Pooling and Suppression in Human Spatial Vision

DAVID JOSEPH HOLMES

Doctor of Philosophy

THE UNIVERSITY OF ASTON IN BIRMINGHAM

September 2003

Synopsis

Aston University
Pooling and Suppression in Human Spatial Vision
David Joseph Holmes
Doctor of Philosophy
2003

A distinct feature of several recent models of contrast masking is that detecting mechanisms are divisively inhibited by a broadly tuned 'gain pool' of narrow-band spatial pattern mechanisms. The contrast gain control provided by this 'cross-channel' architecture achieves contrast normalisation of early pattern mechanisms, which is important for keeping them within the non-saturating part of their biological operating characteristic. These models superseded earlier 'within-channel' models, which had supposed that masking arose from direct stimulation of the detecting mechanism by the mask. To reveal the extent of masking, I measured the levels produced with large ranges of pattern spatial relationships that have not been explored before. Substantial interactions between channels tuned to different orientations and spatial frequencies were found. Differences in the masking levels produced with single and multiple component mask patterns provided insights into the summation rules within the gain pool. A widely used cross-channel masking model was tested on these data and was found to perform poorly. The model was developed and a version in which linear summation was allowed between all components within the gain pool but with the exception of the self-suppressing route typically provided the best account of the data. Subsequently, an adaptation paradigm was used to probe the processes underlying pooled responses in masking. This delivered less insight into the pooling than the other studies and areas were identified that require investigation for a new unifying model of masking and adaptation. In further experiments, levels of cross-channel masking were found to be greatly influenced by the spatio-temporal tuning of the channels involved. Old masking experiments and ideas relying on within-channel models were re-evaluated in terms of contemporary cross-channel models (e.g. estimations of channel bandwidths from orientation masking functions) and this led to different conclusions than those originally arrived at. The investigation of effects with spatio-temporally superimposed patterns is focussed upon throughout this work, though it is shown how these enquiries might be extended to investigate effects across spatial and temporal position.

Key Words: Visual psychophysics; Inhibition; Contrast normalisation; Contrast discrimination; Cortex

Acknowledgements:

My thanks are offered to Dr Tim Meese, for academic supervision. My thanks are also offered to the BBSRC and the James Watt Memorial Foundation, for financial support.

List of Contents:

Chapter One:

General Introduction

1.1 The purpose of vision and methods of its exploration (15)
1.2 Spatially tuned mechanisms in vision (15)
1.3 Early models of masking (18)
1.4 The within-channel model is challenged (20)
1.5 Foley's tests of the within-channel masking model (21)
1.6 Foley's models of human pattern masking (22)
1.7 Other models of divisive inhibition (24)
1.8 Why cross-channel models of masking are important (24)
1.9 Thesis Preview (25)

Chapter Two:

General Methods of Research

2.1 Chapter Introduction (28)
2.2 Stimuli (28)
  Sinusoidal gratings (28)
  Plaids (29)
  Notation for grating and plaid stimuli (29)
  Spatial truncation of sinusoids (29)
  Size of stimuli (29)
  Temporal envelope of stimuli (30)
2.3 Apparatus (30)
2.4 Observers (30)
2.5 Procedure (31)
2.6 Experimental paradigms (31)

Chapter Three:

Simultaneous Masking by Plaids and Gratings: Threshold Elevation and Threshold Facilitation

3.1 Chapter Introduction (33)
3.2 Chapter General Methods (35)
3.3 Experiment 1. Masking and facilitation as a function of test spatial frequency (35)
  Methods (35)
  Results and Discussion (35)
    Masking (35)
    Facilitation (36)
3.4 Experiment 2. Control for adjacent pixel nonlinearities (38)
  Introduction and Methods (38)

Results and Discussion (38)
  Previous reports of facilitation (39)
3.5 Experiment 3. Facilitation as a function of the orientation difference between test and mask (40)
Introduction (40)
Methods (40)
Results and Discussion (40)
3.6 Experiment 4. The effect of mask contrast on masking and facilitation produced by grating and plaid masks (44)
Introduction (44)
Methods (44)
Results and Discussion (45)
3.7 General Discussion (46)
Facilitation: Causes? (46)
Masking: Linear summation of mask components? (47)
3.8 Experiment 5(a). The effect of spatial frequency on mask component summation: Masking as a function of test spatial frequency for 1, 3 and 12cpd grating and plaid masks (49)
Introduction (49)
Methods (49)
Results and Discussion (51)
  (i) The facilitation from remote masks (51)
  (ii) The masking from extremely remote masks (51)
  (iii) The test spatial frequency at which peak masking occurred (52)
  (iv) The difference between the masking produced by the plaid and the grating (a) (53)
  The difference between the masking produced by the plaid and the grating (b) (55)
3.9 Experiment 5(b). The dependency of super-suppression on mask contrast (56)
Introduction and Methods (56)
Results and Discussion (56)
Another finding specific to the use of 3cpd stimuli (60)
3.10 Chapter Summary (60)

Chapter Four:

Linear- and Super-Suppression from Remote Masks

4.1 Chapter Introduction (61)
4.2 Chapter General Methods (61)
4.3 Experiment 1. Linear- and super-suppression when test and mask have the same spatial frequency (62)
Introduction and Methods (62)
Results (62)
  Stimuli with 1cpd (62)
  Stimuli with 3cpd (a) (62)
  Stimuli with 3cpd (b): A control (64)
  Stimuli with 9cpd (64)

Discussion (64)
4.4 Experiment 1(b). Linear-suppression for remote masking (65)
Introduction and Methods (65)
Results and Discussion (65)
4.5 Understanding masking (67)
4.6 Foley's (1994a) models (68)
4.7 Fitting Foley's models to remote masking data (69)
Modelling Methods (69)
Results and Discussion (70)
Experiment 1(c). Fitting Foley's models to more remote masking data (71)
4.8 Revisiting data in chapter 3. Linear-suppression from many stimulus SF combinations (73)
Introduction (73)
Modelling Methods (73)
Results and Discussion (75)
4.9 Experiment 2. Linear-suppression and mask component orientation (76)
Introduction (76)
Methods (76)
Results and Discussion (77)
4.10Chapter Discussion (77)

Chapter Five:

Testing Some Cross-Channel Masking Models

5.1 Chapter Introduction (79)
5.2 Methods (80)
(i) Experimental Methods (80)
(ii) Modelling Methods (80)
5.3 Results and Discussion (82)
The models and their fits (87)
F3 (88)
F2 (88)
Need for an alternative to Foley's (1994a) models (89)
The hybrid model (89)
The compound model (90)
The compound-plus model (90)
Physiological interpretations of the models (91)
The alternative compound-plus model (91)
5.4 Chapter Summary and Discussion (92)

Chapter Six:

A Reappraisal of Orientation Masking

6.1 Chapter Introduction (97)
6.2 General Methods (101)
6.3 Masking as a function of mask orientation (101)
    Methods (101)
    Results and Discussion (104)
6.4 Orthogonal masks don't produce dippers (105)
    Introduction (105)
    Methods (105)
    Results and Discussion (106)
6.5 Orientation masking is phase-insensitive (107)
    Introduction and Methods (107)
    Results and Discussion (108)
6.6 Chapter Summary (109)
6.7 Chapter Discussion (109)

Chapter Seven:

Pattern Adaptation with Pattern Masking

7.1 Chapter Introduction (110)
7.2 Preliminary Experiment. Estimating the amount of forward masking produced with an adaptation paradigm (111)
    Introduction (111)
    Methods (111)
    Results and Discussion (112)
7.3 Attempting to gain insight into linear-suppression using adaptation (113)
7.4 Chapter Discussion (117)

Chapter Eight:
Summary and Further Discussions (121)

References (127)

Appendix One:

Some pilot work and ideas on lateral masking

A1.1 Chapter Introduction (160)
    A brief review of some lateral masking effects (160)
A1.2 (Pilot) Experiment 1. Lateral masking irrespective of mask orientation - as long as the mask is spatially close to the centre of the test (161)
    Introduction (161)
        Rationale (161)

        The pilot work (162)
    Methods (163)
    Results and Discussion (164)
        What is being revealed by the lateral masking (when test and mask are directly abutting)? (165)
A1.3 Experiment 2a. Does masking reveal asynchronies of CRF excitation & suppression latencies? (167)
    Introduction (167)
    Methods (168)
    Results and Discussion (168)
A1.4 Experiment 2(b) Previous masks produce more masking: a product of (i) the latency period for channel pooling or (ii) characteristics of the individual channels excited by the stimuli? (169)
    Introduction and Methods (169)
    Results and Discussion (170)
A1.5 Chapter Summary (171)

List of Figures:

1.1. A schematic illustration showing how the thesis sections are related. (27)

3.1. Masking and facilitation as a function of test spatial frequency. (36)

3.2. Masking and facilitation as a function of test spatial frequency for three naïve observers. (37)

3.3. Effects of test orientation on masking and facilitation. (39)

3.4. Facilitation as a function of test spatial frequency. Different panels are for different test orientations. The mask had a spatial frequency of 1cpd, was oriented -45deg and had a contrast of 8%. (42)

3.5. The test spatial frequency at which the functions shown in Fig. 3.4 cross from masking to facilitation, as a function of the orientation of the test. (43)

3.6. Effects of mask contrast on masking and facilitation produced by gratings and plaids. (45)

3.7. Data from Fig. 3.6. Masking and facilitation as a function of test spatial frequency, produced by grating and plaid masks which had the same stimulus contrasts at either 16% or 24%. (48)

3.8. Masking and facilitation as a function of test spatial frequency, produced by grating and plaid masks matched in component contrast at 24%. Mask spatial frequency was 1, 3 or 12cpd. (50)

3.9. Differences between masking functions produced by plaid and grating masks. Mask spatial frequency was 1, 3 or 12cpd. Ranges of predicted differences have been estimated for Foley's (1994a) model 2 and model 3 with parameters set to reasonable values (shaded patches). Parameter values were selected from those estimated in chapters 4 and 5. (54)

3.10. Masking and facilitation as a function of test spatial frequency. Grating and plaid masks had either 1 or 3cpd and they were matched in either their component or stimulus contrast at 24%. (56)

3.11. Masking and facilitation as a function of test spatial frequency, produced by 3cpd grating and plaid masks matched in component contrast at 8% or 24%. (58)

3.12. The difference between the masking and facilitation produced by a plaid mask and that produced by a grating mask (plaid mask RE grating mask), as a function of test spatial frequency. Masks had 3cpd and a component contrast of either 8% or 24%. The magnitude of the difference in masking produced by the two types of mask is shown for Foley's (1994a) model 2 and model 3 with parameters set to reasonable values (shaded patches). Parameter values were selected from those estimated in chapters 4 and 5. (59)

4.1. Masking as a function of mask stimulus contrast. Masks were gratings or plaids and the stimuli had 1 or 3cpd. (63)

4.2. Masking as a function of mask stimulus contrast. These data are from a control experiment, which addressed some of the results shown in Fig. 4.1. Masks were gratings or plaids and all stimuli had 3cpd. The number of cycles per stimulus patch was matched to that of the 1cpd stimuli used in gathering data in Fig. 4.1. (63)

4.3. Masking as a function of mask stimulus contrast. The test and masks had 9cpd and the masks were gratings or plaids. (63)

4.4. Test detection thresholds as a function of mask stimulus contrast. The test and masks had 1 and 3cpd, respectively, and masks were gratings or plaids. Fits of Foley's (1994a) model 2 and 3 are shown. (66)

4.5. A summary of key results shown in chapter 4 and chapter 3. The masking produced by gratings and plaids, matched in component (CC) or stimulus (SC) contrast, is shown for various stimulus spatial frequency combinations. (67)

4.6. Masking as a function of mask stimulus contrast. The test and masks had 3 and 9cpd, respectively, and the masks were -45° gratings or ±45° plaids. (71)

4.7. Test contrast as a function of the stimulus contrast of grating and plaid masks. Test and mask spatial frequencies were 2 and 6cpd, respectively, or 3 and 9cpd, respectively. Fits of Foley's (1994a) model 2 and model 3 are shown. (72)

4.8. Grating and plaid masking as a function of mask component orientation. The test and masks had 1 and 3cpd, respectively. Mask stimulus contrast was 32%. (77)

5.1. Stimuli used for the experiment shown in chapter 5. (81)

5.2. The affect of 3cpd -45deg grating and ±45deg plaid fixed contrast masks on test contrast, as a function of pedestal contrast, when test and pedestal were vertical 1cpd gratings. The fixed contrast masks had a stimulus contrast of 20%. (82)

5.3. The effect of 3cpd grating and plaid fixed contrast masks on test contrast, as a function of pedestal contrast, when the test and pedestal were 1cpd gratings. The fixed contrast masks had a stimulus contrast of 20%. Data are from DJH. The same data are shown in each panel fit by a different cross-channel masking model. (84)

5.4. The effect of 3cpd grating and plaid fixed contrast masks on test contrast, as a function of pedestal contrast, when test and pedestal were 1cpd gratings. The fixed contrast masks had a stimulus contrast of 20%. Data are from TSM. The same data are shown in each panel fit by a different cross-channel masking model. (85)

5.5. Schematic illustration of Foley's (1994a) model 2. (94)

5.6. Schematic illustration of Foley's (1994a) model 3. (94)

5.7. Schematic illustration of the hybrid model. (95)

5.8. Schematic illustration of the compound model. (95)

5.9. Schematic illustration of the compound-plus model. (96)

5.10. Schematic illustration of the alternative compound-plus model. (96)

6.1. Masking as a function of mask orientation. Stimuli had 1, 3 or 9cpd, and were temporally transient or sustained. Tests were vertical whilst mask orientation was varied between 0 and -90deg. (102)

6.2. Data are from Phillips and Wilson (1984) and show masking as a function of mask orientation. Stimuli had 0.5, 2 or 8cpd, and were temporally transient or sustained. Tests were vertical whilst mask orientation was varied between 0 and -45deg. (103)

6.3. Masking as a function of mask contrast. Stimuli were transient and had 1cpd. Tests were vertical whilst masks were either vertical or horizontal. The dashed lines represent the unmasked test detection thresholds. (106)

6.4. Masking as a function of mask orientation. Stimuli were temporally transient and had 1cpd. The mask was either in-phase or in quadrature-phase with the test. (108)

6.5. Masking as a function of mask orientation. Stimuli were temporally transient and had 3cpd. The mask was either in-phase or in quadrature-phase with the test. (108)

7.1. The effect of pre-trial random-phase forward masks on the detection of tests presented in a 2-AFC paradigm. Mask contrast was either 32% or 99%. Tests were 1cpd vertical gratings with raised cosine envelopes. The full-field, phase-shifted masks were gratings oriented -45deg, which had a spatial frequency of either 1cpd or 3cpd. (112)

7.2. Within-Channel Experiment. The effect of adaptation on masking when test, mask and adapter were all -45deg 3cpd gratings. Adapter contrast was 32% and mask contrast was either 8% or 22.4%. (114)

7.3. Test contrast as a function of mask contrast after adaptation. The test was a vertical 1cpd grating and the mask was a -45deg 3cpd grating or a +45deg 3cpd grating. The observer adapted to either mean luminance ('blank') or a -45deg 3cpd grating; data from which are shown in panel (a) and panel (b), respectively. Panel (c) shows the data of the two previous panels, superimposed. (115)

7.4. The effect of adaptation on masking; the influence of (relative) adapter contrast. Mask contrast was 5.6%, whilst adapter contrast was either 32% or 99%. The test was a 1cpd vertical grating whilst mask and adapter spatial frequencies were matched at 3cpd. Mask gratings were oriented either -45deg or +45deg whilst adapter gratings were always oriented -45deg. Data from conditions using mean luminance ('blank') masks and adapters are also shown. (116)

7.5. The effect of adaptation on masking when all stimuli had a spatial frequency of 1cpd. The influence of (relative) adapter contrast is shown; mask contrast was 5.6% whilst adapter contrast was either 32% or 99%. The test was vertical whilst mask gratings were oriented either -45deg or +45deg and adapter gratings were always oriented -45deg. Data from conditions using mean luminance ('blank') masks and adapters are also shown. (117)

A1.1. Lateral masking effects. Data shown in the three panels were gathered using different types of enveloping between test and mask stimuli. The different columns in each panel represent the effects of different types of mask. (166)

A1.2. Masking from spatially superimposed or surrounding patterns, which were presented either before or after the tests. The spatial frequencies and orientations of test and mask were matched or they differed. (169)

A1.3. Masking as a function of the interval between mask and test onset. Stimuli were spatially superimposed 1cpd vertical Gabors. Masks were either in phase or in anti-phase with the tests. (170)

List of Tables:

4.1. F2 parameter values and goodness of fits for data in Fig. 4.4. (70)

4.2. F3 parameter values and goodness of fits for data in Fig. 4.4. (70)

4.3. F2 parameter values and goodness of fits. Data and fits are shown in the left-hand panels of Fig. 4.7. (72)

4.4. F3 parameter values and goodness of fits. Data and fits are shown in the right-hand panels of Fig. 4.7. (72)

5.1. F2 parameter values and goodness of fits for data in Figs. 5.3 and 5.4. (86)

5.2. F3 parameter values and goodness of fits for data in Figs. 5.3 and 5.4. (86)

5.3. Hybrid model parameter values and goodness of fits for data in Figs. 5.3 and 5.4. (86)

5.4. Compound model parameter values and goodness of fits for data in Figs. 5.3 and 5.4. (86)

5.5. Compound-plus model parameter values and goodness of fits for data in Figs. 5.3 and 5.4. (86)

5.6. Alternative compound-plus model parameter values and goodness of fits for data in Figs. 5.3 and 5.4. (86)

6.1. Unmasked thresholds for each test condition in Fig. 6.1. (102)

A1.1.The diameter and spatial enveloping of the stimuli used in (pilot) experiment 1 of Appendix 1. (163)

Chapter One: General Introduction

1.1 The purpose of vision and methods of its exploration

Vision responds to changes in light intensity over space and time. From luminance changes, the general purpose of early vision is to 'develop useful....descriptions of the shapes and surfaces that form the image' (Marr, 1982, p.41). Luminance contrast is the difference between a pattern's maximum and minimum luminance as a proportion of its mean luminance. Scenes are encoded by vision using this. Neurons in the retina, Lateral Geniculate Nucleus (LGN) and striate cortex (also known as V1 or the primary visual cortex) have receptive fields, which are selectively sensitive to particular patterns of contrasts. Fundamental roles of human pattern vision are performed through feedforward and feedback networks between the LGN and V1, and extra striate areas. From the retina onwards, the organization of the visual system is parallel and hierarchical (Lennie, 1998). Extra-striate areas perform 'higher' visual functions, such as binocular disparity, motion selectivity and colour processing. This thesis is concerned with psychophysical investigations into functions thought to be of V1 and subcortical areas.

Neurophysiology, psychophysics, and more recently, brain imaging, have allowed us to create, test and modify theories concerning the functional architecture of cortical and subcortical areas that determine our visual perceptions. The invasive nature of most physiological studies prevent them from being performed on humans (they are usually performed on monkey or cat), however, psychophysics has permitted detailed models of human vision to be formulated by measuring human behaviour. Brain imaging bridges the gap between the micro world of the neuron and the macro world of psychophysics.

Mechanisms involved with contrast processing in early vision are selectively sensitive to spatio-temporal characteristics of a pattern, such as its frequency, phase, position and spatial orientation. Because of such selectivity, it is popular to use sinusoidal grating stimuli in studies of early vision because their spatio-temporal characteristics are simple and easily controlled, and they are mathematically tractable. Sinusoidal gratings will be simply referred to as gratings from here. They are 1-D sinusoidal luminance modulations across space and their contrasts are defined as the ratio between their amplitude and mean luminance. Their spatial frequencies are defined as the number of periods modulated per unit visual angle (the unit usually applied is cycles per degree) and their orientations are defined as the angle between vertical and the orientation of their 'stripes'. Gratings are often multiplied by a Gaussian function, which localises them in space but spreads them in the Fourier domain. Such a pattern is known as a Gabor function and it closely resembles the qualities of the visual receptive fields of a class of neurons in the visual system (Jones & Palmer, 1987b, and see next section). (Though, strictly speaking, Gabor's are 1-dimensional mathematical functions, vision scientists have extended them to 2-dimensions). Every experiment in this thesis measured the detectability of gratings in, or after, the presence of one or more other gratings.

1.2 Spatially tuned mechanisms in vision

Most of the work in this thesis is concerned with the processing of luminance contrast, spatial frequency, and orientation in early vision (though temporal frequency and phase are addressed in chapter 6). Contrast sensitivity as a function of spatial frequency (measured with grating stimuli) can be approximated by an inverted U on log-log coordinates (Campbell & Robson, 1968), with the function peaking at around 3 to 5cpd. However, for over forty years it has been thought that the overall sensitivity of the visual system is not the result of a single broadband mechanism. Rather, contributing to the visual system's general sensitivity are a number of mechanisms selective to different narrow bands of spatial frequencies and orientations. Since their responses to a given stimulus are dependent upon the spatial characteristics of that stimulus, they are said to be spatially 'tuned'. The tuning functions' shapes and widths have been estimated in many physiological and psychophysical studies. Their tightness is often summarised as a single figure known as 'bandwidth'. With reference to spatial frequency tuning, the bandwidth is usually reported at full-width, half-height in octaves. With reference to orientation tuning, it is usually reported at half-width, half-height in degrees. These terms and measures are used here.

A physiological study concerned with estimating spatial tuning characteristics typically involves recording cell responses with an electrode. The electrode is inserted into a region of the visual pathway of a non-human observer. Measurements are then taken whilst the animal is shown different patterns. Classes of cells in the retina, LGN and striate cortex are receptive to spatial regions of 2-D visual space. These receptive fields have areas that, when stimulated, either increase or decrease the firing rate of the cell from its background level. The size, number and shape of these excitatory and inhibitory regions are a major factor in determining a cell's tuning to spatial frequency and orientation (e.g. Hubel & Wiesel, 1959; 1962, 1968). Initially, the patterns shown to observers in the physiological studies were spots of light and lines (Kuffler, 1953; Hubel & Wiesel, 1962, 1968), but gratings are now favoured (e.g. Enroth-Cugell & Robson, 1966). Maffei and Fiorentini (1973) found that the tuning of cat neurons to the spatial frequency of gratings is progressively narrower, moving up the visual pathway from the retina, to the LGN, and then to the striate cortex. Very similar results were shown by DeValois et al. (1982) for macaque. The responses of cortical neurons to contrast, as a function of spatial frequency, can be approximated as inverted U's on log-log co-ordinates. As the preferred spatial frequency of these mechanisms increases, their bandwidths decrease (Tolhurst & Thompson, 1981; DeValois et al., 1982) from around 2.5 to 1.25 octaves at full-width, half-height (with the average being about 1.4 octaves). With regard to orientation tuning, though it is practically absent in subcortical areas (e.g. see Passaglia et al., 2002 for a review), it is a prevalent characteristic of neurons in the striate cortex (Hubel & Wiesel, 1962; 1968; DeValois et al., 1982; Jones et al., 1987). The orientation bandwidths of these decrease with an increase in their preferred spatial frequency from about ±30° to about ±15° at half-width, half-height (DeValois et al., 1982).

Early physiological investigations into the spatial selectivity of visual mechanisms led some authors to deduce that simple V1 receptive fields serve to encode spatial features directly. It has been found that this is not the case, however. Fourier's theorem shows that any 1-D pattern can be expressed in terms of sine waves with different spatial frequencies, amplitudes and phases. This is also the case for 2-D patterns when the vertical and horizontal components of the sinusoids are analysed. It has become apparent that visual mechanisms are excited by characteristics of these sinusoids, such as limited ranges of their spatial frequencies and orientations. For example, Kelly (1976) noted that the fundamental sinusoids in a checkerboard pattern are oriented ±45° from the edges of the checks. He showed the psychophysical thresholds for checkerboards are determined by these sinusoidal components rather than the edges. Measuring V1 responses to such patterns, DeValois et al. (1979) added support to this. In addition, Maffei et al. (1979) and Pollen and Ronner (1982) measured the responses of V1 neurons to drifting square wave gratings. They found that when the cell's preferred spatial frequency was three times the spatial frequency of the grating, its modulated response had three times the periodicity of the square wave's fundamental frequency. In other words, the cells responded to the third harmonic of the gratings.

In the 1970's and 80's, many psychophysical experiments were performed in which the spatial frequency and orientation selectivities of visual mechanisms were estimated. Typically, these investigated how a pattern's contrast detection threshold was affected if the pattern was shown in the presence of, or after, another pattern. Three of the most widely used psychophysical 'tools' are subthreshold-summation, selective adaptation and simultaneous masking. These are briefly described here, though much of the discussion of adaptation is witheld until chapter 7 whilst masking receives extensive coverage later in this chapter and throughout this thesis.

In subthreshold summation experiments, it is typical for the detection thresholds of two patterns to be measured independently, after which the threshold of one (the test) is measured in the presence of the other pattern which is presented below its own detection threshold. The extent by which the subthreshold pattern lowers the detection threshold of the test pattern indicates the extent to which the two stimuli excite the same detecting mechanism. When spatio-temporally superimposed gratings are sufficiently different in terms of their spatial frequencies (e.g. Campbell & Robson, 1968; Graham & Nachmias, 1971) or orientations (Phillips & Wilson, 1984; Georgeson & Shackleton, 1994; Meese & Williams, 2000), they are each detected at their individual detection thresholds, that is, they do not sum to threshold. This is indicative of the gratings exciting different channels.

Selective adaptation is the term given to the procedure and consequence of exposing an observer to a phase-shifting (adapter) pattern for a prolonged period of time. It is usually measured in terms of the degree to which the contrast sensitivity to a subsequently viewed (test) pattern is altered. It differs from conventional after-images since, by shifting the adapter patterns across observers' visual fields, local light adaptation is averaged out (Jones & Tulunay-Keesey, 1980; Kelly & Burbeck, 1980). The amount by which an adapter grating increases the detection threshold of a test grating tends to reduce as the difference between the two patterns is increased with respect to their spatial frequencies (e.g. Pantle & Sekuler, 1968; Blakemore & Campbell, 1969a; Graham, 1972; Georgeson & Harris, 1984) or orientations (e.g. Gilinsky, 1968; Mayo et al., 1968; Sekuler et al., 1968; Blakemore & Campbell, 1969a; Blakemore & Nachmias, 1971; Movshon & Blakemore, 1973; Greenlee & Heitger, 1988; Greenlee et al., 1988; Maattanen & Koenderink, 1991; Snowden, 1992). Measuring the width of the tuning function produced by plotting test contrast against the manipulated spatial characteristic provides insight into the test detecting mechanism's tuning bandwidths.

The term 'masking' covers a range of experimental paradigms (for review see Breitmeyer, 1984), which may probe quite distinct mechanisms. Typically, it is the name given to the (usually detrimental) interference in detectability of one stimulus (the test) by another (the mask). One reason for the widespread use of masking is that masks can raise test thresholds by a factor of ten or more (by comparison, subthreshold summation produces effects that are in the region of a factor of two). Adaptation and masking differ primarily in time-course and test onset (see chapter 7). Like an adaptor, a mask tends to produce greater effects - the closer it is to the test in terms of spatial frequency (e.g. Greis & Roehler, 1970; Stromeyer & Julesz, 1972; Legge, 1978; Legge & Foley, 1980; Henning et al., 1981; Wilson et al., 1983; Daugman, 1984; Anderson & Burr, 1985; Henning, 1988) or orientation (e.g. Campbell & Kulikowski, 1966; Daugman, 1984; Phillips & Wilson, 1984; Harvey & Doan, 1990). Test contrast plotted as a function of the manipulated spatial characteristic produces masking functions that can be used to estimate mechanism bandwidth.

1.3 Early models of masking

Masking functions do not provide a direct estimate of the test detecting mechanism's tuning bandwidth. Rather, a model of masking is required which allows a gauging of how the magnitude of masking reflects the mechanism's sensitivity to the mask. A central assumption underlying the estimates was that threshold elevation occurs only when test and mask are processed by the same detecting mechanism, which operates independently from, and in parallel to, other mechanisms (i.e. there is a single processing stage). For many years, a widely used model of masking supposed an initial quasi-linear filtering stage followed by a nonlinear transducer and additive noise (Wilson, 1979; Legge & Foley, 1980). The transducer is sigmoidal in that it accelerates for low input levels and becomes compressive for moderate input levels and above. This reflects the contrast response characteristics of a simple cell (e.g. Movshon et al., 1978b; Albrecht & Hamilton, 1982; Sclar et al., 1990). If the test detecting mechanism is substantially responsive to both test and mask stimuli then it is driven into the compressive region of its output characteristic. After this, the detection threshold of the test increases with mask contrast and it is this that is termed masking. The positive correlation between masking and mask contrast is consistent with abundant psychophysical data (e.g. Wilson, 1979; Legge & Foley, 1980; Georgeson & Georgeson, 1987; Snowden & Hammett, 1998). One notable feature is that the widths of the masking functions are much broader than the bandwidths of the underlying mechanisms implied by the modelling. This is because of the effectively compressive region of contrast transduction that each mechanism is considered to have. Nachmias referred to this model in the title of a paper as the 'standard model', which reflects the high regard that it was once given.

In the above model, neural noise is zero mean, unit variance and additive. The transducer's compressive region, however, could be explained if the quantity of noise was positively correlated with signal strength (i.e. multiplicative; Tolhurst et al., 1983). This possibility is consistent with physiology that has shown the standard deviation of the spike count from cat V1 neurons to increase with stimulus contrast (Tolhurst et al., 1981; Tolhurst et al., 1983). A model with this characteristic is equivalent to the one described above in many circumstances (Wilson, 1980; Legge, 1984; Gorea & Sagi, 2001; Kontsevich et al., 2002). Yet another possibility is Pelli's (1985) suggestion that the transducer's accelerative region could be alternatively explained if the observer was considered to be uncertain as to which visual mechanisms to monitor for test detection.

'Channels' and 'filters' are terms used throughout this thesis. The detecting mechanism described above will be referred to as a channel. It is the quasi-linear mechanism that is followed by some form of nonlinearity. The term 'filter' will be used to refer to the quasi-linear mechanism on the front end of a channel and can be conceptualised as reflecting the function of a neuron's receptive field.

The model assuming additive noise and the model assuming multiplicative noise are both 'within-channel' models of masking. That is, masking is a consequence of the mechanism detecting the test also detecting the mask. In these models, masking is due to the masks producing a reduction in the signal to noise ratio for the test. Within-channel masking is discussed in terms of the sigmoidal transducer throughout this thesis, but the arguments apply equally for the model characterised by multiplicative noise.

The transducer functions were estimated by superimposing a test grating on a spatio-temporally similar (or matched) mask grating and measuring test detection at a range of mask contrasts (Nachmias & Sansbury, 1974; Legge & Foley, 1980; Wilson, 1980; Wilson et al., 1983; Phillips & Wilson, 1984). Test detection is typically facilitated by low contrast masks and suppressed by higher contrast masks. In other words, a plot of contrast increment threshold against mask contrast (TvC functions) produces a contrast masking function that is often referred to as a 'dipper function' (Nachmias & Sansbury, 1974, Legge & Foley, 1980), examples of which appear in Figs. 5.2 to 5.4 of chapter 5 and Fig. 6.3 of chapter 6. The upper region of a dipper function where masking occurs is a power function with an exponent typically ≤ 1 (1 corresponding to Weber's law) (Legge & Foley, 1980; Legge, 1981; Swift & Smith, 1983; Zenger & Sagi, 1985). (Though Greenlee & Heitger, 1988, reported the exponent > 1.) The within-channel model of masking typically takes the form:

$$r_t = c_t^p / (z^q + c_t^q) \qquad (1.1)$$

where r_t is the response output of the nonlinearity of the test detecting mechanism and c_t is the contrast response of a linear filter selectively tuned to the test component. The constant z has a value greater than 0 and is usually termed the semi-saturation constant because, for $p = q$, when $c_t = z$, r_t reaches a value that is half of that at which it saturates. (Though p was not constrained to equal q in any of the models presented in this thesis, z is termed the semi saturation constant here for convenience). The slopes of the accelerative and compressive regions of the sigmoid are controlled by the values of p and q respectively. Typical estimates of these parameters are $p > 2$ and $p \geq (q + 0.3)$. Because the upper region of the masking function (the dipper function) is dominated by c_t, Eqn. 1.1 becomes

$$r_t = c_t^{p-q} \qquad (1.2)$$

Discrimination thresholds (Δc_t) are approximately proportional to the reciprocal of the first derivative of r_t with respect to c_t:

$$\Delta c_t \approx (dr_t / dc_t)^{-1} \qquad (1.3)$$

$$dr_t / dc_t = (p - q)\, c_t^{(p-q)-1} \qquad (1.4)$$

therefore:

$$\Delta c_t \approx c_t^{1-(p-q)} \qquad (1.5)$$

In other words, the log-log slope of the 'dipper handle' (the part of the slope where test thresholds are raised) is approximately equal to 1 - (p - q), which is equivalent to about 0.7.

As mentioned, the within-channel model of masking has been used in estimating channel tuning widths. In considering pathways sensitive to the vertical orientation, Wilson et al., 1983, estimated (foveal) vision to be served by at least six pathways tuned to different bands of spatial frequency, with peak sensitivities ranging from 0.7cpd to 15.0cpd. They showed the bandwidth of the channel's tuning functions was dependent on the spatial frequency to which they exhibited peak sensitivity. They ranged from 2.5 octaves at low frequencies to around 1.25 octaves at high spatial frequencies. Phillips and Wilson (1984) used masking and the 'single processing stage' approach in estimations of orientation half-bandwidths which averaged around 20deg and increased with decreasing spatial frequency to which the channels were most sensitive; they were about 30deg for channels tuned to low spatial frequencies (0.5deg) and about 15deg for channels tuned to high spatial frequencies (11cpd).

1.4 The within-channel model is challenged

Around 30 years ago, findings began to emerge from some masking and adaptation studies that were inconsistent with the independent channels hypothesis. For example, using forced choice procedures it was shown that a fundamental mask of moderately high contrast with sub-threshold third harmonic test could be discriminated from the mask alone (Stromeyer & Klein, 1974; Nachmias & Weber, 1975; Tolhurst & Barfield, 1978; Lawden, 1983; Georgeson & Georgeson, 1987; Hess & Pointer, 1983). Adaptation studies, which used complex gratings containing more than one spatial frequency or orientation, also questioned the independence of channels. They showed that less threshold elevation was produced by composite adapting stimuli at either spatial frequency or orientation, than when the individual components were adapted to alone (Tolhurst, 1972; Nachmias et al., 1973; Stecher et al., 1973; Tolhurst & Barfield, 1978; Georgeson, 1980). These two groups of studies showed that the detection of a test component could be markedly affected by the presentation of a component that differed substantially from the test in its spatial characteristics. It was difficult to see how these results could be reconciled with what was known about the spatial bandwidth of detecting mechanisms without acknowledging interactions between these mechanisms.

Within-channel models of masking emerged from the masking data produced when test and mask were similar in their spatio-temporal characteristics. As already discussed, the nonlinear transducer function describing within-channel masking (Eqn. 1.1) is sigmoidal; it has an accelerative region followed by a compressive region. It predicts that as the input of mask contrast to the test detecting mechanism is increased, the mask first facilitates and then suppresses test detection. Such 'dipper' functions have been shown many times psychophysically when the test and mask are very similar. Contradicting the within-channel concept of masking however, are TvC (test contrast versus mask contrast) functions measured when tests and masks have differed in their spatial or temporal characteristics. These have not supported the idea that masking and facilitation are due to a common process. For example, Georgeson and Georgeson (1987) showed that if the test and mask were not presented at exactly the same temporal position (the mask was presented either slightly before or slightly after the test), then though masking remained intact, facilitation was abolished. In addition, when test and mask differ substantially in spatial frequency (Tolhurst & Barfield, 1978; Legge & Foley, 1980; Ross & Speed, 1991) or orientation (Ross & Speed, 1991; Ross et al., 1993; Foley, 1994a;

Zenger & Sagi, 1996; Foley & Chen, 1997, and see next section), masking again remains intact whilst facilitation is abolished.

1.5 Foley's tests of the within-channel masking model

Foley (1994a) (see also Foley & Boynton, 1994) tested the adequacy of the within-channel model at predicting the effects of simultaneous masking on test detection when test and mask differed (widely) in orientation. A 2-AFC simultaneous masking paradigm used grating stimuli that were matched in spatial frequency and temporal duration at 2cpd and 33ms respectively. The effect of a variable contrast full-field mask on the detection of a Gabor test was manipulated by increasing the orientation difference between the two stimuli from 0deg to 90deg. The within-channel model predicts that an increase in this difference causes the dipper function to be shifted horizontally by a multiplicative constant whilst remaining unchanged in shape. Foley found however, that facilitation only occurred when the difference was small (up to around 15deg) and substantial masking occurred at all relative orientations including orthogonal. Further, as the difference was increased, the masking range of the TvC function changed from being slightly concave downward to being shallower and more linear. As discussed in the last section, similar results have also been reported by others (e.g. Ross & Speed, 1991; Ross et al., 1993).

In a second test of the within-channel model, Foley (1994a) superimposed vertical test and variable contrast pedestal mask Gabors with a full-field fixed contrast mask whose orientation was manipulated between 0deg and 90deg. The within-channel model predicts that regardless of orientation, adding the fixed contrast mask causes the TvC function to be shifted horizontally by an additive constant. Because the model combines test and mask inputs by a linear operation, it predicts that the addition of the mask would not change the shape of the function. However, the functions measured with and without the mask were of a different shape, causing them to cross. The fixed contrast mask produced an elevation in test thresholds when pedestal contrasts were low whilst when they were high, thresholds were substantially lowered. The fact that the masking function still had a dip when the fixed contrast mask contributed to it was of particular interest. This dip has been shown by myself (see chapter 5) and by others (e.g. Ross et al., 1993). If the within-channel model of masking was correct, then it would be expected that the fixed-contrast mask would push the operating characteristic of the transducer past its accelerative range and into its compressive range. The consequence of this would be that the masking function would have no dip (since the dip is due to the transducer's accelerative range).

Foley (1994a) showed that though the within-channel model did well at predicting the masking that was produced when test and mask were matched in orientation, its predictions were substantially weaker when their orientations differed. The main reason for the poor fits was the model's feature of linearly combining the inputs from test and mask components in a single channel. Foley (1994a) noted that a better fit could have been achieved with the model if it had been assumed that, as the difference between test and mask orientation was increased, different channels with different parameter values were responsible for test detection (that is, the test detecting mechanism changed even though the test did not). Foley was confident that this was not the case however, because he found no evidence that the percept of the stimuli changed or scallops appeared in the TvC functions as the difference between test and mask orientations was increased.

Rather, his data suggested that the test and mask stimuli were processed by the same mechanism (or a set of similarly tuned mechanisms) and a well fitting model would need to encapsulate channel interactions.

1.6 Foley's models of human pattern masking

As discussed in the previous section, Foley (1994a) showed that mask components remote in orientation to a test component could substantially raise the detection threshold of the test pattern. Physiological studies have also provided support for cross-orientation suppression. Typically in such a study, a neuron's response to superimposed test and mask gratings is measured. The test's orientation is matched to that which is preferred by the neuron, whilst the orientation of the mask is orthogonal to that of the test. Though a typical neuron in the striate cortex is highly responsive to a grating having the orientation preferred by that neuron, it shows no contrast response to a grating orthogonal to that preferred (e.g. Campbell et al., 1968). However, of particular interest is that neuronal response to an optimally oriented test is suppressed by the superposition of an orthogonal mask (Bishop et al., 1973; Morrone et al., 1982; Bonds, 1989; Bauman & Bonds, 1991; Crook & Eysel, 1992; DeAngelis et al., 1992; Sengpiel & Blakemore, 1994; Carandini et al., 1997, Crook et al., 1998; Sengpiel et al., 1998; Allison et al., 2001). Such suppression has also been shown when test and mask are remote in spatial or temporal frequency (Morrone et al., 1982; Bonds, 1989, 1991; Bauman & Bonds, 1991; DeAngelis et al., 1992; Geisler & Albrecht, 1992; Allison et al., 2001) or spatial position (see appendices 1 and 2 for reviews).

Cross-orientation suppression has often been referred to as 'cross-orientation inhibition' (Morrone at al., 1982). This is because it is thought that a V1 neuron, having a preferred orientation, is inhibited by the responses of a 'pool' of V1 neurons with different preferred orientations; the fundamental tenet of this being that channels interact. This mutual inhibition between neurons (Robson, 1988; Bonds, 1989) has been shown to be divisive (Albrecht & Geisler, 1991; Heeger, 1992a), meaning that masks have the same affect as dividing test contrast by a weight which is proportional to the contrast of the masks. In considering interactions between cat striate neurons, Heeger (1992a) formulated the first model characterised by divisive inhibition operating across spatial frequency and orientation selective units. In his normalisation model, each neuron has an accelerating nonlinearity but is also divisively inhibited by a pool of responses of other neurons.

Inspired by Heeger's (1992a) work, Foley (1994a) attempted to describe the processes underlying the orientation masking data that he had gathered. He did this by extending the within-channel model of masking in light of what had been speculated about divisive inhibitory pooling.

The equation describing a within-channel model of masking is repeated here from earlier:

$$r_t = c_t^p / \left(z^q + c_t^q\right) \qquad (1.5)$$

where r_t is the response output of the nonlinearity of the test detecting mechanism and c_t is the contrast response of a linear filter selective for the test component. Z is the semi-saturation constant. A property of this model is that the mask can only produce masking when it excites the detecting mechanism.

The within-channel model is easily extended to accommodate cross-channel effects by introducing further terms γ_i. These represent the contribution of the contrast of the mask components which do not excite the test detecting channel but which do suppress the test. The model can be expressed as:

$$r_t = c_t^p / (z^q + c_t^q + \sum_i(\gamma_i^q)) \qquad (1.6)$$

Note that this and the within-channel model are identical when the mask components excite only the test detecting mechanism. Models having the property that masking can be produced by masks that do not excite the test detecting mechanism are referred to here as 'cross-channel' models of masking. A related, but less formal term that is used here is 'remote masking'. This is used without reference to any specific model, but to describe situations in which mask and test stimuli are so different that it seems unlikely that they would stimulate a common detecting mechanism.

Foley (1994a) formulated two models of cross-channel masking in attempting to gain insight into the data described in the last section. He referred to these as model 2 and model 3. In this thesis they are referred to as Foley 2 and Foley 3 and are discussed and described mathematically and pictorially in chapter 5. In his models, and in their developments (chapter 5), levels of masking are predicted by the Fechnerian notion (e.g. see Boynton & Foley, 1999) that to detect a test in the presence of a mask, the response of the mechanism used for detection must be greater than its response to the mask presented alone by a constant increment (termed k in the modelling). This assumption has been used many times previously to model contrast discrimination (Legge & Foley, 1980; Foley & Yang, 1991; Foley, 1994a; Foley & Chen, 1997; Boynton & Foley, 1999; Boynton et al. 1999; Dannemiller & Stephens, 2000). The models were not designed to account for effects of phase and spatial frequency, or probability summation across mechanisms. They attempt to describe data as a set of mathematical functions and are specified in terms of experimental stimulus components, rather than the outputs of multiple psychophysically or neurophysiologically derived filters. This is unlike some other models of divisive inhibitory interactions, which have quantified their mechanisms by assuming the number and bandwidths of spatial filters (Teo & Heeger, 1994; Watson & Solomon, 1997). Foley's (1994) models postulate that pattern mechanisms have both a linearly summed excitatory input and a divisive inhibitory input. If there is no mask present, the response of the test detecting mechanism is its rectified excitation raised to an exponent (p) and then divided by a divisive inhibitory input limited at low levels by an additive constant (z). Foley 2 and 3 differ in their treatment of test and mask inputs to the divisive inhibitory pool. The inputs are raised to an exponent (q) before summation by Foley 3 (as in Eqn. 1.6) and after summation by Foley 2 (not shown here).

Foley fit his models to data gathered from several pattern masking experiments, which enabled the values of their (free) parameters to be determined. From this analysis it is clear that in his model, the orientation tuning of divisive inhibition is much broader than that of excitation. On the other hand, it has been noted that the breadths of excitatory and inhibitory tuning are approximately the same at high spatial frequencies (Boynton & Foley, 1999). Like Georgeson and Georgeson (1987), Foley concluded that different processes underlie masking and facilitation; he deduced that facilitation comes from masker excitation whilst masking comes from divisive inhibition. This fits well with observations that facilitation is far more dependent on the spatio-temporal characteristics of the stimuli than is masking (assuming that masking pools across many mechanisms).

Foley's models have been used as the basis for describing a wide variety of masking data by Foley and co-workers (Foley, 1994a, 1994b, 1996; Foley & Crites, 1997; Foley & Chen, 1997; Foley & Schwartz, 1998; Boynton & Foley, 1999; Foley & Chen, 1999; Chen et al., 2000a, 2000b), and by others (for example: Watson & Solomon, 1997; Snowden & Hammett, 1998, Itti et al., 2000).

1.7 Other models of divisive inhibition

From the initial proposals (e.g. Blakemore et al., 1970) to quantitative models (see below), lateral (or intracortical) inhibition between cortical neurons is the dominant theory of filter interactions in early pattern analysis. All of the models are similar in that they describe mechanisms as receiving a broadly tuned divisive inhibitory input (or 'pool') with an excitatory input from a linear receptive field. Some differ in their architecture and many are concerned with describing different phenomena.

Foley's (1994) models, which were published after Ross and Speed's (1991) model, are the most developed that have been formulated to predict masking data produced with grating stimuli. The majority of the work in this thesis is concerned with measuring masking using grating stimuli. For this reason, Foley's models are considered in some detail here.

Other models incorporating a divisive inhibitory pool differ from Foley's in where their emphasis is placed. For example, the models of Teo and Heeger (1994) and Watson and Solomon (1997) are image driven, that is, they can accept any grey level image as input. Olzak and Thomas' (1991) predicts spatial discrimination data whilst Itti et al.'s (2000) is concerned with simultaneously predicting contrast and spatial discrimination data. All of these are feedforward models. In contrast, a feedback network was originally considered by Heeger (1991, 1992a) in describing cat neuronal responses (though his mathematical implementation was, in fact, feedforward). Wilson and Humanski (1993) focused on the time course of adaptation and presented what they thought was evidence for a feedback model. However, Hammett and Snowden (1995) pointed to flaws in the stimulus design used by Wilson and Humanski (1993), and have presented convincing evidence that their data is not indicative of feedback architecture.

1.8 Why cross-channel models of masking are important

It can no longer be considered that channels act independently, but why do they interact? Interconnections have been credited for important computations in form and motion vision, texture perception and stereopsis. One purpose of channel interactions that has received a great deal of attention is contrast gain control. An important phenomenon that has driven work in this area is:

- The contrast response functions of cortical neurons saturate but their selectivity remains invariant with contrast (Albrecht & Hamilton, 1982; Sclar & Freeman, 1982; Li & Creutzfeldt, 1984).

The reason for this has been shown to be that:

- Neurons do not saturate at a fixed firing rate determined by the electrical properties of the cell, but at a rate that is stimulus dependent (Maffei et al., 1973; Dean, 1981; Albrecht & Hamilton, 1982).

The effective contrast of stimulus patterns on the visual system is referred to as contrast gain (e.g. Foley & Chen, 1997). Its control means that the limited dynamic ranges of cortical neurons do not compromise contrast vision. That is, 'population codes' are protected despite neuronal saturation. One example of this is the (almost) contrast invariance of orientation tuning (Sclar & Freeman, 1982; Carandini et al., 1997). The models described in the last two sections can explain contrast gain control. They describe how a channel's response to a preferred stimulus can be suppressed (divisively inhibited) by superimposing an additional stimulus. The suppression has been shown to be largely nonspecific to stimulus dimensions (e.g. spatial frequency, orientation etc.). Thus, the excitation of every channel is divided by a quantity that is proportional to the activity of a large number of channels. This dynamic mechanism normalises cortical responsiveness around an area of mean contrast. Studies have shown how adaptation can influence the operation of gain control. For example, adapting to a high contrast grating causes a shift in the semi-saturation of the S-shaped contrast-response function of cortical neurons so that they are centred over higher contrasts (Ohzawa et al., 1982, 1985; Albrecht et al., 1984; Saul & Cynader, 1989a, 1989b; Bonds, 1991). Contrast gain control is thought to originate in V1 because lateral geniculate responses do not reflect its characteristics (Movshon & Lennie, 1979; Ohzawa et al. 1985, Albrecht & Hamilton, 1982; Li & Creutzfeldt, 1984). For example, visual neurons in the LGN do not saturate.

Non-linear normalization or gain control circuits have also been shown to play major roles in refining spatial frequency (e.g. Baumann & Bonds, 1991) and orientation (e.g. Sillito, 1975, 1979) specificities of cat visual mechanisms (for reviews, see Somers, 1995, Vidyasagar et al., 1996, Sompolinsky & Shapley, 1997, and Ferster & Miller, 2000). Moreover, it has been theorised that they may be responsible for many aspects of contrast detection and discrimination in masking and adaptation paradigms, and in other psychophysical tasks, such as texture (Graham, 1991; Graham et al., 1992; Bergen & Landy, 1991) and spatial discriminations (Olzak & Thomas, 1991; 1992a; 1992b and Clifford et al., 2001; though see Westheimer & Gee, 2002).

1.9 Thesis Preview

The initial aim of the work in this thesis was to investigate effects of spatially remote masks. This was to provide insights into how channels tuned to different spatial frequencies and orientations might interact. This led to tests of masking models, each of which feature a divisive inhibitory pool. The range of inputs into this pool and the summation rules applied within it were investigated. Fig. 1.1 shows how the sections of this thesis are related.

The experimental sections are chapters 3 to 7 and appendix 1. Chapter 2 describes the general methods used in the experiments.

Chapter 3 shows several remote masking effects. Of particular interest was a final experiment in which the masking from patterns with two components that were remote in orientation from the test, appeared to be due to the linear sum of their component contrasts. Chapter 4 builds on this finding and shows that the result holds across a range of mask contrasts. Foley 2 fits the data well whilst Foley 3 is rejected. In chapter 5, a new model is introduced which is a hybrid of Foley's two models. This model reconciles Foley's data with those in chapter 4. An experiment is then devised which provides further tests for the models. Important differences between observers' data prompts the development of a series of modifications to the models, which improve the fit to the data.

The aim of chapter 6 was to investigate whether Phillips and Wilson (1984), in their classic study of orientation tuning, may have mis-estimated channel bandwidth because of their use of a within-channel masking model. Orientation masking was measured for a greater range of test-mask orientation differences than had previously been investigated. Substantial phase-insensitive masking occurs at all relative orientations and evidence was provided that is consistent with this being due to cross-channel rather than within-channel masking. The cross-channel masking effects clearly show the need to reappraise the bandwidth estimates made by Phillips and Wilson (1984).

Chapter 7 extended the line of enquiry developed in chapter 4, which was shown there to distinguish between cross-channel models of masking. By devising related experiments using the adaptation paradigm, the aim here was to provide further tests of whether the masking was cross-channel or within-channel. Unfortunately, this line of enquiry delivered much less insight than the masking experiments of the earlier chapters.

Appendix 1 sets out some ideas and pilot work concerning how the enquiries of chapters 3 to 5 could be extended to investigate masking effects across spatial position.

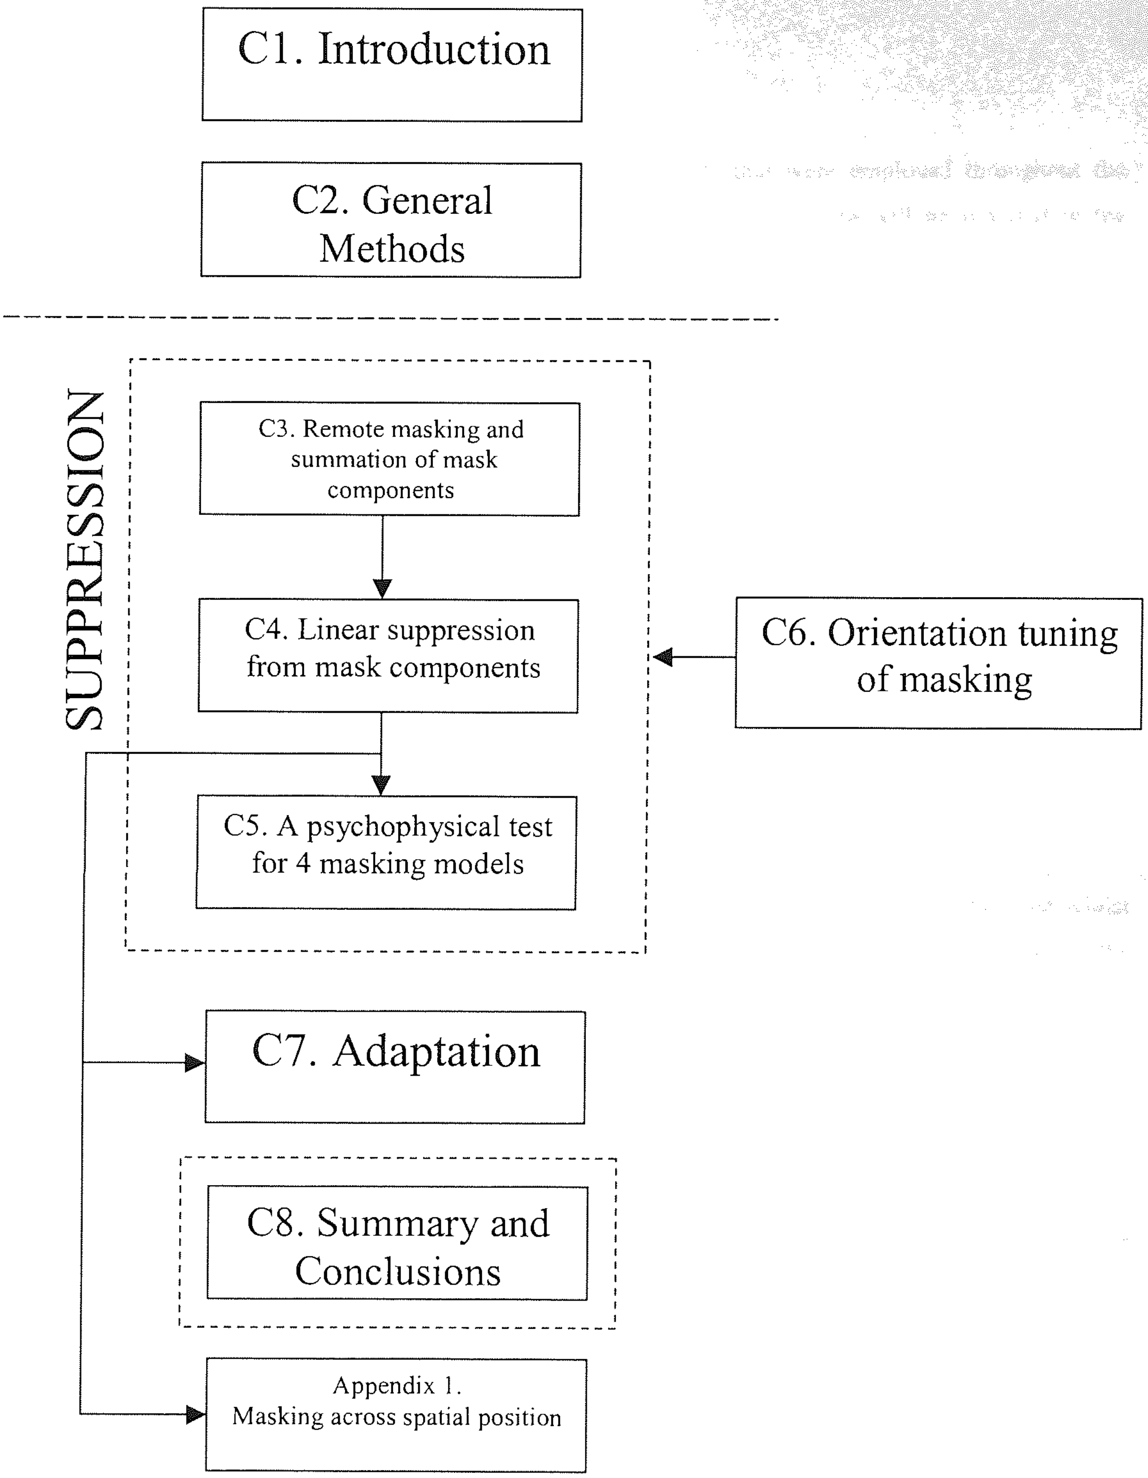

Fig. 1.1. A schematic illustration showing how the sections of this thesis are related.

Chapter Two: General Methods of Research

2.1 Chapter Introduction

This chapter provides a description of the general methods that were employed throughout the experimental work presented in the forthcoming chapters. Deviations from these will be reported in the relevant sections.

2.2 Stimuli

The stimuli were patches of luminance modulation. Each stimulus was the product of a signal and a windowing function. The signal was either a sinusoidal grating, or a plaid made from the sum of two sinusoidal gratings. A fixation point was positioned at the centre of the display and was visible throughout all of the experiments. All of the stimuli were centred on this point and the majority of the sinusoidal patterns used were in sine phase with this point (exceptions to this are noted in the relevant sections). All tests and masks were circular whereas adapters covered the whole of the rectangular display screen. Brief descriptions follow of the types of stimulus pattern and windowing that were used.

Sinusoidal gratings

Sinusoidal gratings are visual patterns whose luminance varies sinusoidally in one direction whilst remaining constant along any line running perpendicular to that direction. The luminance profile of the sinusoid can be formally expressed as:

$$\mathrm{L}(x,y) = \mathrm{L}_0[1 + m.\cos(2\pi fx + \phi)] \qquad (2.1)$$

where x and y are positional coordinates, L is luminance, and L_0 is mean luminance. The variables f and ϕ are the sinusoid's frequency and phase respectively. The number of times a sinusoidal grating's wavelength is repeated in a unit distance is its spatial frequency. Most commonly, the unit of distance is a degree of visual angle, thus spatial frequency is described in cycles per degree (cpd). Spatial phase is the position of a grating's sinusoidal luminance profile relative to a fixed point (origin). When a grating's peak luminance occurs at the origin, the grating is sometimes described as being in cosine phase (with the origin) and when it occurs at $\pi/2$ radians from the origin, it is sometimes described as being in sine phase.

The amplitude (half peak-to-peak) of the sinusoidal grating defined in Eqn. 2.1 is L_0m since its maximum and minimum luminance is $L_0(1 + m)$ and $L_0(1 - m)$ respectively. The Michelson contrast of a sine wave grating is the amplitude divided by the mean luminance (L_0), giving m. Contrasts are expressed either in % or dB where contrast in dB = 20log (100m). In most cases in this thesis, fixed contrasts (e.g. of fixed contrast masks or adapters) are referred to in percent, whilst variable contrasts (e.g. of thresholds or differences in thresholds) are referred to in dB.

Plaids

Whilst gratings are 1-D variations in luminance, plaids are 2-D variations in luminance. The plaid patterns that were used in the following experiments contained two superimposed sinusoidal components that were matched in spatial frequency, phase, contrast and their temporal characteristics. In all but one of the experiments (chapter 4, section 4.9), the components differed in orientation by 90degrees. Tests and adapters were always gratings whilst masks were either gratings or plaids.

Notation for grating and plaid stimuli

The main ways in which the sinusoidal grating and plaid stimuli were varied throughout the experiments of this thesis was in terms of the spatial frequency and orientation of their components. A shorthand has been adopted to describe the grating and plaid stimuli throughout. For example, a grating with a spatial frequency of 1 cycle per degree and oriented at -45degrees is expressed as a -45deg 1cpd grating, whilst a plaid with one component oriented at -45degrees and the other oriented at +45degrees, both with a spatial frequency of 3 cycles per degree, is expressed as a ±45deg 3cpd plaid.

When reference to stimulus spatial frequency is made, this is to the spatial frequency of the sine-wave components prior to them being multiplied by a windowing function. Grating orientations are given as the angle to which the grating is oriented counter-clockwise to vertical. Contrasts refer to the pattern's contrast prior to enveloping.

Spatial truncation of sinusoids

If the windowing of a stimulus patch of sinusoidal grating is too 'sharp' (abrupt or very steep), then at its edge there will be artefacts containing spatial frequencies that are quite distant from the nominal spatial frequency content of the patch. Since it was desirable to avoid spatial frequency artefacts in the stimuli, where practical, the experiments of this thesis used stimuli that were smoothly windowed with either a raised cosine or Gaussian function (for information on these types of windowing, refer to Graham, 1989, pp 47-66). For a small minority of experiments, it was more desirable to use stimuli with abrupt edges; it is noted where this was the case.

Size of stimuli

Generally, stimuli were one of two sizes (discrepancies from these will be described in the relevant sections). The larger stimuli measured 4.4° at full-width, half-height and were curtailed with a raised-cosine envelope that measured 0.3° at each edge (the diameter of the plateaus was 76% of the total patch diameter;

1cpd stimuli filled the plateaus with 3 ½ cycles). The smaller stimuli measured 1.7° at full-width, half-height and were windowed with a gaussian. (With 1cpd, these stimuli had 1.7 cycles at half-height).

Temporal envelope of stimuli

In most of the experiments in this thesis, the temporal enveloping of the test and mask stimuli was matched. The stimuli were usually presented for 100ms and temporally curtailed with a rectangular window function (step-edge on and off). It is noted in the relevant sections where this was not the case. The brief temporal duration of the stimuli means that it is unlikely the observers moved their gaze during stimulus presentation.

Some experimental designs required that other temporal durations and types of temporal windowing were used in attempts to stimulate specific temporal channels (i.e. transient and sustained), or to replicate the methods of a particularly relevant study. Discussion has been provided where these have been used.

2.3 Apparatus

The stimuli were generated using a graphics system that consisted of a Dell P102t Pentium computer with a Cambridge Research Systems VSG 2/3 graphics board, and either an Eizo F553-M, a Sony Trinitron Multiscan 200PS, or a Sony Trinitron Multiscan 20seII monitor. These monitors had mean luminances of 65.8 cd/m^2, 69.8 cd/m^2 and 60.6 cd/m^2, respectively, and their framerate was 120Hz. The VSG was operating in pseudo-12 bit mode. A chin and headrest was used to keep the observer's head steady and at a fixed distance from the display. In the absence of stimuli, the display was uniform except for the small, dark, fixation square (two pixel by two pixel) at the centre. Look-up tables (LUTs) were used to gamma correct the output of the grating generator across the entire contrast range.

2.4 Observers

A complete set of data for all experimental conditions was obtained from the author, DJH. Data were gathered from a further 6 observers. TSM was the author's academic supervisor. Both DJH and TSM were highly practiced psychophysical observers and were aware of the purpose of each experiment. They were familiar with all stimuli and response procedures. ADP was a research fellow in psychophysics; practised in psychophysical observations relating mainly to stereopsis. RJS was a postgraduate in psychophysics, whilst JFT and CEB were postgraduates in other scientific disciplines. GR was an Optometry undergraduate. RJS, JFT, CEB and GR were psychophysically unpractised and, with ADP, were naïve to the purpose of any of the experimental conditions. Each observer had normal or optically corrected to normal vision. The contrast sensitivity function of the principle observer (DJH) was tested and found to be normal.

2.5 Procedure

Observers sat before a monitor screen in a darkened room with their heads comfortably positioned in the chin and headrest. They fixated on the fixation point throughout all of the experimental sessions and all stimuli were viewed binocularly. The distance between the observer and the monitor was 57cm, 114cm or 228cm. The shorter distance was used when particularly large stimuli were required (for example, the annuli described in appendix 1). The larger distance was used when the stimuli's spatial frequency was high; increasing the distance allowed each stimulus cycle to be represented by more pixels.

In all experiments, the detection threshold of a stimulus was measured. The terms 'test' and 'target' are used interchangeably to describe this stimulus. A temporal two-alternative forced-choice technique (Nachmias & Sansbury, 1974; Nachmias & Weber, 1975) was used to measure the test thresholds. The test was presented pseudo-randomly in either of the intervals and, if the effects of a spatio-temporally superimposed mask were being investigated, the mask was presented in both intervals. The task of the observer was to determine which interval contained the test. Intervals within a trial were separated by 500ms, which is a common duration used for this parameter in the literature; also refer to chapter 7, section 7.2 for experiments and discussion related to this. Auditory feedback was one of two different pitched sounds, which signalled whether the observer was correct or incorrect.

A three-down, one-up staircase procedure (Wetherill & Levitt, 1965; Meese, 1995) varied the contrast of the test stimuli in logarithmic steps. Thresholds were obtained using a computer-generated version of the randomized double staircase measurement technique introduced by Cornsweet (1962) in that for each stimulus condition, data were collapsed from a pair of interleaved staircases. Usually, each staircase terminated after 10 or 12 contrast reversals. The thresholds were estimated using probit analysis, as were the standard errors (McKee et al., 1985). After discarding trials from the initial converging stage of the staircases (Meese, 1995), 10 or 12 reversals for 2 staircases usually gave around 80 to 100 trials for the probit analysis. In the rare instance that a threshold measurement had a standard error exceeding 3dB, it was rejected and repeated. This criterion was set before data gathering began.

Conditions were interleaved either within or across experimental sessions. When the conditions were interleaved within the sessions, either: (i) the condition run on each trial was chosen pseudorandomly, or (ii) trials of each condition (i.e. all those from a 'staircase pair') were blocked together and these blocks were interleaved pseudorandomly. Interleaving across sessions was also pseudorandom.

Threshold measurements were usually taken from five experimental sessions (from which the mean was calculated). (It is noted where the number of replications differed from this). The standard error bars in the figures correspond to ± 1 standard error of these means.

2.6 Experimental paradigms

Masking has been used in most of the investigations described in this thesis. In some of the experiments in chapter 7, adaptation was used. The adaptation paradigm is discussed at length in chapter 7. General points that apply throughout the thesis concerning the masking procedure that was used are discussed below.

Except for some of the experiments described in chapter 7 and in the appendices, all of the masking experiments adopted a simultaneous masking paradigm where test and mask were spatio-temporally superimposed. Prior to windowing, the test was always a sinusoidal grating, whilst the masks were either sinusoidal gratings or plaids constructed from two sinusoidal gratings. The term 'pedestal' is used to describe a mask that is matched in spatial frequency, orientation, phase and windowing to a spatio-temporally superimposed test pattern. The contrast of a mask is reported either as its component contrast or its stimulus (peak-to-peak) contrast (which is the sum of the mask's component contrasts). The main reason for this, was that in many of the experiments the masking produced by a grating and a plaid were compared when they were matched either in component or stimulus contrast.

Throughout this thesis, the term 'masking' is used to describe a decrease in the detectability of a test arising from the presence of a mask, whilst 'facilitation' is used to describe an improvement in the detectability of a test arising from the presence of a mask. Functions produced by plotting the effect of masks on test contrast, as a function of some stimulus parameter, are termed masking functions whether they show masking or facilitation. In many of the experiments, test contrast was measured as a function of mask contrast. Such functions are termed throughout this thesis as TvC functions.

Chapter Three:

Simultaneous Masking by Plaids and Gratings: Threshold Elevation and Threshold Facilitation

3.1 Chapter Introduction

There is abundant psychophysical and neurophysiological evidence, which is consistent with the view that human vision employs channels that are each responsive to a limited range of spatial frequencies and orientations (e.g. see chapters 1, 4 and 6, and Graham, 1989, and DeValois & DeValois, 1988, for reviews). One widely used tool that has been used to investigate channel selectivity, is masking. About twenty years ago, masking effects were often described by within-channel masking models. In these models, masks only affect the detectability of the test stimulus if they excite the test detecting mechanism. However, a more recent body of work suggests that at least some masking arises from suppressive input to the detecting mechanism. In these cross-channel masking models (e.g. Foley, 1994a), the orientation and spatial frequency of mask components can be so remote from the detecting mechanism that they provide no excitatory stimulation, but cause only suppression. I refer to the masking that is produced in this situation as 'remote masking'. A primary aim of this chapter is to investigate the extent of remote masking for a range of test spatial frequencies. Although a reinterpretation of masking functions from the 'within-channel' masking literature might go some way towards achieving this goal, these earlier studies tended to explore a fairly narrow range of mask spatial frequencies. For example, in Harvey and Doan (1990) and Daugman (1984), the difference between the spatial frequency of mask and test was never greater than 2 octaves. The range used by Wilson et al. (1983) was somewhat greater than this (up to about 3 octaves) but in that study the orientations of test and mask were fixed. In the experiments shown in this chapter, large ranges of spatial relationships were investigated which have not been explored before. The spatial frequency of the mask was held constant (either at a low, medium or high value) whilst the test and mask spatial frequencies differed by up to 4 ½ octaves. In addition to this, the orientations of test and masks differed by up to 45deg.

Another 'remote masking' study that is of particular relevance to this chapter is that of Derrington and Henning (1984). Using 3cpd test and mask components, they compared the masking effects of a +45deg grating and a ±45deg plaid on a vertical test stimulus. The overall contrast of the plaid was twice that of the grating, though the level of masking that it produced was considerably more than a factor of 2 greater than that of the grating. (The relationship between the contrast of a grating mask and the masking that it produces has been well known for some time and is considerably less than that found by Derrington and Henning when the mask was changed from a grating to a plaid; for example, see Nachmias & Sansbury, 1974; Legge & Foley, 1980; Wilson, 1980; Wilson et al., 1983; Phillips & Wilson, 1984).

Derrington and Henning's study was performed before the publication of experiments and models addressing cross-channel masking. They attempted to describe the substantial masking produced by the plaid in terms of within-channel masking. The motivation for this work was to consider whether nonlinear distortion products could explain the data. Nonlinear distortion products are artifacts that can be introduced by earlier luminance nonlinearities (i.e. prior to orientation tuning) and produce spatial frequency components at twice the pattern's modulation frequency. For the plaid masks, the distortion products would have been vertical but with a higher spatial frequency than that of the mask's components. To investigate whether the substantial masking produced by the plaid could be explained in terms of distortion products, they extended their first experiment by measuring the spatial frequency tuning of the masking produced by the plaid. For various orientations of plaid component they found that the peak masking did not occur when the test spatial frequency was matched to that of the supposed distortion product. They concluded that their effects couldn't be explained in any simple way by such visual nonlinearities.

In summary, the particularly interesting findings of Derrington and Henning are: (i) a grating and plaid mask produced a substantial amount of masking when test and mask were remote in orientation and spatial frequency, (ii) a plaid mask produced far more masking than what was expected from that produced by a grating mask, and (iii) distortion products were not responsible for (i) or (ii).

A great deal more is now known about masking than what was known at the time of Derrington and Henning's work. In light of the evidence for cross-channel masking models, the substantial remote masking appears to be suggestive of channel interactions rather than any kind of within-channel effect. Even in light of cross-channel models however, it is still difficult to see why the plaid would produce so much more masking than the grating. Further, it is unknown how extensive this curious effect is. Is it dependent upon, for example, the spatial frequency of the stimuli or the contrast of the mask? Derrington and Henning (1989) missed the opportunity to measure its dependence on test spatial frequency in their second experiment. They measured masking as a function of test spatial frequency but only for the plaid. If they had measured the masking from the grating as well as that from the plaid, this would have revealed something of the effect's extent.

As described above, the general aim of the experiments in this chapter was to use a large range of stimuli in an attempt to investigate remote masking effects that may have been missed previously. The spatial frequency of test and mask were varied whilst the patterns differed in orientation by 45deg. These experiments provided the opportunity to investigate the extent of Derrington and Henning's effect. This was the main aim of the chapter. The extent of the effect was investigated by measuring the masking produced not only from -45deg gratings, but also from ±45deg plaids.

There are two possible outcomes that have particular interest. These are the situations where:

1. The masking levels produced by one-component (grating) and two-component (plaid) remote patterns are about the same when they are matched in stimulus (peak-to-peak) contrast. This would indicate linear summation of the component contrasts in the cross-channel suppression process.
2. The masking produced by the plaid is greater than what would be expected from the masking produced by one of its components. This is what Derrington and Henning (1989) found. If the grating and plaid were matched in stimulus (peak-to-peak) contrast, the plaid would produce more masking than the grating.

I shall refer to these as linear-suppression and super-suppression, respectively.

3.2 Chapter General Methods

All stimuli were spatially curtailed by a raised-cosine function and temporally curtailed by a square wave. They measured 4.4° at full-width, half-height, and their plateaus measured 4.1°. They had a temporal duration of 100ms.

3.3 Experiment 1. Masking and facilitation as a function of test spatial frequency

Methods

For the main observer (DJH), mask spatial frequency was 1cpd whilst test spatial frequency ranged from 0.5cpd to 22.1cpd. Test gratings had a vertical orientation. Mask type (either no-mask, -45deg grating, or ±45deg plaid) was pseudo-randomly interleaved within experimental sessions. Test spatial frequency was pseudo-randomly varied across sessions. Mask components had 8% contrast.

DJH repeated the experimental conditions eight times and data from the first three replications were discarded. These runs were used as practice since it was the first time that the principal observer in the experiments of this thesis had performed in a psychophysical experiment.

The experiment was repeated with three naïve observers (JFT, ADP and RJS). Methods were the same as those used for DJH, except that a restricted range of test spatial frequencies were used (1.0, 4.0, 5.7 and 8.0 cpd) and the experimental conditions were repeated only six times. The first session for each observer was treated as a practice session and the data that it delivered were discarded. These changes lessened the time required to gather the data.

Results and Discussion

Figs. 3.1 and 3.2 show threshold elevation as a function of test spatial frequency for DJH and the three naïve observers, respectively. Functions for grating (filled symbols) and plaid (open symbols) masks are shown. The data were broadly comparable across observers.

Masking

A broad region of masking extended to spatial frequencies either side of the mask (arrow), consistent with earlier findings in which the orientation of the mask was closer to that of the test (Wilson et al., 1983). A notable feature of the data is that the greatest elevation did not occur when the grating and plaid masks were of the same spatial frequency as the test. See section 3.9 for a discussion of this.

Where masking was present, it was greater for the plaid mask than for the grating mask. This was consistent with the effect that Derrington and Henning (1989) reported a ±45deg 3cpd plaid and a +45deg 3cpd grating mask had on the detectability of a 3cpd vertical grating test. Of great interest however, was that though the elevation produced by the grating masks was similar to that found by Derrington and Henning (about 6dB), the differences between the grating and plaid masking was less in the present study than the 12dB difference found by them. Many experiments of this chapter investigate this difference further.

Facilitation

An unexpected feature of the masking data was the shallow region of facilitation, which appeared to be spatial frequency tuned. It was apparent when the spatial frequency of the test was about 2 to 4 octaves higher than that of the mask. Others have also varied both the orientation and spatial frequency difference between test and mask (e.g. Daugman, 1984; Derrington & Henning, 1989; Harvey & Doan, 1990). However, the range of their manipulations was not as extensive as that investigated here and so it seems likely that their studies would have missed this feature of the masking functions. The facilitation is explored further in the following sections.

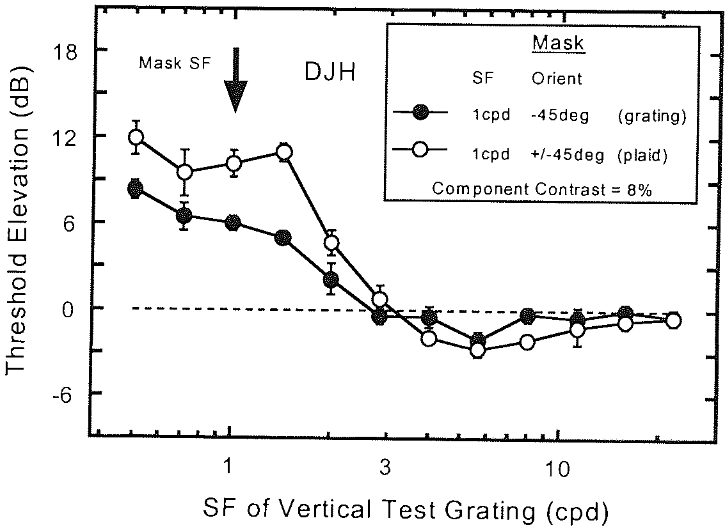

Fig. 3.1. Masking and facilitation as a function of test spatial frequency.

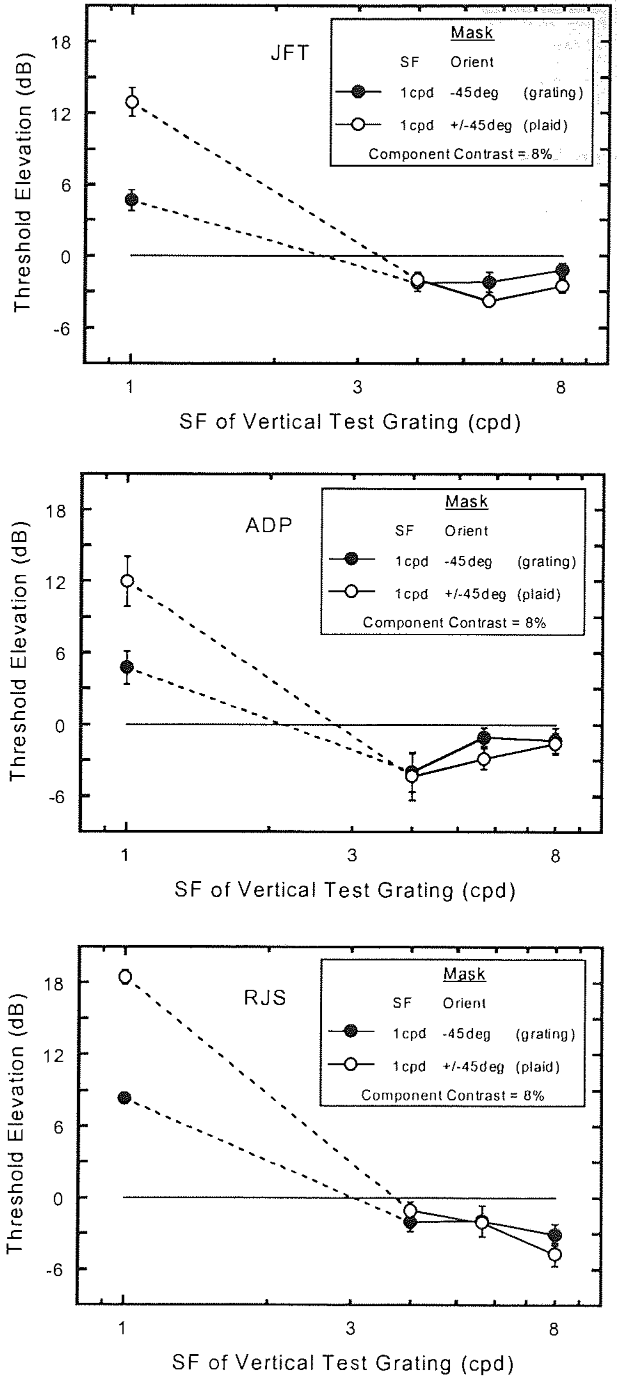

Fig. 3.2. Masking and facilitation as a function of test spatial frequency for three naïve observers.

3.4 Experiment 2. Control for adjacent pixel nonlinearities

Introduction and Methods

A raster monitor display refreshes its luminance in a horizontal direction, that is, the top line of pixels are refreshed first, then the second to top line, and so on. Though the nonlinearity existing between luminance and voltage was easy to rectify by means of look-up tables (i.e. gamma correction) there is another nonlinearity related to the monitor that should be considered which cannot be easily rectified. Adjacent (or spatial) pixel nonlinearity is the term used to describe the phenomenon whereby the luminance at which a pixel is presented is affected by the preceding pixel in the raster flow (e.g., see Lyons & Farrell, 1989; Mulligan & Stone, 1989; Naiman & Makous, 1992, Klein et al., 1996). The effect of this can be seen if adjacent pixels are alternated between low and high luminance. If the alternation is in the raster direction, then the mean luminance of the pattern will be lower than if the alternation is perpendicular to the raster direction. A consequence of this is that a pattern, which is composed of vertical, high spatial frequency, high contrast gratings, may have incorrect mean luminances and thus, distorted contrast.

As the previous experiment (section 3.3) used a raster monitor and vertically oriented high spatial frequency test components, the data may have been contaminated by adjacent pixel nonlinearities. To address this the experiment was repeated with vertical and horizontal tests (and with only plaid masks). (Test orientation was varied across sessions). If adjacent pixel nonlinearities had contaminated the data in the previous experiment, the orientation of the test would be expected to influence the functions when the test spatial frequency was high.

Results and Discussion

Fig. 3.3 shows the effect of plaid masks on the detectability of vertical (filled symbols) and horizontal (open symbols) tests at a range of test spatial frequencies. Whether the test was oriented vertically or horizontally had little effect on the function when test spatial frequency was above 3cpd. This indicates that the results of the previous experiment were not contaminated by adjacent pixel nonlinearities.

At low test spatial frequencies, masking was greater when the test was vertical than when it was horizontal. This may be due to anisotropies in masking or learning effects. Learning seems plausible. Prior to this experiment, the observer was well practiced at masking tasks that used a vertical test (the conditions of the experiment described in section 3.3 were replicated eight times prior to the start of this experiment). Dorais and Sagi (1997) found evidence for learning in masking experiments. They performed a simultaneous masking experiment in which test and mask stimuli were matched in their component spatial frequencies. The masking produced by ±45deg plaids on a horizontal test decreased with practice. When, after substantial practice, the test orientation was changed to vertical, the practice effects did not transfer and levels of masking returned to the level found for a horizontal test prior to practice. Zenger and Sagi (1996) suggested that reduction of masking caused by practice might be due to selective reduction of inhibitory gain in a contrast gain control mechanism.

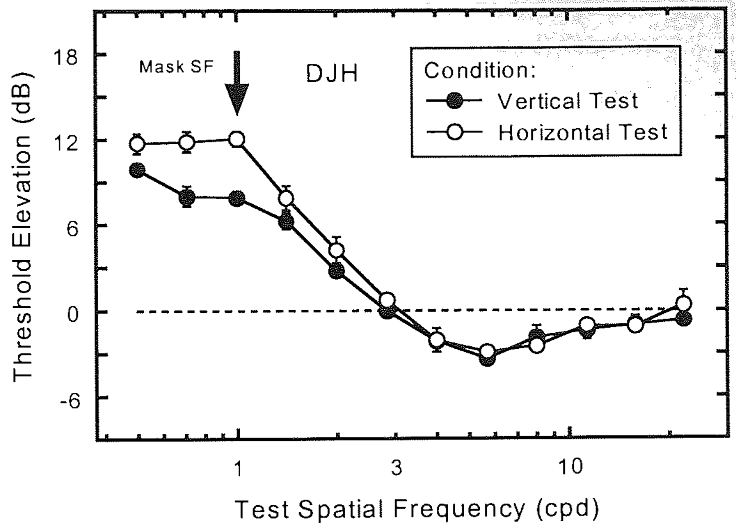

Fig. 3.3. Effects of test orientation on masking and facilitation.

Previous reports of facilitation

A distinct feature of the data from experiments 1 and 2 is the shallow region of facilitation. Phenomena of this type have been reported previously. For example, when test and mask are spatio-temporally superimposed, substantial facilitation occurs at low mask contrasts. This is highly dependent on the test and mask being similar in terms of spatial frequency (Legge & Foley, 1980; Ross & Speed, 1991), orientation (Campbell & Kulikowski, 1966; Ross & Speed, 1991; Foley, 1994a, Foley & Chen, 1997), phase (Foley, 1994b; Zenger & Sagi, 1996), temporal frequency (Boynton & Foley, 1999) and temporal position (Georgeson & Georgeson, 1987). (For an example using plaid masks, see Zenger & Sagi, 1996.) This facilitation is thought to reflect nonlinearities of the channel detecting the test pattern (see chapter 1). Facilitation also occurs at higher mask contrasts when tests and masks are either remote in spatial frequency (e.g. Stromeyer & Klein, 1974; Barfield & Tolhurst, 1975; Georgeson, 1975; Nachmias & Weber, 1975; DeValois, 1977; Tolhurst & Barfield, 1978; Kelly & Burbeck, 1980; Georgeson & Georgeson, 1987) or have adjacent spatial positions (Sagi & Hochstein, 1985; Dresp, 1993; Polat & Sagi, 1993, 1994; Kapadia et al., 1995; Polat & Norcia, 1996; Zenger & Sagi, 1996; Yu & Levi, 1997; Solomon et al., 1999). These occurrences of facilitation have been thought to result from interactions between channels with different spatial tuning; though some authors have interpreted them as being merely due to the masks providing subliminal excitation to the channels detecting the test patterns.

3.5 Experiment 3. Facilitation as a function of the orientation difference between test and mask

Introduction

In the preceding two sections, facilitation was shown to sometimes occur when the test and mask patterns differed in both spatial frequency and orientation. Tolhurst and Barfield (1978) found facilitation when test and mask patterns differed in spatial frequency but had the same orientation. They found it when the test spatial frequency was between one and two octaves higher than that of their 4.25cpd mask. The next experiment questions whether the facilitation shown in the previous sections and that shown by Tolhurst and Barfield are of the same origin.

Methods

Mask type (either no-mask or grating mask) was pseudo-randomly varied within experimental sessions. Grating masks had a spatial frequency of 1cpd, were oriented -45deg and had a contrast of 8%. Tests were single gratings whose orientation and spatial frequency were pseudo-randomly varied across sessions. Test orientation ranged from 0deg to -45deg in 5degree intervals. At each orientation, tests were presented at five spatial frequencies, separated by half-octave intervals.

Results and Discussion

Fig. 3.4 shows facilitation as a function of the orientation (across panels) and spatial frequency (within panels) difference between test and mask gratings. The top-left panel shows facilitation when test and mask differed in orientation by 45deg (like the experiments presented in the preceding sections) and the bottom-right panel shows facilitation when test and mask were matched in orientation (like Tolhurst & Barfield). When test and mask differed in orientation by 45deg, the peak magnitude of facilitation and the test spatial frequency at which it occurred was comparable to that reported in previous sections of this chapter. The facilitation formed a continuum in that it was evident at all orientation differences. Its magnitude consistently peaked at around 3dB, which was similar to that shown by Tolhurst and Barfield (when mask contrast was the same as that used here). Furthermore, as the difference between test and mask orientation decreased, the test spatial frequency at which functions crossed between masking and facilitation also tended to decrease (see Fig. 3.5). This was particularly interesting because Tolhurst and Barfield noted facilitation when test spatial frequency was between 1 and 2 octaves higher than that of the mask. My data show more facilitation in this area when test and mask are matched in orientation (like Tolhurst & Barfield) compared to when they differ in orientation.

In summary, both 1cpd (shown here) and 4.25cpd (shown by Tolhurst & Barfield) masks facilitate the detection of tests that have spatial frequencies over one octave higher than those of the masks. This similarity is consistent with the two instances of facilitation being the result of similar underlying processes.

The facilitation shown with 1cpd stimuli generalises across orientations so it seems likely that that shown by Tolhurst and Barfield does also.

The following section shows the effect of mask contrast on facilitation and masking. Tolhurst and Barfield found that contrast was important for facilitation. If the facilitation that I have found is of the same origin as theirs, then contrast should be important here as well. The effect of contrast on masking was examined to determine whether important differences between my data and Derrington and Henning's (see next section) were due to differences in mask contrast.

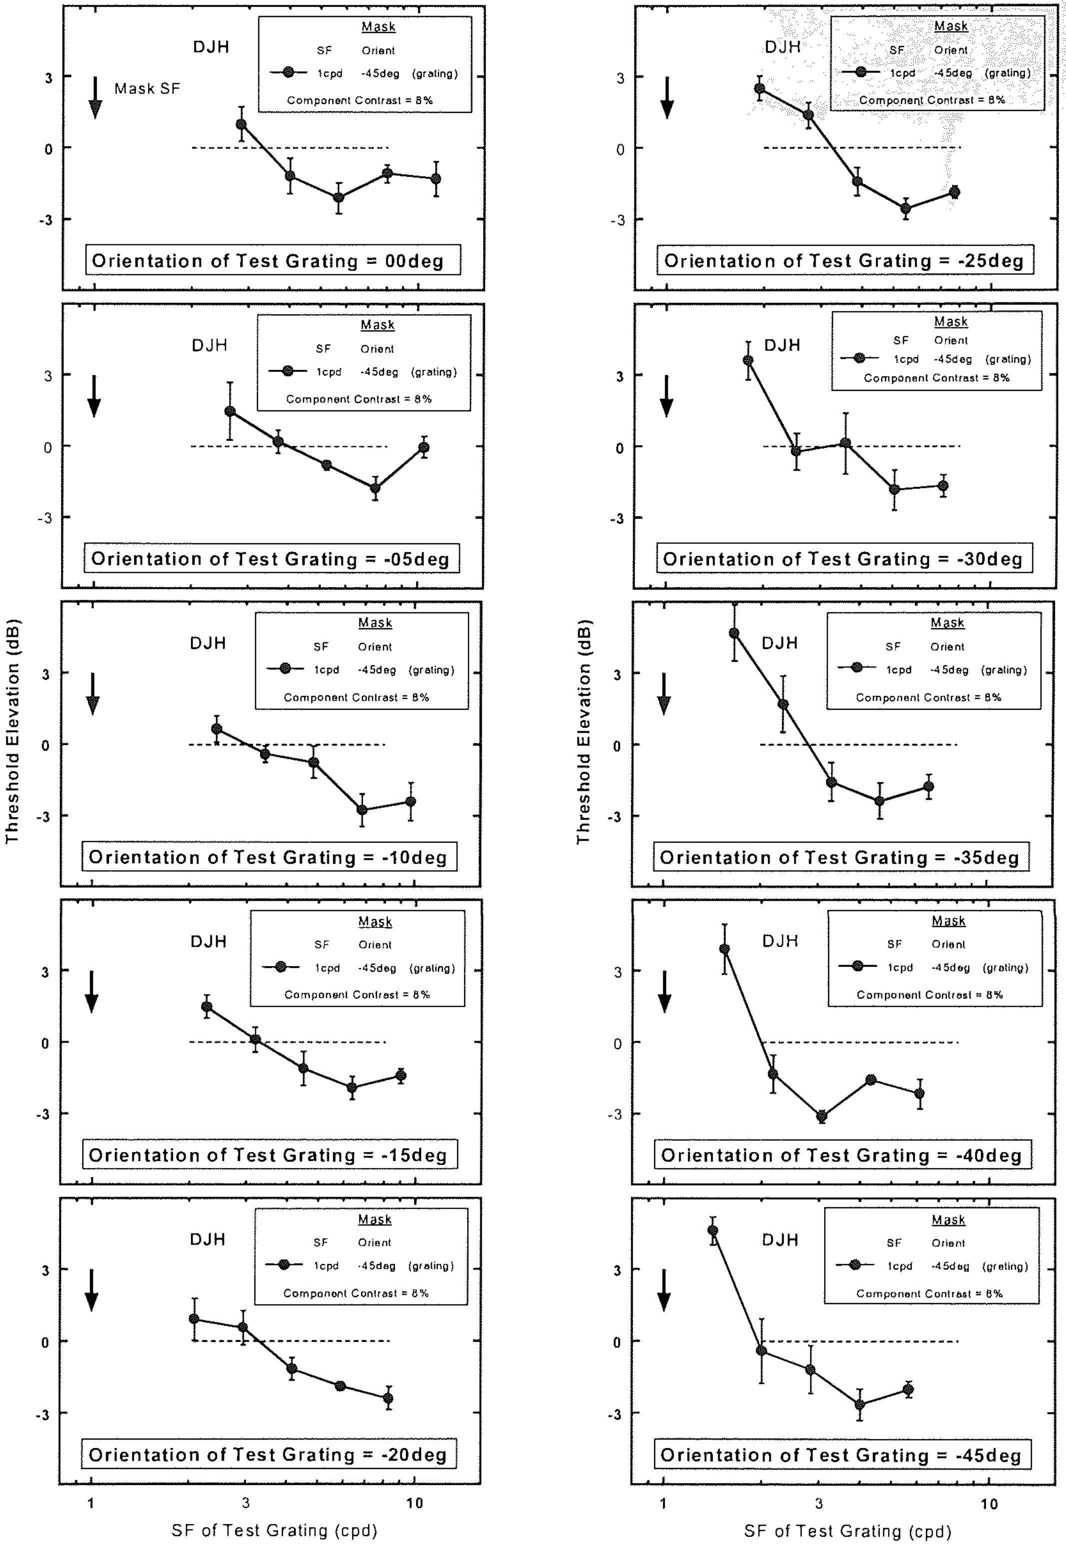

Fig. 3.4. Facilitation as a function of test spatial frequency. Arrows represent the mask spatial frequency (1cpd). Different panels are for different test orientations. The dotted baseline covers the same range (2 to 8cpd) in each panel. Note that the range of test spatial frequencies varies across panels.

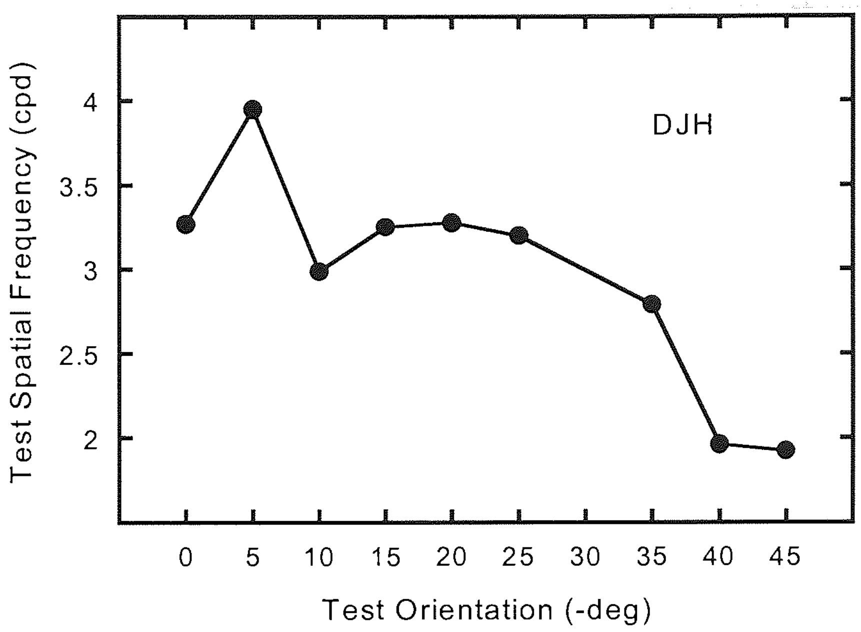

Fig. 3.5. The test spatial frequency at which the functions shown in Fig. 3.4 cross from masking to facilitation, as a function of the orientation of the test. The mask had a spatial frequency of 1cpd, was oriented -45deg and had a contrast of 8%. Data has been omitted because the function gathered with the -30deg test crossed from masking to facilitation twice.

3.6 Experiment 4. The effect of mask contrast on masking and facilitation produced by grating and plaid masks

Introduction

There are two main discrepancies between the data shown in the last few sections and previously published work.

- Discrepancy 1: The facilitation produced by grating and plaid masks, which had stimulus contrasts of 8% and 16% respectively, produced a very similar level of facilitation. So, for this very limited range of mask contrasts at least, it seems that the facilitation is not dependent on mask contrast. However, Tolhurst and Barfield (1978) found that their facilitation was contrast dependent. As mask contrast was changed, the magnitude of their facilitation varied between 0 and 8dB.
- Discrepancy 2: When test and mask were matched in spatial frequency and grating and plaid masks had the same component contrasts, the plaid produced between 3-6dB more masking than the grating. However, Derrington & Henning (1984) showed plaids to produce about 12dB more masking than gratings when the characteristics of their stimuli were matched in this way. One difference between the stimuli used in the two studies was mask component contrast. Those used in the experiments shown in the previous sections were 8%, whilst those used by Derrington & Henning were 20%.

In the next experiment, component contrast was manipulated to address the following two questions.

- Does the facilitation that I have found show any signs of being dependent on mask contrast? Tolhurst and Barfield's was dependent. If the facilitation shown here was not, strong doubts would be raised as to whether the two occurrences of facilitation have the same origin.
- The difference between the masking produced by a plaid and a grating was reported by Derrington and Henning to be much larger than that which I have found. Is this due to the different mask contrasts used in the two studies?

Methods

Tests were vertical gratings, which ranged in spatial frequency from 1 to 8cpd. Test spatial frequency and mask component contrast were pseudo-randomly varied across experimental sessions. Mask components had a spatial frequency of 1cpd and were presented at 8%, 12%, 14%, 16% or 24% contrast. Mask type (either: no-mask, -45deg grating mask or ±45deg plaid mask) was pseudo-randomly varied within sessions.

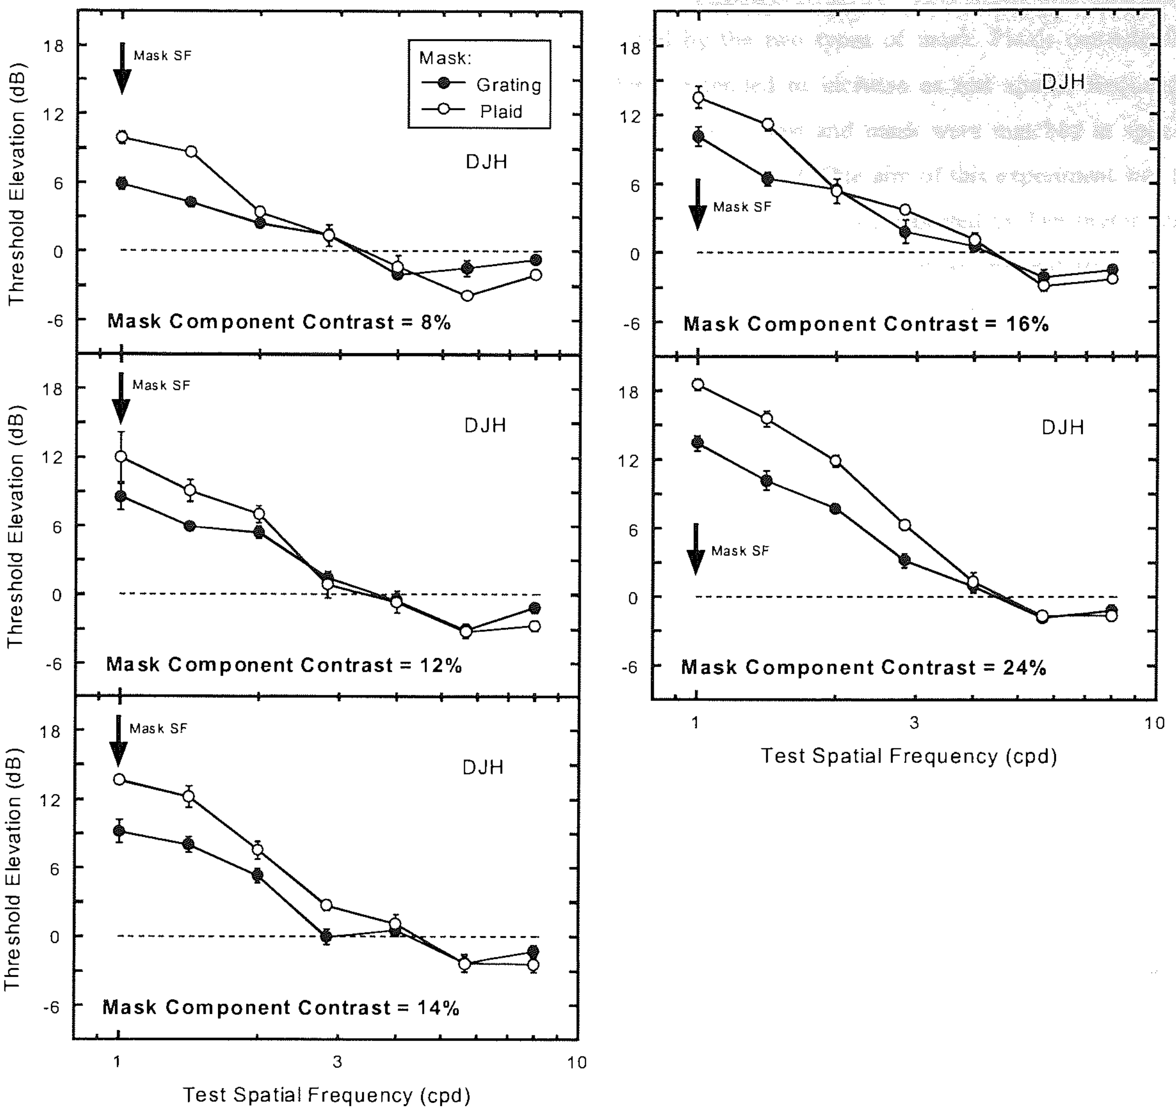

Fig. 3.6. Effects of mask contrast on masking and facilitation produced by gratings (filled symbols) and plaids (open symbols). Arrows indicate mask spatial frequency.

Results and Discussion

Each panel of Fig. 3.6 shows threshold elevation as a function of test spatial frequency for grating (filled symbols) and plaid (open symbols) masks which, as before, were matched in their component contrasts. Different panels show data gathered using masks of different contrasts.

The magnitude of facilitation (seen at the right of the masking functions) was not dependent on mask contrast. This is contrary to the facilitation found by Tolhurst and Barfield. Mask type (grating or plaid) also had no effect on its magnitude. Furthermore, the test frequency at which the facilitation occurred was the same regardless of mask type or contrast. The facilitation and its independence of mask characteristics are discussed in the next section.

At the left of the masking functions it can be seen that increasing mask contrast increased the magnitude of masking, which is consistent with previous work (Nachmias & Sansbury, 1974; Legge & Foley, 1980; Wilson, 1980; Bradley & Ohzawa, 1986; Ross & Speed, 1991; Foley, 1994a). However, it had very little effect on the difference between the levels produced by the two types of mask. Plaids consistently produced greater masking than gratings and this difference tended to increase as test spatial frequency decreased and drew closer to that of the mask components. When test and mask were matched in spatial frequency, the magnitude of the difference remained between 3 and 6dB. One aim of this experiment was to find out whether the very large difference between plaid and grating masking reported by Derrington and Henning was dependent on the contrast of masks that they had used. They used a component contrast of 20%. The present data show the difference produced by masks ranging in component contrast from 8% to 24% was always much smaller than that found by Derrington and Henning. So, the difference between the data here and theirs is not simply due to mask contrast.

3.7 General Discussion

Facilitation: Causes?

This is the first time that facilitation of test detection has been found using masks remote from the tests in both spatial frequency and orientation. It is likely that this is due to the limited ranges of the relative spatial dimensions of tests and masks that have been used in previous investigations (see the chapter introduction).

An interpretation of the facilitation reported by others which occurs when test and mask differ in either spatial frequency or position is that it is caused by channel interactions (see section 3.5 for references). Could this explanation be used for the facilitation shown here? Regarding the present data, a particularly interesting finding is that for the limited range of mask contrasts that were used, the facilitation appeared to be independent of mask contrast. How could this be explained if the facilitation was due to interactions? Perhaps the facilitation produced by the mask saturates at or below a contrast of 8% (which was the lowest mask contrast used). It may be possible to reveal whether this is the case by measuring the facilitation produced by masks at a range of low contrasts.

Another possible interpretation of the facilitation is that the mask may be subliminally exciting the channel most sensitive to the test. The S-shaped contrast response function of channels means that if the mask provided a little excitation to the test channel, facilitation would be expected. This is because the test channel would be pushed into the accelerative region of its operating range. Such an explanation would seem more plausible if the facilitation was found to be phase dependent. (The facilitation of test detection thresholds produced by pedestals is phase dependent. For example, see Foley, 1994b; Zenger & Sagi, 1996).

The facilitation may have a non-sensory cause. For example, it may have been produced by attentional improvement (encouraged by the presence of the mask). The mask may have guided the observer to the test's spatial position and this may have helped them in its detection. Such an explanation could account for why the facilitation showed no effect of contrast and why it was the same for grating and plaid masks. It does seem unlikely however, since the observer was highly trained and was very experienced at identifying where the test was to be presented (the test was also always centred on the fixation point). Moreover, it seems unlikely that the effect should be spatial frequency tuned. Such a cause would be expected to produce a vertical downward translation of the whole masking function but no facilitation (or masking) was present at high test spatial frequencies. Thus, for the account to survive, prior to downward translation a small secondary region of masking would need to occur at high test spatial frequencies. This seems very unlikely.

Masking: Linear summation of mask components?

When grating and plaid masks were matched in their component contrasts, the masking produced by the plaid tended to be greater than that produced by the grating. One possibility (which is returned to later in this thesis) is that the inhibitory influences of the mask components are derived from the linear sum of the mask component contrasts. To test this idea directly, data selected from Fig. 3.6 are replotted in Fig. 3.7 such that grating (filled symbols) and plaid (open symbols) masks were matched in their stimulus (peak-to-peak) contrast at either 16% (top panel) or 24% (bottom panel). With masks matched in stimulus contrast, the masking and facilitatory effects were almost exactly the same. This was the case for both mask contrasts investigated.

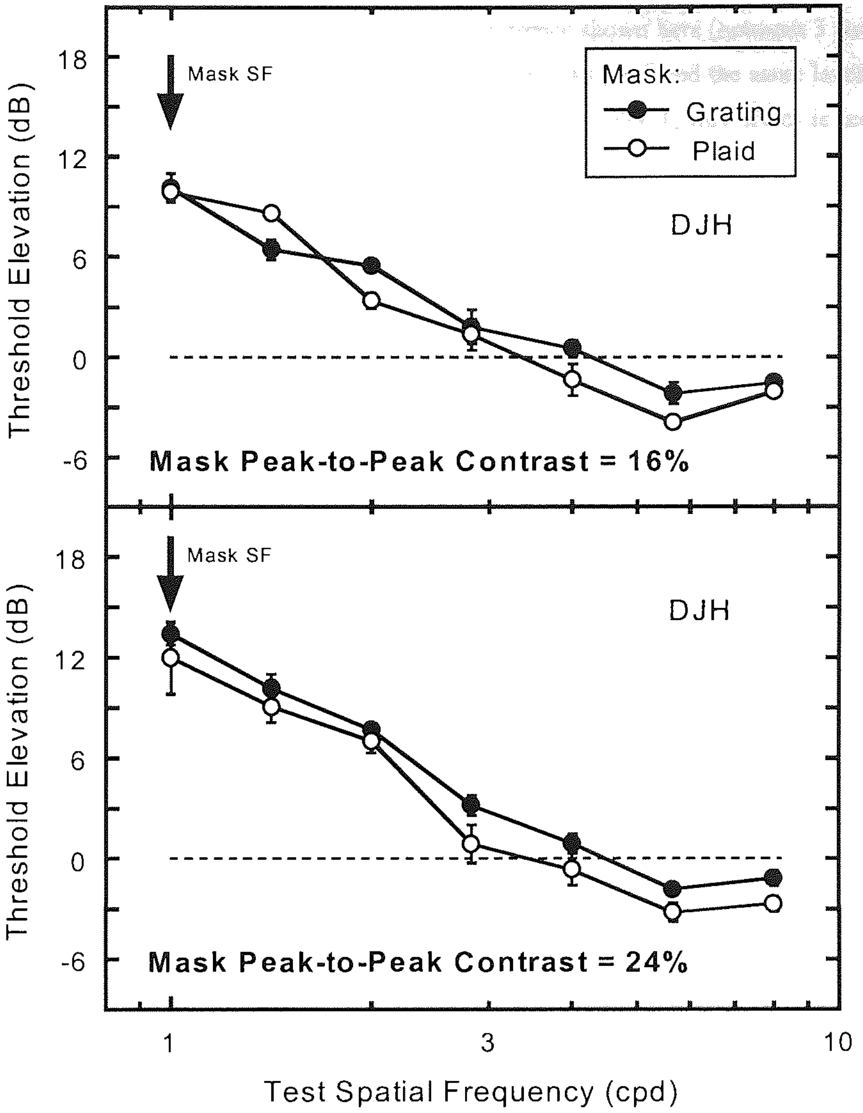

Fig. 3.7. Data from Fig. 3.6. Masking and facilitation as a function of test spatial frequency, produced by grating (filled symbols) and plaid (open symbols) masks which had the same stimulus contrasts at either 16% (top panel) or 24% (bottom panel). Arrows represent mask component spatial frequency.

What follows is a comparison between the masking shown so far in this chapter and what was found by Derrington and Henning:

- New findings: When grating and plaid masks were matched in component contrast, the masking from the plaids was between 3 and 6dB greater than that from the gratings when the test and mask patterns were matched in spatial frequency. When the masks were matched in stimulus contrast, they produced almost exactly the same levels of masking. So, the extra masking produced by the plaid when masks were matched in component contrast, was simply due to the plaid having more contrast than the grating. That is, there was linear-summation of mask component contrasts. As defined in chapter 1, this effect is termed 'linear-suppression'.
- Derrington and Henning's findings: They only measured the effect of grating and plaid masks that were matched in component contrast, and matched in spatial frequency to the test. Like here, they also found

the plaid to produce more masking than the grating. However, they found the difference in masking levels to be around 12dB. This is much larger than the difference shown here (between 3 and 6dB) and is larger than what would be expected if the grating and plaid masks produced the same levels of masking when they were matched in stimulus contrast. As defined in chapter 1, this effect is termed 'super-suppression.'

3.8 Experiment 5(a). The effect of spatial frequency on mask component summation: Masking as a function of test spatial frequency for 1, 3 and 12cpd grating and plaid masks

Introduction

The difference in masking produced by a plaid and a grating was found by Derrington and Henning to be much greater than what has been shown here. They showed super-suppression whilst I have shown linear-suppression. The aim of this experiment was to investigate the reason for this discrepancy.

In investigating why the two sets of findings differed, the main differences between the methods implemented in the two studies are considered. Relative to the stimuli used in the experiments presented in previous sections, those used by Derrington and Henning had a larger diameter and a longer temporal duration. Moreover, they used 3cpd masks whilst I had used 1cpd masks. In this experiment, it was investigated whether the last of these, the spatial frequency difference, was the reason why the findings from the two studies differed.

Derrington and Henning only measured the difference between plaid and grating masking when test and mask were matched in spatial frequency. This experiment extended theirs by (i) measuring the effect of grating and plaid masks at a range of test spatial frequencies, and (ii) examining the effect of mask spatial frequency on the levels of masking.

Methods

Tests were vertical gratings whose spatial frequency was pseudo-randomly varied across experimental sessions. Mask type (either no-mask, -45deg grating mask, or ±45deg plaid mask) was pseudo-randomly interleaved within sessions. Masks had a spatial frequency of 1, 3 or 12cpd and a component contrast of 24%. (This mask contrast was similar to that used by Derrington & Henning.) Sessions testing each mask spatial frequency were blocked together and the order of these blocks was pseudorandom. An exception to this is that in the top-left panel of Fig. 3.8, the data at and between the test spatial frequencies of 1cpd and 8cpd are replotted from the appropriate panel of Fig. 3.6. DJH gathered the data represented by the points at 0.5cpd and 0.7cpd in this panel, independently from the rest of the data.

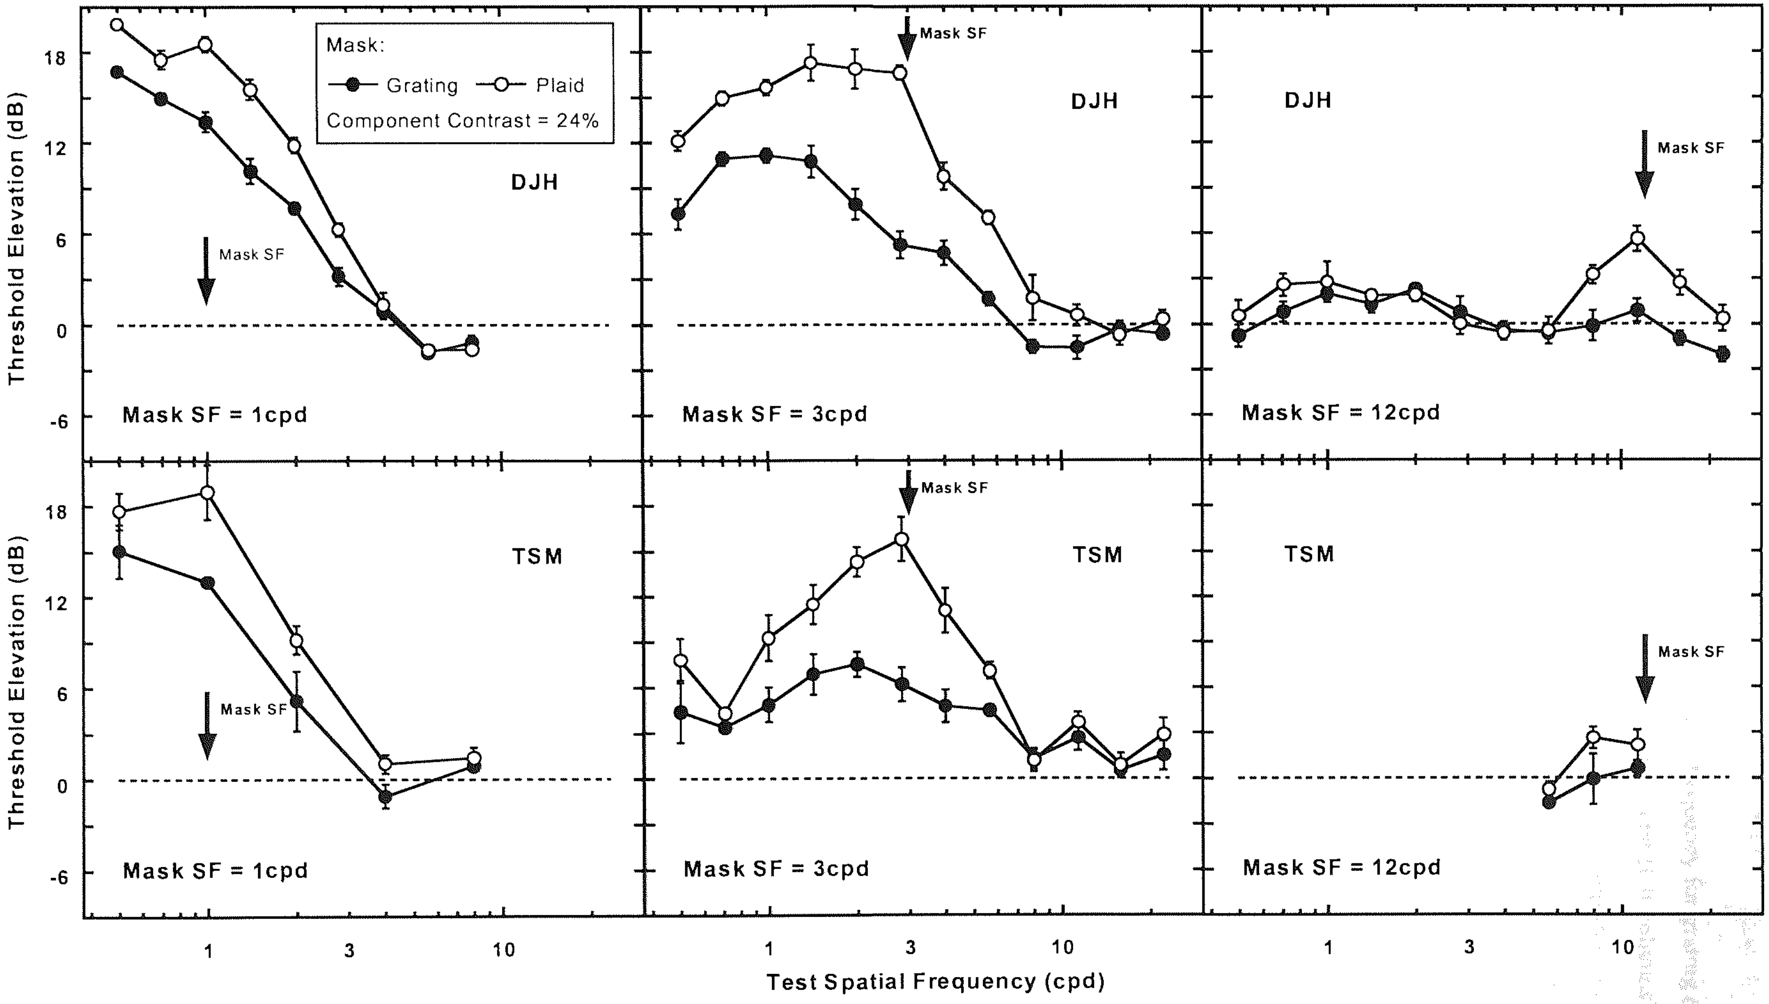

Fig. 3.8. Masking and facilitation as a function of test spatial frequency, produced by grating (filled symbols) and plaid (open symbols) masks, matched in component contrast at 24%. The spatial frequency of the mask components was either 1cpd (left panels), 3cpd (middle panels) or 12cpd (right panels). Arrows indicate mask spatial frequency. The data were gathered from observers DJH (top panels) and TSM (bottom panels).

Results and Discussion

Fig. 3.8 shows threshold elevation as a function of test spatial frequency for grating (filled symbols) and plaid (open symbols) masks. The spatial frequency of the masks was 1cpd (left columns), 3cpd (middle columns) or 12cpd (right columns). The data were gathered from observer DJH (top row) and TSM (bottom row). All of the masks raised test thresholds when the spatial frequency of the test was at or around that of the masks (arrow). The level of masking fell off markedly as mask spatial frequency was increased, consistent with earlier findings in which the orientations of test and mask were similar (Wilson et al., 1983; Phillips & Wilson, 1984) or distant (Phillips & Wilson, 1984).

The general aim of the work in this chapter was to investigate remote masking effects that may previously have been missed. It was considered that such effects might have been missed because of the limited ranges of spatial relationships between test and mask stimuli that have been investigated (i.e. the ranges of the spatial frequencies and orientations of the stimuli). The present experiment was extensive with regards to this and the data show many characteristics that are of interest to the current discussions. These will be described and discussed in turn. They concern: (i) the facilitation from remote masks, (ii) the masking from extremely remote masks, (iii) the test spatial frequency at which peak masking occurred, and (iv) the difference between the masking produced by the plaid and the grating. It is the last of these that directly addresses the original aim of this experiment. The description and discussion of this is withheld until after those of the other characteristics, because it is closely related to the next section.

(i) The facilitation from remote masks

A major discussion point of the last few sections was the remote facilitation that was found. Returning to this briefly, the previous experiments only investigated the facilitation that was produced by 1cpd masks. The data from DJH in Fig. 3.8 show that 3cpd masks may facilitate as well. Further, the difference between the spatial frequencies of the test patterns whose detection is facilitated, and the masks that facilitate, may differ with mask spatial frequency. Expanding on this, the functions produced by 1cpd masks crossed from masking to facilitation when test spatial frequency was around 2½ octaves higher than that of the masks, whilst for 3cpd (grating) masks it occurred when test spatial frequency was around 1½ octaves higher. TSM showed no convincing evidence of facilitation. This was quite surprising because the facilitation from 1cpd masks for DJH and three naïves was very similar (Figs. 3.1 and 3.2). Negligible facilitation was produced by the 12cpd masks (right panels of Fig. 3.8), so, though masks facilitate the detection of tests that are a few octaves higher in spatial frequency, they do not appear to facilitate the detection of tests that are a few octaves lower in spatial frequency.

(ii) The masking from extremely remote masks

It is interesting to note that around 3dB of masking was produced (for DJH) by 12cpd masks when test and mask were remote in orientation by 45deg and remote in spatial frequency by 2 ½ to 4 octaves. This is the most remote masking revealed in this thesis. Within-channel models of masking may account for these data if it is assumed that the channels involved had very broad bandwidths. However, since the spatial frequency and orientation bandwidth of mechanisms sensitive to high spatial frequencies have been estimated to be between 1.25 and 1.5 octaves (Wilson et al., 1983), and 15 and 30deg (Phillips & Wilson, 1984), respectively, such models struggle to answer why there would be any masking at all in this region.

(iii) The test spatial frequency at which peak masking occurred

In the majority of cases, the masking functions reached a peak when the spatial frequencies of test and mask differed. Others have showed this phenomenon. For example, Legge and Foley (1980) and Wilson et al. (1983) showed it when they altered mask spatial frequency whilst keeping the test spatial frequency constant (and the orientations of the stimuli were similar or matched). Harvey and Doan (1990), who measured the effect of an 8cpd mask on the detection of tests at various spatial frequencies, showed it when the stimuli were either matched or different in orientation (see their Fig. 9 in particular), and Derrington and Henning (1989) showed it with 3cpd plaid masks.

Masking functions gathered by keeping the test constant whilst varying the spatial frequency of the mask (e.g. Legge & Foley, 1980; Wilson et al., 1983) have been used to estimate the bandwidth of the test detecting channels. As already discussed, such studies used within-channel models of masking for their estimates. It was considered that a masking function, which peaks when test and mask spatial frequencies are unequal, is indicative of the channel most responsive to the test being most sensitive to a spatial frequency that is (slightly) different from that of the test. This is not necessarily the case, however. Since these papers were published, a great deal of data has been produced which favour interpretations in terms of cross-channel models of masking. This complicates an interpretation of a masking function's peak. It could be revealing peak sensitivities of a channel but it may be revealing some form of interaction that occurs between channels.

Functions gathered by keeping the mask constant whilst varying the spatial frequency of the test, such as in the studies of Derrington and Henning (1989) and Harvey and Doan (1990), and in the experiments performed for this thesis, allow an investigation into how masks effect the operation of the channels that are most sensitive to the test. The feature of the present data described above, could be revealing spatial dimensions of a particular channel but this seems unlikely because of two main empirical observations. The first is the large range of test spatial frequencies where a greater or comparable level of masking was produced to that when test and mask spatial frequencies were matched. Refer particularly to the data of DJH in the top-middle panel of Fig. 3.8. These show that a -45deg 3cpd grating mask produced around 6dB more masking when test spatial frequency was between 1 and 2 octaves lower than that of the mask, than when it was matched to that of the mask. The second observation is that masking functions peaking when test spatial frequency was lower than that of the mask was the common result in the papers mentioned previously in this paragraph and in Fig. 3.8. If the effect were due to test dimensions, the highest level of masking would be expected to occur at test spatial frequencies that were either above, below, or very similar to that of the mask. From these two observations, it appears feasible that channel interactions are producing the effects. Whatever the underlying cause, it differs across observers (refer particularly to the middle panels of Fig. 3.8) and with mask spatial frequency (Fig. 3.8).

(iv) The difference between the masking produced by the plaid and the grating (a)

Derrington and Henning (1989) showed the difference in masking produced by plaids and gratings to be much greater than that what was shown in previous sections of this chapter. One major difference between the studies was that they had used 3cpd masks whilst I had used 1cpd masks. The main aim of this experiment was to investigate whether the difference in results was due to the different mask spatial frequencies that had been used in the studies. This was found to be the case. The data from Fig. 3.8 have been replotted in Fig. 3.9 to show the difference between the masking and facilitation produced by plaid and grating masks. When the spatial frequencies of the test and mask components were both 1cpd, the plaid mask produced between 3 and 6dB more masking than the grating mask. This was a similar difference to what has been shown previously in this chapter. When test and mask were both 3cpd however, the plaid produced around 9 to 12dB more masking than the grating. This difference is very close to that reported by Derrington and Henning. Masks with 12cpd produced much less masking than those of a lower spatial frequency and the difference between the levels produced by the plaid and grating was between 0 and 6dB.

So, the curiously large difference between the levels of masking produced by a +45deg 3cpd grating and a ±45deg 3cpd plaid on the detectability of a 3cpd vertical test pattern, reported by Derrington and Henning (1989), is not general. It is highly dependent on mask spatial frequency, and on whether the test and masks have spatial frequencies that are very similar (within half an octave) (see the middle panel of Fig. 3.9).

It is also interesting to note that though the functions from DJH and TSM produced by 3cpd masks appear quite different (TSM's reaches a peak much closer to the test spatial frequency; see middle panels of Fig. 3.8), replotting these data to show differences between plaid and grating masking for the two observers results in functions that are very similar to each other (middle panel of Fig. 3.9).

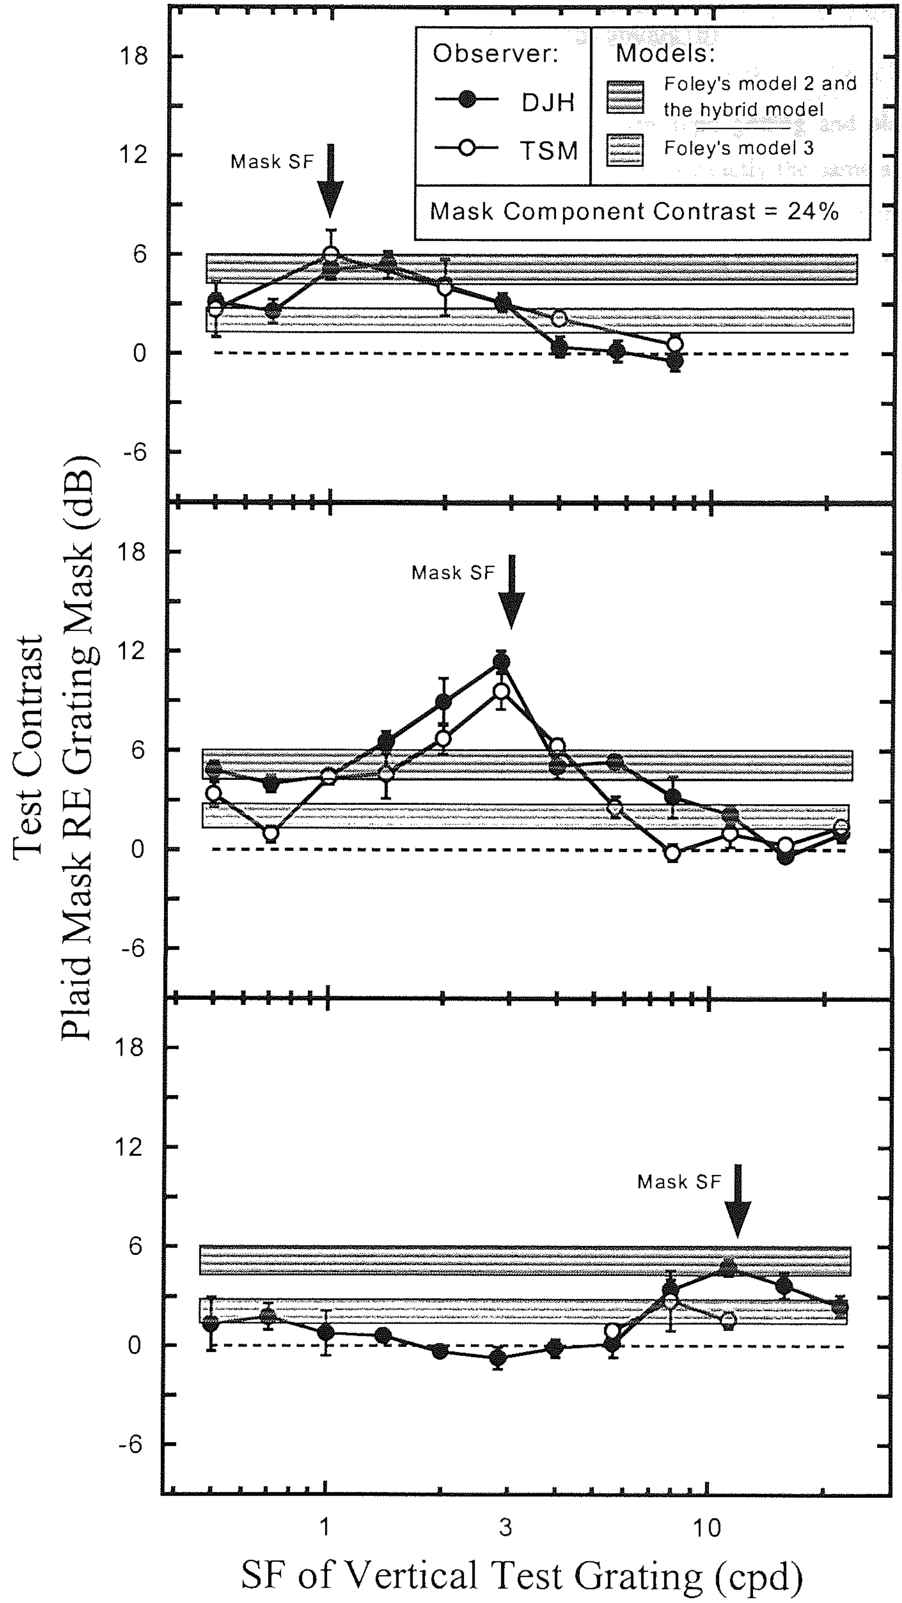

Fig. 3.9. Differences between masking functions produced by plaid and grating masks. Data points > 0 indicate that masking was greater for the plaid than for the grating. Mask spatial frequency (indicated by arrows) was 1cpd (top panel), 3cpd (middle panel) or 12cpd (bottom panel). Data were gathered from DJH (filled symbols) and TSM (open symbols). Ranges of predicted differences have been estimated for some of the quantitative models of masking described in chapter 4. Parameter values were selected from those estimated in chapters 4 and 5. Discussion of these modelling estimates is withheld until chapter 4, section 4.8.

The difference between the masking produced by the plaid and the grating (b)

Results from experiment 4 (see Fig. 3.7) showed that when 1cpd grating and plaid masks were matched in stimulus contrast, the masking that they produced was almost exactly the same at a range of test spatial frequencies. This is indicative of linear-suppression, that is, linear summation of mask component contrasts in the suppression process. When masks were matched in component contrast and spatial frequency, the difference between the masking produced by the plaid and grating was much smaller when the stimuli were 1cpd than when they were 3cpd. This lead to the masking produced by the 3cpd plaid to be termed super-suppression. That is, the plaid masking was considered to be much greater than that predicted from the grating masking, if the difference was simply due to the extra contrast of the plaid. It seemed obvious that if the 3cpd grating and plaid masks were matched in stimulus contrast, the plaid would still produce substantially more masking than the grating. This has not been tested directly, however. It is tested directly here.

Fig. 3.10 compares the masking produced by grating (filled symbols) and plaids (open symbols) as a function of test spatial frequency when mask spatial frequency was either 1cpd (left panels) or 3cpd (right panels) and when the masks were matched in component (top panels) or stimulus contrast (bottom panels) at 24%. Data in the top- and bottom-left panels are replotted from Figs. 3.6 and 3.7, respectively. Data in the right panels have not been shown before. They were gathered using the same methods as the data in the left panels.

When grating and plaid masks were matched in stimulus contrast (bottom panels), they generally produced about the same levels of masking; this is the signature of linear-suppression. As expected however, super-suppression was evident when test and mask had 3cpd; the plaid produced around 6dB more masking than the grating.

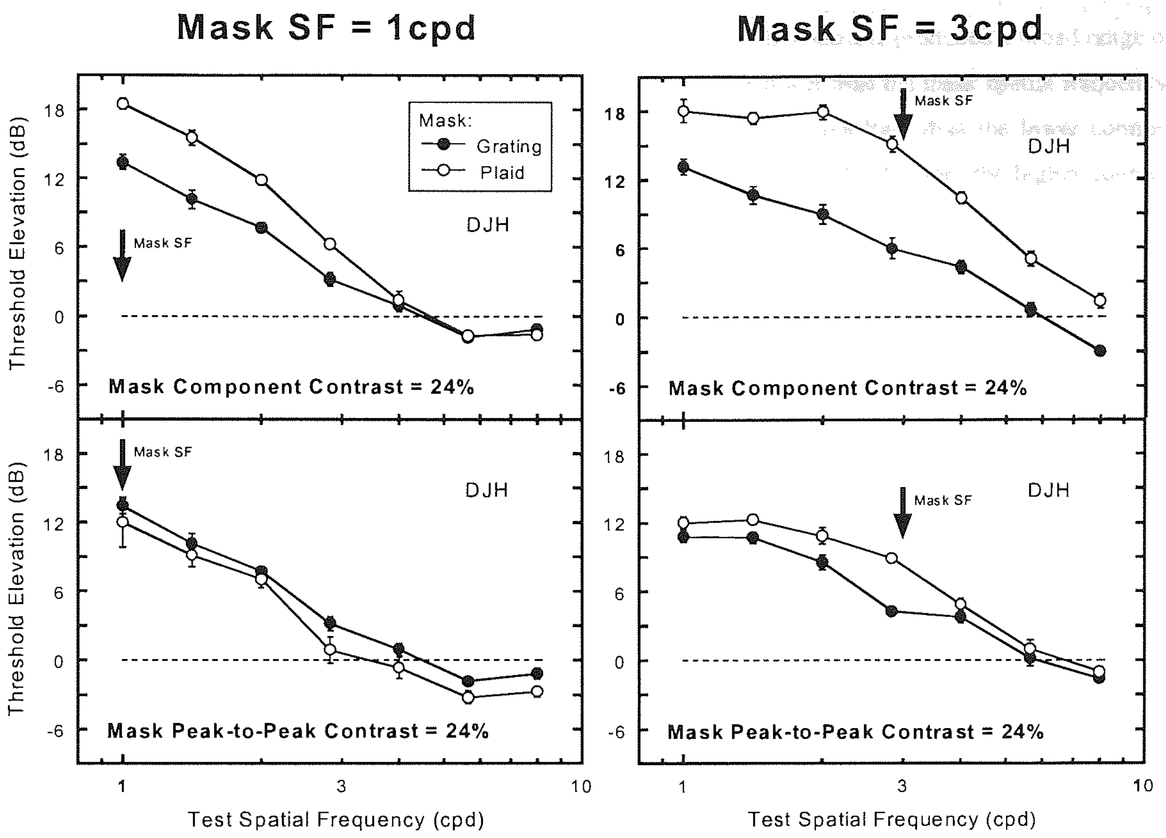

Fig. 3.10. Masking and facilitation as a function of test spatial frequency. Grating (filled symbols) and plaid (open symbols) masks had a spatial frequency of either 1cpd (left panels) or 3cpd (right panels) and they were matched in either their component (top panels) or stimulus (bottom panels) contrast at 24%. Arrows indicate mask spatial frequency.

3.9 Experiment 5(b). The dependency of super-suppression on mask contrast

Introduction and Methods

It was shown in the previous sections that the difference between the masking from 1cpd plaids and gratings is consistent with linear-suppression. This appeared to be the case regardless of mask contrast (see Figs. 3.6 and 3.7). However, when the stimuli had 3cpd, super-suppression was found. The next experiment explored the effect of mask contrast on super-suppression. Would super-suppression be shown to be independent of mask contrast like linear-suppression has been? Mask component contrast was either 8% or 24%. Methods were described in section 3.9.

Results and Discussion

Data shown in the bottom panels of Figs. 3.11 and 3.12 are taken from the top-middle panel of Fig. 3.8 and the middle panel of Fig. 3.9, respectively. Fig. 3.11 shows threshold elevation as a function of test spatial frequency for 3cpd grating (filled symbols) and plaid (open symbols) masks whose component contrast was either 8% (top panel) or 24% (bottom panel). Masks of both contrasts produced a broad range of substantial masking, which peaked at test spatial frequencies that were lower than the mask spatial frequency. As with 1cpd masks (Fig. 3.6), the higher contrast masks produced more masking than the lower contrast masks yet mask contrast did not affect magnitudes of facilitation. Regarding facilitation, the higher contrast masks facilitated the detection of test patterns that were of a higher spatial frequency than those whose detection was facilitated by the lower contrast masks. There is a hint of this in the data gathered with 1cpd masks (see Fig. 3.6).

Fig. 3.12 shows the difference between the masking and facilitation produced by the plaid and that produced by the grating for the data shown in Fig. 3.11. It is particularly striking from this figure that though super-suppression was produced by the higher contrast masks, it was not produced by the lower contrast masks. The difference between the plaid and grating masking produced by the lower contrast mask was between 3 and 6dB. The size of this difference is around that expected if linear-suppression was evident (see experiment 4). To summarise:

1. Super-suppression has been shown only when tests and masks have around 3cpd and the mask is of a high contrast.
2. Linear-suppression appears to be much more widespread than super-suppression. It has been shown with:
   (i) 1cpd masks at a range of test spatial frequencies and mask contrasts
   (ii) 3cpd masks when the spatial frequency of the test is either below or above that of the mask
   (iii) 3cpd masks with a 3cpd test when the mask is of a low contrast

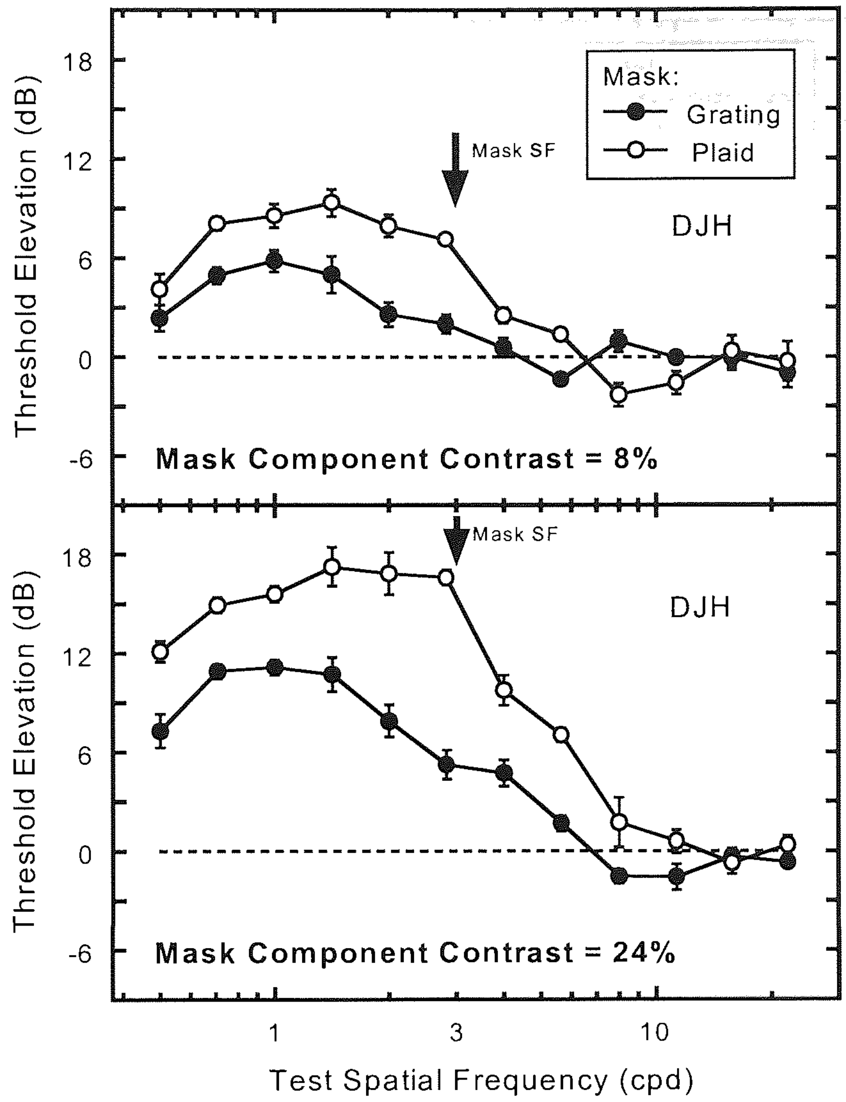

Fig. 3.11. Masking and facilitation as a function of test spatial frequency, produced by 3cpd grating (filled symbols) and plaid (open symbols) masks matched in component contrast at 8% (top panel) or 24% (bottom panel). Data in the bottom panel has previously been shown in the top-middle panel of Fig. 3.8. Arrows indicate mask spatial frequency.

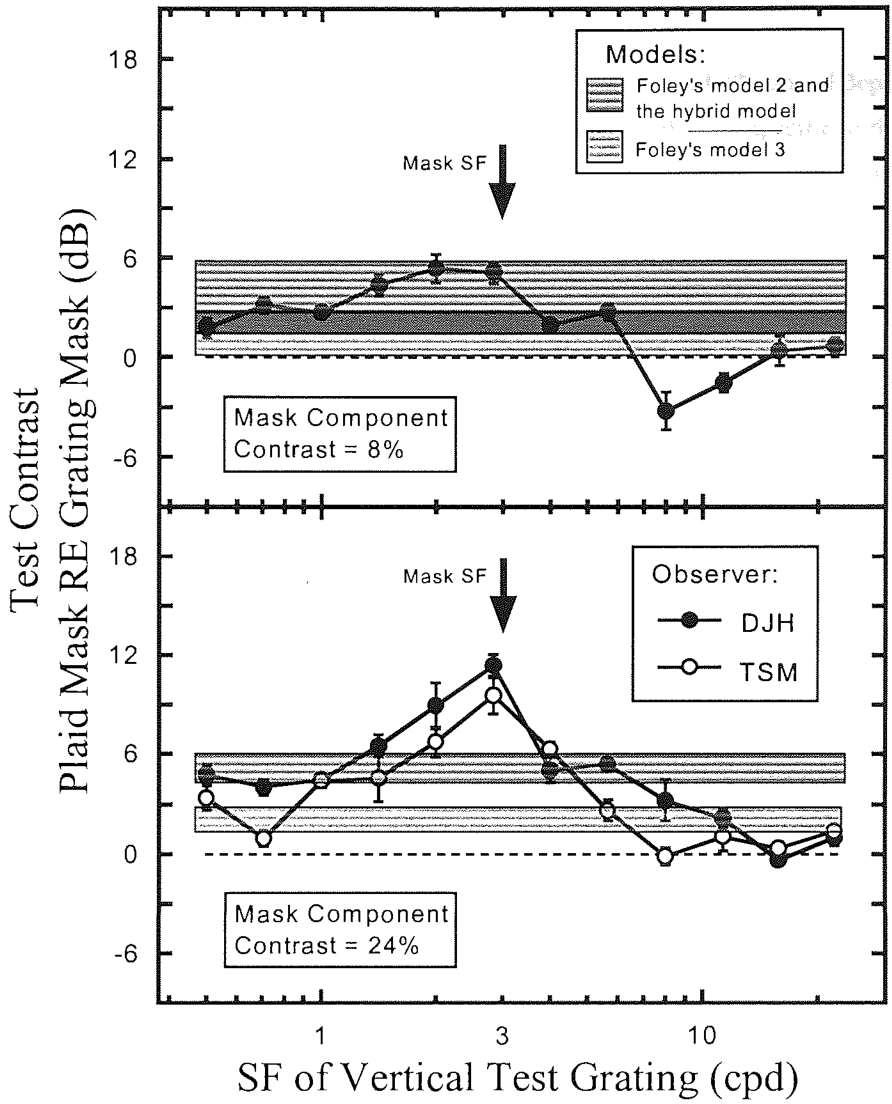

Fig. 3.12. The difference between the masking and facilitation produced by the plaid mask and that produced by the grating mask (plaid mask RE grating mask), as a function of test spatial frequency. Mask component spatial frequency, indicated by an arrow, was 3cpd. Mask component contrast was either 8% (top panel) or 24% (bottom panel). Data in the bottom panel has previously been shown in the middle panel of Fig. 3.9. The magnitude of the difference in masking produced by the two types of mask is shown for two quantitative models of masking with parameters set to reasonable values (shaded patches). Discussion of these modelling estimates is withheld until chapter 4, section 4.8.

Another finding specific to the use of 3cpd stimuli

Super-suppression appears to be specific to when test and mask are both around 3cpd. An obvious question is whether there are any other findings reported in the literature that are specific to the use of 3cpd stimuli. The answer to this is yes. Patterns with certain characteristics tend to induce epilepsy and migraine in photosensitive observers and perceptual illusions in normal observers. Patterns with the following characteristics have been shown to be the most potent (see Wilkins et al., 1979; 1984):

- Stripes, the width of which are the same as the distance between them
- An optimal spatial frequency of 3cpd
- A high contrast. Abnormal brain activity was shown to increase as pattern contrast was increased between 5% and 30%.

So, just like super-suppression, these effects are likely to occur with 3cpd, high contrast gratings.

Using MEG, Adjamian and Holliday (2002) have shown that the above patterns cause widespread excitatory gamma oscillations (30-60Hz) in V1. Such oscillations are produced by a subset of superficial pyramidal cells, called chattering cells (Gray & McCormick, 1996). It is considered that the purpose of this oscillation is to synchronise the neuronal activity responsible for processing a particular visual scene so that a unified perception of it can be achieved (Gray et al., 1989).

3.10 Summary

Masking has been a widely used tool in psychophysics. Channel-bandwidths were estimated from its effects using within-channel masking models. It then enjoyed resurgence when cross-channel effects were observed. Given all the interest in masking, it is important that its range is thoroughly explored. The general aim of the work in this chapter was to explore remote masking effects. The range of the spatial relationships between test and mask was greater than what had been previously explored. The test and mask differed in orientation by 45deg and the relative spatial frequencies of the stimuli were varied across a wide range.

The most striking finding was that substantial masking was produced from a wide range of remote masks. This is consistent with the main feature of cross-channel theories of masking; that a mask, which does not excite the test detecting mechanism, can affect the test's threshold.

A more specific aim of this chapter was to probe an interesting finding of Derrington and Henning's (1989). They showed that the detection threshold of a vertical test was raised 12dB more by a plaid than a grating mask when stimulus spatial frequency was 3cpd and mask contrast was 20% I have provided evidence that this super-suppression is a very specific occurrence. It appears not to be present when test and masks are matched at a lower spatial frequency, when they differ in spatial frequency, or when masks are of a lower contrast. In these circumstances, the far more widespread finding is that plaids and gratings, matched in stimulus contrast, produce about the same level of masking. This is consistent with linear-suppression; the linear summation of mask component contrasts.

A more thorough and systematic investigation into the extent of linear-suppression and super-suppression is shown in the next chapter. Findings of that investigation are used to test some cross-channel models of masking.

Chapter Four:

Linear- and Super-Suppression from Remote Masks

4.1 Chapter Introduction

In the last chapter, the detectability of test patterns was shown to be affected in a variety of ways by spatio-temporally superimposed masks. The stimuli were remote from each other in orientation or in both orientation and spatial frequency. The masking levels produced by gratings and plaids were compared to explore the summation rules for inputs from remote mask components in cross-channel masking. The most interesting effects that were revealed have been termed linear- and super-suppression. The reader is reminded what these are:

(i) Linear-suppression: The masking levels produced by gratings and plaids were about the same when the patterns were matched in stimulus contrast. This indicates linear summation of the component contrasts in the cross-channel suppression process.

(ii) Super-suppression: The masking levels produced by plaids were greater than those produced by gratings when the patterns were matched in stimulus contrast. This indicates that the inputs of the plaid's components are summed nonlinearly in the cross-channel suppression process.

Super-suppression appeared to be far less widespread than linear-suppression. It was only found when stimuli were matched at 3cpd and when the masks were of a 'high' contrast. Linear-suppression was found when: (i) 3cpd stimuli were used but the masks were of a 'lower' contrast, (ii) the stimuli were matched at a lower spatial frequency, or (iii) the stimuli differed in spatial frequency.

There were two aims of the experiments in this chapter. The first was to perform a more systematic and thorough investigation of linear- and super-suppression. The second was to investigate how well some favoured cross-channel suppression models predict the grating and plaid masking investigated here.

4.2 Chapter General Methods

All of the stimuli used for this chapter were spatially curtailed with a raised-cosine function and temporally curtailed with a rectangular window function (step-edge on and off). Apart from the stimuli used for the data in Figs. 4.3, 4.9 and 4.10, they measured 4.4° at full-width, half-height, and had a plateau measuring 4.1°. Their temporal duration was 100ms. Tests were vertical gratings whilst masks were gratings or two-component plaids. Except for data shown in Figs. 4.7 and 4.11, the orientation of the grating mask was -45deg and the plaid components were oriented ±45deg.

4.3 Experiment 1. Linear- and super-suppression when test and mask have the same spatial frequency

Introduction and Methods

This experiment investigated the dependence of linear- and super-suppression on stimulus spatial frequency and mask contrast.

Grating and plaid masks were matched in stimulus contrast. This (ranging from 8 to 48%), and stimulus spatial frequency (1, 3 or 9cpd), were varied across sessions, whilst mask type (no-mask, grating mask or plaid mask) was varied within sessions.

Results

Threshold elevation is shown as a function of mask stimulus contrast for grating (filled symbols) and plaid (open symbols) masks. Stimulus spatial frequency was 1cpd (Fig. 4.1), 3cpd (Fig. 4.1 & 4.2) or 9cpd (Fig. 4.3). As in other studies where the orientation of the test was remote from that of the masks, all of the functions showed masking to increase monotonically with mask contrast (e.g. Ross & Speed, 1991; Foley & Boynton, 1994; Foley, 1994a; Foley, 1996; Foley & Chen, 1997). In the subsections below, the data gathered using stimuli with each spatial frequency are described and discussed.

Stimuli with 1cpd

The masking levels produced by 1cpd gratings and plaids were similar at all mask contrasts (left-hand panels of Fig. 4.1). There was a trend for the grating to produce slightly less masking than the plaid, particularly with GR. For DJH, the mean average of this difference across mask contrasts was 0.9dB (SE = 0.21dB), whilst for GR it was 2.1dB (SE = 0.55dB). At all four contrasts investigated with DJH and at two of the contrasts investigated with GR however, the magnitude of the difference was within experimental error (i.e. the standard error bars of corresponding datum points overlapped). Thus, linear-suppression appears to be evident across a large range of mask contrasts when stimuli had 1cpd.

Stimuli with 3cpd (a)

With 3cpd stimuli, the masking levels produced by plaids were consistently greater than those produced by gratings (right-hand panels of Fig. 4.1). This super-suppression is consistent with Derrington and Henning (1989) and data in chapter 3. Also consistent with data in chapter 3 is that super-suppression was most evident from the higher contrast masks. For DJH, the difference in threshold elevation was 1.3dB and 8.4dB when the contrasts were 8 and 48% respectively, whilst for GR, it was 3.9dB and 8.5dB respectively.

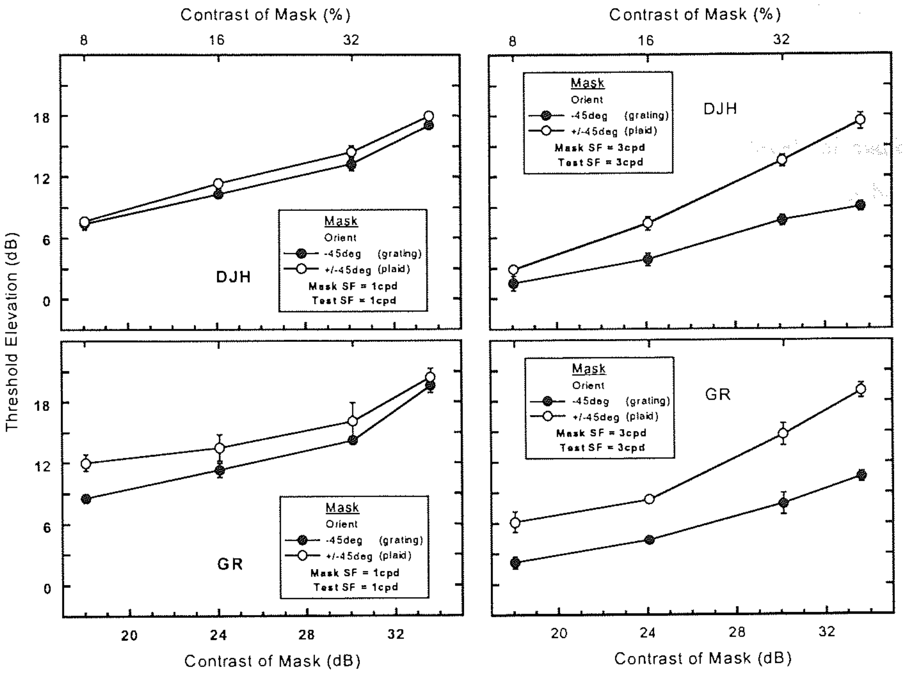

Fig. 4.1. (Above) Masking as a function of mask stimulus contrast. Masks were gratings (filled symbols) or plaids (open symbols) and the stimuli had 1cpd (left panels) or 3cpd (right panels). Data are from DJH (top panels) and GR (bottom panels).

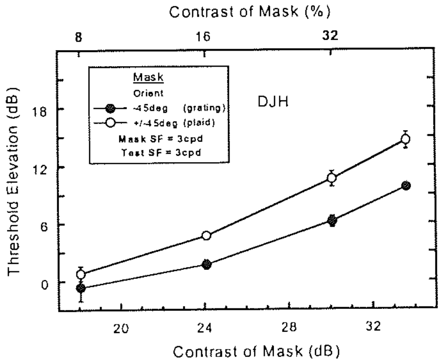

Fig. 4.2. (Above) Masking as a function of mask stimulus contrast. These data are from a control experiment, which addressed the results shown in the right-hand panels of Fig. 4.1 The number of cycles per stimulus patch was matched to that of the 1cpd stimuli (refer to text for details). Masks were gratings (filled symbols) or plaids (open symbols) and the stimuli had 3cpd.

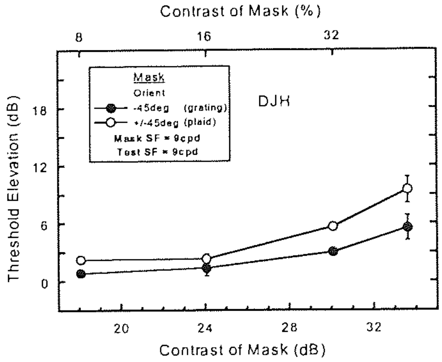

Fig. 4.3. (Above) Masking as a function of mask stimulus contrast. Masks were gratings (filled symbols) or plaids (open symbols). Stimulus had 9cpd.

Stimuli with 3cpd (b): A control

When the stimuli had 1cpd, plaids and gratings produced about the same levels of masking, whilst when they had 3cpd, plaids consistently produced more masking than gratings. The differences between plaid and grating masking levels seem to be spatial frequency dependent. However, there was an extraneous variable between the methods that may have confounded the results; the stimuli were matched in diameter, so those with 3cpd had three times more cycles than those with 1cpd. This might have produced the difference. (One possible reason for this is spatial probability summation. For a review, see Graham, 1989). To examine this, the experiment was repeated with 3cpd stimuli matched to those with 1cpd in terms of the number of cycles per stimulus patch. The data are shown in Fig. 4.2 and are comparable to those in Fig. 4.1 (right-hand panels), showing that the difference between the plaid and grating masking is not reliant on the number of cycles.

Stimuli with 9cpd

The masking produced with 9cpd stimuli (Fig. 4.3) was far less than that produced with the lower spatial frequency stimuli, consistent with the findings of others (e.g. Campbell & Kulikowski, 1966; Daugman, 1984; Phillips & Wilson, 1984; Morrone et al., 1987; Harvey & Doan, 1990; Ross et al., 1993). (See chapter 6 for a further discussion.) Again, the plaid consistently produced greater levels of masking than the grating. The difference at the lower mask contrasts resembled linear-suppression whilst super-suppression was evident at higher mask contrasts.

Discussion

These data are consistent with those in chapter 3. When test and masks were matched in spatial frequency, the 1cpd gratings and plaids produced around the same level of masking irrespective of mask contrast; linear-suppression was evident. With 3cpd stimuli, plaids produced progressively more masking than gratings as mask contrast was increased; super-suppression was evident (supporting Derrington & Henning; 1989). Little masking was produced with 9cpd stimuli.

Has super-suppression been shown with stimuli having spatial frequencies other than 3cpd? There was a trace of it when the 9cpd stimuli were used and to a lesser extent, when the 1cpd stimuli were used. Has super-suppression been reported in the literature when stimuli with spatial frequencies other than 3cpd have been used? Ross and Speed (1991) found it with 2cpd stimuli. They measured masking as a function of mask stimulus contrast when the test was a vertical grating and the mask was either a +45deg grating or a ±45deg plaid. Between mask contrasts of around 2 and 20%, the plaid was found to produce between 3 and 6dB more masking than the grating and (like here with the 3cpd stimuli) this difference increased steadily with mask contrast (see their Fig. 4b). The magnitude of this super-suppression was substantially smaller than that shown here using 3cpd stimuli. So, it seems that some super-suppression occurs when stimuli of spatial frequencies other than 3cpd are used, but, with the limited data that is available, it appears to have the greatest magnitude when stimuli are matched with 3cpd and the masks are of a high contrast.

4.4 Experiment 1(b). Linear-suppression for remote masking

Introduction and Methods

In the last experiment, the effect of spatial frequency on suppression-type was investigated when the test and masks had the same spatial frequency. Here, that work is extended by using tests and masks that differ in spatial frequency as well as, like before, orientation. The method was the same except that the test and masks had spatial frequencies of 1cpd and 3cpd, respectively.

Results and Discussion

The left- and right-hand panels of Fig. 4.4 show the same data fit by two models of masking. Explanation and discussion of these fits is witheld until sections 4.6 and 4.7. The detectability of the 1cpd test pattern is shown as a function of the stimulus contrast of the 3cpd grating (filled symbols) and plaid (open symbols) masks for three observers. Like before, monotonically increasing functions were produced. Linear-suppression was particularly striking in the data.

These results and those in the previous section are consistent with chapter 3. There, grating and plaid masking was shown as a function of test spatial frequency for 1 and 3cpd masks (see Fig. 3.7 and the bottom-left panel of Fig. 3.10, respectively). (The 1cpd masks had 16% or 24% contrast, whilst the 3cpd masks had 24% contrast). Linear-suppression was evident with most stimulus spatial frequency combinations, the most obvious exception being when test and mask both had 3cpd. A summary of key data from this and the last chapter is shown in Fig. 4.5. The main point to take from this figure is that the greatest difference between the masking levels produced by plaids (checked columns) and gratings (lined columns) occurred when the test and masks had 3cpd.

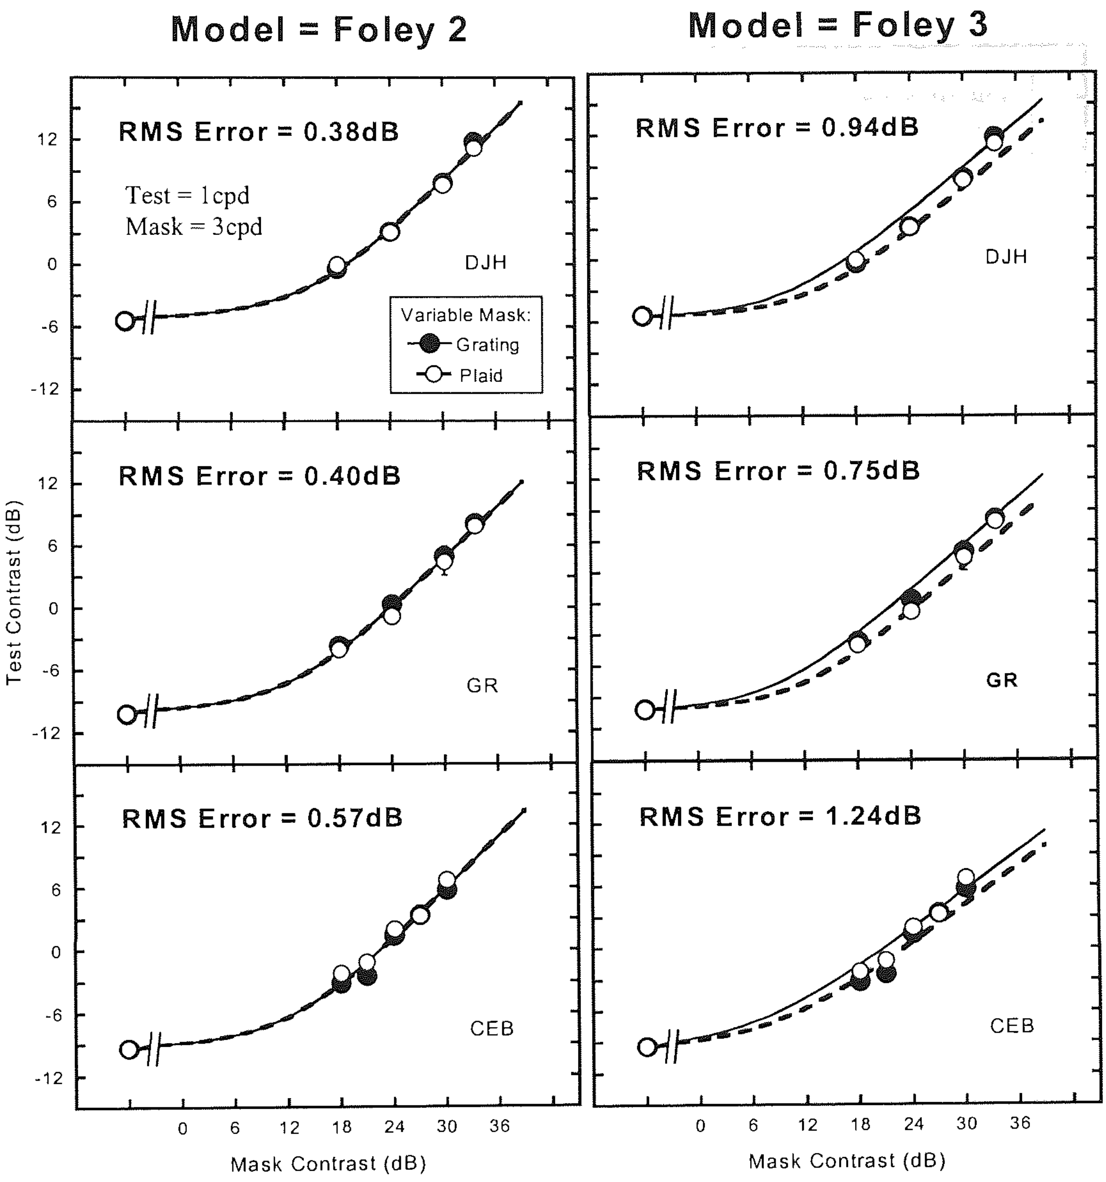

Fig. 4.4. Test detection thresholds as a function of mask stimulus contrast. The test and masks had 1 and 3cpd, respectively. Fits of F2 (left-panels) and F3 (right-panels) are shown. Solid and dashed curves represent fits to the grating (filled symbols) and plaid (open symbols) masking, respectively. Data are from DJH (top panels), GR (middle panels) and CEB (bottom panels).

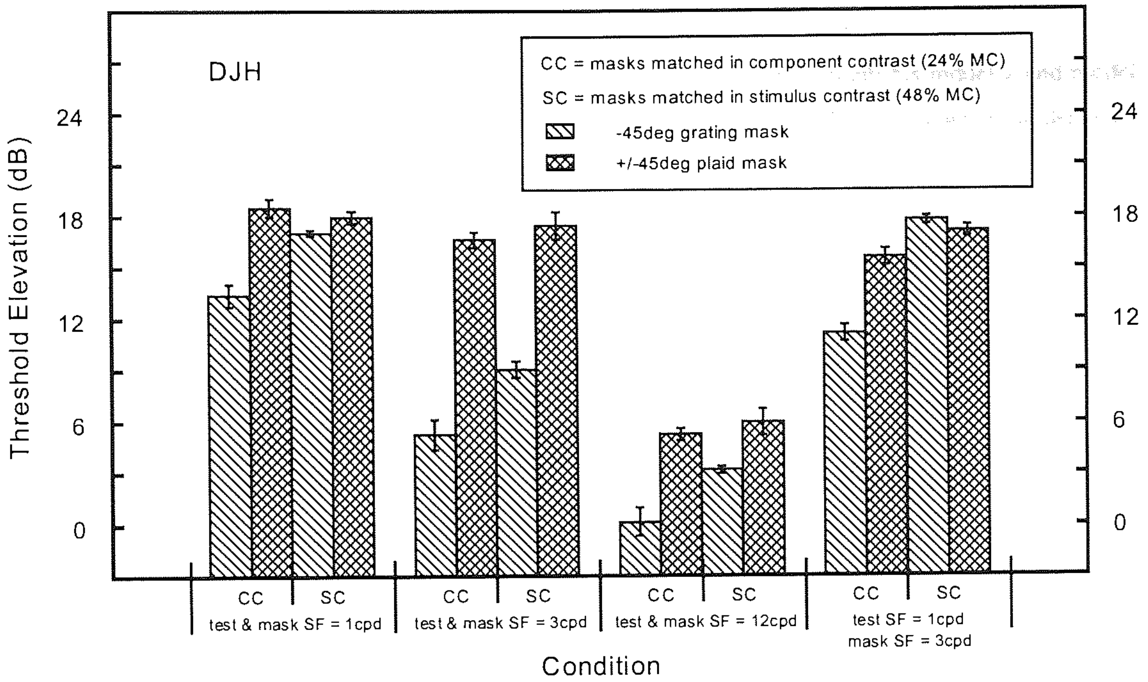

Fig. 4.5. A summary of key results shown so far in this chapter and in chapter 3. The masking produced by gratings (striped columns) and plaids (checked columns), matched in component (CC) or stimulus (SC) contrast, is shown for various stimulus spatial frequency combinations.

4.5 Understanding masking

Suppression has been described in terms of intracortical inhibition in many physiological reports (e.g. Morrone et al., 1982; Bonds, 1989; Bauman & Bonds, 1991; DeAngelis et al., 1992; Heeger, 1992a; Carandini & Heeger, 1994; Sengpiel & Blakemore, 1994; Sengpiel et al., 1995; Carandini, Heeger and Movshon, 1997; Sengpiel et al., 1998; Walker et al., 1998; Allison et al., 2001). Direct evidence for the link came from the pharmacological study of Morrone et al. (1987). They found that antagonizing the γ-aminobutyric acid (GABA) inhibitory neurotransmitter in V1, led to a reduction of the inhibition produced by masks parallel or orthogonal to the prefered orientation of a neuron. Accordingly, many psychophysical models have suggested that suppression from masking and adaptation is due to nonlinear contrast gain control or normalization mechanisms operating in this area (e.g. Geisler & Albrecht, 1992; Heeger, 1992a; Carandini & Heeger, 1994; Foley & Chen, 1997; Watson & Solomon, 1997; Meese & Holmes, 2002). Foley (1994a) reported powerful data consistent with this, which he described with models that are now two of the most widely cited and complete descriptions of pattern masking. (All of the above is discussed in greater detail in chapter 1). The subsequent sections of this chapter investigate how well his models capture the remote masking shown here.

4.6 Foley's (1994a) models

Foley's (1994a) models are referred to here as F2 and F3. (He termed them his model 2 and model 3). Chapter 1 provides an introductory discussion and justification for these models. They were not designed to account for effects of phase, spatial frequency, or probability summation across mechanisms. They attempt to describe data as a set of mathematical functions and are specified in terms of experimental stimulus components, rather than the outputs of multiple psychophysically or neurophysiologically derived channels. Mechanism sensitivity is determined by the experimental data to which they are fit whilst only the most fundamental mechanism properties are considered. It is assumed that channel bandwidths are sufficiently narrow that each stimulus component excites a different channel, thus: (i) the mask provides no excitation to the channel excited by the test, and (ii) the test provides no excitation to the channel(s) excited by the mask. Foley (1994a) showed this assumption to be correct for his study (in which mask and test stimuli were remote from each other) by showing that the within-channel model failed in describing his data. The pathways assigned by the models can be conceptualised as being groups of channels that respond to the particular stimulus components. The implications of this simplifying assumption in the modelling of data from conditions where remote masks would have been expected to provide (some) excitation of the test detecting mechanism (shown in Figs. 3.9 and 3.11) is discussed later (section 4.9).

Foley's models are defined below. A key to their terms follows:

r_t = response of the test detecting mechanism
c_t = equivalent to the contrast(s) of the test (+ pedestal)
γ_i = equivalent to the contrasts of the remote mask components
z = semi-saturation constant
w = weighting factor
p = excitatory exponent
q = inhibitory exponent
k = a criterion amount - see below

F2 can be expressed as:

$$r_t = c_t^p / \{z^q + (\Sigma[w.\gamma_i] + c_t)^q\}, \tag{4.1}$$

where c_t is proportional to the contrast of the test (+ pedestal), and where $\Sigma[w.\gamma_i]$ provides the weighted sum of the contributions to inhibition from the (remote) mask components. One interpretation of c_t appearing on the denominator is that it reflects self-inhibition from the test (+ pedestal) component. The terms of the denominator in F2 and comparable models will be collectively referred to as the divisive inhibitory pool; this describes the effect that they exert upon mechanism excitation (i.e. the terms of the numerator).

The inputs of the test and mask components to the denominator of F2 are raised to q after summation. In F3, these inputs are summed after being raised to q. F3 can be expressed as:

$$r_t = c_t^p \,/\, \{z^q + \Sigma[(w.\gamma_i)^q] + c_t^q\}. \qquad (4.2)$$

Schematic illustrations of F2 and F3 are presented in Figs. 5.5 and 5.6, respectively, of chapter 5. Of particular interest is that F2 predicts linear-suppression because it linearly sums inputs from mask components. F3 on the otherhand, does not.

F2 and F3 share the assumption that a contrast increment in the test is detected when the response of the test detecting mechanism (r_t) to the test and mask is greater than its response to the mask alone by some criterion k. That is, at contrast increment threshold:

$$r_{t(mask+test)} - r_{t(mask)} = k. \qquad (4.3)$$

4.7 Fitting Foley's models to remote masking data

The results of experiment 1(b) (Fig. 4.4) were produced with masks remote from the test in both spatial frequency and orientation. Assuming the stimuli were exciting narrowly tuned channels, it seems likely that this masking was produced by cross-channel effects. F2 and F3 were tested by fitting them to these data. It was expected that F2 would produce the better fits because the mask terms are summed linearly in its divisive inhibitory pool; this linear-suppression is a particularly striking feature of the data.

Modelling Methods

The models were fit by the downhill simplex iterative optimization algorithm (Press et al., 1989) operating on a Pentium PC. They were initialised with 100 pseudo-randomly selected sets of parameters. After each fit, the minimum RMS errors achieved from the procedures were recorded. In all cases, the best fits from the 100 starts were arrived at by many of the initial parameter sets. F2 and F3 were evaluated in terms of their qualitative predictions, and their quantitative fits in terms of RMS error.

When test and mask patterns are spatio-temporally matched, 'dipper' shaped masking functions are produced. The contrast response function of the channel most sensitive to the test pattern is determined from this function (see chapter 1). Since such functions were not measured in the present experiments however, p (the test channel's nonlinearity) and z (the semi-saturation constant) were fixed in the modelling at 2.5 and 1, respectively. These values are approximate to those typically reported where dipper functions have been measured (for example, see Foley's papers and chapter 5).

Results and Discussion

The left- and right-hand panels of Fig. 4.4 show best fits of F2 and F3 respectively, to the contrast of a 1cpd test as a function of the stimulus contrast of 3cpd grating (filled symbols, solid curves) and plaid (open symbols, dashed curves) masks. Tables 4.1 and 4.2 show the parameter values and the goodness of these fits. As expected, F2 fit consistently better than F3.

	k	p	q	z	w	RMS error (dB)	Number of datum points
DJH (Fig. 4.5)	0.17	2.5(fixed)	2.25	1(fixed)	0.11	0.38	9
GR (Fig. 4.5)	0.05	2.5(fixed)	2.14	1(fixed)	0.19	0.40	9
CEB (Fig. 4.5)	0.06	2.5(fixed)	2.16	1(fixed)	0.19	0.57	11

Table 4.1. F2 parameter values and goodness of fits for data in Fig. 4.4.

	k	p	q	z	w	RMS Error (dB)	Number of datum points
DJH (Fig. 4.6)	0.16	2.5(fixed)	1.86	1(fixed)	0.32	0.94	9
GR (Fig. 4.6)	0.05	2.5(fixed)	1.88	1(fixed)	0.37	0.75	9
CEB (Fig. 4.6)	0.06	2.5(fixed)	1.59	1(fixed)	0.53	1.24	11

Table 4.2. F3 parameter values and goodness of fits for data in Fig. 4.4.

Despite the nonlinear summation of mask components in F3, this model achieved a reasonably good fit. The goodness of fit is not the only reason why F2 should be seen to fit the data more accurately than F3, however. The size of the difference between the grating and plaid masking functions, predicted by F3, is dependent upon the magnitudes of, and the relationship between, its exponents p and q. The difference predicted is important because of the linear-suppression evident in the masking data. If the difference is small, the model would appear more viable than if it were large. Foley's models have been fit to larger bodies of data. For example, Foley (1994a) used his models to simultaneously fit masking functions produced when test and mask were spatially matched and when they were remote in orientation. When simultaneously fitting such functions, it was appropriate to allow all parameters to vary. Typically, it was found that $p \approx q + 0.4$. This is because $1 - (p - q)$ describes the slope of the upper region of a dipper function on log-log axes and, with stimuli similar to those used in the present experiments, the gradients of such slopes tends to be in the region of 0.6. (For a review, see Legge, 1981.) As already noted, however, masking was not measured with spatially matched stimuli here, so the parameters were not constrained in this way. Across observers, $(p - q)$ for fits of F2 were around 0.3 and for F3 they were around 0.7. So, F3 predicts a much shallower slope for within-channel masking than what would be expected. Meese and Holmes (2002) discuss the influence of p and q in detail. Their Fig. 4 shows F3's predictions of the difference between plaid and grating masking with different values of p and q. The larger the difference between p and q, then the larger the difference between grating and plaid mask predictions. If $p = 2.5$ and $(p - q = 0.4)$ then the difference would be around 3dB. In the literature p has very rarely taken values lower than 2.5, though it has occasionally been reported with a value of around 3. If $p = 3$ for example, and $(p - q)$ was again 0.4, then the difference between plaid and grating masking would be around 3.5dB. To summarise, RMS values indicate that F2 is better than F3.

Furthermore, resetting p and q to be more consistent with values arrived at when fitting within-channel masking (contrast discrimination dipper) data amplifies this difference.

Experiment 1(c). Fitting Foley's models to more remote masking data

The aim here was to investigate whether F2 also fit other remote masking data better than F3. The methods used in gathering the data in Fig. 4.4 were repeated, but the spatial frequencies of the stimuli were two or three times higher than those used previously (6cpd masks with 2cpd tests or 9cpd masks with 3cpd tests). The numbers of cycles per stimulus patch were the same as before. The data were gathered from DJH and were fit using the modelling methods described previously.

Fig. 4.6 shows masking as a function of grating (filled symbols) and plaid (open symbols) mask stimulus contrast. Test and mask spatial frequencies were 3 and 9cpd, respectively, and the mask components were oriented 45deg from the test. Since there was very little masking regardless of mask contrast, modelling was inappropriate. This prompted a further experiment in which the angle between test and mask components was reduced to 22.5deg.

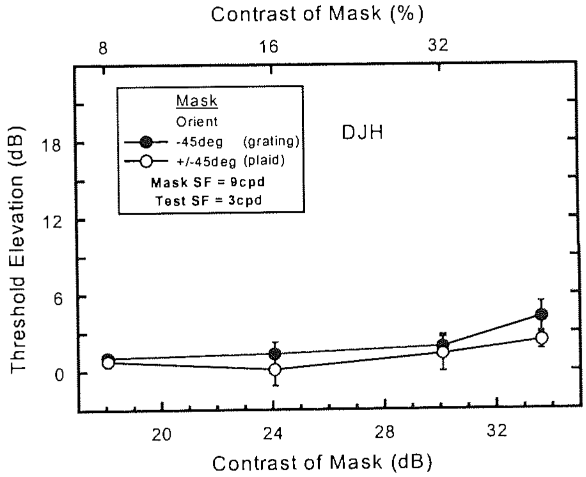

Fig. 4.6. (Above) Masking as a function of mask stimulus contrast. Masks were -45° gratings (filled symbols) or ±45° plaids (open symbols). The test and masks had 3 and 9cpd, respectively.

The left- and right-hand panels of Fig. 4.7 show the same data fit by F2 and F3, respectively. Test contrast is shown as a function of grating (filled symbols, solid curves) and plaid (open symbols, dashed curves) mask stimulus contrast. The top panels show the effects of 6cpd masks on the detection of a 2cpd test when the angle between test and mask components was 45deg. The bottom panels show the effect of 9cpd masks on the detection of a 3cpd test when the angle between the test and mask components was 22.5deg.

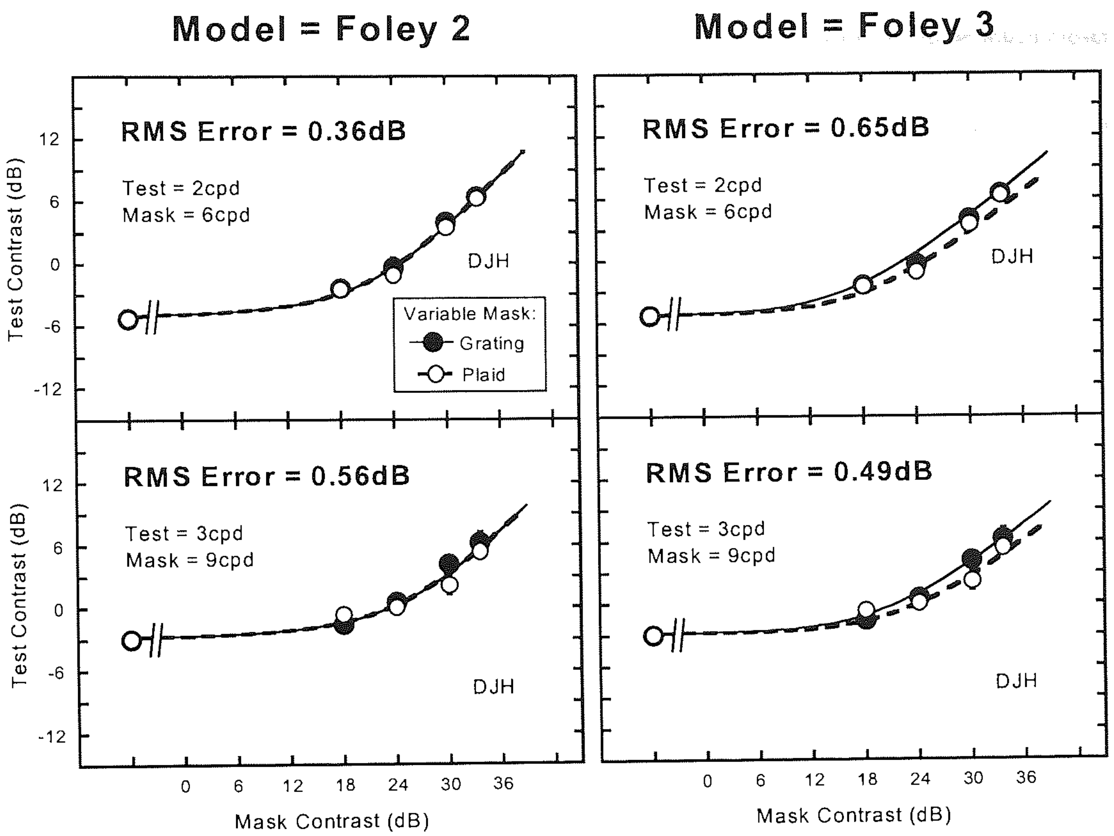

Fig. 4.7. Test contrast as a function of the stimulus contrast of grating (filled symbols, solid curves) and plaid (open symbols, dashed curves) masks. The left- and right-hand panels show fits of F2 and F3, respectively. Test and mask spatial frequencies were 2 and 6cpd, respectively (top panels), or 3 and 9cpd, respectively (bottom panels).

	k	p	q	z	w	RMS error (dB)	Number of datum points
6on2	0.18	2.5(fixed)	2.26	1(fixed)	0.05	0.36	9
9on3	0.30	2.5(fixed)	2.27	1(fixed)	0.03	0.56	9

Table 4.3. F2 parameter values and goodness of fits. Data and fits are shown in the left-hand panels of Fig. 4.7.

	k	p	q	z	w	RMS Error (dB)	Number of datum points
6on2	0.17	2.5(fixed)	1.75	1(fixed)	0.16	0.65	9
9on3	0.28	2.5(fixed)	1.73	1(fixed)	0.11	0.49	9

Table 4.4. F3 parameter values and goodness of fits. Data and fits are shown in the right-hand panels of Fig. 4.7.

The grating and plaid masking functions were not as similar as when the test and masks had 1 and 3cpd, respectively (Fig. 4.4). Particularly with the higher spatial frequency stimuli, the plaids tended to produce less masking than the gratings. Consequently, the fits of F2 and F3 were more alike and one model didn't convincingly fit the data better than the other. Again though, the value of (p - q) is important in assessing the two models. Fits of F2 to both data sets deliver (p - q) values of around 0.25 whilst fits of F3 deliver values of around 0.75. As already discussed, fits to larger data sets gathered with similar stimuli, where all parameter values were free to vary, have consistently shown the (p - q) value to be much closer to 0.25 than 0.75.

In summary, F2 has tended to fit the remote masking data better than F3. However, F2's superiority is not as convincing with data gathered using the higher spatial frequency stimuli. More thorough tests are required to determine which model produces the best fits to remote masking data. Such tests are shown in the following chapter.

4.8 Revisiting data in chapter 3. Linear-suppression from many stimulus spatial frequency combinations

Introduction

The aims of the experiments in this chapter were: (i) to perform a systematic investigation of linear- and super-suppression, and (ii) to use data gathered for (i) to test Foley's models. Linear-suppression was found to be far more prevalent than super-suppression and F2 typically fit data better than F3 since it predicts linear-suppression.

Figs. 3.9 and 3.12 of chapter 3 show differences between plaid and grating masking (TE plaid RE grating) for a large range of test and mask spatial frequency combinations. If the masks had been matched in stimulus contrast, superposition of masking functions would have indicated linear-suppression whilst (plaid masking > grating masking) would have indicated super-suppression. The masks were matched in component contrast, however, so the type of suppression that the masking is consistent with is not directly apparent from the functions. As part of the present investigation, these functions are revisited and modelling has been applied to reveal whether they are consistent with linear-suppression, super-suppression, or some other type of suppression.

In Figs. 3.9 and 3.12, the data are fit by Foley's models. (Also fit is the hybrid model. Discussion of this model is witheld until chapter 5. Note only that its fits are the same as F2's). F2 has been used to estimate how much more masking the plaid would be expected to produce than the grating for linear-suppression. A difference greater than this estimate would indicate super-suppression.

Modelling Methods

In order to simplify the analysis, the contribution from the contrast of the test to the denominator of the models (which was negligible anyway) was set to zero.

F2 and F3 are rearranged below to calculate the difference between plaid and grating masking. The model terms are:

$r_{t(gm)}$ = response of the test detecting mechanism when the test is masked by the grating

$r_{t(pm)}$ = response of the test detecting mechanism when the test is masked by the plaid

$c_{t(gm)}$ = contrast of the test, at detection threshold, when it is masked by the grating

$c_{t(pm)}$ = contrast of the test, at detection threshold, when it is masked by the plaid

z = semi-saturation constant

p = excitatory exponent

q = inhibitory exponent

w = weighting of mask components

γ = contrast of each remote mask component

Consider F2:

When the test is masked by the grating, F2 is given by:

$$r_{t(gm)} = c_{t(gm)}{}^{p} \;/\; \{z^{q} + (w.\gamma)^{q}\}$$

and for plaid masking, F2 is given by:

$$r_{t(pm)} = c_{t(pm)}{}^{p} \;/\; \{z^{q} + (2.w.\gamma)^{q}\}$$

Assuming $r_{t(gm)} = r_{t(pm)}$ at detection threshold, then:

$$c_{t(gm)}{}^{p} \;/\; \{z^{q} + (w.\gamma)^{q}\} = c_{t(pm)}{}^{p} \;/\; \{z^{q} + (2.w.\gamma)^{q}\}$$

Therefore, the ratio of the plaid and grating masking is given by:

$$c_{t(pm)} \;/\; c_{t(gm)} = [\{z^{q} + (2.w.\gamma)^{q}\} \;/\; \{z^{q} + (w.\gamma)^{q}\}]^{1/p}$$

Now consider F3:

For grating masking, F3 is defined as:

$$r_{t(gm)} = c_{t(gm)}{}^{p} \;/\; \{z^{q} + (w.\gamma)^{q}\}$$

For plaid masking, it is defined as:

$$r_{t(pm)} = c_{t(pm)}{}^{p} \;/\; \{z^{q} + (w.\gamma)^{q} + (w.\gamma)^{q}\}$$

Applying the sequence gives:

$$c_{t(pm)} \;/\; c_{t(gm)} = [\{z^{q} + (w.\gamma)^{q} + (w.\gamma)^{q}\} \;/\; \{z^{q} + (w.\gamma)^{q}\}]^{1/p}$$

The contrast of the remote mask components (γ) was 24% for the data in Fig. 3.9 and those in the bottom panel of Fig. 3.12. It was 8% for the data in the top panel of Fig. 3.12. The data do not constrain the model fits sufficiently to allow p and q to be estimated from them. Therefore, the difference between the plaid and grating masking was calculated for many combinations of values for the parameters. The values have been calculated by myself (see chapter 5) and others (see the work of Foley and colleagues) by fitting the models to data sets where all parameters were free to vary. The ranges of parameter values used here were chosen to resemble those:

- $p = 2.25$ to 4.00 (in steps of 0.25)
- $(p - q) = 0.1$ to 0.4 (0.1 steps)
- $z = 1.0$ to 3.0 (0.5 steps)
- $w = 0.2$ to 1.0 (0.2 steps)

Results and Discussion

In Figs. 3.9 and 3.12, the ranges of the estimates of F2 and F3 for (TE plaid RE grating) are represented with light shading and darker shading, respectively. (Note that the darkest shading in the top panel of Fig. 3.12 indicates the region where the estimates of F2 and F3 overlap). The estimates were delivered with the parameters of the models set to reasonable values. The lowest and highest values were as follows:

- When mask component contrast was 8%, the smallest difference predicted by F2 was 1.56dB ($p = 4.0$, $q = 3.6$, $z = 3.0$, $w = 0.2$) and the largest difference was 5.87dB ($p = 4.0$, $q = 3.9$, $z = 1.0$, $w = 1.0$).

- When mask component contrast was 8%, the smallest difference predicted by F3 was 0.17dB ($p = 4.0$, $q = 3.9$, $z = 3.0$, $w = 0.2$) and the largest difference was 2.65dB ($p = 2.25$, $q = 2.15$, $z = 1.0$, $w = 1.0$).

- When mask component contrast was 24%, the smallest difference predicted by F2 was 4.25dB ($p = 2.5$, $q = 2.1$, $z = 3.0$, $w = 0.2$) and the largest difference was 5.87dB ($p = 4.0$, $q = 3.9$, $z = 1.0$, $w = 0.8/1.0$).

- When mask component contrast was 24%, the smallest difference predicted by F3 was 1.33dB ($p = 4.0$, $q = 3.6$, $z = 3.0$, $w = 0.2$) and the largest difference was 2.67dB ($p = 2.25$, $q = 2.15$, $z = 1.0$, $w = 1.0$).

Note that the estimates were calculated with the lowest value of w being 0.2. Very low values of w indicate very little cross-channel masking and consequently, very little difference between plaid and grating masking. (When $w = 0$, F2 and F3 are identical). Thus, data below the range of model predictions (but $\geq$ 0dB) could be predicted if w was set at some value lower than 0.2. Evidence for this was the very low levels of masking sometimes produced. (See Figs. 3.8 and 3.11, from which the currently discussed data were calculated).

Consider the predictions in Fig. 3.9. F2 captures the data well when mask spatial frequency was 1cpd and test spatial frequency was close to this (within about an octave). This is consistent with the linear-suppression evident when masks were matched in stimulus contrast and matched to the test in spatial frequency at 1cpd (Fig. 4.1). When mask spatial frequency was 3cpd, F2 captures the data well when test spatial frequency was about 1cpd or less. This is consistent with the linear-suppression shown when masks were matched in stimulus contrast and mask and test spatial frequencies were 3 and 1cpd, respectively (Fig. 4.4). Fits were of particular interest when masks had 3cpd and the tests had a spatial frequency very close to this. The functions rise above the predictions of both models. Obviously, this super-suppression can't be captured by F2 because that model features linear-suppression. However, the middle panel of Fig. 3.9 clearly shows that F3's fit is even worse than F2's. This is because it predicts plaid masking to be less than grating masking, rather than the other way around (this is also evident in Figs. 4.1 & 4.2). The differences between 12cpd plaid and grating masking were very small. This was due to the low level of cross-channel masking produced with these stimuli (see Fig. 3.8). Moving onto Fig. 3.12, this shows the effect of mask contrast on super-suppression. (Note that the data in the bottom panel of Fig. 3.12 are the same as those in the middle panel of Fig. 3.9). The data and fits reflect what was shown previously when stimuli were matched with 3cpd and the masks were matched at a range of stimulus contrasts (Figs. 4.1 and 4.2). That is, super-suppression is far more evident when the mask is of a high contrast than when it is of a lower contrast. With the lower mask contrast (top panel), F2 predicts the data for a wide range of test spatial frequencies.

In summary, linear-suppression is widespread when test and mask are remote in orientation. Super-suppression is evident only when test and masks have close to 3cpd and the masks are of a high contrast. Whilst F2 predicts linear-suppression, neither F2 nor F3 predict super-suppression.

4.9 Experiment 2. Linear-suppression and mask component orientation

Introduction

Masking from gratings and plaids (matched in stimulus contrast) was almost exactly the same when the test and masks had 1 and 3cpd, respectively (Fig. 4.4). This is the best evidence found so far for linear-suppression; a phenomenon that has important implications for quantitative models of masking. The robustness of the effect was investigated by manipulating plaid angle.

Methods

Grating and plaid mask components were oriented between 0 and -90deg, and 0 and ±90deg, respectively. (The orientation difference between vertical and one of the plaid's components was always the same as that of the other component). Two conditions were run in each session; in one, the mask was a grating, whilst in the other, it was a plaid. Trials of each condition were blocked together and the sequence of the blocks was varied. In each session, the orientation difference between the mask components and vertical was the same for the gratings and plaids. Test and mask components were in cosine-phase with the fixation point. This was different from all of the other experiments in this chapter in which the components were in sine-phase. The reason for this was that the ±90° plaid mask with components in sine phase would have 0% contrast. However, the ±90° plaid mask with components in cosine phase has the same contrast as the other masks. Mask stimulus contrast was 32% and unmasked thresholds were randomly measured across sessions. Data were gathered from DJH and TSM.

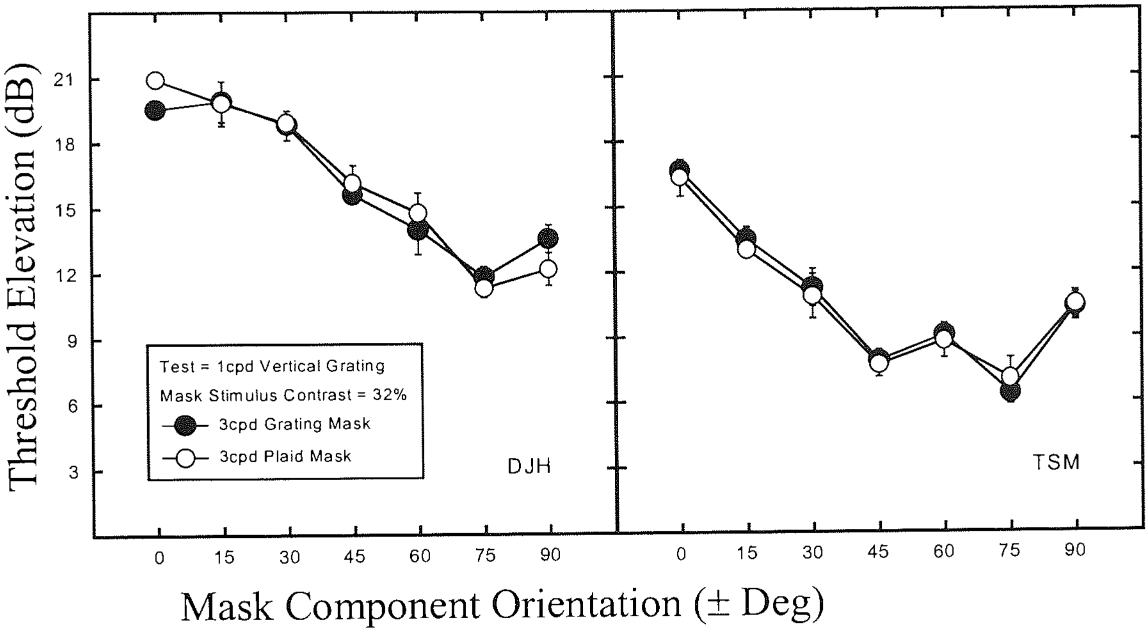

Fig. 4.8. Grating and plaid masking (filled and open symbols, respectively) as a function of mask component orientation. Tests and masks had 1 and 3cpd, respectively. Mask stimulus contrast was 32%. Data are from DJH (left panel) and TSM (right panel).

Results and Discussion

Fig. 4.8 shows the masking of a vertical 1cpd test produced by 3cpd grating (filled symbols) and plaid (open symbols) masks, as a function of mask component orientation. Data are from DJH (left panel) and TSM (right panel). Grating and plaid masking is almost exactly the same. This is strong evidence for linear-suppression across a wide range of test-mask orientation differences. F2 would produce better fits than F3 because F2 predicts linear-suppression whilst F3 does not.

4.10 Chapter Discussion

In these experiments, remote masking was measured to investigate the extent of linear- and super-suppression, and to test predictions of Foley's (1994a) cross-channel suppression models (F2 and F3). Many instances of linear-suppression were found, whereas fewer instances of super-suppression were found, suggesting it is less widespread. Consistent with chapter 3, super-suppression only occurred when stimuli had around 3cpd and the masks were of a high contrast. Concerning the model fits, though the models have been reported in the literature to predict a large body of remote masking data, both fail to predict super-suppression. When Foley's models were fit to data where super-suppression was not evident, F2 produced the better fits because it predicts linear-suppression.

No direct tests were performed to reveal whether the masking was due to within- or cross-channel effects. However, a variety of evidence from the literature suggests that it was due to cross-channel effects.

Examples are channel bandwidth estimates and larger bodies of remote masking data, in which there are telltale characteristics (e.g. Foley, 1994a; Ross & Speed, 1991). These are discussed further in chapter 5 where data providing strong support are shown.

Though F2 produced the better fits in this chapter, it is not generally better than F3 at capturing remote masking data. In the paper in which the models were introduced (Foley, 1994a), data were shown to which F3 produced the better fits. To understand this discrepancy, a consideration is required of the two main differences between my investigations and Foley's. Firstly, whilst he measured the effect from only a one-component remote mask, I compared the effects from one- and two-component remote masks. Secondly, whilst he measured contrast discrimination (with a variable contrast pedestal), I only measured contrast detection. Regarding the first difference, Foley's experiments did not address the issue of how the inputs from more than one remote mask component may be summed. My experiment did and the data that I gathered represented linear-suppression, which is a feature of F2. Regarding the second difference, whilst Foley probed how self-inhibition is summed with inputs from remote mask components, I didn't. Self-inhibition is inhibition from components that excite the test detecting mechanism. In my experiments, self-inhibition was assumed to be only produced by the test component. Because this was at detection threshold, it had a negligible effect on the responses of the models. In Foley's experiments however, both test and pedestal contributed to self-inhibition. Because the pedestal was varied between low and high contrast, self-inhibition had a substantial effect on the model predictions. F2 and F3 treat self-inhibition differently. In F2, self-inhibition is added to the inputs from the remote mask components before the sum is raised to an exponent. In F3, self-inhibition is raised to an exponent before it is added to the inputs from the remote mask components. Foley's data showed that the fixed contrast mask raised test discrimination thresholds at low pedestal contrasts (when self inhibition was small) but lowered them at high pedestal contrasts (when self inhibition was high). F3's treatment of self-inhibition predicted this whilst with no parameter values that I have tested (see next chapter) was this feature displayed by F2. This was the main reason why F3 fit Foley's data better than F2.

In summary, F2's feature of linear-suppression is important in fitting data shown here, whilst F3's non-linear summation of self-inhibition with other inhibitory inputs was important in fitting Foley's (1994a) data. To simultaneously describe the data here with Foley's (1994a), a new model is required which has both of these features. Before attempting to develop such a model, a single experiment was designed that would address summation of both (i) remote mask components and (ii) remote mask components and pedestal. This experiment and the developments of Foley's models are shown in chapter 5.

Chapter Five:
Testing Some Cross-Channel Masking Models

5.1 Chapter Introduction

In the last chapter, it was shown that inputs from remote mask components can sum linearly in the divisive inhibitory pool. Foley has shown that these inputs sum nonlinearly with self-inhibition in this pool. Can a single model be found to accommodate both of these findings? Before attempting to develop such a model, a single experiment was designed that would address summation of both (i) remote mask components, and (ii) remote mask components and pedestal. The experiment, which is an extension of Foley's (1994a), is featured in this chapter. Like his, dipper functions were measured in the presence and absence of a superimposed remote mask of fixed contrast. The experiment extends Foley's by considering effects of both one- and two-component fixed masks. This provides the opportunity to simultaneously investigate how remote components are summed with themselves as well as with self-inhibition in the divisive inhibitory pool.

A less important difference between Foley's (1994a) study and the present study concerns stimulus spatial frequency. Whilst all of Foley's (1994a) stimuli were spatial frequency matched with 2cpd, the test + pedestal used here had 1cpd whilst the fixed contrast masks had 3cpd. There were two main reasons for this change. This experiment was designed to investigate cross-channel masking. Foley used stimuli matched in spatial frequency but he used a fixed contrast mask whose orientation was orthogonal to the test. Estimates of the orientation bandwidth of channels (see chapter 1 and 6 for discussions) suggest that orthogonal patterns excite different channels. As their orientation difference is reduced however, the likelihood of them providing some excitation to the same channel increases. Some excitation (though only a small amount) would be likely when they differ by 45deg. To simplify the modelling, the assumption was made that the stimulus components excite different channels. It was made likely that this was actually the case by: (i) the components of the plaid mask being orthogonal to each other and differing by 45deg from the test, and (ii) making the test and mask spatial frequencies remote (since they could differ in orientation by only as much as 45deg). The second reason for using these stimuli was that they revealed linear-suppression in a contrast detection study (chapter 4). Using them here provided the opportunity to investigate how widespread this very interesting phenomenon is with a contrast discrimination study. Of course, this was central to the design of the experiment, which aimed to show the summation of these mask components with the pedestal.

The results from this new experiment display important characteristics that are not all captured by either F2 or F3. Some alternative cross-channel masking models are introduced to capture these effects.

5.2 Methods

(i) Experimental Methods

The 2-AFC procedure described in chapter 2 was used to measure the contrast discrimination thresholds of test patterns in the presence of variable and fixed contrast spatio-temporally superimposed masks. The stimuli had a duration of 33ms and were temporally curtailed with a square wave. The tests and variable contrast masks (pedestals) were 1cpd vertical gratings, which were spatially curtailed with a circular Gaussian function. The patches had 1.67cycles at full-width, half-height. Pedestal contrast was varied across experimental sessions and ranged from 0% to 31.6%. The fixed contrast masks had 3cpd, were spatially curtailed with a raised cosine envelope and measured 4.4° at full-width, half-height (their plateaus had about 10 ½ cycles). They had a contrast of 20%. There were three fixed mask conditions: no-mask, -45° grating mask and ±45° plaid mask. Examples of the stimuli used are shown in Fig. 5.1. Note that, for clarity, they are shown at a higher contrast to that at which they were presented in the experiments. In a single session, pedestal contrast was fixed and trials from the three mask conditions were blocked together and their sequence was randomised. Pedestal contrasts were selected randomly and each stimulus condition was performed 3 (TSM) or 5 times (DJH).

(ii) Modelling Methods

Modelling methods were the same as those described in chapter 4, section 4.7. In addition to F2 and F3, four additional models of physiologically plausible alternatives were fit. These were termed the hybrid model, the compound model, the compound-plus model and the alternative compound-plus model. The last two of these had six free parameters whilst the others had five. No parameters were fixed in the fitting procedures.

Variable Contrast Mask (Pedestal):

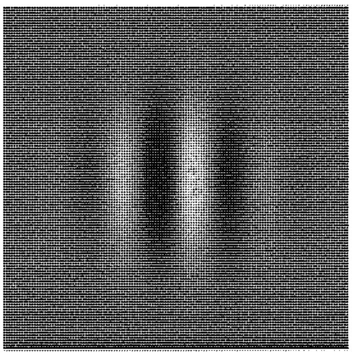

Fixed Contrast Mask:

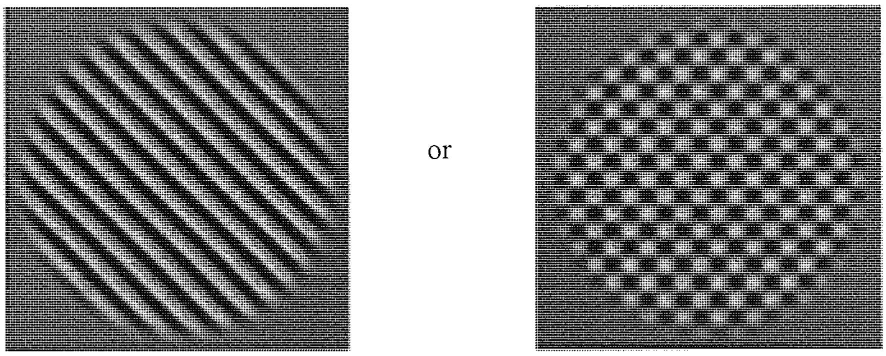

Test:

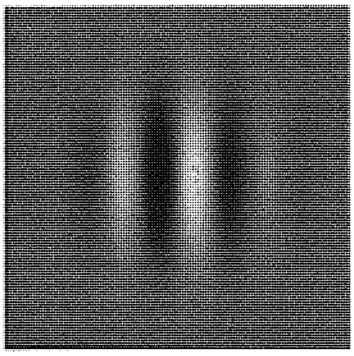

Fig. 5.1. Stimuli. Note there was also a fixed contrast mask of 0% contrast (no-mask condition). For clarity, the stimuli are shown at a higher contrast to that at which they were presented in the experiments.

5.3 Results and Discussion

Discrimination thresholds of a 1cpd test grating in the presence of a 3cpd grating (filled symbols) or plaid (open circular symbols) fixed contrast mask, or no fixed contrast mask (open triangular symbols) are shown in Figs. 5.2, 5.3 and 5.4. Fig. 5.2 presents a line graph of the data for both observers whilst Figs. 5.3 and 5.4 present model fits to the data of DJH and TSM, respectively. The parameter values and goodness of fit of these are shown in Tables 5.1 to 5.6.

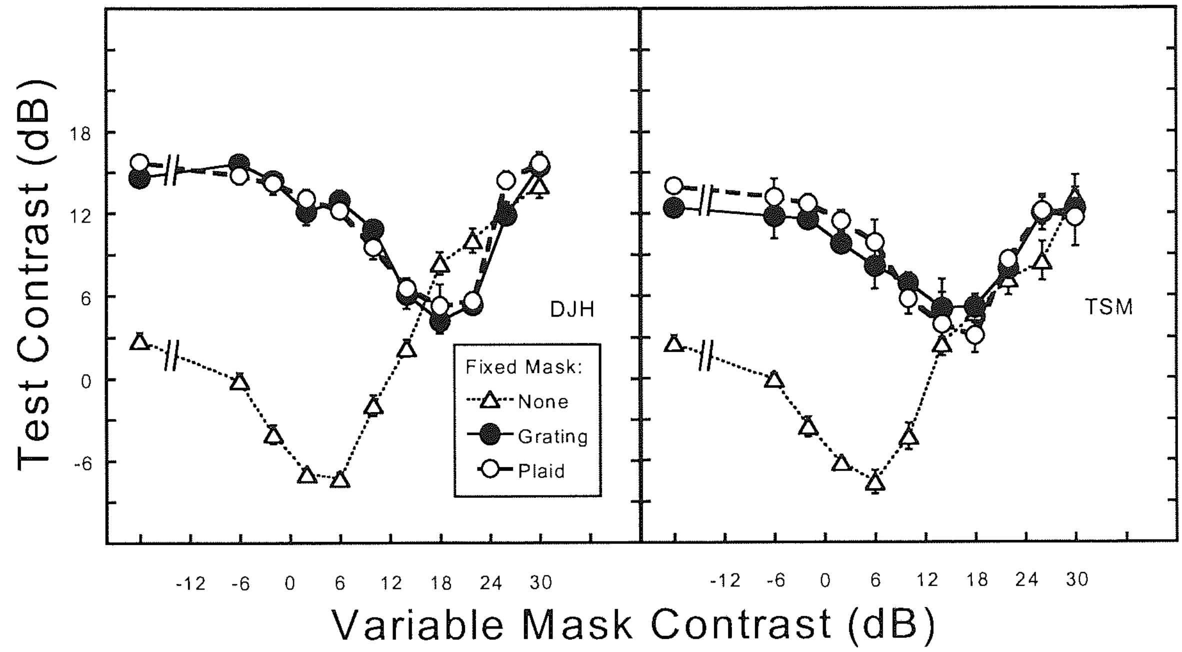

Fig. 5.2. The effect of -45deg 3cpd grating (filled symbols) and ±45deg 3cpd plaid (open symbols) fixed contrast masks on test contrast, as a function of pedestal contrast, when test and pedestal were vertical 1cpd gratings. The fixed contrast masks had a stimulus contrast of 20%. Data are from DJH (left panel) and TSM (right panel).

Contrast discrimination in the absence of a fixed contrast mask took the form of a dipper function. As the contrast of the pedestal was increased, test thresholds were first lowered and then raised. This has been reported many times previously (e.g. Nachmias & Sansbury, 1974, Legge & Foley, 1980; Foley, 1994a). Note that, throughout this chapter, 'dipper function' and 'dipper' are terms used only in reference to the functions produced by the pedestal in the absence of fixed contrast masks.

The dip was preserved with the fixed contrast masks, though the function was transformed such that thresholds were raised in its lower limb. A first approximation of their effect is that of a diagonal translation (upward and rightward) of the dipper. The shifts are not simply oblique, however. Such a shift (with dipper handles of about 0.6) would not cause the functions produced with and without a fixed contrast mask to crossover. Also, it would not result in the convergence of the functions at high pedestal contrasts. The shift would be similar to that described by the fatigue model of adaptation, discussed by Meese & Holmes (2002). The way in which the contrast discrimination functions were transformed by the fixed contrast masks is very good support for cross-channel masking. As pointed out by Foley (1994a), the standard within-channel model of masking (Legge & Foley, 1980) (see chapter 1) predicts that the fixed contrast masks would shift the dipper laterally by an additive constant on log axes. Obviously, this is not the case. Of great importance is the large dip in the functions produced with the fixed contrast masks; a characteristic that is also evident in Foley's (1994a) data. In both within- and cross-channel masking models (see chapter 1 and the last chapter), such a dip is only predicted when the models are in the accelerating region of their output characteristic. It would only be in this region if the fixed contrast masks did not excite the test detecting mechanism. Because the masks raise test thresholds at low pedestal contrasts, they must be producing masking of the test detecting mechanism. Therefore, whilst they do not excite the test detecting mechanism, they do mask it. This is the signature of cross-channel masking models, and cannot be accounted for by within-channel masking models. This point is also of relevance to the discussion of the contrast detection data in the last chapter where cross-channel masking was assumed but no direct evidence of it was shown. The above finding is strong evidence for this earlier assumption.

For DJH, there is a narrow range of pedestal contrasts (18 to 24dB) where the fixed contrast masks actually improved discrimination. This is very important to the model fitting; it is captured better by F3 than by F2 because of their different treatments of self-inhibition. Foley's (1994a) data also show this crossover. However, this is not seen in the data of TSM.

The contrast discrimination functions produced in the grating and plaid mask conditions (circles) were very similar. There is a systematic displacement of these functions for TSM, which is opposite in direction to that predicted by F3 (Fig. 5.4), but this is slight. So, whilst grating and plaid masking functions were shown to be very similar for contrast detection in chapter 4, the same has been shown for contrast discrimination in this chapter. (The test and remote masks used in the two experiments were matched in spatial frequency and orientation). The data shown in chapter 4 were interpreted as being indicative of linear-suppression; the data shown here appears to be further evidence of linear-suppression. This is discussed later in relation to the models and their fits.

To recap, in a single experiment I have found evidence for both linear-suppression (consistent with chapter 4) and non-linear summation (consistent with Foley, 1994a) of the terms in the denominator of Foley's (1994a) models. This suggests that neither F2 nor F3 are correct. A new model is needed that can accommodate the full set of results.

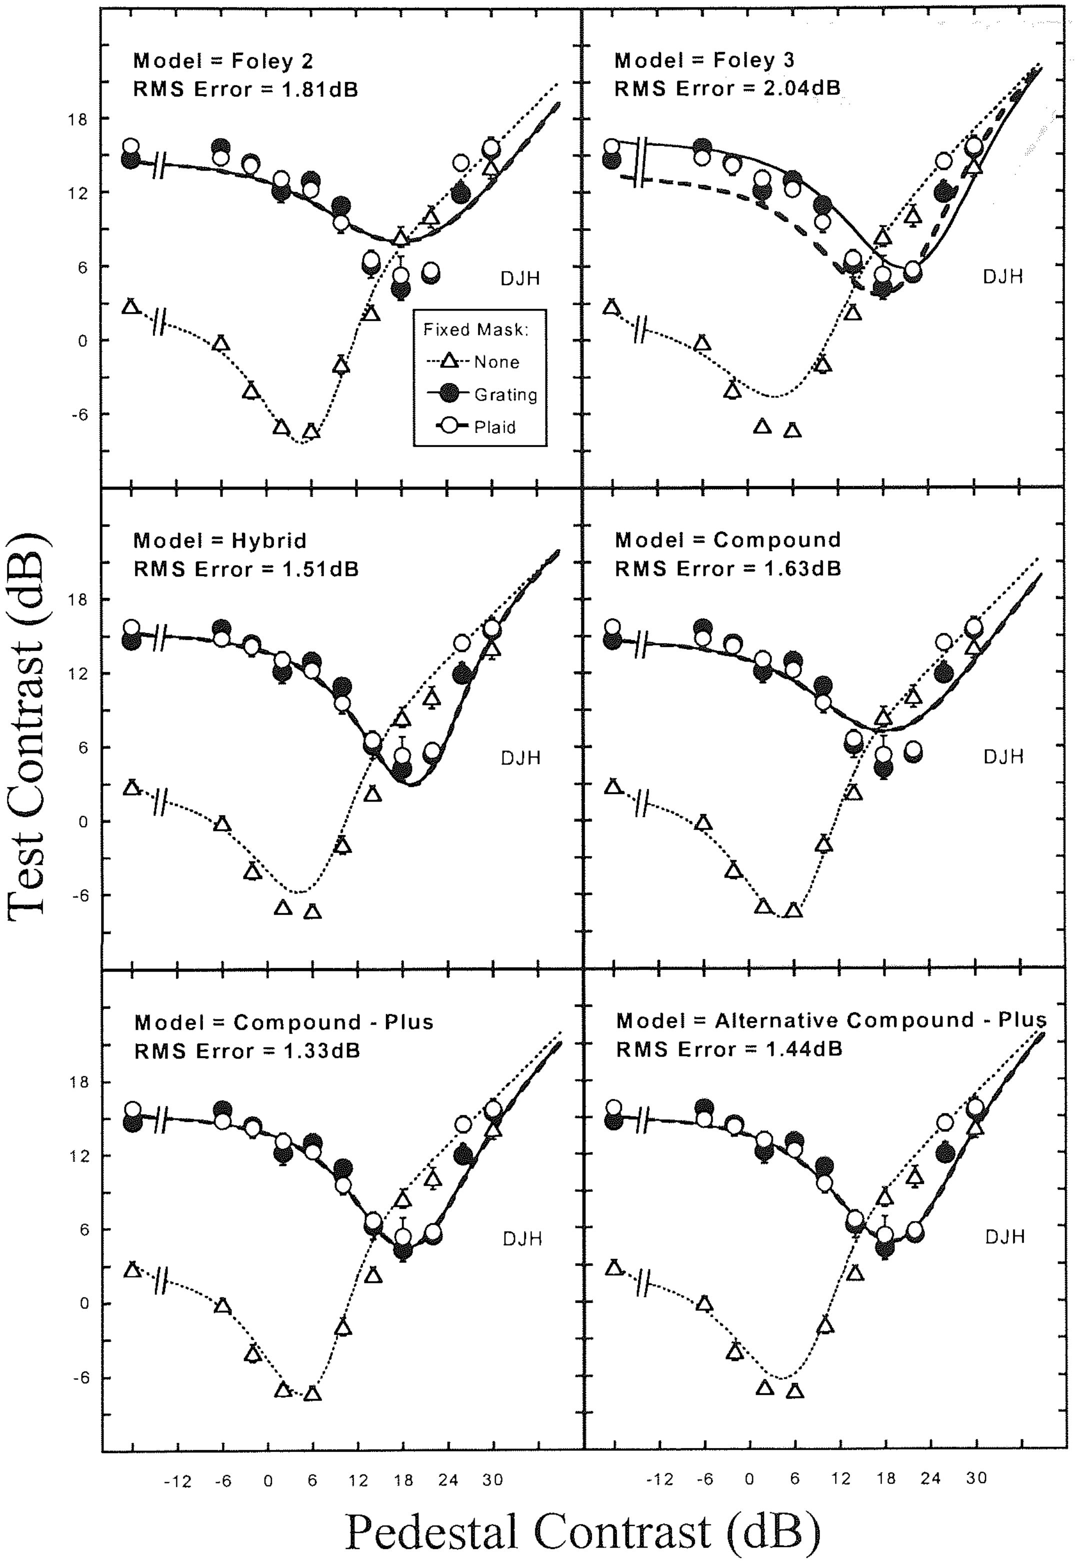

Fig. 5.3. The effect of 3cpd grating (filled symbols) and plaid (open symbols) fixed contrast masks on test contrast, as a function of pedestal contrast, when the test and pedestal were 1cpd gratings. The fixed contrast masks had a stimulus contrast of 20%. Data are from DJH. The same data are shown in each panel fit by a different cross-channel masking model.

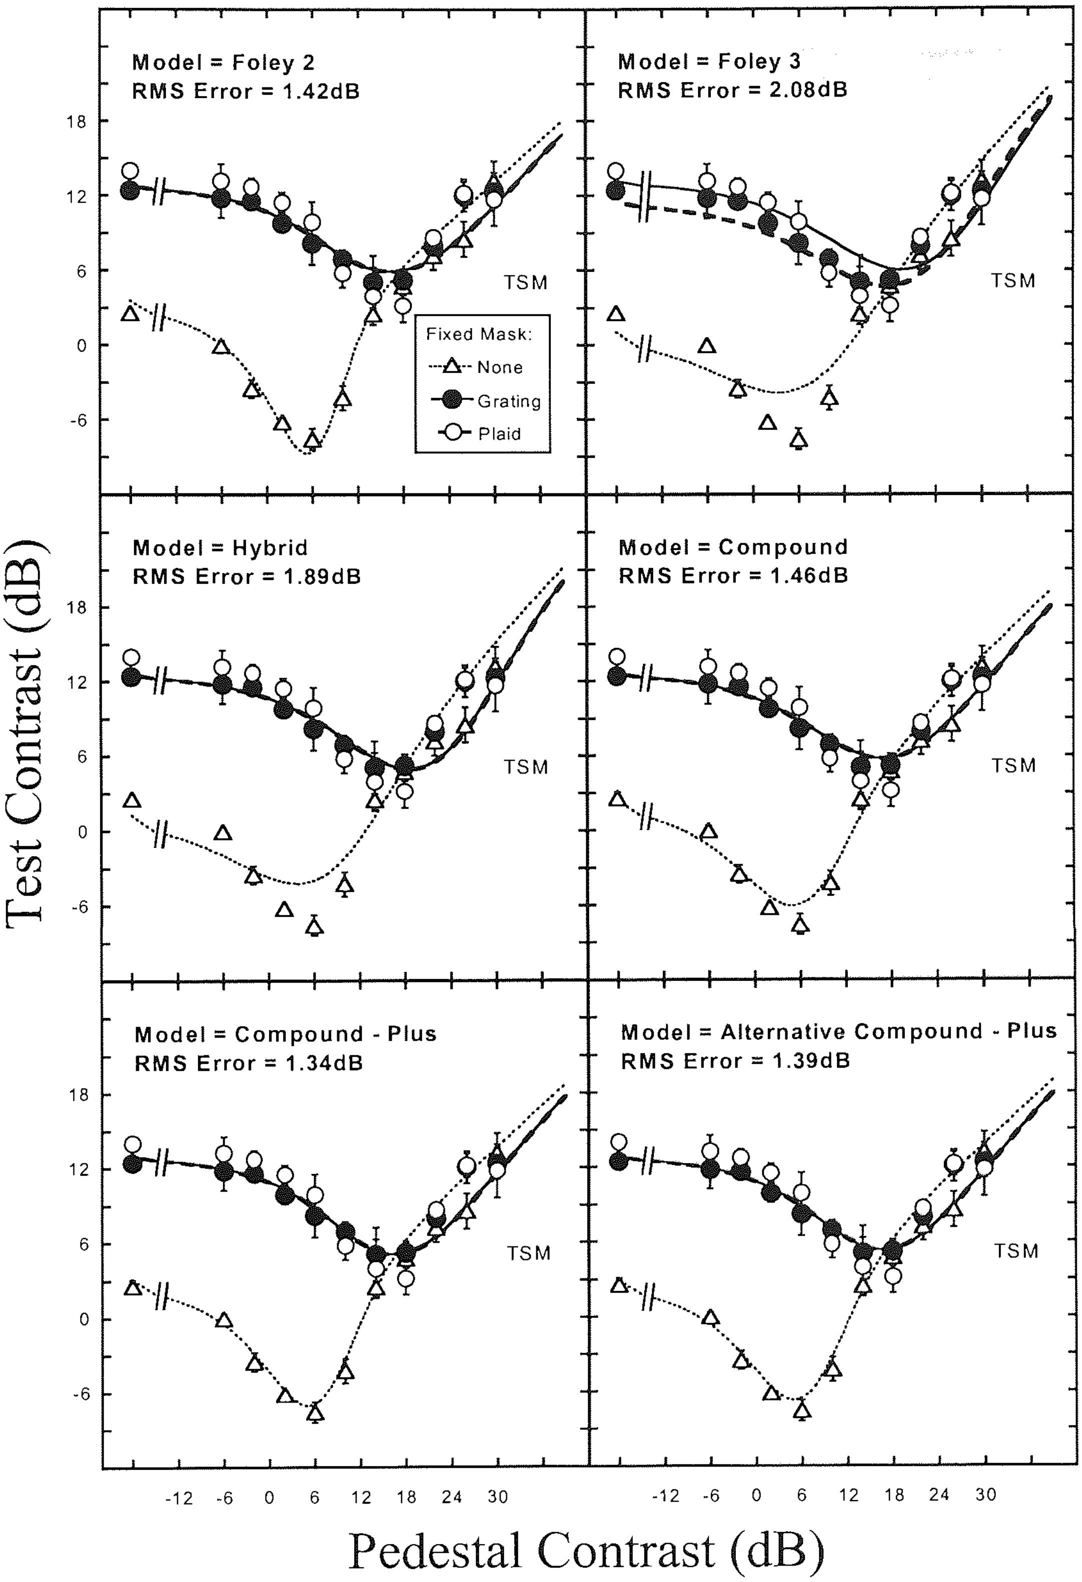

Fig. 5.4. The effect of 3cpd grating (filled symbols) and plaid (open symbols) fixed contrast masks on test contrast, as a function of pedestal contrast, when test and pedestal were 1cpd gratings. The fixed contrast masks had a stimulus contrast of 20%. Data are from TSM. The same data are shown in each panel fit by a different cross-channel masking model.

	k	p	q	z	w	RMS Error (dB)	Number of datum points
DJH (Fig 5.3)	0.24	3.59	3.22	2.08	0.24	1.81	33
TSM (Fig 5.4)	0.27	4.13	3.69	2.09	0.15	1.42	33

Table 5.1. F2.

	k	p	q	z	w	RMS Error (dB)	Number of datum points
DJH (Fig 5.3)	0.28	2.63	2.26	2.13	0.72	2.04	33
TSM (Fig 5.4)	0.21	1.93	1.58	2.71	0.76	2.08	33

Table 5.2. F3.

	k	p	q	z	w	RMS Error (dB)	Number of datum points
DJH (Fig 5.3)	0.30	3.24	2.85	2.02	0.53	1.51	33
TSM (Fig 5.4)	0.21	2.04	1.69	2.62	0.64	1.89	33

Table 5.3. Hybrid model.

	k	p	q	z	w	RMS Error (dB)	Number of datum points
DJH (Fig 5.3)	0.12	3.50	3.13	2.60	0.37	1.63	33
TSM (Fig 5.4)	0.12	2.73	2.33	3.11	0.44	1.46	33

Table 5.4. Compound model.

	k	p	q	z	w	β	RMS Error (dB)	Number of datum points
DJH (Fig 5.3)	0.28	3.70	3.32	2.04	0.43	0.15	1.33	33
TSM (Fig 5.4)	0.27	3.34	2.91	2.15	0.28	0.53	1.34	33

Table 5.5. Compound-plus model.

	k	p	q	z	w	β	RMS Error (dB)	Number of datum points
DJH (Fig 5.3)	0.35	3.21	2.84	1.90	0.42	0.24	1.44	33
TSM (Fig 5.4)	0.41	3.16	2.75	1.85	0.23	0.61	1.39	33

Table 5.6. Alternative compound-plus model.

The models and their fits

The model fits shown in Figs. 5.3 and 5.4 are discussed below. Four of the models that have been fit (hybrid, compound, compound-plus and alternative compound-plus) are new models and are defined and discussed below. These are introduced because neither of Foley's models provided an adequate account of the masking data. Schematic illustrations and mathematical definitions of all the models are given in Figs. 5.5 to 5.9. The alternatives are structured similarly to Foley's in that excitatory test channels are divisively inhibited by a more broadly tuned pool of channel inputs. Although the models are simple mathematical expressions in which operations relate to stimulus components (like Foley's) and not filter outputs, the new models were designed with physiological plausibility in mind. It is important to note that for all of the models it is assumed that mechanism excitation is provided by only the test component or masks that are matched in all dimensions to the test component.

There are three main features of the data, which play a large part in determining how well the models fit. The models differ greatly in how well they capture them. These are:

(i) The crossover: In the data of DJH, the functions produced with and without the fixed contrast masks crossover each other at high pedestal contrasts. There is a pronounced dip in this crossover.

(ii) The linear-suppression: The functions produced with the grating and plaid fixed contrast masks are almost exactly the same.

(iii) The depth of the dippers: The dipper functions (produced in the absence of the fixed contrast masks) had deep dips. That is, the maximum amount of facilitation produced by the pedestal was high.

How well each model captures these characteristics is discussed below.

Definitions of the models' terms follow:

r_t = response of the test channel

c_t = contrasts of the test + pedestal

γ_i = contrasts of the remote mask components ($i > 1$)

z = semi-saturation constant

w = weighting (of the remote mask components)

p = excitatory exponent

q = inhibitory exponent

β = switch

F3

As shown earlier, F3 can be expressed as:

$$r_t = c_t^p \,/\, \{z^q + \Sigma[(w.\gamma_i)^q] + c_t^q\}. \qquad (5.1)$$

Fig. 5.6 shows a schematic illustration of this model. Its fits appear in the top-right panels of Figs. 5.3 and 5.4. F3 raises the inputs of the test + pedestal and the remote masks on the denominator to q before summation. Fits of F3 to Foley's (1994a) data are better than those of F2. The fits of F3 to my data however, were clearly worse than those of F2 (see below). F3 does not capture the depth of the dippers or the linear-suppression. The shallow dippers are not an inherent property of the model, but occur because it also has to simultaneously capture the functions produced with the fixed contrast masks. If F3 was fit only to the data produced with no fixed contrast mask, its fits would be the same as those of F2. F3's failure to capture the linear-suppression is an unavoidable failure of the model. Foley's (1994a) data did not expose this flaw because he only measured the masking produced by a one-component remote mask. F3 is the only model described in this chapter that doesn't capture the linear-suppression. Where there is a difference between the grating and plaid masking functions (for TSM), it is in the opposite direction to that predicted by F3.

Though F3 fits the data poorly, it predicts the crossover particularly well. There was a similar crossover in Foley's (1994a) data and this was the main reason why F3 fit his data better than F2.

F2

Instead of raising the inputs of test and masks on the denominator to the exponent q before summation, like F3, F2 raises them to the exponent after summation. As shown earlier, F2 can be expressed as:

$$r_t = c_t^p \,/\, \{z^q + (\Sigma[w.\gamma_i] + c_t)^q\}. \qquad (5.2)$$

Fig. 5.5 shows a schematic illustration of the model. Its fits, which appear in the top-left panels of Figs. 5.3 and 5.4, are better than those of F3. It is obvious that the main reason for this is that F2 captures the linear-suppression whilst F3 doesn't. (This was the reason why F2 fit the contrast detection data in chapter 4 better than F3.) F2 also does much better at capturing the depth of the dippers. However, F2 does less well than F3 at capturing the crossover. A model is required that predicts all three of these data characteristics.

Need for an alternative to Foley's (1994a) models

F2 described the linear-suppression shown with the contrast detection data in chapter 4; whilst F3 produced the best fits to Foley's (1994a) contrast discrimination data. Neither of Foley's models could produce good fits to both data sets. The data shown in this chapter confirms the need for a hybrid model. This new model has characteristics of both F2 and F3.

The hybrid model

The hybrid model is a natural progression from F2 and F3:

(i) Like F2, the inputs of the remote mask components to the divisive inhibitory pool are summed before they are raised to q. This feature allowed F2 to capture the linear-suppression that was evident both in chapter 4's contrast detection data and in the contrast discrimination data shown here.

(ii) Like F3, the self-inhibitory component in the divisive inhibitory pool is raised to q before being summed with the other components in the pool. This feature allowed F3 to capture the pronounced dip in the crossover that was evident both in Foley's (1994a) contrast discrimination data, and in the contrast discrimination data from DJH shown here.

The hybrid model can be expressed as:

$$r_t = c_t^p \,/\, \{z^q + (\Sigma[w.\gamma_i])^q + c_t^q\}. \qquad (5.3)$$

Its fits are shown in the middle-left panels of Figs. 5.3 and 5.4. Fig. 5.7 shows a schematic illustration of the model. So, are better fits produced by the hybrid model than by either F2 or F3? The hybrid model captures the linear-suppression like F2, and it also captures the pronounced dip in the crossover like F3. From the three models fit so far, this model produces the best fit to DJH with no extra parameters. For TSM, the fit was a little worse than that of F2. The main problem with the fit to TSM is that the model fails to capture the depth of the dipper. So, the hybrid model captured the data from DJH better than F2 or F3, but it didn't help matters for TSM.

The data from DJH showed both linear-suppression and a pronounced crossover. This was why the hybrid model produced the best fits to these. The data from TSM showed only linear-suppression. This was why F2 produced the best fits to these. The difference between the hybrid model and F2 is the treatment of self-inhibition. In the hybrid model, the remote mask components are not combined with self-inhibition before they are raised to q. In F2, they are. So the question is, should the input of self inhibition be added to the remote mask inputs before or after they are all raised to q. DJH's data suggests that they should, whilst TSM's suggests that they shouldn't.

The compound model

From the above, it appears that what is required is a model that can incorporate the characteristics of both F2 (to fit TSM's data) and the hybrid model (to fit DJH's data). The models described in the following two subsections (the compound-plus model and the alternative compound-plus model) incorporate a soft-switch, which allows them to switch between characteristics of the two models. The compound model, described for the first time here, should be considered as an interim model between the hybrid and compound-plus model. This is because, though it doesn't have a switch and so is unable to flip between F2 and the hybrid model, it does have inhibitory features of the two models. The compound model can be expressed as:

$$r_t = c_t^p \,/\, \{z^q + (\Sigma[w.\gamma_i] + c_t)^q + c_t^q\}. \tag{5.4}$$

Fits of this appear in the middle-right panels of Figs. 5.3 and 5.4, whilst a schematic illustration of it is shown in Fig. 5.7.

The self-inhibitory component appears twice in the compound model. Like F2, it is combined with the inputs from the remote mask components before they are raised to q. Like F3 and the hybrid model however, it is also raised to an exponent before it is combined with the other components of the divisive inhibitory pool. This model fits DJH's data better than either F2 or F3, and TSM's better than F3. Fits of this model and F2 to TSM's data are about the same.

The compound-plus model

This model incorporates the divisive-inhibitory pools of F2 and the hybrid model. Both this and the next model have an extra parameter to the previous models. It is termed β and works as a soft switch between the pools of F2 and the hybrid model. The model can be expressed as:

$$\begin{aligned} r_t = c_t^p \,/\, \{z^q + \beta((\Sigma[w.\gamma_i] + c_t)^q) \\ + (1-\beta)((\Sigma[w.\gamma_i])^q + c_t^q)\}. \end{aligned} \tag{5.5}$$

Fig. 5.9 shows a schematic illustration of this. Its fits are shown in the bottom-left panels of Figs. 5.3 and 5.4.

The switch, β, is constrained to have values between 0 and 1. With these extreme values, this model is equivalent to F2 (when $\beta = 1$) and the hybrid model (when $\beta = 0$). It was therefore expected that this model would fit the entire data set at least as well as the earlier models. In fact, it offers an improvement. This model captures the depth of the dippers and the linear-suppression like F2, and the pronounced crossover like the hybrid model.

Physiological interpretations of the models

The models can be interpreted in terms of plausible physiological architecture that they could be describing. A feature common to them all is that the excitation of a channel narrowly tuned to the test is divisively inhibited by a broadly tuned pool which receives contributions from all of the stimulus components. The models are distinct in how inputs are handled by their pools. In F3, they are individually processed by narrowly tuned channels before being summed, whilst in F2 they are summed within the same isotropic channel. With regard to F2, though narrowly tuned channels are prevalent in V1, the pooling is not inconsistent with masking occuring in cortex. This is because a proportion of neurons have been found in this region that have little sensitivity to orientation (i.e. isotropic) or spatial frequency (e.g. Baumgartner et al., 1965; Poggio, 1972; Blasdel & Fitzpatrick, 1984; Sato et al., 1996; Ringach et al., 2002). In the hybrid model's divisive inhibitory pool, a channel narrowly tuned to the test is summed with a channel that processes both remote mask components. Alternatively, the compound model sums a channel narrowly tuned to the test with a channel that processes the test component as well as both remote mask components. These interpretations are illustrated in Figs. 5.5 to 5.8.

Though the compound-plus model is a natural mathematical synergy of F2 and the hybrid model, it is physiologically not as obvious as the others. The most plausible interpretation of it follows. As β approaches 1, there is an increase in the inhibition produced by a channel that is excited by all of the stimulus components. In direct proportion to this, there is a decrease in the inhibition produced by a channel excited only by the remote mask components, and in the self-inhibition produced from the channel narrowly tuned to the test that is monitored for test detection.

The alternative compound-plus model

The alternative compound-plus model was developed to achieve the same goals as the compound-plus model, whilst being physiologically more obvious. It can be expressed as:

$$r_t = c_t^p \,/\, \{z^q + (\Sigma[w.\gamma_i] + \beta.c_t)^q + (1\text{-}\beta)c_t^q\}. \quad (5.6)$$

A schematic illustration is shown in Fig. 5.10. Fits are shown in the bottom-right panels of Figs. 5.3 and 5.4. At the extreme β values, this model and the compound-plus model are equivalent to either F2 (when $\beta = 1$) or the hybrid model (when $\beta = 0$). Physiologically, this model describes two pathways that are weighted differently for the two observers. The weighting (β) determines the proportion of self-inhibition produced by the channel narrowly tuned to the test component that is monitored for test detection, relative to the self-inhibition produced by a broadly tuned channel that is excited by all of the stimulus components.

Fits of the alternative compound-plus model are not quite as good as those of the compound-plus model for either observer. They capture the crossover very well, but they don't quite capture the depth of the dippers. However, the fits are better than those of the other four models.

5.4 Chapter Summary and Discussion

Chapter 4 showed that one- (grating) and two- (plaid) component remote masks, matched at a range of stimulus contrasts, had almost exactly the same effect on contrast detection. With reasonable parameter values (see section 4.7), F2 provided a better fit to this linear-suppression than F3. However, Foley (1994a) showed that the effect of a one-component fixed contrast remote mask on contrast discrimination was better fit by F3 than F2. The experiment here was an extension of that in chapter 4 and Foley's (1994a). Effects of both grating and plaid fixed contrast remote masks on contrast discrimination were measured. Foley's (1994a) models and four alternatives, introduced for the first time here, were fit.

The data are inconsistent with a within-channel model of masking and support a cross-channel model of masking. The various model fits mainly differ in how well they capture three features of the data, these being: (i) the linear-suppression, (ii) the pronounced crossover, and (iii) the depth of the dippers. The main findings were that:

(i) F3 did not capture the linear-suppression.
(ii) F2 and the compound model were poor at capturing the crossover.
(iii) F3 and the hybrid model were poor at capturing the depth of the dippers.
(iv) The compound-plus and alternative compound-plus models produced the best fits, capturing all three features well. (These two models had six free parameters however, whilst the others had five).

Regarding Foley's models, it is apparent from the data that F3 should be rejected. F2 provides a good fit to the data of TSM, but not to the data from DJH nor data in Foley (1994a), suggesting it is a poor general model

These results provide new insights into the work of Foley. The overall fits of F3 in Foley (1994a) were superior to those of F2, but only the masking produced by a one-component remote mask was measured. I found that F2 was consistently better than F3 (across experiments and observers) at simultaneously capturing the relative masking produced by single and multiple component remote masks. The main reason for this was that F2 assumes linear summation of components in the divisive inhibitory pool whilst F3 does not. It is only under certain circumstances, such as those investigated by Foley (1994a), that fits of F3 are better than those of F2.

The experiment here was prompted by the need to develop a new model that had the features responsible for the good fits of F2 to chapter 4's data and F3 to Foley's (1994a) data. The inhibitory pool of this 'hybrid model' sums remote mask inputs linearly (like F2), whilst it sums the self-inhibitory input nonlinearly with the rest of the pool's inputs (like F3). This model did well at capturing DJH's data because these showed the linear-suppression and the pronounced crossover. However, TSM's data were captured better by F2 because, though the data showed the linear-suppression, it didn't exhibit the crossover. So, the presence (DJH) and absence (TSM) of the crossover is obviously an important difference between data sets.

When measuring the effect of a remote mask component on contrast discrimination, Foley (1994a) reported the crossover but Ross and Speed (1991) and Ross et al. (1993) did not find a crossover (i.e. their functions merely converged). In this respect, Foley's (1994a) data were similar to DJH's, whilst the data of Ross and colleagues were similar to TSM's. Foley (1994a) acknowledged the difference between his data and those of Ross and colleagues, but he didn't comment on why they might have differed. One reason why psychophysical findings can be inconsistent across labs is that different experimental procedures and stimulus design are used. I have shown the two characters of behaviour using the same paradigm and stimuli, suggesting that observer differences are the cause.

Therefore, a single model was required which could capture the data of both DJH and TSM. The alternative compound-plus (ACP) model incorporates features of both F2 and the hybrid model and was found to be very successful in capturing data with and without the pronounced crossover. To summarise the progression of the modelling:

(i) Foley's (1994a) data favours F3 over F2.
(ii) The data in chapter 4 favours F2 and the hybrid model over F3.
(iii) The data in this chapter favours the ACP model over F2, F3 and the hybrid model.

Crossover is due to the remote masks facilitating contrast discrimination. Why would this occur with DJH but not TSM? The ACP model allows (via β) alternation between two pathways that are weighted differently for the two observers (refer to Fig. 5.10). This weighting determines the proportion of self-inhibition produced by the channel narrowly tuned to the test component that is monitored for test detection, relative to the self-inhibition produced by a broadly tuned channel excited by all of the stimulus components.

Foley 2

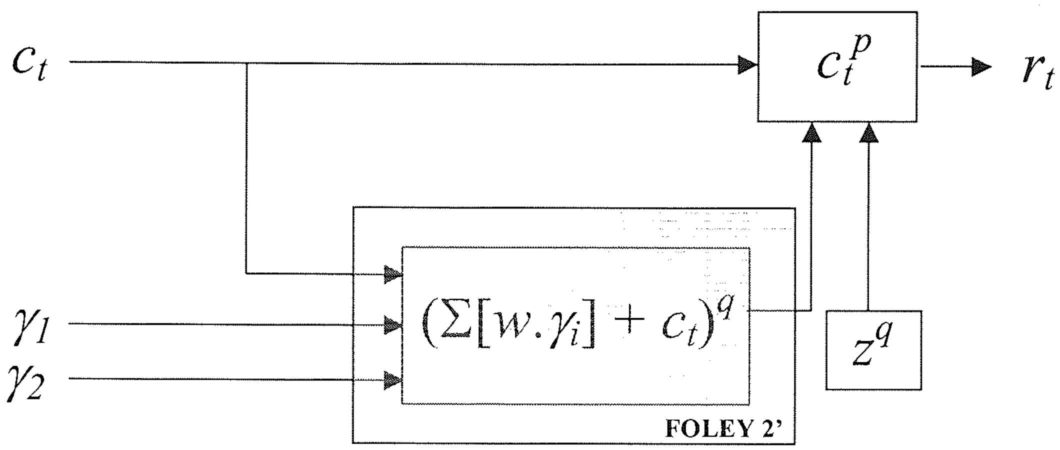

$$\boxed{r_t = c_t^p \,/\, \{z^q + (\Sigma[w.\gamma_i] + c_t)^q\}}$$

Foley 3

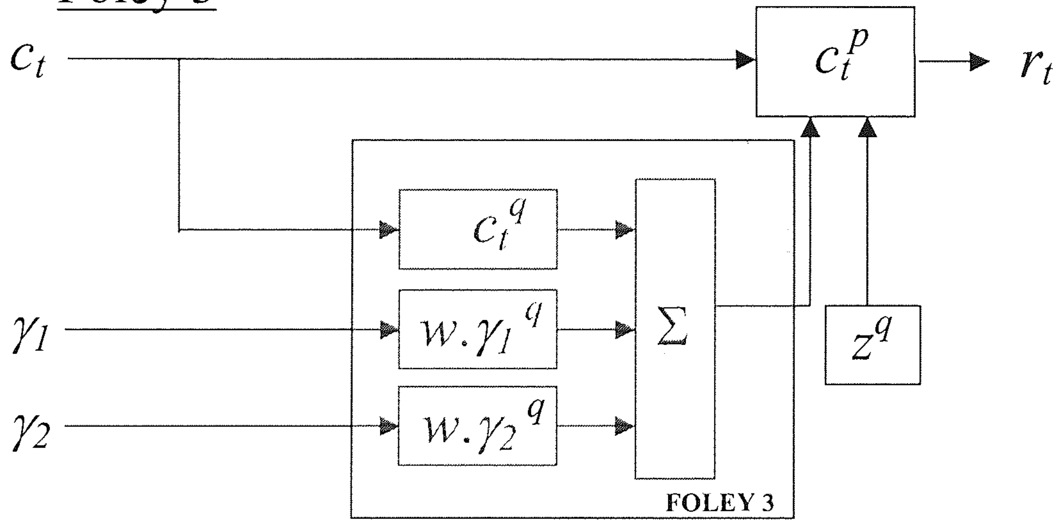

$$\boxed{r_t = c_t^p \,/\, \{z^q + \Sigma[(w.\gamma_i)^q] + c_t^q\}}$$

Fig. 5.5 (top fig.). Schematic illustration of F2. Fig. 5.6 (bottom fig.). Schematic illustration of F3. (Note that mathematical definitions of the models appear below the illustrations).

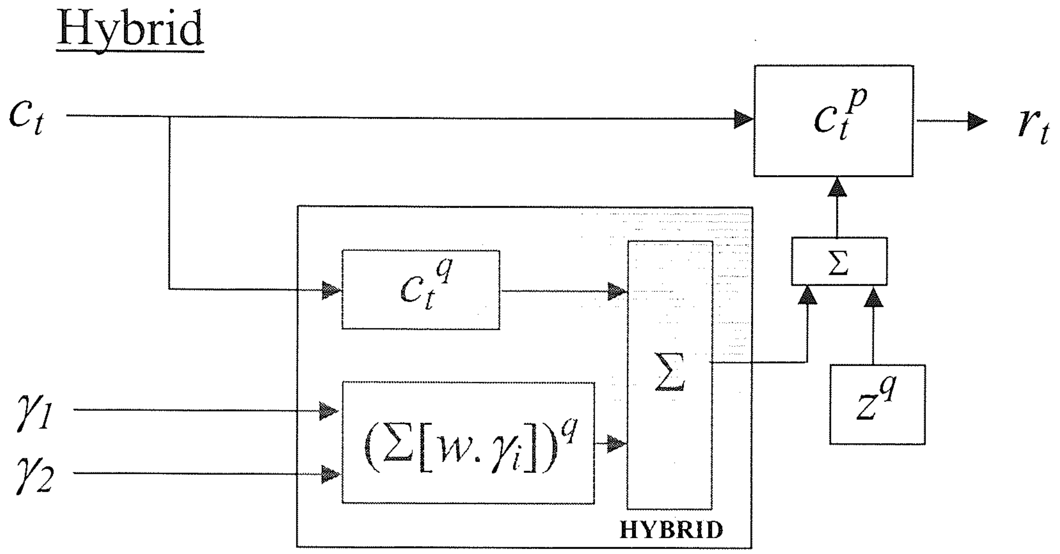

$$\boxed{r_t = c_t^p \;/\; \{z^q + (\Sigma[w.\gamma_i])^q + c_t^q\}}$$

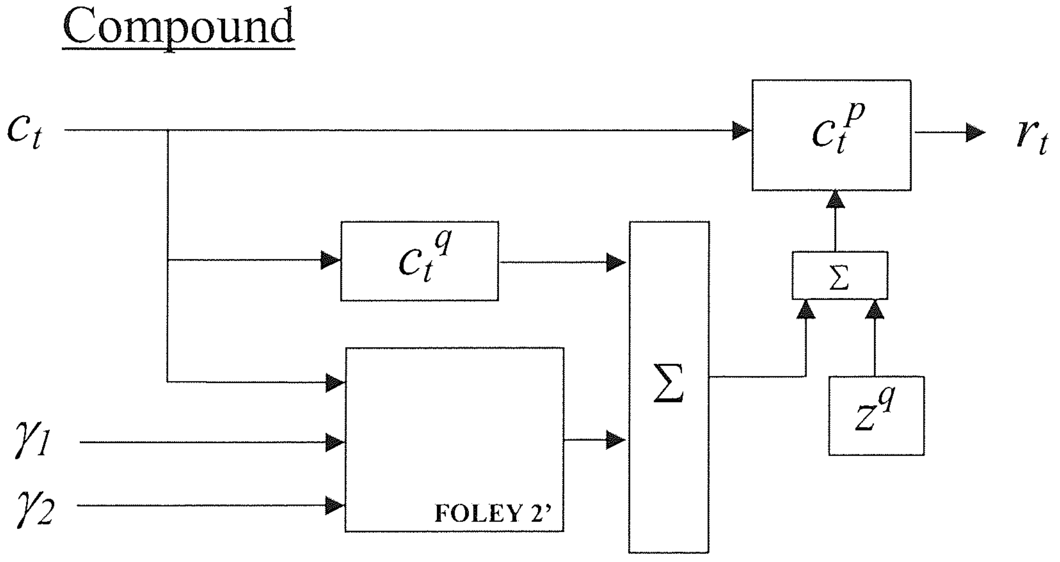

$$\boxed{r_t = c_t^p \;/\; \{z^q + (\Sigma[w.\gamma_i] + c_t)^q + c_t^q\}}$$

Fig. 5.7 (top fig.). Schematic illustration of the hybrid model. Fig. 5.8 (bottom fig.). Schematic illustration of the compound model. (Note that mathematical definitions of the models appear below the illustrations).

Compound-Plus

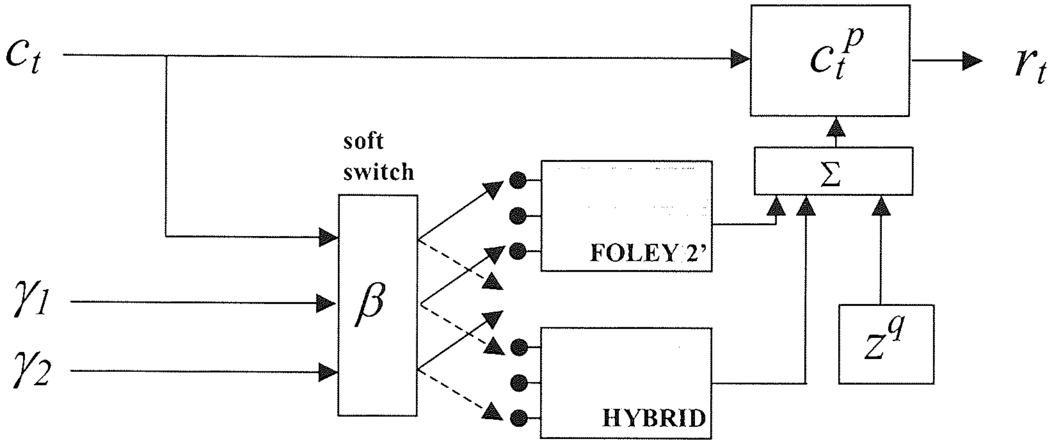

$$r_t = c_t^p \,/\, \{z^q + \beta((\Sigma[w.\gamma_i] + c_t)^q) + (1\text{-}\beta)((\Sigma[w.\gamma_i])^q + c_t^q)\}$$

Alternative Compound-Plus

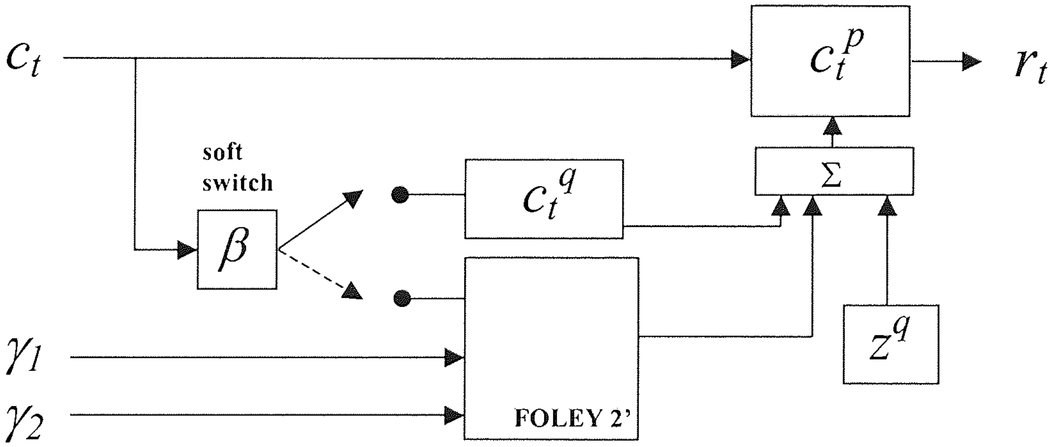

$$r_t = c_t^p \,/\, \{z^q + (\Sigma[w.\gamma_i] + \beta.c_t)^q + (1\text{-}\beta)c_t^q\}$$

Fig. 5.9 (top fig.). Schematic illustration of the compound-plus model. Fig. 5.10 (bottom fig.). Schematic illustration of the alternative compound-plus model. (Note that mathematical definitions of the models appear below the illustrations).

Chapter Six: A Reappraisal of Orientation Masking

6.1 Chapter Introduction

Orientation selectivity is a prominent feature of striate cortex. Single cell recordings (e.g. Hubel & Wiesel, 1962; 1968; 1974; Campbell et al., 1968; Schiller et al., 1976; DeValois, Yund & Hepler, 1982; Jones et al., 1987) have shown cells in this region to have receptive fields with elongated excitatory and inhibitory subregions. Psychophysical evidence for the selectivity has come from studies of: masking (e.g. Sekuler, 1965; Campbell & Kulikowski, 1966; Daugman, 1984; Phillips & Wilson, 1984; Harvey & Doan, 1990), adaptation (e.g. Gilinsky, 1968; Mayo et al., 1968; Sekuler et al., 1968; Blakemore & Campbell, 1969a; Blakemore & Nachmias, 1971; Movshon & Blakemore, 1973; Greenlee & Heitger, 1988; Greenlee et al., 1988; Maattanen & Koenderink, 1991; Snowden, 1991; 1992), subthreshold summation (e.g. Kulikowski et al., 1973; Thomas & Shimamura, 1975; Phillips & Wilson, 1984; Georgeson & Shackleton, 1994; Meese & Williams, 2000), and discrimination at threshold (e.g. Thomas & Gille, 1979; Burr & Wijesundra, 1991). The emphasis in these studies has been on estimating bandwidths of channel selectivity. For example, mask and adapter patterns typically raise the detection threshold of a test pattern by the greatest extent when they share its orientation. Increasing the orientation difference between them tends to reduce threshold elevation. In masking and adaptation studies, threshold elevation functions have been used to estimate bandwidths, which are expressed here as the distance (in degrees) between the mechanism having peak sensitivity, and the mechanism having a sensitivity which is half of this (i.e. its half-width at half-height). (Note that the term 'bandwidth' refers to orientation bandwidth throughout this chapter).

The first attempts of using psychophysics to measure channel bandwidths occurred in the 1960's. Campbell and Kulikowski (1966) considered the width of masking functions to give direct estimates of them and reported that they were around 12 to 15°. Blakemore and Nachmias (1971) and Movshon and Blakemore (1973) used adaptation and the equivalent contrast transform (ECT; a discussion of this is withheld until later) and estimated that bandwidths were around 7 to 8°.

Discrepant estimates were followed by the suggestion that a channel's orientation bandwidth is related to its preferred temporal and spatial frequencies. Channels fall into two main categories of temporal selectivity (e.g. Enroth-Cugell & Robson, 1966; Tulunay-Keesey, 1972; Tolhurst, 1973; Kulikowski & Tolhurst, 1973; Pantle, 1973). They have been distinguished as 'movement' and 'form' analysers (e.g. Sharpe & Tolhurst, 1973), or more commonly, as those that are 'transient' and 'sustained' (e.g. Legge, 1978a). The transient channel is thought to have the greater sensitivity to higher temporal frequencies, but also to lower spatial frequencies (e.g. Kulikowski, 1971; Kulikowski & Tolhurst, 1973; Legge, 1978a; Kelly & Burbeck, 1984; 1987; Hess & Snowden, 1992), indeed, lower spatial frequency gratings can be detected at shorter temporal periods (e.g. Breitmeyer, 1975; Lupp et al., 1976; Vassilev & Mitov, 1976; Harwerth & Levi, 1978). With respect to orientation bandwidths, it has been widely reported that those of transient channels are broader than those of sustained channels (though Phillips & Wilson, 1984, estimated them as being about the same). Single cell recordings by DeValois, Yund and Hepler (1982) have revealed that channels favouring the lower spatial frequencies have the broader orientation bandwidths. Adaptation and discrimination at threshold studies by Snowden (1992) confirmed this and in addition, showed that higher temporal frequency stimuli excite channels with broader orientation bandwidths. Consistent with this, Sharpe and Tolhurst (1973) used adaptation with the ECT in estimating bandwidths at a range of spatial frequencies (1 to 6cpd), and found them to be approximately double when using temporally transient (5Hz flicker- or drift-rate) stimuli, compared to when they used temporally sustained (stationary) stimuli. Further, Anderson et al. (1991) estimated bandwidths of around 15° with a temporally sustained 3cpd pattern and, though bandwidths were not measured at this particular spatial frequency with transient patterns, all of their estimates with these were much broader.

Many studies have probed the effects that suprathreshold gratings have on the detection threshold of test gratings when the two patterns are orthogonal. Changing ideas about the theoretical importance of these measurements, that is, developments in understanding the underlying processes that they are revealing, are discussed below. Firstly, some of the main investigations into this are briefly described. Adapter gratings typically have very little, if any, effect on the detectability of orthogonal test gratings (e.g. Gilinsky, 1968; Mayo et al., 1968; Sekuler et al., 1968; Blakemore & Campbell, 1969a; Blakemore & Nachmias, 1971; Movshon & Blakemore, 1973; Greenlee & Heitger, 1988; Greenlee et al., 1988; Maattanen & Koenderink, 1991). However, in line with the above distinctions between channel selectivities, the adaptation study of Kelly and Burbeck (1987) showed a broadening of orientation bandwidths as stimuli were lowered in spatial frequency and raised in temporal frequency, so that at extreme values there appeared to be no orientation tuning. When test and mask are orthogonal, masking has been shown to be either negligible (e.g. Campbell & Kulikowski, 1966; Daugman, 1984; Harvey & Doan, 1990) or substantial (Ross & Speed, 1991; Foley, 1994a; Foley & Chen, 1997). Dependencies of orthogonal masking on stimulus characteristics were investigated by Burbeck and Kelly (1981). Test gratings were superimposed with fixed contrast orthogonal mask gratings that had the same temporal and spatial frequencies. The observer's task was to adjust the contrast of the test to detection threshold. As these author's have shown with adaptation, test thresholds were elevated the most when the stimuli had low spatial frequencies and high temporal frequencies. Unfortunately, interpretation of these results is complicated because a criterion dependent psychophysical method was used. (It was reported that a 2 interval forced choice, 2IFC, procedure delivered the same results, but the data were not shown). A further complicating factor is the prolonged duration of the stimuli (even with 2IFC they were shown for 13secs), which means that adaptation may have contaminated the data. However, the study was extended by Ferrera and Wilson (1985) who, using a single-interval yes/no procedure, also found substantial orthogonal masking to be produced with low spatial frequency, high temporal frequency stimuli. In addition, they found that this masking is broadly tuned to spatial frequency.

The substantial effects from suprathreshold gratings upon the detectability of remotely oriented test gratings, were originally thought to be due to the operation of channels very broadly tuned, or isotropic, to orientation. This view was based on within-channel masking models in which masking could only be produced when test and mask patterns excited the same channel. However, many of the experiments presented in previous chapters of this thesis and elsewhere (e.g. particularly see chapter 1 for discussions) have shown that such remote masking can be accurately predicted with cross-channel masking models. In these, masking can occur when the test and mask patterns excite different channels. The aim of the work in this chapter is to investigate whether channel bandwidths have been misestimated in the past because interactions between channels were not considered.

Firstly, an example from the adaptation literature is discussed, where the failure to acknowledge channel interactions is thought to have led to the misestimation of channel bandwidths from such functions. It will then be argued that a further oversight of channel interactions has led to orthogonal masking produced by cross-channel suppression being misinterpreted as indicative of channels very broadly tuned to orientation.

The 'equivalent contrast transformation' (ECT) (e.g. Blakemore & Nachmias, 1971; Movshon & Blakemore, 1973; Sharpe & Tolhurst, 1973) was devised to overcome a problem arising from different contrasts of adapter patterns being used across studies. Of course, an adapter's effect is reduced by increasing the orientation difference between the test and adapter patterns; importantly, however, it is also lowered by reducing its contrast. ECT has been used in attempting to collate data from the different studies by describing one of these effects in terms of the other. For example, the reduction in adaptation produced by, say, halving adapter contrast, may be represented in terms of how much the orientation difference between the test and adapter patterns would need to be increased by to produce the same effect. It was thought that this would allow the ECT to deliver the same orientation bandwidths regardless of adapter contrast. However, a requirement of this is that the functions relating adapter contrasts to adaptation levels maintain a constant slope, regardless of the orientation difference between test and adapter patterns. Snowden (1991) found that this assumption was incorrect; the slopes reduced with an increasing orientation difference. This change is consistent with a large body of evidence (e.g. see chapter 1, section 1.8) that shows inhibition between channels to operate in optimising contrast discrimination around the prevalent contrast of a recently viewed or present scene. An example of how adaptation influences this operation is that an adapter which does not excite a neuron excited by a test pattern can produce a horizontal shift of its contrast response function (e.g. Ohzawa et al., 1982).

There has been ample evidence provided by physiologists showing that an adapter pattern can affect the mechanism used for test detection without providing it with excitation (e.g. Ohzawa et al., 1982; 1985; Albrecht et al., 1984; Saul & Cynader, 1989a; 1989b; Bonds, 1991). However, the large number of adaptation studies in which bandwidths have been estimated have ignored this; they assumed that if an adapter affected test detection then it must have been providing excitation to the test detecting mechanism. This oversight, due to the majority of the studies being performed before much was known about channel interactions, would have lead to some misestimates of orientation bandwidth.

The notion of channel interactions has been essential to understanding most of the masking data in the previous three chapters. I, and others (e.g. Ross et al., 1993; Foley, 1994a), have provided strong evidence that masking produced with remotely oriented stimuli is not due to isotropic test detecting mechanisms as supposed by some (e.g. Burbeck & Kelly, 1981; Ferrera & Wilson, 1985). The effect of a remotely oriented mask on contrast discrimination thresholds (chapter 5) is a particularly convincing demonstration that masking with remotely oriented stimuli is a result of divisive inhibitory interactions between channels narrowly tuned to orientation. This has pertinence to bandwidth estimates in orientation masking studies because they, like the adaptation studies mentioned above, assumed that threshold elevation could only occur when the stimuli were exciting the same mechanism.

This chapter is a development of the previous three. As already stated, its aim is to investigate whether cross-channel effects may have contaminated bandwidth estimates produced in studies which only assumed within-channel effects. Further, it will be addressed why orthogonal masking has been found in some studies but not in others. It was previously discussed how estimates of channel bandwidth have varied substantially with stimulus spatial and temporal frequencies. The role of cross-channel suppression in these variations is investigated here. Emphasis is given to Phillips and Wilson's (1984) particularly influential study (which has been cited well over a hundred times), which has few of the weaknesses of other similar studies (such as Burbeck & Kelly's, 1981). Because of the emphasis on this study throughout this chapter, it is described in detail below.

Phillips and Wilson (1984) used a single-interval yes/no masking procedure in which full-field mask gratings were presented at 40% contrast. The tests were spatially localised patterns which had a spatial frequency bandwidth of around 1 octave at full-width, half-height. Test and mask were matched at a range of spatial frequencies (between 0.5 and 12cpd) and there were two temporal conditions. In the 'transient' condition, the mask was an 8Hz sine wave that was presented for 1 second and curtailed with 0.25 second linear ramps, whilst the test was one cycle of an 8Hz square wave. In the 'sustained' condition, the mask was windowed with one cycle of a 1Hz sine wave and the test was windowed by a Gaussian with a duration of one second with full-width at half-height of 0.4secs. It was found that far less masking was produced with sustained stimuli than with transient stimuli, and with high spatial frequency stimuli than with low spatial frequency stimuli (consistent with others; e.g. Burbeck & Kelly, 1981). Estimates of bandwidths were about the same with transient and sustained stimuli at all spatial frequencies investigated; they were around 30° at 0.5cpd, and decreased to around 15° at 10cpd. These were broadly in agreement with estimates by Anderson et al. (1991); using transient stimuli (which drifted at 8Hz) they found bandwidths of around 50° at 0.1cpd that decreased to around 30° at 10cpd.

An experiment similar to Phillips and Wilson's (1984) is shown here. Orientation masking was measured using stimuli at a range of spatial frequencies, which were designed to predominantly stimulate either transient or sustained mechanisms (e.g. see Wilson, 1980b; Moulden et al. 1984; Phillips & Wilson, 1984; Anderson & Burr, 1985; Georgeson, 1987; Hess & Snowden, 1992; Fredericksen & Hess, 1998). This study differs in one important respect, however, from similar studies that have manipulated stimuli in such a way. Whilst Phillips and Wilson (1984) measured orientation masking between 0 and 45deg, and Burbeck and Kelly (1981) only measured orthogonal masking, here it has been measured between 0 and 90deg. It was suspected that this larger range might reveal the contribution of cross-channel effects to the masking levels.

6.2 General Methods

Tests and masks were gratings, matched in spatial frequency, that were spatially enveloped with a Gaussian and a raised-cosine function, respectively. At full-width, half-height, test stimuli had 1.67 cycles and masks had a diameter of 4.4°. The test patterns were always oriented vertically. Temporally transient stimuli were used in all experiments; whilst in the main experiment (section 6.3), sustained stimuli were also used. The transient stimuli were one cycle of a 15Hz square wave whilst the sustained stimuli were linearly ramped on and off, had a total duration of 1sec and had a plateau of 0.5sec. Patterns were in sine phase with the fixation point, except for some conditions of the experiment in section 6.5. For those, the test was in sine phase whilst the mask was in quadrature phase relative to the test (sine phase + 90deg).

6.3 Masking as a function of mask orientation

Methods

Stimuli had 1, 3 or 9cpd. Masks with 40% contrast were oriented between vertical and horizontal. Spatial and temporal frequencies were varied across sessions. Mask orientation was interleaved within sessions and (when all mask orientations were used) across sessions. The latter achieved a reasonable session duration, which was about 30 to 40 minutes when sustained stimuli were used. (In one session, masks were oriented 0, -20, -40, -60 and -80 deg, whilst in another session, they were oriented -10, -30, -50, -70 and -90 deg). For DJH, masks were interleaved across trials. For TSM, trials with masks oriented the same were blocked together and the sequence of the blocks was randomised. (The reasons for this difference are discussed in section 6.4.) Unmasked thresholds were measured in each session.

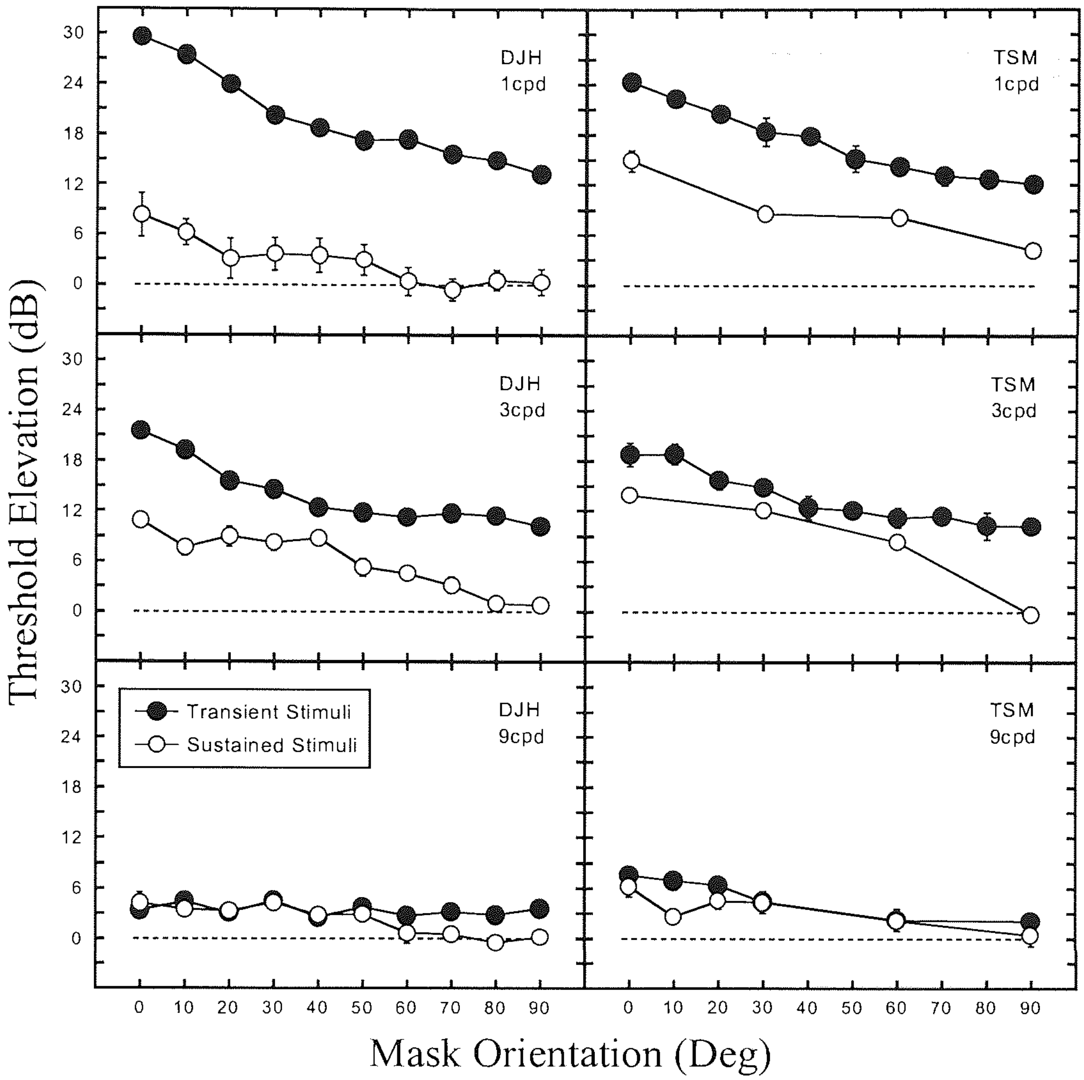

Fig. 6.1. Masking as a function of mask orientation. Stimuli had 1cpd (top panels), 3cpd (middle panels) or 9cpd (bottom panels), and were temporally transient (filled symbols) or sustained (open symbols). Tests were vertical whilst mask orientation was varied between 0 and -90deg. Data are from DJH (left panels) and TSM (right panels).

SF (cpd)	1		3		9	
	T	S	T	S	T	S
DJH	1.40 (0.25)	3.57 (0.79)	8.52 (0.25)	4.73 (0.57)	25.87 (0.18)	15.25 (0.74)
TSM	0.08 (0.28)	-1.05 (1.23)	5.74 (0.55)	0.79 (0.37)	22.28 (0.88)	15.49 (0.60)

Table 6.1. Unmasked thresholds (dB) for each test condition in Fig. 6.1 (above). Stimuli had 1, 3 or 9cpd and were temporally transient (T) or sustained (S). Data are from DJH and TSM. Standard errors are in brackets.

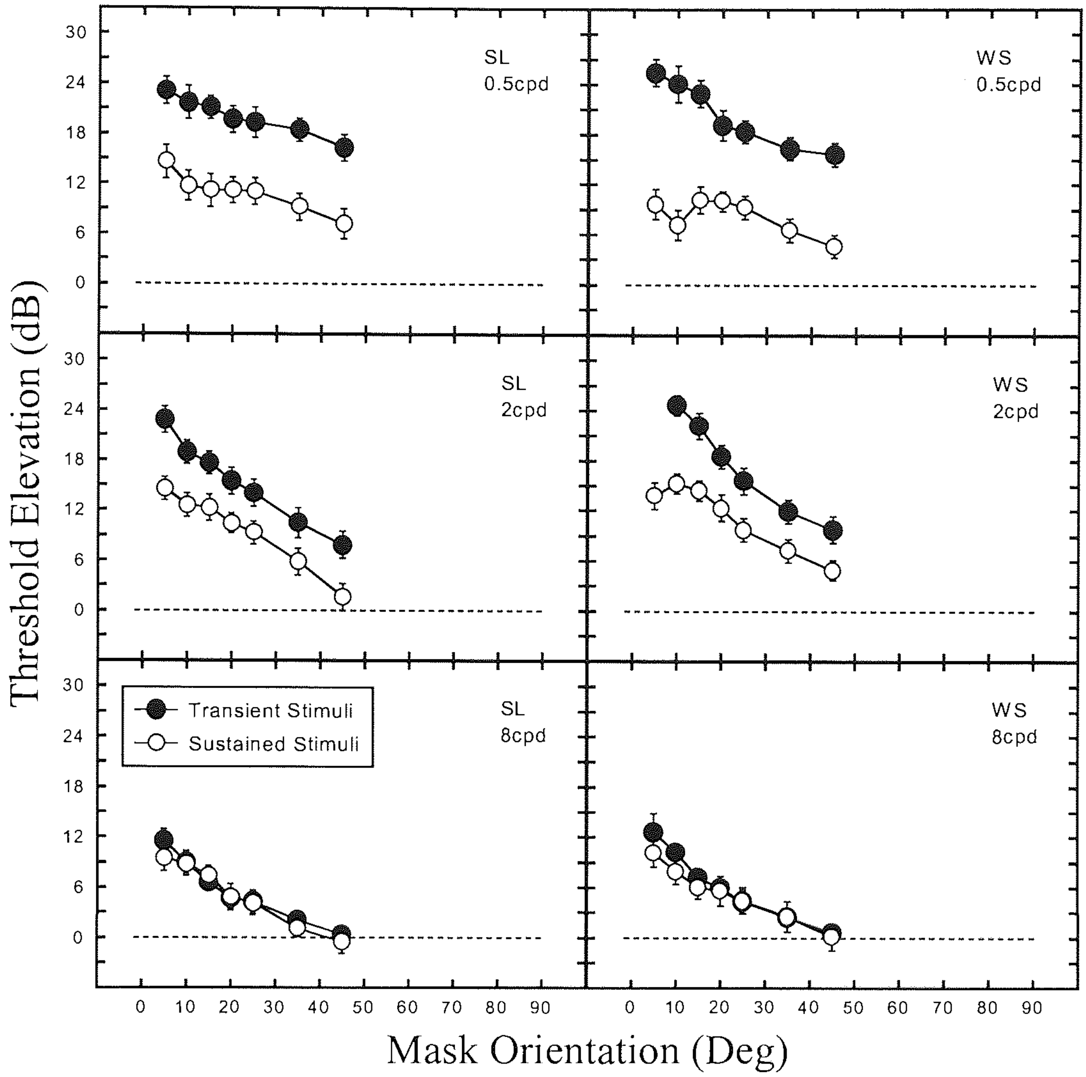

Fig. 6.2. Data are from Phillips and Wilson (1984) and show masking as a function of mask orientation. Stimuli had 0.5cpd (top panels), 2cpd (middle panels) or 8cpd (bottom panels), and were temporally transient (filled symbols) or sustained (open symbols). Tests were vertical whilst mask orientation was varied between 0 and -45deg. Data are from SL (left panels) and WS (right panels).

Results and Discussion

Figs. 6.1 and 6.2 show masking as a function of mask orientation for temporally transient (filled symbols) and sustained (open symbols) stimuli, which had 1 (top panels), 3 (middle panels) or 9cpd (bottom panels). Data are from the present experiment (Fig. 6.1) and Phillips and Wilson (1984) (replotted in Fig. 6.2). Note that 'P&W' is used in reference to Phillips and Wilson's (1984) study. In P&W's data, the influence of cross-orientation suppression could not be judged. However, it was found here that remote masks could elevate test thresholds by up to about four times, which supports the view that substantial cross-orientation suppression contributed to the results of P&W. The contributions would have been greater in some of P&W's stimulus conditions than in others; a great deal of remote masking was produced with transient, low spatial frequency stimuli (3cpd or less), whilst very little (if any) was produced with sustained, high spatial frequency stimuli. This is consistent with differences between the stimuli used in studies that supported the existence of orthogonal masking, and those that supported its absence (particularly see Burbeck and Kelly, 1981). The functions showing substantial remote masking declined quite sharply up to around 30deg, and then declined more gradually. This seems to be indicating two distinct processes underlying the masking functions, with the width of the sharp decline revealing the detecting mechanism's bandwidth, and the slighter decline revealing cross-orientation suppression.

Assuming the cross-orientation suppression elevates entire masking functions, it would have lead to P&W overestimating orientation bandwidths. To explain, P&W's masking functions were due to cross-channel suppression as well as a within-channel effect, but they thought that it was all within-channel masking. Consider the masking function in the top-left panel of Fig. 6.1 produced with transient stimuli (filled symbols). It has been described above how the masking produced when test and mask patterns differed by between 30 and 90deg will mostly be due to cross-channel effects. In this region, the function is only slightly declining and a comparable level of the effect presumably underlies the masking produced when test and mask differ by between 0 and 30deg. The cross-channel effect in this region cannot be identified because the masking there is due to both cross-channel suppression and a within-channel effect. The within-channel effect, which is apparent from the sharper slope in this region, obviously produces far less threshold elevation over a far smaller range of stimulus orientation differences than the masking function (which P&W assumed was all due to a within-channel effect). Though P&W only used data produced with stimuli differing in orientation by between 0 and only 45deg (so an attempt was not made to fit the substantial remote masking), this would obviously have led them to estimate the bandwidths (measured at half-width, half-height) to be wider than if the contribution of cross-channel effects had been recognised. In ongoing collaborative work (e.g. Meese & Holmes, 2003), the functions exhibiting substantial remote masking (in Figs. 6.1 and 6.2) have been fit with Foley's (1994a) suppression models. As expected, these delivered bandwidth estimates that were substantially narrower than P&W's. For example, at 0.5cpd they were about 17deg, whilst P&W's were about 30deg at low spatial frequencies. P&W's bandwidth estimates for transient and sustained channels were equivalent, so acknowledging cross-orientation suppression leads to estimates for (low and medium spatial frequency) transient channels being substantially narrower than those for the sustained channels.

In the present study, mask orientation was varied in only one direction from the test. Others have assumed that the function produced by a pattern oriented in the opposite direction would be the same (e.g. Snowden, 1991; 1992), though asymmetry has been found in the orientation selectivity of cat visual neurons.

Hammond and Andrews (1978), for example, reported that 60% of cells in cat area 17 showed tuning asymmetries in excess of 20% (also see Henry et al., 1974; Rose & Blakemore, 1974). However, if there are as many left bias cells as right bias cells (which seems likely) then this would be expected to have little, if any, impact on psychophysical estimates. In addition, a vertical test was used in the present study whilst both masking (Campbell & Kulikowski, 1966) and adaptation (Gilinsky & Mayo, 1971) studies have shown bandwidths from such stimuli to be slightly narrower than those gathered using oblique tests. For example, Campbell and Kulikowski (1966) found the bandwidth of masking functions to be 12° with vertical test patterns and 15° with oblique test patterns. Neither of these discrepancies detract from arguments made in the present study, however.

6.4 Orthogonal masks do not produce dippers

Introduction

It was shown in the last section that substantial cross-orientation suppression is produced by transient stimuli with a spatial frequency of 3cpd or less. The successes of cross-channel models emphasise the plausibility that this is cross-channel suppression. Alternatively, a within-channel approach would suggest that the remote masking is indicative of isotropic channels (e.g. Burbeck & Kelly, 1981; Ferrera & Wilson, 1985). A telltale sign of test and mask patterns exciting the same channel is facilitation of test detection with low mask contrasts (see chapter 1). The presence of a dipper with the orthogonal masking would be consistent with a within-channel explanation, whilst its absence would be consistent with a cross-channel explanation.

Others have already found that parallel stimuli produce dippers whilst orthogonal stimuli do not (e.g. Foley, 1994a; Foley & Chen, 1997). The functions were measured again because the stimuli used in those studies differed in their spatio-temporal characteristics to the stimuli used here.

Methods

Masking was measured as a function of mask contrast with test and mask components that were (i) parallel or (ii) orthogonal, for the spatio-temporal condition in which the greatest cross-orientation suppression was found (stimuli were transient and had 1cpd). Mask orientation was varied across sessions. Trials in which masks at each individual contrast were shown were blocked together and their sequence was varied within each session. Data are from DJH and TSM.

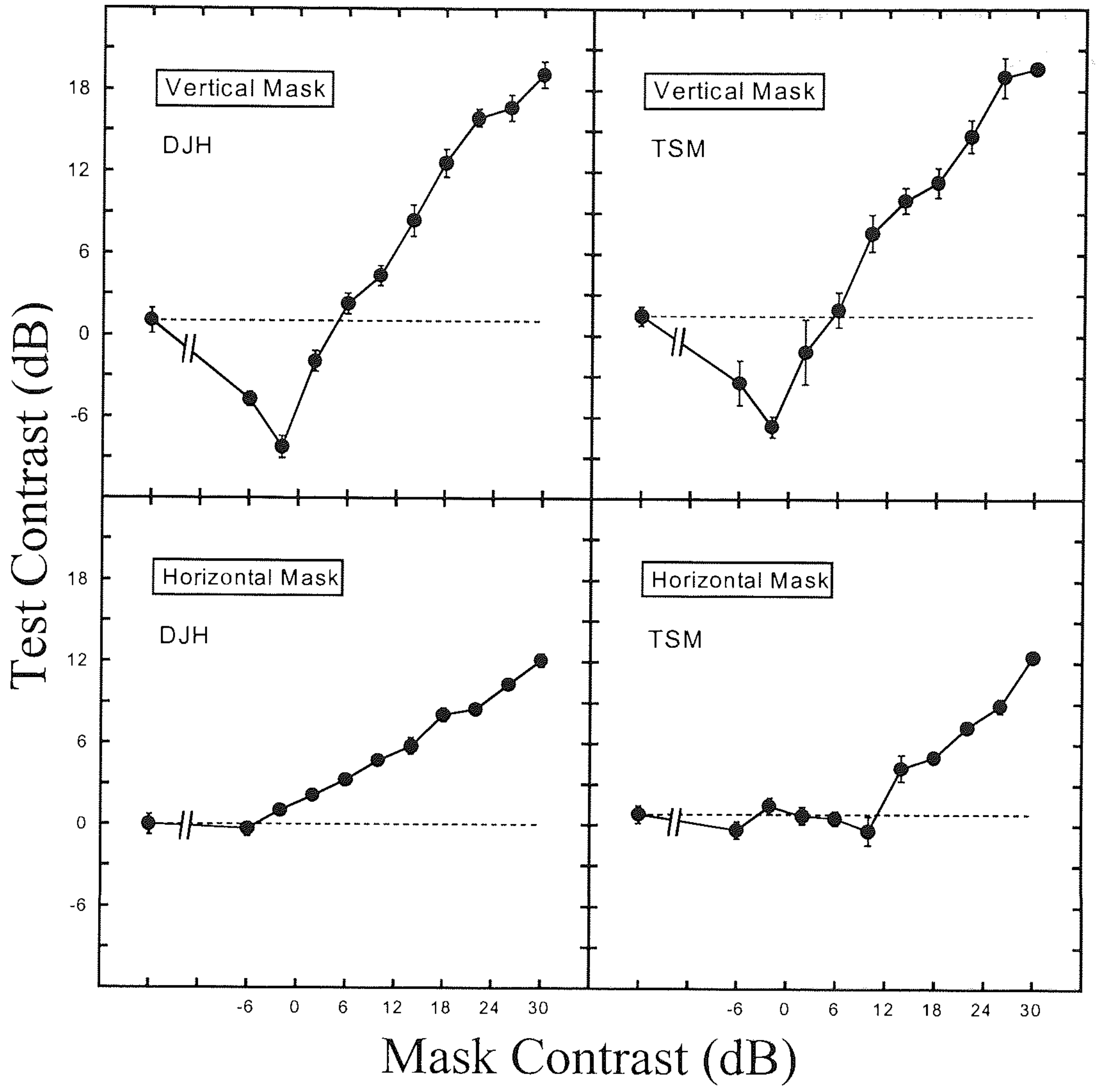

Fig. 6.3. Masking as a function of mask contrast. Stimuli were transient and had 1cpd. Tests were vertical whilst masks were either vertical (top panels) or horizontal (bottom panels). The dashed lines represent the unmasked test detection thresholds. Data are from DJH (left panels) and TSM (right panels).

Results and Discussion

Fig. 6.3 shows masking as a function of mask contrast. A prominent dipper was produced with parallel stimuli (top panels), but it was absent with orthogonal stimuli (bottom panels). This is consistent with the cross-orientation suppression being cross-channel suppression, and the pedestal masking being due to within-channel effects.

It is curious that the parallel mask with 32% contrast produced around 20dB of masking for DJH (top-left panel). In the previous experiment where the same stimuli were used, masks with 40% contrast produced almost 30dB of masking (filled symbols in top-left panel of Fig. 6.1). This substantial difference might be due to differences in method. Here, mask orientation was varied across sessions for DJH, whilst in the previous experiment it was varied within sessions and across trials, which meant that the mask orientation of each trial typically differed to that of the preceding trial. This could have introduced uncertainty which is consistent with elevated masking levels. For TSM, trials with masks oriented the same were blocked together both in this and in the previous experiment. For that observer, the difference in thresholds was not observed; masking was similar for the high contrast mask conditions in both experiments.

6.5 Orientation masking is phase-insensitive

Introduction and Methods

Substantial cross-orientation masking was found using transient stimuli with 3cpd or less (top and middle panels of Fig. 6.1). The aim here was to see whether this suppression is reliant on the relative spatial phase of the test and mask patterns. There are two main reasons why this is pertinent to the present study. The first concerns 'spatial-probability summation'; the theory typically used to explain why increasing the size of a test pattern improves its detectability. An observer's visual space is tiled with spatial filters. It is considered that because larger stimuli excite more of the filters selective to different spatial positions, the probability that one of them will detect the pattern is increased (e.g. Sachs et al., 1971; Quick, 1974; Graham et al., 1978; Legge, 1978a,b; Robson & Graham, 1981; Williams & Wilson, 1983; also see Graham, 1989, for discussions). Regarding the masking, when the test and mask patterns have the same orientation, their phase relationship with these filters is the same. However, changing the mask's orientation alters this; there is less phase alignment when the test and mask have different orientations, compared to when they have the same orientation. If the orientation masking is phase-dependent, then any model used to estimate bandwidths of orientation selective channels from the masking functions would need to account for spatial probability summation. The second reason for probing the phase-specificity of the orientation masking concerns the motivation to identify within-channel affects from cross-channel affects. If, for example, within-channel masking was phase dependent whilst cross-channel masking wasn't, then comparing functions produced with in-phase and out-of-phase stimuli may allow the two to be separated. This would help in estimating bandwidths of oriented channels from masking data.

The transient conditions of the experiment in section 6.3 were repeated with 1cpd and 3cpd stimuli. However, for this experiment test and mask were in quadrature phase; the test was, as before, in sine phase relative to the fixation point, but the mask was in quadrature phase relative to the test (that is, sine phase + 90deg, or anti-cosine phase, relative to the fixation point). The fixation point and the mask being much larger than the test provided the observer with positional information.

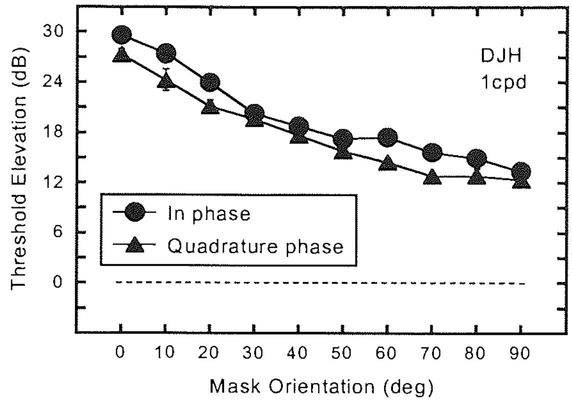

Fig. 6.4. (Above) Masking as a function of mask orientation. Stimuli were temporally transient and had 1cpd. The mask was either in-phase (filled symbols) or in quadrature-phase (open symbols) with the test. Data are from DJH.

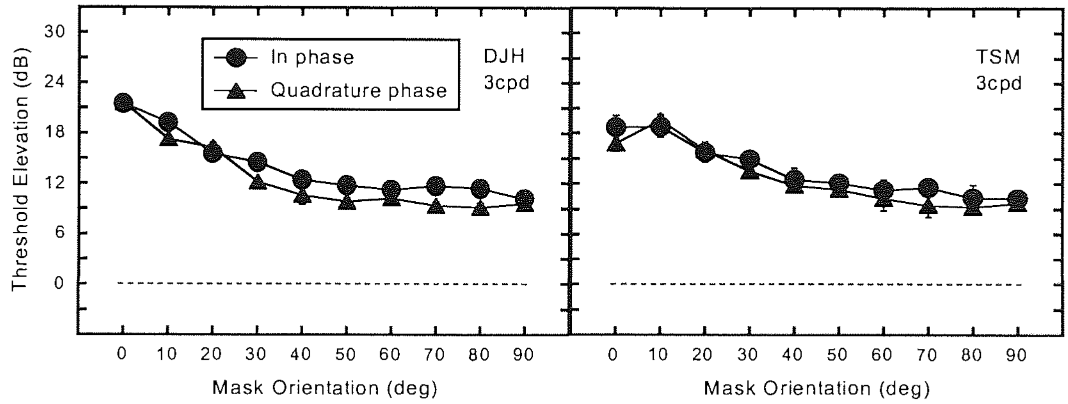

Fig. 6.5. (Above) Masking as a function of mask orientation. Stimuli were temporally transient and had 3cpd. The mask was either in-phase (filled symbols) or in quadrature-phase (open symbols) with the test. Data are from DJH (left panel) and TSM (right panel).

Results and Discussion

Orientation masking functions for 1cpd and 3cpd transient stimuli are shown in Figs. 6.4 and 6.5, respectively. The test and mask were either in-phase (circular symbols) or were in quadrature-phase (triangular symbols). The in-phase data are taken from Fig. 6.1. It is clear that the masking levels were affected very little by stimulus phase. So, when estimating orientation tuning bandwidths, it appears acceptable to ignore spatial probability summation in the models that are used. There have been other reports of masking being insensitive to phase (for a review, see Foley & Chen, 1999). For example, Lawton and Tyler (1994) found this using stimuli matched in orientation and spatial frequency (at a range of spatial and temporal frequencies).

It is a little curious why slightly more masking was consistently produced with DJH when the stimuli were in-phase. He gathered data with these stimuli a period of time before he gathered data with the stimuli that were in quadrature-phase. Practice effects may have slightly lowered the levels of masking that were produced in the later condition. Consistent with this is that by interleaving sessions from the two conditions, TSM produced functions that were more alike.

6.6 Chapter Summary

P&W's reliance on a within-channel masking model would have resulted in them over-estimating the bandwidths of oriented channels. The substantial masking produced by patterns remote in orientation from the test had telltale characteristics of being a cross-channel phenomenon. The insensitivity of the orientation masking to stimulus phase is consistent with it being unnecessary for models used in estimating orientation bandwidth to account for spatial probability summation. It has also been addressed why some previous studies have found orthogonal masking whilst others have not. Consistent with less thorough studies, orthogonal masking was dependent on the stimuli having low spatial frequency and high temporal frequency.

6.7 Chapter Discussion

Why might cross-channel masking be produced with some stimuli and not with others? There is a great deal of evidence supporting the existence of two major anatomical pathways that process visual information. The magnocellular pathway conveys contrast information from stimuli that have high temporal frequencies and low spatial frequencies, and is relatively insensitive to stimulus phase. The parvocellular pathway on the other hand, favours stimuli that have low temporal frequencies and high spatial frequencies, and is phase-sensitive. (Kaplan & Shapley, 1982; 1986; Ungerleider & Mishkin, 1982; Livingstone & Hubel, 1984; Tootell et al., 1988 Hubel & Livingstone, 1990; Merigan & Maunsell, 1990; Schiller et al., 1990a, 1990b; Merigan et al., 1991; also see Milner & Goodale, 1995, for discussions). The magnocellular pathway would favour the stimuli that produced cross-channel suppression, whilst the parvocellular pathway would favour those that didn't.

Chapter Seven:
Pattern Adaptation with Pattern Masking

7.1 Chapter Introduction

One of the most interesting findings in this thesis is the linear-suppression shown in chapters 3 to 5. The clearest examples of it are that the contrast detection (Fig. 4.4) and discrimination (Figs. 5.3 and 5.4) thresholds of 1cpd vertical test gratings, are (substantially) raised to almost exactly the same levels by -45deg 3cpd gratings, as they are by ±45deg 3cpd plaids. Data shown in previous chapters, and by other authors, indicate that processes underlying this occur in (striate) cortex (see chapter 5 for data and discussions). The remote masking was well predicted by cross-channel masking models. These posit that the narrowly tuned channel excited by a test component, is divisively inhibited by the responses of all of the channels excited by the test and mask stimuli. The models of this type that produced the best fits were those in which the inputs of the mask components were linearly summed. There are two obvious physiological interpretations of this. The first is that the mask components are processed by individual channels, narrowly tuned to orientation, whose responses are combined in a way equivalent to linear summation. The vast majority of neurons examined in striate cortex have manifested such selectivity (see chapter 6), which adds credance to this. It is also very well established that adaptation shows the mechanisms involved in contrast detection to have this characteristic (e.g. Blakemore & Campbell, 1969a; Blakemore & Nachmias, 1971; Movshon & Blakemore, 1973). However, an alternative explanation is that the suprathreshold mask components may be processed by the same channel, which would have to be very broadly tuned, or isotropic, with respect to orientation. Though relatively few in number, such neurons have been found in certain areas of striate cortex (for example, in the cytochrome oxidase 'blobs': Livingstone & Hubel, 1984; Horton, 1984). The work in this chapter was an attempt to determine which of these possibilities, if either, underly the linear-suppression.

Adaptation has been found to substantially affect the suprathreshold contrast representation of neurons (e.g. Ohzawa et al., 1982; Albrecht et al., 1984; and see chapter 1, section 1.8) and the perceived contrast of suprathreshold patterns (e.g. Georgeson, 1985). It is plausible that the mechanism providing the input to suppression may be adaptable, and that the efficacy of the masking contrasts that it represents could be manipulated by adaptation. If this is the case, then specific predictions follow. Consider adaptation to a -45deg 3cpd grating, afterwhich the contrast detection threshold of a 1cpd vertical test grating is measured whilst it is being masked by either a -45deg or +45deg 3cpd grating. One possibility is that the adapter will affect the masking produced by the grating that shares its orientation, whilst leaving the masking produced by the orthogonal grating unaltered. This would be consistent with responses of narrowly tuned channels underlying the linear-suppression shown in Fig. 4.4. Alternatively, the adapter may influence the effects of both masks. This would be consistent with the linear-suppression resulting from the response of a very broadly tuned or isotropic channel. The experiment just described is the main experiment here. The aim of this chapter is to use the adaptation-on-masking paradigm in an attempt to uncover processes underlying the linear-suppression.

7.2 Preliminary Experiment. Estimating the amount of forward masking produced with an adaptation paradigm

Introduction

The paradigm adopted for the experiments shown in the next section consists of a long period of adaptation at the start of each session followed by shorter 'top-up' periods between each 2-AFC trial. Such a design attempts to sustain a constant level of desensitisation (e.g. Magnussen & Greenlee, 1986). This experiment addresses the period of time that should exist between the offset of the adapter and the onset of the first interval of each trial. In previous adaptation studies, various durations of this inter-stimulus interval have been used; for example, 100ms (Maattanen & Koenderink, 1991), 200ms (Snowden, 1991), 500ms (Ross & Speed, 1996) and 1 second (Wilson & Humanski, 1993). This duration was considered in designing the procedure used in the following adaptation study. If it were too short, large forward masking effects would be expected from the adapter (Foley & Yang, 1991; Foley & Boynton, 1993), whilst if it were too long, recovery from the desensitisation would be expected to occur. Whether forward masking and adaptation reflect common or distinct mechanisms will not be argued here (though see Foley & Boynton, 1993; Snowden, 1994), but it was considered important that for the following experiments, the effects of forward masking should be reduced to a minimum whilst maintaining substantial adaptation. This is possible because the durations of most forward masking effects are very brief (e.g. Georgeson & Georgeson, 1987; Gorea, 1987; Foley & Boynton, 1993; and see appendix 2, Fig. A.2.1), whilst the time required to recover from adaptation is about the same as the duration of adaptation (Blakemore & Campbell, 1969; Blakemore et al., 1973; Heggelund & Hohmann, 1976; Bodinger, 1978; Rose & Lowe, 1982; Rose & Evans, 1983; Magnussen & Greenlee, 1985).

In this preliminary experiment, the effect of a 100ms forward mask on performance in a 2-AFC test detection task was measured. Varying the duration between the offset of the mask and the onset of the test allows an insight into the temporal distance required between the mask and test patterns for large forward masking effects to be absent.

Methods

Circular tests measuring 4.4° at full-width, half-height were spatially curtailed by a raised cosine envelope. Masks were rectangular patches of grating which had a contrast of either 32% or 99%. They measured 15.2° horizontally by 11.7° vertically and were spatially curtailed by a square-wave. Tests and masks were presented for 100ms and were temporally curtailed by a square-wave. The test-mask combinations investigated were: vertical 1cpd test with -45deg 1cpd mask, and vertical 1cpd test with -45deg 3cpd mask. For each trial, the mask gratings were randomly phase shifted between 0 and 2π radians in the direction of their sinusoidal modulation.

Each trial was initiated with the presentation of the mask, which was followed by an inter-stimulus period (ISI) of mean luminance prior to the trial's first 2-AFC test interval. Test-mask combinations and the duration of the ISI were pseudo-randomly varied across sessions. Mask conditions (either no-mask, 32% contrast mask or 99% contrast mask) were blocked and pseudo-randomly ordered within the sessions.

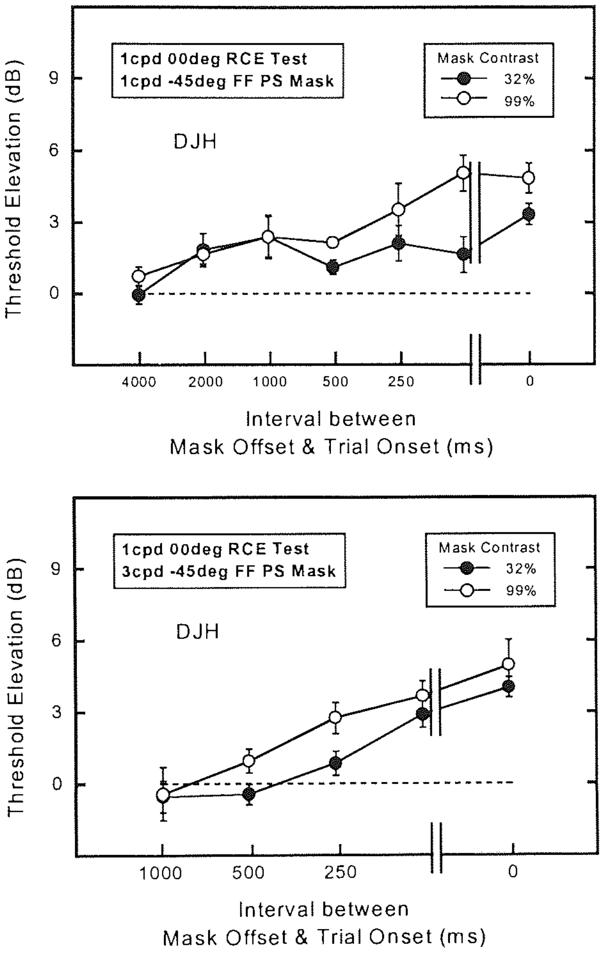

Fig. 7.1. The effect of pre-trial random-phase forward masks on the detection of tests presented in a 2-AFC paradigm. Mask contrast was either 32% (filled symbols) or 99% (open symbols). Tests were 1cpd vertical gratings with raised cosine envelopes (RCE). The full-field, phase-shifted (FFPS) masks were gratings oriented -45deg, which had a spatial frequency of either 1cpd (top panel) or 3cpd (bottom panel). Note the log axes.

Results and Discussion

Fig. 7.1 shows forward masking as a function of the duration (ISI) between mask offset and the onset of the first interval of the 2-AFC trials for two test-mask configurations. Consistent with earlier findings, forward masking increased markedly when the ISI was reduced to within a few hundred milliseconds and typically peaked at the smallest ISI investigated (e.g. Georgeson & Georgeson, 1987; Gorea, 1987; Georgeson, 1988; Foley & Boynton, 1993). Masks had a contrast of either 32% (filled symbols) or 99% (open symbols) and the higher contrast masks consistently produced the greater level of masking. A general decrease in test detectability has been shown previously with increasing contrasts of forward masks (Georgeson, 1988; Foley & Yang, 1991; Foley & Boynton, 1993), simultaneous masks (Legge & Foley, 1980; Wilson et al., 1983; Phillips & Wilson, 1984) and adapters (Stecher et al., 1973; Swift & Smith, 1982; Georgeson & Harris, 1984; Foley & Boynton, 1993). When test and mask spatial frequencies were matched (top panel), masking was produced over a longer ISI than when they differed (bottom panel). When stimulus spatial frequencies were matched and the ISI was 500ms, masking was only about 1½ dB and when they differed it was around 0dB. This ISI was chosen for the following experiments.

An obvious methodological weakness of this experiment is that the forward masks were presented for 100ms. A longer forward mask may produce larger and more sustained effects. However, if longer durations were used, it would have been more difficult to distinguish between forward masking and adaptation effects.

7.3 Attempting to gain insight into linear-suppression using adaptation

A general description and rationale of the main study here was presented in the chapter's introduction. Like most of the experiments performed for this thesis, a 2-AFC simultaneous masking design was implemented and in each trial, 100ms intervals were separated by 500ms. Specific to this study, each session began with a two-minute adaptation period prior to the first trial and, thereafter, 6-second top-up periods of adaptation separated each trial. Such a procedure is commonly used to sustain a high and consistent level of desensitisation. The adaptation periods were temporally separated from the onset of the first interval and the offset of the second interval of each trial by 500ms (see the previous section for a discussion). (Note that the spatial frequencies, orientations, dimensions and enveloping of the test and mask used in gathering the data in the bottom panel of Fig.7.1, were the same as those of the test and adapter, respectively, used here.) To prevent conventional afterimages (local light adaptation), the adapter pattern was spatially phase-shifted every 100ms in the direction of its sinusoidal luminance modulation. The distance of the shift was selected pseudo-randomly and varied between 0 and 2π radians. Tests and masks were circular patches, curtailed with a raised cosine, which measured 4.4° at full-width, half-height. They had a duration of 100ms and were temporally enveloped with a square wave. Adapters measured 15.2° horizontally by 11.7° vertically and had abrupt edges.

A further preliminary experiment was also performed to investigate whether the detection of a test pattern was affected by masking, adaptation, and adaptation-on-masking, using the above methods. It was considered that the probability of achieving effects would be optimised by matching test, mask and adapter patterns in both spatial frequency and orientation. The patterns were -45deg 3cpd gratings whilst adapter contrast was 32% and mask contrast was either 8% or 22.4%. Stimulus combinations were interleaved across sessions. Fig. 7.2 shows that all of the suprathreshold patterns affected the thresholds of the test patterns. Note that iconic representation is used in Figs. 7.2, 7.4 and 7.5. The shading of the bars indicating test contrast reflects the mask pattern that contributed to each level. The level produced with a blank mask is shown as a bar with no shading. The bars showing the levels produced with -45deg and +45deg grating masks are shaded with lines oriented -45deg and +45deg, respectively. Data gathered with different contrasts of grating mask or grating adapter are shown in different panels.

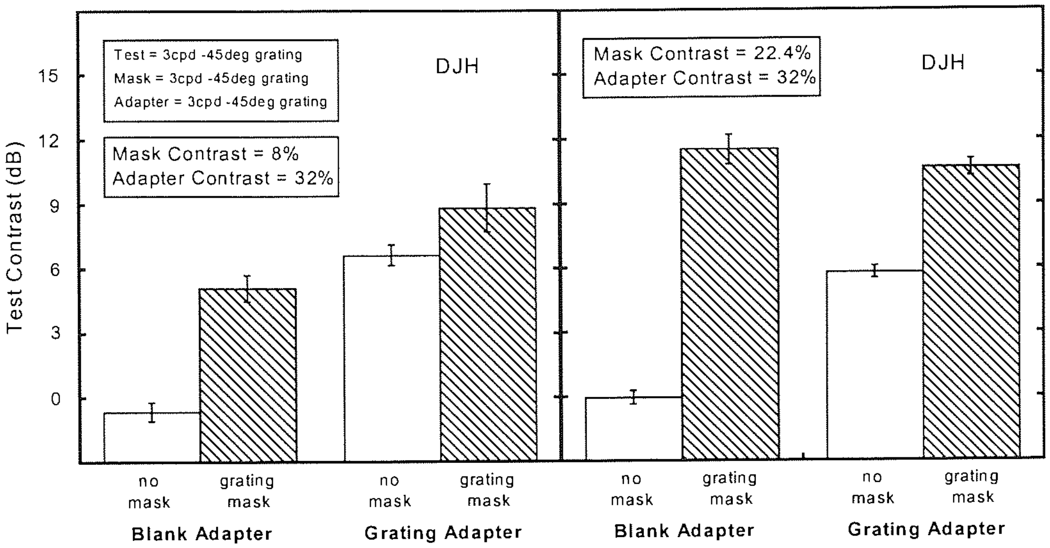

Fig. 7.2. Within-Channel Experiment. The effect of adaptation on masking when test, mask and adapter were all -45deg 3cpd gratings. Adapter contrast was 32% and mask contrast was either 8% (left panel) or 22.4% (right panel).

For the first experiment of the main study, tests were 1cpd vertical gratings. Adaptation conditions (either mean luminance adapter, or -45deg 3cpd grating adapter with 32% contrast) and mask contrast (ranging from 4.0% to 22.4%) were pseudo-randomly interleaved across sessions, whilst mask type (either no-mask, -45deg or +45deg 3cpd grating mask) was pseudo-randomly interleaved within sessions.

Fig. 7.3 shows the effect of -45deg (filled symbols) and +45deg (open symbols) 3cpd grating masks on the detectability of a 1cpd vertical test as a function of their contrast, after adaptation to either a mean luminance field (a) or a -45deg 3cpd grating with 32% contrast (b). Panel (c) shows the data from panels (a) and (b), superimposed. Test threshold as a function of mask contrast was a linearly increasing function on double log axes. However, the adapter had very little effect on test thresholds regardless of mask contrast.

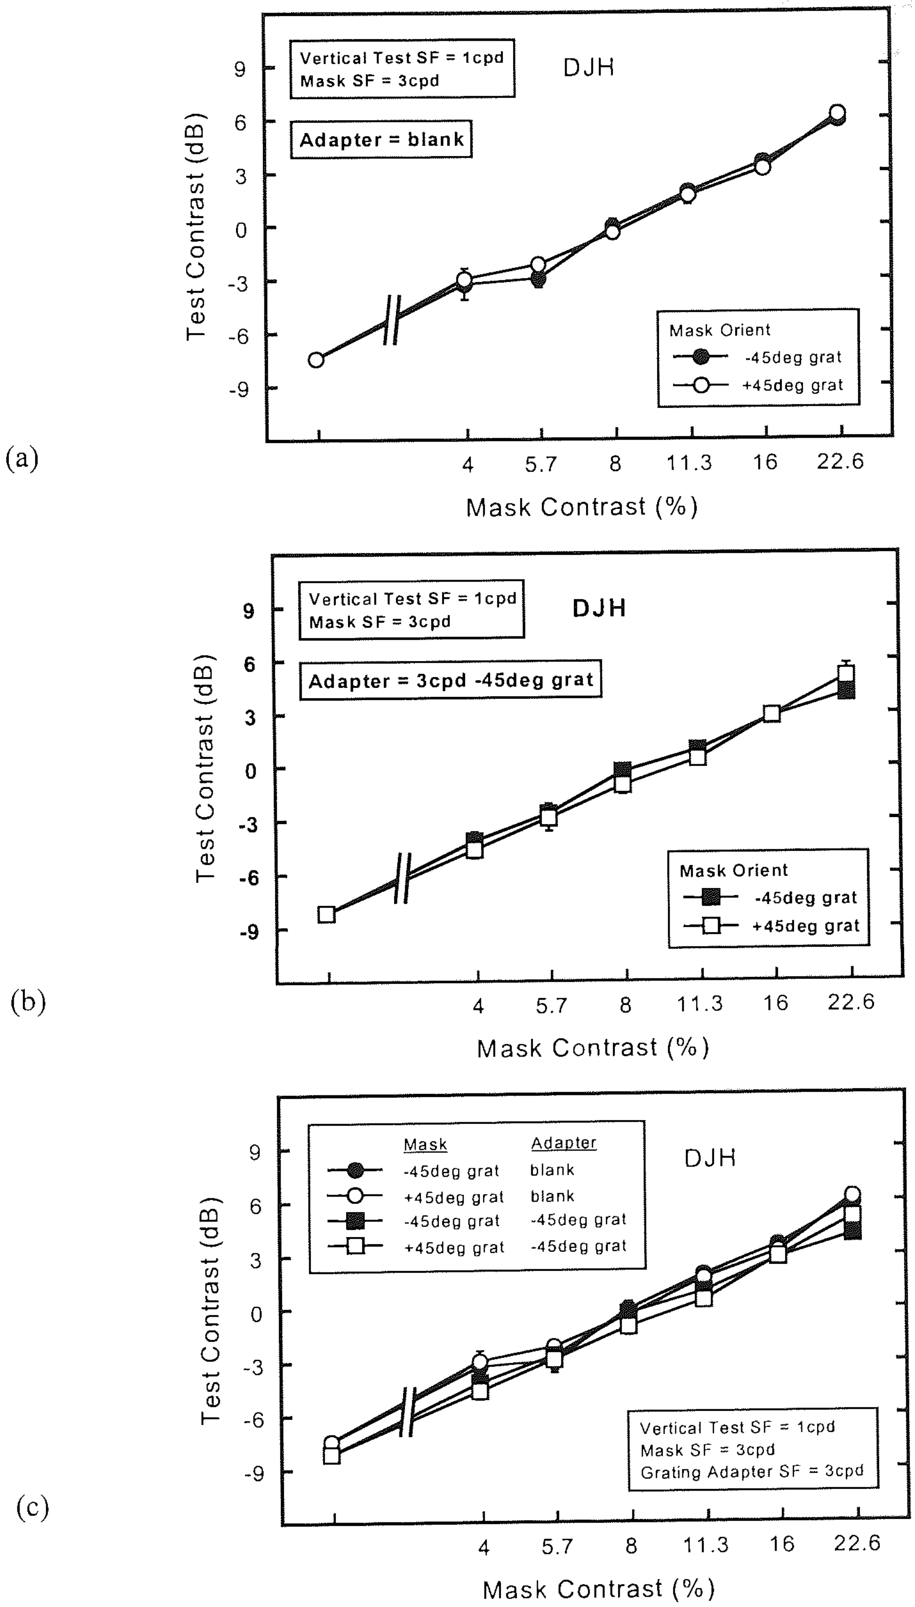

Fig. 7.3. Test contrast as a function of mask contrast after adaptation. The test was a vertical 1cpd grating and the mask was a -45deg 3cpd grating (filled symbols) or a +45deg 3cpd grating (open symbols). The observer adapted to either mean luminance ('blank') or a -45deg 3cpd grating; data from which are shown in panel (a) and panel (b), respectively. Panel (c) shows the data of the two previous panels, superimposed.

So, a grating adapter with 32% contrast has been found not to affect masked thresholds. It is known that increasing an adapter's contrast increases its effect on the detection of a test pattern (Blakemore & Campbell, 1969a; Blakemore & Nachmias, 1971; Stecher et al., 1973; Swift & Smith, 1982; Greenlee et al., 1991; Ross & Speed, 1991). Here, it was investigated whether effects of the adapter on masking levels are revealed when the adapter has a higher contrast. The previous experiment was repeated but using a grating adapter with a contrast of 99% and only one mask contrast (5.6%).

Fig. 7.4 shows the effect of adaptation on masking when adapter contrast was 32% (right panel) or 99% (left panel) and mask contrast was 5.6%. The data in the left panel are from Fig. 7.3. In Fig. 7.4, the unmasked thresholds were slightly lowered by the lower contrast adapter (LCA) and slightly raised by the higher contrast adapter (HCA). The masked thresholds (which were around 6dB higher than the unmasked thresholds) were affected very little. The HCA did slightly raise both masked thresholds, but this probably resulted from the effect that the pattern had on the unmasked threshold.

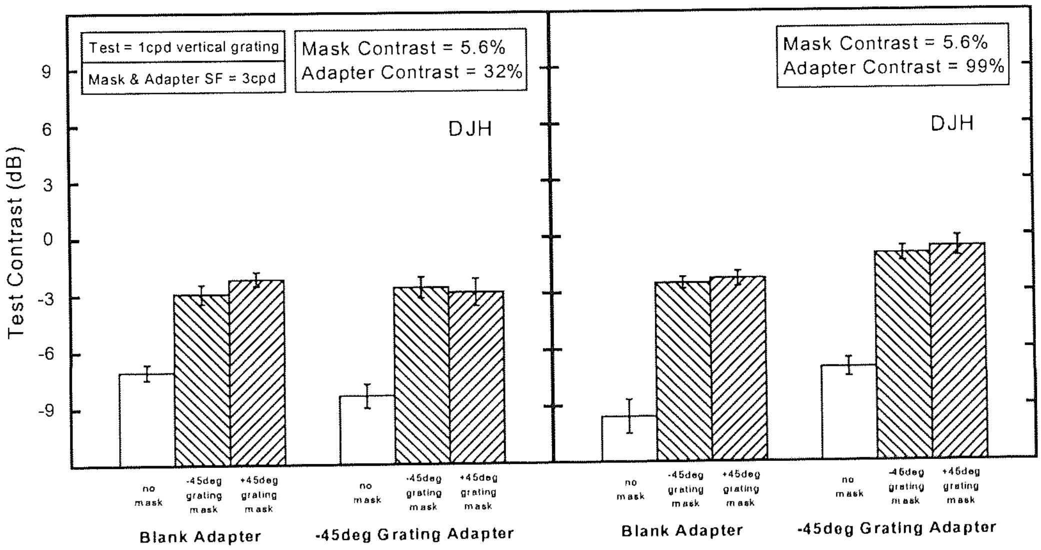

Fig. 7.4. The effect of adaptation on masking; the influence of (relative) adapter contrast. Mask contrast was 5.6%, whilst adapter contrast was either 32% (left panel) or 99% (right panel). The test was a 1cpd vertical grating whilst mask and adapter spatial frequencies were matched at 3cpd. Mask gratings were oriented either -45deg or +45deg whilst adapter gratings were always oriented -45deg. Data from conditions using mean luminance ('blank') masks and adapters are also shown.

To establish how general the null findings in Fig. 7.4 are, the experiment was repeated with all stimulus components (test, mask and adapter) having a spatial frequency of 1cpd. Fig. 7.5 shows that this produced a different pattern of results. Both adapters raised unmasked thresholds, with the HCA producing a larger effect than the LCA. Moreover, the LCA lowered the threshold elevation produced by the mask which shared its orientation with the adapter, yet had very little effect on the threshold elevation produced by the orthogonal mask. In terms of absolute thresholds, these were affected little when the LCA shared its orientation with the mask, but were raised when the patterns were orthogonal. The HCA's effects on masked thresholds were different to these; the levels of threshold elevation produced by both masks were reduced whilst absolute levels were raised.

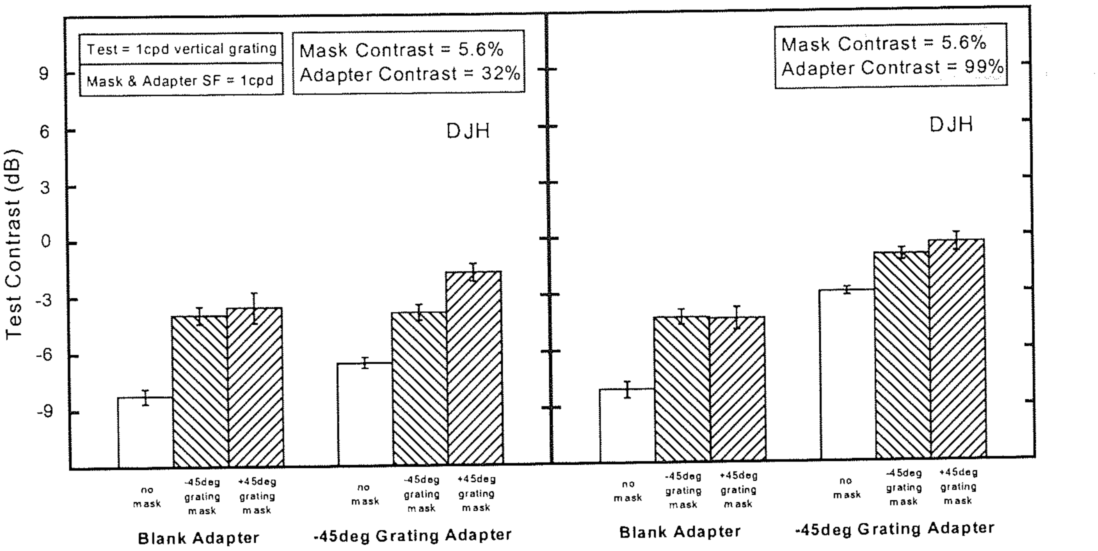

Fig. 7.5. The effect of adaptation on masking when all stimuli had a spatial frequency of 1cpd. The influence of (relative) adapter contrast is shown; mask contrast was 5.6% whilst adapter contrast was either 32% (left panel) or 99% (right panel). The test was vertical whilst mask gratings were oriented either -45deg or +45deg and adapter gratings were always oriented -45deg. Data from conditions using mean luminance ('blank') masks and adapters are also shown.

7.4 Chapter Discussion

In previous chapters, linear-suppression was shown between mask components oriented -45deg and +45deg when they had 3 or 1cpd, and when the test was a vertical grating with 1cpd. Adaptation-on-masking was measured in an attempt to reveal the orientation tuning of channels responsive to these masks. An adapter grating was matched in orientation (and spatial frequency) to one of the mask components, thus being orthogonal to the other. If it affected the masking from the parallel component substantially more than that from the orthogonal component, then this would have been consistent with processing by (linearly summing) oriented channels. Alternatively, if it had substantially affected both, then this would have been consistent with processing by an isotropic channel.

When adapter and masks had 3cpd, the adapter had negligible effects on masking so nothing was revealed about channel orientation tuning. When all of the stimuli had 1cpd, the effects of the adapter were again small, but the lowering of masked thresholds was selective. What these data may be revealing about the tuning of channels, however, is confusing. Whilst the effect of the low contrast adapter is suggestive of oriented masking channels, the effect of the high contrast adapter is suggestive of isotropic masking channels. In the absence of a coherent set of results that bear upon the chapter's aim, only one observer has been tested.

Essential to the outcome of the adaptation-on-masking experiments was the effect of adaptation on the efficacy of mask contrast. That is, the effect that it had on the suprathreshold contrast representation of the mask within the channel that was excited by the mask. If adaptation did affect this, then because the magnitude of masking (for contrast detection and discrimination) is in most cases (at moderate contrasts and above) positively correlated with mask contrast (e.g. see section 3.6 and 5.3), the masked thresholds would be expected to have been lowered. However, masked thresholds were affected little by adaptation. In attempting to understand why this was the case, the effects of adaptation on masking reported by others is reviewed.

Effects when test, mask and adapter are matched in orientation are considered first, followed by effects when mask and adapter are matched in orientation, but have an orientation which is distant from that of the test. Studies investigating the effects of adaptation on contrast detection, contrast discrimination and perceived contrast are considered.

Few studies have probed the effect of adaptation on masking. Most of those that have, used tests, masks and adapters that were spatially matched (or very similar). The effect of adaptation on contrast discrimination, when stimuli are matched, is controversial. Studies into the functional advantages of adaptation were the first to address this (Barlow et al., 1976; Kulikowski & Gorea, 1978; Legge, 1981). It was thought that after adapting to a pattern of a certain contrast, contrast discrimination around the level of the adapter's contrast might be improved. Masks and adapters of similar contrast were therefore used in the studies. Physiological evidence (see below) suggests that adaptation may be advantageous in this way and models of contrast gain control (e.g. Heeger, 1992; Wilson & Humanski, 1993) predict this. However, the above authors produced no evidence that adaptation improved contrast discrimination. If anything, their data were consistent with it having a small destructive influence on such thresholds (though see Abbonizio et al., 2002). More recently, a broader range of relative mask and adapter contrasts has been used. Greenlee and Heitger (1988) fixed adapter contrast at 80% and varied mask contrast between 10% and 80%. Generally, they found that contrast discrimination was hindered by adaptation when mask contrast was low (<50%), and enhanced when it was high (>50%)(that is, when adapter and mask contrasts were alike). Using a similar methodology, Wilson and Humanski (1993) also found the gradient of the contrast-increment threshold function to decrease with adaptation. However, Maattanen and Koenderink (1991) also used a wide range of mask contrasts (with a 50% adapter), but concluded that 'contrast sensitivity is not altered by contrast adaptation'. Investigating possible reasons for the differences between these data sets, Abbonizio et al. (2002) showed discrimination thresholds to be affected substantially more under binocular (like Greenlee & Heitger; 1988) than monocular (like Maattanen & Koenderink, 1991) viewing conditions. (The binocular/monocular manipulation was the same for all stimulus components, i.e. the test, mask and adapter.)(Note that all data shown in this thesis were gathered binocularly). They also found enhancements in contrast discrimination when adapter and mask were similar in contrast, but the effect was small, and varied between observers. (In relation, controversy also surrounds the function relating contrast detection to adapter contrast: for reviews, see Graham, 1989, and Foley & Boynton, 1993).

From the studies above, it would be expected that when test, mask and adapter were matched in spatial frequency and orientation, the adapter would have some notable effect on contrast discrimination. Indeed, this was found here (Fig. 7.2). When planning the experiments in this chapter, effects were expected if the mask and adapter were matched in orientation and spatial frequency, yet differed from the test in terms of one or both of these spatial characteristics. Foley and Chen (1997) provided evidence for this not being the case, however. They investigated the effect of a horizontal adapter on the masking, produced by a horizontal mask, of a vertical test. The stimulus gratings were matched in spatial frequency at 2cpd, adapter contrast was 50%, and the masks ranged in contrast between 0.5% and 32%. They found the adapter to have very little affect on the detection thresholds of the test grating.

So, Foley and Chen (1997) reported little affect from an adapter when the configuration of their stimuli was similar to that used here. Consistent with this, I found that adaptation did not affect masked thresholds when adapter and mask had a spatial frequency of 3cpd, and the test had a spatial frequency of

1cpd. There appeared to be small effects when all stimuli had a spatial frequency of 1cpd. Not much can be said about these data, however, because they are from only one observer. In spite of this, the two adapters, which only differed in their contrasts, had interestingly different influences on the masked thresholds. It would be of theoretical importance to probe whether an adapter's effect on masking is dependent upon its (relative) contrast with such stimulus configurations as was used here; after all, such relations have been reported when all stimuli were matched in spatial frequency and orientation.

It was considered that similar adaptation-on-masking levels in the main experiment here would indicate isotropic channels, whilst different levels would indicate oriented channels. Two major assumptions underlie these predictions. The first has been addressed already, that is, adaptation would influence the efficacy of the masking contrast. Regarding this, there appears to be complex interactions between the contrasts of the stimuli and the dependency of an adapter's effect on their relative orientations and spatial frequencies. Indeed, instances have been reported where adaptation has had very little, if any, affect on perceived contrast or contrast discrimination. This fits, to some degree, with the data shown here. The second assumption is that the affect of an adaptor on the efficacy of a mask's contrast is reliant on whether the adapter is exciting the same channel as the mask. Is this correct? It is a very well established finding that the affect of adaptation on contrast detection is highly orientation specific (Gilinsky, 1968; Mayo et al., 1968; Sekuler et al., 1968; Blakemore & Campbell, 1969a; Blakemore & Nachmias, 1971; Movshon & Blakemore, 1973; Kulikowski, 1972), which is consistent with physiological findings regarding the tuning of striate neurons (e.g. Hubel & Wiesel, 1962; 1968). However, the orientation specificity of adaptation on perceived contrast and contrast discrimination is more controversial. Blakemore et al. (1973) compared the orientation specificity of adaptation on contrast detection thresholds (Blakemore & Campbell, 1969a; Movshon & Blakemore, 1973) and perceived contrasts (their data). They reported equivalent narrow tuning; the adapter was found to have only a negligible effect on perceived contrast when it differed from the suprathreshold test by 45deg. At odds with this are more recent studies, which have found very broad tuning to orientation (Li & Aslin, 1992; Ross, 1992; Snowden & Hammett, 1992; Aslin & Li, 1993; Ross & Speed, 1996). (The effect, which is typically a reduction of perceived contrast, diminishes with increasing test contrast. This trend is similar to that found in the contrast discrimination studies cited above, where thresholds were only elevated at low test contrasts). The widespread effects which adapters have sometimes been found to produce on the perception of suprathreshold tests, suggests that the two patterns may not have to be exciting the same channel(s) for adaptation to occur. The importance of this to the chapter's main experiment is obvious; the reduction in the efficacy of one or both of the mask components is unlikely to provide clues to the orientation specificity of channels if this is the case. A further complicating factor in the interpretation of the discrimination thresholds that I have reported, is the relationship between the magnitude of the adaptation effects and the difference in orientation between the stimuli. Snowden and Hammett (1992) argued that the influence of adapters on the perceived contrast of parallel and orthogonal suprathreshold tests is subtractive and divisive, respectively. They showed that at high test contrasts, the perceived contrast of a test is lowered more by an orthogonal adapter than by a parallel adapter. This is because the function produced with the parallel adapter converges towards the line of no effect as test contrast is increased, whilst that produced with the orthogonal adapter runs parallel to the line of no effect. Ross and Speed (1996) provided no evidence for orthogonal adapters affecting perceived contrast more than parallel adapters, but showed that their effects were consistently less, except at high adapter contrasts (90%) where they were about the same. These two sets

Chapter Eight:
Summary and Further Discussions

The initial aim of the work in this thesis was to investigate masking effects when mask and test patterns were remote in spatial frequency and orientation (so that the masks were outside the pass-band of the test detecting mechanism). Insights emerged into how channels tuned to different spatial frequencies and orientations might interact. Some of the data gathered were used to test and develop cross-channel masking models. A distinct feature of such models is that detecting mechanisms are divisively inhibited by a broadly tuned gain pool of narrow-band spatial pattern mechanisms. These models superseded earlier within-channel models, which had supposed that masking arose from direct stimulation of the mask

Several remote masking effects were shown in chapter 3. Of particular interest was that the masking from plaids, whose components were remote in orientation from the test, appeared to be due to the linear sum of their component contrasts. In chapter 4, this result was shown to hold across a range of mask contrasts. Two widely used masking models proposed by Foley (1994a), which described divisive inhibitory pooling, were tested. Foley's model 2 fitted these data well whilst Foley's model 3 was rejected. In chapter 5, a new model was introduced which is a hybrid of Foley's two models. This model reconciles Foley's data with those in chapter 4. An experiment was devised which provided further tests for the models. Important differences between observers' data prompted the development of a series of modifications to them that improved their fit to the data.

The aim of chapter 6 was to investigate whether Phillips and Wilson (1984), in their classic study of orientation tuning, may have mis-estimated channel bandwidth because of their use of a within-channel masking model. Orientation masking data are shown for a greater range of test-mask orientation differences than had previously been investigated. Substantial phase-insensitive masking occurs at all relative orientations and evidence is provided that is consistent with this being due to cross-channel rather than within-channel masking. The cross-channel masking effects clearly show the need to reappraise the bandwidth estimates made by Phillips and Wilson (1984).

Chapter 7 extends the line of enquiry developed in chapter 4, which is shown there to distinguish between cross-channel models of masking. By devising related experiments using the adaptation paradigm, the aim here was to provide further tests of whether the masking was cross-channel or within-channel. Unfortunately, this line of enquiry delivered much less insight than the masking experiments of the earlier chapters.

Appendix 1 sets out some ideas and pilot work concerning how the enquiries of chapters 3 to 5 could be extended to investigate masking effects across spatial position.

In the majority of the experiments performed for this thesis, the contrast detection or discrimination thresholds of gratings were measured whilst they were being masked by either other gratings or plaids. The tasks allowed the use of a 2-AFC paradigm, which is criterion free and very easy to gather data with. Gratings have been a popular tool since it was realised that visual channels respond directly to limited ranges of spatial frequency and orientation (e.g. Pantle & Sekuler, 1968; Gilinsky, 1968; Blakemore & Campbell, 1969a), rather than to features such as edges (e.g. Coltheart, 1971; Kulikowski & King-Smith, 1973; Stromeyer & Klein, 1974). Their widespread use stems from them being simple and easy to control in terms of their spatio-temporal characteristics, as well as being mathematically tractable. What plaids have to offer is yet to be fully realised, though they have been a key tool in many influential studies (e.g. Adelson & Movshon, 1982; Welch 1989; Georgeson, 1992; Andrews & Blakemore, 1999). In chapters 3 to 5 they were shown to be very useful in probing summation rules of suppressive inputs to divisive inhibitory mechanisms. Working at the limits of the visual system, as in contrast detection and discrimination studies, appears to have been very successful at revealing the dynamics of early visual mechanisms. There are some things to bear in mind with such studies though and these are discussed later.

One of the major consequences of cross-channel masking has been the need to reevaluate all masking studies whose theoretical inferences about channel tuning were based entirely on the assumption of within-channel masking. In chapter 6 it was shown that Phillips and Wilson's (1984) omission of cross-channel effects would have led them to overestimate orientation bandwidths. Well over a hundred studies have cited theirs, many of which have used the estimates to set bandwidth values in models defined to predict various phenomena. It is unlikely, though, that the misestimates would invalidate such models or have any substantial effects on their predictions. Rather, the main value of the investigation in chapter 6 is that cross-orientation suppression was clearly shown to be substantially influenced by the spatio-temporal tuning of the channels involved. This was achieved by using one experimental procedure with wide ranges of stimulus spatial and temporal frequencies. The results go a long way to resolving why some studies have shown substantial orthogonal masking, whilst others have shown none. Of course, it was suspected that this was the case prior to evidence gathered in this thesis. These stimulus frequencies were the most obvious features by which the previous masking experiments differed. However, even though the reason for contrary results may be obvious from previously published data, this does not mean that an experiment in which it is explicitly addressed is unnecessary. There are numerous examples of investigations in which previously published data have been ignored or forgotten, even though the data are relevant and would help in interpreting the findings of those investigations. For example, Itti, Koch and Braun (2000) proposed a single model to simultaneously predict a large body of their masking data, including threshold measurements of contrast detection, discrimination, and orientation masking, and of spatial frequency and orientation discrimination. One aspect of their results is that no orientation masking was found when test and mask differed by more than 60°. They recognised that this was contrary to Foley's (1994a) finding of substantial orthogonal masking (where test thresholds had been raised by about 2.5-fold). The difference was considered to be due to Foley's large masks producing lateral masking as well as cross-orientation masking, whilst their stimuli had been designed to avoid large effects of lateral masking. An alternative explanation, which was not considered but which seems plausible in the light of the data in this thesis, concerns the stimuli's temporal duration and spatial frequency. In Foley (1994a), these were 33ms and 2cpd, respectively, whilst in Itti et al. (2000) they were 250ms and 2.8cpd, respectively. Many studies have produced data that are consistent with stimuli of a higher temporal frequency and lower spatial frequency producing the greater levels of orthogonal masking (see chapter 6). This doesn't undermine Itti's work, since their conclusions appear to be reasonable for the stimuli that they used; the key aspect to their work is that a single model can fit several different tasks with one set of parameters. However, a realisation of the dependency of cross-channel masking on stimulus spatio-temporal frequency is important for Itti's work and similar studies because substantially different masking functions would be expected with changes in these characteristics. The model may be able to cope with such changes by one or more of its parameters assuming different values (such as the weighting of the remote mask inputs), but this is not known until it is investigated. It may be the case that the architecture of their model requires revisions to capture the changes of the threshold levels in all five tasks. Alternatively, it may not be possible to produce good fits to all those data with a single model. So, though Itti et al. (2000) attempted to develop a 'unified account' of spatial vision by showing how one model of divisive inhibition can predict data from a wide range of tasks, small changes to the stimuli might result in poor model fits. Take, for example, their use of gratings with approximately 3cpd. Some effects of masks on the detectability of test patterns have been found to be substantially different when the stimuli have this spatial frequency, compared to when they have lower or higher spatial frequencies (see chapters 3 and 4). The differences found are between the summation rules for remote mask components in divisive inhibitory pooling mechanisms. Though this was not investigated by Itti et al. (2000), the extent of differences across tasks (e.g. discriminations of contrast, spatial frequency and orientation) and stimulus characteristics (e.g. temporal frequency) is not yet known, but may limit the conclusions that can be drawn from their work.

The model of Itti et al. (2000) represents a huge amount of work and should be praised. However, it cannot be taken as a complete model because they used stimuli of only one spatial frequency and one temporal duration. The finding that one model can account for various aspects of spatial vision should be used as a guide for further research. Unfortunately, there are many examples of models based on very restricted data sets that have not been sufficiently tested and developed, but which are regularly cited as being generally correct. Consider Foley's (1994a) models, for example, which are widely used as descriptors for the processes underlying cross-channel masking. Channels are described as being divisively inhibited. Contributing to the inhibition are a 'pool' of responses including the response of the channel that is itself being inhibited (i.e. self-inhibition), and the responses of many channels whose peak sensitivities (to, for example, orientation) cover a broad range of values. How the inputs to this pool are summed is obviously a major aspect of the model. Foley (1994a) investigated how self-inhibition is summed to the contribution from other channels and his models did well at predicting this. However, the ways in which the responses of these other channels are summed together has been ignored prior to the work in this thesis. An investigation into these summation rules revealed the need to develop the models and helped determine the direction of those developments. So, though Foley's (1994a) models have been found to produce good fits to certain data sets, small changes to the stimuli used in gathering these resulted in data that were poorly fit by the models. The work of Foley (1994a) is one example of many where the success of certain models in describing data is not as widespread as is typically thought. Whilst further investigations may reveal important limitations of such models and provide greater insight into visual processes, these are often neglected. It is important that most models like these are seen as being 'work in progress'. Consider, for example, the developments of Foley's (1994a) that were introduced in chapter 5. These models capture the important data characteristics that his do, whilst they also predict linear-suppression by summing the inputs from remote components linearly in the divisive inhibitory pool. Though these models are a step forward from Foley's (1994a), further developments to them are still necessary. For example, the processes underlying the super-suppression shown in chapters 3 and 4 are not understood. Super-suppression is the term given to the instance where a multiple-component mask produces substantially more remote masking than what is expected from the levels produced by a single-component mask. The data in this thesis indicate that it is far less widespread than linear-suppression (it has only been found using stimuli with spatial frequencies of around 3cpd). However, it is not predicted by either of Foley's models or any of those introduced in this thesis. (Foley 3 predicts that the multiple-component mask would produce less masking than the single component mask, whilst all of the other models predict linear-suppression).

Though evidence of super-suppression was found before Foley (1994a) introduced his models (Derrington & Henning, 1989; Ross & Speed, 1991), no mention has been made of it in relation to them (even though only a basic understanding of the models makes it obvious that they don't predict super-suppression). This is one example of where psychophysicists have not identified important gaps in the literature, and haven't performed the investigations that should be done, when they should be done. It has already been discussed that a study having clear limits in its scope can be used as a very valuable guide for further research, but Foley's (1994a) work is an example of where important 'loose-ends' have been overlooked. The view could be taken that this is acceptable because important gaps in understanding can be investigated at a later date, but there are reasons to argue otherwise. Certain times are far better than others to study and gain insight into certain phenomena. To explain, consider the work of Phillips and Wilson (1984) as an example. They measured orientation masking with test and mask gratings that differed by between 0 and 45deg, for a range of stimulus spatial frequencies. A within-channel model of masking was fit to these data to allow the bandwidths of oriented channels to be estimated (see chapter 6 for a discussion of this work). The lower spatial frequency (0.5 and 2cpd) transient stimuli that they used, produced functions indicating substantial masking when test and mask differed by 45deg (see Fig. 6.2). Though the model predicted a systematic and quite rapid reduction of masking as the orientation difference between the stimuli was increased, these high levels were still captured quite well. Phillips and Wilson (1984) didn't show data for orientation differences greater than 45deg. If they had, it would have been clear that the masking functions 'levelled out', that is, the (substantial) masking declined much less between orientation differences of 45 and 90deg, than between 0 and 45deg (as shown in Fig. 6.1). (There are indications in their data that this would be the case, particularly in those from observer WS; see Fig. 6.2). The within-channel model does not predict this feature, so if it had been shown by them then their work would have been less attractive as a published unit of research. However, there would have been several important benefits if this remote orientation masking had been measured by them with the range of stimulus spatial and temporal frequencies that they used. Early evidence would have been found for cross-channel suppression, and for its dependence on spatio-temporal frequency. This would have been important in the movement away from within-channel masking models in the early 1990's. Valuable insight would also have been provided into why some studies of orthogonal masking have revealed substantial effects, whilst others have reported only negligible effects; a discrepancy that has created some confusion since the 1970's.

Foley's (1994a) work provided impetus to investigate summation rules of inputs into the divisive inhibitory pools that he had advocated in his masking models. Phillips and Wilson's (1984) work provided impetus to compare masking across the full range of possible orientation differences between test and mask gratings, including regions where the patterns would not be expected to excite a common detecting mechanism as well as regions where they would. However, such studies had not been carried out prior to the work in this thesis. It was found with experiments in chapters 3 to 5 that rules governing the summation of the inputs from two channels in a divisive inhibitory pool differ substantially with the preferred spatial frequencies of the channels and whether the test pattern excites either of them. Further, it was found with experiments in chapter 6 that whether cross-channel suppression occurs when test and mask patterns have very different orientations depends on their spatial and temporal frequencies. Substantial effects are produced with low spatial frequency stimuli and high temporal frequency stimuli. The work of Foley (1994a) and Phillips and Wilson (1984) are two examples of where important phenomena have been missed because of the very limited ranges of stimuli used. More recent studies also show this lack of attention to the effects of these stimulus characteristics. For example, Schwartz and Simoncelli (2001) used a model of divisive inhibitory pooling in an attempt to describe how natural images might be encoded efficiently, but the dependency of divisive inhibitory pooling on pattern spatial and temporal frequencies was again neglected.

Prior to the modelling developments described above, further work is required to gain a fuller understanding of the processes underlying linear- and super-suppression. Experiments have been planned by Meese and Georgeson to probe the possibility that super-suppression is the result of inputs from second order mechanisms to the divisive inhibitory pool. Regarding linear-suppression, the adaptation study in chapter 7 was designed to investigate whether isotropic channels or linearly summing, narrowly tuned channels may underlie this, but the results from those particular experiments were inconclusive. In considering mechanisms underlying masking, Carandini and colleagues have suggested that the LGN may play a larger role in masking than is widely believed. Masking originating at a subcortical level could explain linear-suppression since the neurons in these areas are known to be very broadly tuned to orientation and spatial frequency (Maffei & Fiorentini, 1973, Levick & Thibos, 1980; DeValois et al., 1982; Rovamo et al., 1982; Fahle, 1986). In attempting to describe how this may be the case despite LGN efferents not being suppressive, Freeman et al. (2002) and Meier and Carandini (2002) have argued that masking is generated partly in the LGN and is supplemented by synaptic depression at the thalamocortical synapse. This theory of synaptic depression allows suppression to originate at the LGN, but manifest itself only at the first synaptic level in the cortex. Linear-suppression would be expected if the amount of synaptic depression produced by a grating was the same as that produced by a plaid when the two patterns were matched in overall contrast. Meese and Georgeson, with Hess, plan to study the role of the LGN in linear-suppression with experiments that probe masking with amblyopic observers and dichoptic masking.

A major part of the support for the idea that masking originates in the LGN, was the finding that cross-orientation suppression (i.e. the suppression of a neuron produced by a grating orthogonal to its preferred orientation) was not affected by adaptation (Freeman et al., 2002). Since adaptation substantially affects neurons in V1 (Maffei et al., 1973; Albrecht et al., 1984; Ohzawa et al., 1985), whilst having very little affect on those in the LGN (Shou et al., 1996; Sanchez-Vives et al., 2000a), it was argued that masking must be occurring in the LGN rather than in the cortex. However, the affect of adaptation on masking is poorly understood. There is a growing body of evidence showing that adaptation sometimes does affect cross-channel masking. Studies discussed in chapter 7 have shown such effects to be dependent upon the contrasts of the mask and adapter patterns (particularly the relationship between them). In addition, data in that chapter are consistent with stimulus spatial frequency also playing a role (e.g. compare Fig. 7.4 and Fig. 7.5).

Masking and adaptation have been widely used as psychophysical tools for over 30 years, but much has still to be learnt about how they affect visual mechanisms. This limits what can be inferred from data produced with them. Further studies of effects from adapters on masking will lead to greater insights into their underlying processes, particularly how the two are related (i.e. how they affect common mechanisms). Work has already begun on simultaneously predicting effects of masking and adaptation on test detection, and effects of adaptation on masking. Foley and Chen (1997) and Meese and Holmes (2002) have developed Foley's (1994a) masking models for this. Their models incorporate desensitisation from adaptation upon divisive inhibitory pooling. Though the data that have been fit with these is very limited, Meese and Georgeson have planned further experiments to test models of this type.

In the experiments of this thesis, effects of masking and adaptation on contrast detection and contrast discrimination thresholds were measured. Some of the resulting data were then used to test masking models. Perceived contrast, measured by contrast matching, is a suprathreshold phenomenon that was not used here, but it has been used elsewhere to test ideas of masking (e.g. Cannon & Fullenkamp, 1991; Snowden & Hammett, 1998; Yu & Levi, 2000; Xing & Heeger, 2001; and see appendix 1) and adaptation (e.g. Blakemore et al., 1973; Snowden & Hammett, 1992; Ross & Speed, 1996; and see chapter 7). Masking experiments using the method have been restricted to the study of lateral masking. It has been avoided with superimposed masking presumably because, rather than being able to judge test contrast, the test and mask contrasts might perceptually integrate (Georgeson & Meese, 1997). However, pilot studies by Meese and Hess have shown that judgements of perceived test contrast are simple to make with stimuli similar to those used in demonstrating linear- and super-suppression (i.e. the test and mask patterns perceptually segment when they are remote in orientation and spatial frequency; see Georgeson, 1998). Meese and Georgeson, with Hess, plan to use contrast matching with such stimuli to study how mask and adapter inputs might be handled by divisive inhibitory pooling mechanisms in suprathreshold viewing conditions.

For all of the masking studies presented in the main body of this thesis, the test and mask patterns were spatially superimposed. Some ideas and pilot work concerning the processes underlying lateral masking are presented in appendix 1.

In summary, it has been discussed here how:

1. People have performed some very good experiments, but the stimulus conditions of these have been limited.
2. This thesis has extended the range of stimulus conditions, and uncovered some valuable new findings. Comparisons of masking levels produced with single- and multiple-component stimuli were particularly fruitful lines of enquiry. They provided insights into the rules governing stimulus contrast summation in contrast gain controls.
3. The work in this thesis has served as a ‘springboard’ for work planned by Meese and Georgeson, with Hess (EPSRC 3year grant ‘Suppression and summation in contrast gain control for human vision’, reference GR/S74515/01).

References

Abbonizio, G., Langley, K., and Clifford, C. W. G. (2002) Contrast adaptation may enhance contrast discrimination. *Spatial Vision*, **16**, 45 - 58.

Adelson, E. H., and Movshon, J. M. (1982) Phenomenal coherence of moving visual patterns. *Nature*, **300**, 523 - 525.

Adjamian, P., and Holliday, I. E. (2002) Response properties of visual cortex to grating patterns using MEG: Evidence for a new subset of superficial pyramidal cells in the primary visual cortex of man. *British Psychophysiological Society, Proceedings of the 30th Annual Scientific Meeting*, pp. 85.

Adini, Y., Sagi, D., and Tsodyks, M. (1997) Excitatory-inhibitory network in the visual cortex: Psychophysical evidence. *Proceedings of the National Academy of Sciences of the United States of America*, **94**, 10426 - 10431.

Adorjan, P., Piepenbrock, C., and Obermayer, K. (1999) Contrast adaptation and infomax in visual cortical neurons. *Review of Neuroscience*, **10**, 181 - 200.

Ahmed, B., Allison, J. D., Douglas, R. J., and Martin, K. A. C. (1997) An intracellular study of the contrast-dependence of neuronal activity in cat visual cortex. *Cerebral Cortex*, **7**, 559 - 570.

Albrecht, D. G., Farrar, S. B., and Hamilton, D. B. (1984) Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. *Journal of Physiology (London)*, **347**, 713 - 739.

Albrecht, D. G., and Geisler, W. S. (1991) Motion selectivity and contrast response function of simple cells in the visual cortex. *Visual Neuroscience*, 7, 531 - 546.

Albrecht, D. G., and Hamilton, D. B. (1982) Striate cortex of monkey and cat: contrast response function. *Journal of Neurophysiology*, **48**, 217 - 237.

Allison, J. D., Smith, K. R., and Bonds, A. B. (2001) Temporal-frequency tuning of cross-orientation suppression in the cat striate cortex. *Visual Neuroscience*, **18**, 941 - 948.

Allman, J., Miezin, F., and McGuinness, E. (1985) Stimulus specific responses from beyond the classical receptive field: neurophysiological mechanisms for local-global comparisons in visual neurons. *Annual Review of Neuroscience*, **8**, 407 - 430.

Anderson, J., Carandini, M., and Ferster, D. (2000) Orientation tuning of input conductance, excitation and inhibition in cat primary visual cortex. *Journal of Neurophysiology*, **84**, 909 - 931.

Anderson, S. J., and Burr, D. C. (1985) Spatial and temporal selectivity of the human motion detection system. *Vision Research*, **25**, 1147 - 1154.

Anderson, S. J., Burr, D. C., and Morrone, M. C. (1991) Two-dimensional spatial and spatial frequency selectivity of motion-sensitive mechanisms in human vision. *Journal of the Optical Society of America, A*, **8**, 1340 - 1351.

Andrews, D. P. (1965) Perception of contours in the central fovea. *Nature (London)*, **205**, 1218 - 1220.

Andrews, T. J., and Blakemore, C. (1999) Form and motion have independent access to consciousness. *Nature Neuroscience*, **2**, 405 - 406.

Andrews, T. J., and Lotto, R. B. (in press) Fusion or rivalry are contingent on the perceptual meaning of visual stimuli. *Perception*.

Aslin, R. N., and Li, A. (1993) Multiple mechanisms of contrast adaptation. *Investigative Ophthalmology and Visual Science Supplement*, **34**, 1362.

Bair, W., Cavanaugh, J. R., and Movshon, J. A. (1999) The time course of contextual modulation in macaque V1 neurons. *Investigative Ophthalmology and Visual Science*, **40 (4)**, 1058.

Barfield, L. P., and Tolhurst, D. J. (1975) The detection of complex gratings by the human visual system. *Journal of Physiology (London)*, 248, 37 - 38.

Barlow, H. B. (1953) Summation and inhibition in the frog's retina. *Journal of Physiology*, **119**, 69 - 88.

Barlow, H. B., MacLeod, D. I. A., and van Meeteren, A. (1976) Adaptation to gratings: no compensatory advantages found. *Vision Research*, **16**, 1043 - 1045.

Bauman, L. A., and Bonds, A. B. (1991) Inhibitory refinement of spatial frequency selectivity in single cells of the cat striate cortex. *Vision Research*, **31**, 933 - 944.

Baumgartner, G., Brown, J. L., and Schulz, A. (1965) Responses of single units of the cat visual system to rectangular stimulus patterns. *Journal of Neurophysiology*, **28**, 1 - 18.

Ben-Av, M. B., and Sagi, D. (1995) Perceptual grouping by similarity and proximity: experimental results can be predicted by intensity autocorrelations. *Vision Research*, **35**, 853 - 866.

Benevento, L. A., Creutzfeld, O., and Kuhnt, U. (1972) Significance of intracortical inhibition in the visual cortex. *Nature*, **238**, 124 - 126.

Ben-Yishai, R., Lev Bar Or, R., and Sompolinsky, H. (1995) Theory of orientation tuning in the visual cortex. *Proceedings of the National Academy of Sciences*, **92**, 3844 - 3848.

Bergen, J. R., and Landy, M. S. (1991) Computational modeling of visual texture segregation. In: Computational Models of Visual Processing, (Eds: Landy, M. S. and Movshon, J. A.), M. I. T. Press, Cambridge, Mass., 253 - 271.

Bishop, P. O., Coombs, J. S., and Henry, G. H. (1971a) Responses to visual contours: Spatio-temporal aspects of excitation in the receptive fields of simple striate neurons. *Journal of Physiology*, **219**, 625 - 657.

Bishop, P. O., Coombs, J. S., and Henry, G. H. (1971b) Interaction effects of visual contours on the discharge frequency of simple striate neurones. *Journal of Physiology*, **219**, 659 - 687.

Bishop, P. O., Coombs, J. S., and Henry, G. H. (1973) Receptive fields of simple cells in the cat striate cortex. *Journal of Physiology (London)*, **231**, 31 - 60.

Bjorklund, R. A., and Magnussen, S. (1981) A study of interocular transfer of spatial adaptation. *Perception*, **10**, 511 - 518.

Blakemore, C., and Campbell, F. W. (1969a) On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology*, *London*, **203**, 237 - 260.

Blakemore, C., and Campbell, F. W. (1969b) Adaptation to spatial stimuli. *Journal of Physiology*, **200**, 11 - 13P.

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

Blakemore, C., Muncey, J. P. J., and Ridley, R. M. (1971) Perceptual fading of a stabilized cortical image. *Nature*, **233**, 204 - 205.

Blakemore, C., Muncey, J. P. J., and Ridley, R. M. (1973) Stimulus specificity in the human visual system. *Vision Research*, **13**, 1915 - 1931.

Blakemore, C., and Nachmias, J. (1971) The orientation specificity of two visual aftereffects. *Journal of Physiology*, **213**, 157 - 174.

Blakemore, C., Nachmias, J., and Sutton, P. (1970) The perceived spatial frequency shift: Evidence for frequency selective neurones in the human brain. *Journal of Physiology*, **210**, 727 - 750.

Blakemore, C., and Tobin, E. A. (1972) Lateral inhibition between orientation detectors in the cat's visual cortex. *Experimental Brain Research*, **15**, 439 - 440.

Blasdel, G. G., and Fitzpatrick, D. (1984) Physiological organization of layer 4 in macaque striate cortex. *Journal of Neuroscience*, **4**, 880 - 895.

Bobak, P., Bodis-Wollner, I., and Marx, M. S. (1988) Cortical contrast gain control in human spatial vision. *Journal of Physiology (London)*, **405**, 421 - 437.

Bodinger, D. M. (1978) The decay of grating adaptation. *Vision Research*, **18**, 89 - 91.

Bonds, A. B. (1989) Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Visual Neuroscience*, **2**, 41 - 55.

Bonds, A. B. (1991) Temporal dynamics of contrast gain in single cells of the cat striate cortex. *Visual Neuroscience*, **6**, 239 - 255.

Bonneh, Y., and Sagi, D. (1999) Contrast integration across space. *Vision Research*, **39**, 2597 - 2602.

Borg-Graham, L. J., Monier, C., and Fregnac, Y. (1998) Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature*, **393**, 369 - 373.

Bowen, R. W., and Wilson, H. R. (1994) A two-process analysis of pattern masking. *Vision Research*, **34**, 645 - 657.

Boynton, G. M., Demb, J. B., Glover, G. H., and Heeger, D. J. (1999) Neuronal basis of contrast discrimination. *Vision Research*, **39**, 257 - 269.

Boynton, G. M., and Foley, J. M. (1999) Temporal sensitivity of human luminance pattern mechanisms determined by masking with temporally modulated stimuli. *Vision Research*, **39**, 1641 - 1656.

Braddick, O., Campbell, F. W., and Atkinson, J. (1978) Channels in vision: basic aspects. In *Handbook of Sensory Physiology. VIII. Perception*, edited by R. Held, H. W. Leibowitz, and H. L. Teuber (Springer, Berlin).

Bradley, A., and Ohzawa, I. (1986) Comparison of contrast detection and discrimination. *Vision Research*, **26**, 991 - 997.

Breitmeyer, B. G. (1975) Simple reaction time as a measure of the temporal response properties of transient and sustained channels. *Vision Research*, **15**, 1411 - 1412.

Breitmeyer, B. G. (1984) *Visual masking: An integrative approach.* New York: Oxford University Press.

Bruce, V., Green, P. R., and Georgeson, M. A. (1996) *Visual Perception: physiology, psychology and ecology (3^{rd} ed.).* Psychology Press: Erlbaum, Taylor and Francis.

Burbeck, C. A., and Kelly, D. H. (1981) Contrast gain measurements and the transient / sustained dichotomy. *Journal of the Optical Society of America*, **71**, 1335 - 1342.

Burr, D. C., and Morrone, M. C. (1987) Inhibitory interactions in the human visual system revealed in pattern-evoked potentials. *Journal of Physiology (London)*, **389**, 1 - 21.

Burr, D. C., Morrone, M. C., and Maffei, L. (1981) Intra-cortical inhibition prevents simple cells from responding to textured visual patterns. *Experimental Brain Research*, **43**, 455 - 458.

Burr, D. C., and Wijesundra, S.-A. (1991) Orientation discrimination depends on spatial frequency. *Vision Research*, **31**, 1449 - 1452.

Campbell, F. W. (1974) The transmission of spatial information through the visual system. In *The Neurosciences Third Study Program*, edited by F. O. Schmitt and F. G. Worden (MIT, Cambridge, MA).

Campbell, F. W., Cleland, B. G., Cooper, G. F., and Enroth-Cugell, C. (1968) The angular selectivity of visual cortical cells to moving gratings. *Journal of Physiology*, **198**, 237 - 250.

Campbell, F. W., Cooper, G. F., and Enroth-Cugell, C. (1969) The spatial selectivity of the visual cells of the cat. *Journal of Physiology*, **203**, 223 - 235.

Campbell, F. W., and Kulikowski, J. J. (1966) Orientational selectivity of the human visual system. *Journal of Physiology*, **187**, 437 - 445.

Campbell, F. W., and Robson, J. G. (1968) Applications of fourier analysis to the visibility of gratings. *Journal of Physiology*, **197**, 551 - 566.

Candy, T. R., Skoczenski, A. M., and Norcia, A. M. (2001) Normalization models applied to orientation masking in the human infant. *Journal of Neuroscience*, **21**, 4530 - 4541.

Cannon, M. W. (1985) Perceived contrast in the fovea and periphery. *Journal of the Optical Society of America (A)*, **2**, 1760 - 1768.

Cannon, M. W., and Fullenkamp, S. C. (1991) Spatial interactions in apparent contrast: inhibitory effects among grating patterns of different spatial frequencies, spatial positions and orientations. *Vision Research*, **31**, 1985 - 1998.

Cannon, M. W., and Fullenkamp, S. C. (1993) Spatial interactions in apparent contrast: Individual differences in enhancement and suppression effects. *Vision Research*, **33**, 1685 - 1695.

Cannon, M. W., and Fullenkamp, S. C. (1996) A model for inhibitory lateral interaction effects in perceived contrast. *Vision Research*, **36**, 1115 - 1125.

Carandini, M., Barlow, H. B., O'Keefe, L. P., Poirson, A. B., and Movshon, J. A. (1997) Adaptation to contingencies in macaque primary visual cortex. *Philosophical Transactions of the Royal Society of London (B, Biological Sciencies)*, **352**, 1149 - 1154.

Carandini, M., and Ferster, D. (1997) A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science*, **276**, 949 - 952.

Carandini, M., and Ferster, D. (2000) Membrane potential and firing rate in cat primary visual cortex. *Journal of Neuroscience*, **20**, 470 - 484.

Carandini, M., and Heeger, D. J. (1994) Summation and division by neurons in visual cortex. *Science*, **264**, 1333 - 1336.

Carandini, M., Heeger, D. J., and Movshon, J. A. (1997) Linearity and normalization in simple cells of the macaque primary visual cortex. *Journal of Neuroscience*, **17**, 8621 - 8644.

Carandini, M., Heeger, D. J., and Senn, W. (2002) A synaptic explanation of suppression in visual cortex. *Journal of Neuroscience*, **22**, 10053 - 10065.

Carandini, M., and Ringach, D. L. (1997) Predictions of a recurrent model of orientation selectivity. *Vision Research*, **37**, 3061 - 3071.

Carlson, C. R. (1978) Thresholds for perceived image sharpness. *Photographic Science and Engineering*, **22**, 69 - 71.

Carlson, C. R., and Cohen, R. W. (1980) A simple psychophysical model for predicting the visibility of displayed information. *Proc. SID*, **21**, 229 - 246.

Carpenter, F. W., and Blakemore, C. (1973) Interactions between orientations in human vision. *Experimental Brain Research*, **18**, 287 - 303.

Chance, F. S., Nelson, S. B., and Abbott, L. F. (1998) Synaptic depression and the temporal response characteristics of V1 cells. *Journal of Neuroscience*, **18**, 4785 - 4799.

Chen, C.-C., Foley, J. M., and Brainard, D. H. (2000a) Detection of chromoluminance patterns on chromoluminance pedestals I: threshold measurements. *Vision Research*, **40**, 773 - 788.

Chen, C.-C., Foley, J. M., and Brainard, D. H. (2000b) Detection of chromoluminance patterns on chromoluminance pedestals II: model. *Vision Research*, **40**, 789 - 803.

Chen, C.-C., and Tyler, C. W. (2002) Lateral modulation of contrast discrimination: Flanker orientation effects. *Journal of Vision*, **2**, 520 - 530.

Chubb, C., Sperling, G., and Solomon, J. A. (1989) Texture interactions determine perceived contrast. *Proceedings of the National Academy of Sciences USA*, **86**, 9631 - 9635.

Clifford, C. W. G., Wyatt, A. M., Arnold, D. H., Smith, S. T., and Wenderoth, P. (2001) Orthogonal adaptation improves orientation discrimination. *Vision Research*, **41**, 151 - 159.

Colonnier, M. (1968) Synaptic patterns on different cell types in the different laminae of the cat visual cortex: an electron microscope study. *Brain Research*, **9**, 268 - 287.

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

Cornsweet, T. N. (1962) The staircase method in psychophysics. *American Journal of Psychology*, 75, 485 - 491.

Cronly-Dillon, J. R. (1964) Units sensitive to direction of movement in goldfish optic tectum. Nature (Lond), **203**, 214 - 215.

Crook, J. M., and Eysel, U. T. (1992) GABA-induced inactivation of functionally characterized sites in cat visual cortex (area 18): Effects on orientation tuning. *Journal of Neuroscience*, **12**, 1816 - 1825.

Crook, J. M., Kisvarday, Z. F., and Eysel, U. T. (1998) Evidence for a contribution of lateral inhibition to orientation tuning and direction selectivity in cat visual cortex: reversible inactivation of functionally characterized sites combined with neuroanatomical tracing techniques. *European Journal of Neuroscience*, **10**, 2056 - 2075.

Curtis, D. R. (1978) GABAergic transmission in the mammalian central nervous system. In: Krogsgaard-Larson, P., Scheel-Kruger, J. and Kofod, H. (Eds.) *GABA neurotransmitters - Alfred Benzon Symposium 12*. Copenhagen, pp. 17 - 27.

Dannemiller, J. L., and Stephens, B. R. (2000) Contrast discrimination under temporally varying contrast conditions. *Perception and Psychophysics*, **62**, 1039 - 1050.

Daugman, J. G. (1984) Spatial visual channels in the fourier plane. *Vision Research*, **24**, 891 - 910.

Dealy, R. S., and Tolhurst, D. J. (1974) Is spatial adaptation an after-effect of prolonged inhibition? *Journal of Physiology (London)*, **241**, 261 - 270.

DeAngelis, G. C., Freeman, R. D., and Ohzawa, I. (1994) Length and width tuning of neurons in the cat's primary visual cortex. *Journal of Neurophysiology*, 71, 347 - 374.

DeAngelis, G. C., Robson, J. G., Ohzawa, I., and Freeman, R. D. (1992) The organization of suppression in receptive fields of neurons in cat visual cortex. *Journal of Neurophysiology*, **68**, 144 - 163.

DeAngelis, G. C., Ohzawa, I., and Freeman, R. D. (1993) Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development. *Journal of Neurophysiology*, **69**, 1091 - 1117.

Derrington, A. M., and Henning, G. B. (1989) Some observations on the masking effects of two-dimensional stimuli. *Vision Research*, **29**, 241 - 246.

DeValois, K. K. (1977) Spatial frequency adaptation can enhance contrast sensitivity. *Vision Research*, **17**, 1057 - 1065.

DeValois, K. K., DeValois, R. L., and Yund, E. W. (1979) Responses of striate cortex cells to grating and checkerboard patterns. *Journal of Physiology*, **291**, 483 - 505.

DeValois, K. K., and Tootell, R. H. B. (1983) Spatial-frequency-specific inhibition in cat striate cortex cells. *Journal of Physiology*, **336**, 359 - 376.

DeValois, R. L., Albrecht, D. G., and Thorell, L. G. (1982) Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research*, **22**, 545 - 559.

DeValois, R. L., and DeValois, K. K. (1988) *Spatial Vision*, Oxford University Press, New York.

DeValois, R. L., Yund, E. W., and Hepler, N. (1982) The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research*, **22**, 531 - 544.

Dorais, A., and Sagi, D. (1997) Contrast masking effects change with practice. *Vision Research*, **37**, 1725 - 1733.

Douglas, R. J., Martin, K. A. C., and Whitteridge, D. (1988) Selective responses of visual cortical cells do not depend on shunting inhibition. *Nature*, **332**, 642 - 644.

Dresp, B. (1993) Bright lines and edges facilitate the detection of small line targets. *Spatial Vision*, **7**, 213 - 225.

Durand, S., Mante, V., Freeman, T. C. B., and Carandini, M. (2002) Temporal properties of surround suppression in primary visual cortex. Abstract 051.6 (poster p306) presented at 3rd Forum of European Neuroscience, Palais des Congres, July 13 - 17, 2002.

D'Zmura, M., and Singer, B. (1996) Spatial pooling of contrast in contrast gain control. *Journal of the Optical Society of America A*, **13**, 2135 - 2140.

Ejima, Y., and Takahashi, S. (1983) Effects of high-contrast peripheral patterns on the detection threshold of sinusoidal targets. *Journal of the Optical Society of America*, **73**, 1695 - 1700.

Ejima, Y., and Takahashi, S. (1985) Apparent contrast of a sinusoidal grating in the simultaneous presence of peripheral gratings. *Vision Research*, **25**, 1223 - 1232.

Enroth-Cugell, C., and Robson, J. G. (1966) The contrast sensitivity of retinal ganglion cells of the cat. *Journal of Physiology*, **187**, 517 - 552.

Fahle, M. (1986) Curvature detection in the visual field and a possible physiological correlate. *Experimental Brain Research*, **63**, 113 - 124.

Felsen, G., Shen, Y.-S., Yao, H., Spor, G., Li, C., and Dan, Y. (2002) Dynamic modification of cortical orientation tuning mediated by recurrent connections. *Neuron*, **36**, 945 - 954.

Ferrera, V. P., and Wilson, H. R. (1985) Spatial frequency tuning of transient non-oriented units. *Vision Research*, **25**, 67 - 72.

Ferster, D., and Miller, K. D. (2000) Neural mechanisms of orientation selectivity in the visual cortex. *Annual Review of Neuroscience*, **23**, 441 - 471.

Field, D. J. (1987) Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, **4**, 2379 - 2394.

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

Finlayson, P. G., and Cynader, M. S. (1995) Synaptic depression in visual cortex tissue slices: an in vitro model for cortical neuron adaptation. *Experimental Brain Research*, **106**, 145 - 155.

Fitzpatrick, D. (2000) Seeing beyond the receptive field in primary visual cortex. *Current Opinions in Neurobiology*, **10**, 438 - 443.

Foley, J. M. (1994a) Human luminance pattern-vision mechanisms: masking experiments require a new model. *Journal of the Optical Society of America A*, **11**, 1710 - 1719.

Foley, J. M. (1994b) Spatial phase sensitivity of luminance pattern mechanisms determined by masking. *Investigative Ophthalmology and Visual Science*, **35**, 4, 1900.

Foley, J. M. (1996) Simultaneous pattern masking: How come threshold elevation bandwidth decreases as stimulus duration increases? *Investigative Opthalmology and Visual Science*, **37**, 4205.

Foley, J. M., and Boynton, G. M. (1993) Forward pattern masking and adaptation: Effects of duration, interstimulus interval, contrast, and spatial and temporal frequency. *Vision Research*, **33**, 959 - 980.

Foley, J. M., and Boynton, G. M. (1994) A new model of human luminance pattern vision mechanisms: analysis of the effects of pattern orientation, spatial phase, and temporal frequency. In T. A. Lawton (Ed.), *Computational Vision based on Neurobiology*, *S. P. I. E. Proceedings*, **2054**, 32 - 42.

Foley, J. M., and Chen, C.-C. (1997) Analysis of the effect of pattern adaptation on pattern pedestal effects: A two process model. *Vision Research*, **37**, 2779 - 2788.

Foley, J. M., and Chen, C.-C. (1999) Pattern detection in the presence of maskers that differ in spatial phase and temporal offset: threshold measurements and a model. *Vision Research*, **39**, 3855 - 3872.

Foley, J. M., and Crites, S. C. (1997) Lateral and superimposed pattern masking: One model describes both. *Investigative Ophthalmology and Visual Science*, **38**, 5.

Foley, J. M., and Schwarz, W. (1998) Spatial attention: effect of position uncertainty and number of distractor patterns on the threshold-versus-contrast function for contrast discrimination. *Journal of the Optical Society of America A*, **15**, 1036 - 1047.

Foley, J. M., and Yang, Y (1991) Forward pattern masking: effects of spatial frequency and contrast. *Journal of the Optical Society of America A*, **8**, 2026 - 2037.

Fredericksen, R. E., and Hess, R. F. (1998) Estimating multiple temporal mechanisms in human vision. *Vision Research*, **38**, 1023 - 1040.

Freeman, T. C. B., Durand, S., Kiper, D. C., and Carandini, M. (2002) Suppression without inhibition in visual cortex. *Neuron*, **35**, 759 - 771.

Fries, W., Albus, K., and Creutzfeldt, O. D. (1977) Effects of interacting visual patterns on single cell responses in cat's striate cortex. *Vision Research*, **17**, 1001 - 1008.

Geisler, W. S., and Albrecht, D. G. (1992) Cortical neurons: isolation of a contrast gain control. *Vision Research*, **32**, 1409 - 1410.

Georgeson, M. A. (1975) Mechanisms of visual image processing: studies on pattern interaction and selective channels in human vision. Doctoral dissertation. University of Sussex.

Georgeson, M. A. (1980) The perceived spatial frequency, contrast and orientation of illusory gratings. *Perception*, **9**, 695 - 712.

Georgeson, M. A. (1985) The effect of spatial adaptation on perceived contrast. *Spatial Vision*, **1**, 103 - 112.

Georgeson, M. A. (1987) Temporal properties of spatial contrast vision. *Vision Research*, **27**, 765 - 780.

Georgeson, M. A. (1988) Spatial phase dependence and the role of motion detection in monocular and dichoptic forward masking. *Vision Research*, **28**, 1193 - 1205.

Georgeson, M. A. (1991) Over the limit: Encoding contrast above threshold in human vision. In Kulikowski, J. J. (Ed.), Vision and visual dysfunction, 5, Cronly-Dillon, J. R. (Series Ed.). London: Macmillan Press.

Georgeson, M. A. (1992) Human vision combines oriented filters to compute edges. *Proceedings of the Royal Society of London B*, **249**, 235 - 245.

Georgeson, M. A. (1998) Edge-finding in human vision: a multi-stage model based on the perceived structure of plaids. *Image and Vision Computing*, **16**, 389 - 405.

Georgeson, M. A., and Georgeson, J. M. (1987) Facilitation of briefly presented gratings: time-course and contrast-dependence. *Vision Research*, **27**, 369 - 379.

Georgeson, M. A., and Harris, M. G. (1984) Spatial selectivity of contrast adaptation: Models and data. *Vision Research*, **24**, 729 - 741.

Georgeson, M. A., and Meese, T. S. (1997) Perception of stationary plaids: The role of spatial features in edge analysis. *Vision Research*, **37**, 3255 - 3271.

Georgeson, M. A., and Shackleton, T. M. (1994) Perceived contrast of gratings and plaids - Nonlinear summation across oriented filters. *Vision Research*, **34**, 1061 - 1075.

Gilbert, C. D., and Wiesel, T. N. (1990) The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Research*, **30**, 1689 - 1701.

Gilinsky, A. (1968) Orientation-specific effects of patterns of adapting light on visual acuity. *Journal of the Optical Society of America*, **58**, 13 - 18.

Gilinsky, A. S., and Doherty, R. S. (1969) Interocular transfer of orientational effects. *Science, N. Y.*, **164**, 454 - 455.

Gilinsky, A. S., and Mayo, T. H. (1971) Inhibitory effects of orientational adaptation. *Journal of the Optical Society of America*, **61**, 1710 - 1714.

Gorea, A. (1987) Masking efficiency as a function of stimulus onset asynchrony for spatial frequency detection and identification. *Spatial Vision*, **2**, 51 - 60.

Gorea, A., and Sagi, D. (2001) Disentangling signal from noise in visual contrast discrimination. *Nature Neuroscience*, **4**, 1146 - 1150.

Graham, N. (1972) Spatial frequency channels in the human visual system: Effects of luminance and pattern drift rate. *Vision Research*, **12**, 53 - 68.

Graham, N. (1989) *Visual Pattern Analyzers*, Oxford University Press, New York.

Graham, N. (1991) Complex channels, early nonlinearities, and normalization in texture segregation. In: *Computational Models of Visual Processing*, (Eds: Landy, M. S. and Movshon, J. A.), M. I. T. Press, Cambridge, Mass., 273 - 290.

Graham, N., and Nachmias, J. (1971) Detection of grating patterns containing two spatial frequencies: A comparison of single-channel and multiple-channel models. *Vision Research*, **11**, 251 - 259.

Graham, N., Robson, J. G., and Nachmias, J. (1978) Grating summation in fovea and periphery. *Vision Research*, **18**, 815 - 825.

Gray, C. M., Engel, A. K., Konig, P., and Singer, W. (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronisation which reflects global stimulus properties. *Nature*, **338**, 334.

Gray, C. M., and McCormick, A. D. (1996) Chattering Cells: Superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science*, **274**, 109 - 113.

Greenlee, M. W., Georgeson, M. A., Magnussen, S., and Harris, J. P. (1991) The time course of adaptation to spatial contrast. *Vision Research*, **31**, 223.

Greenlee, M. W., and Heitger, F. (1988) The functional role of contrast adaptation. *Vision Research*, **28**, 791 - 797.

Greenlee, M. W., Magnussen, S., and Nordby, K. (1988) Spatial vision of the achromat: spatial-frequency and orientation specific adaptation. *Journal of Physiology (London)*, **395**, 661 - 678.

Greis, U., and Roehler, R. (1970) Untersuchungen der subjectiven Detail - Erkenbarkeit mit Hilfe der Orts - frequenzfilterung. *Optica Acta*, **17**, 516 - 526.

Hammett, S. T., and Snowden, R. J. (1995) The effect of contrast adaptation on briefly presented stimuli. *Vision Research*, **35**, 1721 - 1725.

Hammond, P., and Andrews, D. P. (1978) Orientation tuning of cells in areas 17 and 18 of the cat's visual cortex. *Experimental Brain Research*, **31**, 341 - 351.

Hartline, H. K. (1938) The response of single optic nerve fibers of the vertebrate eye to illumination of the retina. *American Journal of Physiology*, **121**, 400 - 415.

Harvey, L. O. Jr., and Doan, V. V. (1990) Visual masking at different polar angles in the two-dimensional Fourier plane. *Journal of the Optical Society of America A*, **7**, 116 - 127.

Harwerth, R. S., and Levi, D. M. (1978) Reaction time as a measure of suprathreshold grating detection. *Vision Research*, **18**, 1579 - 1586.

Heeger, D. J. (1991) Nonlinear models of neural responses in cat visual cortex. In: *Computational Models of Visual Processing*, (Eds: Landy, M. S. and Movshon, J. A.), M. I. T. Press, Cambridge, Mass., 119 - 133.

Heeger, D. J. (1992a) Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, **9**, 181 - 197.

Heeger, D. J. (1992b) Half-squaring in responses of cat striate cells. *Visual Neuroscience*, **9**, 427 - 443.

Heggelund, P., and Hohmann, A. (1976) Long-term retention of the "Gilinsky-Effect". *Vision Research*, **16**, 1015 - 1017.

Henning, G. B. (1988) Spatial-frequency tuning as a function of temporal frequency and stimulus motion. *Journal of the Optical Society of America A*, **5**, 1362 - 1373.

Henning, G. B., Hertz, B. G., and Hinton, J. L. (1981) Effects of different hypothetical detection mechanisms on the shape of spatial-frequency filters inferred from masking experiments. *Journal of the Optical Society of America*, **71**, 574 - 581.

Henry, G. H., Bishop, P. O., Tupper, R. M., and Dreher, B. (1973) Orientation specificity and response variability of cells in the striate cortex. *Vision Research*, **13**, 1771 - 1779.

Henry, G. H., Dreher, B., and Bishop, P. O. (1974) Orientation specificity of cells in cat striate cortex. *Journal of Neurophysiology*, **37**, 1394 - 1409.

Hess, R. F., and Snowden, R. J. (1992) Temporal properties of human visual filters: number, shapes and spatial covariation. *Vision Research*, **32**, 47 - 59.

Holmes, D. J., and Meese, T. S. (1999) Simultaneous masking by plaids and gratings: threshold elevation and threshold facilitation. *Perception*, **28** (suppl.), 88.

Holmes, D. J., and Meese, T. S. (2001) Linear summation for remote masks in a contrast gain pool. *Perception*, **30** (suppl.), 81.

Horton, J. C. (1984) Cytochrome oxidase patches: A new cytoarchitectonic feature of monkey cortex. *Philosophical Transactions of the Royal Society of London (Biology)*, **304**, 199 - 253.

Howell, E. R., and Hess, R. F. (1978) The functional area for summation to threshold for sinusoidal gratings. *Vision Research*, **18**, 369 - 374.

Hubel, D. H., and Wiesel, T. N. (1959) Receptive fields of single neurons in the cat's striate cortex. *Journal of Physiology (London)*, **148**, 574 - 591.

Hubel, D. H., and Wiesel, T. N. (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, **160**, 106 - 154.

Hubel, D. H., and Wiesel, T. N. (1965) Receptive fields and functional architecture in two nonstriate visual areas (18-19) of the cat. Journal of Neurophysiology, 28, 229 - 289.

Hubel, D. H., and Wiesel, T. N. (1968) Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, **195**, 215 - 243.

Hubel, D. H., and Wiesel, T. N. (1974) Sequence regularity and geometry of orientation columns in the monkey striate cortex. *Journal of Comparative Neurology*, **158**, 267 - 294.

Hubel, D. H., and Livingstone, M. S. (1990) Color and contrast sensitivity in the lateral geniculate body and primary visual cortex of the macaque monkey. *Journal of Neuroscience*, **10**, 2223 - 2237.

Ito, M., Westheimer, G., and Gilbert, C. D. (1998) Attention and perceptual learning modulate contextual influences on visual perception. *Neuron*, **20**, 1191 - 1197.

Itti, L., Koch, C., and Braun, J. (1999) A quantitative model relating visual neuronal activity to psychophysical thresholds. *Neurocomputing*, **26 - 27**, 743 - 748.

Itti, L., Koch, C., and Braun, J. (2000) Revisiting spatial vision: toward a unifying model. *Journal of the Optical Society of America A*, **17**, 1899 - 1917.

Iversen, L. L., Mitchell, J. F., and Srinivasan, V. (1971) The release of γ-aminobutyric acid during inhibition in the cat visual cortex. *Journal of Physiology (London)*, **212**, 519 - 534.

Johnson, C. A., Keltner, J. L., and Balestrery, F. (1978) Effects of target size and eccentricity on visual detection and resolution. *Vision Research*, **18**, 1217 - 1222.

Jones, J. P., and Palmer, L. A. (1987a) The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, **58**, 1187 - 1211.

Jones, J. P., and Palmer, L. A. (1987b) An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, **58**, 1233 - 1258.

Jones, J. P., Stepnoski, A., and Palmer, L. A. (1987) The two-dimensional spectral structure of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, **58**, 1212 - 1232.

Jones, R. M., and Tulunay-Keesey, U. (1975) Local retinal adaptation and spatial frequency channels. *Vision Research*, **15**, 1239 - 1244.

Jones, R. M., and Tulunay-Keesey, U. (1980) Phase selectivity of spatial frequency channels. *Journal of the Optical Society of America*, **70**, 66 - 70.

Kapadia, M. K., Ito, M., Gilbert, C. D., and Westheimer, G. (1995) Improvement in visual sensitivity by changes in local context: parallel studies in human observers and in V1 of alert monkeys. *Neurone*, **15**, 843 - 856.

Kaplan, E., and Shapley, R. M. (1982) X-cell and Y-cell in the lateral geniculate nucleus of macaque monkeys. *Journal of Physiology*, **330**, 125 - 143.

Kaplan, E., and Shapley, R. M. (1986) The primate retina contains two types of ganglion cells, with high- and low-contrast sensitivity. *Proceedings of the National Academy of Sciences*, **83**, 2755 - 2757.

Kayser, A., Priebe, N. J., and Miller, K. D. (2001) Contrast-dependent nonlinearities arise locally in a model of contrast-invariant orientation tuning. *Journal of Neurophysiology*, **85**, 2130 - 2149.

Kelly, D. H. (1970) Effects of sharp edges on the visibility of sinusoidal gratings. *Journal of the Optical Society of America*, **60**, 98 - 103.

Kelly, D. H. (1972) Adaptation effects on spatio-temporal sine-wave thresholds. *Vision Research*, **12**, 89 - 101.

Kelly, D. H. (1976) Pattern detection and the two-dimensional Fourier transform: Flickering checkerboards and chromatic mechanisms. *Vision Research*, **16**, 277 - 287.

Kelly, D. H., and Burbeck, C. A. (1980) Motion and vision. III. Stabilized pattern adaptation. *Journal of the Optical Society of America*, **70**, 1283 - 1289.

Kelly, D. H., and Burbeck, C. A. (1984) Critical problems in spatial vision. *CRC Critical Reviews in Biomedical Engineering*, **10**, 125 - 177.

Kelly, D. H., and Burbeck, C. A. (1987) Further evidence for a broadband, isotropic mechanism sensitive to high velocity stimuli. *Vision Research*, **27**, 1527 - 1537.

Kerr, J. L. (1971) Visual resolution in the periphery. *Perception and Psychophysics*, 9, 375 - 378.

Klein, S. A., Hu, Q. J., and Carney, T. (1996) The adjacent pixel nonlinearity: Problems and solutions (Research Note). *Vision Research*, **36**, 3167 - 3181.

Knierim, J. J., and Van Essen, D. C. (1992) Neuronal responses to static texture pattern in area V1 of the alert monkey. *Journal of Neurophysiology*, **67**, 961 - 980.

Koenderink, J. J., and van Doorn, A. J. (1979) Spatiotemporal contrast detection surface is bimodal. *Optics Letters*, **4**, 32 - 34.

Kontsevich, L. L., Chen, C. C., and Tyler, C. W. (2002) Separating the effects of response nonlinearity and internal noise psychophysically. *Vision Research*, **42**, 1771 - 1784.

Kuffler, S. W. (1953) Discharge patterns and functional organisation of mamillian retina. *Journal of Neurophysiology*, **16**, 37 - 68.

Kulikowski, J. J. (1971) Some stimulus parameters affecting spatial and temporal resolution of human vision. *Vision Research*, **11**, 83 - 90.

Kulikowski, J. J. (1972) Orientational selectivity of human binocular and monocular vision revealed by simultaneous and successive masking. *Journal of Physiology (London)*, **226**, 67P - 68P.

Kulikowski, J. J., Abadi, R., and King-Smith, P. E. (1973) Orientation selectivity of grating and line detectors in human vision. *Vision Research*, **13**, 1479 - 1486.

Kulikowski, J. J., and Gorea, A. (1978) Complete adaptation to patterned stimuli: a necessary and sufficient condition for Weber's law for contrast. *Vision Research*, **18**, 1223 - 1227.

Kulikowski, J. J., and King-Smith, P. E. (1973) Spatial arrangement of line, edge and grating detectors revealed by subthreshold summation. *Vision Research*, **13**, 1455 - 1478.

Kulikowski, J. J., and Tolhurst, D. J. (1973) Psychophysical evidence for sustained and transient detectors in human vision. *Journal of Physiology, London*, **232**, 149 - 162.

Kurtenbach, W., and Magnussen, S. (1981) Inhibition, disinhibition, and summation among orientation detectors in human vision. *Experimental Brain Research*, **43**, 193 - 198.

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

Lauritzen, T. Z., Krukowski, A. E., and Miller, K. D. (2001) Local correlation-based circuitry can account for responses to multi-grating stimuli in a model of cat V1. *Journal of Neurophysiology*, **86**, 1803 - 1815.

Lawden, M. C. (1983) An investigation of the ability of the human visual system to encode spatial phase relationships. *Vision Research*, **23**, 1451 - 1463.

Legge, G. E. (1978a) Sustained and transient mechanisms in human vision: temporal and spatial properties. *Vision Research*, **18**, 69 - 81.

Legge, G. E. (1978b) Space domain properties of a spatial frequency channel in human vision. *Vision Research*, **18**, 959 - 969.

Legge, G. E. (1981) A power law for contrast discrimination. *Vision Research*, **21**, 457 - 467.

Legge, G. E. (1984) Binocular contrast summation. 2. Quadratic summation. *Vision Research*, **24**, 385 - 394.

Legge, G. E., and Foley, J. M. (1980) Contrast masking in human vision. *Journal of the Optical Society of America*, **70**, 1458 - 1471.

Lehky, S. R. (1985) Temporal properties of visual channels measured by masking. *Journal of the Optical Society of America A*, **2**, 1260 - 1272.

Lehky, S. R., and Wilson, H. R. (1984) Temporal characteristics of visual channels measured by masking. *Journal of the Optical Society of America A*, **1**, 1273 - 1273.

Lema, S. A., and Blake, R. (1977) Binocular summation in normal and stereoblind humans. *Vision Research*, **17**, 691 - 695.

Lennie, P. (1998) Single units and visual cortical organization. *Perception*, **27**, 889 - 935.

Levick, W. R. (1967) Receptive fields and trigger features of ganglion cells in the visual streak of the rabbits retina. *Journal of Physiology (London)*, **188**, 285 - 307.

Levick, W. R., and Thibos, L. N. (1980) Orientation bias of cat retinal ganglion cells. *Nature (London)*, **268**, 389 - 390.

Levinson, E., and Sekuler, R. (1980) A two-dimensional analysis of direction-specific adaptation. *Vision Research*, **20**, 103 - 107.

Levitt, J. B., and Lund, J. S. (1997) Contrast dependence of contextual effects in primate visual cortex. *Nature*, **387**, 73 - 76.

Li, A., and Aslin, R. N. (1992) Contrast adaptation and perceptual 'filling in' at specific retinal locations. *Investigative Ophthalmology and Visual Science Supplement*, **33**, 1256.

Li, C.- Y., and Creutzfeldt, O. (1984) The representation of contrast and other stimulus parameters by single neurons in area 17 of the cat. *Pflugers Archives*, **401**, 304 - 314.

Li, C. Y., and Li, W. (1994) Extensive integration field beyond the classical receptive field of cat's striate cortical neurons: classification and tuning properties. *Vision Research*, **34**, 2337 - 2355.

Li, W., Thier, P., and Wehrhahn, C. (2000) Contextual influence on orientation discrimination of humans and responses of neurons in V1 of alert monkeys. *Journal of Neurophysiology*, **83**, 941 - 954.

Livingstone, M. S., and Hubel, D. H. (1984) Anatomy and physiology of a colour system in the primate visual cortex. *Journal of Neuroscience*, **4**, 309 - 356.

Lupp, U., Hauske, G., and Wolf, W. (1976) Perceptual latencies to sinusoidal gratings. *Vision Research*, **16**, 969 - 972.

Lyons, N. P., and Farrell, J. E. (1989) Linear system analysis of CRT displays. *Society of information diplay 1989 digest* (Vol. 20, pp. 220 - 223).

Maattanen, L. M., and Koenderink, J. J. (1991) Contrast adaptation and contrast gain control. *Experimental Brain Research*, **87**, 205 - 212.

Maffei, L., and Fiorentini, A. (1973) The visual cortex as a spatial frequency analyser. *Vision Research*, **13**, 1255 - 1268.

Maffei, L., and Fiorentini, A. (1976) The unresponsive regions of visual cortical receptive fields. *Vision Research*, **16**, 1131 - 1139.

Maffei, L., Fiorentini, A., and Bisti, S. (1973) Neural correlate of perceptual adaptation to gratings. *Science*, **182**, 1036 - 1038.

Maffei, L., Morrone, C., Pirchio, M., and Sandini, G. (1979) Responses of visual cortical cells to periodic and non-periodic stimuli. *Journal of Physiology*, **296**, 24 - 47.

Magnussen, S., and Greenlee, M. W. (1985) Marathon adaptation to spatial contrast: Saturation in sight. *Vision Research*, **25**, 1409 - 1411.

Magnussen, S., and Greenlee, M. W. (1986) Contrast threshold elevation following continuous and interrupted adaptation. *Vision Research*, **26**, 673 - 675.

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

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

Marg, E., Adams, J. E., and Rutkin, B. (1968) Receptive fields of cells in the human visual cortex. *Experientia (Basel)*, **24**, 348 - 350.

Marr, D. (1982) *Vision*. Freeman: San Francisco.

Martinez, L. M., Alonso, J. M., Reid, R. C., and Hirsch, J. A. (2002) Laminar processing of stimulus orientation in cat visual cortex. *Journal of Physiology (London)*, **540**, 321 - 333.

Maturana, H. R., and Frenk, S. (1963) Directional movement and horizontal edge detectors in the pigeon retina. *Science*, **142**, 977 - 979.

Maudarbocus, A. Y., and Ruddock, K. H. (1973) Nonlinearity of visual signals in relation to shape-sensitive adaptation responses. *Vision Research*, **13**, 1713 - 1737.

Mayo, T. H., Gilinsky, A. S., and Jochnowitz, A. (1968) Visual masking and tilt aftereffects as a function of tilt and angular separation between test and adaptation patterns. *Proceedings of the 76th Annual Convention, American Psychological Association*, **3**, 97 - 98.

McKee, S. P., Klein, S. A., and Teller, D. Y. (1985) Statistical properties of forced-choice psychometric functions: implications for probit analysis. *Perception and Psychophysics*, **37**, 286 - 298.

Meese, T. S. (1993) *Feature coding in human pattern vision*. Unpublished doctoral dissertation, University of Bristol, England.

Meese, T. S. (1995) Using the standard staircase to measure the point of subjective equality: a guide based on computer simulations. *Perception and Psychophysics*, **57**, 267 - 281.

Meese, T. S., and Georgeson, M. A. (1991) Edge computation in human vision: what controls the combining of oriented filters? *Perception*, **20**, 81.

Meese, T. S., and Holmes, D. J. (2002) Adaptation and gain pool summation: alternative models and masking data. *Vision Research*, **42**, 1113 - 1125.

Meese, T. S., and Holmes, D. J. (2003) Orientation masking: suppression and mechanism bandwidth. *Perception*, **32**, 388.

Meese, T. S., Holmes, D. J., Williams, C. B., and Hess, R. F. (2000) Contrast discrimination and summation. *Perception*, **29**, 121.

Meese, T. S., and Williams, C. B. (2000) Probability summation for multiple patches of luminance modulation. *Vision Research*, **40**, 2101 - 2113.

Meier, L., and Carandini, M. (2002) Masking by fast gratings. *Journal of Vision*, **2**, 293 - 301.

Merigan, W. H., Katz, L. M., and Maunsell, J. H. R. (1991) The effects of parvocellular lateral geniculate lesions on the acuity and contrast sensitivity of macaque monkeys. *Journal of Neuroscience*, **11**, 994 - 1001.

Merigan, W. H., and Maunsell, J. H. R. (1990) Macaque vision after magnocellular lateral geniculate lesions. *Visual Neuroscience*, **5**, 347 - 352.

Michael, C. R. (1972) Visual receptive fields of single neurones in the superior colliculus of the ground squirrel. *Journal of Neurophysiology*, **35**, 815 - 832.

Milner, A. D., and Goodale, M. A. (1995) *The Visual Brain in Action*, Oxford University Press, New York.

Morrone, M. C., Burr, D. C., and Speed, H. D. (1987) Cross-orientation inhibition in cat is GABA mediated. *Experimental Brain Research*, **67**, 635 - 644.

Morrone, M. C., Burr, D. C., and Maffei, L. (1982) Functional implications of cross-orientation inhibition of cortical visual cells. I. Neurophysiological evidence. *Proceedings of the Royal Society of London B*, **216**, 335 - 354.

Moulden, B., Renshaw, J., and Mather, G. (1984) Two channels for flicker in the human visual system. *Perception*, **13**, 387 - 400.

Mountcastle, V. B. (1997) The columnar organization of the neocortex. *Brain*, **120**, 701 - 722.

Movshon, J. A., and Blakemore, C. B. (1973) Orientation specificity and spatial selectivity in human vision. *Perception*, **2**, 53 - 60.

Movshon, J. A., and Lennie, P. (1979) Pattern-selective adaptation in visual cortical neurons. *Nature*, **278**, 850 - 852.

Movshon, J. A., Thompson, I. D., and Tolhurst, D. J. (1978a) Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat's visual cortex. *Journal of Physiology (London)*, **283**, 101 - 120.

Movshon, J. A., Thompson, I. D., and Tolhurst, D. J. (1978) Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *Journal of Physiology*, **283**, 53 - 77.

Mukherjee, P., and Kaplan, E. (1995) Dynamics of neurons in the cat lateral geniculate nucleus: in vivo electrophysiology and computational modeling. *Journal of Neurophysiology*, **74**, 1222 - 1243.

Mullen, K. T., and Losada, M. A. (1994) Evidence for separate pathways for colour and luminance detection mechanisms. *Journal of the Optical Society of America A*, **11**, 3136 - 3151.

Mulligan, J. B., and Stone, L. S. (1989) Halftoning method for the generation of motion stimuli. *Journal of the Optical Society of America A*, 6, 1217 - 1227.

Nachmias, J. (1993) Masked detection of gratings: the standard model revisited. *Vision Research*, **33**, 1359 - 1365.

Nachmias, J., and Rogowitz, B. E. (1983) Masking by spatially-modulated gratings. *Vision Research*, **23**, 1621 - 1629.

Nachmias, J., and Sansbury, R. V. (1974) Grating contrast: Discrimination may be better than detection. *Vision Research*, **14**, 1039 - 1042.

Nachmias, J., Sansbury, R., Vassilev, A., and Weber, A. (1973) Adaptation to square-wave gratings: in search of the elusive third harmonic. *Vision Research*, **13**, 1335 - 1342.

Nachmias, J., and Weber, A. (1975) Discrimination of simple and complex gratings. *Vision Research*, **15**, 217 - 223.

Naiman, A. C., and Makous, W. (1992) Spatial non-linearities of grayscale CRT pixels. *SPIE Proceedings*, **1666**, 41 - 56.

Naka, K. I., and Rushton, W. A. H. (1966) S-potentials from colour units in the retina of fish (cyprinidae). *Journal of Physiology (London)*, **185**, 536 - 555.

Nelson, J. I., Kupersmith, W. H., and Carr, R. E. (1984) A rapid evoked potential index of cortical adaptation. *Electoencephalography and Clinical Neurophysiology*, **59**, 454 - 464.

Ohtani, Y., Okamura, S., Yoshida, Y., Toyama, K., and Ejima, Y. (2002) Surround suppression in the human visual cortex: an analysis using magnetoencephalography. *Vision Research*, **42**, 1825 - 1835.

Ohzawa, I., Sclar, G., and Freeman, R. D. (1982) Contrast gain control in the cat visual cortex. *Nature*, **298**, 266 - 268.

Ohzawa, I., Sclar, G., and Freeman, R. D. (1985) Contrast gain control in the cat visual system. *Journal of Neurophysiology*, **54**, 651 - 667.

Olzak, L. A., and Thomas, J. P. (1986) Seeing spatial patterns. In: *Handbook of Perception and Human Performance*. Eds.: Boff, K., Kaufman, L., and Thomas, J. P.. New York: J. Wiley.

Olzak, L. A., and Thomas, J. P. (1991) When orthogonal orientations are not processed independently. *Vision Research*, **31**, 51 - 57.

Olzak, L. A., and Thomas, J. P. (1992a) Configural effects constrain fourier models of pattern discrimination. *Vision Research*, **32**, 1885 - 1898.

Olzak, L. A., and Thomas, J. P. (1992b) Linear cue summation subserves some compound discriminations. *Investigative Ophthalmology and Visual Science*, **33**, 1343.

Olzak, L. A., and Thomas, J. P. (1999) Neural recoding in human pattern vision: model and mechanisms. Vision Research, **39**, 231 - 256.

Pantle, A. J. (1973) Visual effects of sinusoidal components of complex gratings: independent or additive? *Vision Research*, **13**, 2195 - 2204.

Pantle, A. J. (1983) Temporal determinants of spatial sine-wave masking. *Vision Research*, **23**, 749 - 757.

Pantle, A., and Sekuler, R. (1968) Size-detecting mechanisms in human vision. *Science*, **162**, 1146 - 1148.

Passaglia, C. L., Troy, J. B., Ruttiger, L., and Lee, B. B. (2002) Orientation sensitivity of ganglion cells in primate retina. *Vision Research*, **42**, 683 - 694.

Patterson, R., and Becker, S. (1996) Direction-selective adaptation and simultaneous contrast induced by stereoscopic (Cyclopean) motion. *Vision Research*, **36**, 1773 - 1781.

Pelli, D. G. (1985) Uncertainty explains many aspects of visual contrast detection and discrimination. *Journal of the Optical Society of America A*, **2**, 1508 - 1532.

Phillips, G. C., and Wilson, H. R. (1984) Orientation bandwidths of spatial mechanisms measured by masking. *Journal of the Optical Society of America A*, **1**, 226 - 232.

Poggio, G. F. (1972) Spatial properties of neurons in striate cortex of unanesthetized macaque monkey. *Investigative Ophthalmology and Visual Science*, **11**, 368 - 377.

Polat, U., Mizobe, K., Pettet, M. W., Kasamatsu, T., and Norcia, A. M. (1998) Collinear stimuli regulate visual responses depending on cell's contrast threshold. *Nature*, **391**, 580 - 584.

Polat, U., and Norcia, A. M. (1996) Neurophysiological evidence for contrast dependent long-range facilitation and suppression in the human visual cortex. *Vision Research*, **36**, 2099 - 2109.

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

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

Pollen, D. A., and Ronner, S. F. (1982) Spatial computation performed by simple and complex cells in the visual cortex of the cat. *Vision Research*, **22**, 101 - 118.

Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. (1989) *Numerical Recipes in Pascal: The Art of Scientific Computing.* Cambridge: Cambridge University Press.

Quick, R. F. (1974) A vector-magnitude model of contrast detection. *Kybernetik*, **16**, 65 - 67.

Quick, R. F., and Reichart, T. A. (1975) Spatial frequency selectivity in contrast detection. *Vision Research*, **15**, 637 - 643.

Rabiner, L. R., and Gold, B. (1975) *Theory and Applications of Digital Signal Processing.* Prentice - Hall, Englewood Cliffs.

Regan, D., and Beverley, K. I. (1983) Spatial-frequency discrimination and detection: comparison of postadaptation thresholds. *Journal of the Optical Society of America*, **73**, 1684 - 1690.

Regan, D., and Beverley, K. I. (1985) Postadaptation orientation discrimination. *Journal of the Optical Society of America (A)*, **2**, 147 - 155.

Regan, M. P., and Regan, D. (1987) Orientation tuning and spatial frequency tuning of a phase - independent brain response. *Investigative Ophthalmology and Visual Science*, **28**, S127 [A].

Ringach, D. L., Shapley, R. M., and Hawken, M. J. (2002) Orientation selectivity in Macaque V1: Diversity and laminar dependence. *Journal of Neuroscience*, **22**, 5639 - 5651.

Robson, J. G. (1966) Spatial and temporal contrast sensitivity functions of the visual system. *Journal of the Optical Society of America*, **56**, 1141 - 1142.

Robson, J. G. (1988) Linear and nonlinear operations in the visual system. *Investigative Ophthalmology and Visual Science* (Suppl.), **29**, 117.

Robson, J. G., and Graham, N. (1981) Probability summation and regional variation in contrast sensitivity across the visual field. *Vision Research*, **21**, 409 - 418.

Rose, D., and Blakemore, C. (1974) An analysis of orientation selectivity in the cat's visual cortex. *Experimental Brain Research*, **20**, 1 - 17.

Rose, D., and Evans, R. (1983) Evidence against saturation of contrast adaptation in the human visual system. *Perception and Psychophysics*, **34**, 158 - 160.

Rose, D., and Lowe, I. (1982) Dynamics of adaptation to contrast. *Perception*, **11**, 505 - 528.

Ross, J. (1992) How adaptation and masking alter contrast gain perception. *Perception*, **21**, Supplement 2, 60.

Ross, J., and Speed, H. D. (1991) Contrast adaptation and contrast masking in human vision. *Proceedings of the Royal Society of London B*, **246**, 61 - 69.

Ross, J., and Speed, H. D. (1996) Perceived contrast following adaptation to gratings of different orientations. *Vision Research*, **36**, 1811 - 1818.

Ross, J., Speed, H. D., and Morgan, M. J. (1993) The effects of adaptation and masking on incremental thresholds for contrast. *Vision Research*, **33**, 2051 - 2056.

Rovamo, J., Virsu, V., Laurinen, P., and Hyvarinen, L. (1982) Resolution of gratings oriented along and across meridians in peripheral vision. *Investigative Ophthalmology and Visual Science*, **23**, 666 - 670.

Sachs, M. B., Nachmias, J., and Robson, J. G. (1971) Spatial-frequency channels in human vision. *Journal of the Optical Society of America*, **61**, 1176 - 1186.

Sagi, D., and Hochstein, S. (1985) Lateral inhibition between spatially adjacent spatial frequency channels? *Perception and Psychophysics*, **37**, 315 - 322.

Sanchez-Vives, M. V., Nowak, L. G., and McCormick, D. A. (2000a) Membrane mechanisms underlying contrast adaptation in cat area 17 *in vivo*. *Journal of Neuroscience*, **20**, 4267 - 4285.

Sanchez-Vives, M. V., Nowak, L. G., and McCormick, D. A. (2000b) Cellular mechanisms of long-lasting adaptation in visual cortical neurons *in vitro*. *Journal of Neuroscience*, **20**, 4286 - 4299.

Sato, H., Katsuyama, N., Tamura, H., Hata, Y., and Tsumoto, T. (1996) Mechanisms underlying orientation selectivity of neurons in the primary visual cortex of the macaque. *Journal of Physiology*, **494**, 757 - 771.

Saul, A. B., and Cynader, M. S. (1989a) Adaptation in single units in visual cortex: The tuning of aftereffects in the temporal domain. *Visual Neuroscience*, **2**, 609 - 620.

Saul, A. B., and Cynader, M. S. (1989b) Adaptation in single units in visual cortex: The tuning of aftereffects in the spatial domain. *Visual Neuroscience*, **2**, 593 - 607.

Saul, A. B., and Humphrey, A. L. (1990) Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate nucleus. *Journal of Neurophysiology*, **64**, 206 - 224.

Saul, A. B., and Humphrey, A. L. (1992) Temporal-frequency tuning of direction selectivity in cat visual cortex. *Visual Neuroscience*, **8**, 365 - 372.

Scheibel, M. E., and Scheibel, A. B. (1970) Elementary processes in selected thalamic and cortical subsystems: the structural substrates. In: *The Neurosciences: Second Study Program*, (Ed. Schmitt, F. O.), New York: Rockefeller University Press., 443 - 457.

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

Schiller, P. H., Logothetis, N. K., and Charles, E. R. (1990a) Role of the color-opponent and broad-band channels in vision. *Visual Neuroscience*, **5**, 321 - 346.

Schiller, P. H., Logothetis, N. K., and Charles, E. R. (1990b) Functions of the colour-opponent and broad-band channels of the visual system. *Nature*, **343**, 68 - 70.

Schofield, A. J., and Georgeson, M. A. (2003) Sensitivity to contrast modulation: the spatial frequency dependence of second-order vision. *Vision Research*, **43**, 243 - 259.

Schrater, P. R., and Simoncelli, E. P. (1998) Local velocity representation: evidence from motion adaptation. *Vision Research*, **38**, 3899 - 3912.

Schwartz, O., and Simoncelli, E. P. (2001) Natural signal statistics and sensory gain control. *Nature Neuroscience*, 4, 819 - 825.

Sclar, G., and Freeman, R. D. (1982) Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Experimental Brain Research*, **46**, 457 - 461.

Sclar, G., Maunsell, J. H. R., and Lennie, P. (1990) Coding of image contrast in central visual pathways of the macaque monkey. *Vision Research*, **30**, 1 - 10.

Sekuler, R. W. (1965) Spatial and temporal determinants of visual backward masking. *Journal of Experimental Psychology,* **70**, 401 - 406.

Sekuler, R. W. (1974) Spatial vision. *Annual Review of Psychology*, **25**, 195 - 232.

Sekuler, R. W., and Pantle, A. (1967) A model for after-effects of seen movement. *Vision Research*, 7, 427 - 439.

Sekuler, R. W., Rubin, E. L., and Cushman, W. H. (1968) Selectivities of human visual mechanisms for direction of movement and contour orientation. *Journal of the Optical Society of America*, **58**, 1146 - 1150.

Sengpiel, F., Blakemore, C., and Harrad, R. (1995) Interocular suppression in the primary visual cortex: a possible neural basis of binocular rivalry. *Vision Research*, **35**, 179 - 196.

Sengpiel, F., Baddeley, R. J., Freeman, T. C., Harrad, R., and Blakemore, C. (1998) Different mechanisms underly three inhibitory phenomena in cat area 17. *Vision Research*, **38**, 2067 - 2080.

Sengpiel, F., and Blakemore, C. (1994) Interocular control of neuronal responsiveness in cat visual cortex. *Nature*, **368**, 847 - 850.

Sharpe, C. R., and Tolhurst, D. J. (1973) The effects of temporal modulation on the orientation channels of the human visual system. *Perception*, **2**, 23 - 29.

Shou, T., Li, X., Zhou, Y., and Hu, B. (1996) Adaptation of visually evoked responses of relay cells in the dorsal lateral geniculate nucleus of the cat following prolonged exposure to drifting gratings. *Visual Neuroscience*, **13**, 605 - 613.

Sillito, A. M. (1975a) The effectiveness of bicuculline as an antagonist of GABA and visually evoked inhibition in the cat's striate cortex. *Journal of Physiology (London)*, **250**, 287 - 304.

Sillito, A. M. (1975b) The contribution of inhibitory mechanisms to the receptive field properties of neurons in the striate cortex of the cat. *Journal of Physiology*, **250**, 305 - 329.

Sillito, A. M. (1977) The spatial extent of excitatory and inhibitory zones in the receptive field of superficial layer hypercomplex cells . *Journal of Physiology*, **273**, 791 - 803.

Sillito, A. M. (1979) Inhibitory mechanisms influencing complex cell orientation selectivity and their modification at high resting discharge levels. *Journal of Physiology*, **289**, 33 - 53.

Sillito, A. M., Grieve, K. L., Jones, H. E., Cudeiro, J., and Davis, J. (1995) Visual cortical mechanisms detecting focal orientation discontinuities. *Nature*, **378**, 492 - 496.

Silverman, M. S., Grosof, D. H., DeValois, R. L., and Elfar, S. D. (1989) Spatial-frequency organization in primate striate cortex. *Proceedings of the National Academy of Science, U. S. A.*, **86**, 711 - 715.

Smirnakis, S. M., Berry, M. J., Warland, D. K., Bialek, W., and Meister, M. (1997) Adaptation of retinal processing to image contrast and spatial scale. *Nature*, **386**, 69 - 73.

Smith, M. A., Bair, W., Cavanaugh, J. R., and Movshon, J. A. (2001) Latency of inhibition from inside and outside the classical receptive field in macaque V1 neurons [Abstract]. Journal of Vision, 1(3), 35a, http://journalofvision.org/1/3/35, DOI 10.1167/1.3.35.

Snowden, R. J. (1991) Measurement of visual channels by contrast adaptation. *Proceedings of the Royal Society of London B*, **246**, 53 - 59.

Snowden, R. J. (1992) Orientation bandwidth: the effect of spatial and temporal frequency. *Vision Research*, **32**, 1965 - 1974.

Snowden, R. J. (1994) Adaptability of the visual system is inversely related to its sensitivity. *Journal of the Optical Society of America A*, **11**, 25 - 32.

Snowden, R. J. (2002) Interactions between visual stimuli across the visual field. *Perception*, **31**, 382.

Snowden, R. J., and Hammett, S. T. (1992) Subtractive and divisive adaptation in the human visual system. *Nature*, **355**, 248 - 250.

Snowden, R. J., and Hammett, S. T. (1996) Spatial frequency adaptation: threshold elevation and perceived contrast. *Vision Research*, **36**, 1797 - 1809.

Snowden, R. J., and Hammett, S. T. (1998) The effects of surround contrast on contrast thresholds, perceived contrast and contrast discrimination. *Vision Research*, **38**, 1935 - 1945.

Solomon, J. A., Watson, A. B., and Morgan, M. J. (1999) Transducer model produces facilitation from opposite-sign flanks. *Vision Research*, **39**, 987 - 992.

Somers, D.C., Nelson, S. B., and Sur, M. (1995) An emergent model of orientation selectivity in cat visual cortical simple cells. *Journal of Neuroscience*, **15**, 5448 - 5465.

Sompolinsky, H., and Shapley, R. (1997) New perspectives on the mechanisms for orientation selectivity. *Current Opinion in Neurobiology*, **7**, 514 - 522.

Speed, H. D., Morrone, M. C., and Burr, D. C. (1991) Effects of monocular deprivation on the development of visual inhibitory interactions in kittens. *Visual Neuroscience*, **74**, 335 - 343.

Speed, H. D., Ross, J., Morrone, M. C., and Burr, D. C. (1991) Responsivity to contrast under conditions of masking and adaptation. *Perception*, **20**, 121 - 122.

Spillman, L., and Werner, J. S. (Eds.)(1990)*Visual Perception: the neurophysiological foundations.* Academic Press Ltd.

Stecher, S., Sigel, C., and Lange, R. V. (1973) Composite adaptation and spatial frequency interactions. *Vision Research*, **13**, 2527 - 2531.

Stromeyer III, C. F., and Julesz, B. (1972) Spatial frequency masking in vision: Critical bands and spread of masking. *Journal of the Optical Society of America*, **62**, 1221 - 1232.

Stromeyer III, C. F., and Klein, S. (1974) Spatial frequency channels in human vision as asymmetric (edge) mechanisms. *Vision Research*, **14**, 1409 - 1420.

Stromeyer III, C. F., and Klein, S. (1975) Evidence against narrow-band spatial frequency channels in human vision: the detectability of frequency modulated gratings. *Vision Research*, **15**, 899 - 910.

Swift, D. J., and Smith, R. A. (1982) An action spectrum for spatial frequency adaptation. *Vision Research*, **22**, 235 - 246.

Swift, D. J., and Smith, R. A. (1983) Spatial frequency masking and weber's law. *Vision Research*, **23**, 495 - 505.

Teo, P. C., and Heeger, D. J. (1994) Perceptual image distortion. *Human Vision, Visual Processing and Digital Display V, IS and T/SPIE's Symposium of Electronic Imaging: Science and Technology.*

Thomas, J. P., and Gille, J. (1979) Bandwidth of orientation detectors in human vision. *Journal of the Optical Society of America*, **69**, 652 - 660.

Thomas, J. P., Padilla, G. J., and Rourke, D. L. (1969) Spatial interactions in identification and detection of compound visual stimuli. *Vision Research*, **9**, 283 - 292.

Thomas, J. P., and Shimamura, K. K. (1975) Inhibitory interaction between visual pathways tuned to different orientations. *Vision Research*, **15**, 1373 - 1380.

Tolhurst, D. J. (1972) Adaptation to square-wave gratings: Inhibition between spatial frequency channels in the human visual system. *Journal of Physiology*, **226**, 231 - 248.

Tolhurst, D. J. (1973) Separate channels for the analysis of the shape and the movement of a moving visual stimulus. *Journal of Physiology (London)*, **231**, 385 - 402.

Tolhurst, D. J., and Barfield, L. P. (1978) Interactions between spatial frequency channels. *Vision Research*, **18**, 951 - 958.

Tolhurst, D. J., and Dealy, R. S. (1975) The detection and identification of lines and edges. *Vision Research*, **15**, 1367 - 1372.

Tolhurst, D. J., Movshon, J. A., and Dean, A. F. (1983) The statistical reliability of signals in single neurones in cat and monkey visual cortex. *Vision Research*, **23**, 775 - 785.

Tolhurst, D. J., Movshon, J. A., and Thompson, I. D. (1981) The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Experimental Brain Research*, **41**, 414 - 419.

Tolhurst, D. J., and Thompson, I. D. (1981) On the variety of spatial frequency selectivities shown by neurons in area 17 of the cat. *Proceedings of the Royal Society of London, Series B*, **213**, 183 - 199.

Tootell, R. B. H., Hamilton, S. L., and Switkes, E. (1988) Functional anatomy of macaque striate cortex IV. Contrast and magno-parvo streams. *Journal of Neuroscience*, **8**, 1594 - 1609.

Tootell, R. B. H., Silverman, M. S., Hamilton, S. L., Switkes, E., and DeValois, R. L. (1988) Functional anatomy of macaque striate cortex. V. Spatial frequency. *Journal of Neuroscience*, **8**, 1610 - 1624.

Tulunay-Keesey, U. (1972) Flicker and pattern detection: a comparison of thresholds. *Journal of the Optical Society of America*, **62**, 446 - 448.

Ungerleider, L. G., and Mishkin, M. (1982) Two cortical visual systems. In Ingle, D. J., Goodale, M. A. and Mansfield, R. J. W. (Eds), *Analysis of Visual Behavior* (pp. 549 - 586). Cambridge, Mass.: MIT Press.

Valeton, J. M., and Watson, A. B. (1990) Contrast detection does not have a local spatial scale. *Investigative Ophthalmology and Visual Science (Supplement)*, **31**, 428.

Van Nes, F. L., Koenderick, J. J., Nas, H., and Bouman, M. A. (1967) Spatiotemporal modulation transfer in the human eye. *Journal of the Optical Society of America*, **57**, 1082 - 1088.

Vassilev, A., and Mitov, D. (1976) Perception time and spatial frequency. *Vision Research*, **16**, 89 - 92.

Vautin, R. G., and Berkley, M. A. (1977) Responses of single cells in cat visual cortex to prolonged stimulus movement: neural correlates of visual aftereffects. *Journal of Neurophysiology*, **40**, 1051 - 1065.

Vidyasagar, T. R. (1990) Pattern adaptation in cat visual cortex is a cooperative phenomenon. *Neuroscience*, **36**, 175 - 179.

Vidyasagar, T. R., Pei, X., and Volgushev, M. (1996) Multiple mechanisms underlying the orientation selectivity of visual cortical neurones. *Trends in Neuroscience*, **19**, 272 - 277.

Walker, G. A., Ohzawa, I., and Freeman, R. D. (1998) Binocular cross-orientation suppression in the cat's striate cortex. *Journal of Neurophysiology*, **79**, 227 - 239.

Walker, G. A., Ohzawa, I., and Freeman, R. D. (1999) Asymmetric suppression outside the classical receptive field of the visual cortex. *Journal of Neuroscience*, **19**, 10536 - 10553.

Watson, A. B. (1986) Temporal sensitivity. In: *Handbook of Perception and Human Performance*, **1**. (Eds: Boff, K., Kaufman, L. and Thomas, J.) Wiley, New York.

Watson, A. B., and Robson, J. G. (1981) Discrimination at threshold: Labelled detectors in human vision. *Vision Research*, **21**, 1115 - 1122.

Watson, A. B., and Solomon, J. A. (1997) Model of visual contrast gain control and pattern masking. *Journal of the Optical Society of America A*, **14**, 2379 - 2391.

Webster, M. A., and Mollon, J. D. (1991) Changes in color appearance following post-receptoral adaptation. *Nature*, **349**, 235 - 238.

Webster, M. A., and Mollon, J. D. (1992) Adaptation dissociates different measures of equiluminance. *Investigative Ophthalmology and Visual Science*, **33**, 700.

Welch, L. (1989) The perception of moving plaids reveals two motion processing stages. *Nature*, **337**, 734 - 736.

Westheimer, G., and Gee, A. (2002) Orthogonal adaptation and orientation discrimination. *Vision Research*, **42**, 2339 - 2343.

Wetherill, G. B., and Levitt, H. (1965) Sequential estimation of points on a psychometric function. *British Journal of Mathematical and Statistical Psychology*, **18**, 1 - 10.

Wilkins, A. J., Darby, C. E., Stefansson, S. F., Jeavons, P. M., and Harding, G. F. A. (1979) Television epilepsy: the role of pattern. *Electroencephalography and Clinical Neurophysiology*, **47**, 163 - 171.

Wilkins, A. J., Nimmo-Smith, I., Tait, A., McManus, C., Della Sala, S., Tilley, A., Arnold, K., Barrie, M., and Scott, S. (1984) A neurological basis for visual discomfort. *Brain*, **107**, 989 - 1017.

Williams, D. W., and Wilson, H. R. (1983) Spatial-frequency adaptation affects spatial-probability summation. *Journal of the Optical Society of America*, **73**, 1367 - 1371.

Wilson, H. R. (1979) A four mechanism model for threshold spatial vision. *Vision Research*, **19**, 19 - 32.

Wilson, H. R. (1980a) A transducer function for threshold and suprathreshold human vision. *Biological Cybernetics*, **38**, 171 - 178.

Wilson, H. R. (1980b) Spatiotemporal characterisation of a transient mechanism in the human visual system. *Vision Research*, **20**, 443 - 452.

Wilson, H. R. (1983) Psychophysical evidence for spatial channels. In O. J. Braddick & A. C. Sleigh (Eds.), *Physical and Biological Processing of Images* (pp. 88 - 99). Berlin: Springer.

Wilson, H. R., and Humanski, R. (1993) Spatial frequency adaptation and contrast gain control. *Vision Research*, **33**, 1133 - 1149.

Wilson, H. R., and Kim, J. (1998) Dynamics of a divisive gain control in human vision. *Vision Research*, **38**, 2735 - 2741.

Wilson, H. R., McFarlane, D. K., and Phillips, G. C. (1983) Spatial frequency tuning of orientation selective units estimated by oblique masking. *Vision Research*, **23**, 873 - 882.

Xing, J., and Heeger, D. J. (2000) Center-surround interactions in foveal and peripheral vision. *Vision Research*, **40**, 3065 - 3072.

Xing, J., and Heeger, D. J. (2001) Measurement and modeling of center-surround suppression and enhancement. *Vision Research*, **41**, 571 - 583.

Yu, C., and Levi, D. M. (1997) Spatial facilitation predicted with end-stopped spatial filters. *Vision Research*, **37**, 3117 - 3128.

Yu, C., and Levi, D. M. (2000) Surround modulation in human vision unmasked by masking experiments. *Nature Neuroscience*, **3**, 724 - 728.

Zenger, B., and Sagi, D. (1996) Isolating excitatory and inhibitory nonlinear spatial interactions involved in contrast detection. *Vision Research*, **36**, 2497 - 2513.

Zipser, K., Lamme, V. A. F., and Schiller, P. H. (1996) Contextual modulation in primary visual cortex. *Journal of Neuroscience*, **15**, 7376 - 7389.

Appendix One: Some pilot work and ideas on lateral masking

A1.1 Chapter Introduction

In most of the experiments for this thesis, remote masking effects were measured. The masks differed from the tests in orientation or spatial frequency, or both of these characteristics. Some of the effects have proven to be very useful in testing and developing models of cross-channel masking (particularly see chapters 4 and 5). In these experiments, tests and masks were always spatially superimposed. There is a growing body of literature, which shows that masks spatially adjacent to tests can produce substantial effects on their detectability. An obvious direction to take the work developed in the main body of the thesis is to study remote masking when tests and masks are spatially adjacent. The term 'lateral masking' will be used to describe effects from spatially adjacent masks. Some pilot experiments and preliminary ideas on lateral masking are presented here.

A brief review of some lateral masking effects

There is a great deal of evidence indicating that patterns spatially adjacent to others can affect the visual systems response to them. This is the case even when a pattern is presented outside the classical receptive fields (CRFs) that are excited by the other pattern. Many physiological studies of this 'contextual modulation' (Zipser et al., 1996), have shown facilitatory or inhibitory effects, (for example, see Blakemore & Tobin, 1972; Maffei & Fiorentini, 1976; Sillito, 1977; Fries et al., 1977; Allman et al., 1985; Gilbert & Wiesel, 1990; Knierim & Van Essen, 1992; DeAngelis et al., 1994; Li & Li, 1994; Sillito et al., 1995; Polat & Norcia, 1996; Levitt & Lund, 1997; Fitzpatrick, 2000). The effects are dependent on characteristics of the surround, such as its location (Walker et al., 1999), contrast (Levitt & Lund, 1997; Polat et al., 1998), orientation (Sillito et al., 1995; Polat et al., 1998) and spatial frequency (DeValois & Tootell, 1983; Bauman & Bonds, 1991).

Lateral masking phenomena have also been addressed by many psychophysical studies. For example, adjacent (surround) patterns having a higher contrast than a test patch have been shown to lower the test's perceived contrast (Ejima & Takahashi, 1985; Sagi & Hochstein, 1985; Chubb et al., 1989; Cannon & Fullenkamp, 1991, 1993; Snowden & Hammett, 1998), whilst when the test is equal to or higher in contrast than its surround, the perceptual suppression is reduced or abolished (Cannon & Fullenkamp, 1993; Snowden & Hammett, 1998).

Surround interactions are important because they are likely to contribute to visual grouping (Field et al., 1993; Ben-Av & Sagi, 1995; Itto et al., 1998), texture segmentation (Malik & Perona, 1990) and contrast sensitivity modulation (for example, see Chen & Tyler, 2002).

A1.2 (Pilot) Experiment 1. Lateral masking irrespective of mask orientation - as long as the mask is spatially close to the centre of the test

Introduction

Remote masking effects described in the main body of this thesis have allowed cross-channel models of masking to be tested and developed. Such models (see chapters 4 and 5) describe the mechanism excited by the test as being divisively inhibited by a pool in which test and mask inputs are summed. Remote masking effects have helped to determine the nature of this summation. An effect of particular interest was revealed when (spatially superimposed) tests and masks were remote in spatial frequency and orientation. The masking levels produced by a pattern with two (orthogonal) components appeared to be dependent upon the linear summation of the individual mask component contrasts. This was consistent with masking models whose pools summed mask components before passing them through a nonlinearity, rather than those that passed them through a nonlinearity before summation.

The pilot work presented here tested whether it would be plausible to measure contrast detection thresholds to achieve an insight into rules of summation of lateral masks. The rationale is discussed in further detail below. (Note that this study was not performed because of the reasons highlighted in this section's discussion). Following this, the pilot study is discussed.

Rationale

Consider a circular test patch surrounded with a circular annulus mask. The mask is segmented into four, with upper-left, lower-left, upper-right and lower-right portions. The experiment would have four conditions, which determined the mask type. Masks would be either (i) absent (no-mask), (ii) in the top-right and bottom-left quadrants, (iii) in the top-left and bottom-right quadrants (iv) in all four quadrants. With equal mask contrasts, comparing the levels of masking produced by the 'four quadrant' (iv) and 'two quadrant' (ii and iii) masks might have revealed something about summation rules in lateral masking. For example, with matched mask contrasts, if (ii) and (iii) produced the same level of masking, this would be showing some uniformity of effects from lateral masks. It would probably have been of most interest to compare the levels produced by (ii) and (iii) with those produced by (iv). This might have provided an insight into how masks were summed across spatial position. For example, if (iv) was presented at half the contrast of (ii) and (iii), yet the three (types of mask) produced very similar levels of masking, this would be evidence for linear summation of mask component contrasts across spatial position.

The relative spatial frequencies and orientations (as well as the relative spatial positions) of test and mask could also have been manipulated. For example, a 1cpd vertical test could have been used with -45deg 3cpd masks. This would have allowed an investigation into remote masking effects occurring from lateral masks. It would have been a natural development to the work in chapter 5 where the cross-channel masking produced with such stimuli allowed models of cross-channel spatio-temporally simultaneous masking to be developed.

The pilot work

Before performing an experiment like the one described above, it was necessary to be sure that lateral masks could significantly influence contrast detection thresholds. As described, such masks have been shown to have large effects on perceived contrast but it is far less clear what their effects are on contrast detection. Snowden & Hammett (1998) performed a particularly relevant study. They measured the detectability of a circular test patch in the presence and absence of a directly abutting surrounding mask annulus. (All stimuli were 2cpd vertical gratings). Using square-edged tests and masks that directly abutted, they showed the distance between the inner edge of the mask and the centre of the test to be very important as to whether any masking occurred. Very little or no masking was produced if the test size was 2deg (4 cycles) or greater (they increased the distance up to 4 deg), but beyond this size, the level of masking was raised as the patch size was reduced.

One concern for the present study was whether lateral masking could be produced with stimuli that were of sufficient size and which had appropriate spatial enveloping. It is considered that a test patch should have at least 1½ to 2 cycles at full-width, half-height. Regarding spatial enveloping, for the experiments in this thesis it was common to use stimuli that were gradually curtailed at their edges. This is known to reduce artefacts (occurring at the edges) that contain spatial frequencies and orientations quite distant from the patch's nominal content (e.g. see Kelly, 1970, and Graham, 1989). It was particularly necessary to reduce the occurrence of such artifacts in the experiment described above. This was because lateral masking effects produced with masks remote in spatial frequency and orientation to the test were investigated. If a masking effect was evident with stimuli having square or near-square edges, it may to some extent have been because some of the artefacts from the mask had the same spatial frequency and orientation as the test. If the test and mask were very close, this would make within-channel masking a likely explanation of such effects.

For the pilot study described here, I used tests with a minimum of about 1.8 cycles at full-width, half-height. The tests were circular patches and the masks were circular rings surrounding the patches. The enveloping between the outer edge of the test and inner edge of the masks was one of three types: (i) The two edges were curtailed with a raised cosine. The stimuli did not overlap so there was a point of zero contrast between them (where the two patterns had been curtailed to zero contrast). (ii) The two edges were curtailed with a raised-cosine but they did overlap at half their maximum height. (iii) The two edges were square and directly abutted. Conditions using the enveloping of (i), (ii) and (iii), are referred to in this section as conditions (a), (b) and (c), respectively. (Note that the width of the mask annulus was made very large to enhance the probability of masking, e.g. see Snowden & Hammett, 1998, and Cannon & Fullenkamp, 1991, for a discussion).

It was described above how this pilot study was primarily designed to measure the magnitude of masking produced by spatially adjacent masks with different enveloping between the test and mask stimuli. A second purpose of the study was to investigate whether masks remote in orientation to the tests produced any lateral masking effects. As discussed above, Snowden and Hammett (1998) showed substantial masking to occur with conditions very similar to (iii)(above) if the mask was close to the centre of the test. However, they only used tests and masks that were matched in orientation and spatial frequency (2cpd vertical gratings). Keeping the stimuli similar to those they had used, tests and masks were 2cpd gratings in this pilot study because they had shown such stimuli to produce masking effects under some conditions. As a development to their study however, the masking produced when such stimuli differed substantially in orientation was also investigated. This was because if the tests and masks had remote orientations and were spatially superimposed, it would be considered likely that they would excite different mechanisms. Thus, it would be considered that the masking produced was due to cross-channel suppression. This was important because of the larger study's purpose (described above).

There was reason to question whether there would be masking when test and mask were remote in orientation because lateral masking has been shown to be orientation specific in some cases. For example, Polat & Sagi (1993, 1994) showed the detection of a target Gabor to be enhanced by two collinear high contrast flankers if they had the same orientation as the target. Polat & Sagi (1993) however, showed this to decrease as the difference between the stimulus orientations was increased, culminating in a complete loss of the effect when the stimuli were orthogonal.

Methods

Tests and masks had 2cpd and were centred on the fixation point. They were temporally superimposed, were presented for 100ms, and had abrupt onsets and offsets. The test was a circular patch whilst the mask was a circular ring (annulus) surrounding the test. The mask patterns had a stimulus contrast of 48%. Tests were always vertical gratings but mask type was varied within sessions. The mask was either a grating oriented vertically, horizontally, or at -45deg, or was a plaid with components oriented ±45deg. Trials incorporating each mask type were blocked together whilst measurements of unmasked thresholds were interleaved with measurements of masked thresholds across sessions.

The diameter and spatial enveloping of the stimuli were varied between three conditions (termed here: a, b and c) that were run across sessions. This was done to vary the spectral characteristics of the regions close to the outer edges of the tests and the inner edges of the masks (as described above). For all conditions, the outer edge of the masks was curtailed with a raised cosine and had a diameter of 14.9° at full-width, half-height. The inner hole of the masks and the (outer) edge of the tests were curtailed with either a step edge or a raised cosine function. The distance between peak contrast and no contrast of the latter was 0.25°, which is equivalent to half a cycle of grating. The stimulus diameters and (spatial) envelopes used for the three conditions are listed in Table A1.1.

Condition	a.	b.	c.
Env. of test	Raised cosine	Raised cosine	Step edge
Diameter of test	0.9°	0.9°	1.1°
Env. of masks inner edge	Raised cosine	Raised cosine	Step edge
Width of annulus ring	6.7°	7.0°	6.9°

Table A1.1 The diameter and spatial enveloping of the stimuli used in (pilot) experiment 1

For condition (a), there was a point of zero contrast between test and mask. For condition (b), the raised cosine envelopes of test and mask were superimposed at half-height. For condition (c), the abrupt step edges of test and mask were adjoined.

Results and Discussion

Fig. A1.1 shows how the detection of test patches is affected by surrounding them with mask annuli. The different panels show data gathered using the three types of enveloping described previously. Data in the top, middle and bottom panels are from conditions a, b and c, respectively. The different columns in each panel show the effect of different mask types. (See figure legend for details. Note that the column shading reflects the mask type.) The data are similar to those of Snowden and Hammett (1998). When the stimuli were curtailed with raised cosine functions, little or no effect was produced by any of the masks. When the stimuli were directly abutting, all masks substantially reduced test detectability. The larger experiment described in this section's introduction was not performed because only very small masking effects were produced with stimuli having adequate spatial enveloping.

The level of masking was found to be reliant on the closeness of the mask to the test's centre. Why might this be? Snowden and Hammett (1998) suggested that the centre of a test is used for its detection and the area from which contrast information is gathered for its detection has a diameter of about 2deg. (They noted that this value was similar to previous approximates of foveal stimulation, for example, Howell & Hess, 1978; Ejima & Takahashi, 1983). Both studies agree that masks only have effects if they fall within this area (of summation).

Consider the data from condition (c) where the test and mask had square edges that were directly abutting. Leaving concerns about spectral artefacts to one side for a moment, the data have noteworthy characteristics. For example, they have a similar trend to the orientation masking data in chapter 6 that were gathered using spatially superimposed stimuli. With those, as the orientation difference between test and mask was increased, the masking gradually decreased but was still substantial even when the stimuli were orthogonal. A second noteworthy result shown in the panel is that the level of masking produced by the -45deg grating (middle column) was almost exactly the same as that produced by the (stimulus contrast matched) ±45deg plaid (far-right column). Such results have been shown in previous chapters (3, 4 and 5), which used spatially superimposed rather than spatially adjacent stimuli. Similar masking levels were produced by the two patterns when they had different spatial frequencies, or when they were matched at 1cpd (see chapters 3 and 4 in particular). I have previously discussed at length how such data may reflect linear summation of mask contrasts. However, when stimuli were matched at 2cpd (like in this pilot study), Ross and Speed (1991) showed a plaid to produce more masking than gratings (I have shown a similar result with 3cpd stimuli). To summarise, using 2cpd stimuli I have found evidence of linear summation of mask contrasts with spatially adjacent stimuli whilst Ross and Speed (1991) did not find evidence for this with spatially superimposed stimuli. Importantly though, spectral artefacts produced by the stimuli's sharp edges may be substantially contaminating the data that is shown here (see the introduction of this section for a discussion).

What is being revealed by the lateral masking (when test and mask are directly abutting)?

The stimuli used for this pilot study and by Snowden and Hammett (1998) only produced an effect on test detection when masks were very close to the centre of the test. Other studies of lateral masking have shown effects when masks are comparatively very far from the test's centre. For example, Polat and Sagi (1993, 1994) (using Gabor stimuli) revealed that a pair of co-linear flankers facilitated test detection when the distance between the centre of the test and mask was between 2 and 10 times the test's wavelength. The facilitation was sharply tuned to orientation (it was completely absent when test and mask were orthogonal). It also varied with the distance between the masks and the test, with the greatest level (about 6dB) reported to occur when the distance was about three wavelengths. (Like my result, they showed suppression, which was less specific to orientation, when the distance was less than two wavelengths). It's plausible that whilst Polat and Sagi revealed (facilitatory) interactions between mechanisms tuned to different spatial positions, only a form of simultaneous masking has been shown here (and by Snowden & Hammett, 1998). That is, when the mask is very close to the test's centre, the two patterns produce input to some common mechanism. When the mask is moved away from the test's centre, then its input to this mechanism reduces (whilst the test's remains constant), and so less masking is produced. The data here have similar characteristics to the orientation masking data in chapter 6 that were gathered with spatially superimposed stimuli. Those data were described in terms of cross-channel suppression. The divisive inhibitory pooling mechanism (see chapters 4 and 5) thought to underlie those data may also underlie the data from this lateral masking experiment. In relation, Durand et al. (2002) have provided data consistent with the notion that lateral masking effects from patterns close to and far from the centre of (the CRF excited by) the test result from two distinct underlying processes. They found many cells where surround suppression was absent when the temporal frequency of the annulus was above about 10Hz. However, they also found cells with surround suppression when the temporal frequency of the annulus was as high as 25Hz. Of particular importance to the current discussion, is that this suppression was particularly pronounced when the inner edge of the annulus was close to the centre of the CRF excited by the test.

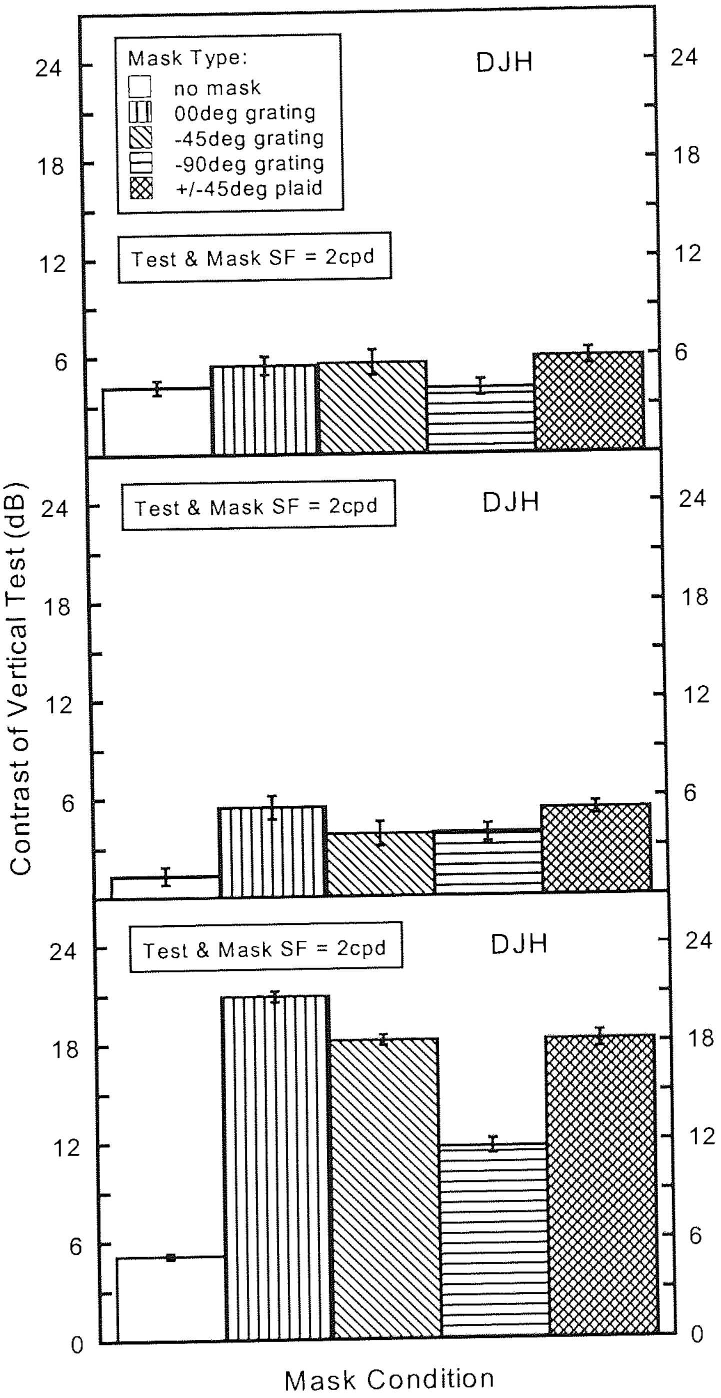

Fig. A1.1. Lateral masking effects. Data shown in the three panels were gathered using different types of enveloping between test and mask stimuli. The enveloping is described in the methods section; data in the top, middle and bottom panels are from conditions a, b and c, respectively. The different columns in each panel represent the effect of different types of mask (see Figure legend). Data is shown from one observer (DJH).

A1.3 Experiment 2(a). Does masking reveal asynchronies of CRF excitation and suppression latencies?

Introduction

Note that this study is not a direct development of that reported in the last section. It is considered here whether a masking paradigm could be used to reveal latencies of suppression produced by patterns placed at different positions of the visual field.

In a physiological study using macaques, latencies of classical receptive field (CRF) excitation and surround suppression were compared (Bair et al., 1999). Stimuli were patches of drifting grating. The test patch was circular and constrained to the CRF, whilst the mask was an annulus surrounding the test outside the CRF. Tests were optimally oriented whilst masks were oriented either the same, or orthogonally to the tests. Irrespective of mask orientation, it was found that the latency of surround suppression was roughly 20ms longer than CRF excitation. Smith et al. (2001) extended this study by comparing the latency of CRF excitation to an optimally oriented stimulus, with that of suppression from masks: (i) of orthogonal orientation constrained to the CRF or (ii) of orthogonal orientation surrounding the CRF. Their novel finding was that cross-orientation suppression from within the CRF had about the same latency as excitation of the CRF. Confirming results from the previous study, suppression from the (orthogonal) surround took about 20ms longer. It was proposed that these results indicate suppression within the CRF occurs within local circuits, whilst suppression from outside it results from long-range connections (within V1) or feedback pathways (from higher visual areas).

The aim here was to search for analogous psychophysical evidence, because this could have important implications for the rest my work. A small test patch was masked by either a spatially superimposed or annulus pattern. The mask was presented either before or after the test. If the data from this experiment were consistent with those of Bair et al. (1999) and Smith et al. (2001), it would be expected that: (i) spatially superimposed masks lower test detectability more, the shorter the stimulus onset asynchrony (SOA), and (ii) lateral masks lower test detectability more when they are presented about 20ms before the test than when the SOA is shorter. Also of interest was that the previous studies reported no effect of the orientation difference between test and mask stimuli on the latency of suppression. To investigate this psychophysically, I used masks whose orientations were either matched or oblique to the test's. I used 1cpd vertical tests and either 1cpd vertical, or 3cpd -45deg masks. I used the latter rather than a horizontal mask, matched in spatial frequency to the test (which would have been closer to the stimuli used in the previous studies). The main reason for this was that data has already been shown (in chapter 5) which supports the notion that the stimuli used here, when spatially superimposed, do not excite the same channel so any masking would be assumed to be due to cross-channel suppression.

Methods

The test was a vertical 1cpd Gabor measuring 1.7° at full-width, half-height. Masks were vertical 1cpd gratings or -45deg 3cpd gratings, presented either as Gabors matched in their spatial dimensions to the test, or as annuli. All stimuli were centred on the fixation point and mask contrast was 48%. The diameters of the inner and outer edges of the annuli, which were curtailed by a square-wave, measured 1.7° and 9.0°, respectively (the width of the rings measured 3.6°). The test was superimposed by the annulus when its amplitude was below or at half of its maximum height. Within an experimental session, trials presenting each mask pattern were blocked together and their sequence was varied. All stimuli were presented for 8.3ms (one frame) and had abrupt onsets and offsets. Whether the test was presented before or after the mask was varied across sessions. The duration separating the onsets of the test and mask, which was between 8.3ms and 133.3ms, was also varied across sessions. To prevent extraneous cues revealing within which interval the test was present, for each stimulus onset asynchrony value investigated, the trial intervals were matched in terms of their duration and the temporal position of the masks. The temporal position at which the test was presented in one interval of each trial was indicated by an audible beep in both intervals of each trial. Whether masked or unmasked thresholds were measured varied across sessions.

Results and Discussion

Fig. A1.2 shows masking produced from superimposed (circular symbols) and annulus (triangular symbols) patterns when the masks were presented either before (left panel) or after (right panel) the tests. Stimuli were matched (filled symbols) or remote (open symbols) in spatial frequency and orientation. An arrow represents the test's absolute detection threshold. Consistent with Smith et al. (2001), a remote mask spatially superimposed on a test generally produced the greatest level of masking at the shortest SOA's.
Consistent with Bair et al. (1999), a remote mask annulus produced most masking when the mask was presented about 20-40ms before the test. However, when test and mask annuli were matched in spatial frequency and orientation, most masking occurred at the shortest SOA's (rather than when the mask was presented a few tens of milliseconds before the test). So, these data offer only partial support for the previous findings.

The masking produced by the spatially superimposed pattern matched in spatial frequency and orientation to the test (filled, circular symbols) is curious. There is a substantial peak in the function when the mask was presented about 30ms before the test. This may be due to the latency of suppression from the mask being longer than that of CRF excitation, though the (much smaller) peak when the mask was presented about 30ms after the test doesn't fit well with this idea. Furthermore, Smith et al. (2001) showed suppression from within the CRF (albeit from an orthogonal mask) to have about the same latency as excitation within the CRF. An alternative possibility is that the peaks may be reflecting some CRF characteristic that is being revealed by the excitation from both test and mask stimuli. An investigation into whether this might be the case is reported in the next section.

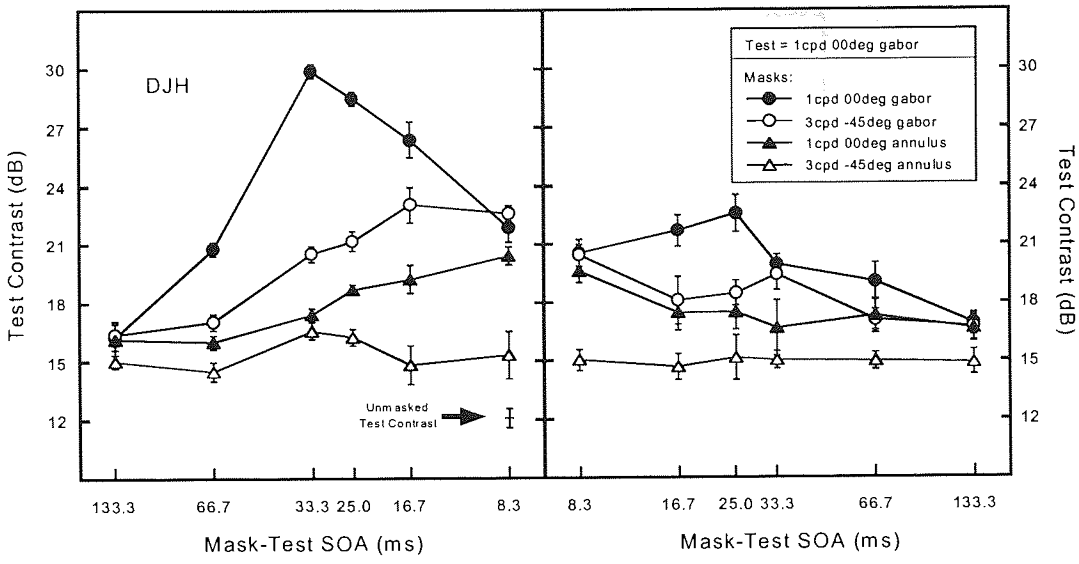

Fig. A1.2. Masking from spatially superimposed (circular symbols) or surrounding (triangular symbols) patterns, which were presented either before (left panel) or after (right panel) the tests. The spatial frequencies and orientations of test and mask were matched (filled symbols) or they differed (open symbols). An arrow in the left panel represents the unmasked threshold. Data are from DJH.

A1.4 Experiment 2(b). Previous masks produce more masking: a product of (i) the latency period for channel pooling or (ii) characteristics of the individual channels excited by the stimuli?

Introduction and Methods

It was shown in the last section that masks presented about 30ms before a test, raised its detection threshold almost 9dB more than a mask presented directly before the test. This occurred when the tests and masks were spatially superimposed 1cpd vertical Gabors. What might the position of this peak masking be due to? In the main body of this thesis (particularly chapter 5), several feed-forward models of channel pooling were shown to produce good fits to masking data. In such models, the channel excited by the test is divisively inhibited by a pool broadly tuned to spatial frequency and orientation, in which test and mask inputs are summed. The duration between mechanism excitation and suppression (apparent in the position of the masking function's peak) might be due to the time taken for such pooling to occur. An alternative possibility is that the position of the peak reflects a characteristic of the CRF. This possibility can be easily tested since CRF's are phase sensitive (see Graham, 1989). If the peak was due to a CRF characteristic, it would be expected to be absent when the test and mask are in anti-phase. Since spatio-temporally superimposed masking has been found to be phase-insensitive (see chapter 6, section 6.5), it would be expected that though the peak in the masking function would be absent, some masking would remain.

The top function in Fig. A1.3 was shown in the previous panel (Fig. A1.2). The experiment used to gather the data in this function was repeated, but the phase of the mask was shifted by 180°

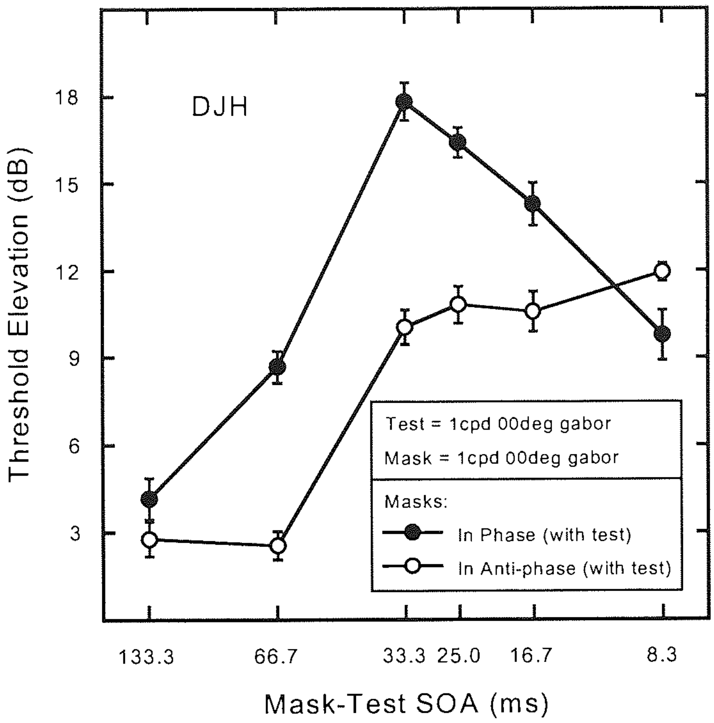

Fig. A1.3. Masking as a function of the interval between mask and test onset. Stimuli were spatially superimposed 1cpd vertical Gabors. Masks were either in phase (filled symbols) or in anti-phase (open symbols) with the tests. Data are from DJH.

Results and Discussion

Fig. A1.3 shows threshold elevation as a function of the time between the onset of the mask and the onset of the test. Tests and masks were spatially either in-phase (filled symbols)(data from last section) or in anti-phase (open symbols). All stimuli were 1cpd vertical Gabors. When they were presented in anti-phase, the prominent masking peak produced with in-phase stimuli was absent, though substantial masking was still produced at a broad range of SOA's. The phase-specificity of these functions suggests that they do not reflect the latency of a feed-forward mechanism, like those described by the models of masking shown in chapters 4 and 5. The phase sensitivity of the peak suggests that it was due to a characteristic of the CRFs that were excited by the stimuli. What might this characteristic be? The most obvious possibility is that it reflects the temporal impulse response function of the stimulated filter (see Watson, 1986, for a comprehensive review). Filters sensitive to medium and high temporal frequencies have bi-phasic temporal impulse response functions, which have an excitatory phase followed by an inhibitory phase. So, the masking peak occurring with in-phase stimuli might be due to the test being presented in the inhibitory phase of the response function produced by the mask. Returning to the data in the left-hand panel of Fig. A1.2, it can be seen that a (smaller) masking peak also occurred when the mask was presented after (rather than before) the test. It occurs at about the same SOA as the larger peak. Though it is unknown whether this too is phase-dependent, it is probably due to a similar underlying process as the larger peak. For example, more masking might occur because the mask is presented in the inhibitory phase of the temporal impulse response function produced by the test.

Like the data shown here, the forward and backward masking data of Georgeson and Georgeson (1987) shows further evidence of masking peaks occuring at SOAs of about 50ms (their Fig. 6).

A1.5 Summary

A major theme of this thesis has been how remote masking effects may provide an insight into cross-channel interactions. The term 'remote masking' has been used to describe masking effects when test and mask were remote in spatial frequency and/or orientation. With these studies, test and mask were always spatio-temporally superimposed. The work presented here should be considered as preliminary ideas and pilot work concerning how those enquiries (particularly of chapters 3 to 5) could be extended to investigate possible interactions occurring across spatial and temporal position.

The first study (experiment 1) was a pilot to an investigation into mask contrast summation across spatial position. It highlighted the importance of the spatial arrangement of the stimuli in invoking measurable effects. A comparison of my data with those of others in the literature shows that there are very different lateral masking effects, which can be induced depending on the stimuli used.

The second study (experiment 2) was an extension to physiological investigations that have shown suppression from spatially adjacent masks to have longer latencies than suppression from spatially superimposed masks. A psychophysical study provided only partial support. Some probable reasons for this were discussed.