These two paradigms focus on the *segregation* of featurally discriminable areas.

Experiments investigating structure gradient (Figure 2, Panel B) examine both segregation and integration processes. Specifically, structure gradient experiments examine the amount of orientation contrast required between different texture patches for observers to perceptually segregate those textures (Landy, & Bergen, 1991; Nothdurft, 1991). As can be seen in Figure 2, Panel B, both perceptual segregation between different surfaces and integration within a surface are required to perceive the texture of the central square as different from the surrounding texture. Studies such as these allow the examination of interactions between segregation and integration mechanisms. Contour completion experiments (Figure 2, Panel C) also address the interplay between segregation and integration mechanisms by asking how observers perceptually integrate elements within a contour yet segregate that contour from a heterogeneous background (Field, Hayes & Hess, 1993; McIlhagga, & Mullen, 1996). Observers are typically asked to detect the presence of a continuous but winding contour composed of gabor segments embedded within an array of randomly oriented noise segments. Since both the target contour and the background segments vary in orientation, the task involves integrating orientation within the target contour while at the same time segregating that contour from the background.

Only relatively recently have there been experiments examining how we process heterogeneous texture information through texture integration paradigms (Figure 2, Panel D; O'Donnell & Raymond, 1996). These paradigms examine how observers code typical features of a locally variable texture, such as orientation, to summarize its characteristics (Dakin & Watt, 1997; Kingdom, Keeble & Moulden, 1995; Keeble, Kingdom, Moulden & Morgan, 1995; Keeble, Kingdom & Morgan, 1997; O'Donnell & Raymond, 1996). Spatial integration has typically been studied in other domains such as pattern integration (Glass, 1969; Glass & Perez, 1973; Glass & Switkes, 1976) and motion integration (Newsome & Pare, 1988; Raymond, 1993; Williams & Sekuler, 1984). These paradigms address how local information such as local dot position in Glass patterns (Figure 2, Panel E), or local motion direction in random dot kinematograms (Figure 2, Panel F) is integrated over space to form a global percept.

In the present thesis, stimuli such as that illustrated in Figure 2, Panel D, were used. Shown is a heterogeneous, partially coherent texture patch that has a signal orientation of horizontal. Measuring the perception of orientation in a "noisy" texture such as this is interesting because it allows us to examine how orientation information is perceptually integrated across space to create a coherent experience of texture despite wide variation. Observers are essentially extracting signal information from noise in these stimuli, a very important process for perception but one that is nonetheless not well understood.

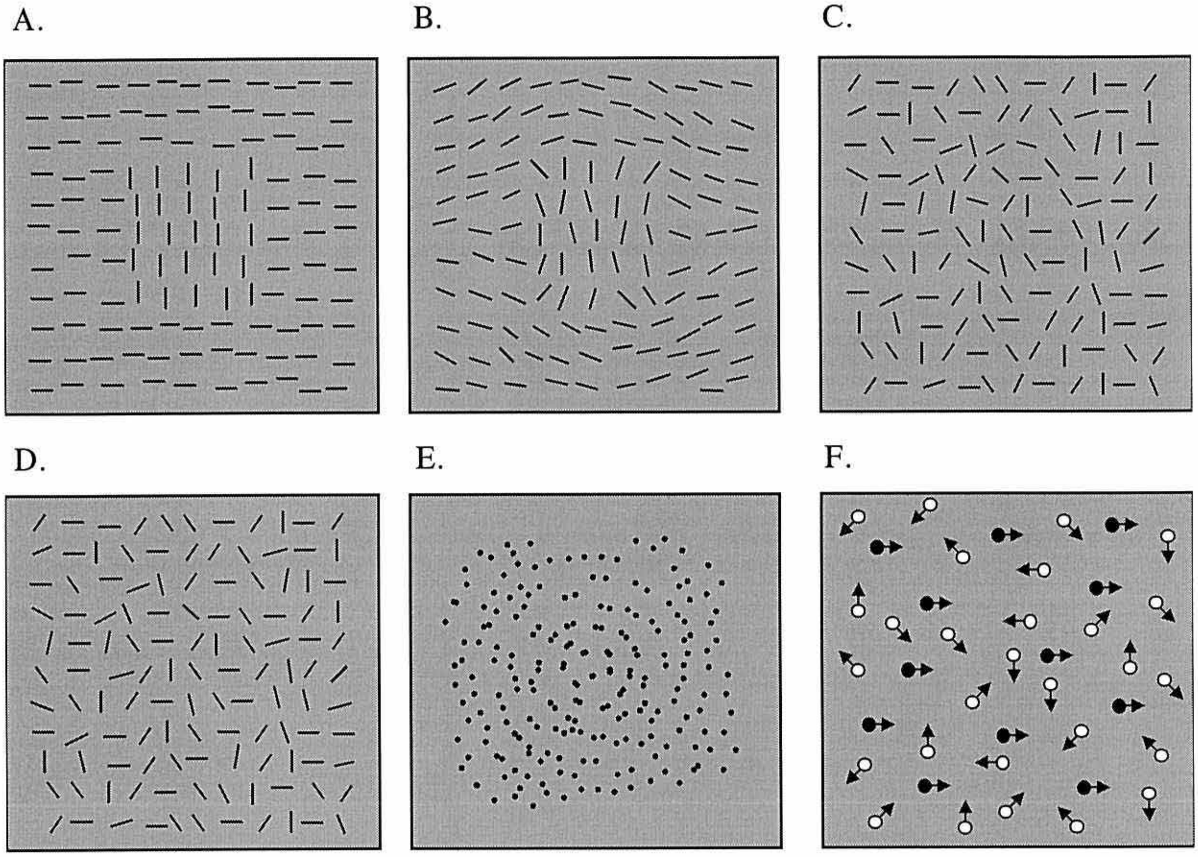

Figure 2. Illustrations of segregation and integration tasks. A. Texture segregation. B. Structure gradient (modified from Nothdurft, 1991). C. Contour completion (modified from Field et al., 1993). D. Texture integration (30% coherent patch, as used in the present thesis and in O'Donnell & Raymond, 1996). E. A concentric Glass pattern (modified from Glass, 1967). F. Motion integration (black dots are signal, white dots are noise; Williams & Sekuler, 1984). See text for further discussion.

Perhaps as a result of the paucity of research specifically examining texture integration mechanisms, at present there is no satisfactory technique for quantifying observers' perceptual sensitivity to complex textures. Of the techniques used, most are non-intuitive and complex arising from more computational or engineering perspectives than perceptual (Kingdom, Keeble & Moulden, 1995; Keeble, Kingdom, Moulden & Morgan, 1995; Keeble, Kingdom & Morgan, 1997). In this thesis, I have developed and tested a simple signal-to-noise technique for quantifying perceptual sensitivity to complex textures. I call this the Texture Coherence Paradigm and it is adapted from analogous techniques used in motion (Nawrot & Sekuler, 1990; Newsome & Pare, 1988) and pattern perception (Wilson, Wilkinson & Asaad, 1997). I use this texture coherence paradigm to examine and quantify how texture integration mechanisms affect our every-day perception of the world and the role of selective attention in such processes.

The remainder of this chapter reviews relevant classic texture perception research as well as recent integration and related paradigms as mentioned above. There are many different terms used in the literature for the structure of textures, however, and some clarification is needed here. In this thesis I use the term texture patch to refer to an area of texture composed of smaller individual objects. The individual objects that make up that texture and define its qualities are called texture elements (these have also been called micro-patterns or texels in the texture literature). For example, in Figure 1, Panel B, the texture patch is the square area of shoreline that is composed of individual stones, or elements. In Figure 2, Panels A-D, each line segment is an element of a larger texture. An element can in turn be made up of components such as multiple line segments in the case of a "**<**" or a "**>**". In this case, the group of components taken together, i.e., the entire "**<**", is referred to as the texture element. Elements have features that define them such as orientation, colour, motion, depth, and shape or form. Texture patches that are made up of elements that are identical or very similar in their features are referred to as homogeneous or uniform, whereas textures made up of elements that are featurally varied (as in Figure 1) are termed heterogeneous or "noisy". Appendix C contains a glossary of terms that are used throughout this thesis.

Texture Segregation, Visual Search and Grouping

Texture segregation and visual search paradigms are closely related and often are used to address similar questions regarding visual mechanisms. Generally, both paradigms are utilized to investigate how observers code features of textures or arrays of elements and represent those features in the brain. Through studying observers' performance in perceptually segregating a different area of texture from a background or detecting a single element that differs in one or more features, both paradigms aim to determine the way features are coded and what errors observers make as a result of this coding. Grouping phenomena are, in a sense, opposite to segregation and visual search processes in that they demonstrate how features of objects cause those objects to perceptually group together rather than segregate. All of these paradigms provide insight into how features of objects are coded within an array of other objects.

- ***Texture Segregation***

The majority of the research to date regarding the perception of texture addresses the subjective segregation of different surfaces (Caelli & Julesz, 1979; Caelli, 1982; Callaghan, Lasaga & Garner, 1986; Enns, 1986; He & Nakayama, 1994; Julesz, 1962; 1975; 1981; Landy & Bergen, 1991; Nothdurft, 1985a, 1985b; Nothdurft & Li, 1985). Thus, many models of texture perception developed from these studies focus on segregation processes. Segregation studies commonly consist of presenting observers with a small, fairly uniform texture patch (e.g., a patch composed of **+**'s) embedded within a similarly textured surround (e.g., composed of **L**'s) and require observers to determine whether the embedded patch was present or absent as is demonstrated in Figure 3 and Figure 4. Results show that under some stimulus conditions (for example, the figure composed of **+**'s on the left panel of Figure 3), observers are able to detect the boundary between even very similar texture areas very quickly and efficiently - seemingly with no effort. This pop-out effect has been termed "pre-attentive" or effortless texture segregation (Julesz, 1975). Under other stimulus conditions (for example, the figure composed of **T**'s on the right panel of Figure 3), longer visual scrutiny is required to make the same discrimination. Numerous researchers have proposed mechanisms underlying effortless and effortful texture segregation.

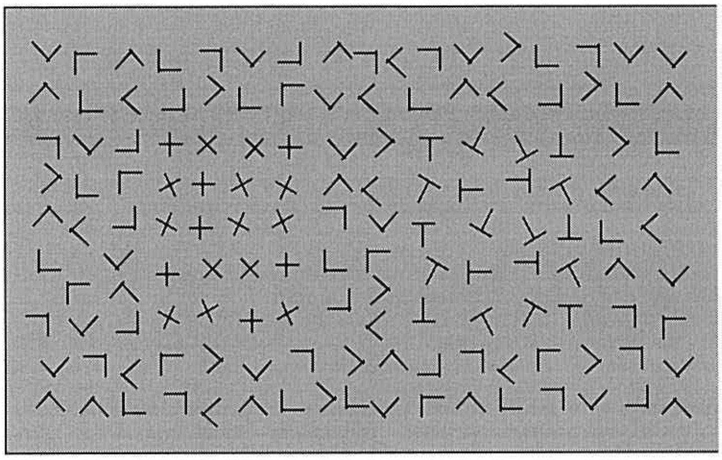

Figure 3. A classic demonstration of "pre-attentive" texture segregation. The area composed of + figures on the left side of the figure segregates easily from the background of L's. However the area on the right side composed of T's does not segregate from the background without scrutiny. All three of the element types are composed of the same line segments and are therefore identical in luminance.

Julesz and colleagues primarily used textures made up of static random dots or random dot micropatterns to address processes of texture segregation. Early work (Julesz, 1962; Julesz, 1975) was concerned mostly with examining the image statistics of such random dot textures. By varying the form of the texture elements, Julesz (1975) independently manipulated the first and higher order statistics of the overall texture. For example, consider a difference in size or luminance between two elements (e.g., ● vs. •). The critical feature of the element, in this case, is size or luminance and is defined at one point in space. Borders between two areas of elements that differ on this continuum are easily perceived as is demonstrated in Figure 4, Panel A. Julesz noted that this border is detectable by examining merely the first-order statistics of the elements. Similarly, in Figure 4, Panel B, the border between elements that differ in orientation (e.g., **|** vs. **/**) is easily perceived. Since the orientation of a line segment element is defined by two points in space, these differences are reflected in the second order statistics of the elements, but not the first. By contrast, the direction of concavity of an element (e.g., **<** vs. **>**) is defined by three points in space. Differences in the direction of concavity among elements does not result in "pre-attentive" segregation and requires close scrutiny for detection, as is demonstrated in Figure 4, Panel C. Only an examination of the third- or higher-order statistics of the elements will result in the detection of the differing area in this case.

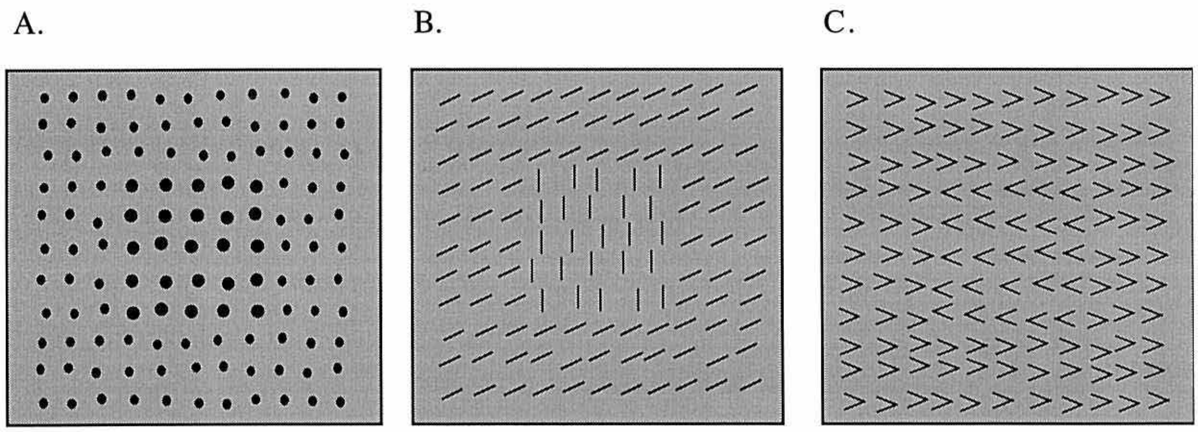

Figure 4. Effortless ("pre-attentive") and effortful texture segregation. A. Texture segregation based on element size or luminance (detectable through differences in first-order statistics) requires little or no focused scrutiny to detect the presence of a different texture patch. B. Segregation on the basis of orientation (detectable through differences in second-order statistics) is also effortless. C. Segregation on the basis of direction of concavity (detectable only through differences in third- or higher-order statistics) requires focused scrutiny and does not segregate pre-attentively.

Julesz and colleagues utilized the finding that some textures pre-attentively segregate from one another and other textures do not, to address the relationship between pre-attentive segregation of texture and the coding of statistical properties of texture images. From the results of such studies, Julesz concluded that the perceptual processes responsible for the pre-attentive segregation of textures are only sensitive to first and second order textural statistics and have no access to higher order statistical information (Julesz, 1975). His conclusion is illustrated in Figure 4, Panel C by demonstrating that a figure composed of **<**'s does not pre-attentively segregate from a background composed of **>**'s.

It should be noted that to create his texture stimuli, Julesz used a method of constraints on the locations of hundreds of dots within a given texture. For example, instead of using line segments as elements, he would constrain two dots to fall a certain distance apart to create a second-order dipole (i.e., • • vs. —). Similarly, with third-order stimuli, three dots were constrained to fall in a **<** type configuration. Interestingly, this approach does not treat the local texture elements as discrete objects with discrete features. Rather, texture elements are treated as a collection of statistics that do not possess separable properties of size or orientation, a distinction that is reflected in Julesz's earlier theories.

Subsequent work emphasized a more feature-processing approach to texture segregation, leading Julesz to suggest that possibly conspicuous local features mediate pre-attentive texture segregation rather than just the image statistics of textures. These local separable features are now considered by Julesz to be the basic elements of pre-attentive perception and have been termed "textons" (Julesz, 1980). Examples of texton classes are an element's colour, orientation, width, and the number of line endings or terminators. For example, since orientation is considered to be a texton class, a patch of vertically oriented line segments placed within a larger patch of obliquely oriented line segments should result in pre-attentive segregation. As demonstrated in Figure 4, Panel B, the inner texture patch does indeed pop-out effortlessly.

Textons have been studied not only with the classic texture segregation task but also with a modified visual search paradigm in which the observer detects the presence of one or more target elements in a display of distractor elements that is presented quickly and is followed by a pattern mask (Julesz & Burt, 1979; Julesz & Bergen, 1983). The measurement of proportion correct responses gives an index of observer performance. Alternately, an observer's reaction time to detect whether a target element is present or absent in the display can be measured (Treisman & Gelade, 1980) in which case, no pattern mask is used and the display is typically left on until the observer has made a response. Both methods provide an index as to how salient a target is in an array of distractors. Presumably, if the target item or texton is pre-attentively discriminable from the distractors, it should be detected accurately and quickly. For example, a target line segment oriented at 90 degrees placed among line segments oriented at 180 degrees should be pre-attentively discriminable and detected quickly and accurately in such a task. However, a line segment oriented at 175 degrees might not be as discriminable when placed among the 180 degree distractors and may take longer to detect.

In Julesz & Burt's (1979) adaptation of a visual search task, numerous target elements were either grouped in a contiguous area or dispersed within the display while the probability of detection was measured. Results showed that when the critical texton belonged to the targets and not the distractors (an additive feature situation), there was no difference in the detection time for dispersed vs. grouped target arrangements. If, however, the critical texton belonged to the distractors and not the target (a subtractive feature situation), the grouped targets were more easily seen than the dispersed targets. This is consistent with the visual search literature in which observers are much faster to detect a target that differs from distractors due to an additive feature than due to the absence of a feature (Treisman & Gelade, 1980). For example, observers are much faster to detect a target ellipse embedded within distractor circles (additive feature: elongation) than they are to detect a target circle among distractor ellipses (subtractive feature) (Treisman & Gormican, 1988).

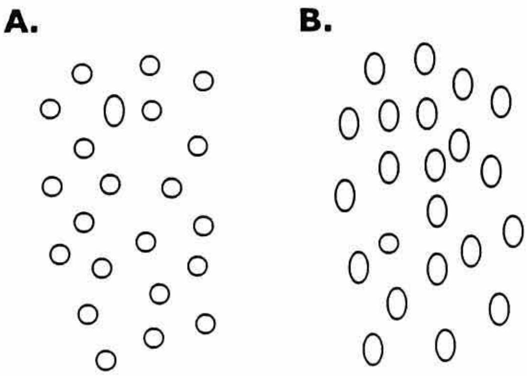

Figure 5. A demonstration of the differences in visual search speed between additive (Panel A) and subtractive (Panel B) feature conditions. Targets are detected more quickly in Panel A compared to Panel B. Adapted from Treisman & Gormican, 1988.

The fact that targets were more easily detected when they were grouped in contiguous areas under subtractive feature conditions led Julesz to propose that a spatial cooperativity process occurs in the texton detector arrays (Julesz & Burt, 1979). A single target with an additive feature discriminating it from the distractors may cause a texton detector array to respond to its presence. If this is the case, in the presence of a target with a subtractive feature, the texton arrays would already be responding to the distractors with the additive feature and thus would require a relatively large area of absent targets for detection to take place. This model of texture processing would predict poor observer performance for any display with a distributed signal, including most heterogeneous or "noisy" textures. In order for any signal to be detected in briefly presented displays, targets would have to be grouped in a sufficiently large contiguous area depending on the resolution of the texton detector arrays. However, as Julesz and Burt's task requires merely the detection of a target, it can still be performed accurately by detecting just one target in isolation and does not require the detection or integration of numerous targets across the field (although this may still occur).

Although this type of model is quite prevalent in the texture perception literature, it has a number of problems. First, there is little mention of the coding of spatial relationships between texture elements which is a very important source of information for many processes such as encoding depth in a visual scene or encoding objects' 3 dimensional form defined by texture (Blake & Marinos, 1990; Witkin, 1981). Second, this model predicts poor performance and a loss of substantial information in a noisy or complex environment where there are many textural cues or features. Such a prediction can be directly tested with the texture coherence paradigm used here. If humans can accurately determine the most predominant signal orientation of a texture patch despite local variations in orientation (noise) and a distributed signal, then it seems Julesz' model requires an additional texture integration stage in which a spatially distributed signal can be effectively used. This stage must also account for the efficient extraction of signal from noise in such displays.

- ***Grouping***

In contrast to the initially low-level segregation approach to texture perception taken by Julesz and colleagues, Beck and his collaborators were interested initially in grouping phenomena and how separate elements group together in a visual scene (Beck, 1966a; 1966b). Inherent in studying grouping phenomena is the assumption that each element is separate at some stage and is processed as a discrete entity. In this view, elements must possess characteristics such as size, shape, colour and orientation that define them. This is in direct opposition to Julesz and colleagues' initial assumptions that textures are defined by their global statistical properties.

Beck proposed that textures are segregated through a series of hierarchical stages of visual processing (Beck, 1982). Initially, the individual elements making up a texture are visually coded by processes similar to those thought to be mediated by complex cortical cells. The features of the textural elements (such as shape, size or orientation) are encoded by these cells and then linked in such a way so as to preserve spatial relationships between elements and spatial texture information across the visual scene as a whole. Element features are subsequently compared and any differences in colour, brightness, orientation or size are encoded. Textural variations over space are detected by overall decision units that compare the magnitude and distribution of stimulation in the feature detectors in neighbouring areas. The resulting output is a spatial map of the differences in textural features that is subsequently used to locate texture boundaries between differing areas of texture.

While the exact properties of the decision units that compare the feature detector outputs in neighbouring areas are not described in detail, Beck may have been alluding to the excitatory and inhibitory interactions among cells sensitive to a particular orientation in primary visual cortex (Gilbert & Weisel, 1989). There is psychophysical evidence supporting this interaction from experiments investigating lateral masking (Polat & Sagi, 1993; 1994) and contour completion (Field et al., 1995). These lateral connections may provide an efficient way to segregate the signal from the noise in heterogeneous textures.

Beck's theory differs from Julesz' texton theory in a number of fundamental ways. First, in Beck's model the assertion that there is a grouping process taking place among the texture elements assumes that each individual element in the display is processed as a distinct entity at the level at which textural characteristics are determined. An interesting issue arising from this theory is that it requires the coding of features to be complete prior to the comparison operations leading to segregation. In light of the hierarchical nature of the theory, this implies that texture segregation is a relatively late visual process compared to Julesz' evidence.

A second difference from texton theory is that the features making up texture elements are pre-attentively coded in Beck's model and not in texton theory. According to Beck's model, since texture elements are processed as discrete objects, all characteristics of the element are coded pre-attentively. Julesz' texton theory holds that only textons are coded pre-attentively and thus only differences in textons can lead to segregation. Examples have been found to show that only certain characteristics of elements (coincidentally those that Julesz classed as textons), cause pre-attentive segregation. It seems then that Beck's model requires some modification in terms of the level of coding texture elements.

A more effective model may come from a theory composed of both Julesz' and Beck's ideas. Possibly, the individual properties of texture elements are coded sufficiently at an early stage to extract the differences between elements but not to the degree in which Beck originally proposed. Features of objects may not have been perceptually combined at this point to create a multi-feature representation, i.e., rectangular, blue, long, oriented horizontally. Rather, features such as orientation and colour may still be separated from one another. Thus, in a heterogeneous texture made up of line segments that differ in orientation, the coding of the orientation of each element would create a spatial map of orientation over space separable from other features, such as colour. This map would subsequently be accessed by the overall decision units that extract the most prevalent orientation. Furthermore, since the elements making up the texture are processed at least to a limited degree, a difference in one isolated element would be detected, rather than requiring a larger contiguous area of differences as Julesz' texton theory proposes. This would lead to relatively good performance in detecting the overall orientation of a noisy texture. This is also consistent with evidence of conjunction errors found in the visual search literature, in which features of objects are re-combined in error. See Chapter 5 for a more in-depth exploration of this idea.

However disparate the starting points of Julesz' and Beck's theories, they seem to have converged on a similar conclusion with respect to texture perception. It seems to be agreed that texture perception is a function that occurs early in the visual pathway and the stimulus parameters required for effortless segregation seem to be defined early in visual processing. Computational models of texture segregation have also provided support for this idea of low level visual processing of texture.

Computational Models of Texture Segregation

The majority of computational models of texture perception describe the texture processing mechanism as a very low level process that occurs at the initial stages of the visual system (Bergen & Adelson, 1988; Caelli, 1985; Gove, Grossberg & Mingolla, 1995; Graham, Beck & Sutter, 1992; Grossberg, Mingolla & Todorovic, 1989; Landy & Bergen, 991; Malik & Perona, 1990; Voorhees & Poggio, 1988). In fact, the inputs to many of these models are the outputs of simulated photoreceptors and subsequent processing reflects the responses of cells in the mammalian striate cortex (Bergen, 1991).

A general simplification of these models is to propose two stages. Initial visual input is received from photoreceptors and passes through a bank of band pass filters each spatially oriented to a different preferred orientation and sensitive to a particular spatial frequency. Numerous filters all sensitive to identical stimuli are receptive to different spatial locations in the visual field. Physiologically, this process may occur in simple cortical cells where receptive fields overlap considerably (Graham et al., 1992). The output from the filters then goes through a rectifying and compressive non-linearity (either full-wave or half-wave) and is subsequently spatially pooled across position. The pooled filter output is receptive to a specific stimulus over the entire visual field and is called a "channel". One channel can be sensitive to vertical orientations over the entire visual field whereas another may be sensitive to rightward oblique orientations. Differences within these channels are pooled and this is the final stage that results in an observer's perceptual decision.

Caelli and colleagues have proposed a complete texture segregation model including all stages from input through to output (Caelli, 1985). An Impletion stage fills in areas of like-texture by causing a spread of activity from highly responding areas to areas of less activity. This essentially causes a suppression of weakly responding areas and is necessary for obtaining uniform regions of texture from stimuli consisting of a few discrete elements with nothing in between. Although a necessary averaging process, this stage can cause undesirable effects of filling in especially for heterogeneous stimuli.

A Correlation stage then follows in which a comparison unit looks for agreement between outputs of different filters at different locations. This stage may reflect the lateral interactions among cortical cells that leads to lateral masking (Polat & Sagi, 1993; 1994) and is implicated in texture integration processes. The final stage is one of Grouping in which pixels are classified as belonging to one region or another. This stage maximizes the correlation within regions and minimizes correlation between regions in order to segregate texture differences.

Within this model there is some confusion as to the scale of comparison used in the stages. For example, it is not clear whether or not the Impletion stage fills in over large distances across the visual field or whether it is very much a local process taking place among a small number of units all with small receptive fields. Similarly, the Correlation stage may compare between a number of local receptors (1 or 2 degrees) or over a larger area (10 or 15 degrees). The possibility that these comparison and filling-in processes may be more wide ranging than is described here is very interesting in the context of integration within a "noisy" texture environment (see Caelli, 1985).

Bergen and Adelson's (1988) computational model of texture segmentation is similar to that of Caelli in that it makes use of a bank of oriented band pass filters to organize the initial input from photoreceptor cells. However, in terms of local feature processing, Bergen and Adelson focus not on the individual elements of texture segregation but rather the more global image representations of the whole texture area.

Specifically, this model is based on the observation that the viewing distance at which the texture stimuli are viewed does not dramatically affect the process of texture segregation. A texture pair that pre-attentively segregates at one viewing distance will most likely segregate at another viewing distance even though the absolute scale of the elements has been changed. Thus texture segregation seems to be insensitive to changes in absolute scale (i.e., changes to texture pairs that are constant for both pairs). However, work by Nothdurft (1985) has shown that changes in *relative* scale (such as changes in spacing without concurrent changes in the size of the elements) have dramatic effects on the pre-attentive segregation of textures. It seems that if the relative relationship between elements is changed, they may no longer segregate pre-attentively.

Bergen and Adelson's model of texture segregation takes element spacing and scale into account by proposing that input is processed though a bank of band pass filters not only sensitive to orientation but also sensitive to spatial frequency. The resulting array of oriented spatial filters thereby takes into account the relative and absolute scales of the texture image.

Outputs from this bank of filters are squared and summed over an area roughly twice that of the filters' receptive field area. An analogy can be drawn between this stage of the model and the complex cells of the striate cortex. Similar to complex cortical cells, the outputs of the model at this stage represent the summation of a number of non-linearly transformed inputs from smaller units that are similar in response but have different positions. Any differences in spatial structure that yield segregating textures would lead to a difference in energy responses in at least one of the filters within this hierarchy. In the interest of neural efficiency, the authors also propose that the model's processes leading to the segregation of differing areas of texture may also mediate other segregation tasks such as those based on colour or brightness.

Models that utilize spatial filtering techniques necessarily need to reduce noise and fill in areas to produce a uniform percept of texture. This process results in effective texture segregation for uniform texture areas but ineffective texture integration in "noisy" heterogeneous environments. Perhaps a specific model more closely based on the physiology of the human visual system will be more promising for processing heterogeneous textures.

Grossberg and colleagues have developed a very complete model of pre-attentive vision that is of particular relevance to the study of texture perception through their examination of textural grouping, perceptual filling in and illusory contour phenomena (Grossberg, Mingolla & Todorovic, 1989; Gove, Grossberg & Mingolla, 1995). This computational model of pre-attentive vision describes the process of perception from receptor input stages to higher-level cooperative and competitive stages of grouping and is very much based on physiological mechanisms that have been found in the mammalian visual cortex.

Input to this model takes the form of retinal ON and OFF ganglion cells with centre-surround antagonistic receptive field structures which in turn project to lateral geniculate nucleus (LGN) ON/OFF cells (Gove et al., 1995). The model's LGN cells receive feedback from higher cortical cells that causes an increase in activation of local areas and an inhibition of activity in surrounding areas. This effectively causes an increase in the ON-OFF contrast of the LGN cell's output to subsequent processing stages. Successive iterations of this cortical feedback loop serve to enhance the neurons that are responding to certain optimal features of an image and suppress those that are nearby.

While the model's use of a cortical feedback loop is similar to mammalian physiology, the brain's cortical feedback to the LGN is not well understood. This feedback loop from cortical cells is thought to play an important role in LGN function as 50% of the total input to the LGN is from the cortex whereas only 20% is from the retina (Robson, 1983). Further, this feedback is strictly topographical causing a correspondence between the locations of the visual fields of both the bottom-up and top-down inputs to a single LGN cell (Dubin & Cleland, 1977; Updyke, 1975; Weber, Kalil & Behan, 1989). The feedback has been found to be excitatory (e.g. Kalil & Chase, 1970) or inhibitory (e.g. Hull, 1968) or even both within the same experiment depending on arousal state and brain stem activity (Funke & Eysel, 1992). Grossberg's model helps to explain these findings to some extent through the model's proposed iterative process for 'sharpening' the visual scene through the excitation of optimal areas and inhibition of sub-optimal or neighbouring areas.

In this model, projections from the LGN synapse with cortical simple cells of primary visual cortex, which make up part of the model's pre-attentive Boundary Contour System (BCS). The BCS is concerned with detecting relatively invariant object boundaries amidst noise and utilizes the interactions between simple, complex and hypercomplex cells in V1 to spatially sharpen neural responses to oriented luminance edges in order to find those borders. The output from hypercomplex cells is integrated through bipole cells that fire provided their two receptive fields are stimulated by appropriately oriented hypercomplex cell input. For example, a bipole cell sensitive to horizontal orientations will be activated by two horizontal line segments falling within its receptive fields but will be inhibited by a vertical line segment within one or both of its receptive fields. In this way, the bipole cells mediate a long-range cooperative process that detects like-oriented line segments separated in space.

In order to extract border information across noisy surfaces, the most prevalent boundary information must be separated from surface noise. The BCS uses a feedback loop between the bipole cells and the hypercomplex cells to accomplish this. The bipole cells initially determine which orientation is receiving the greatest activation from the hypercomplex cells over all the noise in the system and act to further excite them. Through a spatial competition in which strongly responding bipole cells inhibit neighbouring more weakly responding bipole cells, the best location for a boundary is determined. From this competitive and cooperative cycle of feedback, appropriate feature combinations will be grouped and a boundary between different features will emerge from the noise in the image. A subsequent filling in process operates within the boundaries to create a more or less uniform percept within.

Two other subsystems within this model are the Feature Contour System (FCS) and the Object Recognition System (ORS). The FCS is primarily concerned with the surface colour and luminance of an object once its boundaries have been determined.

The Object Recognition System (ORS) mediates higher-level learned object recognition processes and can influence both the BCS and the FCS in an attentive top-down manner. Its effects stem from expectations, experience, recognition, and priming, and may play an important role even in what was once thought to be "low-level" vision.

Grossberg's model provides a very compelling explanation of the perception of texture that is consistent with a global integration process mediating the perception of heterogeneous textures. Through the proposed cooperative and competitive feedback loops in the boundary contour system, a representation of the most prevalent orientation will emerge from the noise in the display allowing observers to perform efficiently on a texture integration task with heterogeneous textures.

Combined Segregation and Integration Processes

Additional evidence for a cooperative texture integration stage composed of excitatory and inhibitory interactions comes from lateral masking and contour integration experiments. Lateral masking is typically demonstrated with two flanking gabor line segments placed on either side of a gabor test segment as illustrated in Figure 6 (Polat & Sagi, 1993, 1994). The luminance contrast of the centre line segment is varied and the minimum contrast required to just detect its presence is taken as the contrast detection threshold. When the two flanking line segments are positioned in close proximity, the contrast detection for the centre segment is higher than if they were not present. This is a lateral masking effect. However, if the flanking segments are positioned slightly farther away from the centre segment, the contrast detection threshold shows a marked decrease, or facilitation due to the flankers' presence (Polat & Sagi, 1993). Perhaps of more interest in terms of texture integration is the finding that if the flanking line segments are oriented end to end with the test segment (Figure 6, Panel A), the facilitation effect is larger than for other orientation relationships (Polat & Sagi, 1994). Facilitation still occurs for line segments placed orthogonally to their global orientation (i.e. "stacked") but to a lesser degree (Figure 6, Panel B).

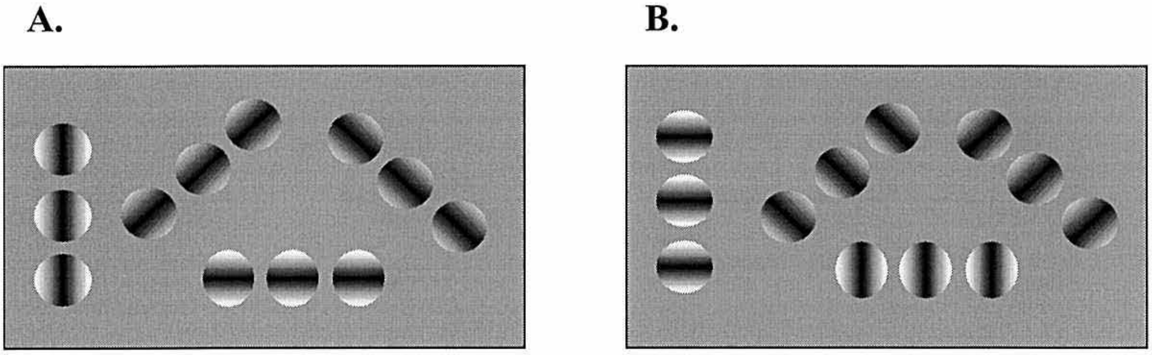

Figure 6. An illustration of the stimuli used in the lateral masking paradigm in which contrast detection threshold for the centre gabor element is reduced by the presence of flanking gabors at moderate proximities. Sensitization effects are orientation-specific resulting in maximal sensitivity for end-to-end placement of gabors (Panel A) and a smaller enhancement of sensitivity for orthogonal orientations of gabors (Panel B). (From Polat & Sagi, 1994)

The finding that the presence of lateral maskers can manipulate the contrast detection threshold of a line segment, and that this effect is orientation-specific, is strong evidence for lateral interactions between orientation selective analyzers in the visual system. Further evidence for these interactions come from contour completion studies.

Contour completion experiments (e.g. Field, Hayes & Hess, 1993; McIlhagga & Mullen, 1996) commonly consist of presenting observers with an array of randomly oriented gabor line segments within which is embedded a trail or contour of line segments (targets), oriented end-to-end, that show good continuation through the array (Figure 7). Observers are asked to detect the presence of such a contour after a brief presentation of the stimulus. Results show that observers can detect a contour quite efficiently provided it has good continuation within the array (Figure 7, Panel A). Perturbations in the orientations of the target lines relative to one another, for example the orthogonal orientation of line segments relative to the global orientation of the contour (Figure 7, Panel B), reduced observers' ability to detect the presence of the path. A local "association field" was proposed to explain the results in which line segments that can be connected with simple curves (those with no sharp inflections) are "associated" and orientation information is integrated among these segments. While this task is essentially a texture segregation task, it does test the effects of a more heterogeneous background than classical texture segregation studies. Furthermore, the association field that is proposed can explain observers' performance in more ecologically plausible situations than the classic texture segregation models.

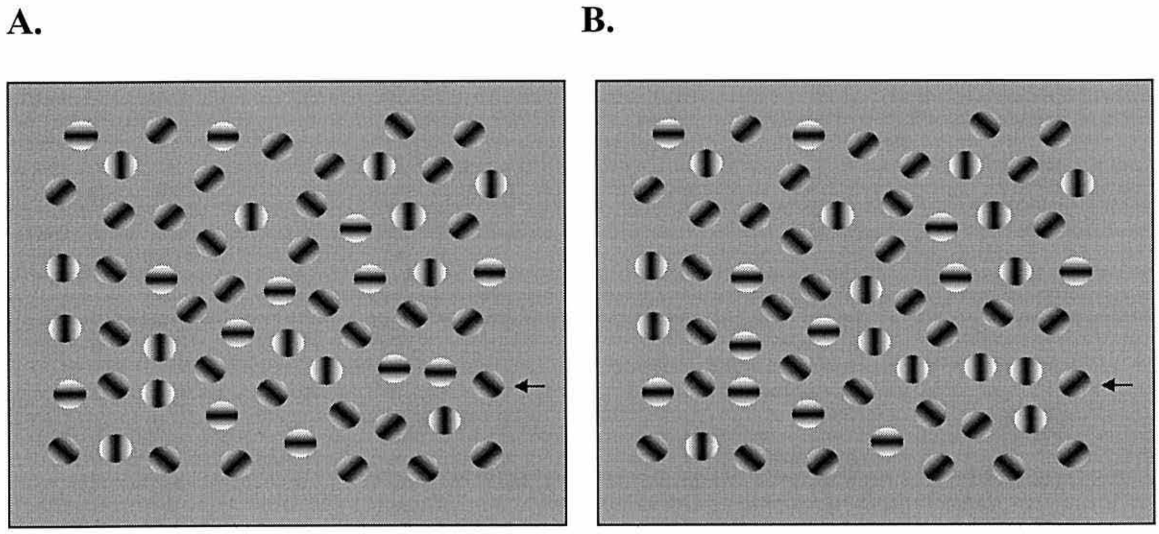

Figure 7. An illustration of the contour integration paradigm in which observers detect the presence of a contour of gabor segments oriented end-to-end that follow a line of good continuation within a patch of random orientation noise (Panel A). The same contour made up of segments with local orientations orthogonal to the global orientation of the contour is more difficult to detect (Panel B). (Adapted from Field, Hayes & Hess, 1993).

In summary, texture perception has historically been examined through tasks requiring the visual segregation of two or more heterogeneous textures. Resulting models emphasize texture perception as a process specialized for such homogeneous images. However, in the real world the majority of scenes are complex, vary locally, and yet apparently are processed efficiently. The evidence to show the presence of such efficient global integration processes that effectively separate signal from noise and produce an overall percept descriptive of a textured surface is described next.

Perceptual Integration Processes

The ability to efficiently integrate complex visual information from the entire visual scene and distill it into an accurate representation of that scene is a remarkable ability. An overwhelming amount of visual information must be processed efficiently for survival that is defined by diverse visual domains, for example texture, form and motion.

- ***Texture Integration***

Only a few recent studies have examined texture integration processes directly (Dakin & Watt, 1997; Keeble, Kingdom, Moulden & Morgan, 1995; Keeble, Kingdom & Morgan, 1997; Kingdom, Keeble, & Moulden, 1995). These paradigms have an advantage over segregation paradigms in that they examine perceptual sensitivity to heterogeneous textures using more realistic natural images with visual noise.

Keeble and colleagues (Keeble, Kingdom, Moulden & Morgan, 1995; Keeble, Kingdom & Morgan, 1997) apply Fourier methods to the orientation domain of texture perception to examine how observers can perceive global structure in texture using oriented line segment stimuli. By varying the probability density function (PDF) of orientation in texture displays, they were able to strictly control the amount of orientation information available to observers. The orientation PDF is the probability that a line segment of a given orientation will occur in the texture image. This function can be varied in form, i.e., sinusoidal, notch or comb functions, and also in amplitude and spatial frequency. For example, Figure 8 illustrates sinusoidal PDFs and the corresponding texture stimuli (modified from Keeble et al., 1995). The task for observers is to discriminate orientationally random textures from those that have been systematically modulated. In this example, the number of line segments of a given orientation is defined by a sinusoidal function through 180° (the cycle repeats for orientations greater than 180°). The resulting texture image has a given number of predominant orientations depending on the frequency of the function. For example, a sinusoidal function with a frequency of 1 cycle/180° with a peak at 45° has most line segments oriented obliquely with the number of those oriented in the remaining orientations decreasing in a sinusoidal fashion (Figure 8, Panel A). The amplitude of the sinusoidal modulation is varied to measure the minimum systematic modulation required for discrimination from an orientationally random texture. All line segments are randomly located in the texture arrays requiring participants to determine the overall orientation by perceptually integrating all orientations over space. Thus in Figure 8, Panel A, the predominant orientation is oblique in the presence of orientation noise.

Texture images resulting from sinusoids with higher frequencies, for example, a frequency of 2 cycles/180° have two predominant orientations (Figure 8, Panel B). As the number of peaks in the sinusoid per cycle (frequency) increases, the number of predominant orientations increases. A randomly oriented stimulus has no orientation modulation, i.e., an orientation modulation amplitude of zero is a flat function (Figure 8, Panel C). In these experiments, the amplitude of the orientation modulation required for observers to correctly discriminate an array of signal (Figure 8, Panel A or B) from that of noise (Figure 8, Panel C) was measured. Results indicate that participants were most sensitive in discriminating textures made up of one predominant orientation (1 cycle/180°) from textures made up of orientation noise. Thresholds were elevated for textures composed of greater than 2 cycles/180° and for textures composed of less than 1 complete cycle. Their results indicate that texture perception is most sensitive for a single global orientation and displays a band-pass pattern of sensitivity around this point.

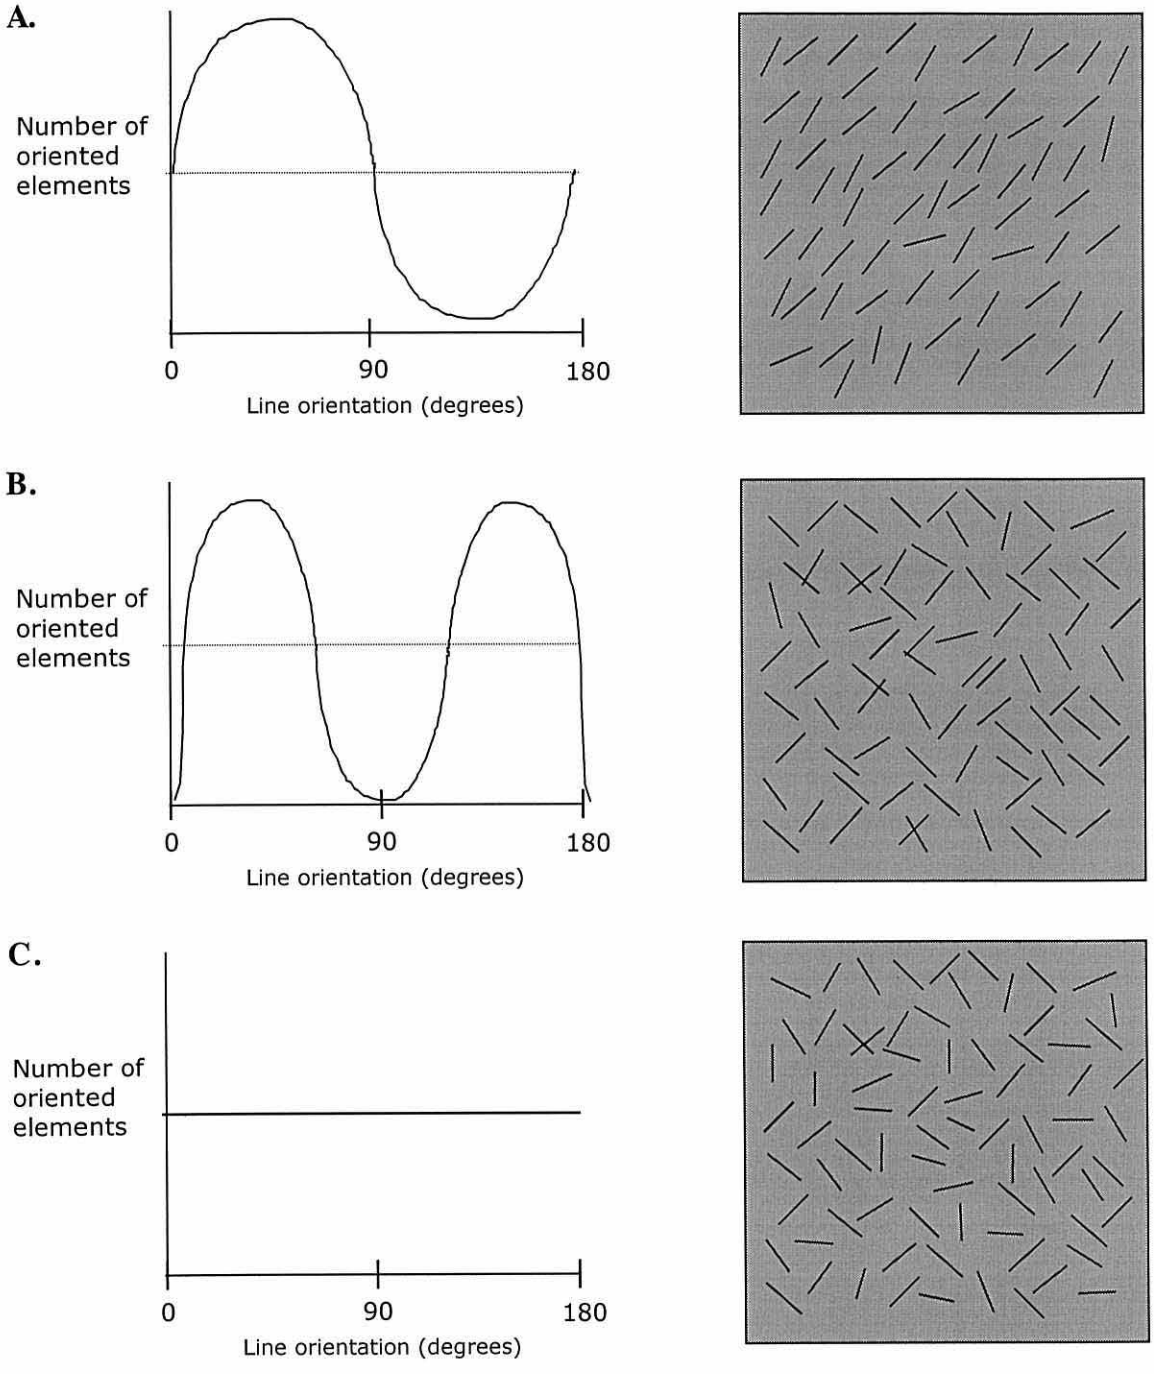

Figure 8. Sinusoidal functions defining orientation modulation and an illustration of the corresponding texture arrays. Panels A & B. Shown are frequencies of 1 and 2 cycles/180° respectively. Panel C. No systematic orientation modulation, i.e., noise (Modified from Keeble et al., 1995).

Based on these results, the authors propose that the visual coding of orientation information in these displays can be modeled using numerous wide-band

(34° half-height, full-width) linear orientation-selective filters. Their linear filter model accurately predicted participants' performance on a number of similar orientation discrimination tasks with different orientation PDFs (i.e., notch and comb functions). Such results indicate that the initial coding of orientation information in heterogeneous textures can be accounted for by simple linear orientation filters even in the presence of orientation noise (Keeble et al., 1995).

In a separate set of experiments that used the same techniques, Keeble and colleagues (Keeble et al., 1997) investigated observer sensitivity for discriminating unimodal and bimodal textures defined by sinusoidal and comb functions. They conclude that the peaks in the bimodal texture PDFs must be separated by at least 13 degrees of orientation in order to be discriminable from unimodal textures with the same variance. They suggest that this task taps a spatially integrative function as the orientation resolution is substantially reduced compared to the better than 1 degree orientation resolution found at the fovea for single line discriminations (Makela, Whitaker & Rovamo, 1993). Further, a mechanism based on only a few central line segments would also be expected to have a lower threshold than 13 degrees indicating that there is integration of orientation occurring which reduces resolution. The authors propose that this integration function occurs at a later stage than at initial coding and could be mediated by lateral interactions between like-oriented analyzers or alternately through grouping processes. Their main concern however is not the texture integration process and they do not expand on this point.

Using similar methods, Kingdom and colleagues more directly examine the spatial integration stage that is required to bind outputs of local filters (Kingdom, Keeble, & Moulden, 1995). Kingdom and colleagues use PDFs to examine orientation modulation across space and to explore spatial integration processes. In these experiments, participants were briefly presented with arrays of line segments varied sinusoidally in orientation across space. Illustrations of such stimuli are shown in Figure 9. Participants were asked to discriminate the modulated line segment array from an array of random orientation noise with a given bandwidth. Since the orientation of a single line element within the modulated array could vary within a given bandwidth, the orientation of one element was not necessarily representative of the global orientation of the array at that point in the sinusoidal function. Thus, the observer had to perceptually integrate numerous line elements over space to determine whether the array was systematically varied or random in orientation. The amplitude of orientation modulation required for correct discrimination was measured.

A.

B.

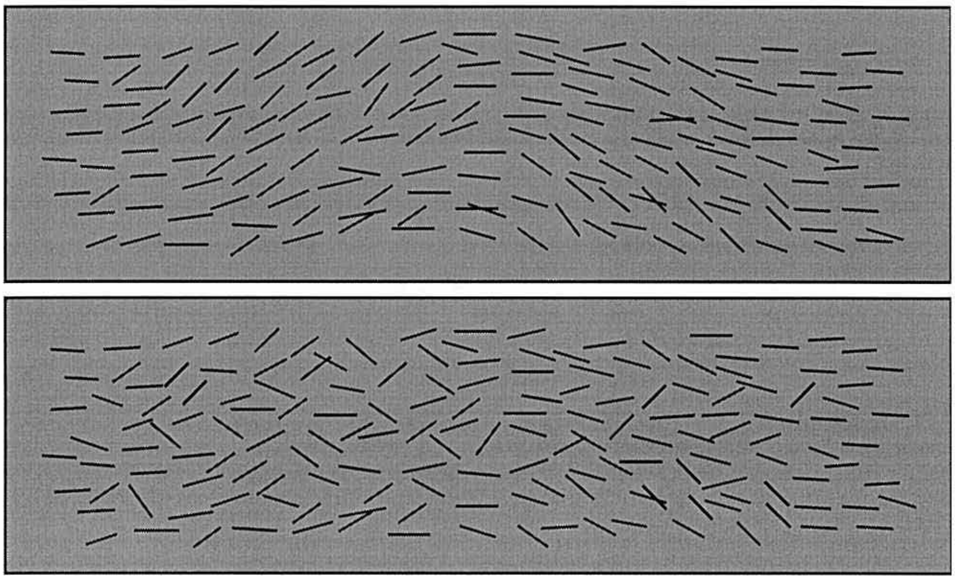

Figure 9 A. An illustration of the sinusoidally modulated line segment array and B. the random noise array used in Kingdom et al., 1995. Orientation modulation is exaggerated in this figure for illustrative purposes.

Results showed that participants required very little modulation to discriminate between the two arrays particularly for low spatial frequencies of modulation. This suggests that there is an integration process operating over extended regions of space that is sensitive to even small systematic modulations in orientation embedded within noise. With the addition of larger amounts of orientation noise (a wider bandwidth within which the orientation of each line segment could vary), observer performance rapidly declined in a non-linear fashion. This non-linear worsening of performance suggests the process mediating global orientation integration is a non-linear statistical process that is not merely a result of linear spatial filters creating a veridical map of all the line orientations in the display (Kingdom et al., 1995).

The authors propose that these non-linear results are due to a cooperative process operating over space through inhibitory and excitatory interactions among neighbouring orientation selective neurons (Kingdom et al., 1995). This cooperativity of like-oriented detectors among neighbouring areas would result in an integration system sensitive to gradual or continuous modulations of orientation (consistent with the results showing best performance for lower spatial frequencies) but one that breaks down in the presence of orientation noise. A system such as this would also show an advantage for patterns showing good continuation or those containing smooth curves of oriented elements which is consistent with the local "association field" proposed by Field, Hayes and Hess (1993). Kingdom and colleagues propose a tentative three-stage process that could mediate such integration processes. Initially, individual element detectors are activated by the texture image, the outputs from which pass through a cooperative stage of inhibitory and excitatory interactions that then passes on a weighted map of local orientations to a final stage. Detection of spatially distributed orientation modulation occurs at this final stage. Such a model is consistent with evidence from Keeble and colleagues (Keeble et al., 1995; Keeble et al., 1997) and closely resembles that proposed by Grossberg in his computational model of texture perception (Grossberg, 1989; Gove et al., 1995). Both models remain to be extensively tested.

While Keeble, Kingdom and colleagues examined observers' abilities in discriminating random and systematically modulated textures; Dakin and Watt (1997) explore specifically, the orientation statistics computed by the human visual system from texture images and how the central tendency of orientation is represented in heterogeneous textures. Using orientation PDFs, Dakin and Watt varied the orientation mean, variance, skew and bimodality of the image function. Using a variety of judgment tasks, Dakin and Watt found that in order to discriminate a bimodal texture from a unimodal texture, observers required at least 10° of orientation difference between two peaks of the PDF of the bimodal texture. This is consistent with the 13° threshold that Keeble et al. (1997) found for a similar task.

Dakin and Watt also found that while variance information is available to texture integration mechanisms, skew information is not. They added skew to the orientation PDF function and modeled human judgments of orientation using various measures of central tendency: zero crossings, threshold edge, peak and centroid. The skew added to the function caused different predictions from each measure and results showed that human orientation judgments were most closely predicted by the centroid of the orientation distribution which is essentially a mean orientation calculated to take into account the cyclical nature of the orientation distribution (see Dakin (1997) for centroid equations). Interestingly, when the centroid model is applied to a bimodal texture in which one set has only a few elements, it completely smoothes over the smaller distribution. In other words, this model is "blind" to small distributions of uniquely oriented segments. Obviously our visual systems are not "blind" to these embedded figures as is shown by "effortless" texture segregation (Caelli & Julesz, 1979; Caelli, 1982; Callaghan et al., 1986; Enns, 1986; He & Nakayama, 1994; Julesz, 1981; Landy & Bergen, 1991; Nothdurft, 1985) and the "pop-out" of unique targets in visual search (Julesz & Bergen, 1983; Sagi, 1990; Treisman & Gelade, 1980; Treisman & Gormican, 1988). These findings suggest that a more complex representation than just the centroid is available.

Taken together, these studies propose that the initial coding of orientation information can be accounted for by numerous linear orientation selective filters with wide-band sensitivity (Keeble et al., 1995). This local orientation information is then integrated on a larger spatial scale. This integration is thought to be a non-linear cooperative process based on lateral interactions between orientation sensitive analyzers and is sensitive to smooth curves and collinear elements (Field et al., 1993; Kingdom et al., 1995; Polat & Sagi, 1993; 1994) but is adversely affected by the addition of noise to the image (Kingdom et al., 1995). Such a cooperative stage allows the determination of the global orientation of a texture image amidst noise.

Specifically, the centroid, variance and bimodality of texture images are available to integration mechanisms to create a representation of a texture. Orientation resolution is lost somewhat as a result, and peaks in orientation functions must be separated by at least 10 to 13 degrees to be discriminated (Dakin & Watt, 1997; Keeble et al., 1997).

The methods that have been used to investigate texture integration processes to date (i.e., PDF functions and Fourier transforms) are complex and not intuitive, though may be useful for such applications as computer vision. A simple technique is needed to easily quantify perceptual sensitivity to complex textures that can be used for addressing such diverse questions as those asked above. My goal here was to develop such a technique that would be easily implemented but robust enough to give experimenters control over varied stimulus factors and experimental designs. Simple signal-to-noise ratio techniques have been used to address integration processes in pattern perception using Glass patterns (Glass, 1969) and in motion perception using random dot kinematograms (Williams & Sekuler, 1984; Newsome & Pare, 1988) and have proven to be valuable. I explore these paradigms next.

- ***Motion Integration***

The motion coherence paradigm was originally developed to address the interactions of multiple vectors of motion presented within one moving stimulus (Williams & Sekuler, 1984). An array of random dots was presented to an observer, each dot either assigned to the "signal" condition or the "noise" condition. Signal dots were systematically displaced in the same direction and noise dots were randomly repositioned within the array. Assignments to signal and noise conditions were randomly changed between motion displacements. The strength of the motion signal was varied by changing the relative percentages of signal and noise dots present. A motion stimulus composed solely of noise dots looked like a "snowstorm" of random motion (Figure 10, Panel C) whereas a partially coherent stimulus looked similar but a predominant direction of motion could be extracted from the noise with surprisingly little coherent motion (Figure 10, Panels A & B). Observers were required to respond whether there was a coherent flow of motion within the stimulus or whether it was random noise.

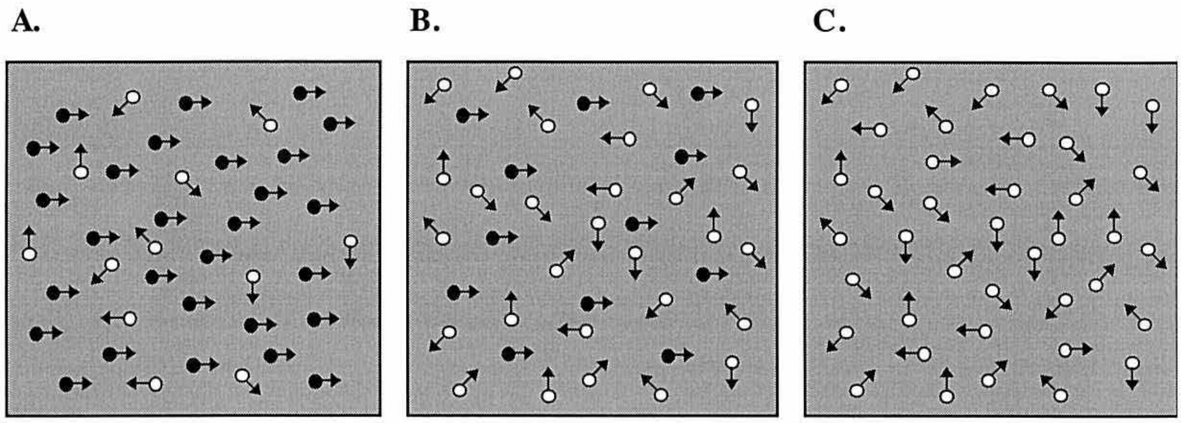

Figure 10. An illustration of the motion coherence paradigm. Shown in panels A and B are two random dot kinematograms that have a rightward signal direction and coherence values of 70% and 30% respectively. Panel C shows a 0% coherence noise stimulus that would resemble a "snowstorm" of motion. Black dots illustrate signals and white dots illustrate noise. In the actual stimuli, dots are all the same colour.

A major advantage of this paradigm is that the local motion vector of an individual dot does not provide reliable information about the global motion since the assignments to signal and noise conditions are randomly changed between displacements. The resulting stimulus requires spatial and temporal integration of motion information in order for correct responses to be made. Another advantage is that a degraded motion stimulus, i.e., low coherence, has the same physical luminance, velocity and spatial frequency of a higher coherence stimulus. Noise can be added to a stimulus without changing other features of that stimulus. Observers are able to judge motion coherence efficiently (i.e., thresholds of 5 to 15% coherence depending on stimulus factors for central foveal presentations; Raymond, 1994) and various authors propose that local motion information is independently detected and responses are then pooled over time and space to create a global percept of directional motion. This coherence paradigm has also been used to study motion perception in special populations such as individuals with glaucoma (Joffe, Raymond & Crichton, 1997), cerebral lesions (Barton, Sharpe & Raymond, 1995; 1996) and dyslexia (Raymond & Sorensen, 1998).

- ***Pattern Integration***

An analogous paradigm has been used to measure the perception of global structure in arrays of systematically displaced static dots in displays called Glass patterns (Glass, 1969;Glass & Perez, 1973; Glass & Switkes; 1976). Glass patterns are composed of two identical arrays of randomly placed dots in which one array is systematically displaced from the other in a radial, linear or concentric manner. The resulting percept of global structure stems solely from the positional relationships among dots and can only be constructed perceptually by integrating dot positions over space.

Wilson, Wilkinson and Asaad (1997) investigated the strength of this integration process and were able to quantify it by adding percentages of positional noise to the dot pairs. A percentage of noise dots was randomly repositioned amidst the systematically positioned signal dots and the maximum amount of noise tolerable in the perception of the Glass pattern was measured. Figure 11 illustrates this paradigm. In Panel A, a typical concentric Glass Pattern is shown in which the dot pairs have been rotated with respect to each other. Panel B shows the same pattern with the addition of positional noise. The perception of the Glass Pattern is degraded but still resolvable. Panel C shows the effects of the addition of more noise in which the concentric structure of the pattern is further degraded.

A. B. C.

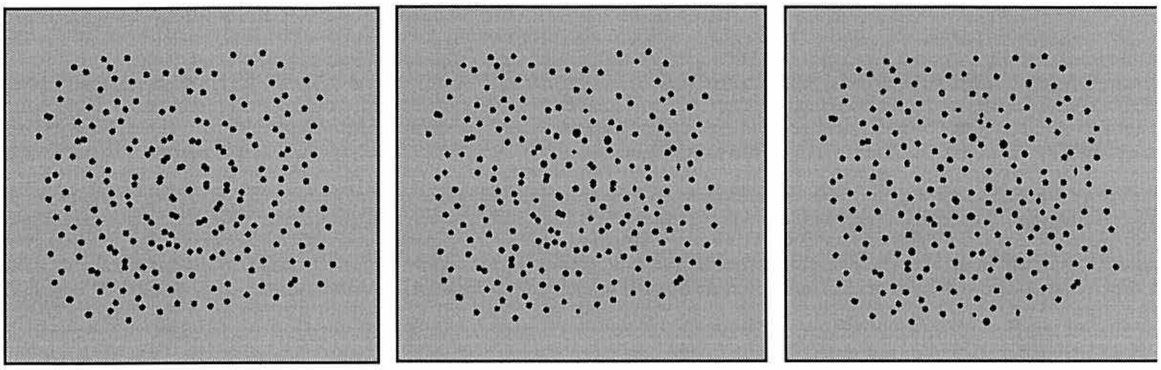

Figure 11. Concentric Glass Patterns with varying degrees of added noise. Panel A. No noise added. This is a typical concentric Glass Pattern. Panels B & C. Approximately 25% and 50% noise added to dot positions respectively.

Observers were asked to discriminate degraded Glass patterns from random dot patterns in a temporal two alternative forced choice (2AFC). Results showed that observers required only 11.6% systematically repositioned dots to correctly select the Glass pattern if that pattern was concentric in structure. Observers had higher thresholds for other Glass pattern structures: radial, 24.1%; hyperbolic, 28.7%; parallel, 56.5%. The authors propose that a simple oriented filtering model, consistent with the physiology of the form vision pathway in primates (specifically V4), can account for the results. They go on to conjecture that units in V4 may form an important link between local coding of orientation in V1 and global face perception in IT.

Indeed, human neuropsychological patients have shown deficits in such tasks involving global integration of local information. Rentschler and colleagues (Rentschler, Treutwein & Landis, 1994) present a patient (KD) with an infero-medial occipito-temporal lesion to the right side who displays difficulties in global visual tasks, for example, difficulties in recognizing handwriting and familiar faces (prosopagnosia). The authors presented her with a discrimination task using Glass patterns (Glass, 1969) in which positional noise had been added. KD showed difficulties in discriminating the global form of the patterns but performed almost perfectly on local micropattern tasks. The dissociation of KD's performance on local and global tasks taken together with the site of her lesion suggests that global integration processes in form vision take place in higher extrastriate areas of the ventral processing stream.

Implications for Measuring Texture Integration

The motion coherence and Glass pattern paradigms are quite analogous to one another in the manner in which they are generated and have two main advantages. First, in these paradigms, random noise is added to the attribute that defines the global percept itself, i.e., random motion noise in motion defined stimuli and random position noise in Glass patterns that are defined by dot positions. A second advantage is that the percept of global structure, i.e., typical motion direction or concentric structure, is not defined locally in these patterns. Rather, a spatial integration process must operate to make global structure visible. Thus, effective discrimination must stem from integration mechanisms.

To examine texture integration processes in a straightforward manner, I have developed the Texture Coherence Paradigm for quantifying texture integration thresholds that is based on motion coherence and Glass Pattern thresholding techniques. Chapter 2 introduces the Texture Coherence Paradigm, demonstrates the perceptual sensitivity of naive observers to heterogeneous textures and addresses the role of embedded contours within textures.

Subsequent chapters utilize the texture coherence paradigm to address how texture integration processes affect our normal every-day perception of textures. Chapter 3 investigates whether there are differences in perceptual sensitivity as a function of texture orientation as may be predicted from results showing oblique effects in contrast sensitivity for grating patterns. Chapter 4 addresses the minimum stimulus duration required for texture integration completion and the role of pattern masking. Chapter 5 investigates whether local element polarity is accessible to texture integration processes and discusses the role of grouping on the basis of local features.

Chapters 6 and 7 investigate the question of temporal interactions of texture integration. Does integration persist in time? Chapter 6 demonstrates the presence of a texture contrast effect that varies in strength as a function of the stimulus onset asynchrony (SOA) between successively presented textures. Chapter 7 explores the role of selective attention in the texture contrast effect. Is texture contrast an "automatic" process or does it show attentional modulation?

The general discussion in Chapter 8 integrates the findings of this thesis into a coherent whole and suggests avenues for further research into texture integration processes.

Chapter 2 - The Texture Coherence Paradigm

Abstract

Visual textures in the natural world are largely "noisy", i.e., the elements comprising them are featurally variable. How are features of the most predominant elements in such textures perceived? I briefly presented line segment textures in which most elements were randomly oriented but a small percentage (signals) had the same orientation. Observers were asked to identify the most common element (i.e., the signal) orientation. The proportion of correct responses increased as percentage of signal elements (coherence) increased. In another experiment, a large number of observers were asked to determine whether the texture was "coherent" or "random" i.e., a signal present vs. signal absent discrimination. Again, proportion coherent responses varied as a function of stimulus coherence. For both tasks, observers were more sensitive to horizontally oriented textures than to vertical. To provide a metric of orientation sensitivity, orientation coherence thresholds (minimum coherence for just correct responding) were computed from both tasks. These are analogous to motion coherence thresholds measured with dynamic dot displays. Mean orientation coherence thresholds for observers ranged from 16% to 23% depending on task and signal orientation. These data indicate a remarkable human capacity to efficiently integrate the complex information inherent in noisy textures.

Since the world is full of featural heterogeneity, I examined here how observers combine features to efficiently construct an accurate representation of a texture. These experiments addressed two main questions. First, can observers efficiently make an orientation judgment about a texture that requires the integration of orientation information over space? Further, can such an integration process be quantified using a simple paradigm? To address these questions, I modified the motion coherence paradigm (Nawrot & Sekuler, 1990; Newsome & Pare, 1988; Williams & Sekuler, 1984) for use with texture stimuli. This allowed me to quantify observers' sensitivity to textures by measuring texture coherence thresholds.

My goal was to quantify texture perception in heterogeneous textures in which there are multiple orientations present, thus the motion coherence paradigm provided an appropriate framework for an analogous texture coherence paradigm. As such, observers were presented with a texture patch composed of oriented line segments each assigned to a signal or noise condition. The percentages of signal lines (oriented horizontally or vertically) and noise lines (randomly oriented in these and other orientations) were varied and observers were asked to either make a judgment as to whether the most predominant orientation within the patch was horizontal vs. vertical (Experiment 1) or to determine if the patch was "coherent" or "random" in orientation (Experiment 3). Percent correct responses were measured for each coherence value. To anticipate, the results showed that observers performed well on this integration task at a level consistent with the motion coherence threshold. These data further indicate that the texture coherence paradigm can be used to quantify texture integration processes. This paradigm is used throughout this thesis and the general methods section provides a description of the methods used that are typical for most experiments in this and subsequent chapters. Exceptions to these methods are outlined separately for each experiment.

General Methods

Observers

All participants were healthy adults and had normal or corrected to normal acuity. None had a prior history of ophthalmologic or neurological disorder. Informed consent was obtained prior to participation.

Apparatus

Texture stimuli were generated by a Cambridge Research Systems Visual Stimulus Generator 2/3 graphics card within a Dell P133 MHz computer, operating customized software written by me in the Delphi Object Pascal programming language. Stimuli were displayed on a gamma corrected Eizo Flexscan (T562-T) 17" monitor with a vertical refresh rate of 100 Hz and a resolution of 2.78 pixels/mm. Observers were supplied with a chin rest to stabilize head position and viewed the display binocularly from a distance of 92 cm in a normally lighted room.

Texture Stimuli

Texture patches, illustrated in Figure 12, were composed of black line segments distributed across a borderless square area and viewed against a uniform gray field (Michelson luminance contrast of 50%). Line segments subtended 1.34 x 13.4 arc min and could be oriented around their centre point in one of 16 possible orientations ranging from 0° to 168.75° separated by 11.25° (using Cartesian coordinates). The position of each line segment was non-random: each was constrained to fall within one cell of an 11 x 11 square grid, but the centre point of the line segment was "jittered" randomly from the cell's centre point by x and y values ranging from -5 to +5 pixels. The positioning grid was used to prevent intersection of line elements. Jittering was used to minimize vertical or horizontal alignments that can influence perception of texture elements (Field et al., 1993; Polat & Sagi, 1993; 1994).

A. B. C.

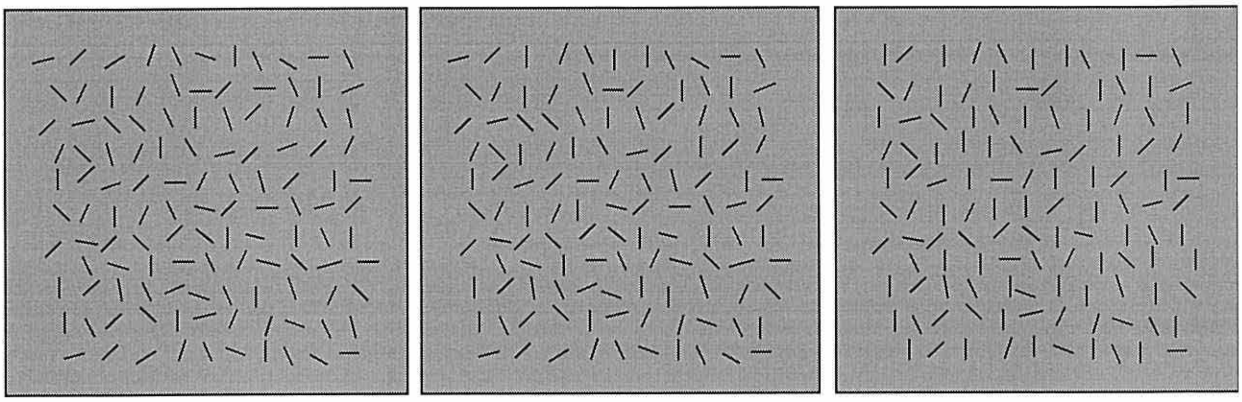

Figure 12. Texture coherence task: Is the typical orientation of the texture horizontal or vertical? Coherence values of 0%, 20% and 40% with a vertical signal direction are illustrated in panels A, B and C respectively.

For each test texture patch, a percentage (referred to here as percent coherence) of randomly chosen line segments was designated as "signal" and all lines within this subset were oriented identically, either vertically or horizontally. All other lines were designated as "noise" elements and randomly assigned one of the 16 possible orientations (which included both signal orientations). Thus, at 0% coherence, the texture patch was composed of randomly oriented line segments, and at 100% coherence, the patch was completely homogenous and composed of lines all oriented in the same orientation.

Pattern Masks

The pattern mask used here was identical in global shape to the test patches. It was composed of black "plus signs" each made from an intersecting horizontal and vertical line segment identical to those used in the test patches. The rules governing the position of each mask element were the same as those used for test textures and their locations were independently calculated. Mask duration was typically 500 ms.

A Typical Trial

A trial, initiated by a key press, began with a 1500 ms presentation of a central white fixation cross (identical to the crosses used in the mask patch). A test patch was then presented and followed immediately by the pattern mask (as illustrated in Figure 13). The observer's task was to judge the predominant orientation of the elements in the texture test patch (horizontal vs. vertical) and to press an appropriate key at the end of the trial. Observers were encouraged to guess if they were unsure and no feedback was provided during the experiment.

Experiment 1 - Orientation Coherence Thresholds

Observers

Nine healthy adults (3 male, 6 female, mean age = 32.4 years, SD = 7.3 years) took part in the experiment. All were inexperienced psychophysical observers and were naive to the study's purpose except for two observers who were my thesis advisor and I.

Stimuli & Procedure

Texture stimuli used here are illustrated in Figure 12 and a typical trial is represented in Figure 13. Test patch duration was 200 ms. Signal orientation was vertical on half of trials and horizontal on remaining trials. Observers judged the predominant orientation of each test patch from these two alternatives. From trial to trial, signal coherence was varied between 0% and 30% in 5% steps. Each coherence value for each signal orientation was presented five times in each block in a pseudo-random order resulting in 70 trials per block (that lasted about three minutes). Observers participated in three blocks within a single experimental session and all observers completed one practice block prior to testing to familiarize themselves with the task.

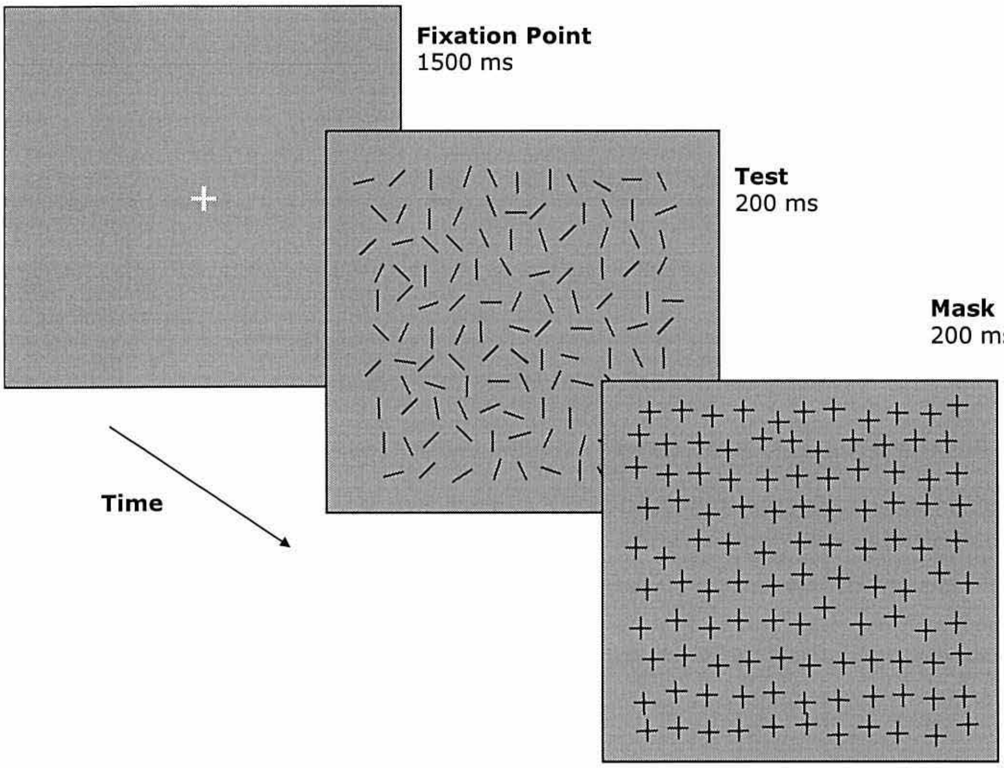

Figure 13. A typical trial was initiated by a key press and followed by a fixation point for 1500 ms. The test was 200 ms in duration and was followed immediately by a mask patch for 200 ms. Shown is a 20% coherence stimulus with a vertical signal orientation.

Results & Discussion

Observers' proportion correct scores were plotted as a function of percent coherence and orientation in Figure 14 to examine group mean psychometric functions. The proportion of correct orientation responses varied as a function of texture coherence. Proportion correct rose from chance levels with 0% coherence to nearly 90% correct for 30% coherence in the texture patch.

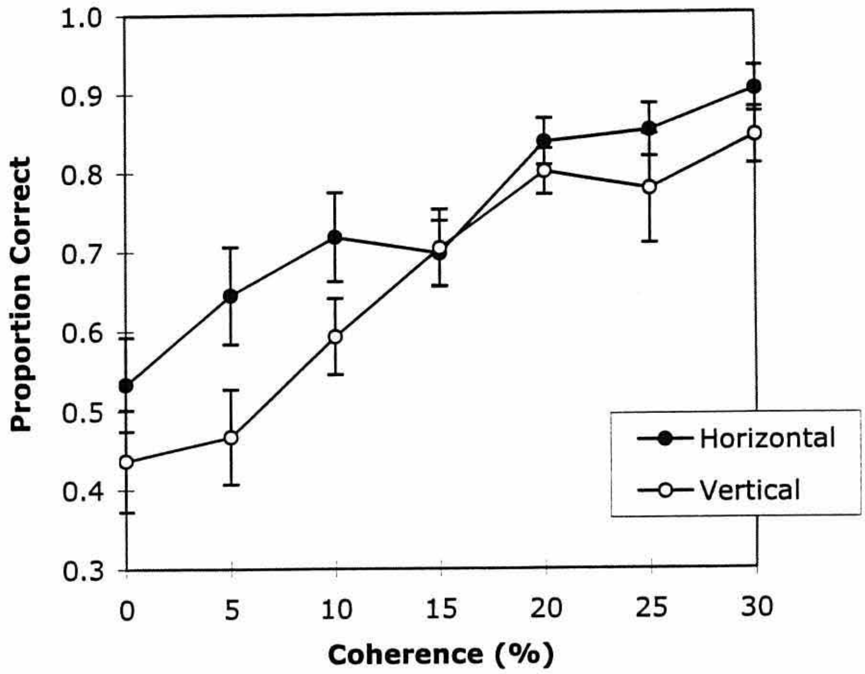

Figure 14. Mean psychometric functions plotting proportion correct as a function of percent coherence and signal orientation for nine observers. Observers were required to judge the orientation of the texture as horizontal or vertical. Error bars represent ± 1 s.e.

A repeated measures analysis of variance (ANOVA) was performed on the proportion correct scores using percent coherence and texture orientation as within-subject factors. Results showed a significant main effect of coherence, $F(6,56) = 19.45$, $p < .001$, a significant main effect of orientation, $F(1,56) = 10.44$, $p < .01$, and a non-significant interaction. A subsequent t-test revealed that observers gave significantly more correct answers in response to horizontally oriented textures compared to vertically oriented textures, $t(1) = 51.7$, $p < .05$.

In order to address what the minimum percentage of coherence was that observers required to reliably determine orientation in these stimuli, I calculated the group mean coherence threshold as a function of orientation. This was done by converting proportion correct scores, as a function of coherence, to standardized Z-score values essentially making the psychometric curve a straight-line function. The percent coherence corresponding to the 75% correct point was then interpolated using a least squares method (see Appendix A for a more detailed explanation of this interpolation technique). The fit of the least squares lines was very good as indicated by derived r^2 values of .98 and .97 for horizontal and vertical, respectively. The group mean thresholds derived in this way were 14.8% coherence for horizontal and 20.3% coherence for vertical indicating that observers required approximately 14 to 20% of the lines oriented in the signal direction to reliably judge the correct orientation of the texture.

One issue in plotting the psychometric functions is whether coherence should be plotted on a linear or logarithmic scale. Fitting least squares lines to the group mean Z-score values plotted as a function of either linear or log coherence results in better fits for linear (r^2 = .97) than log (r^2 = .80) scales. Better fits with linear scales of percent coherence have also been reported for motion coherence (Raymond & Braddick, 1996).

The above results suggest that observers were more sensitive to horizontally oriented textures than to vertically oriented textures given the significant main effect for orientation. Whether this is a true difference in sensitivity as a function of stimulus orientation or merely due to experimental factors cannot be determined with this experiment. The fact that the responses required of observers (horizontal vs. vertical) were not independent of each other may have allowed an observers' response bias to affect the results. For example, any bias in an observer's tendency to report one orientation (i.e., pressing the horizontal response key more frequently) would appear as though he or she was more sensitive to that orientation over the other. Rather than taking the present results as evidence for differential sensitivity to textures as a function of orientation (which may indeed be the case), I further examined the possibility of an anisotropy in sensitivity in Experiment 3.

A further concern raised by this experiment was that since no constraints were placed on the locations of signal or noise lines, there was the possibility in these stimuli that a number of signal lines may have been "clumped" together in the display. This may have affected an observers' response on any given trial enough to cause resulting psychometric functions to be unreliable or misleading. In order to minimize the possibility of observers basing their judgments on a clump of similarly oriented signal lines, I developed an algorithm to distribute the signal throughout the display. Experiment 2 investigates this issue.

Experiment 2 - Accidental Contours

Since, in these texture stimuli, line segments were randomly placed within the "jittered" grid, there was a possibility of accidental contours or small clusters being formed that may have influenced orientation judgments. Indeed, studies of contour completion (Field et al., 1993) and lateral masking (Polat & Sagi, 1993; 1994) indicate that end-to-end or parallel placement of line segments may significantly affect their detectability. To address this issue, three observers from Experiment 1 participated in an additional set of conditions in which the locations of the signal elements were constrained to fall in a grid cell that was not immediately above, below, to the left or right of another signal element. If this "contour-control" algorithm detected that two signal elements were adjacent to one another in the texture such that they created an end-to-end contour or were parallel to one another, the algorithm randomly chose another grid cell in which to place the next signal element and repeated the checking procedure. This contour-control algorithm was only effective for coherence values below 50% as it was limited by the number of available grid cells in which to place signal elements. Using the algorithm resulted in a widely distributed signal within the patch, free of any contours or "stacks" of line segments. No attempt was made to control for the positions of noise elements.

Methods

Observers

Myself and two naïve but experienced observers participated. All had participated in Experiment 1.

Stimuli

Texture stimuli were generated as in Experiment 1 except that the contour control algorithm was used. Stimulus duration was 200 ms and stimuli were masked.

Results

Proportion correct responses (collapsed over signal orientation) were plotted as a function of percent coherence to view psychometric functions. As in Experiment 1, proportion correct increased with increasing percent coherence. Mean texture coherence thresholds for each observer were computed (interpolated as in Experiment 1) and compared to those obtained in Experiment 1 using the uncontrolled displays. A t-test showed that contour-controlled thresholds (group mean = 12.1% coherence, SD = 3.3%) did not differ significantly from uncontrolled thresholds (group mean = 14.4% coherence, SD = 6.1%). This indicates that any contours occurring by chance in the uncontrolled stimuli of the initial experiments were unlikely to have had any significant effect on texture coherence threshold. Nonetheless, the contour-control algorithm was included in all subsequent experiments to ensure that the signal was sufficiently distributed throughout the display.

Experiment 3 - Coherent vs. Random Judgments

I was interested here to examine observer performance at integrating texture information across space and to measure how efficient observers are at this task. To establish that this spatial integration process is stable and measurable across tasks and in different individuals, I also asked observers to do a different judgment as to whether texture patches were "random" or "coherent" which is essentially a signal present vs. absent discrimination. This coherence judgment paradigm has two main advantages. First, the responses of "random" and "coherent" are independent of orientation, i.e., an increase in responding "coherent/signal present" for one signal orientation does not concurrently decrease the frequency of responding "coherent/signal present" for the other signal orientation. This does occur for the orientation judgment, i.e., if an observer increases her frequency of responding "horizontal" then there will be an apparent decrease in sensitivity for "vertical". A second, related advantage is that the coherent vs. random judgment allows the calculation of d' that is, according to Signal Detection Theory, considered to be a measure of sensitivity that is independent of and separable from observer criterion shifts (Green & Swets, 1966).

There were three other important differences between Experiments 1 and 3. The stimuli used here were the same as in the previous experiments however, test patches were not followed by a pattern mask. Subsequent experiments involving stimulus duration and masking (see Chapter 4) demonstrated that with a duration of 200 ms, the presence of the mask makes little difference to observer sensitivity. The second difference was that the contour control algorithm was implemented in Experiment 2 and all subsequent experiments. It was not used in Experiment 1.

A third important difference was that a formalized training phase (of about 64 trials) took place in the present experiment to familiarize the observers with the task. The stimuli used during this training were either 0% or 40% coherent to enable the observers to establish the difference between signal present and absent. In Experiment 1, only about 30 trials of practice were conducted on average depending on the confidence of the observer and the experimental stimuli were used, i.e., a range of coherences from 0 to 30% rather than a coherent/random or horizontal/vertical dichotomy. Given the total number of trials and the relatively easy discrimination tasks used here, this difference in practice had a negligible effect.

Methods

Observers

Fifty healthy adults (16 male, 34 female, mean age = 25.18 years, SD = 7.36 years) took part in the experiment. All were inexperienced psychophysical observers and were naive to the study's purpose. These observers also participated in Experiment 14 investigating texture contrast effects. The current data set was the baseline single task condition from that experiment.

Stimuli & Procedure

Texture stimuli used here were the same as in Experiment 1 except where specified. Texture patches were again oriented either vertically or horizontally and were presented for 200 ms. Patches were not pattern masked. Observers judged whether the texture patch was "random" or "coherent" and indicated their responses with a key press at the end of each trial.

Training phase: All observers completed at least two practice blocks prior to testing in which horizontal or vertical texture patches of either 0% or 40% coherence were presented eight times each for a total of 32 trials. A tone provided feedback in response to either a false alarm (i.e., responding "coherent" to a 0% coherence patch) or a miss (i.e., responding "random" to a 40% coherence patch). Observers had to complete a training block with no more than two false alarms and two misses to move on to the experimental trials. All observers reached this criterion within three practice blocks.

Experimental phase: From trial to trial, signal coherence was varied between 0% and 40% in 10% steps. Each coherence value for each signal orientation was presented seven times in each block in a pseudo-random order resulting in 70 trials per block (about three minutes). Observers participated in two blocks within a single experimental session so that performance on 14 trials was used to determine proportion coherent responses for each point on the psychometric function.

Results

Proportion coherent responses varied as a function of texture coherence. Proportion coherent responses rose from 14% with 0% coherence to 90% for 40% coherence in the texture patch. Figure 15 shows group mean psychometric functions as a function of texture orientation. A repeated measures ANOVA was performed on proportion coherent responses with percent coherence and texture orientation as within-subjects factors. Results showed a significant main effect of coherence, $F(4,45) = 207.4$, $p < .0001$, a significant main effect of orientation, $F(1,45) = 41.9$, $p < .001$, and a significant interaction, $F(4,45) = 2.9$, $p < .05$.

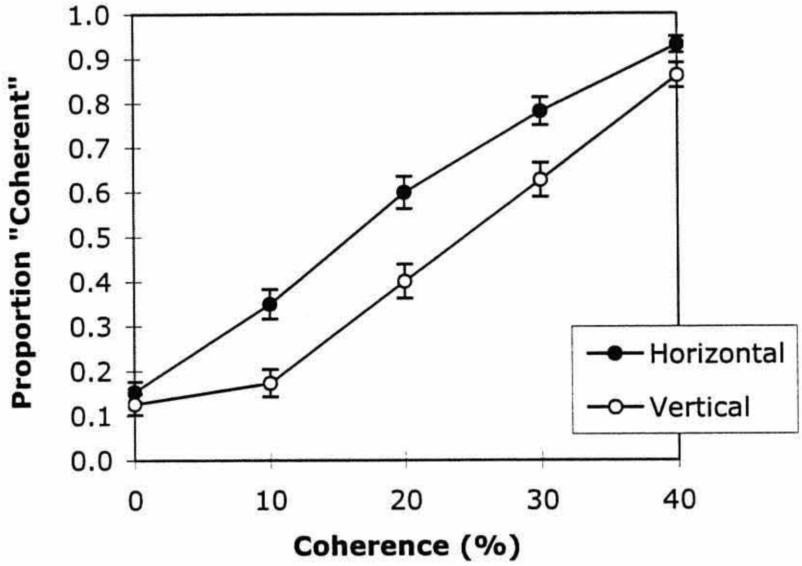

Figure 15. Mean psychometric functions for 50 observers plotted as a function of texture orientation. Observers were asked to judge whether the texture was "random" or "coherent". Vertical bars represent ± 1 s.e.

To address the minimum percent coherence required to just reliably respond "coherent", mean orientation coherence thresholds were calculated. Thresholds were calculated by converting proportion coherent responses to standardized Z-score values and interpolating percent coherence corresponding to the 50% coherence point using a least squares method. The derived r^2 values for the least squares lines were 1.0 and 0.98 for horizontal and vertical respectively, indicating extremely good fits. The group mean thresholds derived in this way were 16.69% coherence for horizontal and 23.34% coherence for the vertical signal orientation. Observers thus required 16 to 23% of signal lines coherently oriented (depending on signal orientation) to reliably judge the texture as coherent rather than random.

A significant interaction between orientation and coherence, i.e., a difference in slopes of the psychometric functions, often reflects a difference in task difficulty. In this case, the significant interaction between orientation and coherence may have indicated that observers found the coherent vs. random judgment easier for patches of horizontal signal orientation over those of vertical signal orientation. However, as is evident from Figure 15, the significant interaction here is most likely due to the convergence of the psychometric functions above and below thresholds rather than due to task difficulty, i.e., function slope. To rule out differences in task difficulty as a factor, I compared the slopes of the lines of best fit and found that they were not different (slope of horizontal function = 16.06, slope of vertical function = 16.88). Thus the difference between horizontal and vertical psychometric functions is most likely due to differential *sensitivity* as a function of orientation rather than differential *difficulty*.

I also calculated mean d' scores from proportion correct for each level of coherence (see Appendix B for a discussion of Signal Detection Theory and d' equations). Mean d' scores are plotted as a function of percent coherence in Figure 16. d' is a measure of observer sensitivity that is independent of shifts in observer criterion because it is based on the *relationship* between false alarm rates and hit rates, rather than on hit rates alone. Because changes in an observer's criterion should affect all stimuli equally, d' scores are relatively impervious to changes in criterion. See Appendix B for a demonstration of this. A repeated measures ANOVA was performed on this d' data using coherence and orientation as within-subjects factors. Results showed a significant main effect of coherence, $F(4,45) = 159.8$, $p < .001$, a significant main effect of orientation, $F(1,45) = 34.8$, $p < .001$, and a non-significant interaction. Using group mean d' scores to interpolate orientation coherence threshold corresponding to the d' value of 1.2 resulted in a threshold of 15.34% coherence for horizontal and 23.26% coherence for vertical [1]. These results are very similar to those found using proportion correct scores.

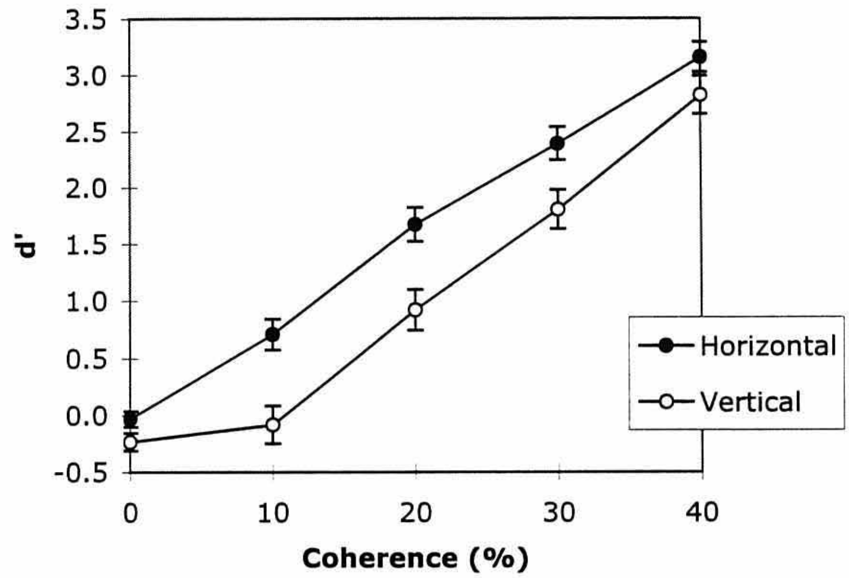

Figure 16. d' psychometric functions for horizontal and vertical signal orientations. Observers completed a coherent vs. random judgment task. Vertical bars represent ± 1 s.e.

[1] Note that the selection of a d' value of 1.2 for threshold calculations is fairly arbitrary here as is the selection of 50% coherent. Given a false alarm rate of .14, the d' value of 1.2 corresponds to a coherence value of 54% coherent responses which is reasonably analogous to other comparisons.

Discussion

In the previous three experiments, healthy human observers were asked to judge the orientation typical of line segments in static textures made from heterogeneously oriented elements. Results showed that observers require only about 16% coherence for just accurate orientation judgments and that this depends on the orientation of the signal lines. Observers were more sensitive to horizontally oriented textures than to vertical (see Chapter 3 for a further investigation and discussion of this finding). A contour control algorithm was then added to the texture generation program to ensure that clumps and contours were not being created by the signal lines and that these were not affecting observers' responses. Experiment 2 demonstrated that thresholds were not significantly influenced by the possible presence of accidental contours created within the texture patterns. In Experiment 3, observers were asked to determine whether the texture patches were "random" or "coherent" rather than to report the orientation of the texture. Results showed a very similar pattern of responding and comparable sensitivity between the two tasks. The fact that a large number of observers can complete a similar perceptual task using different behavioural responses and show very similar sensitivities suggests that there is a stable and pervasive underlying texture integration process. This process is both reliable and measurable in different observers and across different tasks.

What mechanism might mediate the remarkable sensory ability to abstract the typical orientation from textures containing a large proportion of noise information? A two-stage model analogous to that widely used to account for integration and noise reduction in motion perception (Wilson, et al., 1992) is a possibility (Wilson, Wilkinson & Asaad, 1997). Broadly speaking, the first stage of such a model consists of a set of low-level, local operations that extract first and second order information about luminance distributions in the image, possibly occurring early in visual processing (e.g., in V1 or V2). Such a stage would code information such as local motion or orientation and is thought to be perceptually inaccessible. The second stage is comprised of a set of global operations that integrate and reduce noise in dimension-relevant information (e.g., movement direction or orientation), probably occurring late in visual processing (e.g., at extrastriate sites such as MT or V4 for motion and orientation, respectively). Activity at this stage may be perceptually accessible and is undoubtedly linked with activity in the parietal and temporal lobes leading to action and object recognition.

Indeed, current models of texture perception, largely based on texture segmentation studies, suggest that local orientation information in line segment textures is first processed rapidly and in parallel by a bank of oriented spatial frequency filters followed by full wave nonlinear rectification (e.g., Beck, 1982; Bergen & Adelson, 1988; Caelli, 1982; Graham, Beck & Sutter, 1992). Both these operations (oriented filtering and full wave rectification) comprise the first stage because both are probably low level, automatic processes occurring early in visual processing most likely in V1 or V2 (Graham et al., 1992). The results of my experiments and results from a number of other studies investigating integration of orientation information support the idea that a second, higher order stage operates on the outputs of the first stage (Dakin & Watt, 1997; Field et al., 1993; Keeble et al., 1997; Kingdom et al., 1995; Polat & Sagi, 1993; 1994; Wilson et al., 1997). In this second stage, orientation information from spatially discrete elements is pooled over larger areas of the visual field (Kingdom, et al., 1995; Field et al., 1993) via orientation analyzers with crude, wide-band filtering (about 34 deg full width at half height for brief displays; Keeble et al., 1995). Studying global pattern recognition in Glass patterns (Glass, 1969; Glass & Perez, 1973), Wilson et al. (1997) suggested that the second stage may be mediated by activity in V4.

Dakin & Watt (1997) proposed that the collective activity of second stage orientation analyzers provides perceptual information about the central tendency of orientation in the texture's elements. The results presented here are consistent with this view and indeed could be explained by a mechanism that reflects the statistical average. However, simple averaging is unlikely to be the mechanism used as this would tend to "smooth over" smaller but distinct areas of texture that may be useful for processing and render the visual system "blind" to smaller distributions of orientations (as was suggested in Chapter 1). Keeble et al. (1997) showed that observers can discriminate a texture with two (albeit widely separated) peaks in its orientation distribution from one with a single peak at the same mean orientation. As has been suggested for motion, perhaps the output of the second stage analyzers competes for access to perceptual awareness. Only analyzers significantly more active than analyzers immediately adjacent on the orientation continuum would be successful. With the wide band noise employed here, the entire bank of orientation analyzers would be stimulated but only the analyzer most sensitive to the signal orientation would be activated more than the others and would therefore successfully gain control over behaviour. In this view, the orientation coherence threshold reflects the minimum signal-to-noise ratio needed to access perceptual awareness. However, this alone may not be enough to explain observer performance.

A cooperative system of orientation coding may, in turn, operate on the outputs of the second stage. In a cooperative system, the elements composing that system interact to create global behaviour. In this system, the responding of analyzers with a given preferred orientation would facilitate other like-oriented analyzers and would inhibit differently oriented analyzers. The resulting cooperative-competitive network of orientation detectors would be similar to that proposed in motion perception (Chang & Julesz, 1984; Snowden & Braddick, 1989a, 1989b; Watamaniuk, Sekuler & Williams, 1989; Williams, Phillips & Sekuler, 1986; Williams & Phillips, 1987) and would be especially suited for enhancing signal in noisy environments. Such a network would result in the efficient selection of the most numerous orientation present in a heterogeneous texture. A cooperative process is consistent with previous work in texture perception (Dakin & Watt, 1997; Keeble et al., 1997; Kingdom et al., 1995) lateral masking and contour completion (Field et al., 1993; Polat & Sagi, 1993; 1994), Glass Patterns (Wilson et al., 1997) and also with the perceptual model proposed by Grossberg and colleagues (Grossberg, Mingolla &

Todorovic, 1989; Gove, Grossberg & Mingolla, 1995). The role of a non-linear cooperative process operating on the sensitivity of individual decision units is very interesting in terms of the selection and coding of elements. A mechanism by which specific elements that are similar on some continuum, i.e., orientation, can be selected for further processing would be most beneficial in the parsing of a visual scene. Chapter 5 investigates the coding of individual element features and the role of attention in the selection process. Chapter 7 also investigates the role of selective attention in texture perception and addresses how the locus of attention modulates an observer's perception of textures.

In conclusion, the texture coherence paradigm is a simple, straightforward and effective way to objectively measure perceptual texture integration processes. In subsequent chapters, this paradigm is used to investigate numerous other aspects of texture perception such as orientational isotropy, featural coding of elements and the persistence of integration processes through time.

Chapter 3 - Oblique and Global Precedence Effects in Oriented Textures

Abstract

Sensitivity to high spatial frequency gratings is reduced for oblique orientations relative to the cardinal orientations of vertical and horizontal. This orientation-specific anisotropy in sensitivity is called the oblique effect. Since this anisotropy in sensitivity may persist for the perception of supra-threshold oriented textures, I asked whether sensitivity for obliquely oriented complex patterns was different from that for cardinally oriented patterns. I also addressed whether global precedence has an effect on texture perception, in light of the significant sensitivity differences between vertical and horizontal orientations in Chapter 2. Typically in the global precedence effect, the global contours of an object affect the perception of the elements making up that object. To examine possible global precedence effects of patch shape on sensitivity to oriented textures, I used square, diamond-shaped and octagonal texture patches in which each signal orientation either had congruent global edges or did not. Both orientation judgment tasks and coherent vs. random judgment tasks were used to measure sensitivity. Results showed that there were advantages in sensitivity for cardinal over oblique textures but only when the global contours of the patch were parallel with those cardinal orientations, i.e., for square patches. When diamond-shaped or octagonal patches were tested, the orientation-specific sensitivity differences disappeared. This indicates that while there are significant global precedence effects that operate on these texture patches, the orientation anisotropy seen in contrast sensitivity for grating patterns does not persist in heterogeneous texture stimuli. Furthermore, texture integration mechanisms are isotropic with respect to orientation.

Oblique and Global Precedence Effects

In previous experiments I have presented here, signal lines were oriented in one of two cardinal orientations (horizontal or vertical) and only noise lines were oriented obliquely. Although the assumption was made that observers would show orientation isotropy for these supra-threshold texture stimuli, the issue of whether orientation-specific analyzers are indeed similarly sensitive among different orientations was not directly addressed. The previous results from Experiments 1 and 3 that observers were more sensitive to horizontal over vertical orientations is quite analogous to findings in motion perception. Observers are typically more sensitive to motion that is horizontal (i.e., left or right) than they are to motion that is vertical (i.e., up or down) (Raymond, 1994). There is also evidence that line segments of different orientations are not coded equally. For example, psychophysical studies examining the oblique effect (Appelle, 1972; Graham, 1989) have demonstrated that observers are typically more sensitive to vertical and horizontal gratings versus oblique gratings. Investigations into the global precedence effect have also shown that the global shape of a figure has significant effects on the perception of its composite elements (Navon, 1977). These findings prompted three experiments investigating the effects of differential sensitivity to oblique line segments versus cardinally oriented (vertical and horizontal) line segments in these texture stimuli.

Contrast sensitivity and acuity for high spatial frequency grating patterns has been found to be generally worse for oblique orientations than for cardinal orientations (Appelle, 1972; Graham, 1989; Mansfield, 1974). This anisotropy of contrast sensitivity is known as the oblique effect. There are numerous explanations for the oblique effect ranging from the predominant exposure to vertical and horizontal contours in our urban carpentered environment (Annis & Frost, 1973) to an innate anisotropy in the distribution of orientation selective cells in the primate visual cortex (Mansfield, 1974). Mansfield (1974) found that there is indeed a non-uniform distribution of orientation-selective cells in the primate visual cortex and observed a distinct bias in favour of the principal or cardinal axes (vertical and horizontal) over the obliques. Furthermore, this bias was most pronounced for cells at or near the fovea. The oblique effect has also been found in young infants only a few months old (Leehey, Moskowitz-Cook, Brill & Held, 1975), which supports the idea that this effect may be hard-wired in the brain from initial development. The mediating mechanisms underlying the oblique effect have yet to be determined but the fact remains that the effect is a reliable and pervasive phenomenon in human observers.

The oblique effect could have influenced texture integration processes in the previous experiments at two levels of processing, thus changing observers' apparent sensitivity to heterogeneous textures. First, at local levels of perceptual processing, individual line segments that were oriented in cardinal orientations may have been more easily detected than oblique line segments. If this was the case, then greater detectability of signal line elements over noise line elements might contribute to orientation coherence thresholds, thus inflating the apparent sensitivity of texture integration mechanisms. This would predict that orientation coherence thresholds for obliquely oriented signal lines should be higher than those for cardinally oriented signal lines.

Second, since past texture integration research (Dakin & Watt, 1997; Keeble et al., 1995; 1997; Kingdom et al., 1995) supports the idea that there are numerous orientation-selective higher level integrators that combine local orientation information over space, then whether these integrators are similar in sensitivity as a function of orientation remains to be tested. In other words, if there are numerous "channels" that integrate texture information, each sensitive to a slightly different orientation, are integrators that code oblique orientations comparable in sensitivity to those that code cardinal orientations? If they are not comparable in sensitivity, observers should show different threshold sensitivity as a function of signal orientation. Experiments 1 and 3 have already demonstrated a difference in sensitivity for horizontal and vertical signal orientations.

Another possible factor in the perception of oblique and cardinal lines in texture patches is the global precedence effect (Navon, 1977). When observers view a briefly presented letter composed of many small letters, they are faster to identify the local, i.e., small, letters if they are congruent with the global letter than if they are not. For example, when presented with a global "S" composed of small "S's", an observer will be faster to identify the local letters than if the global "S" was composed of small "F's" as illustrated in Figure 17. Thus, the global shape of the figure affects the perception of the components of that figure. In Experiments 1 and 3, the texture patch was square, providing a global edge congruent with cardinal orientations only and may have provided facilitating global cues for cardinal orientations and not for oblique orientations.

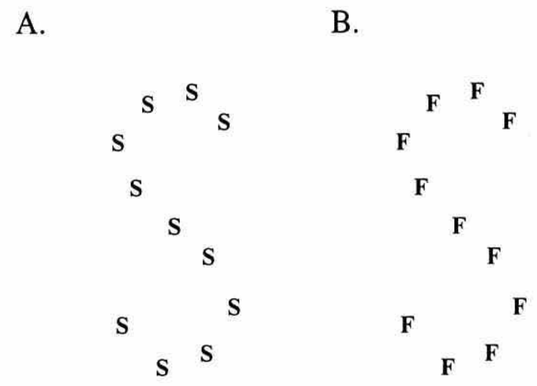

Figure 17. A demonstration of the Global Precedence Effect. Observers are faster to identify the small letters making up a figure when they are consistent with the global structure of that figure (Panel A) compared to when they are not consistent (Panel B).

A third possibility, which is perhaps less likely, is that the different sensitivities found in Experiments 1 and 3 were due to artifacts caused by the computer monitor. If the display pixels, for example, were slightly rectangular rather than completely square, a cardinal (specifically horizontal) line may have appeared longer than other lines. This may have caused the apparent advantage for horizontal signal orientations.

To address these possibilities, in Experiment 4, I assessed texture orientation coherence functions as before, except that I used oblique right (45°) and left (135°) signal orientations in addition to horizontal and vertical, creating a four-alternative forced-choice orientation task (4AFC). In these square stimuli, each of the outer edges of the square patch was parallel with the horizontal and vertical signal line segments but not with the oblique segments. If observers were affected by global precedence, this situation should result in a sensitivity advantage for cardinal orientations. I also tested a subset of these observers with the same stimuli but displayed the computer monitor tilted at 45°. This caused the outer contours of the "diamond-shaped" patch to be parallel with the oblique segments but not with the cardinal signal lines. If observers are truly more sensitive to horizontal and vertical lines over obliques, then this manipulation should make no difference and observers should still show an advantage for cardinal orientations in this situation. If the horizontal advantage seen previously was due to the global precedence of the texture patch contours, then the advantage here should be for oblique orientations.

In Experiments 5 and 6, I presented observers with an octagonal texture patch shape in which two outer contours were parallel with each of the four signal orientations. In the same observers, contrast sensitivity to sinusoidal grating patterns was measured for these four orientations to determine if anisotropies seen here could predict any anisotropy found with the coherence measure (Experiment 5). In Experiment 6, I used the same texture stimuli but asked observers to make a two alternative (2AFC) coherent vs. random judgment.

Experiment 4 - Square and Diamond Shaped Patches

Observers

Four female (including the author) and two male healthy observers participated (mean age = 28.6 years, SD = 5.39 years). All observers had participated in Experiment 1 and were therefore experienced at the texture task but were naïve to its purpose (except the author). In the "monitor tilted" section of this experiment, four of these observers participated [three females (including the author) and one male, mean age = 28.0 years, SD = 5.5 years].

Stimuli

Stimuli used were similar to those used in Experiments 1 through 3 except that two signal orientations of 45° and 135° were included in addition to vertical and horizontal. Illustrations of stimuli can be seen in Figure 18. In the "monitor up-right" condition, the monitor and stimuli were presented as in Experiment 1 (illustrated in Figure 18, Panel A). In the "monitor tilted" condition, the monitor was placed on a wooden stand such that it was tilted 45° counter-clockwise from its usual up-right position. Therefore, horizontal line segments became 45° line segments, 45° segments became vertical, vertical segments became 135° and 135° became horizontal (illustrated in Figure 18, Panel B). A dark cardboard mask with a circular hole around the texture patch was positioned in front of the monitor to cover the contours of the monitor. Coherence values ranged from 0 to 30% in 5% steps.

A. B.

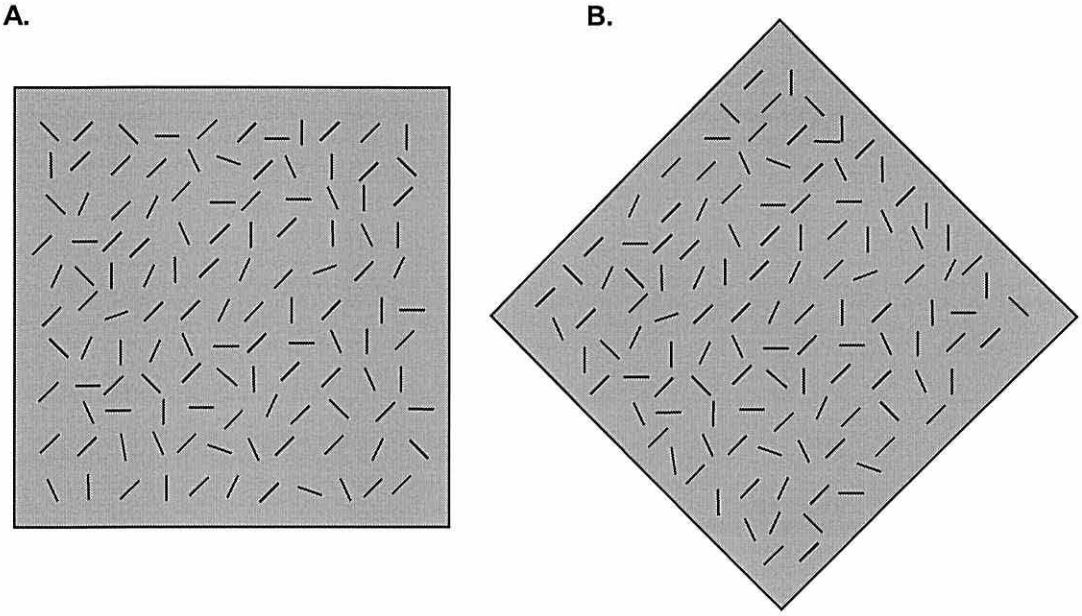

Figure 18. Illustrations of the "monitor up-right" (Panel A) and "monitor tilted" (Panel B) conditions of Experiment 4. In both examples, the signal orientation is 45° and coherence is approximately 40%. Observers were asked to make a 4AFC orientation judgment.

Procedure

Observers were asked to judge the predominant orientation of the texture patch from four alternatives, 0°, 45 °, 90 ° or 135 ° and to press the appropriate key at the end of each trial. A trial was initiated by a key press, followed by the 1500 ms presentation of the fixation point, a 200 ms texture patch and a 500 ms pattern mask. The positions of the elements making up the pattern mask were calculated separately to the test patches and the mask was always the same global shape as the test texture.

Each of four orientations and seven coherence values were presented ten times each for a total of 280 trials per block presented in a pseudo-random order. Observers completed two of these blocks for the monitor up-right condition and a subset of observers also completed three blocks for the monitor-tilted condition in a separate session. Note that the computer program was identical in the monitor up-right and tilted conditions. The same keys were used for each orientation and only the data analysis changed to incorporate the tilted monitor. Each session was completed in less than one hour.

Results & Discussion

Mean proportion correct scores were plotted as a function of coherence to examine psychometric functions for each orientation. These functions for both monitor conditions are illustrated in Figure 19. Proportion correct scores varied from chance (25%) to 90% correct as a function of coherence. Two repeated measures ANOVAs were performed on the monitor up-right and monitor tilted proportion correct scores using percent coherence and signal orientation as within-subject factors. For the up-right condition (illustrated in Figure 19, Panel A), results showed a significant main effect of coherence, $F(6,140) = 30.1$, $p < .001$, a significant main effect of orientation, $F(3,140) = 10.2$, $p < .05$, and a non-significant interaction. A one-way simple effects ANOVA was performed on the data using orientation (cardinal versus oblique) as a within-subjects factor. (A comparison between horizontal and vertical signal orientations revealed no significant differences, as did a comparison between the two oblique conditions.) Results showed that cardinal signal orientations have significantly higher proportion correct scores than do the oblique signal orientations, $F(1,12) = 5.34$, $p < .05$, demonstrating the oblique effect.

For the monitor tilted condition (illustrated in Figure 19, Panel B), results of the repeated measures ANOVA showed a significant main effect of coherence, $F(6,84) = 16.9$, $p < .001$, a non-significant main effect of orientation, and a non-significant interaction. Thus, there were no significant differences in proportion correct scores as a function of orientation, and no oblique effect.

The result that observers were more sensitive for cardinal signal orientations in the monitor up-right condition, compared to oblique orientations, could indicate one of two things. Either there was an oblique effect present that affected the perception of texture orientation such that observers showed significantly increased sensitivity for cardinal orientations regardless of the global shape of the texture patch. Or, texture orientation perception was influenced by the global precedence of the outer contours of the texture patch itself. Results from the monitor tilted condition, showing that there was no oblique effect when the outer contours were parallel to the oblique signal orientations, suggests that in the monitor up-right condition, observers were using the horizontal contours of the square patch on which to base their judgments and this resulted in the cardinal advantage. When the patch was diamond-shaped, there were no such cues and the oblique effect disappeared. These results suggest that the differences in sensitivity are due to the global precedence effect and not to an oblique effect.

If observers were using the outer contours of the texture patch as a cue, there should however, have been a corresponding advantage for the oblique signal orientations when viewing the diamond-shaped texture patch and there was no evidence of this. This finding suggests that there may be a weak advantage for cardinal orientations that causes them to be more sensitive to the outer contours of a patch over oblique signal orientations. In order to test for the presence of the oblique effect with a more neutrally shaped stimulus, Experiments 5 and 6 use an octagonal texture patch in which two outer contours are parallel with each of the four signal orientations. Experiment 5 also demonstrates the presence of the classic oblique effect in contrast sensitivity for oriented gratings.

One other important result to note here is that in Figure 19, Panel A, for the monitor up-right position, the function for the 0°, i.e., horizontal, signal orientation had the highest proportion correct scores which indicates that observers were most sensitive to these stimuli compared to stimuli of other orientations. Worthy of note however is the fact that in Figure 19, Panel B, for the monitor tilted condition, the function that shows the highest proportion correct and thus highest apparent sensitivity is the 45° orientation (although there was no statistically significant difference). These two functions correspond to the same physical stimuli on the monitor since the monitor was rotated counter-clockwise. Oblique effect explanations would not predict this pattern of results though global precedence may. Of concern is whether this pattern of results is caused by a monitor-specific artifact such as non-square pixel shape, as was mentioned above. If the pixels were rectangular, thereby causing horizontal lines to appear longer, this would explain both the apparently increased sensitivity for 45° line segments in the monitor tilted case and for 0° line segments in the monitor up-right case. This is actually a fascinating possibility if it were the case that perturbations in line features (as small as one or two pixels) were expressed as changes in texture integration sensitivity and could be measured through changes in proportion correct scores. However, a component of the calibration procedure that was completed before each experiment was a measurement of pixel "squareness". For this calibration, a 100 x 100 pixel figure was displayed on the monitor and the dimensions were measured using a straight edge. Monitor settings were adjusted accordingly until the figure was perfectly square. Given that this calibration was done before each experiment and that the orientation sensitivity changes seen in the last experiment can also be explained by the effects of global precedence, I don't believe pixel shape is a major issue, but is one worth keeping in mind. As far as the accuracy of the rest of the experiments in this thesis, at no time were horizontal, vertical, or oblique thresholds directly compared to one another without some sort of analysis that collapsed over orientation. For example, proportion correct scores in Chapter 6 were collapsed into "match" and "mismatch" groups that were completely crossed combinations of orientations. I therefore feel that this is an issue worthy of examination but one that does not significantly change the conclusions made for these experiments.

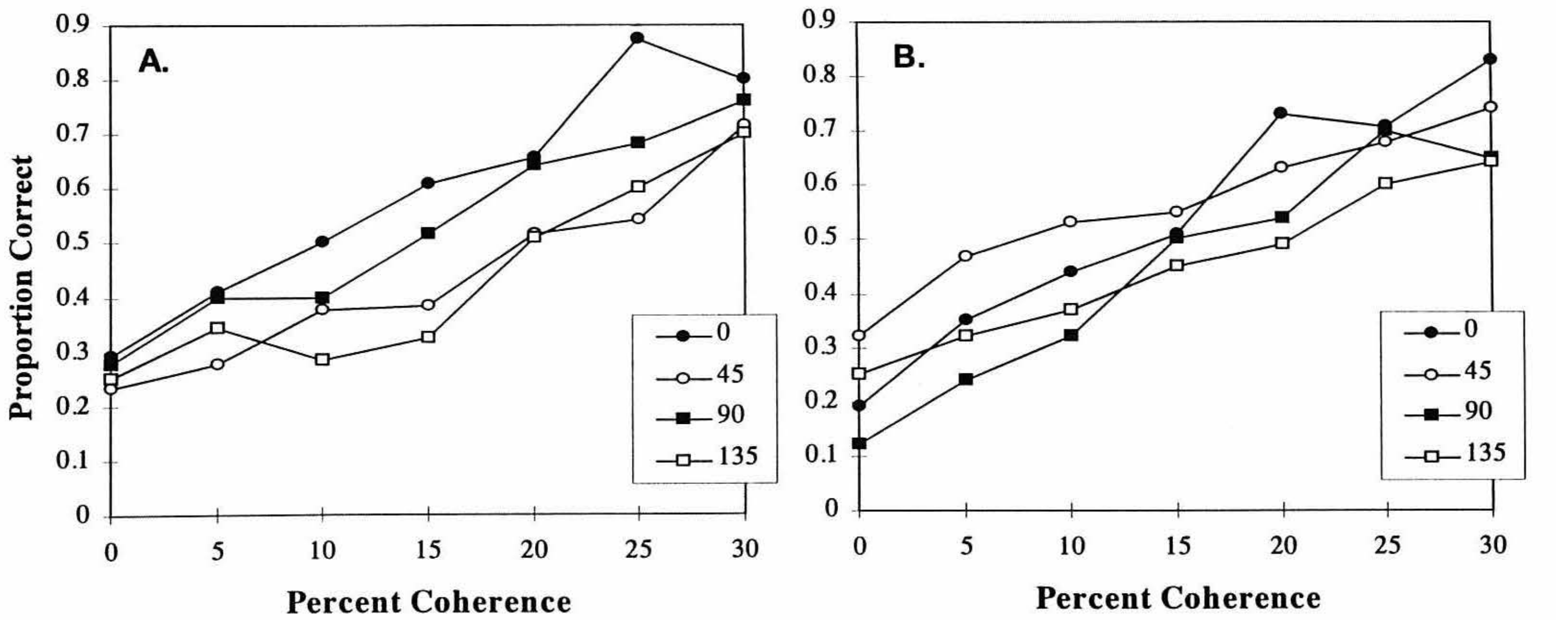

Figure 19. Mean psychometric functions for the monitor up-right (Panel A) and the monitor tilted (Panel B) conditions as a function of signal orientation. Observers showed an increased sensitivity for cardinally oriented signals (filled symbols) in the monitor up-right condition (A) but not for the monitor tilted condition (B). Texture patches were square (A) and diamond shaped (B) in this experiment.

Experiment 5 - Octagonal Patches with Orientation Judgments

Observers

Five healthy females participated (mean age = 19.2 years, SD = 1.3 years). All were naive to the experiment's aim and were inexperienced at both psychophysical tasks.

Stimuli

Contrast sensitivity test: A 2.2 deg circular sine wave grating pattern oriented at 0°, 45°, 90° or 135° (where 0° is horizontal) was presented on the same monitor as used for the texture tasks and was viewed from a distance of 3.6 m. The spatial frequency of the grating was 23 c/deg and the Michelson contrast was 1, 2, 4, 8, 16 or 32%.

Texture test: Signal lines were oriented at 0°, 45°, 90° or 135°. Coherence values were 0, 25, 30, 35, 40, 45 or 50% [2]. Texture patches were generated as described previously but were arranged in an octagonal patch with a diameter of 5.8 deg and composed of 129 elements as is illustrated in Figure 20. An octagonal "plus sign" pattern mask was generated using the same rules as the test patch. Texture stimuli were viewed at a distance of 92 cm, as was the case for all other texture stimuli.

[2] Coherence values used here were chosen to maximize the sampling of responses around threshold. Since the 4AFC judgment was hypothesized to be more difficult than the 2AFC judgment, higher coherence values were used including 0% for comparison purposes.

A. B.

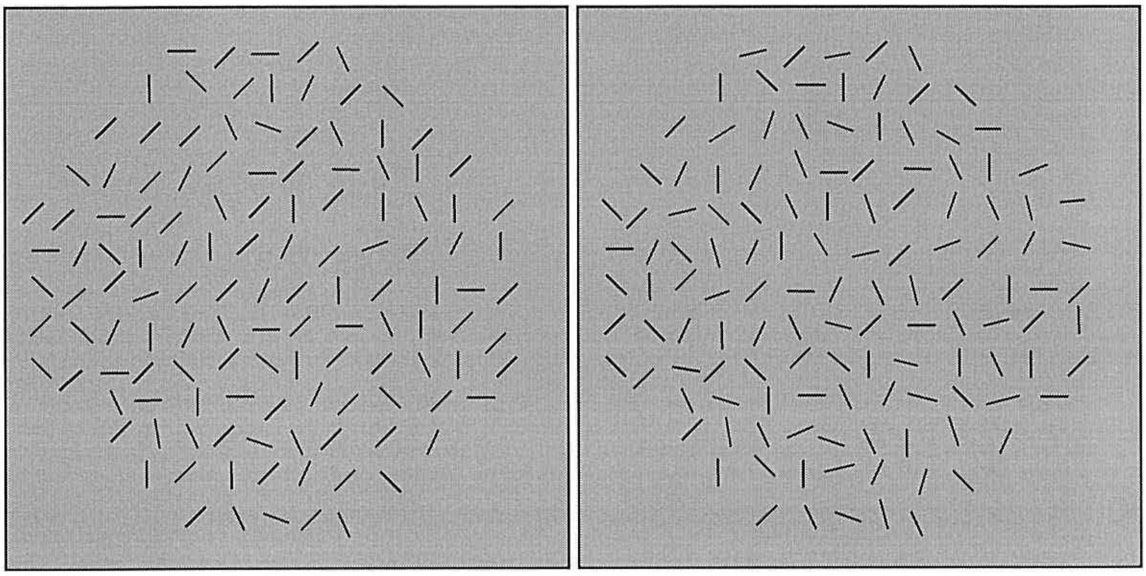

Figure 20. Illustrations of the octagonal stimuli used for the 4AFC oblique effects experiment. Shown in Panel A is a 40% coherent stimulus with a signal orientation of 45° (oblique right). In Panel B, the stimulus is 0% coherence.

Procedure

Contrast sensitivity test: Upon each presentation of a grating pattern, the observer judged its orientation, guessing when unsure. Stimuli were displayed until the participant responded. Each combination of orientation and stimulus contrast was presented five times, for a total of 120 trials presented in a pseudo random order. Testing was completed in ten minutes. Observers were given 15 practice trials with feedback to familiarize them with the task although no feedback was provided during the actual test.

Texture test: After completing the contrast sensitivity test, observers participated in the texture test. In all respects the testing procedure was the same as in Experiment 1 except a 4AFC was used as in Experiment 4. A trial was initiated by a key press followed by the presentation of the fixation point for 1500 ms, a stimulus presented for 200 ms, immediately followed by a pattern mask for 500 ms. Observers were required to determine the orientation of the texture patch from the four alternatives and press the appropriate key on the keyboard at the end of the trial.

Each combination of signal orientation and coherence was presented ten times in a pseudo-random order for a total of 280 trials. Observers completed at least 30 practice trials (with feedback) prior to testing to familiarize themselves with the task. No feedback was provided during the test phase and observers completed one block of each test in a 45-minute session.

Results

Contrast sensitivity test: Proportion correct for each observer was plotted as a function of log contrast for each grating orientation as illustrated in Figure 21. For oblique orientations, the psychometric functions were shifted rightward along the horizontal axis relative to those obtained with cardinally oriented gratings, indicating reduced sensitivity for oblique gratings. A repeated measures ANOVA was performed on these data using log contrast and orientation as within-subject factors. Results revealed a significant main effect of signal orientation, $F(3,96) = 9.69$, $p < .001$, a significant main effect of contrast, $F(5,96) = 36.02$, $p < .001$ and a significant interaction, $F(15,96) = 2.11$, $p < .05$. A t-test revealed that proportion correct scores were significantly higher for cardinally oriented textures than for obliquely oriented textures, $t(4) = 38.14$, $p < .001$, demonstrating the classic oblique effect.

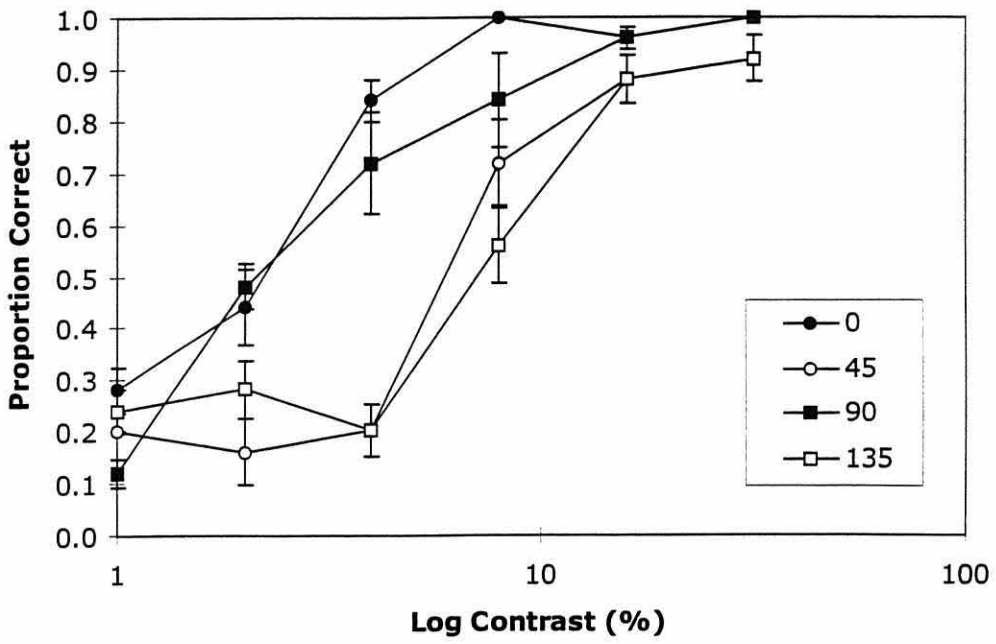

Figure 21. Psychometric functions for contrast sensitivity to oriented gratings. Observers were more sensitive to cardinally oriented gratings (filled symbols) over obliquely oriented gratings (open symbols) demonstrating the classic oblique effect (Appelle, 1972; Graham, 1989; Mansfield, 1974).

Texture test: Proportion correct for each observer was plotted as a function of percent orientation coherence as illustrated in Figure 22. Scores varied from chance (25% correct) at 0% coherence to over 80% correct for 50% coherence. I conducted a repeated measures ANOVA on proportion correct scores using signal orientation and coherence as within-subjects factors and found a marginally significant main effect of signal orientation, $F(3,112) = 6.4$, $p < .07$, a significant main effect of coherence, $F(6,112) = 13.2$, $p < .01$, and a non-significant interaction.

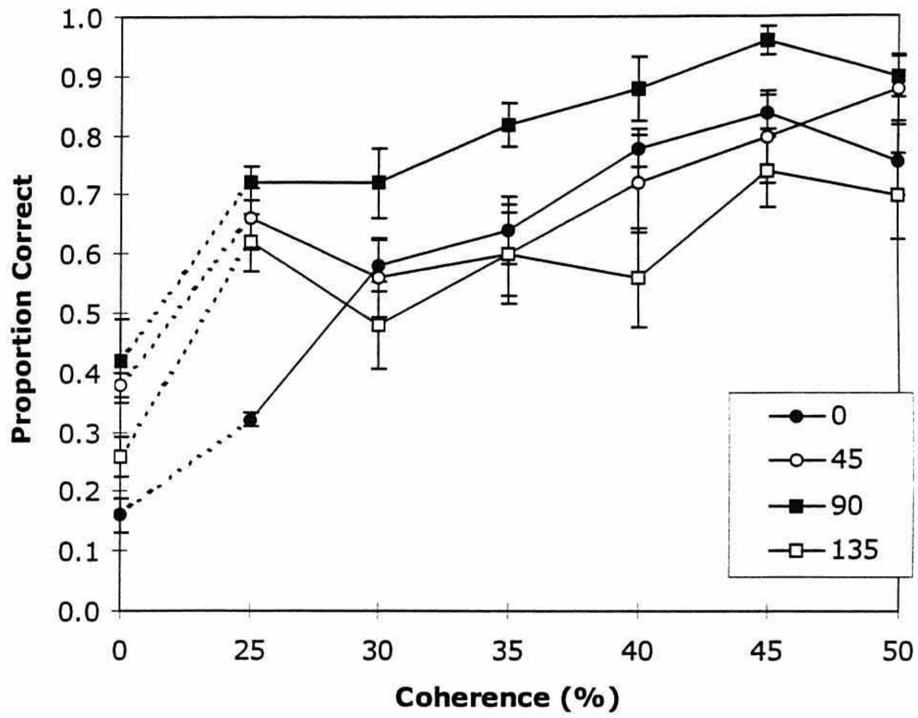

Figure 22. Illustrated are psychometric functions for a 4AFC orientation judgment (where 0° is horizontal).

Although the coherence values in this experiment were chosen to maximize the sampling immediately above and below threshold (and to provide a 0% coherence comparison point), observers performed better than expected on this task and as a result, the functions are quite flat for higher coherence values. The marginally significant main effect of orientation most likely results from the 25% coherence point on the horizontal orientation function that is significantly lower than the other functions. Whether the entire function for horizontal is significantly lower than the other functions is hard to determine here. To gain a more accurate measure of the psychometric functions, Experiment 6 sampled coherence values from 0 to 50% in 10% steps and used a 2AFC coherent vs. random judgment task.

Experiment 6 - Octagonal Patches with Coherence Judgments

The previous experiment demonstrated the classic oblique effect using contrast sensitivity in oriented grating patterns but showed ambiguous results for oriented textures. This was due to the supra-threshold coherence values that were used to test these observers. Also, because these observers were tested with a 4AFC orientation judgment, the responses they gave to each orientation were not independent of the other orientations. In this situation, any response bias for reporting any one (or two) orientation(s) is confounded with the apparent sensitivity for those orientations. This may have been what caused the psychometric function for the horizontal orientation to be displaced from the rest. In the present experiment, I test sensitivity for the same texture stimuli, but ask observers to use a coherent vs. random judgment that effectively separates the responses for each orientation. This results in each orientation having an independent measure of sensitivity and makes their sensitivity measurements comparable. This method also allows calculation of d' values which are a more pure measure of sensitivity changes, independent of criterion shifts (Green & Swets, 1966; see Appendix B for a discussion and demonstration of d' values).

Method

Observers

Ten healthy adults participated (8 males, 2 females, mean age = 25.9 years, SD = 5.17 years). All had previously participated in a psychophysical texture task but were naive as to the aim of the present experiment.

Stimuli

Signal lines were oriented at 0°, 45°, 90° or 135° (where 0° is horizontal). Coherence values were varied from 0 to 50% in 10% steps. Texture patches were generated as described previously and were arranged in an octagonal patch with a diameter of 5.8 degrees and composed of 129 elements. An octagonal "plus sign" pattern mask was generated using the same rules as the test patch (see Figure 20 for an illustration of the stimuli used here).

Procedure

Training phase: All observers completed at least two practice blocks prior to testing in which 0% or 50% coherent texture patches were oriented in one of four orientations (0°, 45°, 90° or 135°) and were presented four times each for a total of 32 trials. A tone provided feedback in response to either a false alarm (responding "coherent" to a 0% coherence patch) or a miss (responding "random" to a 50% coherence patch). Observers had to complete a training block with no more than 2 false alarms and 2 misses to move on to the experimental trials. All observers reached this criterion within three blocks of practice.

Experimental phase: In all respects the testing procedure was the same as in Experiment 3 (coherent vs. random judgment). A trial was initiated by a key press that displayed the fixation point for 1500 ms. Texture patches were then presented for 200 ms and immediately masked for 500 ms. Each combination of signal orientation and coherence was presented ten times in a pseudo-random order for a total of 240 trials per block. Observers completed two blocks in one testing session and at least 30 practice trials prior to testing to familiarize themselves with the task. Testing was completed in 45 minutes.

Results

Proportion coherent responses for each observer were plotted as a function of percent coherence and are illustrated in Figure 23. Scores varied from 14% at 0% coherence to over 96% coherent responses for 50% coherence. I calculated d' values for each observer using their mean proportion coherent responses for the 0% coherence point as the baseline false alarm rate. There was no significant difference in false alarm rates among the orientation conditions. Mean d' values for all observers for each texture orientation are plotted as a function of orientation coherence in Figure 23. As can be seen in the figure, the d' functions for each orientation are not different. I performed a repeated measures ANOVA on the d' values using texture orientation and coherence (0 to 40%) as within-subject factors. The 50% coherence point was excluded from this analysis because sensitivity for all orientation conditions had reached ceiling (as can be seen in Figure 23). The ANOVA showed a significant main effect of coherence, $F(4,180) = 94.88$, $p < .001$, a non-significant effect of orientation and a non-significant interaction. This pattern of results indicates that the oblique effect does not significantly affect sensitivity to complex oriented textures.

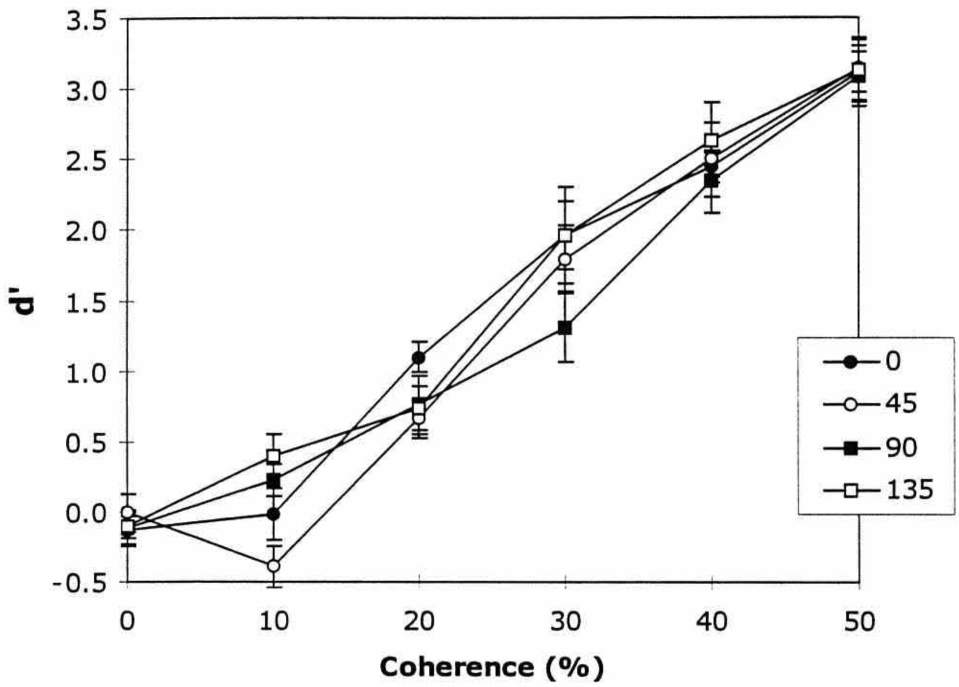

Figure 23. Mean d' values for ten observers making coherent vs. random judgments of a texture patch in one of four possible orientations (where 0° is horizontal).

Discussion

The last three experiments tested whether the oblique effect seen in contrast sensitivity for cardinal and oblique gratings would persist for supra-threshold texture patterns. In Experiment 4, square texture patches were used to demonstrate that if viewed up-right, the oblique effect is present for cardinal signal orientations over oblique signal orientations. If, however, the monitor is tilted such that the texture patch becomes a diamond shape, i.e., if global contours are parallel to the obliques rather than the cardinal orientations, the oblique effect disappears. This suggests that global precedence has a significant effect on orientation perception in textures and that there is no, or very little oblique effect in the absence of those parallel global contours.

Experiment 5 demonstrated the presence of the classic oblique effect with oriented gratings and tested sensitivity for textures in those same observers. Unfortunately, the coherence values that were used to test these observers were chosen to maximally sample around a higher threshold than actually resulted and thus were somewhat inconclusive. Experiment 6 used similar stimuli and a different distribution of coherence values, and observers were asked to complete a coherent vs. random judgment on the textures. Resulting d' scores showed that there were no significant differences in sensitivity for textures as a function of orientation, thus no oblique effect.

The fact that previous experiments in this thesis have demonstrated differential sensitivity as a function of orientation is interesting. In Experiments 1 and 3, observers were more sensitive to horizontal compared to vertical textures when the texture patches were square in global shape. This pattern of results does not conform to that usually seen with the global precedence effect, i.e., vertical signal lines also had a congruent global contour. There was a small possibility that the horizontal superiority may have been due to monitor-specific effects, as discussed above, such as a non-square pixel shape that would cause horizontal lines to look longer than other orientations, or a monitor refresh artifact in which horizontal lines gained some advantage as a result of the vertical refresh direction of the monitor. Part of the calibration procedure completed before testing included measuring whether a square figure on the monitor was indeed square, limiting the possibility that non-square pixels were the cause of horizontal superiority. If indeed monitor artifacts were the cause of these effects, they should persist through all experiments and all conditions. In fact, the horizontal superiority effect did not persist through all experiments (i.e., there is no orientation anisotropy in Experiments 4 or 6) and thus is a less likely cause.

In contrast to results showing differential sensitivity as a function of orientation, when the texture patch was diamond-shaped or octagonal, as in Experiments 4 to 6, there were no orientation-specific differences in sensitivity. It appears that observers were instead basing their orientation decisions on the global contours of the square stimuli, resulting in a global precedence effect, albeit for the horizontal plane only. With the octagonal stimuli, the cardinal global edges were significantly smaller, composed of only 5 elements compared to 11 in the square stimuli. This may have reduced the "strength" of the cardinal cues in the edges of the patch and in turn reduced the global precedence effect. While this partial global precedence effect may seem unlikely, the data is consistent with this pattern. It may be that there is a weak advantage for cardinal orientations that makes them more sensitive to global contours compared to oblique orientations. Interestingly, a similar pattern of results has been found in motion perception (Raymond, 1994) in which observers were more sensitive to horizontal (both right and left motion) compared to vertical motion when presented with square random dot motion stimuli.

For the octagonal stimuli, the absence of changes in sensitivity as a function of orientation indicates that texture integrators working to combine orientation information over space are isotropic with respect to orientation. There seems to be no advantage for cardinally oriented textures over 45° obliques and this suggests that the anisotropies seen in the oblique effect for contrast sensitivity to gratings, and demonstrated in Experiment 5, do not persist for supra-threshold stimuli such as those used here. This is consistent with other integration studies in the texture literature that do not report anisotropies (Keeble et al., 1995; Keeble et al., 1997; Kingdom et al., 1995). In conclusion, while there may appear to be an oblique effect in the perception of oriented textures, this is most likely due to the global precedence effect of outer patch contours. When a neutral patch shape is used, such as an octagon, texture integration mechanisms are isotropic with respect to orientation.

Chapter 4 - Stimulus Duration and Masking

Abstract

I have demonstrated that observers can accurately judge the most predominant texture orientation in a heterogeneous texture with 200 ms of exposure duration. What is the minimum duration that observers require to do this task? I addressed this question by varying the stimulus onset asynchrony between the test patch and the pattern mask from 90 to 200 ms and measured sensitivity in six observers. Observers' thresholds were higher for shorter durations but reached a minimum with a presentation approximately 180-200 ms in duration. This indicates that texture integration mechanisms require at least 180 ms to reach optimal performance in spatial combination of texture information over space and to determine the most predominant orientation in locally heterogeneous textures. It also suggests that higher brain areas may mediate the integration of orientation over space.

Stimulus Duration and Masking

Generally, as a perceptual task gets more and more complex, additional perceptual processing is required and therefore a longer exposure to that stimulus may be necessary for reliable and accurate responding. Why is examining the stimulus duration required for the completion of perceptual processing an interesting question to pursue? By investigating minimum stimulus durations, we can indirectly examine the levels of processing that an input has been through to reach its final processing state. For example, consider the situation in which an extremity touches a hot element on a stove. In response to pain, the limb will retract quickly from the source of the pain. Such a reflex arc is typically composed of just a few components; a touch receptor, a sensory nerve, an interneuron, a motor neuron and a muscle, that complete the circuit between the extremity and the spinal cord. Since responding to pain is such a basic reflex and requires little processing of input for a response to be made, it follows that the initial impulse should be processed quickly by the sensory system. By contrast, in order for a person to perform complex fine motor tasks, e.g., writing a passage with a pen and paper, information must travel to the brainstem and on to the somatosensory and motor cortices where a large amount of cortical processing must occur. With increases in the amount of processing and interconnections comes an increase in the processing time required.

While the mechanisms behind reflex arcs and fine motor tasks are not directly analogous to visual processing, the basic idea is similar. Hypothetically, a simple stimulus processed entirely in V1, e.g., a disc defined by luminance that is displayed on a uniform background, would require little or no processing at higher levels of the visual system and would therefore be processed in a fairly short duration. A similar disc that is defined only by the orientation of the composite texture and is superimposed on a heterogeneous randomly oriented background texture requires many more processing stages and interconnections in the processing pathway. If local orientation is processed initially in area V1, the integration and comparison among outputs of local orientation detectors must be done at a higher stage to receive the necessary outputs from the previous stages. Another set of processors must in turn compare among all orientations within the visual display to extract the presence of the area of differing texture. The extra processing would cause an increase in processing duration just as it did in the fine motor task example.

The actual duration required for a given perceptual task is difficult to predict. Simple perceptual judgments such as the detection of a light could be completed as quickly as it takes impulses to travel from retinal photoreceptors through subcortical areas to the cortex itself. More complex perceptual tasks, would require similar low level processing but may also require integration of information among numerous processing streams and would therefore require longer exposures. Generally, it seems that the more complex a task is, the longer the minimum required duration should be for processing.

According to this hypothesis, a complex task such as reading a word should require an extended exposure due to the amount of processing involved in resolving the letters, putting them in sequence and searching the lexicon for such a word. The Word Superiority Effect demonstrates that actually very little time is required for the identification of a word compared to unrelated letters with brief exposure durations (Cattell, 1886). Indeed, studies examining temporal attention using rapid serial visual presentation (RSVP) of letter and word streams have demonstrated that individual letters (Raymond, Shapiro & Arnell, 1992) and words (Shapiro, Caldwell & Sorenson, 1997) can be reported when presented for only about 15 ms and immediately masked. A similar finding to the word superiority effect that involves pattern perception is Configural Superiority (Pomerantz, 1981). Pomerantz demonstrated that a differing complex figure can be detected within a group of other complex figures with relative ease, compared to the detection of a simpler figure amidst a group of other simple figures. It appears as though certain complex figures have an advantage in perception over more simple figures, as do words over unrelated letter groupings.

Because of these types of superiority effects, the duration required for completion of a perceptual task is difficult to predict.

In the present experiment, I was interested in the effects of stimulus duration on texture integration processes. Experiments 1 and 3 demonstrated that observers could accurately judge the typical orientation in a heterogeneous stimulus with 16 - 20% coherence when displays were presented for only 200 ms and immediately followed by a pattern mask. In this experiment, I decreased the stimulus duration (i.e., the stimulus onset asynchrony, SOA, between test and mask) to determine whether similar levels of performance could be achieved with briefer intervals.

Experiment 7 - Stimulus Duration and Masking

Observers

A subset of six healthy adults from Experiment 1 (mean age = 31 years, SD = 8.7 years, 4 females, 2 males) participated, including myself. Myself and one naïve but experienced observer also completed these texture tasks without a pattern mask, for comparison purposes.

Stimuli

Texture stimuli were generated as in Experiment 1. Signal elements were oriented vertically or horizontally. The pattern masks used were created using the same rules as the test texture patches, however were calculated independently. Masks were composed of crosses made up of vertical and horizontal line segments identical to that used in the test patch. The positions of these crosses were independent of the positions of the signal lines but were still constrained to the same basic grid. Pilot experiments demonstrated that these crosses provided sufficient masking (as compared to similarly sized asterisks or small circles) and their effectiveness is also demonstrated in this experiment.

Procedure

Observers were asked to judge the orientation of the texture within the patch as horizontal or vertical. Test textures were presented for 90, 120, 150, 180 or 200 ms, followed immediately by a pattern mask for 200 ms. Stimulus duration (i.e., SOA) was blocked and the order of these blocks was randomized among observers. Within each block of trials, seven coherence values (0, 5, 10, 15, 20, 25 and 30%) were presented five times each for the two signal orientations in a pseudo-random order resulting in 70 trials per block. Observers completed 5 blocks, one for each duration, per session. Three such sessions were completed on different days. Two observers completed three additional sessions in which the pattern mask was not presented. The stimulus durations presented to these two observers were 30, 60, 90, 120, 150, 180 and 200 ms.

Results & Discussion

Proportion correct was plotted as a function of coherence in order to examine psychometric functions for each stimulus duration. Scores varied from chance (mean = .519, SD = .04) at 0% coherence to 85% correct (mean = .85, SD = .057) with 30% coherence (collapsed over stimulus duration). Mean psychometric functions are plotted in Figure 24 as a function of stimulus duration. Functions for the shorter 90 and 120 ms stimulus durations are displaced to higher coherence values in relation to the functions for the longer durations indicating differences in observer sensitivity. I performed a repeated measures ANOVA on proportion correct scores using stimulus duration and coherence as within-subject variables. Results showed a significant main effect of duration, $F(4,175) = 5.9$, $p < .01$, a significant main effect of coherence, $F(6,175) = 63.4$, $p < .001$ and a non-significant interaction.

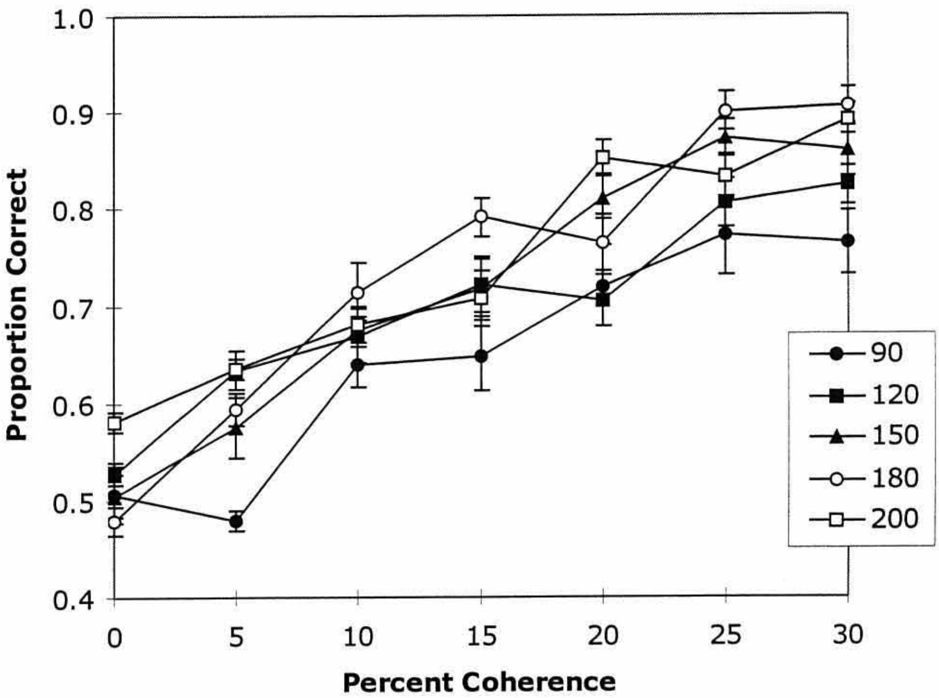

Figure 24. Mean psychometric functions for six observers plotted as a function of stimulus duration (ms). Error bars represent +/- 1 s.e.

To more closely examine observers' performance for the different stimulus durations, group mean coherence thresholds were interpolated that corresponded to the minimum percent coherence required for at least 75% correct responding. This was done by converting group mean psychometric functions, for each stimulus duration, to z-scores and fitting a line of best fit to the z-score function (all r^2 values were greater than 0.9 indicating extremely good fits). The percent coherence that corresponded to 75% correct was then interpolated using the equation of the line. See Appendix A for a discussion of this technique. Interpolated coherence thresholds are plotted as a function of stimulus duration in Figure 25. Thresholds were approximately 24% coherence for a stimulus duration of 90 ms and were reduced to 15% coherence for the longer stimulus durations of 180 and 200 ms.

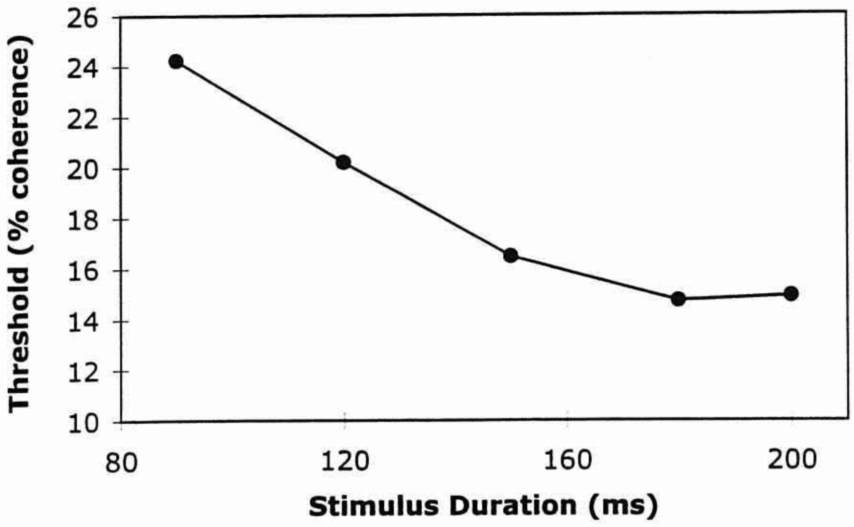

Figure 25. Interpolated coherence thresholds plotted as a function of stimulus duration. Thresholds reached a minimum for 180 to 200 ms stimulus duration.

For comparison purposes, two observers viewed the same stimuli without pattern masks and their thresholds are plotted as a function of stimulus duration in Figure 26. As can be seen in Figure 26, threshold functions for the no mask case for both observers were basically flat. Observers did not improve with longer durations due to the fact that they were as good as they could get even with a stimulus duration of 30 ms. This is most likely due to the persistence of the texture image in the visual system in the absence of a pattern mask following the stimulus. When a mask was present, however, thresholds were higher and did improve with increased stimulus duration. Points plotted off the chart represent stimulus conditions for which observers did not reach 75% correct (and therefore a threshold was not interpolated). As can also be seen from the elevated thresholds in Figure 26, the patches composed of crosses were effective pattern masks for this type of texture stimuli for the shorter stimulus durations. However, for longer durations such as 180 to 200 ms, it made little difference if there was a mask present or not, most likely because processing was complete by 180 ms.

For both observers, the minimum stimulus duration required to accurately determine texture orientation could be estimated as the point at which performance with a mask equals performance without a mask. For example, with the presentation of a mask, observer GH showed thresholds much higher than those without a mask for stimulus durations less than 150 ms. For a duration of 150 ms, performance approached that without a mask and with a presentation of 200 ms, his sensitivity was equal for both mask conditions. For observer HO (an experienced observer and the author), the stimulus duration required was much shorter. Her performance with a mask was similar to that without a mask at about 60 ms of exposure. For this observer, the minimum stimulus duration required for optimal responding was about 120 ms.

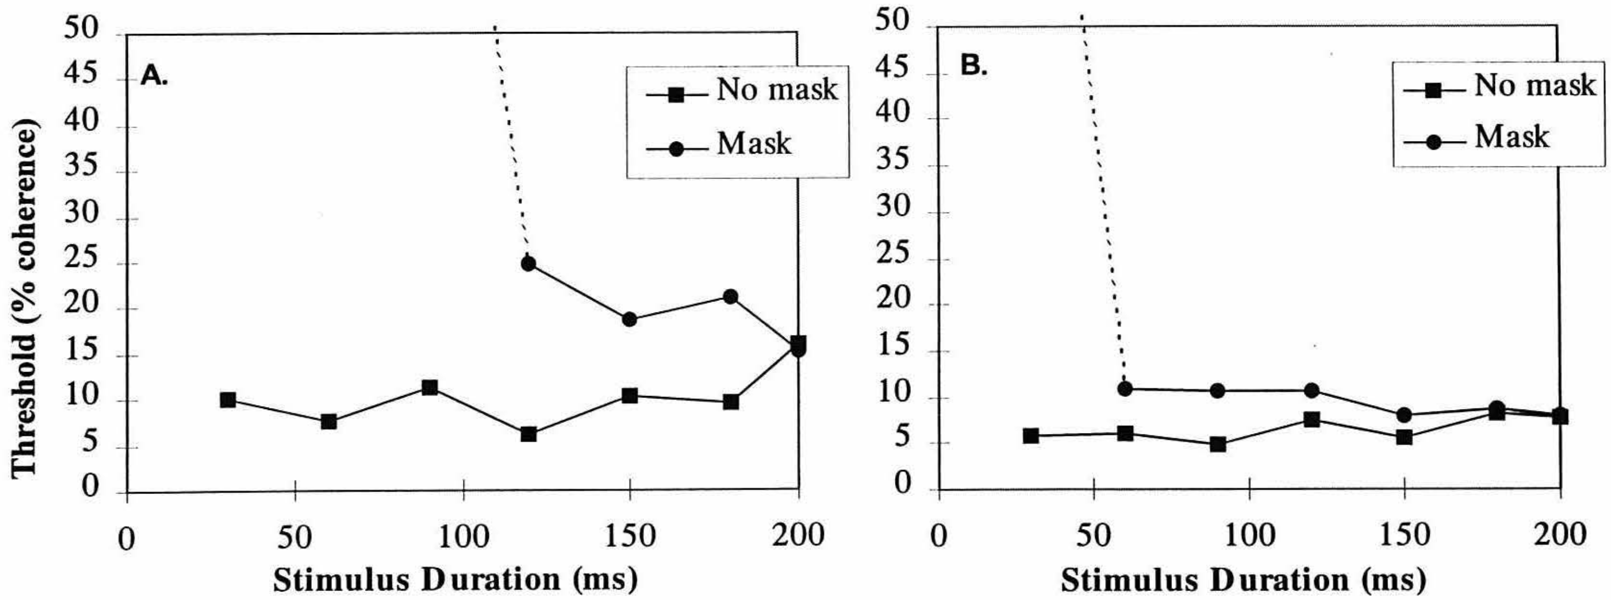

Figure 26. Coherence thresholds for two observers, with and without a pattern mask immediately following the texture patch. Minimum duration required for this task was approximately 200 ms for observer GH (Panel A) and 120 ms for HO (Panel B). Points plotted off the chart represent those points for which observers did not reach criterion (75% correct) at any coherence value.

Generally, observers required 180 to 200 ms to reach their optimal performance and minimize their thresholds. This is a relatively long time interval compared to the duration required for simple orientation discriminations of a line segment. A pilot experiment was conducted in which two observers (including the author) were asked to determine the orientation (horizontal or vertical) of a 100% coherent texture patch presented for 50 ms then followed immediately by a pattern mask. This task could be completed most easily and accurately by reporting the orientation of one line segment only and thus results are taken to be representative of that strategy. Both observers were 100% correct on each of the two blocks of 70 trials that were completed. The fact that a simple orientation judgment takes less than 50 ms to complete and a texture integration task takes at least 180 to 200 ms supports the idea that integration of orientation over space requires a higher level of processing in the visual hierarchy than simple orientation discriminations.

Other investigations into integration processes have demonstrated that observers can determine, with an exposure of 105 ms, which of two texture patches (presented at the same time) contained systematic orientation manipulation and which was random (Keeble et al., 1995). Observers were also given an exposure of 1000 ms with which they performed marginally better. Observers could also effectively determine which texture patch contained orientation modulation using a temporal forced-choice paradigm in which both exposure durations were 107 ms (Kingdom, et al., 1995). Similar results were found for a temporal forced choice task in which observers were to determine which of two texture patches had two main orientations as opposed to one. Results showed that observers performed the task very well with exposures of 105 ms but performed better with 1000 ms (Keeble et al., 1997). Since none of these stimuli were masked in these experiments [3] and since the

[3] In the temporal forced choice paradigm, the presentation of the second stimulus can be considered to be masking the first stimulus. However, the second stimulus itself is typically not masked at all. Temporal interactions between stimuli may also occur using this paradigm (see Experiments 12 to 15 for a demonstration of this).

manipulation of stimulus duration was quite coarse (or wasn't manipulated at all), it can only be concluded that with a stimulus duration of approximately 100 ms, observers can perform integration tasks but have better sensitivity for slightly longer durations. This is consistent with the findings presented here. Observers could perform the integration task and achieve a threshold of 24% coherence with an exposure of 90 ms, however, required a duration of at least 180 ms to minimize thresholds to about 16% coherence.

A study using Glass Patterns demonstrated that observers could efficiently judge, with a stimulus duration of 167 ms, which pattern in a temporal 2AFC presentation contained the degraded Glass Pattern and which was a random pattern (Wilson et al., 1997). The perception of the global structure in Glass Patterns also requires extensive integration of form information over space and is thought to be mediated in the V4 area of the ventral visual pathway (Wilson et al., 1997; Rentschler, Treutwein & Landis, 1994). Other tasks that require integration of information over space and require relatively long stimulus durations include the perception of illusory contours, object recognition, and face perception. Indeed, there is evidence that illusory contours, such as a vertical line demarcated only by the endpoints of rows of horizontal lines, are processed in Area V2 in the rhesus monkey (von der Heydt & Peterhans, 1989). Complete processing of a stimulus such as this would take longer than a stimulus coded earlier in the visual hierarchy. Other complex visual tasks can take an extended period of time depending on the task. For example, the identification of common objects such as toasters, frogs, or scissors, 150 ms (Olds & Engel, 1998); the perception of stereoscopic depth, 1000 ms (Tam & Stelmach, 1998); determining the magnitude of 3-dimensional slant, 500-8000 ms (van Ee & Erkelens, 1999) or identifying facial expressions, 500-1250 ms (Bradley, Mogg, Falla & Hamilton, 1998).

The result that observers require durations of about 180 to 200 ms to complete the texture task at hand is consistent with the multi-stage model of processing described in the introduction to this thesis and proposed by numerous perceptual researchers. No doubt the initial coding of orientation by banks of filters is relatively fast, it is the pooling of the orientation information across space and the competitive, cooperative interactions between orientation channels that takes time. In order to determine the most prevalent orientation in a heterogeneous texture, orientation information from all over the visual field must be integrated together. This is why the minimum duration required for this task is relatively long.

Chapter 5 - Luminance Polarity of Local Texture Elements

Abstract

Although visual textures in the natural world are largely comprised of featurally heterogeneous elements, humans can make rapid assessments of features descriptive of an average or *typical* element in such displays. Using an orientation coherence technique (observers judge the typical orientation of line elements in briefly presented textures with heterogeneously oriented elements), I probed the level of processing mediating such integrative processes. Using textures made from black and white lines, I asked whether orientation coherence was affected with mixed polarity displays and whether observers could selectively judge orientation coherence of just black or just white lines. Observers' orientation coherence thresholds (defined as the minimum percentage of coherently oriented lines necessary for just accurate orientation identification) were unaffected by mixed polarity displays or reductions in texture density. Surprisingly, observers performed above chance but below a level expected if the luminance polarity of each line element was available to awareness. I then presented observers with textures composed of competing signal orientations and asked them to selectively attend to either the white lines or the black lines and to determine signal orientation. Observers were again imperfect at this task, however, required 60% coherence of the competing orientation to reduce performance to chance levels. This indicates that selecting on the basis of one feature and making a judgment on another feature is imperfect in briefly presented heterogeneous textures possibly due to incomplete binding of features but that selective attention can assist in this selection.

Luminance Polarity of Texture Elements

The everyday visual world is largely composed of objects with heterogeneous textured surfaces. In such "noisy" textures, perception of features of the most predominant element may be critical for identification of the object. As the scene gets filled with more and more information, the perceptual *combination* of numerous predominant features is critical in finding a target object in amidst others. For example, in order to find a red apple amongst green apples and red tomatoes, one must be able to combine the coding of "Red" and "Apple". At what level this featural combination takes place and what featural information is available to perception are the issues investigated here.

Previous work of mine (O'Donnell & Raymond, 1997), and others (Dakin & Watt, 1997) has measured psychophysically the process of perceptual integration of individual textural features and provides evidence that humans can indeed accurately judge predominant features of elements in noisy textures using an integrative visual process. The general question I address here concerns what information about local elements is lost in such integration processes and what remains perceptually accessible. Specifically, I focus on orientation perception of "noisy" line segment textures and element luminance polarity and ask at what stage this featural information is combined to form a more complete representation of the element.

Figure 27 is a simplified representation of the hypothesized processing stages that take place in perception (adapted from Wilson, Wilkinson & Asaad, 1997). Current models of texture perception, largely based on texture segmentation studies, propose that local orientation information in line segment or Gabor micropattern textures is first processed in parallel by a bank of oriented spatial frequency filters as is illustrated in Panel A. Such processing is then followed by full-wave nonlinear rectification (e.g., Beck, 1982; Bergen & Adelson, 1988; Caelli, 1982; Graham, Beck & Sutter, 1992; Landy & Bergen, 1991; Malik & Perona, 1990). Both oriented filtering and full-wave rectification are probably low level, automatic operations occurring early in visual processing, most likely in V1 or V2 (Graham et al., 1992). Since full-wave rectification acts to eliminate luminance polarity information, the output of these early stages of processing should be independent of whether elements are darker or brighter than their backgrounds. Panel C represents further filtering by units sensitive to the orthogonal orientation, giving rise to an array of complex "end stopped" cells.

Studies investigating texture and pattern integration suggest that after these low level processes, a higher order set of mechanisms operates to group information (Panel D; Keeble, et al., 1995; Keeble et al., 1997; Kingdom et al., 1995; Wilson et al., 1997) and reduce any featural "noise" in the texture (Dakin & Watt, 1997; O'Donnell & Raymond, 1997). Information from spatially discrete elements may be pooled over large areas of the visual field (Kingdom, et al., 1995; Field, Hess & Hayes, 1993) via orientation analyzers with wide-band filtering (Keeble et al., 1995; Keeble et al., 1997). This yields a perceived "typical" orientation descriptive of the whole texture (Dakin & Watt, 1997; O'Donnell & Raymond, 1997) and shows a threshold function of sensitivity. Other features of the texture (such as polarity and colour) are coded separately and are combined at this later stage to result in a more complete representation of the target object (Panel E).

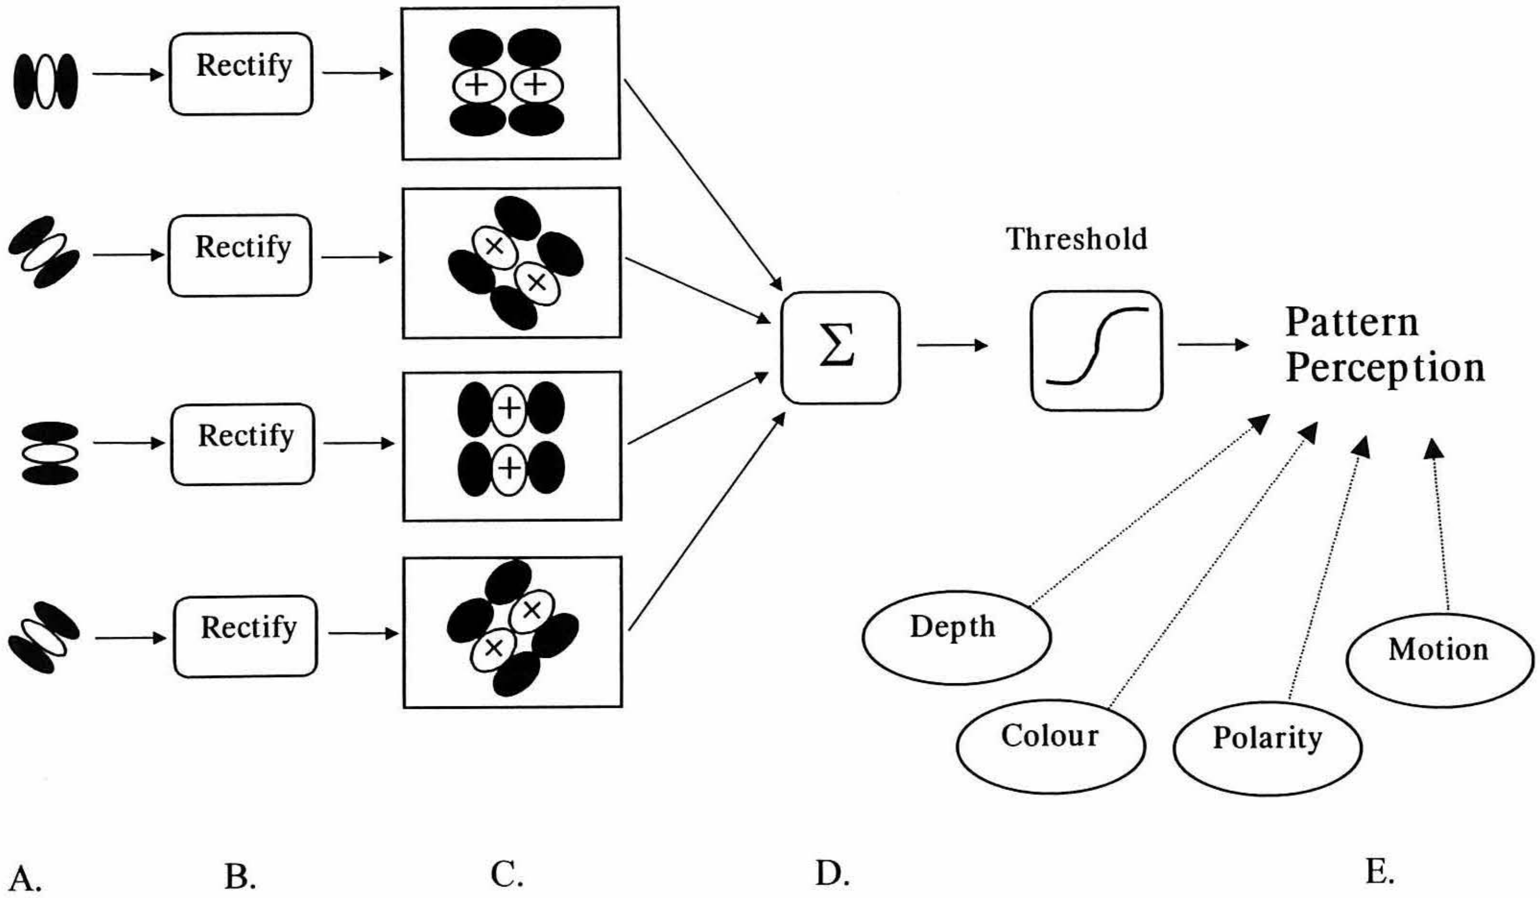

Figure 27. A representation of the perceptual stages of texture and pattern perception (adapted from Wilson et al., 1997). Processing progresses from Panel A initial oriented filtering to Panel E pattern perception. See text for details.

The combination of element information, represented in Panel D, probably contributes to the rapid perceptual emergence of global structure in certain spatially organized textures, e.g., Glass patterns (Glass, 1969; Glass & Perez, 1973; Glass & Switkes, 1976). Glass patterns are an array of dots in which the second dot in a large number of dot pairs making up a texture patch is systematically positioned relative to the first. Depending on the spatial relationship between the two dots, global structures, e.g., concentric circles, are perceived. Wilson, Wilkinson & Asaad (1997) asked whether the global structure in Glass patterns was perceptually constructed before or after full-wave rectification of dot stimuli, i.e., early or late in visual processing. They generated Glass patterns on a grey field in which half the dots were black and half were white. In some patterns, only white dot pairs had relative positions capable of promoting global structure (signal pairs) and all black dots were randomly positioned (noise). Although able to detect global structure in these briefly presented Glass patterns, participants were unable to discriminate patterns containing all white signal pairs and black noise from patterns with the reversed pattern of polarity. This result indicates that mechanisms capable of integrating element information to produce global structure cannot access luminance polarity information about individual elements. The implication is that such integration occurs after oriented filtering and perhaps as late in processing as extrastriate cortex V4, as Wilson et al. (1997) suggest. In four experiments I investigated whether the texture integration processes used to judge typical element features were similarly insensitive to luminance polarity.

Previous experiments have demonstrated that texture orientation judgments require at least 180 ms to reach maximal sensitivity for most naïve observers (see Chapter 4). This supports the idea that a more central integration mechanism mediates texture orientation judgments. If so, then texture models predict that such processes occur after full-wave rectification and would be therefore insensitive to differences in element luminance polarity. Wilson et al.'s (1997) observation that integration mechanisms contributing to *global structure* perception are insensitive to luminance polarity information supports this prediction. However, the polarity of individual elements in their mixed polarity Glass patterns is readily perceptible with a little scrutiny and it is possible that luminance polarity information may be coded differently when applied to global structure representations than when it is applied to element representations, especially when attention is directed specifically at element features. The following experiments attempted to answer these questions. To anticipate, their outcomes indicate that luminance polarity information is *not* lost during texture integration processes and that selective attention can influence this process.

Experiment 8 - Local Element Polarity and Patch Density

In the first experiment, I sought to determine simply if the presence of mixed polarity displays would elevate coherence threshold. If element polarity information is completely discarded prior to global integration, then thresholds should be unaffected by mixed polarity displays. However, if polarity information causes an increase in computational load for integration mechanisms then thresholds might show an elevation.

Methods

Observers

Four female observers (mean age = 27.5 years, SD = 9.88 years) participated in Experiment 8. All were naïve to the experiments' aims except for one who is an experienced psychophysical observer.

Texture Stimuli

I presented observers with three types of texture patterns in separate blocks (see Figure 28 for schematic representations of the stimulus types and Figure 29 for illustrations of stimuli). In all cases, the lines of different types were spatially distributed throughout the display.

1. ***Black & White, full-density*** - Half of signal line segments were white and half were black. Half of remaining noise lines were white, the rest black, creating a spatially mixed black and white texture. The percentage coherence of the black lines only (or the white lines only) was equivalent to the percentage coherence of the entire texture patch.

2. ***All Black, full-density*** - All lines were black. This is essentially the same stimulus as presented in Experiments 1 to 3.

3. ***All Black, half-density*** - Half the lines were black as in the first condition except that remaining lines were the same grey as the background, rendering them invisible. Density was therefore half of that in the first two conditions.

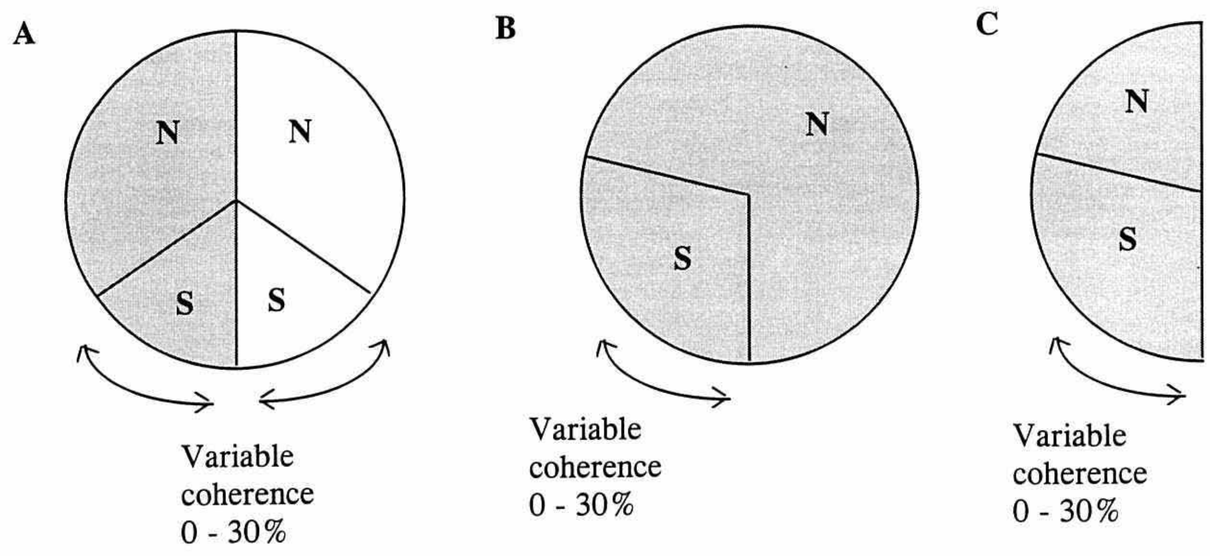

Figure 28. Pie charts depicting the signal and noise element distributions in the three stimulus types for Experiment 8. Shaded and non-shaded areas represent the proportion of line segments that were black and white, respectively. Sections of the pie containing an N were noise lines and sections containing an S were signal lines. Panel A. *Black & White, full density* - Half of all lines were black, the other half were white. Of the black (and white) lines, a proportion (0 - 30%) were signal lines and the rest were noise. Panel B. *All Black, full density* - All line segments are black and a variable proportion (0 - 30%) were signal, the rest were noise. Panel C. *All black, half density* - Half of the line segments were black and the rest were the same grey as the background rendering them invisible. Of the black lines, a proportion (0 - 30%) was signal and the rest were noise.

Procedure

Seven different coherence values (0, 5, 10, 15, 20, 25 or 30%) for each of two signal orientations (vertical and horizontal) were presented 5 times each in a pseudo-random order within a block of trials. For each observer, two such blocks for each texture type were conducted in a pseudo-random order yielding 20 trials for each point on the individual psychometric function (proportion correct as a function of percentage coherence). Testing was completed in about 50 minutes. Participants were asked to judge the "typical" orientation of the patch as vertical or horizontal and to ignore line element colour.

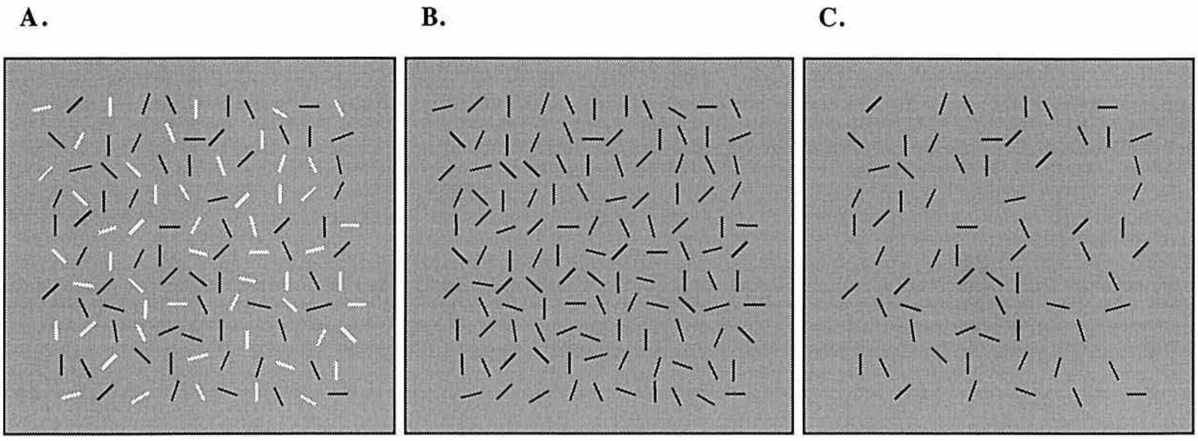

Figure 29. An illustration of sample stimuli from Experiment 8. Panel A. Black & White, full density. Panel B. All Black, full density. Panel C. All Black, half density.

Results

Proportion correct orientation identification varied as a function of texture coherence, replicating previous results (O'Donnell & Raymond, 1997). As can be seen clearly in Figure 30, the group mean psychometric functions obtained for each texture type are highly similar. A repeated measures analysis of variance (ANOVA) performed on proportion correct scores, using coherence and polarity condition as within-subject factors (collapsed over orientation), indicated that there was a significant main effect of coherence, $F (6,63) = 28.7$, $p \leq .001$, no significant main effect of polarity condition, and no significant interaction. I estimated orientation coherence threshold for each observer by converting their proportion correct to standardized Z-score values and interpolated percent coherence corresponding to the 75% correct point using a least squares method [4]. (The derived r^2 values were .69 or better in all cases). Coherence thresholds obtained with mixed polarity textures (mean = 14.32%, SD = 3.87%) were not significantly different from those obtained with uniform polarity textures of the same (mean = 13.44%, SD = 4.13%) or different (mean = 13.67%, SD = 4.23%) density. The performance of observers viewing these mixed

[4] I plotted Z-scores as a function of coherence using both linear and log axes and observed better line fits with linear than log scales.

polarity displays is very comparable to performance with baseline texture coherence tasks viewing single polarity displays, as in Experiments 1 through 3. It seems that having elements of different polarity within a texture has no effect on sensitivity to orientation (Treisman, 1988).

Additionally, this experiment demonstrates that element density had no effect on coherence threshold, at least within the range tested. This is a critical demonstration for further investigations. However, since there was no density effect, it is not clear from this experiment whether observers were selecting only the black (or only the white) elements in the black/white condition, or responding to the texture without regard to element polarity. Since the percentage of signal elements was the same for the whole patch as it was for each polarity subset, thresholds would be the same for either strategy.

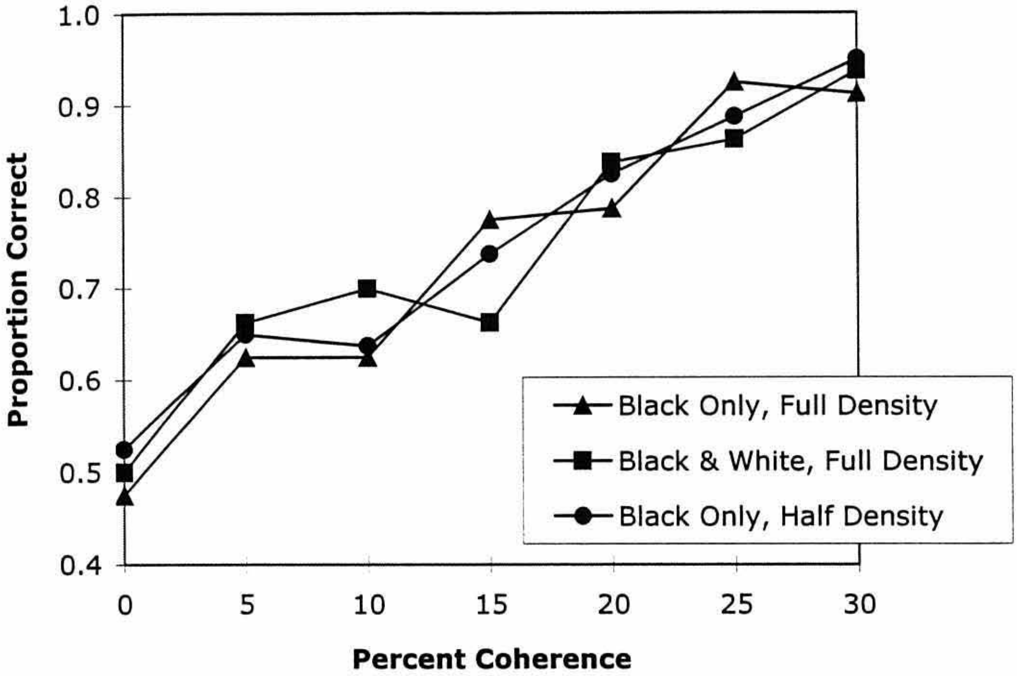

Figure 30. Group mean proportion correct as a function of percent orientation coherence for the three texture conditions of Experiment 8. Triangles represent data obtained with black-only elements, squares represent data from half the elements white and the other half black. Circles represent the condition with black only elements but half the density of the previous two conditions.

Experiment 9 - Black and White Competing Orientations

Since there was no significant effect of luminance polarity or density on sensitivity (within the range tested), the goal in this experiment was to determine if observers could selectively attend to elements of a single polarity and if this selection could improve typical orientation identification accuracy for mixed polarity, low coherence textures. If observers have no access to polarity information in these textures during the integration process, then the instruction to disregard lines of one polarity should offer no processing advantage over the case where no attention instruction was given. In other words, performance on the present tasks should be the same as performance in the previous experiment. Here, I presented four types of textures (each with a fixed percentage of coherence, i.e., one point on the psychometric function) and asked observers to judge the typical orientation of either the black lines only, or, in separate blocks, the white lines only. For clarity, I describe here the texture types used when observers were asked to judge the black lines only. (Replacing the word "black" with "white" and "white" with "black" in the following sentences provides the description for the "judge white only" conditions.)

For comparison purposes, I was able to use the psychometric functions from the previous experiment to predict the approximate proportion correct performance corresponding to a given percent coherence in the present experiment. For example, at 20% coherence, observers were consistently 80% correct in the previous experiment, regardless of polarity or density condition. This provided a comparison point for performance between the experiments.

Methods

Observers

Four naïve and one experienced observer participated (mean age = 24 years, SD = 9.57 years).

Texture Stimuli

There were four different texture stimuli types used in this experiment. All are represented schematically in Figure 31 and three example stimuli are illustrated in Figure 32.

1. ***Uniform polarity/non-competing signal*** *(schematically represented in Figure 31A and illustrated in Figure 32A)* - All line segments were black (a replication of the All Black, full-density condition of Experiment 8). The percentage coherence was fixed at 20% and was predicted to produce correct responses on 80% of trials.

2. ***Mixed polarity/non-competing signal*** *(schematically represented in Figure 31B and illustrated in Figure 32B)* - Half of lines were black and half were white. Unlike the mixed polarity condition in Experiment 8, signal lines were always black and noise lines were white or black. Of the black lines, 40% were oriented in the signal orientation and the rest were randomly oriented (noise). All white lines were noise. Thus, disregarding luminance polarity, the coherence of the entire texture was 20%, as in Condition 1. If observers integrated orientation information from black lines alone and successfully ignored white lines (yielding a coherence of 40%), then they should be correct on almost all trials (based on data from Experiment 8). If observers were unable to selectively attend to one polarity over another, (the coherence of the entire patch disregarding polarity was 20%), performance should have been 80% correct.

3. ***Mixed polarity/competing signal*** *(schematically represented in Figure 31C and illustrated in Figure 32C)* - In this condition, textures were also comprised of half black lines and half white lines. In this case, 20% of the black lines were oriented in the signal orientation and the rest were noise. Of the white lines, 20% were assigned the *perpendicular* signal orientation and the rest were noise. For example, if black signal lines were horizontal, then white signal lines were vertical. The percentage coherence of the black lines was 20% and if white lines could be effectively ignored, performance should be about 80% correct. If however, all lines were processed similarly regardless of polarity, then there would be an equal number of vertical and horizontal lines, and observers should perform at chance levels.

4. ***Mixed polarity/noise*** *(schematically represented in Figure 31D)* - This texture had a global coherence of 0%. Half of elements were black and half were white. This was used to provide a baseline for chance performance.

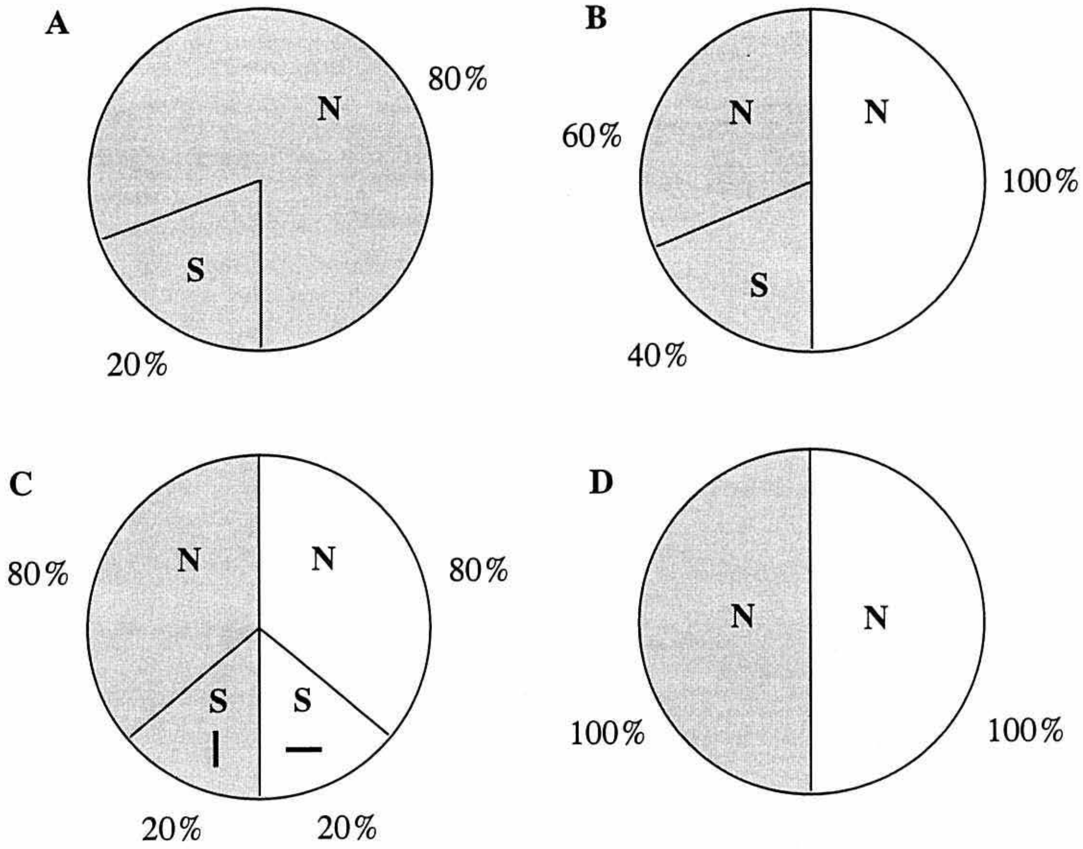

Figure 31. Pie charts depicting the relative signal and noise element distributions for the four stimulus types. See text for a detailed description.

A block of trials consisted of twenty repeats of each of these four polarity conditions with nominal signal orientation as vertical or horizontal making a total of 160 trials presented in a pseudo-random order. Observers completed four blocks, two in which they were asked to judge the typical orientation of the black elements only and to ignore the white elements and two in which they were asked to do the opposite. The order of the initially attended colour was counterbalanced among participants and "attend black" and "attend white" blocks were interleaved.

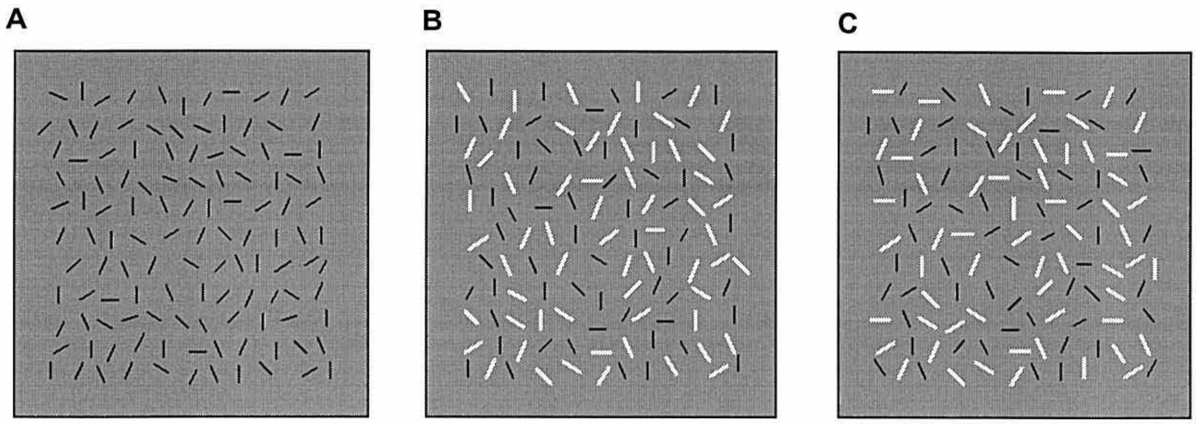

Figure 32. Illustrations of three of the stimulus types. Panel A. Uniform polarity, non-competing condition, 20% coherence. Panel B. Mixed polarity, non-competing orientations, 20% coherence. Panel C. Mixed polarity, competing orientations, 20% coherence.

Results

Mean proportion correct scores for five observers are shown in Figure 33 for each polarity condition. No statistically significant differences in proportion correct scores were found for observers attending to white vs. black line segments so data was averaged over attended signal polarity. This yields a total of 160 observations per condition per observer. Group mean proportion correct for Condition 1 (*Uniform polarity/non-competing signals*) was .77 (SD = .06), a value that corresponds well with that found with the same stimuli in the previous experiment (20% coherence, mean = .79, SD = .10).

In Condition 2 (*Mixed polarity/non-competing signal*), the overall coherence in the stimulus was 20%, the same as in Condition 1. The group mean proportion correct was .83 (SD = .09), a value that was non-significantly different from performance in the uniform polarity condition. As indicated by the asterisk in Figure 33, the performance level expected if observers had been able to selectively attend to elements of the target polarity (which had 40% coherence) was much higher than the observed performance. A one sample t-test showed that observed proportion correct scores for this condition were significantly lower than the predicted value of 1.0, t (4)

= -4.32, $p < .01$. It appears that in this condition observers were unable to select on the basis of polarity and therefore behaved as if all elements had a single polarity.

Group mean proportion correct for Condition 3 (*Mixed polarity/competing signal*) was .59 (SD = .04). Although the coherence for elements of the target polarity was 20% in this condition and should have produced about 80% correct responses if polarity information was available, performance fell well below this value. A one sample t-test showed that observed proportion correct scores were significantly lower than the predicted value of .80, $t(4) = -16.35$, $p < .01$. Again, it seems that observers were unable to use polarity information in making orientation decisions.

Group mean proportion correct for Condition 4 (*Mixed polarity, 0% coherence*) was at chance, as expected (mean = .50, SD = .02).

Although these findings seemed to suggest that observers were unable to select element orientation on the basis of luminance polarity, careful examination of the data indicates that this was not the case. I compared performance in Condition 3 (*Mixed polarity/competing signal*) with that obtained in Condition 4 (*Mixed polarity, 0% coherence*). In the latter condition, observers' performance was at chance, as expected. If, in Condition 3, observers were completely unable to use polarity information, performance should have been at chance because the number of horizontal lines was equal to the number of vertical lines. Without luminance polarity information to "tag" the target orientation, judgments of typical orientation should be evenly split between the two alternatives and performance should have been at chance. However, results of a one sample t-test revealed that performance in Condition 3 was significantly higher than chance, $t(4) = 4.95$, $p < .01$. It seemed that the observers in this condition were not performing as well as expected if they could use the polarity information in the stimuli to make their orientation decision, but were also not performing as poorly as expected if they could not use *any* of the polarity information for their orientation decision.

I calculated that the number of true signal lines (i.e., target-colour signal lines) divided by the number of noise lines defined by colour *or* orientation, (i.e., the sum of target-colour noise lines plus *all* distractor-colour lines) yielded a coherence value of 10% in this condition. Data obtained in Experiment 8 (Figure 30), predicts that with this coherence value, observers should be correct on about 63% of trials, a value not significantly different from the measured 59% correct (as indicated by a t-test). This finding indicates that rather than being able to completely separate signal from noise on the basis of polarity, the texture integration system may treat all other elements, which are not the target polarity, as noise. Experiment 10 further examines this possibility.

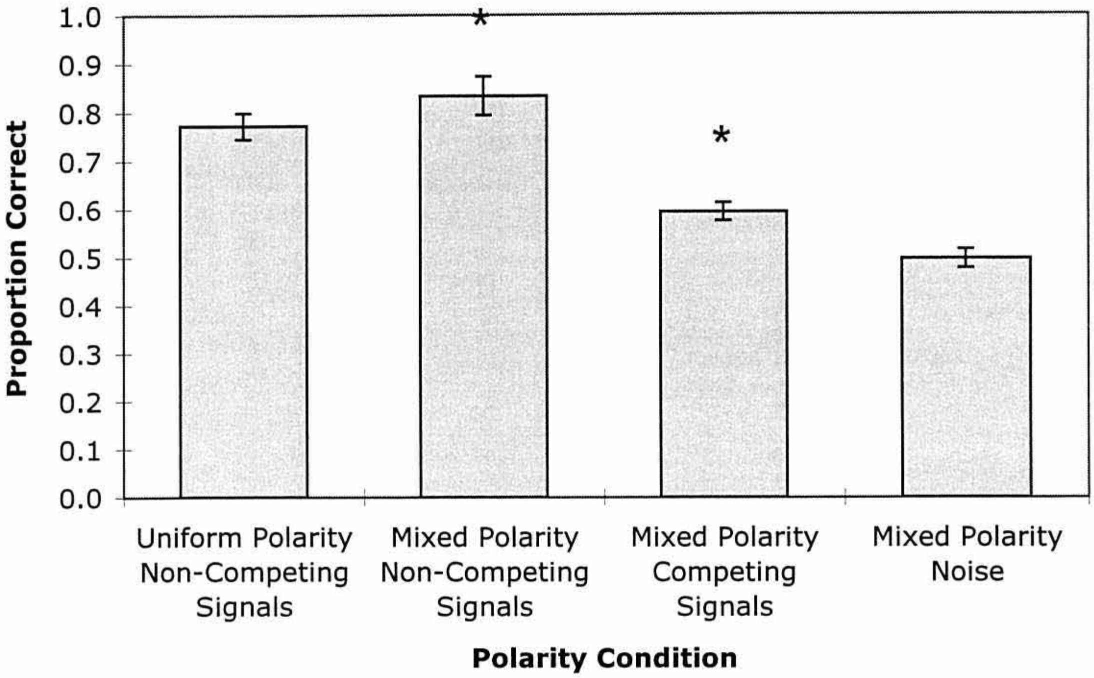

Figure 33. Group mean proportion correct for each of the four texture stimuli presented in Experiment 9. Three of the stimuli used are illustrated in Figure 32. The asterisks indicate the proportion correct predicted if observers were able to effectively select signal elements on the basis of luminance polarity. Vertical lines represent ± 1 s.e.

Experiment 10 - Red and Green Competing Orientations

The findings of the previous experiment indicate that observers do not have complete access to polarity information during texture integration processes, but also demonstrate that such processes are not completely insensitive to polarity. To further examine the role of selecting the orientation of a subset of lines on the basis of an orthogonal feature, I replicated Experiment 9 using red and green line segments that had the same polarity (i.e. both brighter than the background). I used colour here because colour information is typically thought to be processed independently from form and there is evidence of re-combination of this information at later processing stages (Treisman & Gormican, 1988; Prinzmetal, 1981). All conditions were the same as in Experiment 9 and observers were asked to either report the orientation of the red elements or the green elements in separate blocks.

Methods

Observers

Six naïve observers participated (mean age = 19.7 years, SD = 1.21 years). None of these observers had participated in the previous experiment.

Texture Stimuli

All conditions were the same as in Experiment 9 but line segments were red or green. Colours were visually matched in luminance contrast and were both brighter than the background grey (i.e., of the same polarity).

Procedure

Observers were asked to attend to the green lines and ignore the red lines or to do the opposite in separate blocks. Their task was to determine the predominant orientation of the texture patch. The order of the initially attended colour was counterbalanced among participants and "Attend Red" and "Attend Green" blocks were interleaved.

Results

No significant difference was found for proportion correct scores in the "Attend Red" versus "Attend Green" conditions, therefore proportion correct was averaged over signal colour. Mean proportion correct scores for the four conditions showed the same pattern as in Experiment 9 as can be seen in Figure 34.

Group mean proportion correct for Condition 1 (*Uniform polarity/non-competing signals*) was .79 (SD = .03), a value not significantly different from the same stimuli in the previous experiment (mean = .77, SD = .06). In Condition 2 (*Mixed polarity/non-competing signal*), the group mean proportion correct was .85 (SD = .04) a value that was also non-significantly different from performance in the previous experiment (mean = .83, SD = .09). As before, a one sample t-test showed that observed proportion correct scores in Condition 2 were significantly lower than the predicted value of 1.0, $t(5) = -4.23$, $p < .01$, indicating that the use of colour to discriminate signal from noise is imperfect. Rather, observers behaved as if they were treating all elements as the same colour (as they did the black and white elements).

Group mean proportion correct for Condition 3 (*Mixed polarity/competing signal*) was .65 (SD = .06). As before, the coherence for elements of the target colour was 20% in this condition and should have produced about 80% correct responses. A one sample t-test showed that observed proportion correct scores were significantly lower than the predicted value of .80, $t(5) = -5.76$, $p < .01$. This seems to indicate again that observers could not use element colour to help their orientation decision, however when performance was compared to chance levels, it seems that this was not the case.

I compared performance in Condition 3 with that obtained in Condition 4 (*Mixed polarity, 0% coherence*). In the latter condition, observers' performance was at chance (mean = .51, SD = .03), as expected. If, in Condition 3, observers were completely unable to use colour information, performance should have been at chance. However, performance in Condition 3 was significantly higher than chance, $t(5) = 5.914$, $p < .01$.

As in Experiment 9, the number of true signal lines (i.e., target-colour signal lines) divided by the number of noise lines defined by colour *or* orientation, (i.e., the sum of target-colour noise lines plus *all* distractor-colour lines) yielded a coherence value of 10% in this condition. Data obtained in Experiment 8 (Figure 30), predicts that with this coherence value, observers should be correct on about 63% of trials, a value not significantly different from the measured 65% correct.

It seems that while observers are not using colour or polarity information efficiently in these judgments, texture integration mechanisms still have access to that information under conditions of selective attention. Condition 3 (*mixed polarity/competing signal*) provides a critical comparison to address the role of selective attention to polarity and colour and Experiment 11 further explores this competing signals idea.

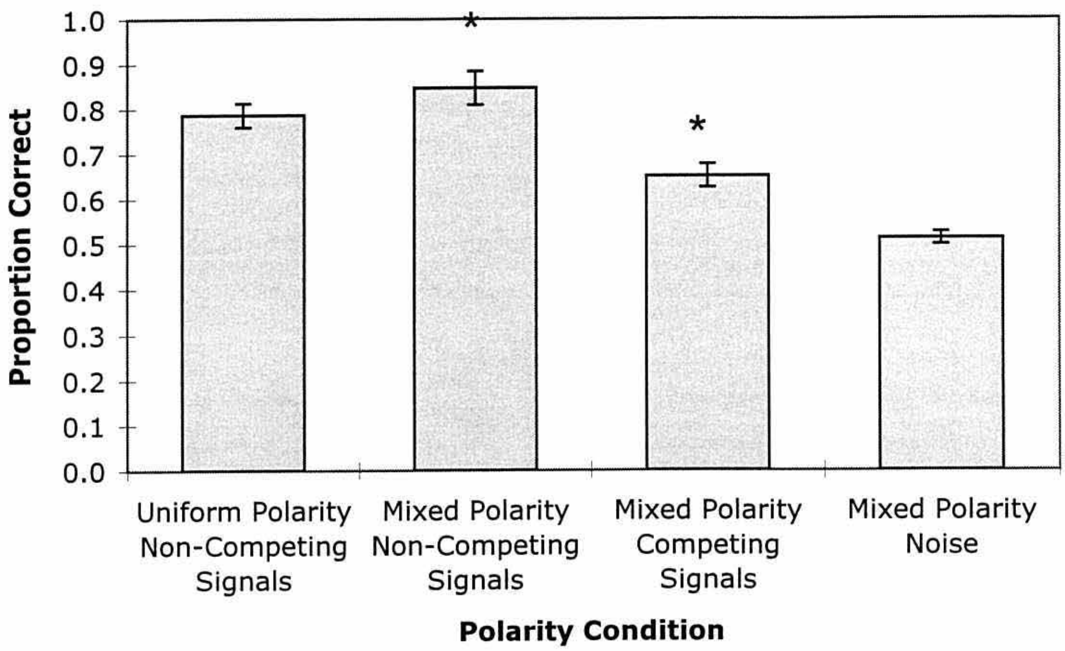

Figure 34. Group mean proportion correct scores for the Red/Green Competing signal experiment. Asterisks indicate expected proportion correct scores assuming observers can effectively select on the basis of colour. See text for more details. Vertical lines represent ± 1 s.e.

Experiment 11 - Grouping of Competing Distractors

The better-than-chance performance in Condition 3 (*mixed polarity/competing signal*) clearly indicates that luminance polarity and colour information is indeed available to texture integration processes but the observation that performance is worse than 80% correct shows that this information is inefficiently applied. A possible mechanism known to be erratic, especially with brief stimulus durations, is that of visual binding. *Binding* is used to describe central processes that allow featural information (e.g., colour, shape, size, depth, orientation) from a given region in spatio-temporal space to be associated, thereby yielding coherent object representations (Treisman & Gormican, 1988). Binding errors are also called illusory conjunctions: when an observer is presented with, say, a red X and a green O and reports a green X to be present, this is called an illusory conjunction error (Prinzmetal, 1981). Perhaps in the present context, orientation coherence mechanisms are not "blind" to luminance polarity or colour but rather are inefficient in combining this information with orientation to group signal and noise separately.

If grouping is a factor in texture integration processes then I should be able to examine the ability of the observer to group signal and noise by varying the amount of orientation similarity within those groups. In Condition 3 (*mixed polarity/competing signal*) above, the amount of signal (20% coherence) was exactly matched by the amount of noise in the perpendicular direction (which can be considered a balanced amount of competing signal, or "anti-signal"). As the amount of similarity among the noise lines increases, it may become easier to group together (Treisman, 1982; Egeth, Virzi & Garbart, 1984; Duncan & Humphreys, 1989) and therefore discard. Since the noise is in a competing orientation, however, performance may decrease as the competing signal increases in strength. If this is the case, the effectiveness of selectively attending to a given stimulus polarity can be measured by increasing the amount of competing signal information of the "to-be-ignored" polarity, and examining at what point selection breaks down and performance approaches chance.

To test the effectiveness of selection under these conditions, I presented textures always with 20% coherent signal in one orientation and variable amounts of competing "anti-signal" in the noise distribution. I asked observers to selectively attend to either the black lines or the white lines in separate blocks. Observers were naïve to the relationship between the signal orientation and the "anti-signal" orientation and were told to ignore lines of the other polarity. Here I describe only the "attend black" condition for clarity (substituting "white" for "black" and "black" for "white" below will describe the "attend white" condition).

Methods

Observers

Six naïve observers participated in this experiment (mean age = 23.5 years, SD = 12.03 years).

Texture Stimuli

Observers were presented with six different texture conditions. In all conditions, 50% of the lines were black and 50% were white. Here I describe the "attend black" conditions. These stimuli types are schematically represented in Figure 35.

1. ***Noise*** *(Figure 35, Panel A)* - All line segments were randomly oriented. This provided a baseline measure of chance performance for our observers.

2. ***Easy Judgment*** *(Figure 35, Panel B)* - 80% of the black lines were oriented in the signal orientation, the rest of the black lines and all of the white lines were randomly oriented (noise) resulting in an overall patch coherence of 40%. This was a relatively easy orientation judgment and as well as providing an easy response for the observers, allowed me to demonstrate that each observer could judge the predominant orientation given enough signal information.

In Conditions 3 - 6 (*Competing Signals*), 20% of the black lines were always oriented in the signal orientation and the rest of the black lines were noise. A given percentage of the white lines were oriented *perpendicularly* to the black signal lines.

3. ***0% Competing Signals*** *(Figure 35, Panel C)* - 20% of the black lines were oriented in the signal orientation and the rest of the black lines were noise. The white lines were all randomly oriented (noise).

4. ***20% Competing Signals*** *(Figure 35, Panel D)* - 20% of the black lines were oriented in the signal orientation and the rest of the black lines were noise. 20% of the white lines were oriented perpendicularly and the rest of the white lines were noise. This condition corresponds to Condition 3 in the previous two experiments.

5. ***40% Competing Signals*** *(Figure 35, Panel E)* - 20% of the black lines were oriented in the signal orientation and the rest of the black lines were noise. 40% of the white lines were oriented perpendicularly and the rest of the white lines were noise.

6. ***60% Competing Signals*** *(Figure 35, Panel F)* - 20% of the black lines were oriented in the signal orientation and the rest of the black lines were noise. 60% of the white lines were oriented perpendicularly and the rest of the white lines were noise.

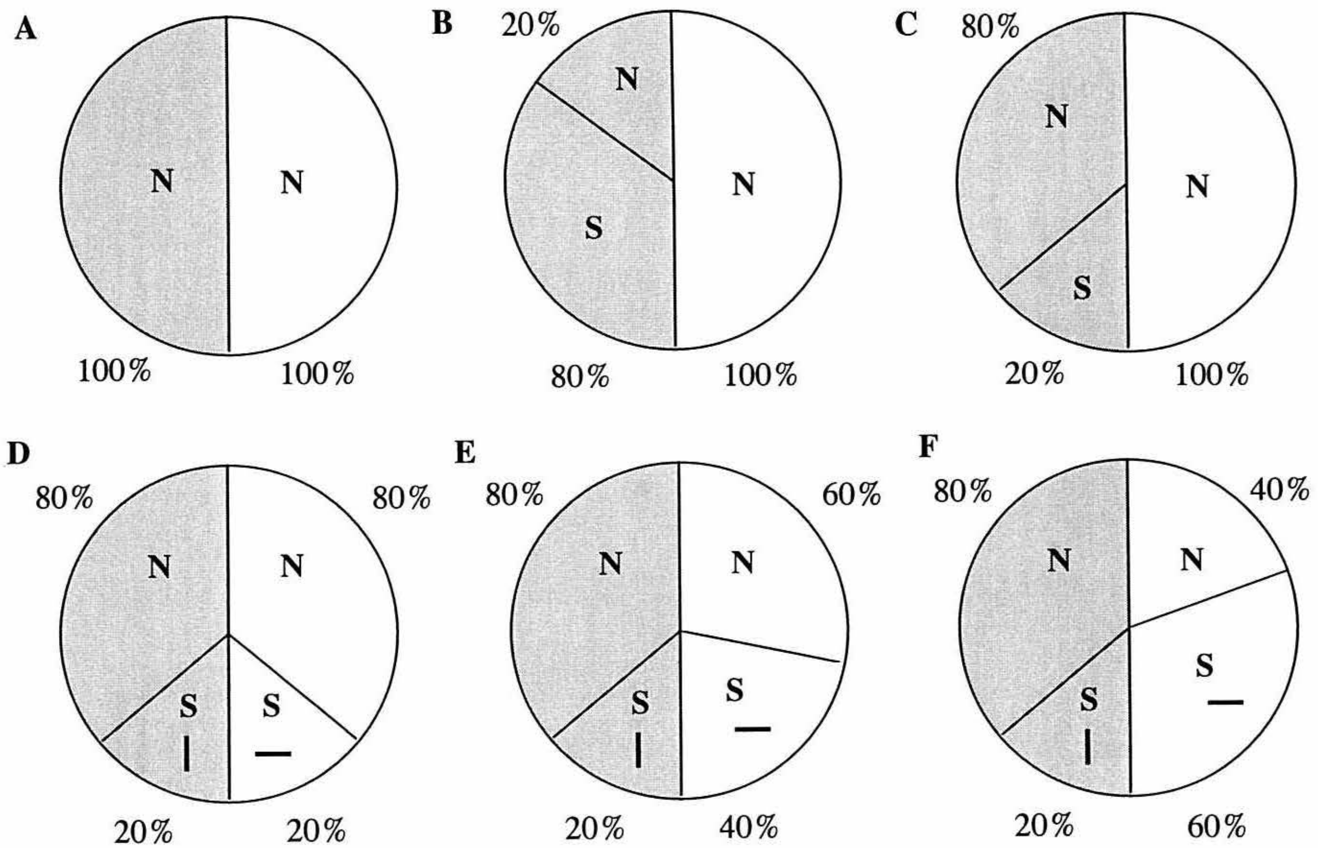

Figure 35. Pie charts depicting the relative signal and noise element distributions for the six conditions in Experiment 11. A. Noise condition. B. Easy Judgment Condition. C. 0% Competing Orientations. D. 20% Competing Orientations. E. 40% Competing Orientations. F. 60% Competing Orientations. See text for detailed descriptions of these conditions.

Procedure

Each of the six texture conditions was repeated 20 times for each of two signal orientations for a total of 240 trials presented in a random order. Observers were asked to judge the predominant orientation (vertical or horizontal) of just the black lines or just the white lines in separate blocks. One block of "Attend Black" and one block of "Attend White" was completed and the order of presentation was counterbalanced between participants.

Participants were also asked to complete two short (3 min) blocks of the *Black & White, full-density* task from Experiment 8. This was done to measure a standard psychometric function for each observer viewing mixed polarity patches. A block consisted of five repetitions of seven coherence values (0, 5, 10, 15, 20, 25, 30%) for two signal orientations (vertical and horizontal) for a total of 70 trials. Observers were asked to judge the predominant orientation of the patch ignoring element polarity. All testing was completed in 40 minutes.

Results

Mean proportion correct scores are shown in Figure 36 for six observers as a function of orientation condition. Since there was no significant difference in the proportion correct scores for "Attend Black" and "Attend White", these scores were averaged resulting in 80 trials for each of the six conditions for each observer.

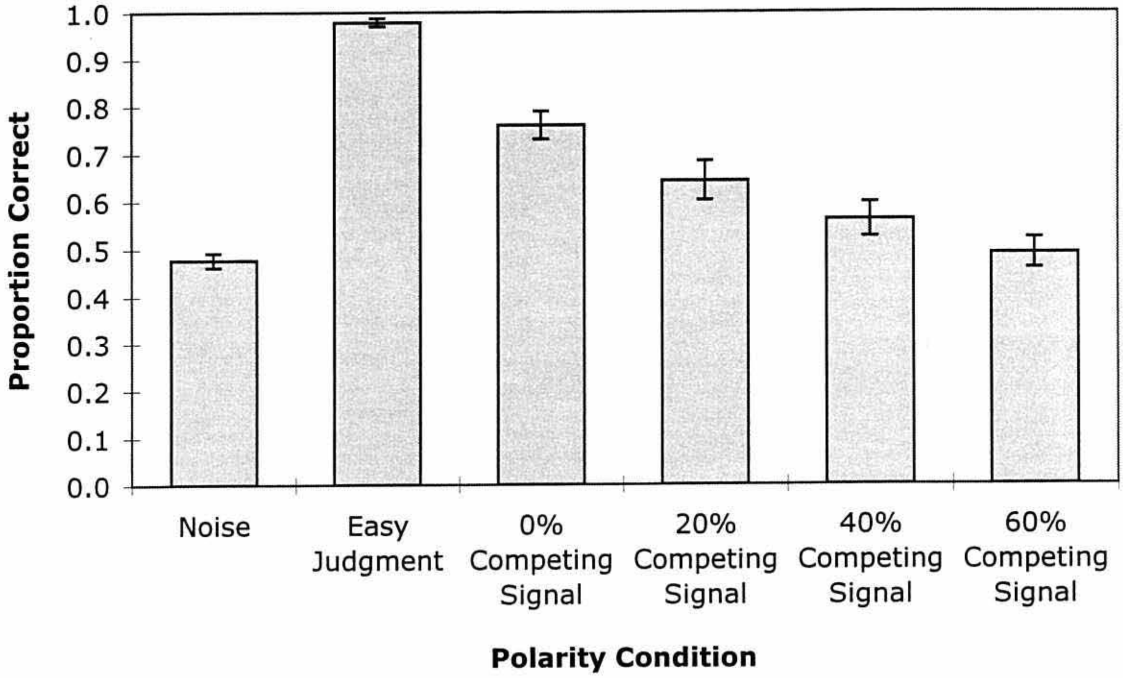

Figure 36. Group mean proportion correct for six texture stimuli conditions from Experiment 11. See text for details of conditions. Vertical lines represent ± 1 s.e.

In Condition 1 (*Noise*), mean proportion correct was at chance (mean = .47, SD = .04) as expected. For Condition 2 (*Easy judgment*), mean proportion correct was .98 (SD = .02) indicating that none of the observers had trouble in determining the most predominant orientation of a noisy texture with mixed polarity.

Conditions 3 - 6 (*Competing signals*) show an interesting trend. In all of these cases, there was a constant amount of signal information (20% coherence) for line segments of the "to-be-attended" polarity. Only the strength of the competing noise orientation changed from 0% to 60%. I calculated the following expected proportion correct scores for given coherence values based on each individual observer's psychometric functions resulting from the baseline (*Black & White full density*) polarity task [5].

[5] I fitted a best fit line using the method of least squares to each observer's psychometric function (r^2 > .71 in all cases) and interpolated the expected proportion correct corresponding to a given coherence value. See Appendix A for an explanation of this technique.

Different proportion correct values could be expected for Conditions 3 - 6 based on a number of assumptions about binding mechanisms. First, if I assume that observers can effectively attend to a subset of elements based on polarity and can judge the orientation of those selected lines (i.e., perfect binding), then the coherence of the "to-be-attended" lines is 20% regardless of the orientations of the noise lines. This would predict a mean proportion correct of .89 (SD = .06) for Conditions 3 - 6 and as can be seen in Figure 36, our observers perform consistently below this level. Obviously, selection is not perfect here indicating that the binding of two feature dimensions, i.e., orientation and polarity, is not totally perfect.

Second, I could assume that featural binding is incomplete, but that observers can detect signal lines of the "to-be-attended" polarity and treat every other line segment as noise (including the noise lines of the "to-be-attended" polarity and all lines of the other polarity). This corresponds to a total coherence of 10% in these stimuli regardless of the orientation of the competing signals and predicts a proportion correct of .65 (SD = .11) for Conditions 3 - 6. In Experiments 9 and 10, I suggested this might be the case on the basis of one coherence point in which the signal and "anti-signal" were matched in strength (i.e., Condition 3 - *Mixed polarity/competing signal*). This prediction may explain the point of matched competition but does not take into account the downward trend of proportion correct in response to increased competing signal strength. If this hypothesis was true, then performance should not change as a function of increasing competing signal strength.

A third possibility is that possibly texture integration mechanisms do not bind featural information at all and that binding takes place at a separate stage (consistent with results from the Glass Pattern polarity experiments of Wilson et al., 1997). Thus, without binding of multiple features, the orientation judgment task must be done on the basis of only one feature, i.e., orientation, rather than the conjunction of two, i.e., polarity and orientation. I therefore interpolated for each observer (using psychometric functions as described above) the expected proportion correct for the total coherence of the texture patch disregarding element polarity to emulate a "no selective attention" case. Figure 37 shows mean expected (circles) and obtained (squares) proportion correct values for six observers as a function of percent competing signal. Expected proportion correct scores below 50% (chance) indicate conditions where the "anti-signal" strength exceeded the signal strength and the observer would be expected to report the "anti-signal" orientation (for example, 60% "anti-signal" combined with 20% signal results in a total coherence of 40% "anti-signal"). I have assumed here, for comparison purposes, that a signal and an "anti-signal" cancel each other completely and in a linear fashion. As can be seen in Figure 37, observed proportion correct scores (squares) are much higher than predicted (squares). The effectiveness of selective attention in this situation can be defined as the difference in the two functions. These results indicate that polarity information is available to texture integration processes but such processes require selective attention to bind features together for discrimination purposes.

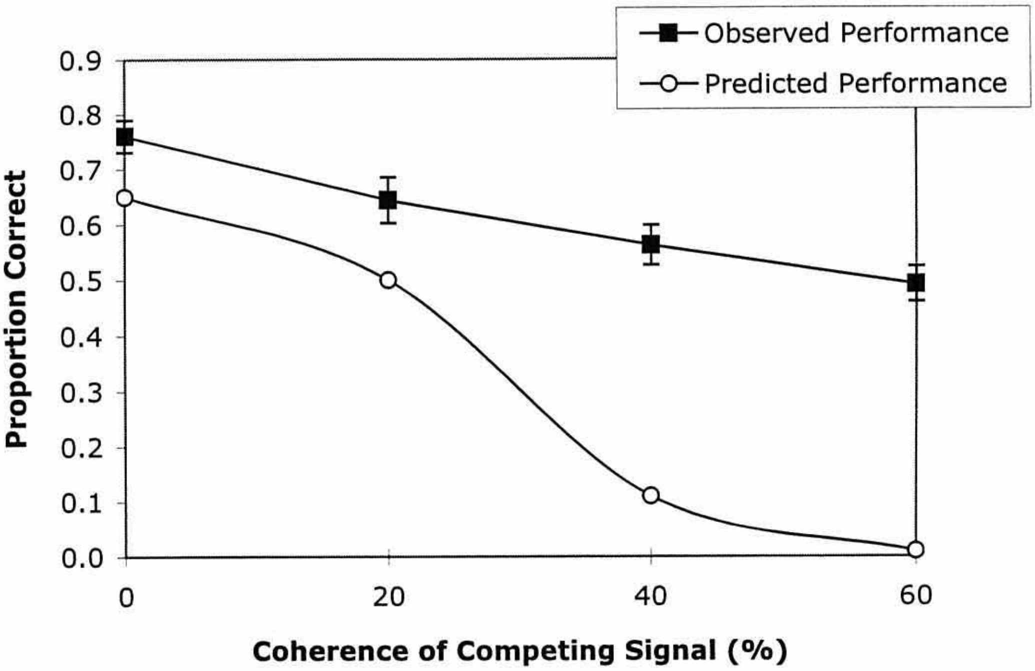

Figure 37. Group mean observed and expected proportion correct scores as a function of competing distractor strength for Experiment 11. Expected scores were calculated by interpolating proportion correct from each observer's psychometric function given each coherence value. Values falling below 50% represent responses in the competing (noise) orientation. Vertical lines represent ± 1 s.e.

Discussion

I demonstrated in Experiment 8 that texture orientation coherence thresholds were not affected by mixed polarity textures (when observers were asked to ignore polarity) or textures of different densities (at least within the range tested). I went on to examine in Experiment 9 whether texture integration processes have access to polarity information even though they are thought to occur after full wave rectification processes that effectively discard this information. Close examination of the data revealed that polarity information is indeed available but is used in an imperfect fashion. Experiment 10 showed similar results for red and green stimuli indicating that it is not polarity *per se* that is inefficiently coded, but rather that features (colour, polarity, orientation) of objects are inefficiently bound together with the stimulus durations used here and result in imperfect selection. Experiment 11 explored this binding process further and showed that polarity and orientation information is indeed accessible to texture integration processes but requires focused selective attention to bind these features together.

The idea that an object's features are coded separately at early stages of vision is widely held in many theories of early visual processing (Duncan & Humphreys, 1989; Humphreys & Muller, 1993; Prinzmetal, 1981; Treisman & Gelade, 1980; Treisman & Gormican, 1988; Wolfe, Cave & Franzel, 1989). Later in processing, these features are bound back together with the help of selective attention, possibly in the inferotemporal (IT) area of the brain (Treisman & Gormican, 1988; Wilson et al., 1997). With short stimulus durations, or with taxing selection tasks, the binding is not totally completed and the result is binding errors or illusory conjunctions. This may be what is happening in Experiment 11. Possibly, featural binding is not complete and since they cannot reliably select the target subset of lines, observers report the competing orientation. The effectiveness of selective attention in this case is remarkable, however, as it requires at least 60% competing signal in the presence of 20% real signal to reduce performance to chance levels. This indicates that selective attention can filter out large amounts of competing information.

In the experiments of Wilson and colleagues (Wilson et al., 1997), their observers may have been successfully coding both the spatial relationships of the partner dots and the polarity of the dots in the mixed polarity Glass patterns, but may have been unable to bind together the two features. Their task was more complex than the one used here due to the perceptual calculations required to determine the spatial structure of the Glass pattern itself and the added task of determining the polarity of the signal dots with an even shorter stimulus duration (167 ms) than the one used here. The fact that their observers could not discriminate efficiently between arrays with white signal/black noise and the reversed pattern of polarity is completely consistent with the idea that feature binding is incomplete with short stimulus durations and taxing selection tasks. Interestingly, observers' performance actually appears to be just *above* chance in Wilson and colleagues' experiment, though they do not report any statistics on this result as to whether it is significantly different than chance. If observers were above chance, this supports the hypothesis that the present results are due to imperfect feature binding processes and that selective attention assists in this function.

Featural similarity among stimuli and perceptual grouping may also have affected observers' performance in the present selection task. Duncan and Humphreys (1989) suggest that the similarity of targets and non-targets is what dictates the efficiency of selection and search. As target-nontarget (N-T) similarity increases, and nontarget-nontarget (N-N) similarity decreases, the effectiveness of search decreases. In the case of the heterogeneous textures used here, targets were the horizontal and vertical lines of a given polarity and all other lines (of both polarities) were distractors. For example, if observers were to select only the black lines and to determine which of the horizontal or vertical lines were more predominant, distractors could be all lines oriented in non-cardinal orientations plus those oriented cardinally of the opposite polarity. Similarity between targets and nontargets in this situation is high as is similarity among nontargets (i.e., they differ only by orientation and polarity but have the same form). This situation predicts difficulties with efficient selection.

Grouping of similar distractors should have assisted observers in the competing signal conditions of Experiment 11. As similarity between distractors increases, as would happen with increasing distractor coherence, the distractors should group more easily and be easily discarded. It has been proposed that observers are able to quickly group areas of similar distractors together in homogeneous groups and search in parallel within those groups thereby increasing search efficiency (Duncan & Humphreys, 1989; Egeth, Virzi & Garbart, 1984; Humphreys, 1993; Treisman & Gelade, 1980; Treisman, 1982). As the homogeneity of distractors increases, therefore, performance should improve. In Experiment 11, performance actually deteriorated as the coherence of the competing distractors increased. This could be due to the fact that the orientation of the distractors was in direct competition with, i.e., on the same continuum as, the orientation of the signals. Distractors could not be discarded as easily as they could if a separable feature from the signals defined them.

While the discussion presented here has focused more on featural binding and grouping processes, the results are nonetheless totally consistent with the ideas of the texture perception literature. As an extension of the present series of experiments, it would be interesting to test a range of stimulus durations to examine where perceptual selection does in fact become complete. I would expect observers to show no effect of competing orientations at the point where featural binding is totally completed. This may give a more direct measure (as opposed to Reaction Time in classic visual search experiments) of the processing time for featural binding. It would also be interesting to examine different featural dimensions such as depth and polarity to see if they show a similar pattern of processing and combination.

The texture coherence paradigm is well suited for examining featural binding and grouping processes in this manner. First, target and distractor similarity can be varied along a continuum (i.e., orientation can be changed continuously over 180°) rather than being constrained to discrete steps. Second, the strength of the distractor stimuli can be varied in a continuous fashion by changing distractor coherence without concurrently changing the overall appearance (or statistics) of the image. Third, the effects on selective attention can be quantified in "stimulus units", i.e., the same stimulus units (percent coherence) can be used for quantifying the effects as are used to define the stimuli themselves. For example, in classic visual search experiments the stimuli are commonly measured in numbers of distractors and targets and the effects are measured in reaction time to make a response. Using the coherence method, both the strength of the stimulus and the size of the attentional effect can be defined in percent coherence, see Chapter 6 - Texture Contrast Effect for an example of this. Texture coherence should be useful for future attention research, as well as perception.

Chapter 6 - Stimulus Interactions over Time: The Orientation Contrast Effect

Abstract

After prolonged viewing of a row of obliquely oriented lines (inducer), observers commonly report that a row of vertically oriented lines appear tilted in the opposite direction to the previously viewed inducer. This is known as the tilt aftereffect (Gibson & Radner, 1937) and is thought to reflect adaptation or fatigue processes of orientation selective visual neurons (Campbell & Maffei, 1971; Coltheart, 1971). I have previously shown that orientation sensitivity can be quantified using a texture coherence paradigm similar to that used to study sensitivity to global motion in random dot kinematograms. Here, I asked whether pre-exposure to 100% coherently oriented textures would cause alterations to orientation sensitivity consistent with the tilt aftereffect. Observers were presented with a 200 ms texture patch (inducer) composed of short line segments oriented horizontally or vertically. After a variable inter-stimulus-interval (ISI), a similar but partially-coherent texture (test) was presented for 200 ms. The percentage of coherently oriented segments in the test patch was varied from trial to trial and observers reported the global orientation of the inducer and whether the test texture was coherent or random (i.e., signal present vs. absent). Mean orientation coherence thresholds were significantly elevated when test orientations matched that of the inducer and were reduced when test and inducer were mismatched. This orientation contrast effect was not present at short ISIs (100 ms) and increased in magnitude for moderate ISIs (200 - 700 ms). These data demonstrate that viewing briefly presented homogeneously oriented textures dramatically changes subsequent global orientation sensitivity to similar, partially coherent textures depending on the orientation relationship between them. The result that the orientation contrast effect is not immediately present at shorter ISIs and that it develops over time suggests that it has an attentional component.

Successive Stimulus Interactions

Visual stimuli rarely occur in isolation in real world visual scenes and tend to interact with one another perceptually. Successive stimulus interaction effects commonly involve observers viewing an initial inducing stimulus that affects the perception of a second, test stimulus causing visual illusions, aftereffects or changes in sensitivity. An inducing stimulus (also called an adapting or priming stimulus) can cause an ambiguous stimulus to look either more similar to the inducer (assimilation or capture effects), or more different (contrast effects). I have chosen to call the initial texture stimulus an "inducer" here because the terms "adapting stimulus" and "prime" have many assumptions attached to them in the visual adaptation and attention literatures respectively that may not hold to be true in this context.

There are numerous examples of successive perceptual interactions including the tilt aftereffect (Campbell & Maffei, 1971; Gibson, 1937; Gibson & Radner, 1937), the size aftereffect (Blakemore & Sutton, 1969), and the motion aftereffect. Other examples are attentional in nature and include repetition blindness (Kanwisher, 1987), the attentional blink (Raymond, et al., 1992), negative priming (Tipper, 1985) and inhibition of return (Posner & Cohen, 1984). All of these temporal interactions seem to be caused by the initial coding of stimulus features affecting subsequent coding of other stimuli.

The tilt aftereffect, for example, occurs when observers are exposed to a row of parallel tilted lines for an extended period and then are asked to judge the orientation of a row of vertical lines presented immediately afterward. Observers commonly perceive the line orientation as being tilted in the opposite direction to that of the initial stimulus (Campbell & Maffei, 1971). Similarly with the size aftereffect, when an observer has viewed a row of thick lines and then judges a row of medium thickness lines, he or she commonly perceives the subsequently presented lines to be thinner than they really are (Blakemore & Sutton, 1969). These effects are a type of stimulus contrast in that the features of the subsequent test stimulus contrast with those of the initial inducing stimulus.

Stimulus interactions using motion stimuli have been shown through studies on motion adaptation (e.g. Raymond, 1993; Raymond & Braddick, 1996), motion aftereffect (e.g. Barlow & Hill, 1963), motion priming (e.g. Raymond, O'Donnell & Tipper, 1998) and motion contrast (e.g. Raymond & Isaac, 1998; Raymond & O'Donnell, 1996). Typically, a motion stimulus is presented for a given interval and is followed by a test stimulus that probes subsequent illusory motion or motion sensitivity losses and gains. After prior exposure to a motion stimulus, results typically show reduced sensitivity for a subsequent stimulus moving in the *same* direction as the initial stimulus and enhanced motion sensitivity for the direction *opposite* to the initial stimulus (Raymond & Isaak, 1998; Raymond & O'Donnell, 1996).

Evidence in the attention literature for temporal interactions in the coding of features also demonstrates changes in sensitivity to subsequently presented stimuli. Repetition blindness for example, is demonstrated using rapid serial visual presentation (RSVP) of words where observers fail to detect the presence of a word when it is presented a short time after the same word, even if other words were interleaved between the two presentations (Kanwisher, 1987). The repetition blindness effect is still present when the two words differ in case (i.e., capitals vs. lower case letters) or even if the two words are presented in a grammatically correct sentence (i.e., When she spilled the *ink*, there was *ink* all over) and the omission creates a grammatically incorrect sentence (i.e., When she spilled the *ink*, there was all over)(Kanwisher, 1987). The coding of the first word affects the coding of the second word.

Since the coding of stimuli in many visual domains (i.e., tilt, size, motion, and attention) all show temporal interactions, it follows that global texture orientation perception may also show a temporal effect. Coltheart (1971) proposed that the tilt aftereffect is induced by adaptation of orientation-specific analyzers in the visual system. Hubel and Weisel (1968) originally found that orientation-selective units are maximally sensitive to a given orientation and will fire at a slower rate as the similarity in orientation between the presented stimulus and their optimal stimulus decreases. Analyzers will also decrease their firing rate in response to protracted stimulation. After extended stimulation, the orientation analyzer enters a state of reduced sensitivity from which it recovers over time. Coltheart proposed that the perceived orientation of a line is defined by the average of the preferred orientations of all the units that respond to it, each weighted by the extent to which the unit responds above baseline. Prolonged viewing of an oriented stimulus will fatigue a given group of analyzers and will bias the weighted average response to subsequent oriented stimuli. This causes negative tilt aftereffects.

Blakemore and colleagues (Blakemore, Carpenter & Georgeson, 1970; 1971), however, proposed that the tilt aftereffect stems from lateral inhibition between oriented analyzers that persists in time. Lateral inhibition results from the center-surround organization of orientation selective analyzers that have been found in the primate visual cortex (Hubel & Weisel, 1968). An orientation-selective analyzer is excited when an appropriately oriented stimulus falls within the central area of its receptive field. If such a stimulus falls on the outer area of that receptive field, the cell's activity is suppressed or inhibited. Center-surround organization works to sharpen the orientation selectivity of such cells by strictly limiting the region of receptors that are activated. When two oriented stimuli are presented in close proximity and the receptive fields of numerous cells overlap, the inhibitory and excitatory activity is summed and the perception of the orientation of one line will be displaced away from the orientation of another, compared to when the line is presented in isolation. Such lateral inhibition may persist in time and cause successively presented line segments to appear to have a greater difference in orientation than they actually do, resulting in the classic tilt aftereffect.

Such low-level perceptual hypotheses have tended to dominate tilt aftereffect research (Carney, 1982; Grabowska, 1987; Harris & Calvert, 1988; Magnussen &

Kurtenbach, 1980; Magnussen & Johnsen, 1986). While low-level explanations may account for many aspects of visual aftereffects, such explanations rarely address the role of higher level processing and may be limited by the simple and static stimuli that are often used in testing them. Since our visual worlds are filled with much more complex stimuli such as textures, I asked here whether more complex and transient oriented stimuli can produce analogous texture orientation interaction effects as compared to those produced with simple stimuli.

To reflect this, the present experiment differs from classic tilt aftereffect experiments in three main ways. First, I used more complex heterogeneously oriented textures to induce temporal interaction effects rather than single line segments. Second, tilt aftereffects are commonly induced using long exposure durations to both the inducing stimulus (typically seconds or minutes) and the test stimulus (presented during an adjustment procedure or until the observer responds). I used relatively brief inducer (200 ms) and test durations (200 - 250 ms). Moderate tilt aftereffects have previously been induced in line segment orientation adjustments with very brief exposures to oriented grating pattern inducers (5 ms; Sekuler, & Littlejohn, 1974) and test stimuli (10 to 600 ms; Wolfe, 1984). Motion aftereffects have also been induced using brief prime exposures (Raymond & Isaak, 1998). Thus, I should be able to induce a change in sensitivity to test textures by briefly presenting an oriented inducer texture at the same spatial location shortly before the test.

A third important difference is that in classic experiments investigating temporal interactions, observers are told to passively view the inducing stimulus. No perceptual task is required in response to the inducer and no attempt is made to control the allocation of attention. Here, observers were required to make an orientation judgment about the inducer and required a criterion performance on the inducer task of at least 85% correct for inclusion in the experiment. Further, only the trials for which the inducer response was correctly reported were included in the analysis. Through this manipulation the attentional state of the observers is more closely monitored and controlled.

Observer sensitivity, after the presentation of the inducer stimulus, was measured with the thresholding technique developed in Chapter 2. An analogous thresholding technique has previously been used successfully to measure changes in sensitivity due to the motion aftereffect (Paradiso, Shimojo & Nakayama, 1989; Raymond, 1993; Raymond & Isaak, 1998; Raymond & O'Donnell, 1996). The basic rationale for using this paradigm is that if there are changes in perception that produce the experience of an aftereffect, then a sensitive way to quantify the aftereffect itself is by measuring the changes in subsequent sensitivity that give rise to the aftereffect. Using the thresholding paradigm, I was able to do this. Here, the inducing stimulus was a 100% coherent texture patch oriented either horizontally or vertically. This inducer was followed by a blank inter-stimulus-interval and then a test texture patch of variable coherence was presented oriented horizontally or vertically. Observers were required to report both the orientation of the inducer patch and the test patch.

I expected orientation sensitivity to a partially coherent test stimulus to be affected by the previous presentation of a coherently oriented inducer patch resulting in a texture contrast effect consistent with the tilt, size and motion aftereffects. Specifically, for test stimuli that matched the previous inducer in orientation, thresholds should have been elevated from baseline. For test stimuli that mismatched the inducer, thresholds should have been lowered from baseline.

Experiment 12 - Successive Texture Integration: An Orientation Contrast Effect

Methods

Observers

Eight healthy naive participants took part in this experiment (mean age = 22.75 years, SD = 3.14 years). Four observers took part in each of two test duration conditions.

Texture Stimuli

Inducer stimuli were 100% coherent texture patches with horizontal or vertical signal orientations. Test stimuli had horizontal or vertical signal orientations and coherence was varied from 0% to 30% in 5% steps.

Procedure

A trial (illustrated in Figure 38) was initiated by a key press and was composed of a fixation cross presented for 1500 ms, an inducer patch for 200 ms, a blank inter-stimulus-interval for 200 ms, a test patch for 200 or 250 ms followed immediately by a pattern mask for 100 ms. The two test durations (i.e., 200 and 250 ms) were used because it was thought that the shorter test duration may have been too difficult a task for the observers. (Results showed that there was no difference in the proportion correct scores for these two conditions and thus results were collapsed over test duration).

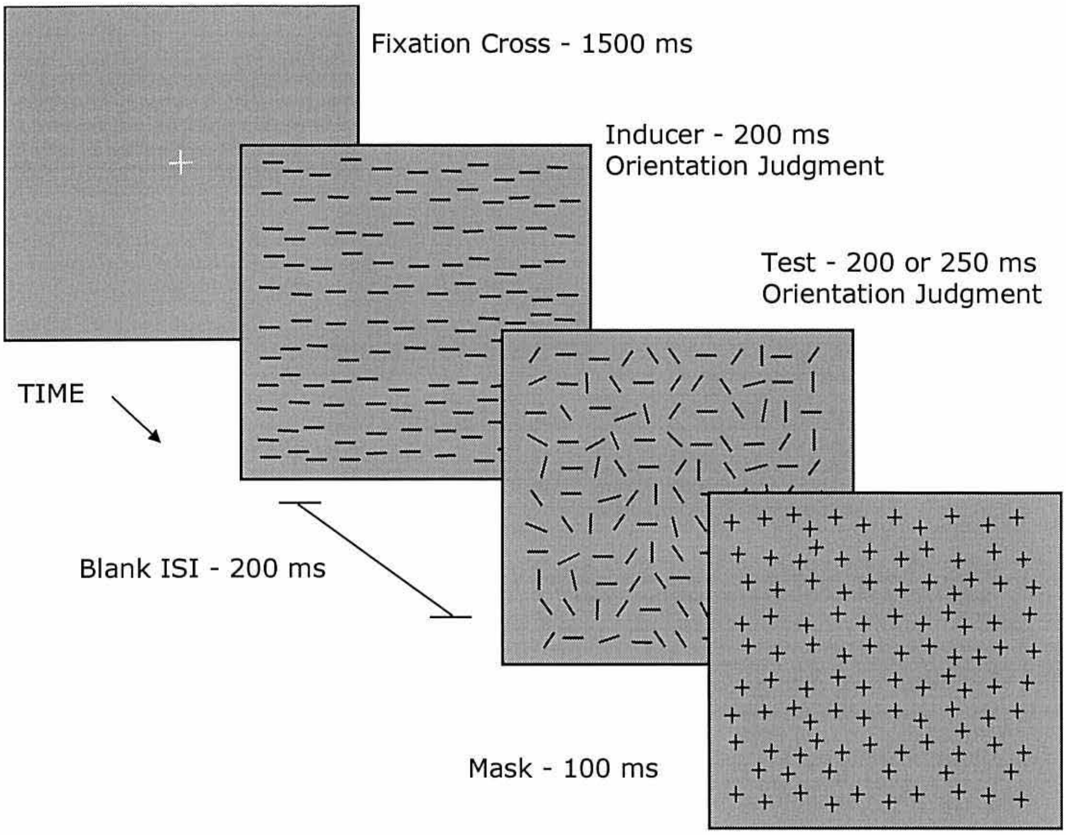

Figure 38. An illustration of a trial using the texture contrast paradigm. Observers were asked to determine the orientation of both the inducer and the test textures. Shown is a "matched" trial in which the orientation of the inducer (horizontal) matches the orientation of the test (horizontal, 30% coherence).

Observers were asked to judge the global orientation of both the inducer and test patches and indicate these judgments with two key presses at the end of the trial. Each of the test orientations (horizontal and vertical) was preceded by each inducer orientation (horizontal and vertical) resulting in four orientation conditions. For each of the orientation conditions, seven test coherence values were repeated five times each resulting in 140 trials per block. Observers completed three blocks of one test patch duration in one session of testing that was completed in 30 minutes. No feedback was given after an initial practice phase of 30 trials in which the observers were told the orientation of both texture patches.

Results & Discussion

Only those trials in which the inducer orientation was correctly reported were included in the analysis to ensure that attention was indeed directed at the inducer. Observers were, on average, 98% correct (SD = 1.8%) in reporting the inducer orientation and errors were not systematic with respect to orientation condition. All observers achieved at least 95% correct on the inducer orientation judgment task.

I averaged proportion correct scores for those trials in which the inducer orientation was the same as the test and called this the *matched* orientation condition. The *mismatched* orientation condition was the average proportion correct scores for trials in which the inducer mismatched the test. This was done for each coherence value and resulted in the average of 30 data points representing each point on the psychometric function for each observer. I then plotted mean psychometric functions for the 200 and 250 ms test durations and found no significant difference. Since I previously demonstrated that sensitivity does not change for durations longer than 180-200 ms (Chapter 4), I therefore averaged over test duration. Mean psychometric functions for match and mismatch orientation conditions are shown in Figure 39.

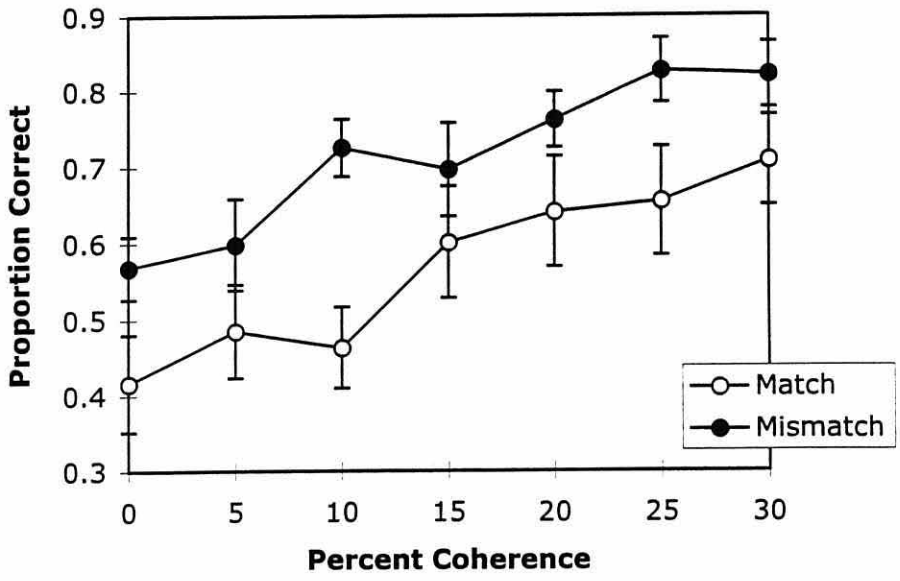

Figure 39. Group mean psychometric functions plotted for match and mismatch orientation conditions. Error bars represent ± 1 s.e. of the mean.

As is shown in Figure 39, the psychometric function for the match condition is shifted to higher coherence values relative to the mismatch condition. I performed an analysis of variance (ANOVA) using orientation condition and coherence as within subject factors. Results showed a significant main effect of orientation condition, $F(1,98) = 8.92$, $p < .05$, and coherence, $F(6, 98) = 25.96$, $p < .001$, and a non-significant interaction. Thus, observers were significantly more sensitive to a texture when it was *mismatched* in orientation to the preceding inducer than when it was *matched*. This demonstrates a texture orientation contrast effect analogous to the contrast effects shown in tilt, size and motion.

Another characteristic of contrast effects is the presence of illusions in neutral test stimuli induced by the preceding stimulus. In the motion aftereffect, motion stimuli with no predominant signal direction (i.e., random motion) appear to stream in the opposite direction than the inducer. Similarly, in the tilt aftereffect, "neutral" vertical test lines often appear tilted away from previously presented inducer lines (Gibson & Radner, 1937). In this experiment, there was the possibility that an illusion of coherent orientation would be created in an otherwise neutral or random test stimulus, i.e., 0% coherent patches. Interestingly, for the 0% test coherence condition, in which an inducer was followed by a test patch with no predominant signal, the functions for mismatch and match do not converge on chance performance. Since the computer program assigned a nominal signal orientation of either horizontal or vertical to all 0% coherence stimuli, even though there was no signal physically present in the textures, the orientation that was reported with a neutral stimulus could be examined.

If there was no illusion of coherent orientation and observers were responding at chance levels, proportion correct should converge on 50% for the match and mismatch conditions. However, if there was a contrast effect that created the illusion of coherence in the opposite orientation, then observers should report more mismatches than matches. The results show that they do just that. Using a matching to sample t-test [6], the results show that the mean proportion correct for mismatch is significantly higher than chance, $t(7) = 2.34$, $p < .05$, and mean proportion correct for match is significantly lower than chance, $t(7) = -1.86$, $p < .05$. Both results indicate that observers experienced an illusion of coherence induced by the coherent inducing stimulus.

While these results suggest an illusion of coherence in neutral textures, there is the possibility that this pattern of results may be caused merely by an observer bias to report mismatches. Similarly, the apparent difference between match and mismatch at mid-range levels of coherence may be due to observer response bias rather than a real perceptual aftereffect. Unfortunately, this experiment cannot dissociate these possibilities because the responses that were required of the observer for inducer and test were not independent of one another. For example, an observer may have demonstrated a bias for rarely reporting matches (i.e., a bias to

[6] A t-test was used here because I was testing whether the *mean* of the sampled distribution was different than 50%. A Chi-square test would have been more appropriate had I been addressing the *variance* of the sample compared to a standard. The use of a t-test was more appropriate in this case.

always press two different keys). The next experiment addresses more closely the issue of response bias.

Experiment 13 - Texture Orientation Contrast with Orthogonal Response Continua

The previous experiment showed that observers had differential sensitivity to successively presented texture patches depending on the orientation relationship between them. There were some limitations to the design of this experiment, however, and they are addressed in Experiment 13.

Orthogonal Response Continua

Since observers were asked to judge the *orientation* of both textures, the apparently enhanced sensitivity to mismatched patches may have been due to a systematic response bias to report mismatches. In this experiment, orthogonal response continua were used for inducer and test to reduce the chance of response bias. Instead of requiring an orientation judgment for the test patch, I asked observers to judge whether the patch was coherently (either vertical or horizontal) or randomly oriented. This two alternative forced choice (2AFC) judgment is analogous to a signal present vs. signal absent judgment and still allows psychometric functions to be plotted but removes explicit orientation judgments from the sensitivity measurement. By using orthogonal response continua, observers should be less biased in their responses if this indeed was a component of the previous results.

In addition, the coherent vs. random (or signal present vs. absent) judgment allows the implementation of the Theory of Signal Detection (Green & Swets, 1966) that takes into account observer criterion shifts and biases. d' values can be calculated from observer's false alarm rates and proportion correct scores which allow the discrimination between shifts in sensitivity and shifts in criterion (see Appendix B for a more detailed discussion of Signal Detection Theory and d' calculations). In this experiment, observers were therefore asked to report the orientation of the inducing patch and whether the test patch was coherent or random.

A Question of Baseline

A second limitation of Experiment 12 was that there was no effective baseline condition with which to determine whether the changes in performance were due to increased sensitivity in the mismatched condition or decreased sensitivity in the matched condition. A "neutral" blank square inducer condition was added in this experiment to measure a baseline threshold for the dual task. This condition was thought to be neutral because it still produced a transient visual event before the test texture, as did the dual texture cases, but did not contain texture information. In addition, all horizontal and vertical orientation information was equal in the square outline figure. Observers were still required to respond with a key press to this blank inducer thereby minimizing the differences between trial types.

Orientation Contrast

If texture contrast effects result from differences in textural orientation, how large do these orientation differences between inducer and test have to be? Are differences smaller than 90° sufficient to produce a tilt contrast effect? In the last experiment, only two conditions of orientation difference were presented; no difference at all (as in the match condition) and the maximum orientation difference possible (90°, as in the mismatch condition). To examine the orientation difference required between inducer and test to cause a change in test sensitivity, obliquely oriented texture inducers were included in this experiment. These textures were oriented 45° clockwise from vertical or 45° counterclockwise of vertical and were effectively "half way" between matched (0° difference in orientation) and mismatched (90° difference in orientation) conditions. If a texture contrast effect requires only a small change in orientation to change subsequent texture sensitivity maximally, then sensitivity to the test after viewing oblique texture inducers should show the same pattern as the mismatched textures. If however, there is a gradient of sensitivity changes corresponding to a gradient of orientation differences in successively presented textures, then the sensitivity function for oblique inducers should lie between the functions for matched and mismatched textures.

Methods

Observers

Three male and four female healthy naive participants took part in this experiment (mean age = 25.71 years, SD = 7.30 years).

Texture Stimuli

Inducing stimuli were composed of 100% coherent texture patches oriented vertically, horizontally, 45° clockwise, or 45° counterclockwise from vertical. On 20% of trials, the inducer was a square black outline the same size as the texture patch and filled with the same blank grey as the background field. The five inducer conditions are illustrated in the top panel of Figure 40. Test texture stimuli were varied in coherence from 0% to 40% in 10% steps.

Procedure

Training phase: Observers completed a training phase in which they judged whether a series of horizontal or vertical texture patches of 0% or 40% coherence were random or coherent. Practice patches were presented for 200 ms and were matched in every other way to test patches in the experimental phase. Responses were indicated by a key press and feedback was given in the form of a beep for either a false alarm (responding "coherent" to a 0% coherence patch) or a miss (responding "random" to a 40% coherence patch). A block consisted of 32 trials (about two minutes) in which the two coherence values in each of the two orientations were repeated eight times. Observers completed at least 2 blocks of practice to reach a criterion of no more than 2 false alarms and 2 misses. All observers reached this criterion within 3 practice blocks.

Test phase: After criterion had been reached for the coherent vs. random judgment, observers completed 30 practice trials of the experimental dual task in which feedback of the orientation (in the case of the inducer) and coherence information (in the case of the test) was provided verbally. A trial (illustrated in the bottom panel of Figure 40) was initiated by a key press and consisted of the presentation of a fixation cross for 1500 ms, an inducer patch for 200 ms, a blank inter-stimulus-interval for 200 ms, followed by a variable coherence test patch for 200 ms. None of the stimuli were masked. Observers were asked to report the orientation of the inducing patch (vertical, horizontal, 45° clockwise or 45° counterclockwise) and to report whether the test patch was coherent or random. Observers indicated their responses by pressing two keys on the computer keyboard at the end of each trial. In the case of a blank square inducer, observers pressed any of the four keys used for the oriented inducers. No feedback was given after the initial 30 practice trials.

The two test orientations (vertical and horizontal) were presented following each of the five inducer conditions (vertical, horizontal, 45° clockwise, 45° counterclockwise and blank). For each of these ten orientation conditions, five coherence values were each repeated five times resulting in a total of 250 trials per block presented in a pseudo-random order. Testing was completed in 45 minutes in one session.

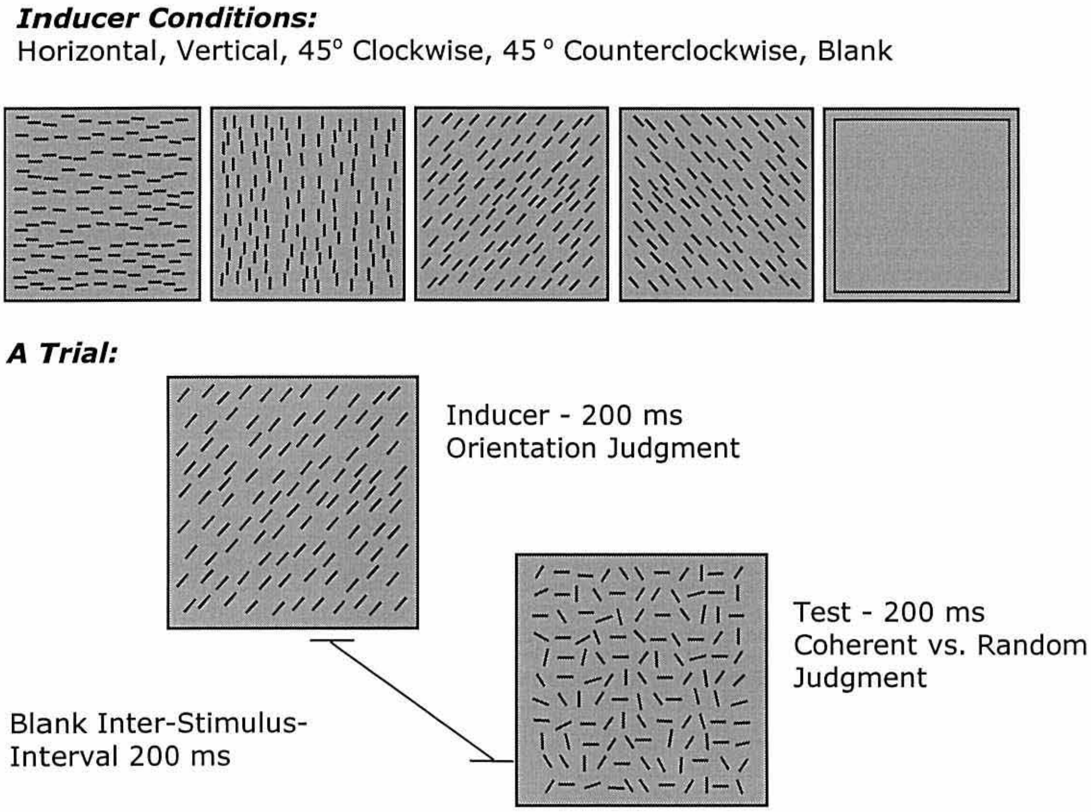

Figure 40. An illustration of the orientation contrast paradigm using orthogonal response continua. Five inducer conditions and an example trial are illustrated. Observers were asked to determine the orientation of the inducer and to determine whether the test was coherent or random.

Results

One observer was excluded from this analysis for not reaching criterion [i.e., poor performance (less than 85% correct) on the inducer orientation task]. All other observers achieved greater than 94% correct and only those trials in which the inducer orientation was correctly reported were included in the analysis. Mean inducer performance was 97.3% (SD = 2%) and errors were not systematic with respect to the orientation condition.

Group mean proportion "coherent" responses were plotted as a function of percent coherence to show psychometric functions as can be seen in Figure 41. As in the previous experiment, proportion "coherent" responses were averaged into each orientation condition, i.e., match, mismatch, oblique and blank. This resulted in the average of 10 data points representing each point on the psychometric function for each observer, except for the oblique condition in which responses for the two oblique inducers were collapsed giving the average of 20 data points.

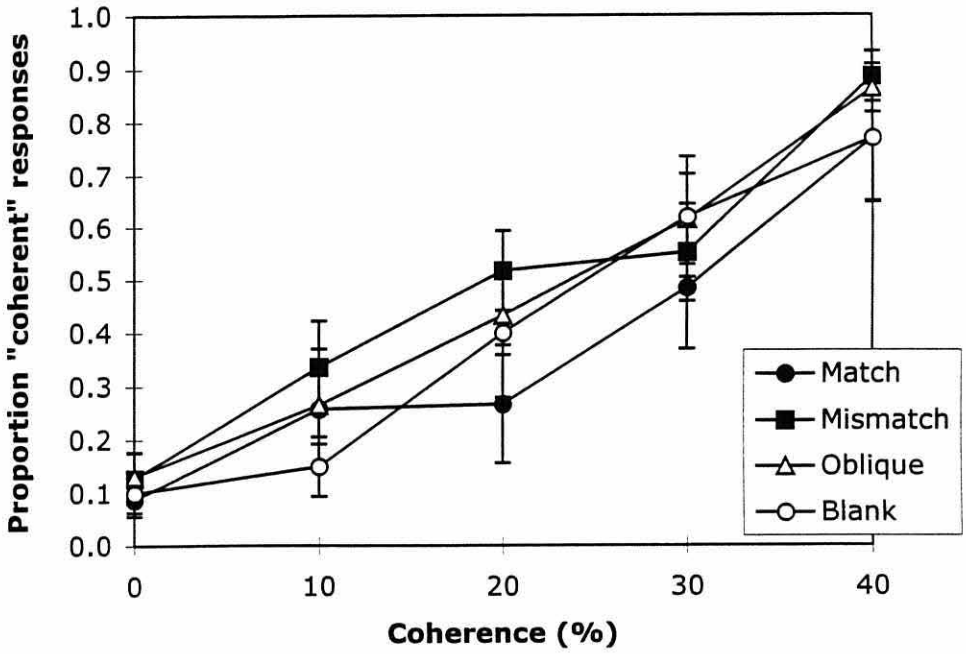

Figure 41. Proportion "Coherent" responses plotted as a function of percent coherence. The resulting psychometric functions show differences in sensitivity for each orientation condition. Error bars represent ± 1 s.e.e.

The false alarm rate for 0% coherence (averaged over orientation condition) was not greater than .3 for any observer (mean false alarm rate = .11, SD = .096). A repeated measures ANOVA showed that there was no difference in false alarm rates among orientation conditions. Mean false alarm rates are shown in Figure 42 for the four orientation conditions. This result is contrary to the previous finding of a coherence illusion in 0% coherent textures. Subsequent experiments will further address this issue.

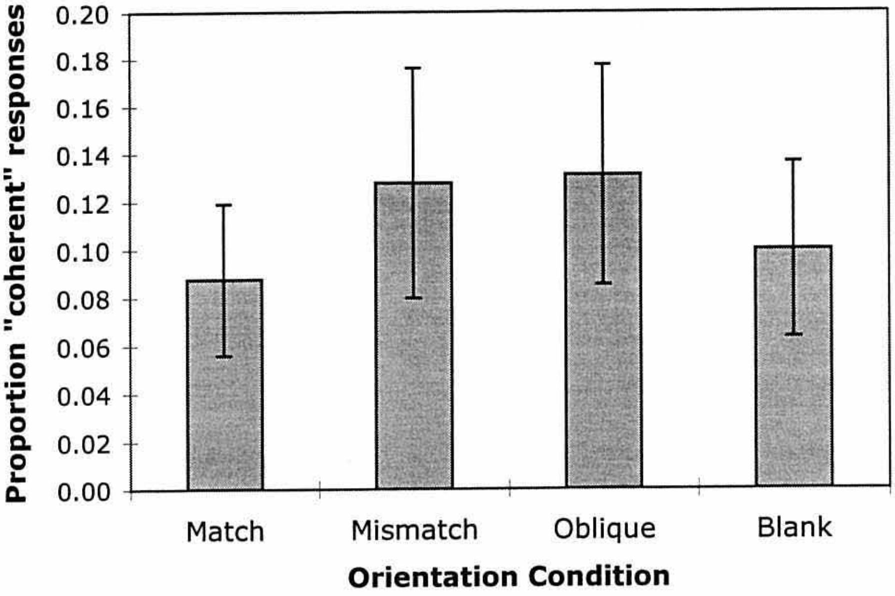

Figure 42. False alarm rates for four orientation conditions. A repeated measures ANOVA showed that there were no significant differences among conditions. Error bars represent ± 1 s.e.e.

d′ values were calculated for the four orientation conditions (see Appendix B for a more detailed description of Signal Detection Theory and d′ calculations). The false alarm rate used was the proportion of "coherent" responses for 0% coherence following a blank inducer for each observer. This baseline condition was chosen because it was thought to be a "neutral" condition, identical in temporal onsets and offsets to the oriented inducers but with no oriented texture information in the blank square. A d′ value was calculated for each coherence value for each observer based on their false alarm rate. The resulting group mean d′ values for each orientation condition are plotted in Figure 43 as a function of percent coherence. I performed a repeated measures ANOVA on the d′ prime scores using orientation condition and percent coherence as within subject factors. The ANOVA showed a significant main effect of coherence, $F(6,168) = 53.27$, $p < .001$, a significant main effect of orientation condition, $F(3,168) = 3.2$, $p < .05$, and a non-significant interaction.

The absence of an interaction raises an important point. Any differences in task difficulty among the inducer conditions would appear as differences in the slopes of the functions, i.e., a more difficult discrimination between random and coherent would produce a flatter slope (more guesses) whereas an easier discrimination would produce a steeper slope (more precisely defined responses of signal present vs. signal absent). Significant differences in difficulty among orientation conditions (i.e., slope) would produce an interaction between coherence and orientation condition. Since there was no significant interaction here, there was no significant difference in task difficulty among inducer conditions. The differences must stem from the sliding of the functions to the left and right on the abscissa indicating heightened or lessened sensitivity to the test.

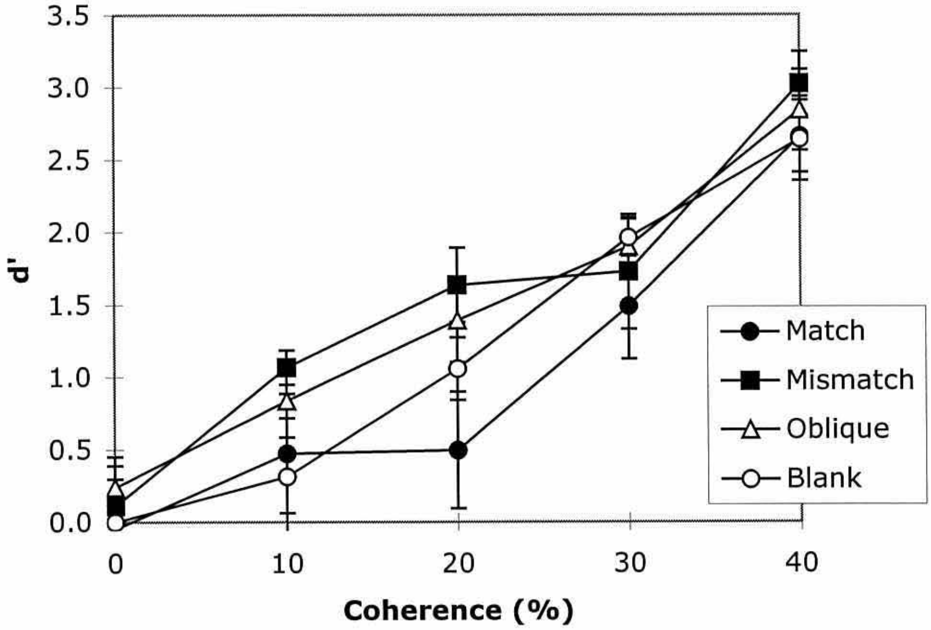

Figure 43. Group mean d' scores plotted as a function of percent coherence and orientation condition. Error bars represent ± 1 s.e.e. of the mean.

A planned simple comparison between the d' psychometric functions for match and mismatch (as a function of coherence) showed that observers reported significantly more "coherent" responses for the mismatch condition than the match condition, $F(1,9) = 5.94$, $p < .05$. Thus, observers were more sensitive to test textures that mismatched the inducer in orientation, again demonstrating the texture contrast effect.

To more closely address the magnitude of the texture contrast effect, orientation coherence thresholds were interpolated for each orientation condition using group mean d' scores. Fitting a least squares line to each psychometric function, the coherence value corresponding to a d' score of 1.5 was interpolated for each orientation condition. Given that the mean false alarm rate was .11, this criterion corresponds to a proportion "coherent" score of approximately .60. Using this method, the texture coherence thresholds for the four orientation conditions were calculated and are shown in Figure 44.

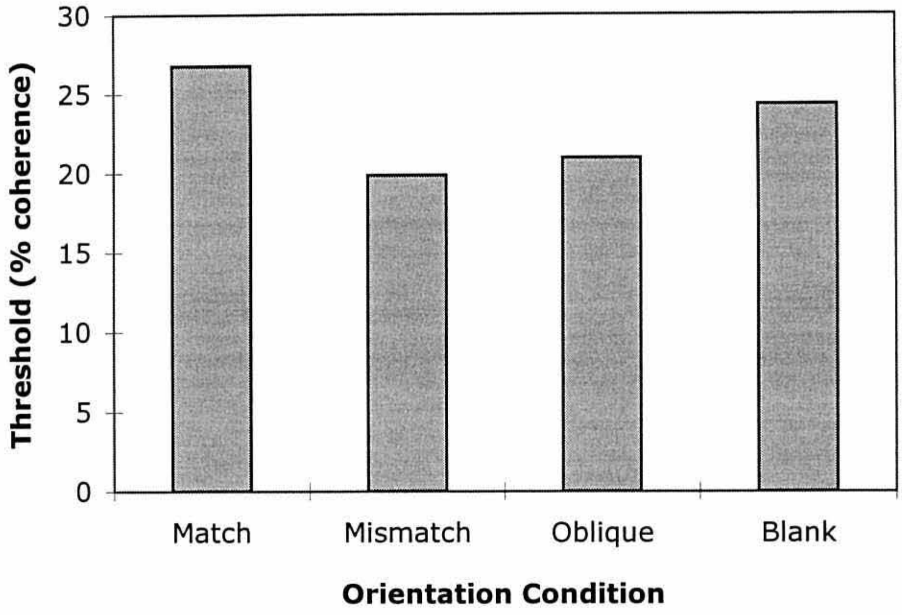

Figure 44. Group mean interpolated coherence thresholds as a function of the orientation condition between inducer and test.

The first outcome to note is that the threshold for mismatch (threshold = 19.83% coherence) was lower than that for match (threshold = 26.76% coherence). This again illustrates the texture contrast effect. The difference in sensitivity between match and mismatch here is 6.93% coherence and can be considered the magnitude of the contrast effect, i.e., observers required about 7% more coherently oriented line segments in the match condition to reach the same performance as the mismatch case.

Does this effect result from an increase in sensitivity for mismatched patches or a decrease in sensitivity for matched patches? Examining the threshold for the blank inducer condition, it seems to be a combination of both. The mean coherence threshold for the blank inducer condition (threshold = 24.31%) was lower than that of the match condition (threshold = 26.76%), and higher than the mismatch condition (threshold = 19.83%) and the oblique condition (threshold = 20.92%). These results indicate that the sensitivity change in the tilt contrast effect is a combination of an *increase* in sensitivity for mismatched textures (both mismatch and oblique conditions) and a *decrease* in sensitivity for those matched to the inducer. This is consistent with the motion aftereffect which is thought to be due to both decreased sensitivity for motion episodes matched in direction and increased sensitivity to motion episodes mismatched in direction (Raymond & Isaak, 1998; Raymond & O'Donnell, 1996).

A second outcome to note here is the threshold for the oblique inducer condition. As can be seen in Figure 44, the coherence threshold for the oblique condition (threshold = 20.92% coherence) is very similar to the threshold for mismatch. Based on the threshold for match, the magnitude of the tilt contrast effect with oblique inducers is 5.83% coherence. It seems that a difference in orientation of 45° between inducer and test is sufficient to produce a change in sensitivity that is almost as large as that produced with a maximal (90°) difference in orientation. This result supports the notion that texture contrast is an all-or-nothing effect that does not change in magnitude as a function of orientation differences. However, smaller differences in orientation between the inducer and the test, i.e., less than 45°, could be examined to clarify this relationship.

Discussion

In summary, observers were asked to view two successively presented texture patches and determine the orientation of the first patch and whether the second patch was coherent or random. There were three main findings:

1. Sensitivity for texture patches matched in orientation was significantly lower than patches that were mismatched in orientation. This demonstrated a texture orientation contrast effect.

2. Based on performance with a blank inducer condition, this contrast effect is due to both an increase in sensitivity to mismatched patches and a decrease in sensitivity to matched patches.

3. Sensitivity for texture patches that followed oblique inducers was similar to that for the mismatched condition indicating that for orientation differences greater than 45°, the texture contrast effect does not change magnitude and shows no gradient of strength as a function of orientation differences.

What mechanism could mediate this orientation-specific temporal contrast effect? According to an adaptation explanation of orientation aftereffects, changes in sensitivity to the test texture, resulting from previous viewing of an oriented inducer texture, are due to fatigue of orientation specific analyzers at early stages of visual processing (Coltheart, 1971). The resulting neuronal fatigue results in the observer's perception of negative aftereffects. This explanation suggests that there may be decreased sensitivity for test stimuli matched in orientation to the inducer but there should be no sensitivity change (compared to baseline) for 90° mismatched orientations (i.e., a vertical line in the presence of a horizontal line should not make the horizontal line look less horizontal, due to their perpendicularity; Morant & Harris, 1965). It also suggests that, because orientation varies along a continuum, the contrast effect should show a gradient in magnitude as a function of orientation differences between inducer and test (Coltheart, 1971; Morant & Harris, 1965). For example, an inducer quite similar in orientation to a test stimulus should show a different size of contrast effect than an inducer quite different to the test. In the present experiment, a test that follows an oblique (45°) inducer should have shown a moderate texture contrast effect when in actuality it showed a contrast effect just as large as the 90° mismatch case. The data presented here does not support the predictions of an adaptation explanation.

Lateral inhibition between oriented analyzers that persists in time has also been proposed as an explanation for orientation aftereffects (Blakemore, Carpenter & Georgeson, 1970; 1971). This inhibition among line segments has the effect of making line segments appear more different in orientation than they actually are. While this explanation is useful for tilt aftereffects that are small in orientation magnitude, only very small distortion effects are seen for lines differing by more than 30°. Furthermore, lateral inhibition may cause small perturbations in orientation that would cause test patches to appear less coherent than they actually are. These perturbations, however, would not be so orientation-specific as to cause a consistent texture contrast effect. Thus, for the stimuli used here, this explanation is limited.

Another explanation for the contrast effect is that it is mediated by selective attention mechanisms. Ecologically, there is a distinct advantage for an organism that can detect new objects (predators or prey) entering a visual scene quickly and efficiently. One possible mechanism for this optimization is the attentional inhibition of previously coded objects at the point of their offset (Raymond et al., 1998). If all objects present in a scene are assigned an inhibitory "tag" once they have been processed, any new object entering the scene has a greater chance of being detected than if all items are coded equally. Furthermore, this inhibitory tag will reduce the chance of the same stimulus being coded repeatedly (causing perseveration of action: MacKay, 1987; Mussler & Hommel, 1997) to the decrement of another, novel stimulus that may hold crucial information. In the previous two experiments, observers may have encoded the features of the inducing stimulus and applied inhibition to those features, including orientation. This inhibition would be subsequently reflected through reduced sensitivity for those test patches that matched the inducer in orientation. Conversely, novel or different stimuli would have had an advantage over those previously coded and this would be reflected as an enhancement of sensitivity for mismatched textures. This hypothesis is consistent with the reduced detection rate of repeated stimuli in repetition blindness (Kanwisher, 1987), with the decreased sensitivity for similarly directed motion episodes in the motion aftereffect (Raymond, et al., 1992) and may contribute to the apparent negative aftereffects of tilt and size (Blakemore & Sutton, 1969; Campbell & Maffei, 1971; Gibson, 1937; Gibson & Radner, 1937).

For the texture orientation contrast effect, a selective attention explanation predicts a decreased sensitivity for matched orientations and an enhanced sensitivity for any mismatched orientations. The data presented in the last two experiments support this prediction. This account would also predict that whatever the orientation difference between inducer and test, there should be a contrast effect provided that the orientation difference is large enough to be resolved. The oblique inducer condition supports this, although further experiments testing smaller orientation differences would make this relationship more clear.

If the contrast effect is indeed mediated by an inhibitory selective attention mechanism, then some predictions can be made as to its effect in different situations. First, the time course of the contrast effect should reflect a gradual build up of inhibition in response to a coded stimulus. Specifically, as the time interval between inducer and test increases, so should the magnitude of the contrast effect and it should extend for relatively long periods of time before dissipating. Second, diverting attention away from the inducer should effectively "switch off" the contrast effect if it is indeed mediated by selective attention. Experiments 14 and 15 respectively investigate these predictions.

Experiment 14 - Inducer-Test Time Interval and Orientation Contrast

An inhibitory attentional mechanism that involves building inhibition after initial perceptual coding may take time to develop and may persist for relatively long time intervals. Here the time interval between inducer and test stimuli was varied to measure the temporal extent of the texture orientation contrast effect. All methods here were the same as in the previous experiment unless specified.

Methods

Observers

Fifty healthy naive participants took part in this experiment (mean age = 25.18 years, SD = 7.36 years), ten in each of the five inter-stimulus-interval (ISI) conditions (100, 200, 300, 400, 700 ms). There were no significant differences among the group mean ages of the five ISI conditions.

Texture Stimuli

Inducing stimuli were composed of 100% coherent texture patches oriented vertically or horizontally. On one third of trials, the inducer was a square black outline the same size as the texture patch and filled with the same blank grey as the background field. Test texture stimuli were varied in coherence from 0% to 40% in 10% steps.

Procedure

Training Phase: Observers completed a training phase identical to that in the previous experiment.

Dual Task Phase: After criterion had been reached for the coherent vs. random judgment, observers completed 30 practice trials of the experimental dual task. A trial was initiated by a key press and consisted of the presentation of a fixation cross for 1500 ms, an inducer patch for 200 ms, a blank inter-stimulus-interval for 100, 200, 300, 400 or 700 ms, followed by a variable coherence test patch for 200 ms. None of the stimuli were masked. Observers were asked to report the orientation of the inducing patch (vertical, horizontal or blank) and to report whether the test patch was coherent or random. Observers indicated their responses by pressing two keys on the computer keyboard at the end of each trial. In the case of a blank square inducer, observers pressed either of the two keys used for the oriented inducers. No feedback was given after the initial 30 practice trials.

The two test orientations (vertical and horizontal) were presented following each of three inducer conditions (vertical, horizontal and blank). Five coherence values were presented for each of these six orientation conditions and each trial was repeated eight times resulting in a total of 240 trials per block presented in a pseudo-random order.

Single Task Phase: Observers also completed two blocks of a single episode texture in which a variable coherence texture patch was presented for 200 ms. This single episode was the same as the test patch in the dual episode texture task and observers were to judge whether the patch was coherent or random and indicate their responses using the same keys. Patches were vertical or horizontal and five coherence values (0, 10, 20, 30 and 40%) were presented ten times each for a total of 70 trials presented in a pseudo-random order. Testing was completed in 45 minutes in one session.

Results

Only those trials for which the observer correctly reported the inducer orientation were included in this analysis. All observers were correct on more than 87% of inducer orientations (mean = 95.9%, SD = .03%) and the errors were not systematic with respect to orientation condition. Each observer's proportion "coherent" responses were plotted as a function of percent coherence to examine psychometric functions for match, mismatch and blank conditions for each inter-stimulus-interval. The function for the single texture task case was also included for comparison.

False alarm rates were calculated and are plotted in Figure 45 as a function of inducer condition and ISI. Each observer's proportion "coherent" responses for 0% coherence in each orientation condition were used as false alarm rates (see Appendix B for a discussion of d' calculations). Since at 0% coherence, match and mismatch conditions are the same (i.e., there is no signal in the test texture, therefore it cannot match or mismatch the inducer) the mean false alarm rate for the two conditions was used as a baseline for both. One of the first things to notice in Figure 45, is that the false alarm rates for match and mismatch are much higher than for blank and single in the 100 ms ISI condition. Since t-tests showed that match and mismatch were not significantly different from each other and blank and single were not significantly different from each other, the two groups were compared to each other (i.e., textured inducer vs. blank or no inducer, collapsed over coherence). Results showed that false alarm rates for the textured inducer conditions (i.e., match and mismatch) were significantly higher, $t(49) = 3.94$, $p < .01$, than those for blank and single. This result may reflect an illusion of coherence that is induced in the neutral stimuli when immediately preceded by a coherently oriented texture, as was seen in Experiment 12. This illusion is not present for the longer ISI intervals, which is consistent with the absence of this effect in Experiment 13. In fact, there was no significant difference in false alarm rates for ISIs longer than 100 ms.

A second effect to note in Figure 45 is that false alarm rates for the blank and single conditions do not change as a function of ISI. An ANOVA on these false alarm rates using ISI as a between subjects factor indicated no significant main effect of ISI for the blank and single conditions. This demonstrates that there were no differences in base false alarm rates, with neutral (or no) inducers, among the groups of different observers.

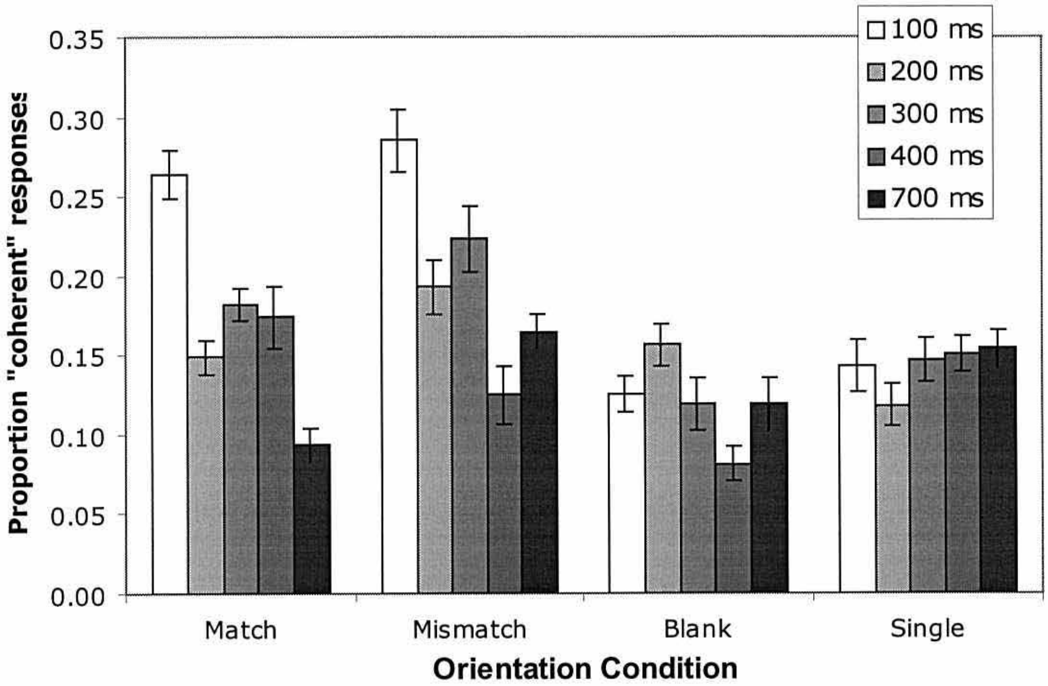

Figure 45. False Alarm rates for four inducer conditions plotted as a function of inter-stimulus-interval. False alarm rates for match and mismatch were significantly different than those for the blank and single conditions with an ISI of 100 ms between inducer and test. Error bars represent ± s.e.

d′ values were calculated for the three orientation conditions for all ISIs and group mean psychometric functions are plotted in Figure 46. Since a d′ score defines the *relationship* between the base false alarm rate and the hit rate, differences in base false alarm rates among conditions (i.e., the 100 ms ISI condition here has a higher false alarm rate than the other conditions) will not substantially affect relative d′ scores. See Appendix B for a demonstration of this.

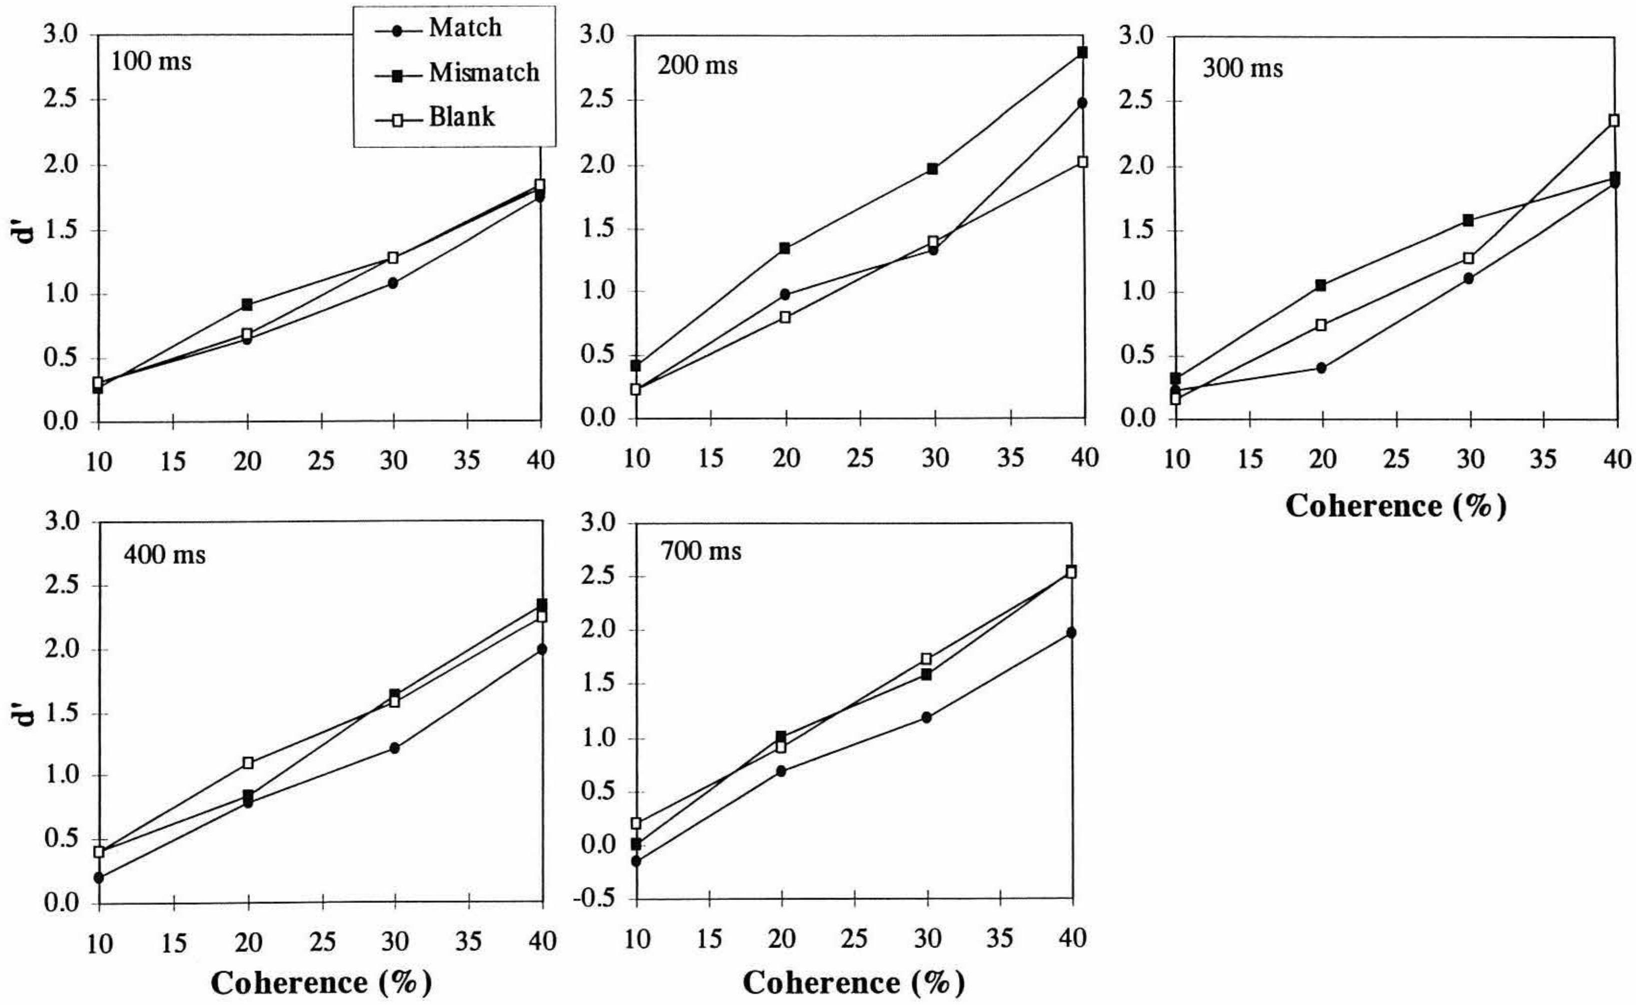

Figure 46. Group mean psychometric functions for 5 ISI conditions as a function of orientation condition.

An overall mixed design ANOVA was performed on the d' scores using ISI as a between subject factor and using orientation condition and percent coherence (10% to 40% coherence) as within subject factors [7]. The results of this analysis showed a non-significant main effect of ISI, a significant main effect of orientation, $F(2, 108) = 11.69$, $p < .001$, and a significant main effect of coherence, $F(3, 108) = 514.9$, $p < .001$. The interaction between coherence and ISI was also significant, $F(12, 135) = 2.33$, $p < .01$. No other higher order interactions were significant.

First, the significant interaction between coherence and ISI needs to be addressed. This analysis indicates that there was a significant difference in the slope of the psychometric functions as a function of ISI that often indicates a difference in task difficulty [8]. While it makes sense that the observers' task was more difficult for shorter ISIs, this difference may have contributed systematically to the magnitude of the texture contrast effect over time independent of changes in sensitivity. The slopes were calculated for each of the mean psychometric functions (collapsed over orientation condition) for each ISI and are presented graphically in Figure 47. The significant interaction was most likely due to the flat slope of the 100 ms condition compared to the other ISIs. Observers may have found this condition difficult and their responses seem to reflect this. Interestingly, there was no systematic upward trend through the slopes as ISI increased. If observers found the task to be increasingly easy as the ISI increased, this would have been the result. Instead, the slopes for ISIs of 200 to 700 are quite similar. If there was indeed a systematic trend here between slope and ISI and it corresponded to the trend between the magnitude of the texture contrast effect and ISI, it would be of concern. In fact, this effect seems to be present only for the 100 ms condition and may be because ISI was a

[7] The 0% coherence point was excluded from this analysis because each condition served as it's own baseline, thus the 0% coherence d' values were all zero or very close to zero in every case.

[8] Differences in task difficulty (reflected by the slope of the psychometric function) should be distinguished from differences in observer sensitivity (reflected by the position of the function along the ordinate).

between subject variable and there were different groups of observers in each ISI condition. For this reason, individual repeated measures ANOVAs were performed on each of the five ISI conditions separately using orientation condition and percent coherence (10 to 40%) as within subject factors. A summary of the results of these analyses is shown in Table 1.

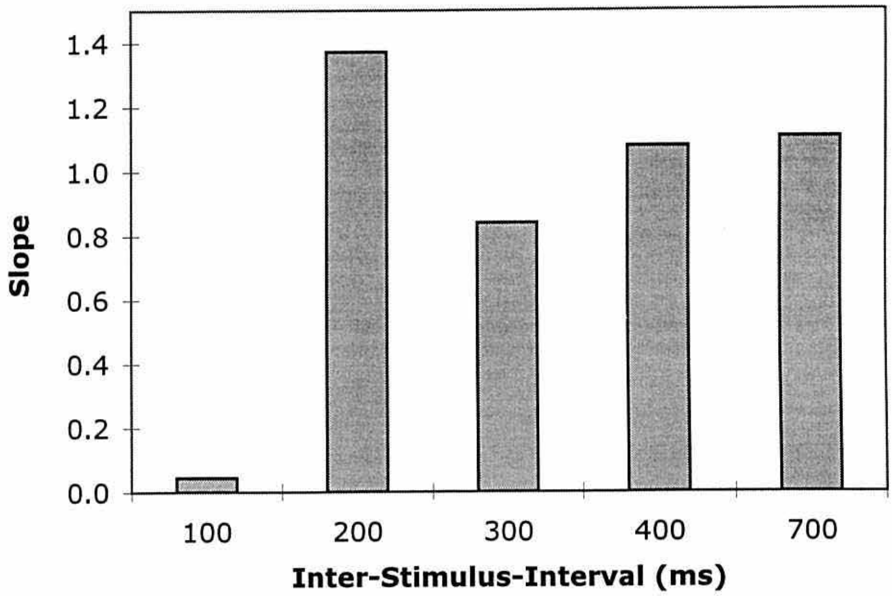

Figure 47. The slopes of mean psychometric functions for 5 ISI conditions. Psychometric functions were collapsed over orientation condition.

For the 100 ms condition, a repeated measures ANOVA showed a significant main effect of orientation condition, a significant main effect of coherence, and a non-significant interaction. A planned simple comparison between match and mismatch, collapsed over coherence, showed a non-significant difference in d' values. The main effect of orientation was due rather to the difference in d' scores between the blank and single task conditions which will be discussed in a later section. Thus, there was no significant orientation-specific change in sensitivity for a test patch when it followed an oriented inducer patch by 100 ms.

A repeated measures ANOVA performed on the 200 ms condition resulted in a significant main effect of orientation condition, a significant main effect of coherence, and a non-significant interaction. A simple comparison between match and mismatch, collapsed over coherence, showed that d' scores for mismatch were significantly higher than those for match, indicating the presence of an orientation contrast effect.

ANOVAs for the 300 and 700 ms conditions showed similar results to the 200 ms condition and are summarized in Table 1. Significant main effects of orientation condition and coherence and non-significant interactions were found. Comparisons between match and mismatch showed that d' scores for mismatch were significantly higher than those for match for the 300 and 700 ms conditions. For the 400 ms condition, the main effect for coherence was significant, the main effect for orientation was marginally significant and there was no significant interaction. A simple comparison was performed between match and mismatch (collapsed over coherence) and showed significantly higher d' scores for mismatch compared to match.

Table 1. Significance values for the main effects of coherence and orientation and the results from simple comparisons between match and mismatch. A significant difference in the match vs. mismatch column indicates the presence of a texture contrast effect. Results from five ISI conditions are summarized.

ISI	Coherence	Orientation	Match vs. Mismatch
100	$F(3,54) = 12.24, p < .001$	$F(2,54) = 43.73, p < .001$	n.s.
200	$F(3,54) = 52.94, p < .001$	$F(2,54) = 9.52, p < .001$	$t(9) = 3.88, p < .01$
300	$F(3,54) = 64.69, p < .001$	$F(2,54) = 2.68, p < .05$	$t(9) = 2.66, p < .05$
400	$F(3,54) = 55.46, p < .001$	$F(2,54) = 1.46, p > .05$	$t(9) = 2.48, p < .05$
700	$F(3,54) = 74.21, p < .001$	$F(2,54) = 5.96, p < .001$	$t(9) = 5.13, p < .001$

These results show that the texture contrast effect is not present with a 100 ms ISI between inducer and test, but is present for ISIs between 200 and 700 ms in duration. To examine the magnitude of the texture contrast effect for each ISI, coherence thresholds were interpolated by fitting a line of best fit, using the least squares method, through the psychometric functions for each orientation condition. I interpolated the coherence value corresponding to a d′ value of 1.5 for each orientation condition for each ISI [9]. These texture coherence thresholds are illustrated in Figure 48. The mean coherence threshold for all observers in the single task case (mean threshold = 23.7% coherence, s.e. = 1.57% coherence) is illustrated as a dotted line. As can be seen in Figure 48, thresholds for the match condition are higher than that of the mismatch condition for all ISI durations.

The issue of whether these orientation-specific changes in sensitivity are due to increased sensitivity for mismatch or decreased sensitivity for match can be examined using the baseline blank inducer dual task texture. If we assume that the blank dual task texture condition is a measure of the typical performance of each observer for the dual task situation, then we can use it to evaluate performance on other inducer conditions. We can also compare this to an observers' optimal performance in the single task case (i.e., only a single judgment task and only one visual event).

---

[9] The maximum standard error of estimate (s.e.e.) for these lines of best fit was 3.7% coherence.

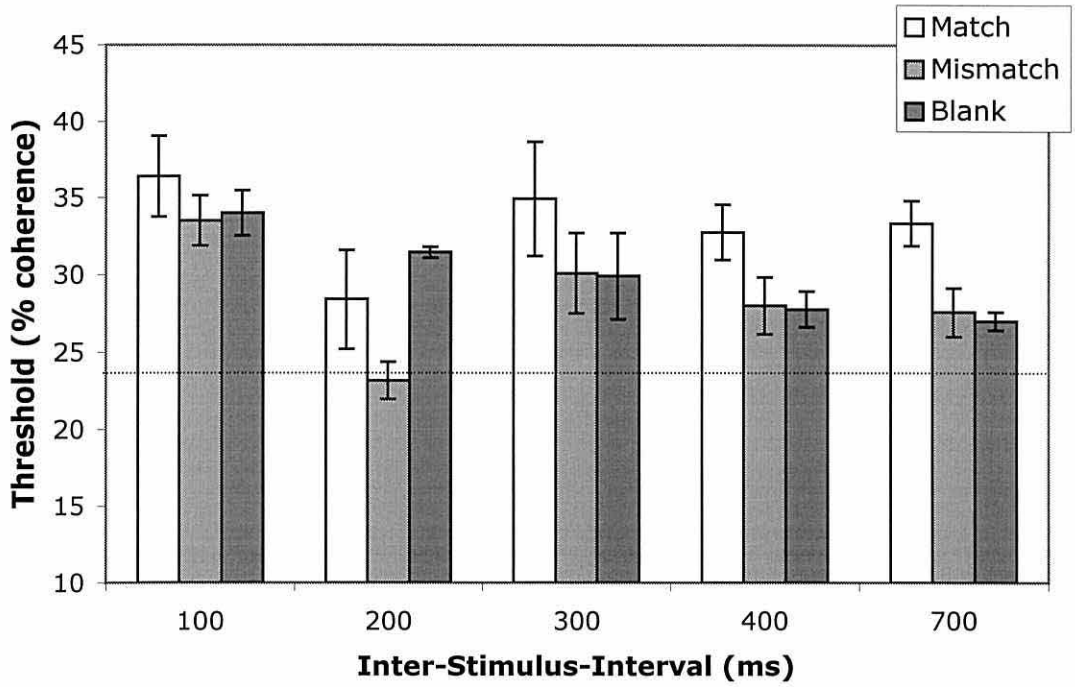

Figure 48. Coherence thresholds for four orientation conditions are plotted as a function of ISI. Mean threshold for all observers in the single texture task is illustrated as a dotted line (mean = 23.8% coherence).

One of the first effects to notice here is the reduced thresholds for the 200 ms ISI condition. Thresholds were much lower for the match and mismatch orientation conditions (as good as the single task condition), than for the other ISI conditions though the relationship between match and mismatch was preserved. Since ISI was a between subjects factor in this experiment, it seems that the difference may be due to the fact that this group of observers was more sensitive than the other groups although there is no obvious reason for this. This difference was not statistically significant, however, in the overall analysis of variance (i.e., the main effect of ISI) and so I will focus on the *relationship* between the match and mismatch conditions rather than absolute thresholds. For all ISI conditions greater than 100 ms, thresholds for match are elevated with respect to mismatch. This orientation contrast effect is not presented for the 100 ms condition.

A second effect to notice is that thresholds for mismatch are as good as the blank condition in all cases except for an ISI of 200 ms. If we assume that the blank condition is a performance measure of the typical sensitivity in a dual task case, this indicates that the orientation contrast effect is characterized by a decrease in sensitivity for the match case. Actually, there is evidence that the blank condition itself is elevated for the short ISI durations (see below) and therefore the decrease in sensitivity for the match condition may be larger than it appears to be in Figure 48.

As mentioned above, thresholds for the blank condition are elevated for short ISIs and decrease systematically as ISIs get longer. In Figure 49, I have re-plotted the thresholds for the blank condition along with thresholds for the single task condition for the same observers. Thresholds for the blank condition are elevated for ISIs of 100 - 200 ms but are similar to those of the single task condition for ISIs longer than 300 ms. These conditions are perceptually very similar in that there is no texture information in the blank inducer, only a blank square area. These large differences in threshold are unlikely to be caused merely by the difference between a dual task and a single task, or they would persist for the longer ISI conditions. Rather, the differences are most likely due to the perceptual differences in the figures between inducer and test. The observers' perception of the blank inducer was of an opaque square figure separable from the background rather than a neutral black transparent outline of a square. The temporal persistence of some sort of perceptual filling-in process of the surface of this square may explain the elevations in thresholds for the subsequent texture patch seen here.

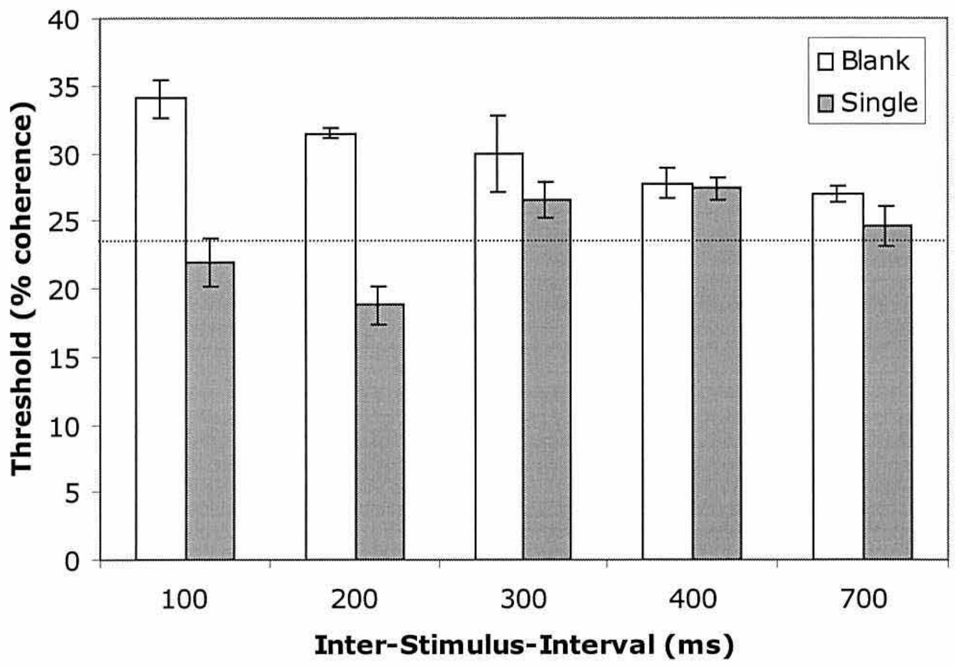

Figure 49. Group mean coherence thresholds for the blank and single conditions as a function of ISI. These conditions are essentially the same, i.e., there is no texture information in the blank inducer, yet there is a large difference in thresholds for the shorter ISIs. The dotted line is the mean coherence threshold for the single task. Error bars represent ± 1 s.e.e.

A more clear way of examining the time course of the texture contrast effect is to calculate the magnitude of the contrast effect by subtracting the interpolated threshold for mismatch from the threshold for match. This value reflects the extra percentage of coherently oriented line segments that observers required in the match condition to respond "coherent" as often as in the mismatch condition. It allows a comparison of the magnitude of the texture contrast effect among groups of different subjects as a function of ISI and is represented graphically in Figure 50. As can be seen from Figure 50, the texture contrast effect is minimal (and is not statistically significant) at the shortest ISI of 100 ms. For longer ISIs, the change in sensitivity between match and mismatch increases in magnitude to 700 ms. The texture contrast effect does not disappear within the range tested here.

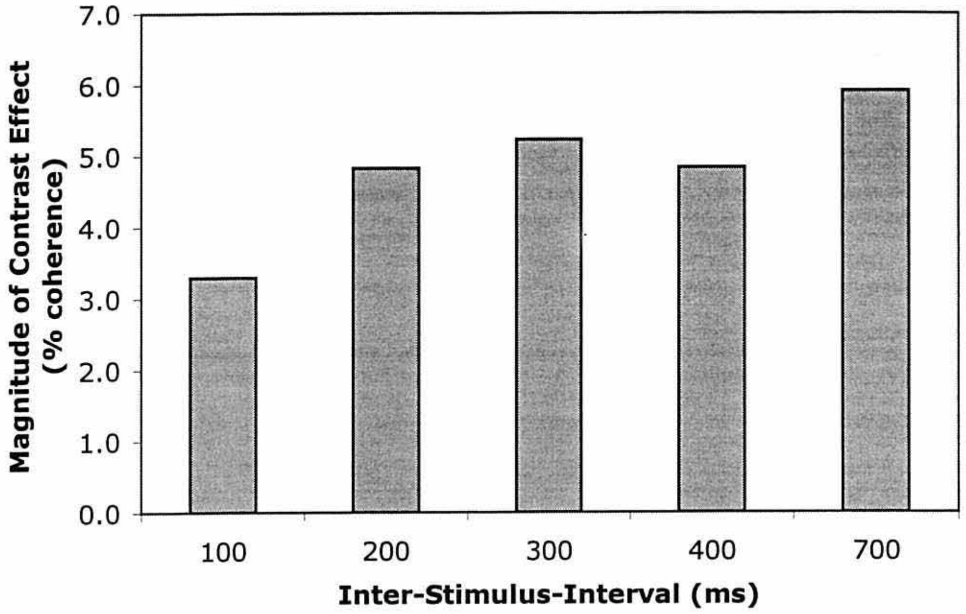

Figure 50. Magnitude of tilt contrast effect as a function of ISI between inducer and test. Texture coherence thresholds were interpolated for match and mismatch conditions. The difference in sensitivity between the two is the magnitude of the tilt contrast effect.

Discussion

I have demonstrated in this series of experiments that:

1. Successive orientation contrast effects can be demonstrated with briefly presented complex oriented textures. Sensitivity for texture patches mismatched in orientation to a previous inducer was significantly higher that that for textures matched in orientation.

2. Orientation contrast effects take time to develop and persist for extended lengths of time (700 ms) suggesting an attentional component.

3. Observers experience an illusion of orientation coherence in otherwise random textures when they are preceded by a more coherently oriented texture (for short ISIs only).

4. The texture coherence paradigm can be used to test changes in sensitivity caused by successive stimulus interactions. Such changes in sensitivity can be discriminated from changes in observer criterion by using Signal Detection Theory d' calculations.

These results suggest that everyday sensitivity to naturally occurring complex textures can be affected by prior exposure to other textures and that this effect may be mediated by a combination of low-level perceptual and higher-level attentional mechanisms. The perceptual hypotheses of Blakemore, Coltheart and colleagues effectively explain adaptation-type effects in perceptual terms. For the fairly long-lasting contrast effects demonstrated here attentional hypotheses may offer an alternative explanation. These relatively long-lasting effects on sensitivity make sense ecologically, giving new and novel stimuli an advantage over older, previously coded stimuli in a visual scene. Animals with this attentional "filtering" mechanism would surely have an advantage over those without. The relatively slow development of the contrast effect, at least 200 ms from the offset of the initial stimulus, would allow the animal to code the basic features of a stimulus and either identify it for further scrutiny or disregard it as an item of little interest. There are, however, situations in which this may not be adaptive, for example, if an animal that has not been fully coded moves in and out of sight. In this situation, attentional filtering causing adaptation-like effects for non-novel stimuli may be detrimental to performance. Perhaps the benefits of the increased sensitivity to novel stimuli outweigh the costs of the decreased sensitivity to non-novel stimuli.

One of the other interesting effects found here is the illusion of coherence that is induced in a random texture when it follows a more coherently oriented texture.

This effect was demonstrated in two of the three experiments in this series by showing that observers reported significantly more "coherent" responses to a 0% coherence stimulus than expected. Interestingly, this effect only occurred for the shortest ISI condition (100 ms). It seems that the perception of orientation persists in time, but as was demonstrated in the orientation judgments of Experiment 12, this illusion causes the perception of coherence in the *opposite* orientation than the inducer. This illusion does not persist for longer ISIs as was reported in Experiment 13. A similar illusion is often reported when coherent motion is perceived in directionally random stimuli after observers have previously viewed a coherently moving stimulus. This illusory coherent motion seems to stream in the *opposite* direction to the inducer. Both of these coherence illusions could be explained in terms of attentional coding. After the offset of the inducing stimulus, the orientation or motion direction just coded is inhibited to reduce the redundant, repetitive coding of the same stimulus over again. Stimuli that differ in orientation or motion direction will therefore have an advantage over others and this may cause the illusion of coherence in an otherwise neutral stimulus.

The finding that thresholds were elevated for textures following the blank square condition over that of the single task condition is another interesting finding reported here. This effect may be caused by the persistence of a perceptual "filling-in" process during the processing of the square figure. Since the visual system is preferentially sensitive to edges and borders, in order to perceive continuous surfaces between those borders, there must be a perceptual filling-in process by which that surface is constructed. Evidence from studies investigating the properties of Kanisza figures indicates that this surface is indeed constructed between inducers and behaves like a surface perceptually (Gove, Grossberg & Mingolla, 1995). This filling-in process has also been incorporated into Grossberg's computational model explaining illusory surfaces and other effects. In the present contrast effect experiment, the process of perceptually filling-in the blank grey area within the black outline to create the percept of a surface could persist in time and may compete with the texture integration process to result in elevated thresholds.

At this point, the fact that the texture orientation contrast effect takes time to develop and persists in time is the primary evidence for the role of attention. If this contrast effect is indeed mediated by selective attention mechanisms, then any re-direction of attentional resources during exposure to the inducer should reduce or eliminate the effect. Experiment 15 investigates this prediction and establishes more evidence supporting the role of selective attention in this orientation contrast effect.

Chapter 7 - The Role of Selective Attention in Texture Integration

Abstract

I have previously demonstrated that texture integration processes persist in time and cause changes in sensitivity for test textures presented after an oriented inducer texture. This texture contrast effect takes time to develop and persists for at least 700 ms after the inducer. Contrast effects could be due to "low-level" visual mechanisms such as fatigue of orientation-selective analyzers or lateral inhibition that persists in time, as is thought to cause the classic tilt aftereffect. However, recent experiments suggest that texture contrast effect may have an attentional component based on the persistence and development of the effect over time. To test this, I asked whether changing the locus of selective attention away from the inducing stimulus would have an effect on it's magnitude. I presented observers with a 100% coherent inducing texture patch with a low contrast "E" embedded at the centre followed after 300 ms by a variable coherence test texture to test subsequent sensitivity. Observers were asked to either attend to and report the orientation of the inducer texture ("attend texture" condition) or to attend to and report the orientation of the central "E" ("attend letter" condition). By presenting the identical stimuli in both conditions, any differences in the magnitude of the texture contrast effect may be attributed to how selective attention is allocated to objects or surfaces. Results showed a texture contrast effect when observers attended to the inducer texture, but no contrast effect when attending to the letter. These results suggest that the texture orientation contrast effect is meditated by selective attention and can be "switched off" by re-directing that attention.

Selective Attention in Integration Tasks

Attention can be defined as a selection mechanism that acts on perceptual input to preferentially process areas of that input that require sensory focus for more intense processing and to inhibit areas of irrelevant stimulation. It is generally assumed that we have no attentional control over low level visual processing such as that in primary visual cortex. Thus, perceptual tasks that are sensitive to attentional manipulation are considered to be relatively "high-level" in nature, and probably mediated by extra-striate cortex. Since the pattern of results from the last few experiments leaves open the possibility that attention may mediate contrast effects, it makes sense to test whether the contrast effect requires selective attention to the inducer stimulus.

The effects of manipulating the locus of attention have been examined in the motion aftereffect literature, producing some interesting results (Chaudhuri, 1990; Shulman, 1991; 1993). The motion aftereffect occurs when observers view a moving stimulus, often an array of moving dots or a grating pattern, for a given period of time and experience an illusion of motion in the opposite direction when the motion is stopped. A study by Chaudhuri (1990) examining the classical motion aftereffect demonstrates that an aftereffect once thought to be a low-level visual illusion is in fact susceptible to attentional manipulation. Studies of inter-ocular transfer (Barlow & Brindley, 1963; Faveau, 1976), dichoptic stimulation (Anstis & Moulden, 1970), and motion aftereffects contingent upon other visual parameters such as colour (Mayhew & Anstis, 1972, Potts & Harris, 1975) and orientation (Mayhew & Anstis, 1972), imply that the motion aftereffect is generated at the level of the visual cortex. It has been hypothesized that direction-selective motion analyzers are fatigued by the adapting stimulus and this changes the subsequent equilibrium among direction-selective analyzers resulting in a negative aftereffect (Barlow & Hill, 1963; Hammond, Mouat & Smith, 1988; Sekuler & Ganz, 1963). This hypothesis predicts that because the motion aftereffect is mediated by relatively "low-level" mechanisms, any changes in the locus of attention during adaptation should have little or no effect on the motion aftereffect. Recent evidence, however, suggests that V1 itself may be attentionally modulated (Lamme, 1995; Zipser, Lamme & Schiller, 1996).

Chaudhuri (1990) presented observers with a random dot motion stimulus with a superimposed alphanumeric display in the center. Observers were either required to monitor the alphanumeric display and press a key whenever a numeral was presented, or to passively view the display using the stream of central characters as a fixation point. After 60 seconds of viewing the inducer, the stimulus became stationary and observers were required to press a key when the motion aftereffect ceased. Results showed that the measured duration and rating of strength of the motion aftereffect was significantly reduced for observers attending to the central alphanumeric display than for those passively viewing the display. Chaudhuri noted that if indeed the motion aftereffect is due to depressed activity in a population of direction-selective analyzers then the attenuation of the motion aftereffect may occur because non-attended adapting patterns are less effective in driving those analyzers in the first place. This in turn would produce an attenuated aftereffect experience. This implies that either attentional mechanisms have considerable influence over processing in V1 where the motion aftereffect is thought to be generated, or that the motion aftereffect has an extrastriate contribution that may be modulated by attention. Since there is little evidence to support the influence of attention as early in the visual hierarchy as V1, the latter possibility seems more likely (but see Lamme, 1995; Zipser, Lamme & Schiller, 1996 for evidence of attentional modulation of V1).

Unfortunately, the Chaudhuri study has a number of weaknesses. Observers were asked to press a key when the motion aftereffect had ceased and were also asked to rate the strength of the aftereffect on a subjective rating scale. These methods are subject to inaccuracies in measurement as well as observer criterion shifts. A more objective measure of the magnitude of the motion aftereffect should have been used. Also, there was considerable difference in task difficulty between the digit detection task and the passive viewing task. Little attempt was made either to control for this task difficulty difference or to control where attention was focused during the passive viewing phase.

Shulman (1991) examined the effects of attentional manipulation on the three-dimensional motion aftereffect (reported by Petersik, Shepard & Malsh, 1984) which is produced in a similar way to the classic linear motion aftereffect but uses rotation in depth rather than motion in a single plane. Shulman asked observers to attend to one of two superimposed squares that were defined by a dot at each vertex and that rotated in opposite directions in depth (i.e., rightwards and leftwards). Observers were told to first attend and report any small perturbations in the movements of the dots defining either the small or large adapting square and then judge the motion direction of a subsequent ambiguous two-dimensional test square. Results showed that observers were more likely to report that the ambiguous test square rotated in the opposite direction to the attended adapting square. In fact, the motion aftereffect was reversed in sign depending on which adapting square was attended. The author concluded that mechanisms mediating the perception of rotation in depth are modulated by attention.

Shulman (1993) also examined the two-dimensional motion aftereffect using a slightly different distractor paradigm to divert attention from the adapting stimuli. Observers viewed four circular apertures (filled with grating patterns) that rotated clockwise or counterclockwise around the display. At the center of the display, a stream of rapidly changing letters was presented. The task during the adaptation period was to either attend to the rotating apertures and report any changes in grating orientation, or to attend to the central letter stream and detect any digits that appeared. The subsequent test stimulus was again an ambiguous motion stimulus and a directional judgment was required from observers. Results showed a motion adaptation effect when observers attended the moving apertures and a reduced aftereffect when observers attended the central letter stream. Interestingly, the motion adaptation effect was reduced but not eliminated even when observers were performing the taxing (approximately 75% correct) digit detection task at the center of the display.

The above results taken together indicate that an aftereffect that was previously thought to be mediated by low-level mechanisms instead can be modulated by attention. In the previous chapter, I demonstrated a contrast effect for oriented textures. After viewing an oriented inducer, sensitivity to a test texture changed depending on the orientation relationship between inducer and test. The question remained as to whether this orientation contrast effect required focal attention to the orientation of the inducing texture and whether the aftereffect could be "switched off" by redirecting attention to a central letter display embedded in the inducer. Given the evidence from the previous experiments of an attentional role, it seemed likely that there may be an attentional modulation of the orientation contrast effect.

To test the role of attention, I utilized a similar paradigm to that of Chaudhuri (1990) and Shulman (1993). Observers were asked either to attend to and report the orientation of a central letter embedded in the central cell of the inducer texture stimulus, or to attend to and report the orientation of the entire texture inducer. The thresholding technique was used, as in the previous chapter, to measure changes in sensitivity to a subsequent test stimulus. As before, only those trials for which the inducer or central letter orientation was correctly reported were used in the analysis (unlike Chaudhuri, 1990 and Shulman, 1993) to ensure that attention was indeed directed to the required location. The central letter of the inducer patch was fixed at a lower contrast than the texture to make the task attentionally demanding so that observers had little attentional resources left to allocate to the inducer orientation (but see discussion section). An inter-stimulus-interval of 300 ms was used between inducer and test to demonstrate the presence of a orientation contrast effect and still make the task demanding for observers.

If the orientation contrast effect is mediated by an attentional mechanism, then it may be eliminated or significantly reduced in magnitude when attention is directed away from the texture inducer. If, however, the orientation contrast effect is mediated primarily by a more low-level mechanism, then the contrast effect should occur regardless of whether focal attention is directed to the inducer or directed to a central letter.

Experiment 15 - Does the Orientation Contrast Effect Require Focal Attention?

Methods

Observers

Eight healthy naive participants took part in this experiment (mean age = 24.88 years, SD = 4.6 years).

Texture Stimuli

Inducer stimuli were 100% coherent texture patches with horizontal or vertical signal orientations. The central line segment of each was replaced by a lower contrast (Michelson luminance contrast of 30%) black "E" that was oriented either with it's long axis vertically oriented (i.e., as it would be read in text) or horizontally oriented (i.e., on it's "back"). The E was composed of four line segments identical to those used in the texture (10 pixels long, 1 pixel wide). Thus, the E was 13.4 x 13.4 arc min in size. Examples of the stimuli are illustrated in Figure 51. All combinations of inducer and central E orientations were presented on separate trials in a random order. Test stimuli had horizontal or vertical signal orientations and coherence was varied from 0% to 50% in 10% steps.

A Trial:

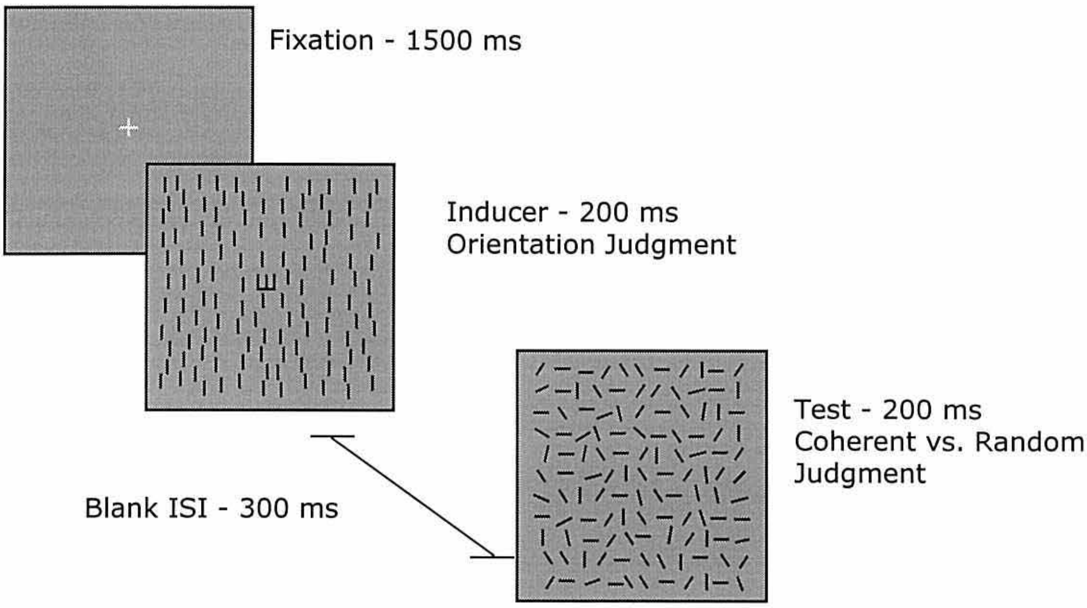

Figure 51. An illustration of a sample trial. Observers were asked to either attend to and report the orientation of the inducer texture or the orientation of the inducer letter and then to judge whether the test stimulus was coherent or random. Shown here is a mismatched trial where the inducer texture is vertical and the test texture is 30% coherent & horizontally oriented. The central letter of the inducer is horizontal.

Procedure

Training phase: Observers completed a training phase in which they judged whether a series of horizontal or vertical texture patches of 0% or 50% coherence were random or coherent. Practice patches were presented for 200 ms and were matched in every other way to test patches in the experimental phase. Responses were indicated by a key press and feedback was given in the form of a beep for either a false alarm (responding "coherent" to a 0% coherence patch) or a miss (responding "random" to a 50% coherence patch). A block consisted of 32 trials (about two minutes) in which the two coherence values in each of the two orientations were repeated eight times. Observers completed at least two blocks of practice to reach a criterion of no more than two false alarms and two misses. All observers reached this criterion within three practice blocks.

Experimental phase: A trial was initiated by a key press and was composed of a fixation cross presented for 1500 ms, an inducer patch for 200 ms, a blank inter-stimulus-interval for 300 ms and a test patch for 200 ms. Observers completed two blocks of the dual task, one "attend texture" and one "attend letter" the order of which was counterbalanced between observers. In the "attend texture" condition, observers were asked to judge the orientation of the texture inducer and to judge whether the test patch was coherent or random. In the "attend letter" condition, observers judged the orientation of the central E and whether the test patch was coherent or random [10]. Observers indicated these judgments with two key presses at the end of the trial. No feedback was given after an initial practice phase of 30 trials in which the observers were told the correct responses verbally.

Every combination of inducer orientation (horizontal and vertical) and central E orientation (horizontal and vertical) was presented for each of the test orientations (horizontal and vertical) resulting in eight orientation conditions. For each of the orientation conditions, six test coherence values were repeated three times each resulting in 144 trials per block. Certain conditions were collapsed in the subsequent data analysis to result in more than three data points per condition. Orientations were collapsed, as in the texture contrast experiments, resulting in two orientation conditions: match and mismatch. The congruency between the inducer letter and inducer texture was collapsed within the match-mismatch distinction. This resulted in the average of 12 points per observer for every coherence value being used in psychometric functions (for each attention condition).

Observers also completed two blocks of a single texture task in which only the test stimulus was presented to use as a baseline measure of texture perception

[10] In a pilot study, all "attend texture" and "attend letter" trials were randomly interleaved and observers were given a cue word at the beginning of each trial instructing them to attend to either the orientation of the central letter or the orientation of the inducer texture. This proved to be very difficult for observers to do and produced very variable results even though observers had no problems in a similar motion experiment using this paradigm (Raymond, O'Donnell & Tipper, 1998). I decided that the blocked design was sufficient as long as the order of blocks was counterbalanced among observers to reduce order and practice effects.

sensitivity. This was the same as the test stimulus in the dual task case and a coherent vs. random judgment was required. Vertical or horizontal patches with one of five coherence values ranging from 0% to 40% in 10% steps were presented five times each for a total of 70 trials. All testing was completed in one 50-minute session.

Results

Only those trials for which the observer correctly determined the inducer or letter orientation were included in this analysis. Observers were on average 97% correct (s.e. = 1%) on the inducer task and errors were not systematically distributed over orientation condition. Performance on the texture orientation task (mean = 97.7%, s.e. = 0.8%) was not different from performance on the letter orientation task (mean = 95.6%, s.e. = 1.4%). All observers achieved greater than 87% correct on both inducer tasks.

I plotted proportion "coherent" responses for match (i.e., inducer and test match in orientation) and mismatch (i.e., inducer and test mismatch in orientation) orientation conditions for each level of coherence for the "attend texture" and "attend letter" conditions to examine psychometric functions. This resulted in the average of 12 values representing each point on the function. I then calculated d' values for match and mismatch for the two attention conditions. To calculate d' values, I used as a base false alarm rate the mean proportion of "coherent" responses for a 0% coherence stimulus for each attention condition[11]. There was no significant difference between the false alarm rates for the two attention conditions. A repeated measures ANOVA on these d' values using percent coherence, orientation (match vs. mismatch) and attention condition (attend letter vs. attend texture) as within subject factors showed a significant main effect of percent coherence, $F(5,168) = 51.01$, $p < .001$, and a significant main effect of orientation, $F(1,168) = 6.83$, $p < .05$.

[11] The false alarm rate for the single task case could have been used here, but would not have reflected as accurately the false alarm rate for the dual task case.

Interestingly, there was no significant main effect of attention location. This indicates that there was no significant difference in sensitivity to the test patch due to differences in inducer task difficulty. Had there been a difference in task difficulty between the attend texture and attend letter conditions, then this main effect would have been significant.

The critical comparison that addressed whether attention modulates the texture orientation contrast effect was the interaction between attention condition and orientation. This interaction was significant, $F(1,168) = 10.72$, $p < .05$. As shown in Figure 52, observers show an orientation contrast effect when they attend to the inducer texture but do not show an orientation contrast effect when they attend to the central letter. Planned comparisons between match and mismatch for the "attend texture" condition (collapsed over coherence) resulted in a significant contrast effect, in which observers were significantly more sensitive to mismatch over match, $t(7) = 46.5$, $p < .001$. For the "attend letter" condition, there was no significant difference between sensitivity for mismatch and match, thus no orientation contrast effect.

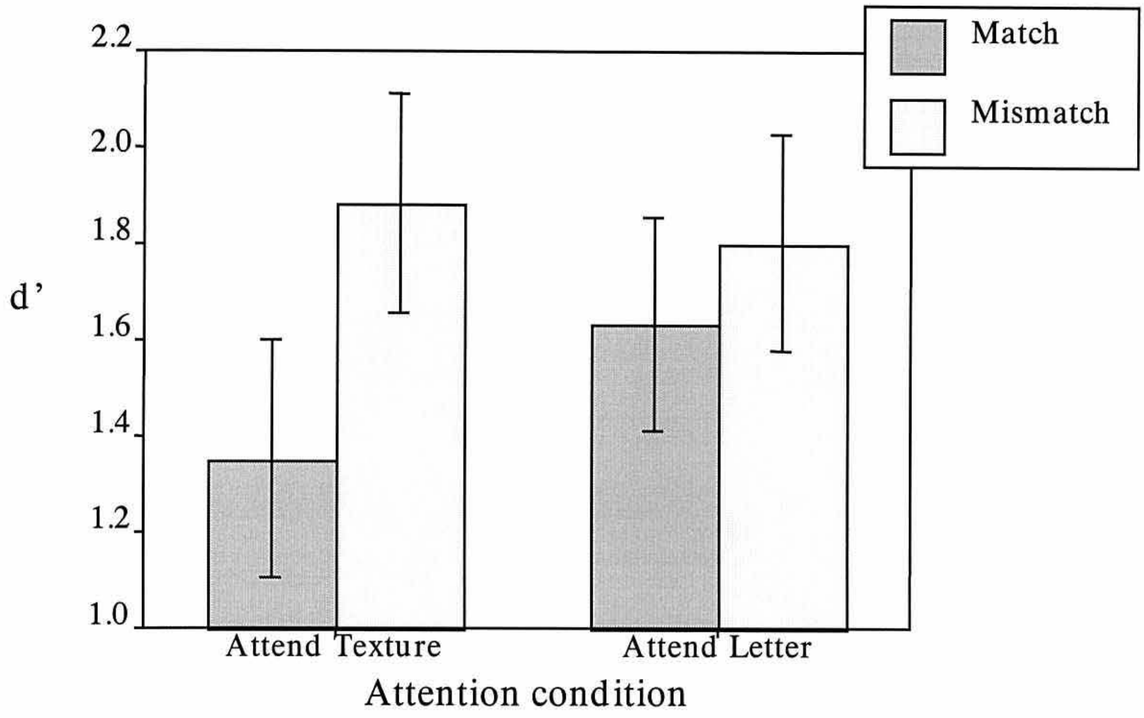

Figure 52. d' values for match and mismatch as a function of attention condition. The difference between match and mismatch in the "attend texture" condition demonstrates an orientation contrast effect. No orientation contrast effect was found for the "attend letter" condition.

The relationship between the orientation of the inducer texture and the central letter (i.e., incongruent vs. congruent) may have played a role in the orientation contrast effects reported here. Observers may have found that the orientation of the inducer texture or central letter was easier to determine when the orientations of the two were congruent vs. when they were incongruent. This difference may have had an orientation-specific effect on sensitivity to the test stimulus. However, as is evidenced by the very high proportion correct for both inducer tasks (attend texture: mean proportion correct = 97.7%, s.e. = 0.8%, attend letter: mean = 95.6%, s.e. = 1.4%), there was no apparent effect on the observers' performance in reporting the orientation of the inducer texture and letter. Any errors that were made were not systematically distributed for congruent or incongruent trial types. Similarly, there was no apparent effect of inducer congruency on test performance. Proportion "coherent" responses were averaged into congruent and incongruent trial types for each coherence value and attention condition so that psychometric functions could be examined. For example, a congruent trial was one in which the orientation of the E matched the orientation of the inducing texture. Functions for congruent and incongruent for both attention conditions were very similar and a repeated measures ANOVA using congruency, coherence and attention condition as within subjects variables showed that there were no significant differences in the functions and no significant interactions. Thus, the role of orientation congruency within the inducer stimulus is minimal.

Discussion

I have shown here that the texture orientation contrast effect can be modulated by selective attention. By presenting *identical* stimuli to observers in the "attend texture" and "attend letter" conditions and varying only the locus of selective attention, I have demonstrated that sensitivity to textures can be significantly changed and the orientation contrast effect can be "switched off".

A similar finding was recently reported in the motion perception literature in which observers viewed a transparent motion stimulus composed of two planes of motion and were cued as to which plane to report on every trial (Raymond, O'Donnell & Tipper, 1998). Depending on which plane of motion was attended during the inducer, sensitivity to a subsequent test stimulus was significantly affected in very specific ways. Results such as these, indicating that aftereffects, priming and contrast effects can be "switched off" or modulated so easily with the re-direction of attention, suggest that the aftereffects themselves may have a large attentional component. It follows then, that if contrast effects are mediated by a large attentional component that any perturbations in the focus of that attention would have large effects on the size of the contrast effect. This is consistent with results from the previous chapter.

There was an additional difference between the texture and letter inducer tasks that may have contributed to the effects reported here. For the "attend letter" task, observers had to shift their attentional focus between the inducer (attend centrally) and the test (attend globally). This local-to-global attention shift was not present in the "attend texture" condition given that both are global tasks. If observers found the attentional shift from local to global significantly more difficult than the no-shift condition, and if this difference had an effect on sensitivity to the test, the main effect for the attention condition would have been significant in the original analysis. It was not significant, indicating that task difficulty played a minimal role in this experiment.

Because there was no difference here between overall performance on the "attend texture" task and the "attend letter" task, the results suggest that the differences in sensitivity to the test stimulus were due to attentional modulation of the orientation contrast effect. However, since mean performance on the "attend letter" condition was over 95% correct, compared to 75% correct in the Shulman (1993) study, the question remains as to whether the central distractor task was difficult enough to tax attentional resources to the point where there was little left to allocate to the orientation of the texture. I feel confident that the task was taxing enough due to the short duration of the inducing stimulus (200 ms) compared to that of the Shulman study (20 sec) and the reduced contrast of the central "E". Furthermore, if there had been "left-over" attentional resources to direct to the orientation of the texture during the "attend letter" condition, then it would have acted to inflate the orientation contrast effect in the "attend letter" condition. The present results therefore tend to be conservative. Possibly, the attentional effects would have been larger had the central task been more demanding. In this case, the contrast effect may have been eliminated altogether for the attend letter condition.

The design of the present study has several advantages over that used in many perceptual aftereffect studies. This experiment was able to test sensitivity using an objective measure that is independent of subjective ratings of the duration or strength of an illusion. Furthermore, the coherence thresholding paradigm can directly quantify the magnitude of sensitivity changes and express it in stimulus units (i.e., percent coherence) something that other paradigms cannot do. This paradigm promises to be a useful tool in future examinations into perceptual aftereffects.

It has been demonstrated here that attention modulates the orientation contrast effect and may modulate aftereffects originally thought to be low-level in nature. This is not to say that aftereffects themselves are solely a high-level phenomenon, (where attentional mechanisms have been traditionally thought to mediate sensitivity), this experiment cannot conclude that, but rather that attentional modulation may also occur at lower levels of vision. Further experimentation is needed to make the distinction between whether attention is modulating low-level visual mechanisms, high-level mechanisms, or a combination of both. These results also suggest that the locus of attention needs to be considered and significantly controlled in perceptual experiments or results may be misleading. Attention may play a larger role in perceptual interactions than was first thought.

Chapter 8 - General Discussion

The series of experiments discussed in this thesis have demonstrated that human observers can take complex, heterogeneous texture information from a visual scene and efficiently distill an accurate representation of the most predominant orientation within a textured area. This ability to integrate heterogeneous texture information into an efficient representation is imperative in making sense of a visually complex world. The challenge of objectively measuring this ability was addressed by the development of the Texture Coherence Paradigm, a psychophysical procedure for easily and objectively quantifying texture integration processes.

Measuring Texture Integration Processes

In Chapter 2, I introduced the Texture Coherence Paradigm and tested it using both an orientation judgment and a signal present vs. signal absent judgment (allowing Signal Detection Theory d' calculations; Green & Swets, 1966). Observers were presented with an array of line segments in which a percentage of the lines were oriented identically (signal) and the remaining lines were oriented randomly (noise). Observers performed well on both tasks, showing that only 16 to 20% coherence was required for naïve observers to reliably determine the predominant orientation of the texture (or the presence of a coherent signal) in briefly presented displays. This level of performance is comparable to that of the analogous Motion Coherence Paradigm using similarly constructed motion stimuli (Newsome & Pare, 1988; Raymond, 1993; Williams & Sekuler, 1984). The fact that a large number of observers can complete a similar perceptual task using different behavioural responses and show very similar sensitivities suggests that there is a stable and pervasive underlying texture integration process. This process is both reliable and measurable in different observers and across different tasks.

Observers also showed that they were more sensitive to horizontal textures than to vertical textures, a finding that has also been found for motion stimuli (Raymond, 1993). This apparent orientation anisotropy, taken in isolation, suggested that orientation mechanisms vary in sensitivity as a function of signal orientation, however subsequent experiments in Chapter 3 explored this idea further and found that this was not the case. Orientation isotropy and related issues are discussed below.

Initial results from Chapter 2 were consistent with a two-stage perceptual model of the type proposed by numerous other researchers (Beck, 1982; Bergen & Adelson, 1988; Caelli, 1982; Graham, Beck & Sutter, 1992; Wilson, et al., 1992; Wilson, Wilkinson & Asaad, 1997). Generally, these models propose that local image information is initially coded by a bank of oriented band pass filters each preferentially sensitive to a specific orientation range and spatial frequency. This stage occurs fairly early in visual processing, i.e., V1 or V2, codes local information such as orientation or motion and is thought to be perceptually inaccessible (however, see below for further discussion of this point). In the second stage, the output from these filters is rectified and pooled over space via orientation analyzers with crude, wide band filtering (about 34 deg full-width at half-height for brief displays; Keeble et al., 1995). This spatial pooling of orientation information results in numerous "channels" receptive to specific orientations over the entire visual field and evidence from Glass pattern studies suggests that the second stage may be mediated by activity in extrastriate area V4 (Wilson et al., 1997).

Results presented in the present thesis are consistent with the idea that the outputs of the second stage are in turn pooled in a cooperative network and that a perceptual decision is made on the basis of this third stage. A cooperative process is consistent with previous work in texture perception (Dakin & Watt, 1997; Keeble et al., 1997; Kingdom et al., 1995) lateral masking and contour completion (Field et al., 1993; Polat & Sagi, 1993; 1994), Glass patterns (Wilson et al., 1997) and also with the perceptual model proposed by Grossberg and colleagues (Grossberg, Mingolla & Todorovic, 1989; Gove, Grossberg & Mingolla, 1995). In a cooperative system, the elements composing that system interact to create global behaviour and act to make observers more sensitive to signal embedded within noise than they would be with a purely linear perceptual system, effectively enhancing the signal-to-noise ratio. Initially, the input to the network is made up of an array of different distributed orientations and the network acts to reduce noise by increasing the excitation of like-oriented analyzers within a channel, and decreasing the excitation of other analyzers in response to the image. Excitatory connections between like-oriented analyzers and inhibitory connections between differently oriented analyzers create a spatial pattern of excitation representative of the orientations within the input. Through feedback loops, iterative repetitions cause the orientation information to be "sharpened" as information is processed further and as excitation and inhibition increase (Gove, Grossberg & Mingolla, 1995; Grossberg, Mingolla & Todorovic, 1989).

The output to the decision unit takes the form of a map of orientations in which the most common or representative orientation is the most highly activated against the background noise. This mechanism does not necessarily smooth over differently oriented singularities on a spatial basis, i.e., an area of differently oriented lines does not have to reach a certain size before they are detected (as Julesz proposed). Rather, the system determines the most common orientation and selects on the basis of frequency rather than spatial extent. This leads to the prediction that an area of like-oriented line segments composed of only two to three lines would be detected and incorporated into an orientation judgment. The data presented here is consistent in that observers were unaffected by the addition of a "contour control" algorithm to the texture generation programs that ensured a distributed signal (Experiment 2). This algorithm eliminated the possibility that signal lines would be placed in close proximity to one another (although no positional restrictions were placed on the noise distribution which included the signal orientations).

An interesting extension to the present experiments would be to create clusters and contours in texture stimuli and examine how these clusters of like-oriented line segments dispersed among differently oriented segments might change an observer's response to a texture and thereby their texture threshold. Nothdurft and colleagues have done some interesting work in this vein examining structure gradient (Nothdurft, 1985a; 1985b; 1991; Nothdurft & Li, 1985). Other interesting work in this area examines contour completions within textures (Field, Hayes & Hess, 1993; McIlhagga & Mullen, 1996) and lateral masking among line segments and gabors (Polat & Sagi, 1993, 1994).

Another interesting prediction that results from the perceptual model described above is that observers should require a relatively extended duration in which to integrate orientation information over space. Observers should require sufficient time for orientation information to be processed through the series of iterative loops that "sharpen" and reduce noise in the image (Gove, Grossberg & Mingolla, 1995; Grossberg, Mingolla & Todorovic, 1989). This extended duration requirement is consistent with the perception of other complex images such as Glass patterns and illusory contours. Stimulus duration was investigated in Chapter 4.

A further question regarding texture integration processes is at what stage of processing texture integration takes place. Julesz and Beck would no doubt disagree on this point, but more recent evidence from Glass pattern studies suggests that pattern integration takes place after full wave rectification (Wilson, Wilkinson & Asaad, 1997). Investigating what featural information is available to integration processes should elucidate at what stage processing occurs relative to other perceptual processes such as feature coding. Since the features of an object, such as colour, form, texture and motion, are thought to be coded somewhat separately at early stages in processing and are subsequently re-combined to form a more complete representation of that object later in processing (Duncan & Humphreys, 1989; Humphreys & Muller, 1993; Prinzmetal, 1981; Treisman & Gelade, 1980; Treisman & Gormican, 1988; Wolfe, Cave & Franzel, 1989), the question asked here was does this re-combination occur before or after texture integration processes are complete? Chapter 5 investigated this idea.

The Texture Coherence Paradigm has been useful for quantifying the perception of textural features that vary along a continuum as well as for examining the perceptual combination of features (see Chapter 5). Given its simplicity, the paradigm could also be used to investigate other features that vary along a continuum such as depth, form, colour, or any combination of these features as it has in motion perception. In addition, by using a coherent vs. random judgment, d' calculations can be made resulting in a criterion-free measure of sensitivity, according to the Theory of Signal Detection (Green & Swets; 1966). Some specific avenues for future research could include testing the perception of depth with this technique. Observers could be presented with a three-dimensional cloud of dots and asked whether they perceive the dots to be mostly in a single plane of depth or whether they are randomly distributed throughout many planes, i.e., random vs. coherent. The percent coherence of dot positions in a given plane could be easily varied. A more complex version of this idea could consist of presenting observers with line segments oriented in three dimensions. These types of stimuli, though quite complex, would more closely emulate real-world perception while the experimenter would easily retain control of stimulus factors. Texture Coherence promises to be a useful tool for future research.

Texture Integration is Isotropic but Shows Global Precedence Effects

The series of experiments presented in Chapter 3 addressed whether perceptual orientation channels are equal in sensitivity to one another as a function of orientation. As seen in the Oblique Effect, contrast sensitivity to oriented gratings is better for cardinal orientations (vertical and horizontal) than for obliquely oriented gratings. I investigated here whether this anisotropy of sensitivity persists for texture integration mechanisms. A related issue is whether Global Precedence affects the perception of composite textures. Results from Chapter 2 suggested that observers were differentially sensitive given different signal orientations, i.e., observers were more sensitive to horizontal signals than to vertical when presented with square texture patches. Often, the form of the global figure itself affects an observer's perception of elements making up that global figure. By using differently shaped texture patches and by varying the responses required of observers, I determined that orientation channels are comparable in their sensitivity for texture integration tasks, i.e., channels are orientationally isotropic, for stimuli with neutral global shapes. However, Global Precedence does play an important role in that sensitivity for certain orientations was significantly affected if there were global contours congruent to those orientations and not others.

This proclivity towards Global Precedence indicates that along with local orientation detectors, our visual systems are sensitive to a much larger scale of orientation that encompasses the entire texture patch. It also suggests that this greater scale can significantly alter sensitivity at smaller scales of analysis, so much so that coherence thresholds for certain orientations were affected more than others. Indeed, there is evidence suggesting that area V2 replicates the structure of area V1 but at a larger spatial scale (Kisvárday, Tóth, Rausch & Eysel, 1995). The representation of orientation at larger scales in V2 or higher areas may mediate an orientation-specific Global Precedence effect for textures. Specifically, in the present studies observers were much more sensitive to textures oriented horizontally over others when there was a congruent global contour present, e.g., with square patches. However, there was no horizontal advantage when viewing octagonal patches that consisted of two global contours parallel to each signal orientation, including obliques, and a much shorter horizontal contour than in the square patches.

Interestingly, when using square patches there was no concurrent advantage for vertical signal lines as is demonstrated in the classic oblique effect, just an advantage for horizontal. This is a puzzling result. Could it be that there is a Global Precedence effect operating here, but that it is only for the horizontal plane, and not for vertical? While this may seem counterintuitive, our every day environment may hold a clue. Since the horizon is always horizontal when we are standing upright and moving through the environment, scanning motions parallel to that plane may be over-learned, perhaps giving rise to an advantage for the horizontal orientation. Further, very common tasks such as reading consists of horizontal scanning motions that are required (in many languages) when reading a passage. These suggestions are also consistent with the motion perception literature in which observers are much more sensitive to motion in the horizontal plane than to vertical motion.

A logical extension to this chapter would be to induce global precedence effects and examine their characteristics. In the experiments presented here, the focus was on determining what was causing the sensitivity differences and controlling those factors to minimize the effects. Some interesting further investigations could focus on inducing larger global precedence effects and examining what stimulus parameters maximize and minimize the precedence of the outer contours. Pilot studies (not reported in this thesis) conducted using rectangular texture patches resulted in stronger global precedence effects than that presented here (using square patches). These preliminary results indicated that manipulating the ratio between the length and width of the rectangular texture area would change the size of the precedence effect. The size and density of the texture patch may also make a difference to global precedence, although no density effect was found within the very small range tested in Chapter 5, Experiment 8. Further experiments investigating more closely the properties of global precedence and texture density would be interesting.

In terms of orientation isotropy, these experiments indicate that the orientation-specific "channels" that code heterogeneous textures are isotropic in nature. This suggests that the mechanisms that code orientation in low contrast spatial frequency gratings, i.e., those mechanisms that mediate the oblique effect, are distinct, or are at a separate stage, from those that code orientation in heterogeneous texture stimuli. In the interest of neural efficiency, local orientations in textures are most likely coded by the mechanisms that also code orientation in grating patterns. This means that at some point in the visual hierarchy, the advantage seen for cardinally oriented gratings over obliques is overcome to create an isotropic system at the integration stage.

Integration Requires Extended Stimulus Durations

Naïve adult observers required approximately 180 to 200 ms of exposure to accurately judge the orientation of a heterogeneous texture. If indeed texture integration mechanisms are mediated by a cooperative-competitive feedback system as discussed above, this extended period of time would be needed for the system to complete the iterative repetitions required to extract signal from noise and to determine the most prevalent orientation present. This is consistent with perception of other complex stimuli such as Glass patterns in which the spatial relationships between component dots must be calculated before perceptual integration into a coherent Glass pattern can be completed (Glass, 1969; Glass & Perez, 1973; Glass & Switkes, 1976). Similarly with illusory contours, a line defined only by the endpoints of other lines must be perceptually constructed using spatially distributed information, which takes time (Gove, Grossberg & Mingolla, 1995). The extended durations seen here for texture integration is consistent with other texture integration tasks in which observers perform much better given a longer exposure to the stimulus (Dakin & Watt, 1997; Keeble et al., 1995; Keeble et al., 1997; Kingdom et al., 1995).

Another possibility is that focused selective attention is required for observers to maximize their effectiveness at integration tasks. With brief exposures to a stimulus, observers are able to complete the orientation judgments but require higher levels of coherence, i.e., much more signal strength, to do so and therefore show elevated coherence thresholds. With longer durations, selective attention can be focused on the task and performance improves, however observers require at least a 180 to 200 ms exposure to show this level of sensitivity. Relatively extended durations are required for the build up of inhibition and excitation within the attentional system. Evidence supporting an attentional explanation comes from Chapter 6 in which the texture contrast effect was not present until the inter-stimulus-interval between the inducer and test was at least 200 ms. Further evidence from Chapter 7 supports the attentional hypothesis in that without attention directed at an inducing texture, the texture contrast effect is effectively "switched off". The effectiveness of selective attention in perceptual tasks was also demonstrated in Chapter 5 in which features of textures that varied along different continua were integrated into a more complete representation with the assistance of attention.

It may be that the iterative competitive/cooperative loops discussed above ***are*** the initial stages of the attentional mechanism. Grossberg and colleagues suggest that the pre-attentive perceptual grouping processes that occur at fairly low levels of processing (i.e., LGN, V1, V2) are themselves their own attentional primes (Grossberg, 1998; Grossberg, in press). Through "folded feedback" circuits, excitation and inhibition among the layers of LGN, V1 and V2, generate perceptual representations of images and can explain not only "normal" perception but also illusions such as illusory contours. Thus what we have been calling attention may be a collection of properties already built into the perceptual system.

While the performance of observers was very individualized with respect to the stimulus duration required, as is demonstrated with two observers in Experiment 7 (mask and no mask conditions), perhaps other, more systematic factors are involved. The age of the observers, for example, may significantly affect thresholds. In the present experiments, observers were all between the ages of 18 and 45, but children between the ages of 7 and 10 have been tested with this task and perform very well, often displaying thresholds of 10% coherence with an exposure duration of less than

500 ms (pilot data, not reported here). Thresholds of adults, by contrast, seem to asymptote at 16 to 20% coherence and longer stimulus durations seem to make little difference. A more systematic investigation of the role of observer age would be interesting, to see if it is consistent with the findings that children are also much more sensitive to motion integration tasks compared to adults (Raymond & Sorensen, 1998). Similarly, elderly individuals may have a "slowed" integration/attentional system compared to younger adults and may require even longer stimulus durations with which to make an orientation decision. A more simple texture orientation task, i.e., 80 to 100% coherence, could also be presented to these groups of individuals to distinguish between a generalized loss of acuity and a deficit in integration processes. If present, such a deficit would no doubt be detrimental to complex tasks in every-day life such as driving.

The evidence reported here, that naïve adults require relatively extended stimulus durations to effectively integrate orientation information, indirectly supports the idea of an extra-striate contribution to integrative texture perception. Interestingly, it has been proposed that higher levels of extrastriate cortex, possibly V4 of the infero-temporal visual pathway, mediate other integrative tasks, such as Glass pattern perception (Wilson et al., 1997). This area is thought to mediate functions that require perceptual integration over space and feeds into higher-level areas involved in form vision such as the perception of faces. More in-depth investigations of required stimulus durations in normal observers as well as patient populations might provide some more direct evidence regarding where in the brain texture integration processes take place. Human neuropsychological patients have shown a dissociation of deficits in tasks involving global integration of local information. Rentschler and colleagues (Rentschler, Treutwein & Landis, 1994) present a patient (KD) with an infero-medial occipito-temporal lesion to the right side who displays difficulties in global visual tasks but performs well on local micropattern tasks. This dissociation of function suggests that ventral extra-striate areas may mediate global integrative tasks. Further studies such as these could provide converging evidence as to where in the brain integration tasks such as texture integration take place.

Features of Textures are Bound Incompletely Without Selective Attention

In this series of experiments, I examined at what point in perceptual processing texture integration takes place in relation to featural combination. To this end, I tested observer's abilities in selecting line segments on the basis of one feature and reporting another feature. I was interested specifically in the combination of polarity and orientation and whether texture integration processes have access to polarity information regarding local individual elements, i.e., does texture integration take place before or after featural binding at the element level has been completed. If featural integration is completed before texture integration processes, then observers should have no problems selecting a subset of line segments on the basis of polarity and then reporting orientation. If, however, featural binding takes place after or at a similar stage to texture integration, observers should have problems with these tasks. I presented observers with textures made up of lines of two polarities and asked them to determine the orientation of a subset of those lines selected on the basis of polarity. Interestingly, observers seemed unable to effectively do these tasks suggesting that polarity information is unavailable to texture integration mechanisms. Further investigations indicated that polarity information was not totally lost to perception (as was hypothesized by Wilson et al., 1997, with a similar task), but rather the two features of polarity and orientation were incompletely bound, and that selective attention could assist in this selection task.

It seems that texture integration occurs at a similar stage as does feature integration. Researchers have proposed that features are bound back together with the help of selective attention, possibly in the inferotemporal (IT) area of the brain

(Treisman & Gormican, 1988; Wilson et al., 1997). Evidence presented here is consistent with that idea.

The idea of using competing distractors allowed the demonstration of exactly how much extraneous and distracting information the attentional system could filter out. It took about 60% competing orientation information to reduce observers' performance to chance when viewing a 20% coherent stimulus. This objectively quantifies the effectiveness of selective attention in this situation and could be applied to other situations to quantify the strength of attention independent of the properties of the features used.

Texture Contrast Effects Indicate that Integration Persists in Time

Past research has demonstrated that integration tasks persist in time and can affect perception of and sensitivity to subsequent stimuli causing perceptual aftereffects and illusions, for example, aftereffects of tilt, size and motion. I tested whether briefly viewing an oriented inducing texture stimulus would significantly affect an observer's sensitivity to a subsequent variable coherence texture. Results showed a texture orientation contrast effect in which sensitivity to a texture matched in orientation to the inducer was reduced and sensitivity to a mismatched texture (both 90° and 45° differences in orientation) was increased. This contrast effect was not present 100 ms after the inducer but persisted for 200 to 700 ms after the offset of the inducer. Due to the prolonged presence of the effect, an attentional component was suggested.

To reduce the effects of perseveration (MacKay, 1987; Mussler & Hommel, 1997) and to optimize available perceptual resources, observers may have been initially coding the features of the first stimulus and then applying inhibition to that representation after initial coding was complete so that the same object would not be coded repeatedly (Raymond et al., 1998; Kanwisher, 1987; Posner & Cohen, 1984;

Tipper, 1985). This application of inhibition would allow new and possibly novel stimuli entering the visual field, i.e., the test, to be coded immediately and to be discriminated from the inducer. Those stimuli that are different to the preceding inducer would have an advantage over those more similar. This would result in an elevated threshold for test stimuli oriented the same way as the inducer (match) and a reduced threshold for test stimuli that are mismatched in orientation to the inducer; consistent with the results reported here. This system makes sense ecologically in that an animal that is more sensitive to detecting and identifying new, novel objects entering it's field of view, whether predator or prey, would have a distinct advantage over those animals that either repeatedly code objects already in the field of view, or show no discrimination between these new objects requiring identification and old ones that are already coded and identified.

An interesting prediction that comes out of this hypothesis is that all features of a previously coded stimulus should have inhibition applied to them, whether they are relevant to the present task or not. For example, the inducer texture had many other features that were not directly relevant to the task at hand such as outer contour shape, luminance, colour, size, and density, but that nonetheless may have inhibition associated with them. If observers were asked to judge a feature on another seemingly irrelevant dimension, this inhibition may be measured. Such inhibitory effects between successively presented stimuli are demonstrated in studies examining inhibition of return (Posner & Cohen, 1984), repetition blindness (Kanwisher, 1987) and negative priming (Raymond, O'Donnell & Tipper, 1998; Tipper, 1985; Tipper & Driver, 1988; Tipper, MacQueen & Brehaut, 1988).

Observers also experienced an illusion of coherence in noise stimuli following a brief exposure to a coherent inducer. This caused observers to respond that they had actually seen the opposite orientation in the test than was presented in the inducer. This effect is analogous to the motion aftereffect in which observers view a moving stimulus for a given duration and then while viewing a subsequent stationary (or random noise) stimulus, experience the illusion of motion in the opposite direction. An interesting aspect of this texture coherence illusion was that, unlike the motion aftereffect, it was fairly short-lived. This illusion was only present with the shortest ISI between inducer and test (100 ms) and observers viewing stimuli with longer ISIs did not experience the illusion. This may be a side effect of the coding of the inducer stimulus and the application of inhibition to its representation. Possibly, it takes at least 100 ms after the offset of the stimulus to code its features and to apply inhibition to that representation. This would explain why the observers viewing stimuli with ISIs longer than 100 ms did not experience this coherence illusion. The relatively long exposures required for more complete featural coding is also consistent with results from Chapter 4 in which I investigated stimulus durations.

Some interesting experiments could be done to extend the understanding of texture contrast effects. How similar do the inducer and test have to be to produce a contrast effect? Can an inducer made up of a sinusoidal grating induce a contrast effect in a line segment texture? Also, inducers most likely do not have to be 100% coherent to produce contrast effects. If not, what is the minimum coherence required to produce these effects? Is there a gradient of contrast strength depending on the coherence of the inducer? Do inducers and tests have to be in the same spatial location to create contrast effects? While there most likely would be apparent motion perceived between the two if they were successively presented in different spatial locations, it would be interesting to see if a contrast effect could be demonstrated with the inducer at one location and the test at another. This would convincingly rule out any adaptation explanations for contrast effects.

Texture Contrast Effects can be "Switched Off" by Redirecting Selective Attention

Chapter 7 demonstrated that the texture orientation contrast effect can be "switched off" by redirecting selective attention away from the inducing stimulus.

Observers were presented with a 100% coherent inducer texture followed 300 ms later by a partially coherent test texture. A low-contrast letter that varied in orientation was displayed at the center of the inducer texture. Observers were asked either to attend to the central letter and report its orientation ("attend letter") or attend to the inducer texture and report its orientation ("attend texture"). All observers then reported whether the test texture was random or coherent. Results showed a significant orientation contrast effect when attention was directed at the inducer texture, and a contrast effect that was markedly reduced in size when attention was directed at the central letter. It seems that merely giving observers a task that redirected their attention away from the inducer was sufficient to switch off the aftereffect.

Judging by the 96% correct performance on the letter inducer task, observers did not find the orientation judgment of the low contrast "E" to be very taxing. Perhaps a harder distractor task (i.e., lower contrast or smaller "E") would have eliminated the contrast effect in the attend letter condition altogether. Observers may have had sufficient attentional resources left over to partially attend to the inducer texture. A harder inducer task for the attend letter condition however, may have caused texture thresholds to be generally higher than for the attend texture inducer task, making results harder to interpret. In the present experiment, there was no significant difference in performance between the attend letter and attend texture tasks.

The fact that the attend letter task required a central locus of attentional fixation and the attend texture task required a distributed locus of attention may be an issue for some, despite the fact that there was no difference in performance between these conditions. The purpose of this experiment was to see the effects of a re-direction of selective attention. The logical place to do this re-direction without requiring eye movements was at the centre of the texture patch. Had the attend letter condition been some other task, for example, attend to the colour of the outer border of the inducing patch, there would have been too many differences between the attention conditions. In the present experiment, observers were asked to do the same task in both attention conditions, i.e., determine orientation, and the only condition that changed was the locus of attention. I feel that this control over task differences outweighs the concerns about global vs. local focus of attention.

Some Speculation

The underlying mechanism that mediates heterogeneous texture perception is one that is of great advantage in making sense of the crowded and busy perceptual world. Such a mechanism can detect patterns of similarities, i.e., collinear line segments, which are separated by relatively large spatial extents and super-imposed on noisy backgrounds. We do know that the brain is extremely good at imposing perceptual structure on both systematically modulated and un-modulated displays, thus supposing the presence of a well-developed spatially integrative system makes sense. In the interest of neural efficiency, the mechanism underlying texture perception would no doubt mediate other forms of pattern detection. Global integration of a scene should occur at a stage at which pattern processing from many different "modules", i.e., motion, colour, form, depth, is fairly close to completion, in order for the integrated percept to be informative. However, initial percepts of stimuli are also informative, though are much more prone to errors in feature binding. Proposing such a late stage for integration naturally begs the question does attention assist in integration? The answer in light of the present evidence is yes; attention plays a large role in integration and feature binding processes.

Traditionally, attentional modulation of perception has been thought of as something reserved for extrastriate areas such as V4, IT, or V5/MT. However, recent evidence has been reported that there is significant attentional manipulation as early in the visual hierarchy as V1 (Grossberg, 1998; Grossberg, in press; Grossberg & Raizada, 1999; Motter, 1993; Roelfsema, Lamme & Spekreijse, 1998; Watanabe, Sasaki, Nielsen, Takino & Miyakawa, 1998; Zipser, Lamme, & Schiller, 1996), an area that was until relatively recently thought to be perceptually inaccessible. This is fascinating research and while it is still in it's infancy, seems to provide some fairly strong evidence supporting these claims. The ideas presented in this thesis are consistent with the idea of attentional modulation of striate area V1, though texture integration processes themselves most likely take place at a later stage. Attentional modulation of V1 would cause a competitive/cooperative integrative system to be more efficient than it would be without such modulation. Essentially, the steps previously thought to be mediated by extrastriate, attentionally modulated areas can now be "brought forward" to an earlier processing stage thus completing integration more rapidly than first thought.

This idea also predicts that an integrative system would also be more efficient in terms of the volume of processed information. If processes in V1 are able to use horizontal connections to reduce the volume of information processed by higher stages, this would not only be more efficient in terms of temporal processing but also more efficient in terms of attentional load. In other words, a smaller attentional load for higher visual areas results in less "effortful" perceptual tasks and more attention left over for other tasks. This, no doubt, would be an incredible advantage.

Final Thoughts

The texture coherence paradigm used here to quantify the perception of heterogeneous textures has proven to be a useful and easily implemented tool for perception research. There are many more questions still to be answered with regards to texture integration processes and I hope that the research reported here has helped to generate some direction as well as a useful tool for future investigations.

Appendices

Appendix A - Interpolation of Coherence Thresholds

In most experiments presented here, I plotted proportion correct as a function of coherence resulting in a classic psychometric function. I estimated observers' thresholds by converting proportion correct to Z scores (a linear trend), fitting a line of least squares to the Z score function and calculating the coherence value corresponding to a criterion proportion correct score. This appendix describes the interpolation procedure in detail and is represented graphically in Figure 53.

Depending on the experiment, criterion was defined differently. In Experiments 1, 2, 4, 5 and 7-12, criterion was the point at which observers were just able to correctly determine texture orientation (i.e. 75% correct was chosen because it is significantly above 50% baseline performance for a 2AFC). In Experiments 3, 6 and 13-15, I use a "coherent" vs. "random" judgment and so criterion was the point at which observers responded "coherent" 50% of the time (i.e. 50% "coherent" judgments is significantly above the baseline of 0% "coherent" responses). The example demonstrated here uses a criterion of 75% correct for a 2AFC orientation judgment, however, the steps to interpolate thresholds are the same for the 50% case, only the criterion value needs to be changed.

Conversion to Z-scores:

To accurately interpolate coherence threshold, I needed to transform the typical S-shaped psychometric function into a linear trend that could be expressed in a simple equation. When proportion correct scores are converted into Z-scores (by means of a look up table) and plotted as a function of stimulus values (i.e., coherence), the resulting function is linear and is expressed simply by the equation of a line.

Line of Least Squares:

I then fitted a line of best fit through the Z-score function using the method of least squares in which the squared deviations from the fit are minimized. This line can be defined simply by:

$$Y = mx + b$$

where Y is the proportion correct score, x is the coherence value, m is the slope of the line and b is the intercept of the line when x is 0. I wanted to interpolate the coherence value corresponding to the 75% correct point and thus, the equation was re-arranged to solve for x:

$$x = \frac{Y - b}{m}$$

Substituting the criterion proportion correct score of .75 for Y, and substituting b and m as defined by each observer's line of best fit, I calculated the coherence value that corresponded to the 75% correct point on each observers' function. In all cases, the fit resulting in the highest r^2 value was used.

Interpolating proportion correct:

In Chapter 5 (Polarity), I interpolated the expected proportion correct score given a coherence value and based on an observer's psychometric function. This was done by using the equation of the line of best fit to solve for Y:

$$Y = mx + b$$

and substituting the required coherence value into x. Again, the fit that maximized r^2 was used in all cases.

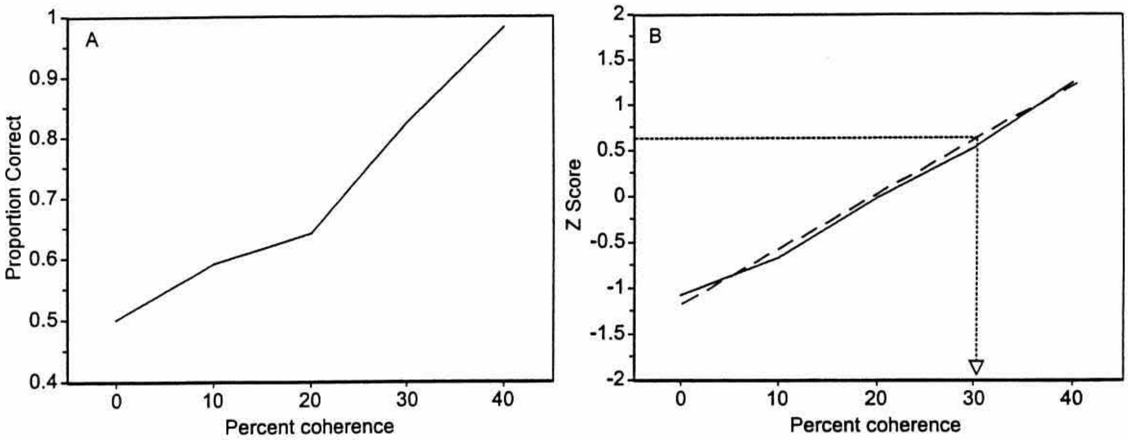

Figure 53. A. Psychometric function for a 2AFC orientation judgment. B. The psychometric function in A has been converted to Z scores using a look up table. The dotted line represents the line of least squares fitted to the Z score function. From this line, the coherence value that corresponds to the 75% correct value (Z = .67) can be interpolated. In this example, coherence threshold is approximately 30% coherence.

Appendix B - Signal Detection Theory

In classical psychophysical experiments, observers' thresholds are by nature a function of both stimulus detectability and the location of the observers' subjective criterion. Thus, supposed changes in observer sensitivity is confounded with changes in criterion. d' calculations, from the Theory of Signal Detection, separate changes in sensitivity from changes in observer criterion. This appendix describes Signal Detection Theory and explains d' calculations.

Signal Detection Theory is based on the assumption that all physical stimuli are perceived against a background of activity, or noise, which can be either external to the observer (i.e., from the environment) or internal to the observer (i.e., the spontaneous firing rate of the nervous system). A stimulus must be strong enough to exceed this background noise to be detected by the observer. Thus we can plot the noise distribution (N, signal absent) and the signal plus noise distribution (SN, signal present) as in Figure 54. As the difference in the means of the two distributions gets larger, i.e., as the two distributions move father apart, the observers' discrimination gets easier because the stimulus is sufficiently strong to rise above the background noise. As the two distributions move closer together (as they would if the signal was very weak), the discrimination becomes more difficult until the point is reached at which the Signal + Noise distribution (SN) cannot be distinguished from the Noise distribution (N). At this point, the observer would perform at chance on the perceptual task.

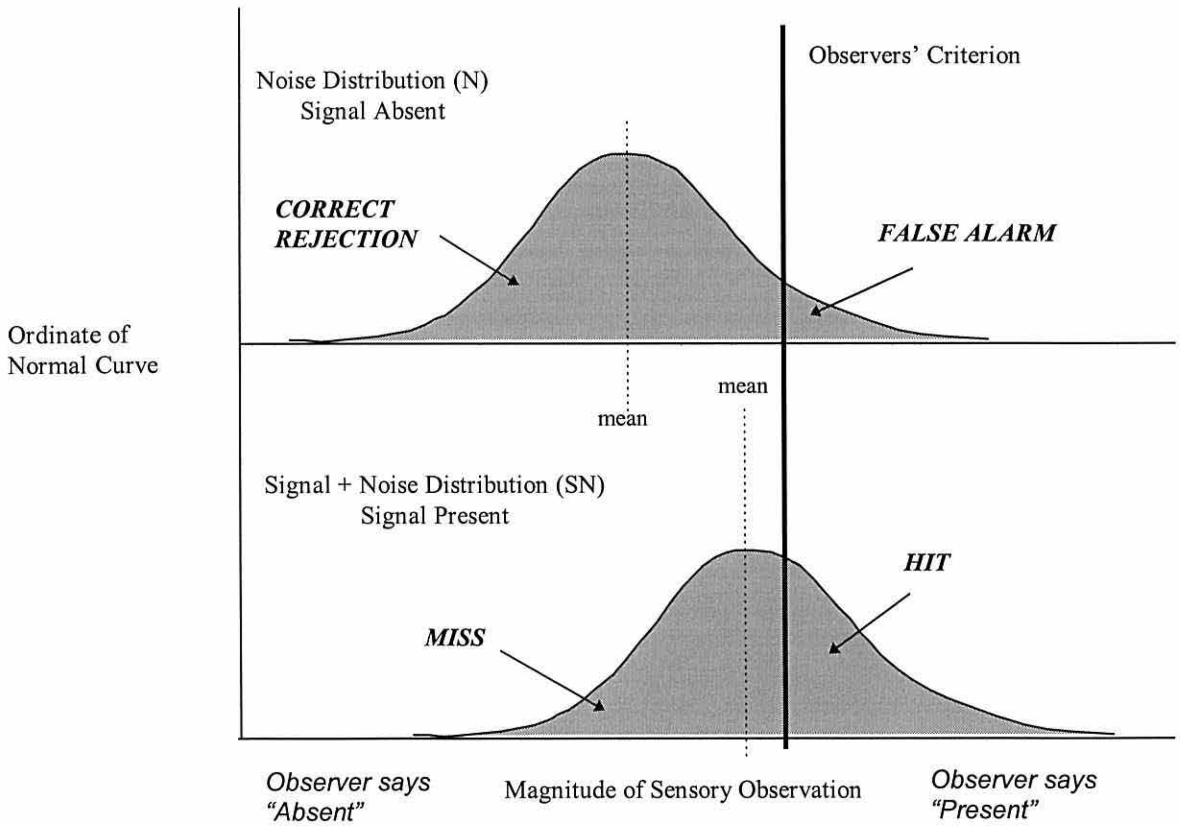

Figure 54. Signal Absent (N) and Signal Present (SN) distributions. As the difference in the means of the two distributions gets larger, the more detectable the Signal Present situation is from the background noise.

The value of d′ is a measure of detectability and is equal to the difference between the means of the SN and N distributions. Because the location of the SN distribution in relation to the N distribution is entirely a function of stimulus intensity and properties of the sensory system, d′ is a pure index of stimulus detectability that is uncontaminated by the location of the observer's criterion (Green & Swets, 1966).

There are four possible situations that arise from the combination of the observer's response and the condition of the stimulus as summarized in Table 2. A Hit occurs when an observer correctly reports that the signal was present. A False Alarm occurs when an observer responds "signal present" when the signal was not actually present. A Miss occurs when the observer reports "signal absent" when in fact the signal was present. A Correct Rejection occurs when the observer correctly reports that the signal was not present.

Table 2. Definitions of response conditions for Signal Detection Theory

		Observer Response	
		"Signal Present"	"Signal Absent"
Stimulus Condition	Signal Present	***HIT***	***MISS***
	Signal Absent	***FALSE ALARM***	***CORRECT REJECTION***

The false alarm rate is defined as the rate at which the observer reports that there was a signal present when in fact there was not. In the experiments I present here, I used a "coherent" vs. "random" judgment, therefore the false alarm rate was the proportion of "coherent" responses to a stimulus with 0% coherence.

d′ is calculated by subtracting the false alarm rate (to a neutral stimulus) from 1.0 and converting this value to a Z-score. This Z-score represents the area under the normal distribution curve between the distribution mean and the observer's criterion. In the example shown in Figure 55, this value is 2.05.

1.0 – p (false alarm) -> Zn
Zn = ***2.05***

The hit rate (for a given coherence value) is also subtracted from 1.0 and converted to a Z-score. The resulting value is 0.39.

1.0 – p (hit) -> Zsn
Zsn = ***0.39***

Since d′ is the difference between these distributions, d′ is equal to:

d′ = Zn – Zsn
d′ = 2.05 – 0.39
d′ = ***1.66***

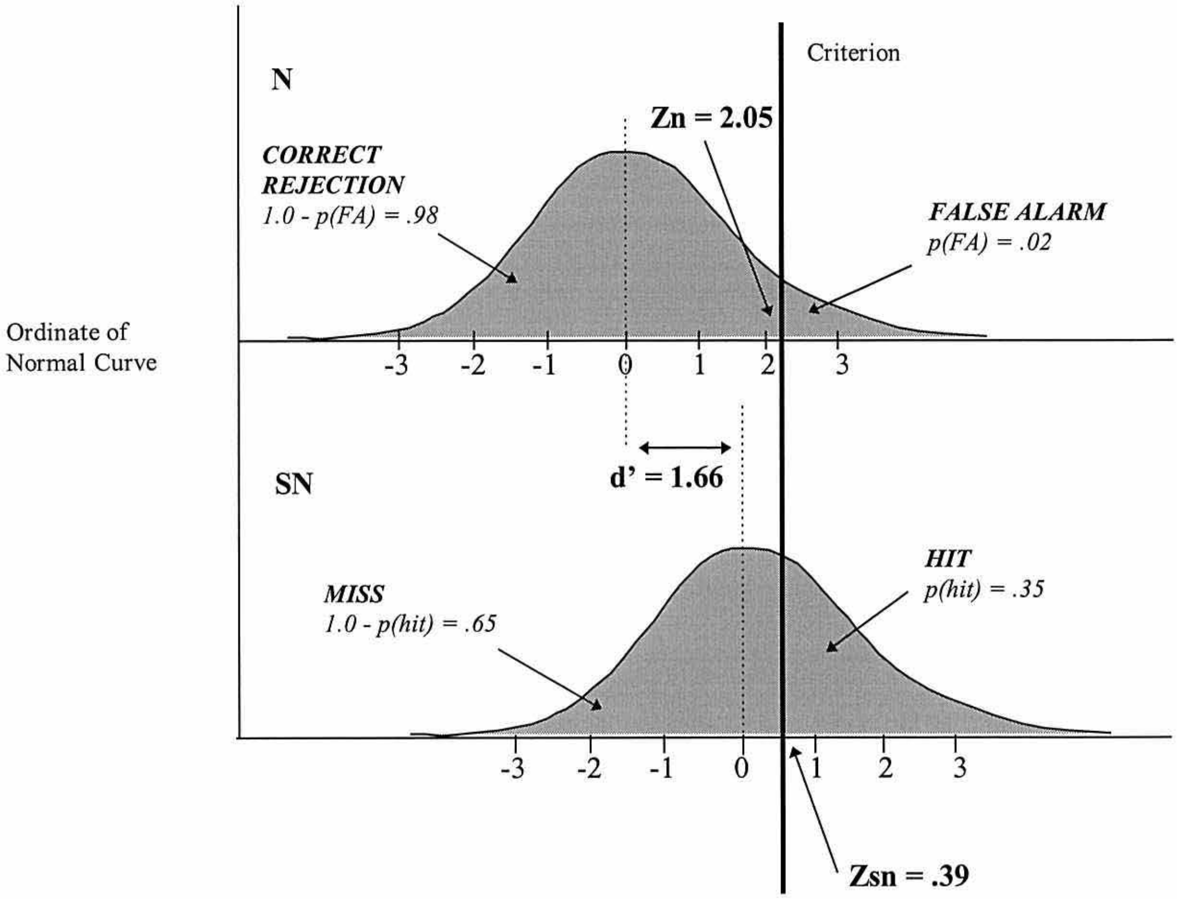

Figure 55. An illustration of the distributions of noise (N) and signal plus noise (SN) expressed in Z-scores. The location of the criterion on the noise distribution is found by subtracting the false alarm rate from 1.0 and converting this value to a Z-score. The location of the criterion on the signal plus noise distribution is found by subtracting the hit rate from 1.0 and converting this value to a Z-score. The value of d', a measure of the observer's sensitivity to the signal, is found by subtracting Zsn from Zn (d' = Zn - Zsn). From Green & Swets (1966).

The result of these calculations is a function of d' scores that resembles a psychometric function but is independent of the observer's criterion as illustrated in Figure 56. Consider the following example. An observer views a series of stimuli repeatedly and her criterion changes with each of the three series presented, however, her underlying sensitivity does not change. In the first series, she responds normally with a false alarm rate of .10. For the second and third series, she becomes less conservative and her false alarm rates increase to .30 and .40 respectively. Panel A of Figure 56 illustrates the psychometric functions that would result from this situation.

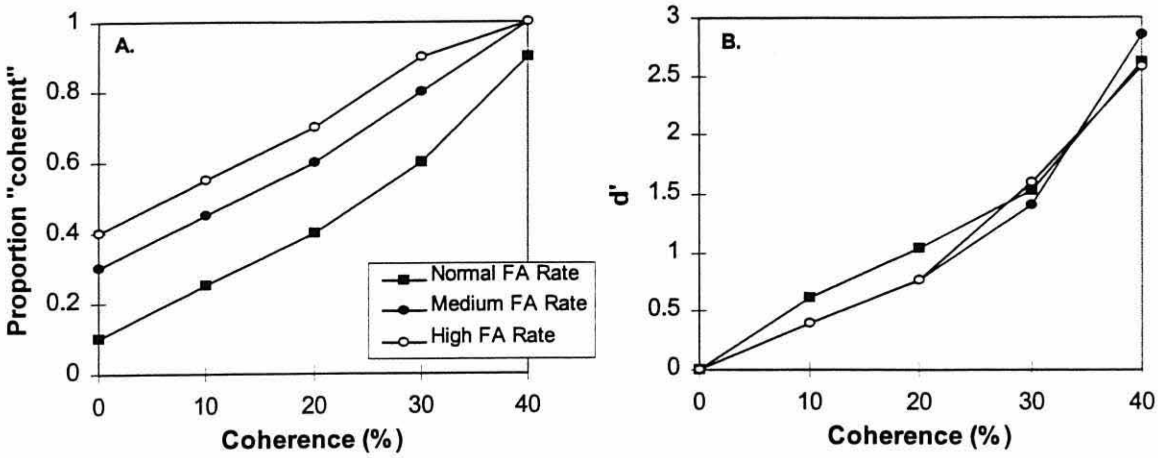

Figure 56. Changes in criterion do not necessarily mean changes in sensitivity. d' calculations allow the distinction between shifts in sensitivity vs. shifts in criterion. Panel A illustrates proportion "coherent" psychometric functions for three subjective criteria. Panel B illustrates the similarity of the d' functions for those criteria.

d' scores are calculated from the false alarm rates of each function and proportion "coherent" responses and are plotted in Panel B of Figure 56. As can be seen, there is little difference in the d' functions among the criterion conditions. Since d' scores are a relative value comparing false alarm rates and hit rates, they can effectively "filter out" changes in criterion and reflect only underlying changes in sensitivity.

Appendix C - Glossary of Terms

Coherence - The percentage of signal lines oriented in the signal orientation amongst noise orientations. For example, 20 vertical signal lines amongst 80 randomly oriented noise lines is 20% coherence.

Contrast Effect - A temporal interaction effect among successively presented stimuli that shows a decrease in sensitivity for stimuli preceded by a similar stimulus and an increase in sensitivity for non-similar stimuli. This applies to many domains, i.e., tilt, size, motion, attention, and not just to visual aftereffect paradigms.

Element - A texture element is the basic unit of a texture. A texture is composed of repetitions of elements, for example, line segments or dots.

Features - Elements have features which define them such as orientation, colour, size, shape, depth or motion. An element can be defined by numerous features, for example, a signal element can be white, vertical, and 10 pixels long X 1 pixel wide.

Heterogeneous - A texture that is made up of featurally variable elements. For example, a patch of randomly oriented line segments is a heterogeous texture.

Homogeneous - A texture that is made up of elements that are featurally identical or that vary only slightly. For example, a patch of line segments oriented vertically or all within a few degrees of vertical is a homogeneous texture.

Integration - A perceptual task that requires spatial combination of local information in order to perceive a global structure or to determine a global feature. Typically, local featural information is too unreliable or not sufficiently informative to allow a correct judgment in an integration task.

Noise - A subset of line segments within a texture that is randomly oriented in one of 16 possible orientations. Noise elements are randomly positioned within the texture patch.

"**Noisy**" - see Heterogeneous.

Patch - A patch is an area of texture composed of repetitions of texture elements. In most of the experiments presented here, the texture patch is an 11 X 11 square array of elements.

Pre-attentive - A task that requires little focused scrutiny or effort to complete. A pre-attentive task is thought to be mediated by relatively low-level visual mechanisms and can be completed quickly and accurately with brief stimulus durations.

Signal - A subset of line segments within a texture that are oriented in the same signal orientation (usually horizontal or vertical). Signal elements are randomly positioned within the texture patch.

Texture - An area of space composed of and defined by numerous smaller objects or elements. See also heterogeneous and homogeneous.

Appendix D - Table of Figure Captions

Figure 56. Changes in criterion do not necessarily mean changes in sensitivity. d′ calculations allow the distinction between shifts in sensitivity vs. shifts in criterion. Panel A illustrates proportion "coherent" psychometric functions for three subjective criteria. Panel B illustrates the similarity of the d′ functions for those criteria. 201

Appendix E - Table of Table Captions

References

Annis, R. C. & Frost, B. (1973). Human visual ecology and orientation anisotropies in acuity. *Science, 182,* 729-731.

Anstis, S. M. & Moulden, B. P. (1970). Aftereffect of seen movement: Evidence for peripheral and central components. *Quarterly Journal of Experimental Psychology, 22*, 222-229.

Baker, C. L., Hess, R. F. & Zihl, J. (1991). Residual motion perception in a "motion blind" patient, assessed with limited lifetime random dot stimuli. *Journal of Neuroscience, 11(2)*, 454-461.

Barlow, H. B. & Brindley, G. S. (1963). Inter-ocular transfer of movement after-effects during pressure blinding of the stimulated eye. *Nature, 200*, 1347.

Barlow, H. B. & Hill, R. M. (1963). Evidence for a physiological explanation of the waterfall illusion and figural aftereffects. *Nature, 200,* 1345-1347.

Barton, J. J. S., Sharpe, J. A., & Raymond, J. E. (1996). Directional deficits in pursuit and motion perception in humans with unilateral cerebral lesions, *Brain, 119*, 1535-1550.

Barton, J. J. S., Sharpe, J. A., & Raymond, J. E. (1995). Retinopic and directional defects in motion discrimination in humans with cerebral lesions, *Annals of Neurology, 37*, 665-675.

Beck, J. (1966a). Perceptual grouping produced by changes in orientation and shape. *Science, 154*, 538-540.

Beck, J. (1966b). Effect of orientation and shape similarity on perceptual grouping. *Perception & Psychophysics, 1*, 300-302.

Beck, J. (1982). Textural segmentation. In *Organization and Representation in Perception*. ed. Beck, J. Hillside, NJ: Erlbaum.

Bergen, J. R. (1991). Theories of Texture Perception. In D. Regan (ed). *Vision and Visual Dysfunction: Spatial Vision (vol 10)*. Macmillan Press: London.

Bergen, J. R. & Adelson, E. H. (1988). Visual texture segmentation and early vision. *Nature, 333*, 363-364.

Biederman, I. (1985). Human image understanding: Recent research and a theory. *Computer Vision, Graphics and Image Processing, 32*, 29-73.

Blake, A. & Marinos, C. (1990). Shape from texture: Estimation, isotropy and moments. *Artificial Intelligence, 45,* 323-380.

Blakemore, C., Carpenter, R. H. S. & Georgeson, M. A. (1970). Lateral inhibition between orientation detectors in the human visual system, *Nature, 228*, 37-39.

Blakemore, C., Carpenter, R. H. S. & Georgeson, M. A. (1971). Lateral thinking about lateral inhibition, *Nature, 234*, 418-419.

Blakemore, C. B. & Sutton, P. (1969). Size adaptation: A new aftereffect. *Science, 166*, 245-247.

Caelli, T. (1982). On discriminating visual textures and images. *Perception & Psychophysics. 31(2)*, 149-159.

Caelli, T. (1985). Three processing characteristics of visual texture segmentation. *Spatial Vision, 1*, 19-30.

Caelli, T. & Julesz, B. (1979). Psychophysical evidence for global feature processing in visual texture discrimination. *Journal of the Optical Society of America, 69(5)*, 675-678.

Callaghan, T. C., Lasaga, M. I. & Garner, W. R. (1986). Visual texture segregation based on orientation and hue. *Perception & Psychophysics, 39(1)*, 32-38.

Campbell, F. W. & Maffei, L. (1971). The tilt aftereffect: A fresh look. *Vision Research, 11*, 833-840.

Carney, T. (1982). Direction-specific tilt illusion: With and without gaze fixation. *Perception, 11(5)*, 529-533.

Cattell, J. M. (1886). The inertia of the eye and brain, *Brain, 8*, 295-312.

Chang, J. J. & Julesz, B. (1984). Cooperative phenomena in apparent movement perception of random-dot cinematograms. *Vision Research, 24*, 1781-1788.

Chaudhuri, A. (1990). Modulation of the motion aftereffect by selective attention, *Nature, 344*, 60-62.

Coltheart, M. (1971). Visual feature-analyzers and after-effects of tilt and curvature, *Psychological Review, 78(2)*, 114-121.

Culham, J. C. & Cavanaugh, P. (1994). Motion capture of luminance stimuli by equiluminous colour gratings and by attentive tracking. *Vision Research, 34*, 2701-2706.

Dakin, S. C. (1997). The detection of structure in glass patterns: Psychophysics and computational models, *Vision Research, 37(16)*, 2227-2246.

Dakin, S. C. & Watt, R. J. (1997). The computation of orientation statistics from visual texture, *Vision Research, 37(22)*, 3181-3192.

Dubin, M. W. & Cleland, B. G. (1977). Organization of visual inputs to interneurons of lateral geniculate nucleus of the cat. *Journal of Neurophysiology, 40*, 410-427.

Duncan, J. & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review, 96(3)*, 433-458.

Enns, J. (1986). Seeing textons in context. *Perception & Psychophysics, 39(2)*, 143-147.

Favreau, O. E. (1976). Motion Aftereffects: Evidence for parallel processing in motion perception. *Vision Research, 16*, 181-186.

Field, D., Hayes, A. & Hess, R.F. (1993). Contour integration by the human visual system: Evidence for a local "association field", *Vision Research, 33(2)*, 173-193.

Funke, K. & Eysel, U. T. (1992). EEG-dependent modulation of response dynamics of cat dLGN relay cells and the contribution of corticogeniculate feedback, *Brain Research, 573,* 217-227.

Gibson, J. J. (1937). Adaptation, after-effect and contrast in the perception of tilted lines: II. Simultaneous contrast and the areal restriction of after-effect. *Journal of Experimental Psychology, 20,* 553-569.

Gibson, J. J. & Radner, (1937). I. Quantitative studies. *Journal of Experimental Psychology, 20,* 453-467.

Gilbert, C. D. & Weisel, T. N. (1989). Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *Journal of Neuroscience, 9,* 2432-2442.

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

Glass, L. & Perez, R. (1973). Perception of random dot interference patterns, *Nature, 246,* 360-362.

Glass, L. & Switkes, E. (1976). Pattern recognition in humans: Correlations which cannot be perceived, *Perception, 5,* 67-72.

Gove, A., Grossberg, S. & Mingolla, E. (1995). Brightness perception, illusory contours and corticogeniculate feedback. *Visual Neuroscience, 12,* 1027-1052.

Grabowska, A. (1987). Visual field differences in the magnitude of the tilt after-effect. *Neuropsychologia, 25(6),* 957-963.

Graham, N., Beck, J. & Sutter, A. (1992). Nonlinear processes in spatial-frequency channel models of perceived texture segregation: Effects of sign and amount of contrast. *Vision Research, 32(4),* 719-743.

Green, D. M. & Swets, J. A. (1966). *Signal Detection Theory and Psychophysics.* New York: Wiley.

Grossberg, S. (1998). How does the cerebral cortex work? Learning, attention and grouping by the laminar circuits of visual cortex. Spatial Vision.

Grossberg, S. (in press). The complementary brain: A unifying view of brain specialization and modularity. Trends in Neuroscience.

Grossberg, S., Mingolla, E. & Todorovic, D. (1989). A neural network architecture for preattentive vision. *IEEE Transactions of Biomedical Engineering, 36(1),* 65-85.

Grossberg, S. & Raizada (1999). Contrast-sensitive perceptual grouping and object-based attention in the laminar circuits of primary visual cortex (Technical Report No. CAS/CNS TR-99-008, Forthcoming in Vision Research). Boston: Boston University, Department of Cognitive and Neural Systems.

Hammond, P., Mouat, G. S. V. & Smith, A. T. (1988). Motion aftereffects in cat striate cortex elicited by moving texture. *Experimental Brain Research, 72,* 1-20.

Harris, J. P. & Calvert, J. E. (1988). Contrast, spatial frequency and test duration effects on the tilt aftereffect: Implications for underlying mechanisms. *Vision Research, 29(1),* 129-135.

He, Z. J. & Nakayama, K. (1994). Perceiving textures: Beyond filtering. *Vision Research, 34(2)*, 151-162.

Humphreys, G. W. & Muller, H, J. (1993). Search via recursive rejection (SERR): A connectionist model of visual search. *Cognitive Psychology, 25,* 43-110.

Hull, D. H. (1968). Corticofugal influence in the macaque lateral geniculate nucleus. *Vision Research, 8*, 1285-1298.

Joffe, K. M., Raymond, J. E., & Crichton, A. (1997). Motion coherence perimetry in glaucoma and glaucoma suspects, *Vision Research, 97*, 955-964.

Julesz, B. (1962). Visual pattern discrimination. *IRE Trans. Inf. Theory, IT-8*, 84-92.

Julesz, B. (1975). Experiments in the visual perception of texture. *Scientific American, 232*, 34-43.

Julesz, B. (1978). Visual texture discrimination using random-dot patterns, *Journal of the Optical Society of America, 68*, 268-70.

Julesz, B. (1980). Spatial nonlinearities in the instantaneous perception of textures with identical power spectra. *Philos. Trans. R. Soc. Lond. (Biol.), 290*, 83-94.

Julesz, B. (1981). Textons, the elements of texture perception and their interactions. *Nature, 290*, 91-97.

Julesz, B. & Burt, P. (1979). *Bulletin of the Psychonomics Society, 12,* 242.

Julesz, B. & Bergen, J. R. (1983). Parallel versus serial processing in rapid pattern discrimination. *Nature, 303*, 696-698.

Kalil, R. E. & Chase, R. (1970). Corticofugal influence on activity of lateral geniculate nucleus neurons in the cat. *Journal of Neurophysiology, 33,* 459-474.

Kanwisher, N. C. (1987). Repetition blindness: type recognition without token individuation. *Cognition, 27,* 117-43.

Keeble, D. R. T., Kingdom, F. A. A. & Morgan, M. J. (1997). The orientational resolution of human texture perception. *Vision Research, 37(21),* 2993-3007.

Keeble, D. R. T., Kingdom, F. A. A., Moulden, B., & Morgan, M. J. (1995). Detection of orientationally multimodal textures. *Vision Research, 35(14),* 1991-2005.

Kingdom, F. A. A., Keeble, D. & Moulden, B. (1995). Sensitivity to orientation modulation in micropattern-based textures. *Vision Research, 35(1)*, 79-91.

Kisvárday, Z. F., Tóth, E., Rausch, M. & Eysel, U. T. (1995). Comparison of lateral excitatory and inhibitory connections in cortical orientation maps of the cat. *Society for Neuroscience Abstracts, 360.11*, 907.

Lamme, V. A. (1995). The neurophysiology of figure/ground segregation in primary visual cortex. *Journal of Neuroscience, 15,* 1605-1615.

Landy, M. S. & Bergen, J. R. (1991). Texture segregation and orientation gradient. *Vision Research, 31(4)*, 679-691.

Leehey, S. C., Moskowitz-Cook, A., Brill, S. & Held, R. (1975). Orientational anisotropy in infant vision, *Science, 190,* 900-902.

Livingstone, M. S., Robsen, G. D., Drislane, F. W. & Galaburda, A. M. (1991). Physiological and anatomical evidence for a magnocellular deficit in developmental dyslexia. *Proceedings of the National Academy of Science, 88*, 7943-7947.

Lovegrove, W. J., Martin, F., Bowling, A., Blackwood, M., Badcock, D. & Paxton, S. (1982) Contrast sensitivity functions and specific reading disability, *Neuropsychologica, 20*, 309-315.

MacKay, D. G. (1987). *The Organization of Perception and Action*. New York: Springer.

Magnussen, S. & Kurtenbach, W. (1980). Adapting to two orientations: Disinhibition in a visual aftereffect. *Science, 207(4433),* 908-909.

Magnussen, S & Johnsen, T. (1986). Temporal aspects of spatial adaptation: A study of the tilt aftereffect. *Vision Research, 26(4),* 661-672.

Makela, P., Whitaker, D. & Rovamo, J. (1993). Modelling of orientation discrimination across the visual field, *Vision Research, 33*, 723-730.

Malik, J. & Perona, P. (1990) Preattentive texture discrimination with early vision mechanisms, *Journal of the Optical Society of America A., 7(5)*, 923-932.

Mansfield, R. (1974). Neural basis of orientation perception in primate vision. *Science, 186,* 1133-1135.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.

Martin, F. & Lovegrove, W. J. (1987). Flicker contrast sensitivity in normal and specifically disabled readers. *Perception, 16*, 215-221.

Mayhew, J. E. W. & Anstis, S. M. (1972). Movement aftereffects contingent on color, intensity and pattern. *Perception & Psychophysics, 12*, 77-85.

McIlhagga, W. H. & Mullen, K. T. (1996). Contour integration with colour and luminance contrast. *Vision Research, 36*, 1265-1279.

Merigan, W. H. & Maunsell, J. H. (1993). How parallel are the primate visual pathways? *Annual Review of Neuroscience, 16*, 369-402.

Michaels, C. F. & Turvey, M. T. (1979). Central sources of masking: Indexing structures supporting seeing at a single brief glance. *Psychological Research, 41*, 1-61.

Moran, J. & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science, 229*, 782-784.

Morant, R. B. & Harris, J. R. (1965). Two different aftereffects of exposure to visual tilts. *American Journal of Psychology, 78,* 218-226.

Motter, B. C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2 and V4 in the presence of competing stimuli, *Journal of Neurophysiology, 70,* 909-919.

Mussler, J. & Hommel, B. (1997). Blindness to response compatible stimuli. *Journal of Experimental Psychology: Human Performance and Perception, 23*, 861-72.

Nawrot, M. & Sekuler, R. (1990). Assimilation and contrast in motion perception: Explorations in cooperativity. *Vision Research, 30(10)*, 1439-1451.

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology, 9*, 353-383.

Newsome, W. T. & Pare, E. B. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *Journal of Neuroscience, 8(6)*, 2201-11.

Nothdurft, H. C. (1985a). Sensitivity for structure gradient in texture discrimination tasks. *Vision Research, 25(12)*, 1957-1968.

Nothdurft, H. C. (1985b). Orientation sensitivity and texture segmentation in patterns with different line orientation. *Vision Research, 25(4)*, 551-560.

Nothdurft, H. C. (1991). Texture segmentation and pop-out from orientation contrast. *Vision Research, 31(6)*, 1073-1078.

Nothdurft, H. C. & Li, C. Y. (1985). Texture discrimination: Representation of orientation and luminance differences in cells of the cat striate cortex, *Vision Research, 25(1)*, 99-113.

O'Donnell, H. L., & Raymond, J. E. (1996). Global perception of orientation in 'noisy' textures. *Investigative Ophthalmology and Visual Science, 38(4)*, 636.

O'Donnell, H. L. & Raymond, J. E. (1997). A new slant on the tilt aftereffect: Orientation aftereffects on global perception of heterogeneous line-segment textures. *Perception*.

Paradiso, M. A., Shimojo, S. & Nakayama, K. (1989). Subjective contours, tilt aftereffects, and visual cortical organization. *Vision Research, 29(9)*, 1205-1213.

Petersik, A., Shepard, A., & Malsh, R. (1984). A three-dimensional motion aftereffect produced by prolonged adaptation to a rotation simulation. *Perception, 13*, 488-497.

Polat, U. & Sagi, D. (1993). Lateral interactions between spatial channels: Suppression and facilitation revealed by lateral masking experiments. *Vision Research, 33(7)*, 993-999.

Polat, U. & Sagi, D (1994). The architecture of perceptual spatial interactions. *Vision Research, 34(1)*, 73-78.

Pomerantz, J. R. (1981). Perceptual organization in information processing. In M. Kubovy and J. R. Pomerantz (eds.), *Perceptual Organization,* Hillsdale, NJ.: Erlbaum. Pp. 141-180.

Prinzmetal, W. (1981). Principles of feature integration in visual perception. *Perception and Psychophysics, 30*, 330-340.

Posner, M. I. & Cohen, Y. (1984). Components of visual orienting. In Bouma, H. & Bouwhuis (Eds), *Attention and performance X: Control of language processes.* Hillsdale, NJ: Lawrance Erlbaum.

Potts, M. J. & Harris, J. P. (1975). Movement aftereffects contingent on the colour or pattern of a stationary surround. *Vision Research, 15(11)*, 1225-1230.

Ramachandran, V. S. (1987). Interaction between colour and motion in human vision. *Nature, 328,* 645-647.

Raymond, J. E. (1993). Movement directional analyzers: Independence and bandwidth. *Vision Research, 33,* 767-775.

Raymond, J. E. (1994). Directional anisotropy of motion sensitivity across the visual field. *Vision Research, 34 (8),* 1029-1037.

Raymond, J. E. & Braddick, O. (1996). Responses to opposed directions of motion: Continuum or independent mechanisms? *Vision Research, 36(13),* 1931-1937.

Raymond, J. E. & Isaak, M. (1998). Successive episodes produce direction contrast effects in motion perception. *Vision Research, 38,* 579-90.

Raymond, J. E. & O'Donnell, H. L. (1996). Motion Capture vs. Contrast by Prior Events. *Investigative Ophthalmology and Visual Science, 37 (3),* 744.

Raymond, J. E., O'Donnell, H. L. & Tipper, S. (1997). Attention modulates global motion sensitivity. *Investigative Ophthalmology and Visual Science, 38(4)* 369.

Raymond, J. E., O'Donnell, H. L., & Tipper, S. P. (1998). Priming reveals attentional modulation of human motion sensitivity, *Vision Research, 38,* 2863-2867.

Raymond, J. E., Shapiro, K. L. & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: an attentional blink? *Journal of Experimental Psychology: Human Perception and Performance, 18,* 849-860.

Raymond, J. E., & Sorensen, R. E. (1998). Visual motion perception in children with dyslexia: Normal detection but abnormal integration, *Visual Cognition, 5,* 389-404.

Rentschler, I., Treutwein, B. & Landis, T. (1994). Dissociation of local and global processing in visual agnosia, *Vision Research, 34(7),* 963-971.

Robson, J. A. (1983). The morphology of corticofugal axons to the dorsal lateral geniculate nucleus in the cat. *Journal of Comparative Neurology, 216,* 89 - 103.

Roelfsema, P. R., Lamme, V. A. & Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. *Nature, 395,* 376-381.

Sagi, D. (1990). Detection of an orientation singularity in gabor textures: Effect of signal density and spatial frequency, *Vision Research, 30,* 1377-1388.

Sekuler, R. W. & Ganz, L. (1963). Aftereffects of seen motion with a stabilized retinal image. *Science, 139,* 419-420.

Shapiro, K. L., Caldwell, J. I., & Sorensen, R. E. (1997). Personal names and the attentional blink: The 'cocktail party' effect revisited., *Journal of Experimental Psychology: Human Perception and Performance, 23,* 504-514.

Shulman, G. L. (1991). Attentional modulation of mechanisms that analyze rotation in depth. *Journal of Experimental Psychology Human Perception and Performance, 17(3),* 726-737.

Shulman, G. L. (1993). Attentional effects on adaptation of rotary motion in the plane. *Perception, 22,* 947-961.

Snowden, R. J. & Braddick, O. J. (1989a). The combination of motion signals over time. *Vision Research, 29(11),* 1621-1630.

Snowden, R. J. & Braddick, O. J. (1989b). Extension of displacement limits in multiple-exposure sequences of apparent motion. *Vision Research, 29(12),* 1777-1787.

Tipper, S. P. (1985). The negative priming effect: Inhibitory effects of ignored primes. *Quarterly Journal of Experimental Psychology, 37A,* 571-590.

Tipper, S. P. & Driver, J. (1988) Negative priming between pictures and words: Evidence for semantic analysis of ignored stimuli. *Memory and Cognition, 16,* 64-70.

Tipper, S. P., MacQueen, G. M. & Brehaut, J. C. (1988) Negative priming between modalities: Evidence for the central locus of inhibition in selective attention, *Perception & Psychophysics, 43,* 45-52.

Treisman, A. M. (1982). Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 194-214.

Treisman, A. M. (1988). Features and objects: The fourteenth Bartlett memorial lecture. *Quarterly Journal of Experimental Psychology, 40,* 201-237.

Treisman, A. M. & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12(1),* 97-136.

Treisman, A. M. & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review, 95,* 15-48.

Ungerleider, L. G. & Haxby, J. V. (1994). 'What' and 'Where' in the human brain. *Current Opinion in Biology, 4,* 157-165.

Updyke, B. V. (1975). The patterns of projection of cortical areas 17, 18 and 19 onto the laminae of the dLGN in the cat. *Journal of Comparative Neurology, 163,* 377-396.

von der Heydt, R., Peterhans, E. & Baumgartner, G. (1984). Illusory contours and cortical neuron responses, *Science, 224,* 1260-1262.

Voorhees, H. & Poggio, T. (1988). Computing texture boundaries from images. *Nature, 333,* 364-367.

Watamaniuk, S. N. J., Sekuler, R. & Williams, D. W. (1989). Direction perception in complex dynamic displays: The integration of direction information. *Vision Research, 29(1),* 47-59.

Watanabe, T., Sasaki, Y., Nielsen, M., Takino, R., & Miyakawa, S. (1998). Attention-regulated activity in human primary visual cortex, *Journal of Neurophysiology, 79,* 2218-2221.

Weber, A. J., Kalil, R. E. & Behan, M. (1989). Synaptic connections between corticogeniculate axons and interneurons in the dorsal lateral geniculate nucleus of the cat, *Journal of Comparative Neurology, 289,* 156-164.

Williams, D. W. & Phillips, G. (1987). Cooperative phenomena in the perception of motion direction. *Journal of the Optical Society of America A, 4(5),* 878-885.

Williams, D. W., Phillips, G. & Sekuler, R. (1986). Hysteresis in the perception of motion direction as evidence for neural cooperativity. *Nature, 324,* 253-255.

Williams, D. W. & Sekuler, R. (1984) Coherent global motion percepts from stochastic local motions. *Vision Research, 24*, 55-62.

Wilson, H. R., Wilkinson, F. & Asaad, W. (1997) Concentric orientation summation in human form vision, *Vision Research, 37(17)*, 2325-2330.

Witkin, A. (1981) Recovering surface shape and orientation from texture. *Artificial Intelligence, 17*, 17-47.

Wolfe, J. M. (1984). Short test flashes produce large tilt aftereffects, *Vision Research, 24(12),* 1959-1964.

Wolfe, J. M., Cave, K. R. & Franzel, S. L. (1989). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 419-433.

Wolfe, J. M. (1992). "Effortless" texture segregation and "parallel" visual search are not the same thing, *Vision Research, 32(4)*, 757-763.

Wolfe, J. M. (1994). Guided search 2.0: A revised model of visual search. *Psychonomic Bulletin and Review, 1(2)*, 202-238.

Zipser, K., Lamme, V. A. & Schiller, P. (1996). Contextual modulation in primary visual cortex. *Journal of Neuroscience, 16,* 7376-7389.