INFORMATION TO USERS

This material was produced from a microfilm copy of the original document. While the most advanced technological means to photograph and reproduce this document have been used, the quality is heavily dependent upon the quality of the original submitted.

The following explanation of techniques is provided to help you understand markings or patterns which may appear on this reproduction.

1. The sign or "target" for pages apparently lacking from the document photographed is "Missing Page(s)". If it was possible to obtain the missing page(s) or section, they are spliced into the film along with adjacent pages. This may have necessitated cutting thru an image and duplicating adjacent pages to insure you complete continuity.

2. When an image on the film is obliterated with a large round black mark, it is an indication that the photographer suspected that the copy may have moved during exposure and thus cause a blurred image. You will find a good image of the page in the adjacent frame.

3. When a map, drawing or chart, etc., was part of the material being photographed the photographer followed a definite method in "sectioning" the material. It is customary to begin photoing at the upper left hand corner of a large sheet and to continue photoing from left to right in equal sections with a small overlap. If necessary, sectioning is continued again — beginning below the first row and continuing on until complete.

4. The majority of users indicate that the textual content is of greatest value, however, a somewhat higher quality reproduction could be made from "photographs" if essential to the understanding of the dissertation. Silver prints of "photographs" may be ordered at additional charge by writing the Order Department, giving the catalog number, title, author and specific pages you wish reproduced.

5. PLEASE NOTE: Some pages may have indistinct print. Filmed as received.

Xerox University Microfilms

300 North Zeeb Road
Ann Arbor, Michigan 48106

75-29,286

MEZRICH, Joseph Jacob, 1945-
SPATIAL FREQUENCY PROCESSING FOR THRESHOLD VISIBILITY AND FOR PICTURE RECOGNITION.

The University of Michigan, Ph.D., 1975
Psychology, experimental

Xerox University Microfilms, Ann Arbor, Michigan 48106

THIS DISSERTATION HAS BEEN MICROFILMED EXACTLY AS RECEIVED.

SPATIAL FREQUENCY PROCESSING FOR THRESHOLD VISIBILITY AND FOR PICTURE RECOGNITION

by

Joseph Jacob Mezrich

A dissertation submitted in partial fulfullment
of the requirements for the degree of
Doctor of Philosophy
(Psychology)
in The University of Michigan
1975

Doctoral Committee:

Associate Professor Robert G. Pachella, Chairman
Associate Professor Daniel G. Green
Professor Wilfred M. Kincaid
Professor J.E. Keith Smith

ACKNOWLEDGMENTS

I would like to express sincere gratitude to my thesis chairman: Robert G. Pachella, and to my thesis committee: Daniel G. Green, Wilfred M. Kincaid, and J.E. Keith Smith for their valuable advice and assistance throughout the course of the research reported here. Appreciation is also expressed for the generosity of Carol M. Cicerone in allowing me the use of her laboratory, and for providing general assistance. Thanks are extended to Jack Di Guisseppe for making his spatial filtering computer programs available to me, and to Frank Goode for computer programming assistance. Important guidance was also provided by David H. Krantz, this is gratefully acknowledged.

Clyde H. Coombs was responsible for my coming to the University of Michigan, and for providing financial support throughout my graduate studies in psychology as a Mathematical Psychology trainee. I am grateful to him for this.

I thank George D. Cody, Istvan Gorog, and Roger W. Cohen of RCA Laboratories, Princeton, N.J., for providing me with hard copy of the filtered noise and for the grating slides, and for supporting my writing of this report while on the staff of RCA Laboratories.

The typing of this report, and its earlier versions, was done by Arlene F. Grier. She did an expert job, and I thank her for it.

A special acknowledgement is given to Livia and Aliza, my wife and daughter. They made the difficult times comfortable, and the comfortable times enjoyable.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS ii

LIST OF TABLES iv

LIST OF FIGURES v

LIST OF APPENDICES vi

CHAPTER 1: INTRODUCTION 1

CHAPTER 2: SINGLE FREQUENCY VISIBILITY PROCESSING 9

CHAPTER 3: PICTURE RECOGNITION WITH NARROW-BAND NOISE MASKING 39

CHAPTER 4: SINGLE FREQUENCY VISIBILITY PROCESSING AGAINST A NON-BLANK BACKGROUND 58

CHAPTER 5: SUMMARY AND COMMENT 77

APPENDIX 84

REFERENCES 96

LIST OF TABLES

Table		Page
1	MODULATION OF TEST GRATINGS	18
2	VALUES OF k_L AND k_Q IN THE RELATION $\log s = k_L \log t + k_Q (\log t)^2 + \text{constant}$	25
3	EXPERIMENTAL SUMMARY FOR NOISE MASKED PICTURE RECOGNITION	50
4	TEST-NOISE COMBINATIONS FOR NON-BLANK BACKGROUND GRATING VISIBILITY.	62
5	AVERAGED - log C for S.B.	65
6	AVERAGED - log C for B.B.	65
7	AVERAGED - log C' for S.B.	66
8	AVERAGED - log C' for B.B.	66
A1	RESULTS SUGGESTIVE OF CONTRAST HOLE EFFECTS . . .	92
A2	MIDDLE FREQUENCY NOISE SESSION RESULTS	94

LIST OF FIGURES

1	GRATING VISIBILITY THRESHOLDS FOR VARIOUS EXPOSURE DURATIONS, FROM SCHOBER AND HILTZ	10
2	CONTRAST SENSITIVITY AS A FUNCTION OF EXPOSURE DURATION, FROM NACHMIAS	12
3	SCHEMATIC OF OUR OPTICAL ARRANGEMENT	16
4	LOG CONTRAST-SENSITIVITY AS A FUNCTION OF EXPOSURE DURATION FOR B.B.	26
5	LOG CONTRAST-SENSITIVITY AS A FUNCTION OF EXPOSURE DURATION FOR S.B.	27
6	LOG CONTRAST-SENSITIVITY FOR B.B.	32
7	LOG CONTRAST-SENSITIVITY FOR S.B.	33
8	CONTRAST-SENSITIVITY FUNCTIONS, FROM NACHMIAS	36
9	GEOMETRY OF THE ANNULUS FILTER	41
10a	THE EARLY NOISE TRIAL	47
10b	THE LATE NOISE TRIAL	48
11	LOG CONTRAST-SENSITIVITY AS A FUNCTION OF BACKGROUND FREQUENCY FOR B.B.	68
12	LOG CONTRAST-SENSITIVITY AS A FUNCTION OF BACKGROUND FREQUENCY FOR S.B.	69
13	LOG MASKING EFFECT OF A BACKGROUND FREQUENCY FOR B.B.	70
14	LOG MASKING EFFECT OF A BACKGROUND FREQUENCY FOR S.B.	71
A1	AMPLITUDE OF SPECTRUM FOR FINITE EXTENT GRATING	85
A2	LOW FREQUENCY AND HIGH FREQUENCY NOISE PATTERNS	88

LIST OF APPENDICES

APPENDIX 1: SUFFICIENT VISUAL ANGLE TO AVOID FREQUENCY DISTORTION IN FINITE EXTENT GRATINGS 84

APPENDIX 2: PICTURE-STIMULI FOR THE RECOGNITION EXPERIMENT 86

APPENDIX 3: PICTURE DISTORTION BY FILTERED NOISE SUPERPOSITION: A SOURCE OF CONTRAST HOLES AND PLUGS. 87

APPENDIX 4: MIDDLE FREQUENCY SPATIAL NOISE MASKING IN PICTURE RECOGNITION 94

CHAPTER 1

INTRODUCTION

The past decade has witnessed increasing focus on the examination of spatial vision as spatial frequency analysis. Interest in this issue stems from the convergence of a number of diverse investigations upon the fact that for threshold stimuli the spatial frequency sensitivity of the visual system is derived from independent channels each sensitive to a restricted range of frequencies. The bandwidth of these channels is generally concluded to be about an octave for frequencies beyond about 3 cycles/degree (e.g. a channel with maximum response at 10 cycles/degree is essentially insensitive to frequencies higher than 20 cycles/degree or lower than 5 cycles/degree). Very low frequencies apparently all report to the same channel (e.g. Blakemore and Campbell, 1969b). In a loose sense one can say that for threshold signals achromatic spatial vision is derived from about five distinct channels.

The evidence leading to the independent channels hypothesis comes from the many demonstrations of frequency dependent adaptation effects and from the direct examination of the detectability of sinusoidal modulations of luminance (simple gratings) as compared to the detectability of the superposition of sinusoidal luminance modulation of different frequencies (complex gratings). In the detection experiments of Sachs, Nachmias, and Robson (1971) an elegant analysis of the psychometric functions generated by simple and complex gratings revealed independent channels of bandwidth less than an octave. The same general conclusion follows from the adaptation experiments. For example, Blakemore and Campbell (1969b) showed

that adapting to a suprathreshold sinusoidal grating of a given frequency subsequently raises the threshold for frequencies within an octave of the adapting frequency, but is generally ineffective in modifying sensitivity to frequencies outside this range.

The validity of the independent channels hypothesis is critical if spatial frequency analysis by the visual system is to be taken seriously. Various investigators since Schade (1956) have examined the contrast threshold for sinusoidal luminance modulation. The value of the derived transfer function relating stimulating spatial frequency and psychophysical response as a descriptive tool cannot be denied. The utility of linear systems theory as a modeling device is well known. However the set of conceivable orthogonal expansions admitted by linear systems theory is by no means restricted to the Fourier series. One might just as well claim that the function of spatial vision is a decomposition of patterns into a set of Legendre polynomials. Without the findings indicating independent channels there is no reason to attribute any particular significance to a Fourier decomposition of patterns as far as visual processing is concerned.

For this reason it is important to note that two studies provide some trouble for the otherwise neat scheme of independent channels. Tolhurst (1972) examined the sensitivity reduction for spatial frequencies following adaptation to suprathreshold square wave patterns. The energy distribution of such a pattern contains discrete frequencies at the fundamental frequency, and at odd multiples of the fundamental. The frequency components of significance are at the fundamental and three times the fundamental. If these frequency

components are independently detected, i.e. processed by distinct channels, it would follow that sensitivity reduction following adaptation to a square wave luminance distribution should be equivalent to the sensitivity reduction following separate adaptation to the Fourier components of the square wave. However, adaptation to separate Fourier components yields greater sensitivity reduction than is achieved by the square wave. Tolhurst concludes that for suprathreshold signals there is considerable cross-talk; channel independence is not maintained above threshold.

An even stronger result was found by Nachmias et.al. (1973). They found, using a forced choice procedure, that adapting to a square wave is equivalent to adapting to a sine-wave as far as spatial frequency sensitivity reduction is concerned. Square wave adaptation does not yield sensitivity reduction for the harmonics of the square wave. Evidence of Fourier decomposition for distinct channel processing is absent for suprathreshold patterns.

Studies of Fourier processing, either from the linear systems viewpoint or from the independent channels viewpoint, have typically been concerned with the examination of thresholds for pattern visibility. With one notable exception, consideration has not been given to the issue of pattern recognizability. While the connection between steady state pattern visibility and the processing of such patterns for recognizability seems obvious, it deserves to be made explicit. At the very least the results of Tolhurst (1972) and of Nachmias et.al. (1973) cited above cause one to wonder whether the suprathreshold signal detection that is involved in pattern recognition is properly described by threshold visibility measurements.

A fairly dramatic demonstration of the likely relevance of the channels hypothesis to pattern recognizability has recently been presented by Harmon and Julesz (1973). By superimposing computer generated noise restricted to a two-dimensional band of spatial frequencies on a picture it was shown that picture recognizability suffered considerably if the noise band overlapped the frequencies that contained picture information, but did not affect recognizability if the noise was located beyond the frequency regime of picture information. Unfortunately, this was a demonstration in the strictest sense of the word.

The authors provide the noise cum picture presentations they have produced for public submission to an interocular trauma test. Inspection of their presentation, indeed, persuades one to accept the qualitative validity of the channels notion (the effect does hit one between the eyes). However, a systematic investigation of the degree of recognition degradation induced by filtered noise is not attempted. A completely convincing argument for the applicability of the concept of Fourier processing by channels for visual recognition remains to be made.

The thrust of the study reported here is to examine the spatial frequency issue for recognition from an information processing point of view. The broad outline of the recognition process proposed by Neisser (1967) suggests a particularly interesting area for consideration.

Neisser's hierarchical scheme for recognition contains three basic components: a preattentive stage that performs a crude, global assessment of the scene to be perceived and sets up the organism for

the act of recognition; a focal attention stage used for scrutinizing local scene segments for detail; and a figural synthesis stage used for aggregating the assimilated local information and constructing a unified percept.

The figural synthesis stage is compelling, but is beyond the scope of the issue at hand. It will not be considered here.

The concepts that Neisser labels preattentive processes and focal attention are not new. For example, one finds mention of similar ideas in Hebb's work (1949) and in earlier work of the Gestalt psychologists. What distinguishes Neisser's statement is his emphasis that the inclusion of these processing stages is requisite for any non-trivial recognition, in man or machine. (See Narasimhan, 1970, for a presentation of comparable ideas in the context of machine recognition).

While detailed functioning of the preattentive stage is difficult to specify with the typical feature alphabets, i.e. alphabets constructed from line and curve segments, it is possible to be somewhat less vague about this stage if spatial frequency analysis by channels is the mechanism for assimilation of the environment by the visual system.

How is it possible to perform a crude global analysis prior to the local analysis that leads to the unified global percept? In a Fourier decomposition crude global information, such as general object size, is conveyed by the low spatial frequencies, while local detail, such as precise contour information, is conveyed by the higher spatial frequencies (see Rosenfeld, 1970, for pictures that illustrate this point). Phase information is important primarily for

object separation in multi-object scenes. We will disregard phase.

It seems that a tentative match can be made between the information conveyed by the lower spatial frequencies and the information processed by the preattentive stage; likewise a tentative match can be made between the information conveyed by the higher frequencies and the information processed by the focal attention stage.

On the other hand, some have argued that object recognition depends entirely on the perceptual realization of contour information at the object boundary. That is, object specification is conveyed exclusively by step function luminance gradients. Smoothly varying luminance gradients, i.e. low spatial frequency information, apparently has a role in esthetics, but none in recognition. Proponents of this point of view have typically been concerned with the recognition of silhouettes (e.g. Attneave & Arnoult, 1956) or alphanumeric material (e.g. Rumelhart, 1970, Rumelhart & Siple, 1974) and refer to supporting physiological evidence reported by investigators who exclusively use step function luminance gradients in their stimulating patterns (e.g. Hubel & Wiesel, 1959, 1962) or to successful machine recognition procedures developed for contour defined objects (e.g. Selfridge, 1959). If attention is restricted to patterns without important low frequency content, it is not surprising that such information is deemed irrelevant. Broad areas of consideration, such as physionogmic perception, are necessarily excluded from such a framework. Nevertheless, this school of thought provides a useful counterpoint to the picture of differential spatial frequency processing rates that can be framed in Neisser's hypothesis.

We have suggested that Neisser's sequential stages dichotomize spatial frequencies with respect to their function in the recognition process. Low spatial frequencies are dealt with first to provide a context for the later processing of higher frequencies, i.e. pre-attentive processing precedes focal attention. In terms of processing rates, low spatial frequencies should be processed more rapidly than high spatial frequencies.

The competing hypothesis suggests the reverse ordering of high and low frequency processing rates. If low spatial frequencies only serve the function of fleshing out an achieved percept, one would expect such information to be a relatively late component of visual analysis.

The issue at hand is not whether either of these hypotheses survives a critical test, but whether it is meaningful to pose the hypotheses as stated. If spatial vision is basically spatial frequency analysis, then one should be able to identify models of visual recognition by statements they make about the processing of spatial frequencies. Testing the validity of a specific model is then straightforward. The relative processing rates for a set of spatial frequencies spanning some range of interest are examined for consistency with an appropriate prediction derived from the model being entertained.

The specific issue we will concern ourselves with here is an assessment of the degree of correspondence between the relative spatial frequency processing rates measured for threshold visibility situations and the observed processing rates for an experimental task involving picture recognition. An examination of the relative visibility

benefit derived from increased exposure for various spatial frequencies is our first goal.

CHAPTER 2

SINGLE FREQUENCY VISIBILITY PROCESSING

Surprisingly, a survey of the literature reveals only two studies that include measurements of threshold visibility for periodic patterns of different spatial frequency as a function of exposure duration. Both studies used square wave, rather than sinusoidal luminance modulation. The choice of square wave patterns was a practical necessity for these investigators. Unfortunately this leads to some difficulties in the interpretation of data, as will be pointed out shortly.

The first investigation into this matter was that of Schober and Hiltz (1965). The curves in Figure 1, taken from their paper, provide the best summary of their results. The two panels in this figure display measurements taken for two observers. To read their spatial frequency scale in more conventional units, use the fact that 0.01 lines/min. of arc is .6 cycles/deg as a conversion factor. We can immediately conclude from these curves that exposure duration affects sensitivity to different frequencies differently. If all frequencies benefited identically from exposure duration, then it would necessarily follow that the family of contrast threshold curves for a given observer should have the property that exposure duration affects only the vertical displacement of the threshold function, but not its shape. This is clearly not the case. More detailed inference from these curves is difficult.

A more careful study of the effect of exposure duration was carried out by Nachmias (1967). In addition to examining contrast threshold function changes with exposure duration changes, as was

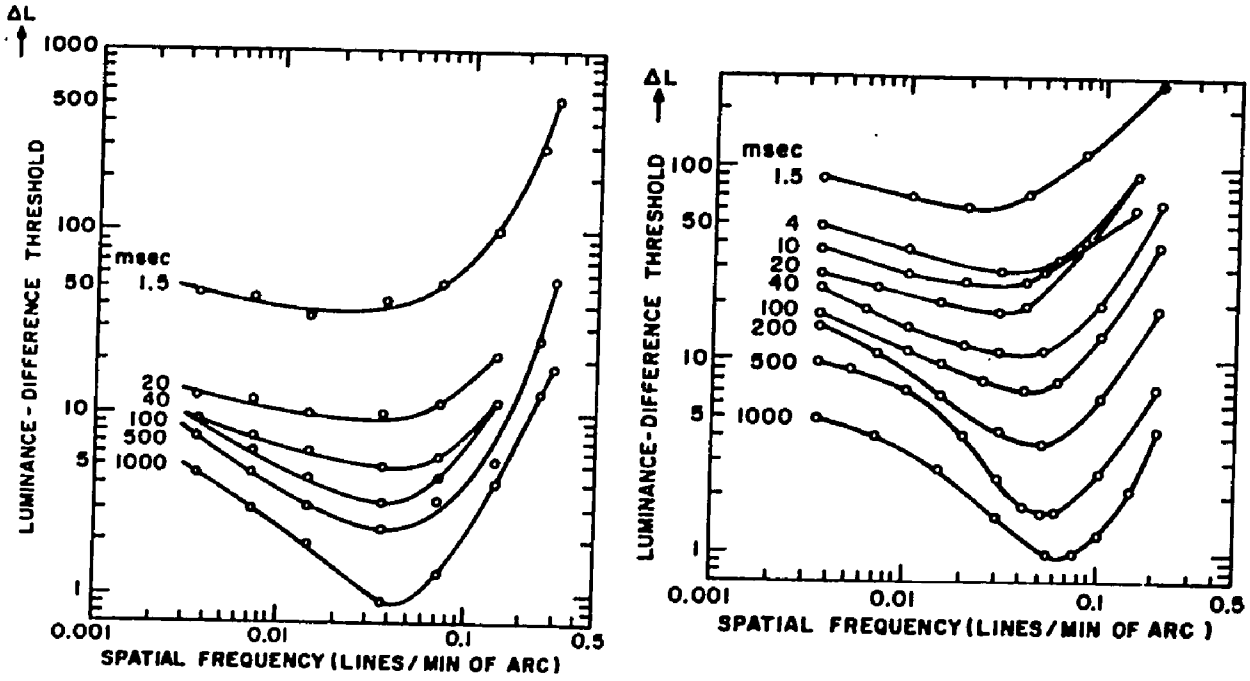


FIGURE 1. Visibility threshold as a function of spatial frequency for various exposure durations. Results for two observers in the study by Schober and Hiltz (1965).

done by Schober and Hiltz (1965), he carried out detailed measurements of exposure duration effects for two spatial frequencies, 0.7 cycles/deg and 17.5 cycles/deg. Mean luminance was 10 ml, one third the value of the Schober & Hiltz experiment. Nachmias' results are displayed in Figure 2. These curves indicate that the low frequency grating reaches its plateau sensitivity at a far shorter exposure duration than does the high frequency grating. Whereas sensitivity to the high frequency grating is still increasing at 300 msec exposure, the low frequency grating does not greatly benefit from increasing the exposure beyond about 50 msec.

We may therefore conclude from these two experiments that for square wave gratings different spatial frequencies are processed by the visual system at different rates and, if a generalization from two spatial frequencies is appropriate, low frequencies seem to be processed more rapidly than high frequencies.

There is, however, an essential difficulty in using results obtained with square wave gratings to make statements about spatial frequency processing that is particularly critical in the situation examined by Nachmias. The low frequency grating used in his study was 0.7 cycles/deg. Decomposing a 0.7 cycles/deg square wave into its Fourier components yields a fundamental at 0.7 cycles/deg and a first harmonic at 2.1 cycles/deg at one third the amplitude of the fundamental. There are also higher harmonics at lower amplitudes. The present discussion does not suffer if they are disregarded. From data presented by Nachmias, one finds that the peak sensitivity to square wave gratings with an average luminance of 10 ml. is somewhere between 2 and 3 cycles/degree. Let us accept this as evidence that

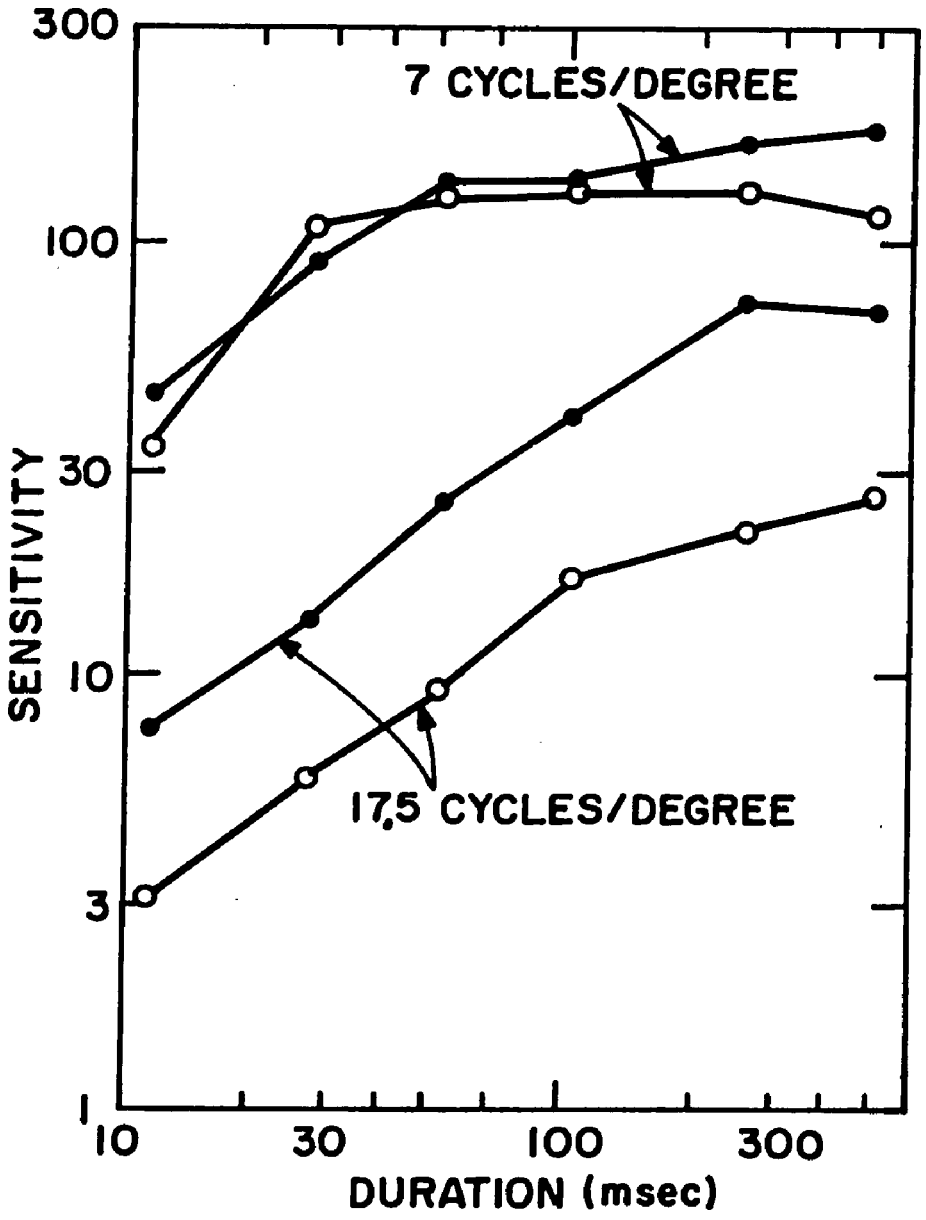


FIGURE 2. Dependence of contrast sensitivity on exposure duration for two spatial frequencies. Empty circles and filled circles represent results for different observers. Taken from Nachmias (1967).

the peak sensitivity to sinusoidal luminance modulation would be in the same neighborhood. For the sake of argument, assume that the peak sensitivity is, in fact, at 2.1 cycles/degree. Furthermore, assume that the sensitivity increase as we go from very low frequencies to the frequency of peak sensitivity is linear with frequency. This is a typical result. Then the sensitivity at 0.7 cycles/deg should be one third of that at 2.1 cycles/degree. It was pointed out that measuring the detection threshold for a 0.7 cycle/deg square wave is roughly equivalent to determining the threshold for the superposition of a 0.7 cycle/deg sine wave and a 2.1 cycle/deg sine wave at one third the amplitude of the fundamental. We see, however, that the sensitivity difference between these two frequencies compensates for the amplitude difference. Consequently, the threshold for 0.7 cycle/deg square wave visibility is not a true indication of 0.7 cycle/deg frequency sensitivity. In particular, one would expect that sensitivity measurements for low frequency square waves would tend to over-estimate the true frequency sensitivity. This would be revealed by a negligible low frequency decline in the contrast sensitivity functions for square waves. This surprising feature is evident in the results reported by Nachmias.

The problem is even further complicated when exposure duration is the parameter of interest. It seems to be the case that changes in exposure time benefit different frequencies differently. There is no reason to expect, for example, that 0.7 cycle/deg and 2.1 cycle/deg information are processed at the same rate. Consequently the estimation error in using 0.7 cycle/deg square wave sensitivity

as a measure of frequency sensitivity is a function of exposure duration. Without prior information concerning how the sensitivity rate differs for the fundamental and first harmonic of the square wave used, it is difficult to evaluate the true sensitivity as a function of exposure increase at the frequency of interest. It should be noted that Nachmias was careful to point out that the true shape of the contrast sensitivity function cannot be properly inferred from tests with square wave gratings.

Summarizing the above, we can say that it is quite clear from the studies of Nachmias and of Schober and Hiltz that the sensitivity increase that accompanies exposure increase is not the same for all spatial frequencies. The results presented by Nachmias lead one to suspect that low spatial frequencies are processed more rapidly than high spatial frequencies; however for the reasons cited above this conclusion requires some caution.

To accurately assess the change in sensitivity for different spatial frequencies as exposure duration is increased it is necessary to use sine wave gratings as stimuli. Furthermore, it is desirable to use a sample of spatial frequencies that would afford us a bit more precision in characterizing relative processing rates than is possible by testing exposure effects using one low frequency grating and one high frequency grating. For these reasons it was deemed worthwhile to perform a variation of Nachmias' experiment using sine wave gratings with more than two frequencies. The spatial frequencies selected for study were 1.25 cycles/deg, 2.5 cycles/deg, 5 cycles/deg, 10 cycles/deg, and 20 cycles/deg. The choice of octave separation

between these frequencies was motivated by the channels hypothesis discussed earlier. Presumably we will be stimulating a distinct channel with any one of these frequencies and possibly reveal different time constants for different channels. It should be clear that this experiment, by itself, can say nothing about the existence of channels.

Method

Apparatus:

The experimental set-up used was a modification of the optical arrangement designed by Cicerone (1974). Our version of this arrangement presented the observer with a $4.7^{o} \times 4.7^{o}$ square adaptation field which he could replace with a $3.7^{o} \times 2.5^{o}$ sine wave grating, stripes vertical, optically superimposed on a geometrically congruent, blank, dilution field, at the press of a button. The duration of this replacement was preset foreach exposure, as was the luminance of the grating field and of the dilution field. At the termination of the grating field-dilution field presentation, the adaptation field immediately reappeared. A small centered dot in the adaptation field was used for fixation. The mean luminance of the adaptation field, 3.4 ml, was maintained for each grating field-dilution field exposure. A schematic diagram of the apparatus is shown in Figure 3.

The light source, S, was an automobile lightbulb (tungsten filament with monitored 2.4 A, 60 Hz input). The light from this single source passed through all three channels of the system: the adaptation channel, channel A; the blank, dilution field channel, channel B; and the channel containing the test grating, channel T. Identical

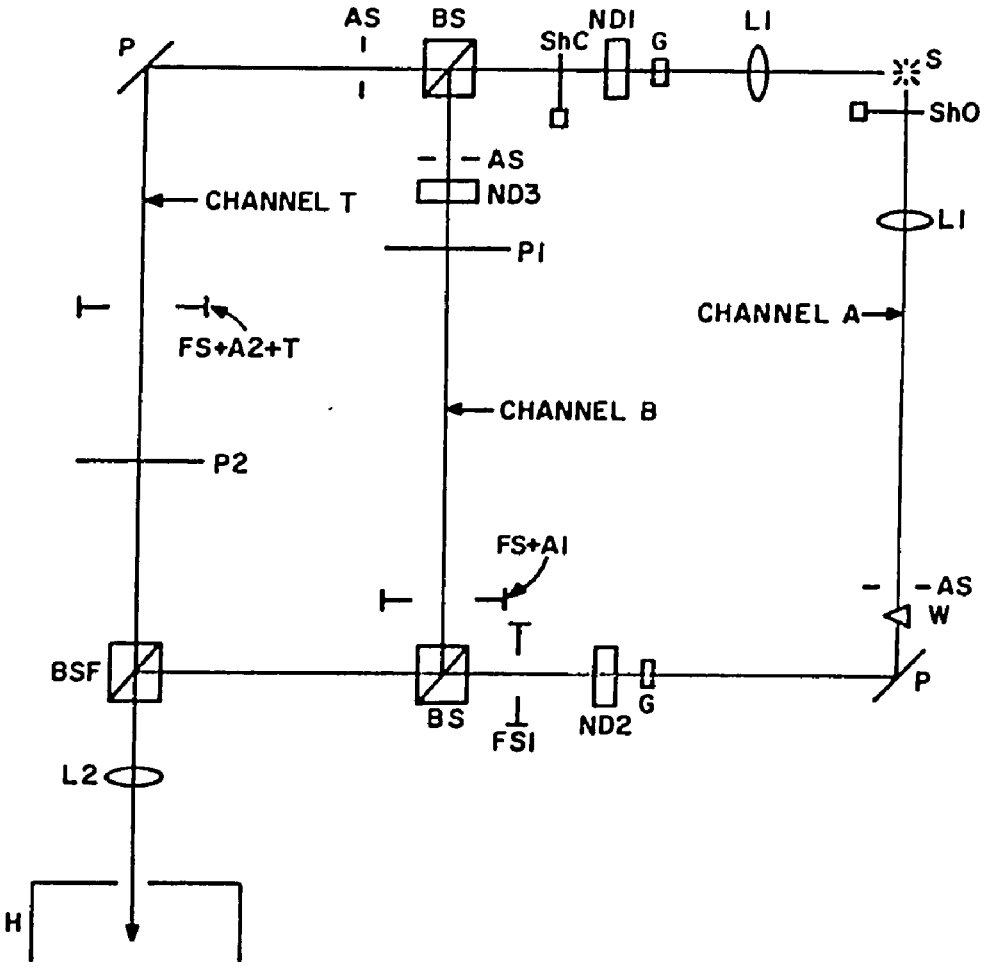


FIGURE 3. Schematic of our optical arrangement. See text for details.

lenses, L_1, imaged the lightbulb filament on the aperture stops, AS. These three images were passed through the various optical components, including the prisms, P, and the beam splitters, BS. After final recombination by beam splitter BSF, the beams passed through the Maxwellian lens, L_2, which imaged the aperture stops as 1 mm diameter, congruent circles in the plane of the observer's pupil. Use of an artificial pupil was unnecessary, since 1 mm is smaller than the natural pupil diameter even for stronger illumination than used here.

The field stop in channel A, FS1, was a 4.7^{o} x 4.7^{o} square containing a transparent sheet of glass with a small opaque dot in its center for use as a fixation point. Channels B and T contained identical 3.7 x 2.5^{o} rectangular field stops. Glass encased, high quality polarizing sheets (Polaroid HN22) were permanently taped onto these field stops for use as analyzers in conjunction with polarizers P1 and P2. Polarizers P1 and P2 were also glass encased and were situated in rotatable, ball-bearing mounted, large apertures that had engraved, 1^{o} increment, circular scales (Ealing Optical Co.). The appropriate test grating, T, for a given session was taped onto the field stop of channel T. The neutral density filters, ND1, 2, and 3, as well as the wedge filter W, were selected for equal luminance in all channels. The Kodak Wratten filters, G, colored the light green.

Two electromagnetic shutters, ShO and ShC, (Uniblitz model 225, Vincent Associates) were used. ShO passed light through channel A, and ShC passed light through channels B and T. The normally closed shutter, ShC, and the normally open shutter, ShO, were controlled by separate timing and drive circuits (Vincent Associates Model 300 for ShC, and Model 310 for ShO). A button press triggered both timing

circuits, closing ShO and opening ShC simultaneously and for the same duration.

The observer's head was held steady by biting on a dental impression. An occluder was used over the unused eye, and a hood, H, was used to eliminate stray light.

By proper use of P1 and A1 in channel B, and P2 and A2 in channel T, we set the luminance of Channel B to 3.4 (1-α) ml and the luminance of channel T to 3.4 α (1 + m sin $2\pi fx$)ml., where α is any value between one and effectively zero, f is the frequency of the test grating, and m is the modulation depth of the test grating. Channels B and T are optically superimposed with a resulting luminance 3.4 (1 + αm sin$2\pi fx$) ml. The contrast of the observed grating is αm, so contrast changes are accomplished by changing α.

The gratings were slides, reduced from 35 mm format negatives of CRT sine wave grating photographs. The fidelity of the sinusoidal modulation was checked, and the modulation depth was measured for each slide from microdensitometer recordings. The modulation depth for each grating slide is given in Table 1.

Table 1. Modulation of Test Gratings

Frequency	Modulation Depth
1.25 cycles/degree	.82
2.5 cycles/degree	.84
5 cycles/degree	.80
10 cycles/degree	.75
20 cycles/degree	.53

For practical reasons, a grating is restricted to 3.7^{o} visual angle in our set-up. When a sine wave grating is confined to a finite visual angle it is no longer truly periodic. It is, therefore, important to have some assurance that the angle is large enough that we can accept the finite grating as a reasonable approximation to the desired sinusoid. Analysis is presented in Appendix 1 to show that 3.7^{o} is adequate for the range of spatial frequencies employed here. The recent empirical study of this issue by Hoekstra et.al. (1974) also indicates that we have chosen a sufficiently large angle.

Since we used crossed polarizers to attenuate channels B and T it was necessary to avoid the use of white light, as provided by the unfiltered source, S. The degree of attenuation of a polarizer-analyzer pair is a function of wavelength, so the use of white light introduces color differences between the three channel fields that depend on the contrast setting. By consulting the transmission profiles provided by the manufacturer for the polarizers, and by visual inspection, it was decided that a Kodak #61 green Wratten filter provided excellent color uniformity. The location of these green filters (G) is indicated in Figure 3.

The visual angle subtended by the adaptation field was chosen to be larger than that of the grating field and of the dilution field, and the geometry of the adaptation field was made different from that of the other fields to eliminate any potential apparent motion cues at the exposure onset or offset. If all three fields were geometrically identical, and a very slight misalignment occurred

between the three fields, conditions would be right for apparent motion of the field boundaries at the onset or offset of an exposure. This would have the appearance of a jump of the viewed field that would depend on the contrast setting. This potential problem was avoided by our use of a large, square adaptation field.

A photometer (Photovolt Model 502m), fitted with a photopic correction filter, was used to measure the luminance level by situating the photosensitive face of the tube a fixed distance from the Maxwellian image. Prior to this the photometer with filter had been calibrated with a standard source. Visual photometry was used to insure that the grating channel and the blank channel had identical transmission properties. For this measurement the blank channel field stop was rotated 180^{o} about the axis perpendicular to the optical path, and a grating was placed out of focus, near the prism, P, of channel T, rather than taped on the channel T field stop. This created a convenient bipartite field with the blank field on one side and the defocussed grating field on the other side. With the polarizer-analyzer pairs in each channel set for 50% transmission, an up-down procedure was implemented to achieve a luminance match by changing the neutral filter, ND3, in channel B. Both the experimenter and one of the subjects (B.B.) participated in this procedure, with resulting agreement for both observers. This was repeated for every grating.

For practical reasons, a different visual procedure was used to check that the adaptation field luminance was not noticeably different from that of the other fields. After the tests for

grating field-blank field equivalence were completed, the blank channel field stop was returned to its original orientation, the polarizer-analyzer pair for this channel was set for maximum transmission, and the grating channel was extinguished. Using the shutters, the adaptation field and blank field were alternated and the observer was asked to note any apparent luminance change around the location of the fixation point as filter ND2 of channel A was changed.

Observers:

Two paid university undergraduates served as observers. Subject S.B. was a 22 year old female and subject B.B. was a 24 year old male. Both were well corrected myopes. The subjects were trained in the task to be performed for approximately 20 hours prior to the experimental sessions.

Procedure:

Before an experimental session the subject dark adapted for about three minutes. Longer dark adaptation periods were tried during the training sessions, however the shorter adaptation periods were found to be adequate. The process of bite-bar adjustment, checking of test-field alignment with blank-field, and the warm-up period gave the subject approximately ten minutes to light adapt to 3.4 ml. A subject was typically involved in two one hour sessions back to back. There was a five minute break between sessions. After the break the subject was re-adapted for slightly shorter periods than used for the first session of the pair.

In a given session one spatial frequency was tested for all exposure durations of interest; 10 ms, 30 ms, 50 ms, 100 ms, and 300 ms.

A trial consisted of two sequential exposures of the same duration, one containing the test grating the other being blank. The subject's task was to indicate which exposure was not blank.

After being given a verbal ready signal, the subject pressed a button which initiated the first exposure of the trial. The subject then waited for another ready signal, following resetting of the polarizers, before initiating the second exposure. The delay between exposures was less than 10 seconds. The subject responded after viewing the pair of exposures by ringing a bell once or twice. Responses were recorded by the experimenter. There were no visual or auditory cues associated with polarizer adjustment. Presentation of the grating in the first or second exposure of a trial was equally likely. All exposure durations were equally likely for a particular trial if the preceeding trial resulted in a miss, or if the preceeding two trials resulted in hits. Otherwise, the same exposure duration was used for a consecutive pair of trials. This psuedo-random exposure selection was used to avoid undue complication in data collection by the experimenter using the algorithm described below. Since the subject was unaware of the algorithm being used, and since no response feedback was given, the exposure duration used on any given trial was effectively random.

Wetherill's procedure (Wetherill & Levitt, 1965), was used to estimate contrast thresholds. In a two-alternative forced-choice task, as was employed, the threshold is taken to be the stimulus strength that yields response accuracy in the neighborhood of 75%. If we determine the stimulus level at which the probability of two

consecutive correct responses equals the probability of the complementary outcome we have $P^2(\text{correct}) = 1 - P^2(\text{correct})$, or $P(\text{correct}) = \sqrt{1/2} = .707$. The threshold measurements were effectively a search for the contrast level where this condition was met for each exposure setting. This contrast threshold search for a particular exposure duration proceeds as follows. Responses to trials are examined in pairs. If two consecutive trials result in hits for a given contrast level, the contrast is decreased by one step for the next pair. All other response configurations for a pair of trials indicate increasing the contrast by one step for the next trial pair. If the first response for a trial pair was a miss we increase the contrast by one step without bothering to measure the second response for that pair, since two consecutive hits cannot occur. This procedure was carried out independently for the five exposure durations that were randomly interleaved. Based on pilot work, a .2 $\log_{10}$ unit change in contrast was chosen as the step size.

If one records contrast level for an exposure trial as a function of trial number for all presentations of that exposure within a session, using this procedure, one encounters local minima and maxima, i.e. changes in direction from increasing contrast to decreasing contrast and vice versa. These are referred to as peaks and valleys. The mean of the n peaks and valleys provides the appropriate estimate of the 70.7% accuracy level.

A run is defined as a strictly increasing or decreasing sequence of contrast changes between a peak and a valley. There are n runs corresponding to the n peaks and valleys for a given exposure duration within a session. If the sequence of runs is segmented into pairs we

can calculate the mean of the corresponding averages of the peak and valley in each pair. Of course, this gives the same estimate of the 70.7% level as found by simply taking the average of the n peaks and valleys. However, as Wetherill points out, this also enables us to form an estimate of the variance associated with our threshold estimate. This segmentation procedure was used in analyzing the experimental data.

A session was completed when 6 or 8 peaks and valleys were recorded for all five exposure durations. A decision about the number of peaks and valleys was made prior to the session. Only data taken after the first valley was used in establishing threshold estimates to avoid bias introduced by starting point selection.

Two or three sessions were used for sensitivity measures for a particular spatial frequency. Each replication was separated by about a week.

One further comment about the contrast step size should be made. A contrast step of .2 $\log_{10}$ units is fairly easy to achieve using our polarizer arrangement for all except the very lowest contrast values. In this extreme situation the very rapid change of polarizer-analyzer attenuation as a function of polarizer angle necessitated a somewhat larger contrast step. This could introduce a slight overestimation of contrast sensitivity under certain pathological conditions. To guard against such a bias contaminating our results, Wetherill estimates were checked against psychometric function estimates. Both procedures yielded estimates within one standard error of each other, and systematic deviations were not evident. The likelihood of bias in the lowest contrast situation can be disregarded.

Results:

The change in contrast sensitivity as a function of exposure duration for both observers is displayed in Figures 4 and 5. The points shown are the means for all sessions. Regression analysis performed for each spatial frequency made use of the segmented runs described above. The resulting curves are shown in these figures.

For both observers, the relationship between log-sensitivity (log s) and log-exposure (log t) is linear within the range of exposures studied for all frequencies except 2.5 cycles/degree. For this frequency, the relationship between log s and log t had a significant quadratic component ($p < .05$) for both observers, indicating saturation of the sensitivity function. The regression coefficients are of particular interest for our consideration of relative processing rates. They are summarized in Table 2.

Table 2

Values of k_L and k_Q in the relation

$\log s = k_L \log t + k_Q (\log t)^2 + \text{constant}$

including 95% confidence limits for linear cases

frequency	S.B. k_L	S.B. k_Q	B.B. k_L	B.B. k_Q
1.25 cycle/deg	.51 ± .19	0	.37 ± .14	0
2.5 cycle/deg	1.64	-0.34	1.34	-0.24
5 cycle/deg	.42 ± .12	0	.33 ± .11	0
10 cycle/deg	.57 ± .18	0	.75 ± .12	0
20 cycle/deg	.85 ± .23	0	.61 ± .26	0

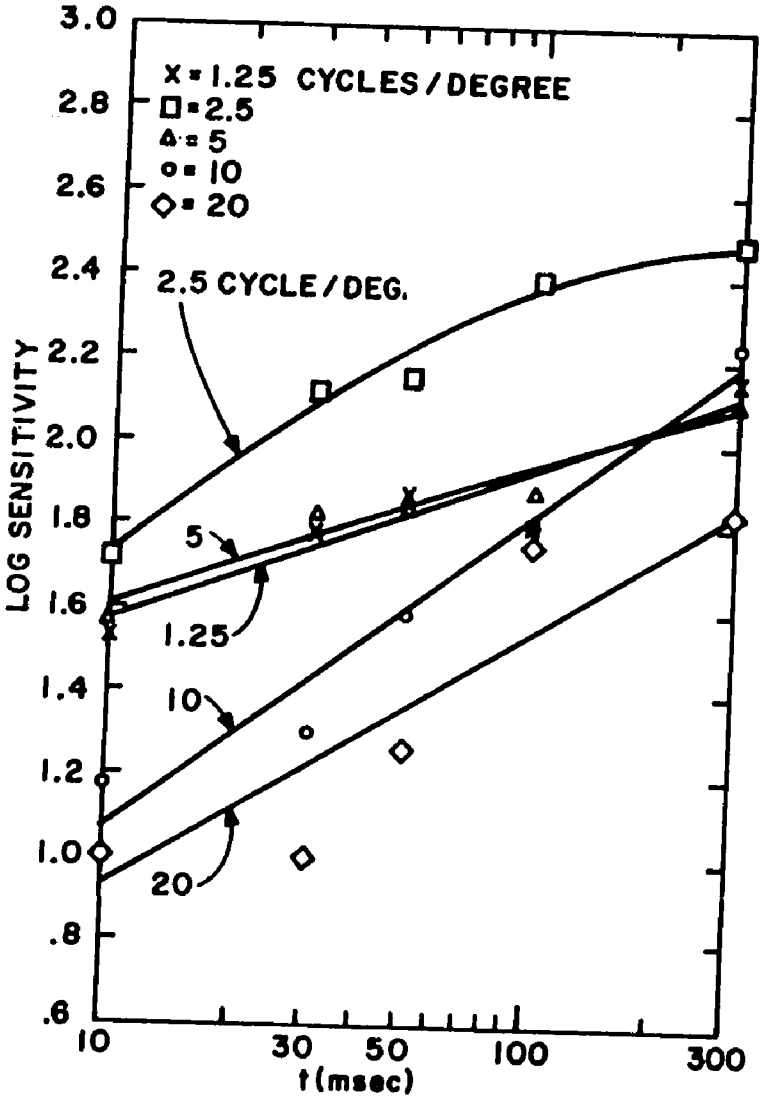


FIGURE 4. Log contrast-sensitivity as a function of exposure duration for five spatial frequencies. Results for observer B.B.

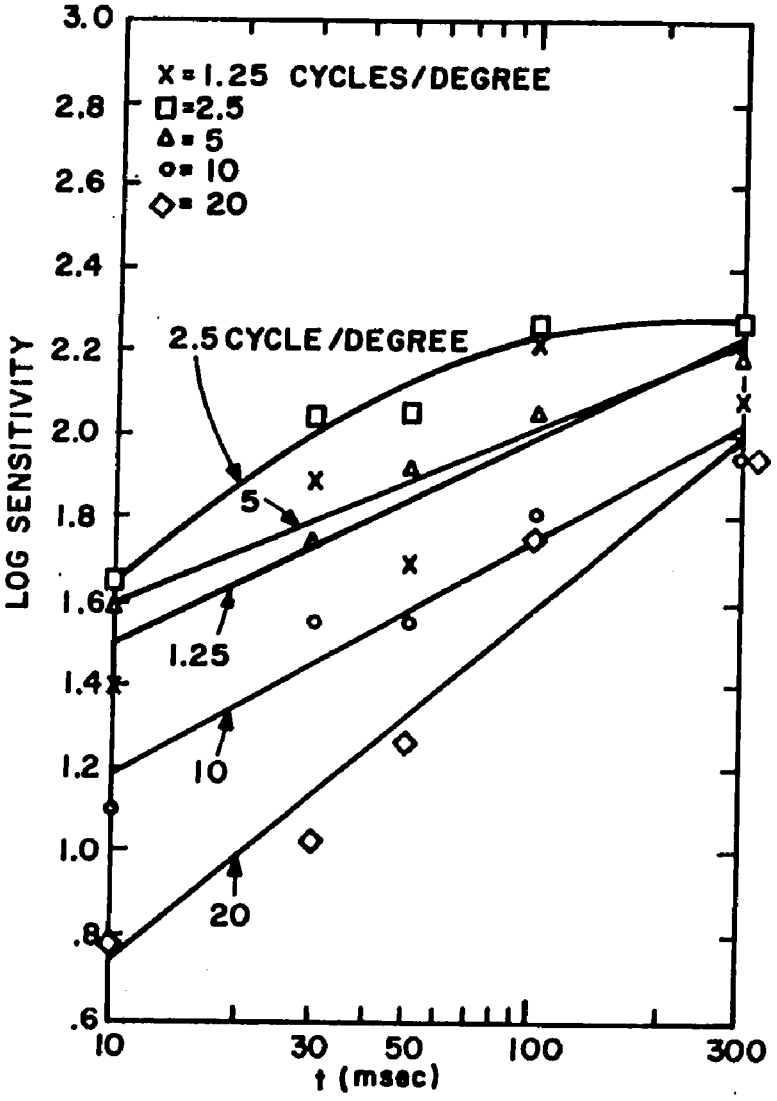


FIGURE 5. Log contrast-sensitivity as a function of exposure duration for five spatial frequencies. Results for observer S.B.

At first glance it might seem that a linear relation between log s and log t, as occurs in all but the 2.5 cycle/degree case, is simply an instance of Bloch's law. That is, for sufficiently brief exposures one should be able to trade off grating contrast and exposure duration so that Ct = constant without affecting grating visibility. However, this would require log (Ct) = log (t/s) = constant, which would constrain the k_L values for the linear cases of Table 2 to be unity. Except possibly for the rather noisy 20 cycle/degree k_L estimate for observer S.B., we find k_L values strictly less than one. This agrees with the results of Nachmias (1967) in that the high frequency log s - log t slope for both of his observers was about .75 for exposures less than 200 msec. The range of exposures for which Bloch's law holds for grating visibility still needs to be firmly established, however it is not altogether unlikely that it is appropriate only for exposures briefer than those measured, i.e. for less than 10 msec., so that a non-negative log s - log t slope less than one would be found for longer exposures (Nachmias, personal communication).

The main conclusion we may derive from our results is the fact that 2.5 cycles/deg is the most rapidly processed spatial frequency of those studied. For exposures not exceeding 300 msec the sensitivity of 2.5 cycles/deg shows signs of reaching an asymptote, whereas there is no evidence of saturation of the sensitivity-exposure relation for the other frequencies.

Inference about the relative processing rates for frequencies

other than 2.5 cycles/degree can be made by comparing the values of k_L in Table 2.

Consider, for example, the 5 cycle/deg and 20 cycle/deg slopes in the log s - log t relation of observer S.B., i.e. consider their two values of k_L in Table 2. The slope of the 20 cycle/deg curve is significantly greater than that of the 5 cycle/deg curve (their 95% confidence intervals do not overlap). This fact is equivalent to the statement that an increase in exposure beyond 10 msec benefits the detection of a 20 cycle/deg grating more than it benefits detection of a 5 cycle/deg grating. In other words, given a very biref grating exposure, say 30 msec long, the visual system transfers a greater percentage of the amount of information that could be obtained in a 300 msec exposure for a 5 cycle/deg grating than for a 20 cycle/deg grating. Assuming the processing of such information proceeds in a generally continuous manner, i.e. temporal quantization is not employed, we can conclude that frequency information in the neighborhood of 5 cycles/deg is processed earlier than frequency information in the neighborhood of 20 cycles/degree. Thus, a relatively large linear slope in the 30-300 msec range is taken to indicate a less complete processing for brief exposures. The concept of relatively less complete processing will be taken here to be equivalent to the concept of relatively slow processing rate.

It can be argued that our k_L estimates are too noisy to make completely convincing statements about the fine detail of processing rates from comparisons of k_L values, regardless of the statistical validity of such conclusions. This point of view does have merit,

nevertheless it does no harm to indicate what is implied by the data provided appropriate caution is exercised when giving a general interpretation to these results.

Continuing with this approach we find for observer S.B. that 1.25 cycle/deg, 5 cycle/deg, and 10 cycle/deg gratings are not processed at significantly different rates, however all three of these gratings are processed more rapidly than a 20 cycle/deg grating.

For observer B.B. we can conclude that 1.25 cycle/deg and 5 cycle/deg gratings are processed at effectively the same rate, both faster than a 10 cycle/deg grating. The slope estimate for the 20 cycle/deg grating for B.B. was somewhat noisy, as revealed by the broad confidence interval. It was necessary to compare the k_L value for this frequency with the other k_L estimates by pooling the variance for each comparison and submitting the k_L difference to a t test. This yielded a 20 cycle/deg slope not significantly different from the 10 cycle/deg slope, and significantly larger than either low frequency slope. The 20 cycle/deg grating is therefore processed at the same rate as the 10 cycle/deg grating, and slower than the 1.25 cycle/deg and 5 cycle/deg gratings for observer B.B.

The general trend of spatial frequency processing rates for our observers seems consistent with the findings of Nachmias (1967) cited above. A very low frequency grating can be said to be processed more rapidly than a very high frequency grating. The fine detail of processing rates across the frequency spectrum does not seem so easily captured by our results, however we do not feel justified in rejecting the hypothesis that processing rate is an increasing

function of spatial frequency. Even so, we will still have to be a bit careful in any attempt at establishing a corespondence between the rate of single frequency processing for visibility and frequency processing for recognition. This is the goal of our second experiment, which we will turn to shortly.

In this experiment we have used two independent variables, spatial frequency and exposure duration, and one dependent variable, contrast sensitivity. For each curve in Figures 4 and 5 we kept one of the independent variables, spatial frequency, constant and determined the best least squares fit for the relationship between the remaining two variables. This makes good sense, since this is how the experiment was performed. It is of interest to examine what these derived relationships suggest about the relation between spatial frequency and contrast sensitivity with exposure duration held constant. In doing this we will not have as precise a statement as could be obtained by performing a polynomial regression on frequency with exposure held constant, or by performing the even more complicated multiple regression of the sensitivity dependence on both exposure and frequency. However, the contrast sensitivity functions directly implied by the curves of Figures 4 and 5 should be no less informative than such functions determined by the more conventional procedure of applying a piece-wise linear fit to a log-log plot of the data on either side of the sensitivity peak.

The separate contrast sensitivity curves for 10, 30, 50, 100, and 300 msec exposures derived from the curves of Figures 4 and 5 are displayed in Figures 6 and 7. Included with each point estimate

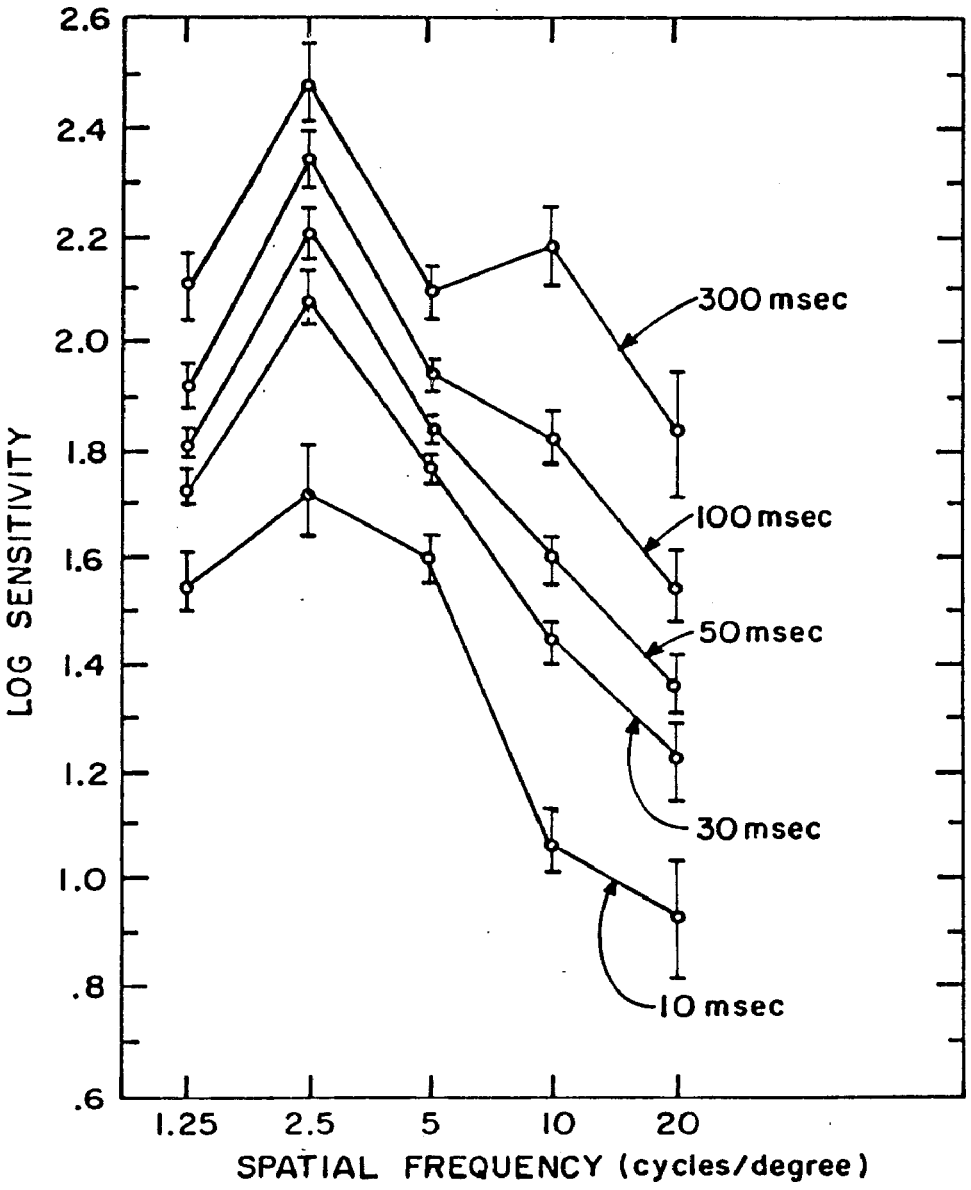


FIGURE 6. Log contrast-sensitivity as a function of spatial frequency for five exposure durations. These curves are derived from the results displayed in Fig. 4 for B.B. See text for details.

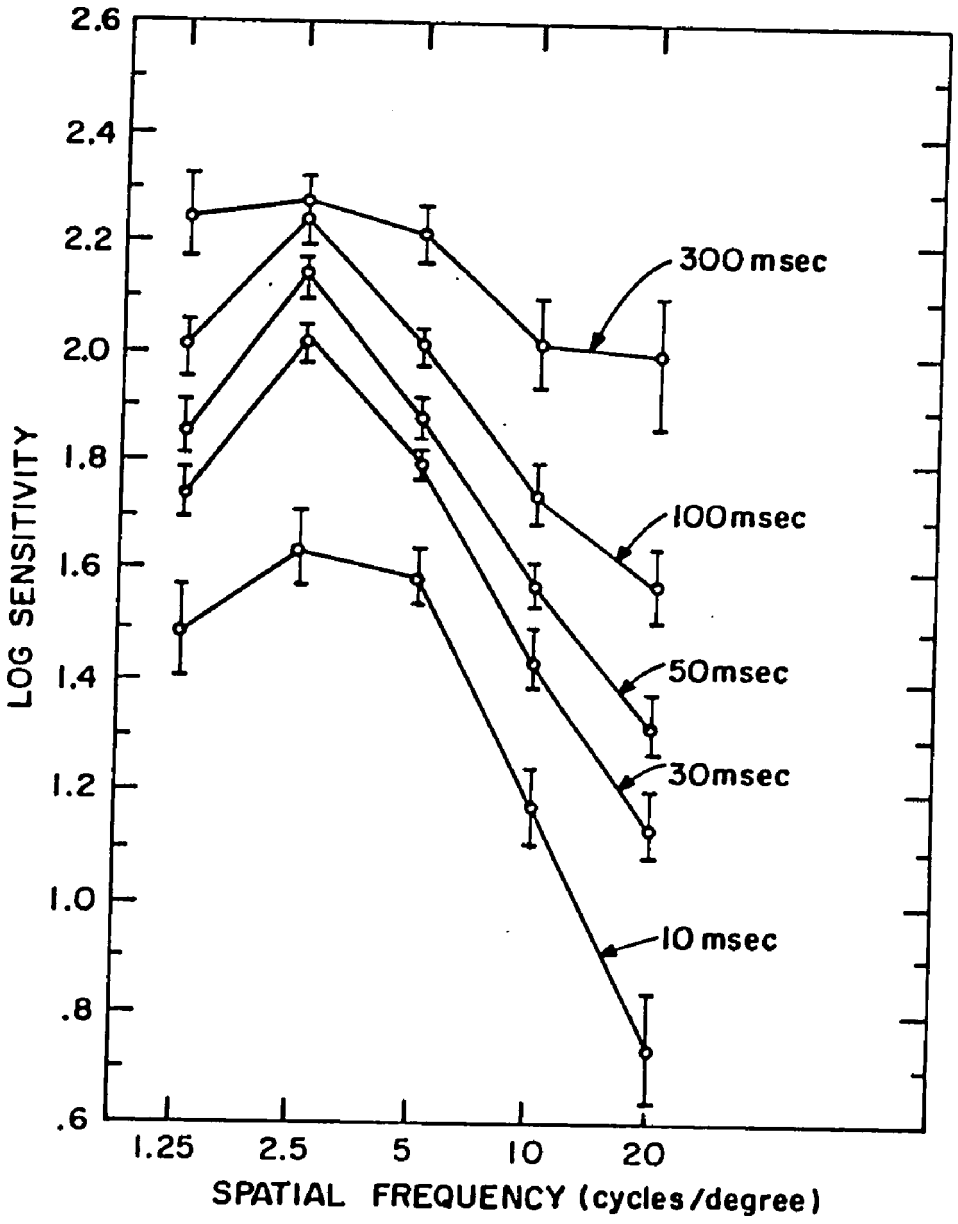


FIGURE 7. Log contrast-sensitivity as a function of spatial frequency for five exposure durations. These curves are derived from the results displayed in Fig. 5 for S.B. See text for details.

for these curves are the appropriate standard error bars (± 1 S.E.).

The derived contrast sensitivity functions for both observers are similar, except for the 300 msec branches of the families of curves. This could reflect real differences between these observers, however it might simply suggest that one of the other procedures mentioned above, such as multiple regression, would yield a better rendition of the 300 msec curves. In any event, a general comparison between these sets of curves and the contrast sensitivity functions for a variety of exposure durations found by the previous investigators can be made.

The most apparent feature of the curves in Figures 6 and 7 is the existence of a well defined peak and low frequency roll-off. Within the sensitivity of our five frequency measurements, the location of the peak does not change with exposure duration, but shows signs of becoming absent as exposure is decreased below 10 msec.

We have displayed the corresponding results of Schober and Hiltz (1965) in Figure 1 above. They plotted contrast thresholds, with a spatial frequency scale in lines/min. of arc. Inverting their curves and applying the conversion factor 0.01 lines/min of arc equals .6 cycle/deg is the appropriate transformation for comparison with our results (aside from an additional scale factor on the sensitivity axis). The usual sensitivity peak found in the transformed curves begins to disappear for exposures less than about 20 to 40 msec. An additional effect that is most prominent in the data displayed in the right panel of Figure 1 is a migration of the

location of the peak as exposure is decreased. For this observer the peak drifts from about 4.2 cycles/deg for 1 second exposures to about 2.1 cycles/deg for the briefest exposures.

These two attributes are also found in the contrast sensitivity functions reported by Nachmias (1967). These results are displayed in Figure 8. The migration of the location of the peak is most evident for the results displayed in the right panel of this figure. The observer's data displayed in the left panel also shows the disappearance of the sensitivity peak for very brief exposures.

Our results share the characteristic diminution of the peak in the sensitivity curve for very brief exposures found for some of the observers examined by Nachmias and by Schober and Hiltz. However we find no evidence of the peak migration found by these investigators . This difference might simply be due to a lack of sensitivity in our measurements because of our rather small sample of spatial frequencies tested. On the other hand, there are two important differences between our stimuli and those of Nachmias and of Schober and Hiltz which could be the source of this difference. We used green illumination, whereas the other investigators used white light; and we used sine wave gratings.

The issue of color is potentially significant for brief exposures, Different time constants for the different color mechanisms have been documented for steady state flicker (e.g. Green, 1969, Kelley, 1974) and for single pulses (Cicerone, 1974). Furthermore, it is known that the contrast sensitivity peak for spatial frequencies occurs at different locations for the different color mechanisms

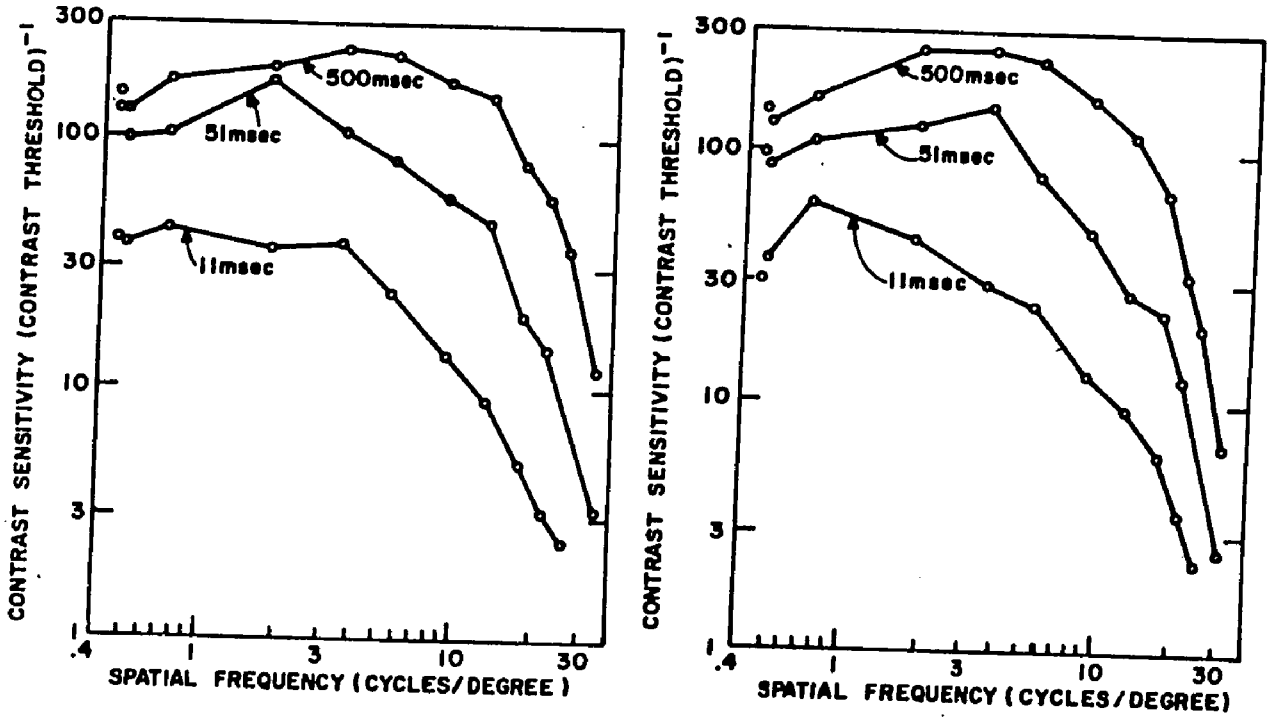


FIGURE 8. Contrast sensitivity as a function of spatial frequency for three exposure durations. Results for two observers in the study by Nachmias (1967).

(e.g. Green, 1968, Kelley, 1974). Hence the possibility emerges that the shift in the peak of the contrast sensitivity function for brief achromatic grating flashes is due to differences in the processing rates of color mechanisms that are unimportant for longer exposures. We would, therefore, not observe such a shift in our measurements, since only one color mechanism is tapped for all exposures. However, conditions of chromatic adaptation are usually required for the color mechanisms to isolate themselves (a notable exception is the Benham's top illusion and associated effects reported by Festinger, et.al., 1971). Since chromatic adaptation was not used in the research we have cited, it seems safe to disregard a color mechanisms explanation of the peak shift.

We have commented earlier about the multiple confounding that exists in spatio-temporal sensitivity measurements using square wave gratings. Since, as has been demonstrated, the harmonics that go into a square wave are not necessarily processed at the same rate and, as been pointed out, the detection of the first harmonic is potentially as important as detection of the fundamental for low frequency square wave threshold determination, it is difficult to evaluate the significance of a shift of the peak in contrast sensitivity functions for such patterns. Hence, a likely source of the lack of a peak migration in our results and the occasional existence of such an effect in the results of Nachmias and of Schober and Hiltz is that we used sine wave gratings rather than square wave gratings.

It must be remarked, however, that we cannot state with complete confidence that the peak migration does not occur for sinusoidal

gratings until finer measurements with a larger sample of spatial frequencies are made.

Our intent in deriving the contrast sensitivity functions shown in Figures 6 and 7 directly from the sensitivity-exposure curves displayed in Figures 4 and 5 was to provide a simple check on the accuracy of the log s - log t relations. Since the resulting sensitivity curves seem generally reasonable we can assume that our log s - log t relations are self-consistent.

The aim of this experiment was to establish the visibility processing rates for a number of spatial frequencies so that a test could be made of the applicability of threshold visibility measurements for spatial frequencies to studies of pattern recognition. The relation between processing rate and spatial frequency may be a bit more complicated than we would have hoped for, but a useful test can still be made if we focus our attention on well chosen sets of frequencies. This test is the topic of our next experiment.

CHAPTER 3

PICTURE RECOGNITION WITH NARROW-BAND NOISE MASKING

The demonstration of the potential significance of spatial frequency processing for picture recognition presented by Harmon and Julesz (1973) has been discussed in the introduction of this report. The thrust of their demonstration was to show that the ability of two-dimensional band-limited noise to degrade picture recognizability depends on the location of the noise band in frequency space. An extension of this technique into the time domain makes it a suitable approach for our area of concern. An examination of the temporal properties of the recognizability degradation induced by noise bands well separated in frequency space should yield results consistent with the observed processing rates for threshold visibility of spatial frequencies if single frequency threshold measurements are relevant for an understanding of information processing for recognition.

Suppose we restrict our attention to two spatial frequencies, A and B cycles/degree, such that A cycles/deg is processed more rapidly than B cycles/deg as far as the property of threshold visibility is concerned. Now, suppose that within a brief exposure of a picture stimulus we distort the spatial frequency information in the neighborhood of A cycles/deg, by the addition of band limited noise, and remove distortion for the later part of the exposure. Compare this to the complementary situation, i.e. distort frequency information in the neighborhood of A cycles/deg only for the later part of the exposure. If it is the case that more information is obtained about A cycles/deg in the early portion of an exposure than in the later

portion, then one would expect performance to be better for late distortion than for early distortion. The opposite result should hold for distortion in the neighborhood of B cycles/deg.

From the results of the preceeding experiment it seems that 2.5 cycles/deg and 10 cycles/deg are in good neighborhoods. These two frequencies are clearly processed at different rates for threshold visibility, the lower frequency being processed more rapidly.

It might be argued that temporal integration for brief exposures would render the situations of early distortion and late distortion equivalent as far as visual processing is concerned. It must be pointed out, however, that such an argument assumes the validity of the notion of temporal quantization. It has been demonstrated by Baron (1971), and others, that this notion is not correct.

Materials:

As noted, the distortion of frequency information is accomplished by the addition of band-limited noise to a picture stimulus. The desired noise bands were computer generated.

A master 256 x 256 matrix was produced with each entry being randomly assigned one of 256 equally likely values. These 256 values correspond to a linear scale of transmission. A scene produced from such a matrix is statistically white, i.e. it has a flat energy spectrum. This noise matrix was processed by passing it through two-dimensional pass-band annulus filters. (See Bracewell, 1965, for background material in Fourier analysis and spatial filtering). The filter geometry is displayed in Fig. 9. The filter has unity value in the shaded area and zero value outside the shaded area. The annulus width, Δf, was kept

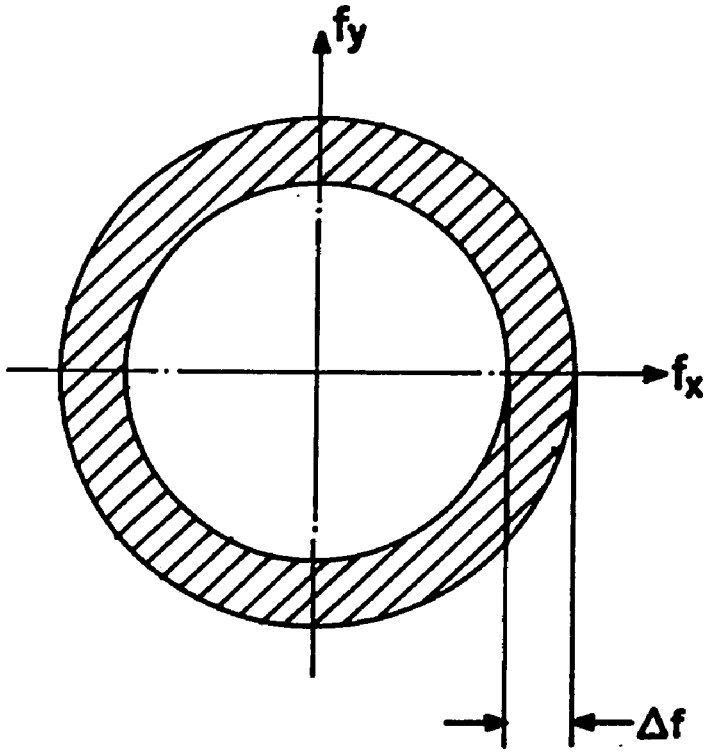


FIGURE 9. Geometry of the annulus filter used in the preparation of two-dimensional, narrow-band spatial noise patterns by digital processing. The filter amplitude is unity in the shaded region, and zero elsewhere. Narrow-band noise was produced by passing two-dimensional white noise through the filter.

constant and the inner radius was varied to place the band in different regions of the Fourier plane.

By using this technique we were able to obtain two-dimensional noise of bandwidth Δf and any desired center frequency. Since we are using this noise to mask pictures which have frequency information at all orientations, the use of two-dimensional noise is desirable. By imposing circular symmetry we are disregarding the possibility that the ordering of spatial frequency processing rates is not the same for all orientations. Such a possibility seems unlikely. Given our ignorance about orientation effects, circular symmetry seems a desirable property.

Hard copy of the noise patterns was realized by reading the data for these noise scenes on magnetic tape into a programmed flying-spot scanner. Hard copy was in the form of slide-transparencies.

The noise bandwidth, Δf in Figure 9, was set at 3.7 cycles/degree for the viewing conditions used. The noise band in the neighborhood of 2.5 cycles/degree extended from 0 to 3.7 cycles/deg (centering the band slightly below 2.5 cycles/deg), and the noise band in the neighborhood of 10 cycles/deg extended from 7.4 to 11.1 cycles/deg (centering the band slightly below 10 cycles/deg). An intermediate band of 3.7 to 7.4 cycles/deg was also prepared for additional investigation.

The noise slides were square. Under the experimental viewing conditions each side of the slide subtended 3.9^{o} visual angle.

The annulus filters employed gave each noise pattern the same bandwidth. The geometry of such a filter necessarily assigns

greater spectral energy to higher frequency noise patterns. Index the continguous, concentric filters we have used so that n = 0 denotes the 0-3.7 cycle/deg filter, n = 1 denotes the 3.7-7.4 cycle/deg filter, etc. Then the area of filter n is $A = (2n+1)\pi(\Delta f)^2$, where Δf is the common bandwidth. Hence the 7.4-11.1 cycle/deg noise band has five times the spectral energy content as the 0-3.7 cycle-deg band. It should be kept in mind that spectral energy has nothing to do with the average transmission of a noise slide, they were all equal in this respect.

This difference in energy between the noise patterns is a direct consequence of the equal bandwidth constraint. In a sense, this is a desirable consequence. As we go to higher spatial frequencies the sensitivity of the visual system decreases. The greater energy of the high frequency patterns tends to compensate for the presumed lower sensitivity to such patterns.

The stimuli were eight picture slides of common objects selected from a large population of candidate slides. The selection procedure involved recognizability performance under conditions of brief exposure for four observers. The pictures chosen were approximately equally recognizable under these conditions. The picture information was centered on each slide, the remainder of the slide being uniformly transparent. Further details on these slides are found in Appendix 2.

The fact that low frequency information is contained in these slides, regardless of quality, is guaranteed by the facts that they are not opaque and that they are not simply silhouettes or line drawings. The presence and importance of higher frequency information

was informally checked as follows. A diffusing sheet in front of a pattern acts as a low-pass spatial frequency filter. By placing each of the produced noise patterns in the location of the picture stimulus and interposing a diffusing sheet at a fixed location between the eye and the pattern one can roughly determine the cut-off frequency of the diffusing sheet in this arrangement. In this manner we chose a diffusing sheet that effectively removed spatial frequencies higher than about 2.5 cycles/degree. When the low-pass filtering was performed picture quality was severely degraded, and temporal thresholds for recognition were roughly three times what they were in the absence of filtering. We were, therefore, confident that the pictures used had considerable, significant high frequency information.

A comment about our repeated use of the same noise patterns to mask picture information is in order. Statistical correlation between distinct spatial regions is absent in a white noise pattern, however non-zero spatial correlations are introduced by spatial filtering. In fact, the operation of spatial filtering is isomorphic to the operation of introducing spatial correlation. Of course filtered noise is still noise, but a question exists as to whether the idiosyncracies of a particular, consistently used noise pattern with non-zero spatial correlation can be used to some advantage that would not be available if the noise pattern was different for every presentation. There is some evidence this concern is justified for very practiced observers, but is not a problem for inexperienced observers, as used in this experiment. Support for this assertion is found in Appendix 3.

Procedure and Apparatus:

By superimposing the projection of a noise slide onto the projection of a picture-stimulus we accomplish the desired addition of noise to a specific region of spatial frequencies of the picture. Replacing the noise slide with an appropriate density neutral filter yields a projected picture with the same average luminance and picture contrast as in the presence of noise. Using such a switching procedure we were able to alter the signal-to-noise ratio in a desired region of the spatial frequency spectrum for a desired fraction of the stimulus exposure without altering average luminance or contrast.

A three channel tachistoscope was used for this purpose (Scientific Prototype, Model GB). Viewing was binocular. One of the channels contained a picture slide, one contained a noise slide, and the remaining channel contained a white, diffuse slide with a small centered dot for fixation. The noise channel and the fixation channel were set at equal luminance both with the photometer described in the preceeding experiment and by flicker photometry. Flicker photometry was used throughout the course of this experiment to insure that this equal luminance condition was maintained. The luminance in the picture channel was set quite low to lengthen the time necessary for recognition under tachistoscopic presentation. The mean luminance of the picture plus noise, or picture plus fixation scene was 3.4 ml. It will be recalled that this was the luminance used in the preceeding experiment.

An experimental trial was composed of one of the two sequences of events depicted in Figures 10a & 10b. For both types of trials the subject viewed the fixation scene and, when ready, pressed a button which initiated one of the two sequences. For the early noise situation, depicted in Figure 10a, the button press brought the picture into view for a predetermined exposure duration τ, and replaced the fixation scene with the noise pattern for $t \in [0, \tau/2]$. Except for this $\tau/2$ interval, the fixation scene was always displayed. For the late noise situation, depicted in Figure 10b, noise replacement occurred for $t \in [\tau/2, \tau]$. The two types of trials differed only in when the noise replacement occurred.

The subject's task was to identify the picture presented from among the eight alternatives. The response was verbally indicated at the termination of the exposure and recorded by the experimenter. The subject was provided with a typed list of the stimuli to use as an aid when decisions were difficult. The selection of a particular picture for any trial was random with the constraint that presentation of the same stimulus on two successive trials was avoided. No response feedback was provided.

Before the experimental sessions began a subject participated in one or two practice sessions to become able to recognize the pictures at brief exposures. Twenty to forty warm-up trials immediately preceeded each experimental session, as well. For the practice and warm-up trials the fixation field was constantly in view; no noise was used.

Each experimental session was devoted to the examination of the effects of a particular noise pattern on picture recognition. For

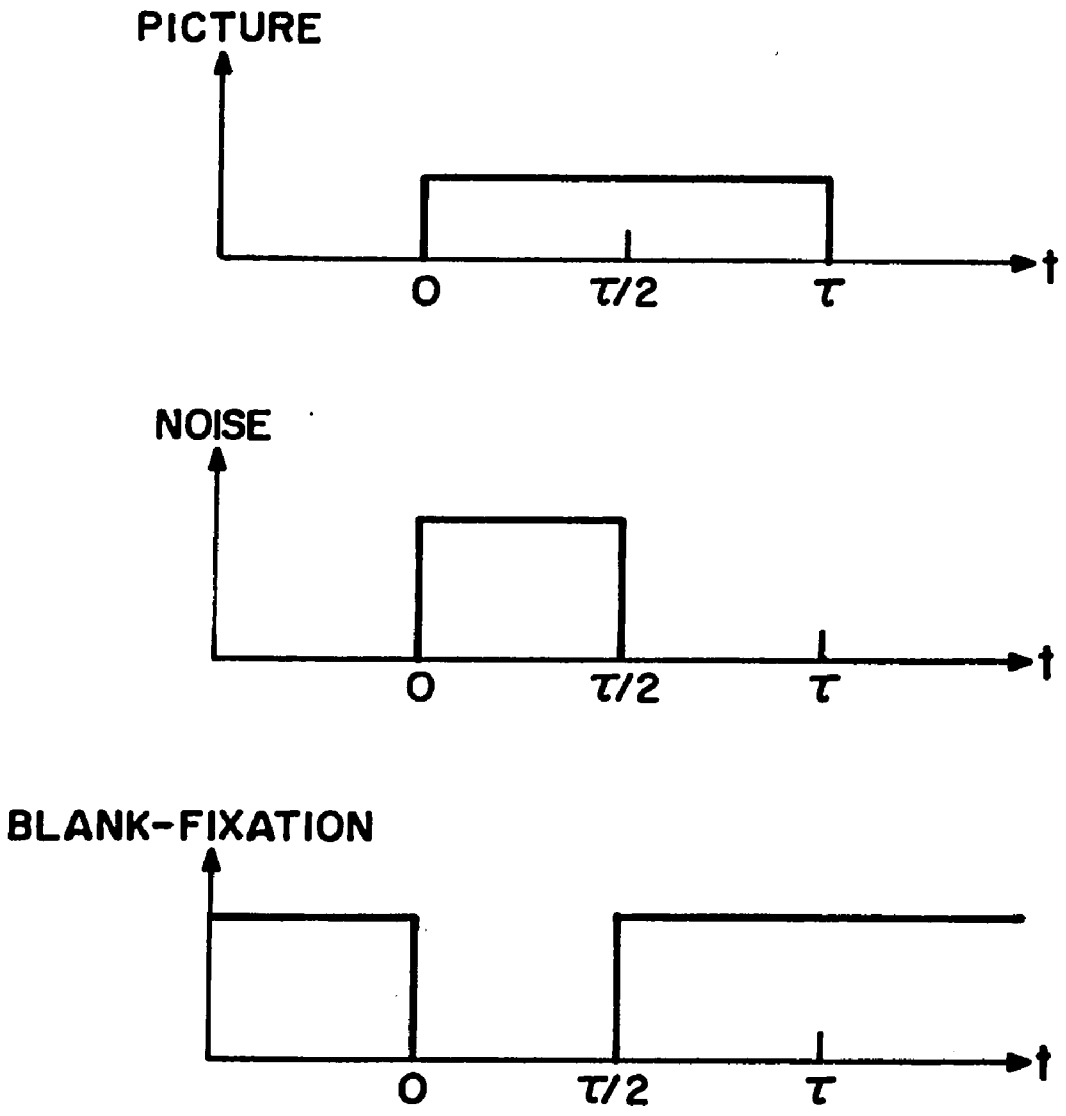


FIGURE 10a. The early noise trial. While viewing the blank-fixation field, the subject pressed a button at time zero which brought the stimulus picture into view for a duration τ, and replaced the fixation field by a noise field for a duration $\tau/2$.

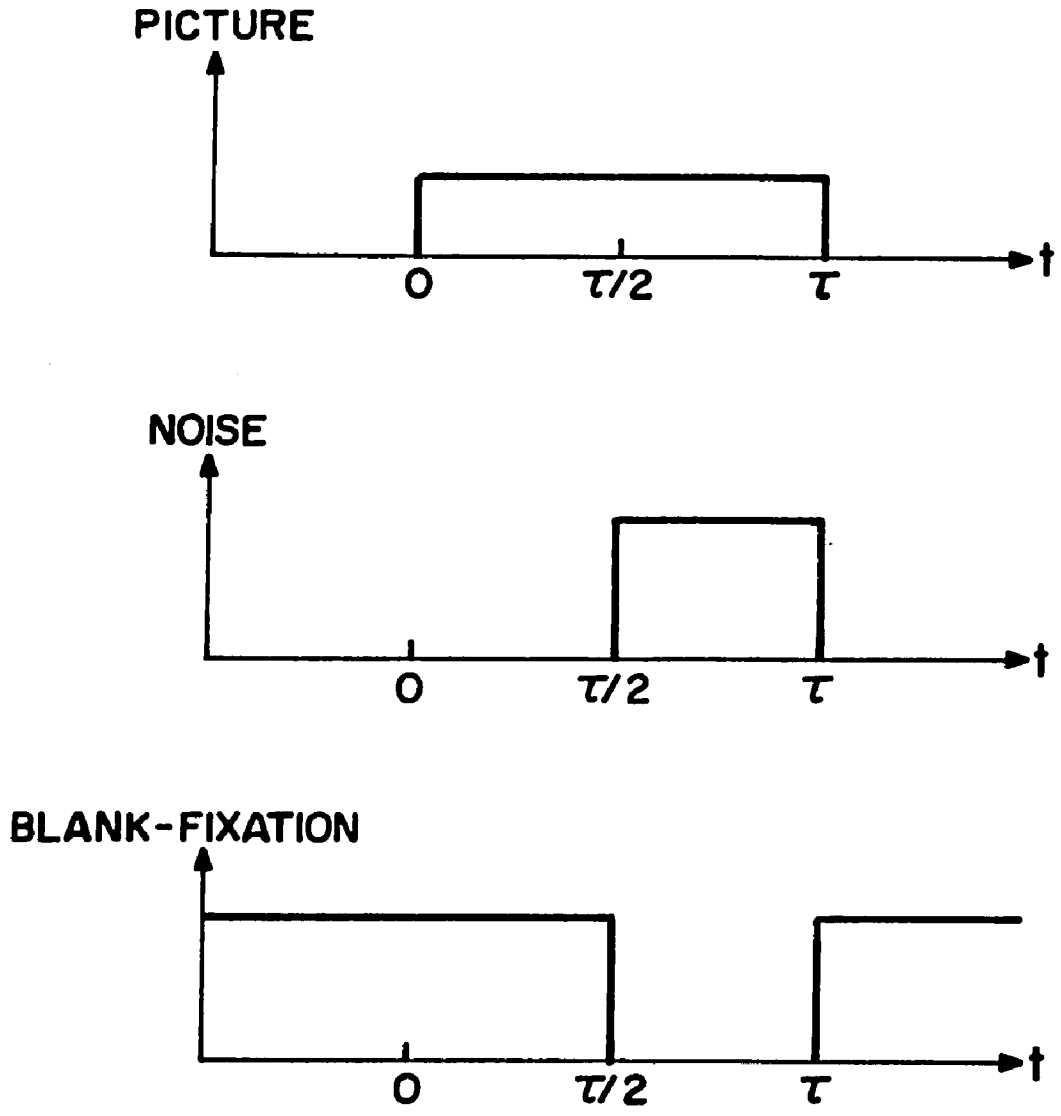


FIGURE 10b. The late noise trial. While viewing the blank-fixation field, the subject pressed a button at time zero which brought the stimulus picture into view for a duration τ. At time $\tau/2$ the fixation field was replaced by a noise field for a duration $\tau/2$.

the first six subjects only early noise and late noise trials were considered. There were forty stimuli of each trial type, hence there were eighty experimental trials in a session. For the remaining twelve subjects it was decided to include forty trials without noise, i.e. the type of trials used in the practice sessions. For these subjects there were, obviously, 120 trials in an experimental session.

Because of the difficulty and length of time associated with resetting the tachistoscope for the various noise presentation schemes, noise trial types were randomized in blocks of four trials. There were ten blocks of early noise, ten blocks of late noise, and, for twelve subjects, ten blocks of no noise, all randomly interspersed in a session.

As we have indicated, each session began with a series of no-noise warm-up trials. Then, after a brief break, the subject was shown the set of pictures slowly for examination following which she was tested for the exposure duration necessary to achieve roughly 71% recognition accuracy in the eight-alternative forced choice task for no noise. The technique used for exposure testing was the temporal equivalent of the Wetherill estimation procedure used in the preceeding experiment. After the appropriate exposure duration was determined, the subject was allowed to examine the pictures again before data collection trials began. Following presentation of half the trials for a session, the subject was given a five minute break. A typical complete session took about an hour.

Subjects were tested in separate sessions, with either the 0-3.7 cycle/deg or 7.4-11.1 cycle/deg noise pattern being used in a

session. An additional session for the 3.7-7.4 cycle/deg noise pattern was included for the first twelve subjects. The results with this pattern were generally uniformative. To save time, this condition was deleted for the remaining six subjects. Further comment and results for this condition are found in Appendix 4.

The type of noise used in a particular session was counterbalanced among subjects.

Subjects:

Eighteen females, ages 18-24, were employed for this experiment. Twelve of these subjects were recruited from the subject pool of the Human Performance Center of the University of Michigan, and six of the subjects were recruited from the general population of Princeton, New Jersey, where the experiment was concluded.

Results and Discussion:

Results of low and high frequency noise effects for the eighteen subjects are summarized in Table 3.

Table 3

Experimental Summary

	Early Noise	Late Noise	No Noise	$\bar{\tau}$ (Average Exposure in msec)
Low Frequency Noise	73.2%	73.1%	69.8%	28
High Frequency Noise	61.4%	69.1%	67.7%	27

The entries indicate response accuracy for the experimental task. The high and low frequency no-noise conditions represent identical experimental situations, i.e. picture identification with picture

flash superimposed on the fixation field. However, the exposure duration used for a particular subject was typically different for different sessions, i.e. for different noise types. It therefore makes sense to make a distinction between the corresponding no-noise results.

It was pointed out that no-noise trials were not included for six of the subjects. Following an arcsine transformation of the data collected for all eighteen subjects, a least squares estimate of the missing no-noise data was obtained (following the procedure outlined by Cochran and Cox, 1957). All data analysis for this experiment made use of the variance stabilizing arcsine transformation.

For the high frequency noise condition, analysis of variance indicated significant treatment effects ($F(2,28) = 3.81$, $p < .05$), after correction due to bias introduced by missing data estimation (see Snedecor & Cohran, 1967, for correction procedure). It may be noted that the bias correction was very small. The Newman-Keuls test was used to examine particulars about differences between treatment effects (Winer, 1971). From this analysis it was determined that there was no significant difference between the effects of late noise and the effect of no-noise, i.e. late high frequency noise is an ineffective mask. However, the effect of early presentation of high frequency noise is important in that the recognition accuracy under this condition is significantly less than either no-noise or late noise ($p < .05$ in both situations).

Analysis of variance for the low frequency noise condition revealed a lack of significance for treatment effects. Half exposures of low frequency noise have no observed effect on the

recognizability of briefly viewed pictures for the situation we have explored.

Our results are somewhat surprising on two counts. There appears to be no indication from single frequency sensitivity data of the preceeding experiment that sensitivity should be greater for the 7.4-11.1 cycle/deg noise pattern than for the 0-3.7 cycle/deg noise pattern (cf. Figs. 6 & 7). However we find a masking effect for the high frequency noise and none for the low frequency noise. Furthermore, the temporal locus of the high frequency masking effect does not seem consistent with what is indicated by the single frequency processing rates derived in the preceeding experiment. These processing rates indicated that spatial frequency information in the neighborhood of 10 cycles/deg is a relatively late component of visibility processing. If it is the case that picture recognition involves processing of all spatial frequencies contained in the picture, and if the relative processing rates for threshold visibility are accurate indicators of processes involved in recognition, one would have expected that the disruptive effects of high frequency noise would be more pronounced in the later part of the stimulus viewing period. However, the observed disruptive effect occurs for the early part of an exposure.

A fairly strong linearity assumption is implicit in using data on sensitivity to single frequencies in one-dimension to make predictions about the masking effects of a two-dimensional band of noise. For example, this assumption is necessarily involved in making the very reasonable statement that a low frequency noise pattern should mask low frequency signals at least as well as a high

frequency noise pattern masks high frequency signals. Pollehn and Roehrig (1970) present evidence that this is decidely not the case.

Using filtered noise of bandwidth 1.3 cycles/deg, they found that a high frequency band of noise is quite effective for masking sine wave gratings within the frequency range spanned by the noise band, and also has a considerable masking effect for low frequency gratings. In this situation, the term masking refers to an induced relative threshold elevation. The optimal masking occurred for noise centered at about 8 cycles/deg. Centering the noise band at lower spatial frequencies yielded negligible masking. Interestingly, ineffective noise masking held even when the band was centered near what they determined to be the region of peak grating sensitivity for their viewing conditions, 2.5 cycles/deg.

These effects reported by Pollehn and Roehrig are consistent with our finding that picture recognition suffers more from high frequency masking than from low frequency masking. It seems to be the case that when zero-bandwidth stimuli, i.e. simple sine-wave gratings, are used visual sensitivity to spectral regions is very different than for non-zero bandwidth, suprathreshold stimuli. In our view, disinhibition is a likely source of this difference.

A two-dimensional band of noise contains energy at all orientations and over a large set of neighboring frequencies (theoretically the set is infinite, but in practice this is not accomplished). If inhibition among neighboring detectors is operative, one would expect the degree of inhibition to be monotonic with the degree of excitation of the inhibiting detector. A strong possibility, then, is that since low frequency detectors are relatively quite sensitive, the

simultaneous excitation of a neighborhood of these detectors would cause considerable lateral inhibition. There would result a net disinhibition to the effect that the masking of a single frequency in low frequency noise would be less than one might have expected in the absence of inhibition. Similarly, since high frequency detectors are individually less sensitive the degree of disinhibition would be less pronounced, and masking would be more effective. We have discussed evidence for disinhibition between spatial frequency detectors in adaptation experiments earlier (i.e. Tolhurst, 1972). It is not altogether unlikely that such an effect is revealed in the experiment of Pollehn and Roehrig, and in ours.

Unfortunately, there seems to be a flaw in the particular procedure used by Pollehn and Roehrig to generate bandlimited noise that could be a source of concern in evaluating their results. Two-dimensional signals were displayed on a CRT face by laying down one-dimensional scan lines from the top to the bottom of the screen. This is, of course, the standard technique. One gathers that filtered noise was displayed by passing the output of a white noise source through a band-pass filter and laying down this processed noise on the consecutive scan lines. For the orientation parallel to the scan lines the resulting luminance distribution has the spatial frequency properties desired by the authors. However, at any other orientation the situation is quite different. For example, at the orientation perpendicular to the scan lines the noise is effectively white since the luminance values for any two distinct points on a line at this orientation are uncorrelated. The distribution of spatial frequencies for some orientation that is neither perpendicular nor parallel to

the scan lines is neither white nor restricted to the intended band.

The test stimuli used for the masking studies by Pollehn and Roehrig were one-dimensional sine-wave gratings with luminance information parallel to the scan lines. If one is willing to postulate that only noise at the orientation of the grating is effective in masking, then their results can be accepted without qualification. This orientation independence premise is valid for threshold stimuli (e.g. Gilinsky, 1968, Blakemore & Campbell, 1969a), however the masking noise was suprathreshold. Orientation cross-talk is, therefore a potential artifact. It is not at all clear how the presence of such an artifact could be the source of negligible low frequency masking while maintaining considerable high frequency masking. Nevertheless, the possibility of orientation cross-talk indicates caution in citing this work as supportive of our results.

However, the work of these researchers does support our contention that a linearity assumption deserves closer scrutiny. Modification of spatial frequency sensitivity due to interaction of neighboring frequency detectors or neighboring orientation detectors means that linearity does not hold. In this case we need to know much more than the processing attributes of single spatial frequencies for visibility to accurately characterize the effects of a set of suprathreshold frequencies, such as bandlimited noise.

To avoid possible confusion, it should be pointed out that we are using the term linearity in the sense that it is used to define a linear system. In systems theory the property of linearity is

defined as follows: if some input A produces an effect A' and some other input B produces an effect B', then the simultaneous input of $\alpha A + \beta B$ should produce the effect $\alpha A' + \beta B'$, where α and β are arbitrary constants. We do not claim that the modification of sensitivity by the interaction of neighboring detectors cannot be described by a set of linear relations. Of course this can be done, the modeling of lateral inhibition in the limulus being a prime example (e.g. Ratliff, 1966). We are simply taking linearity to mean the absence of significant interactions.

This issue is brought into sharper focus by a closer examination of the connection between the single frequency processing study in our first experiment and the picture recognition task of the current experiment. In measuring the contrast threshold for a particular frequency grating for a particular exposure duration, as in experiment 1, we are in essence asking for the minimum signal-to-noise ratio necessary for detection, where the signal is represented by the grating and the noise is represented by a blank field. In picture recognition, however, the situation is a bit different. The signal is some complex of frequencies that signify important picture characteristics, and the noise is some different complex of frequencies that provide no discrimination information. It is not strictly necessary that the properties of signal extraction from a blank background bear a close resemblance to the properties of signal extraction from a non-blank background.

Suppose, for example, we asked how the results of experiment 1 would be affected if the observer was not asked to indicate the

presence of a grating against a blank background, but to indicate the presence of a grating against the background of a different frequency grating. Qualitative differences between performance in such a task and the blank background task would provide some insight into the apparent inconsistency between the results of our first two experiments.

CHAPTER 4

SINGLE FREQUENCY VISIBILITY PROCESSING AGAINST A NON-BLANK BACKGROUND

The luminance distribution viewed on every exposure in experiment 1 was of the form

$$L = k\,[1 + \alpha\, M_a \sin \omega_a x]$$

$$= k\,[\alpha\,(1 + M_a \sin \omega_a x) + (1-\alpha)] \qquad (1)$$

The proportionality constant, k, indicates the absolute luminance level. To simplify matters we will assume that L is normalized so that k = 1 in the following discussion.

The luminance distribution of the test channel was $\alpha(1+M_a \sin\omega_a x)$, and the luminance distribution of the blank channel was $(1 - \alpha)$ in that experiment. For exposures where no signal was present α was set to zero. The experiment was performed to find the minimum value of α that allowed an observer to discriminate between the luminance distribution of the test exposure and the luminance distribution of the blank exposure. The values of M_a and ω_a were fixed in an experimental session. Determination of the threshold value of α, α_t, gave us the contrast threshold, $\alpha_t M_a$.

In the present experiment we examine a straightforward generalization of the situation in experiment 1. The test channel luminance distribution is $\alpha(1 + M_a \sin \omega_a x)$, as before. The non-test channel luminance is now, however, $(1 - \alpha)\,(1 + M_b \sin \omega_b x)$. As M_b goes to zero we have the situation of the previous experiment. So, for this experiment the normalized luminance for any exposure is of the form

$$L' = \alpha(1 + M_a \sin \omega_a x) + (1 - \alpha)(1 + M_b \sin \omega_b x). \qquad (2)$$

For exposures when the test grating is absent α is set to zero, and the observer views a single grating of frequency ω_b. The experiment was performed to determine the minimum value of α that enabled the observer to detect the presence of the grating of frequency ω_a. It should be noted that changing α does not change the mean luminance.

It is worth observing that equation (2) is simply a special case of the expression

$$L'' = \sum_i \alpha_i (1 + M_i \sin \omega_i x), \quad \sum \alpha_i = 1, \quad M_i \in [0, 1]$$

$$= 1 + \sum_i \alpha_i M_i \sin \omega_i x$$

With proper choice of the M_i we have the Fourier expansion of any odd periodic function. In this experiment we are considering the case $M_i = 0$, $i \neq 1, 2$. In experiment 1 we set $M_i = 0$, $i \neq 1$. It is clear, therefore, that this experiment is a logical extension of the previous one.

The test grating contrast threshold, $\alpha_t M_a$, is an appropriate visibility measure for examining the hypothesis that grating detectability is not critically affected by the presence of a background grating, particularly if the test and background are separated in frequency by more than an octave. However, if we succeed in demonstrating interaction between test and background gratings by uncovering systematic variations in the threshold measure, $\alpha_t M_a$, with background, we will also have demonstrated that a more accurate visibility measure would be one that explicitly takes into account

the presence of the non-blank background. In this circumstance the test-background signal-to-noise ratio would be more appropriate than the contrast threshold. The signal-to-noise ratio can be viewed as the proper generalization of simple grating contrast for this type of compound grating situation.

Let:

$$S = \alpha M_a \sin \omega_a x$$

$$N = 1 + (1 - \alpha) M_b \sin \omega_b x$$

so that $L' = S + N$, where S denotes the signal and N denotes the noise. This is a reasonable specification, since the signal detection task is to identify exposures for which M_a is non-zero both for the current experiment and for the previous, $M_b = 0$, experiment. Comparison with the first line of expression (1) above may be helpful in convincing the reader that this signal and noise separation is the most sensible. If we let A denote the common width of the gratings, where A is assumed to be an integral multiple of the grating period, we have:

$$\frac{S}{N} = \sqrt{\frac{\frac{1}{A}\int_o^A S^2 \, dx}{\frac{1}{A}\int_o^A n^2 \, dx}} = \frac{1}{\sqrt{2}} \left(\frac{\alpha M_a}{\sqrt{1 + 1/2 \, M_b^2 \, (1 - \alpha)^2}} \right)$$

or

$$\sqrt{2} \, \frac{S}{N} = \frac{\alpha M_a}{\sqrt{1 + 1/2 \, M_b^2 \, (1 - \alpha)^2}}$$

When $M_b = 0$ we have the situation examined in experiment 1. In this case we also have $\sqrt{2} \, \frac{S}{N} = \alpha M_a$, which is the contrast value we would have using the standard definition of contrast for simple gratings.

The signal-to-noise ratio, $\sqrt{2}\,\frac{S}{N}$, is an appropriate measure of generalized contrast, C':

$$C' = \frac{\alpha M_a}{\sqrt{1 + \frac{1}{2} M_b^{\,2} (1 - \alpha)^2}} \qquad (3)$$

We can meaningfully compare the value of C' determined from the threshold value of α for the current, non-zero M_b situation to the threshold contrast determined previously for $M_b = 0$.

The empirical distinction between the generalized contrast of equation (3) and simple contrast, contained in the numerator of equation (3), is not likely to be important. Even in the most extreme condition of $M_b = 1$ and $\alpha \to 0$ the two threshold measures differ by less than .09 $\log_{10}$ units. Therefore, experimental conclusions should not depend significantly on which of the two formulations of contrast is used. Nevertheless, these expressions of grating visibility are arrived at from different theoretical considerations. For the sake of completeness, results will be reported using both types of contrast measure.

Method:

The experimental set-up of experiment 1 was used for this experiment with one change, the 10 cycle/deg grating used in the previous experiment was taped onto the blank channel field stop. The neutral filters previously used in the blank channel to insure constant mean luminance regardless of contrast for the various test gratings were suitably replaced to insure constant mean luminance with the 10 cycle/deg grating in what was the blank channel and the various gratings used in the other channel. The procedures used to

select the appropriate density filters for this purpose in experiment 1 were used again.

The test-noise grating combinations we chose to study are given in Table 4.

Table 4

Test-Noise Combinations

		Noise		
		10 cycles/deg	5 cycles/deg	2.5 cycles/deg
Test	10 cycles/deg	—		
	5 cycles/deg		—	—
	2.5 cycles/deg		—	—

A dash in this table indicates a combination not examined.

By taping either the 2.5 cycle/deg grating or the 5 cycle/deg grating onto the field stop of what was previously the test channel we could study any of the four combinations in this table. For example, with the 5 cycle/deg grating in place and using the polarizers as before, we had the 5 cycle/deg test, 10 cycle/deg noise situation. Using the same arrangement we could also examine the 10 cycle/deg test, 5 cycle/deg noise situation by reversing the sense of the polarizer adjustments and redefining the task for the observer.

All grating pairs were in phase, as determined by microdensitometer recordings of grating transmission. The recordings were those used to determine the modulation depth of the gratings for experiment 1, as reported in Table 1 earlier.

Attention was restricted to two exposure durations for this experiment, 30 msec and 300 msec. This was deemed adequate to expose

potential qualitative differences in spatial frequency processing characteristics between the current $M_b \neq 0$ situation and the previous $M_b = 0$ situation. This will become more apparent in the following discussion of experimental results.

Experimental sessions were identical to those of experiment 1, except that the observer's task was now to report on exposures that contained the test grating, and only two exposure duration types were tested in a session. The use of two, rather than the previous five, exposure durations made it feasible to use complete randomization of the exposure setting for each trial. It was pointed out, in discussing procedures for experiment 1, that the cumbersome bookkeeping required to keep track of five interleaved exposure settings made it necessary to occasionally avoid changing the exposure setting until a pair of trials was completed. Bookkeeping was far simpler in this experiment, so randomization constraints were unnecessary.

Each session was devoted to measurements on one of the test-noise combinations of Table 4. Each session represented by an entry in this table was recplicated three times for each observer. All table entry combinations were tested before a particular test-noise combination was put to its next replication, however the order of progression through Table 4 was different for each replication.

As before, threshold estimates of α were derived using the sequential technique of Wetherill & Levitt (1965). It turned out that the psychometric function check used to convince ourselves that the non-uniformity in contrast step size for the lowest contrast polarizer settings did not bias our threshold estimates for the

previous $M_b = 0$ case, was not necessary here. Thresholds were in the equal step-size range. The relatively higher thresholds found for $M_b \neq 0$ are an important feature of the results to be discussed.

To convince ourselves that no trivial artifact, such as the fact that only two types of exposure were examined in a session rather than five, could be the source of differences between the results of this experiment and those of experiment 1, one of the observers was submitted to an additional replication of this experiment, with the noise field replaced by the blank field of the previous experiment using proper adjustment for constant mean luminance. This was done after completing the three replications of Table 4. The results of this test were entirely consistent with the results of experiment 1. Observed differences between $M_b = 0$ and $M_b \neq 0$ results should, therefore, not be due to any obvious procedural artifacts.

Observers

The same observers used in experiment 1 were employed for this experiment.

Results and Discussion

Threshold estimates for simple test grating contrast, $C = \alpha M_a$, and generalized contrast, C' of expression (3), were determined for both observers for the various test-noise combinations. The logarithms of the reciprocals of these values, i.e. the log-sensitivities, are recorded in Tables 5 through 8. Tables 5 and 6 contain the - log C estimates and Tables 7 and 9 contain the - log C' estimates. The following notation is used: n/m denotes n cycles/deg test with m cycles/deg mask; n/B is the situation studied in the earlier

experiment, n cycles/deg test with blank mask. The n/B entries are taken from the regression results of experiment 1. For the sake of brevity, the following discussion will refer only to the simple grating sensitivity measures given in the first pair of tables, however all conclusions hold for the generalized sensitivity measures as well.

Table 5

Averaged - log C for S.B.

test/mask	Exposure		Δ
	30 msec	300 msec	
10/5	.94	1.00	.07
10/2.5	1.02	1.20	.18
5/10	1.22	1.29	.07
2.5/10	1.39	1.57	.18
10/B	1.45	2.02	.57
5/B	1.80	2.22	.42
2.5/B	2.02	2.27	.25

Table 6

Averaged - log C for B.B.

test/mask	Exposure		Δ
	30 msec	300 msec	
10/5	.97	.95	-.02
10/2.5	.94	1.25	.31
5/10	1.50	1.55	.05
2.5/10	1.78	1.99	.21
10/B	1.44	2.19	.75
5/B	1.77	2.10	.33
2.5/B	2.08	2.48	.40

Table 7

Averaged - log C'for S.B.

test/mask	Exposure 30 msec	300 msec	Δ
10/5	.98	1.05	.07
10/2.5	1.08	1.26	.18
5/10	1.27	1.34	.07
2.5/10	1.44	1.62	.18
10/B	1.45	2.02	.57
5/B	1.80	2.22	.42
2.5/B	2.02	2.27	.25

Table 8

Averaged - log C'for B.B.

test/mask	Exposure 30 msec	300 msec	Δ
10/5	1.01	.99	-.02
10/2.5	.99	1.31	.32
5/10	1.55	1.60	.05
2.5/10	1.83	2.04	.21
10/B	1.44	2.19	.75
5/B	1.77	2.10	.33
2.5/B	2.08	2.48	.40

The columns labeled Δ in these tables indicate the difference in log-sensitivity between the 300 msec and 30 msec exposures of a particular row entry. Statistical evaluation of the various Δ values made use of the segmented Wetherill estimates. The segmentation procedure was described in experiment 1. t-tests were used for the data of this experiment (our unequal sample sizes made these

tests the most convenient for comparing values - see Snedecor and Cochran, 1967, pg. 115).

In referring to particular values of Δ the notation $\Delta(n/m)$ is useful. $\Delta(n/m)$ indicates the value of Δ for n cycle/deg test and m cycle/deg mask, and $\Delta(n/B)$ denotes the value of Δ for the corresponding blank mask condition.

Let us consider, first, what is revealed for the detection of a 10 cycle/deg grating against various backgrounds. A graphical representation of these results is displayed in Figures 11-14. Figures 11 and 12 depict the log-sensitivities recorded in Tables 5 and 6 for the 10 cycle/degree test grating as a function of noise background frequency on a linear scale.

For both observers, t-test analyses yield: $0 \approx \Delta(10/5)$, $\Delta(10/5) < \Delta(10/2.5)$ ($p < .001$ for B.B., $p < .05$ for S.B.), $\Delta(10/2.5) < \Delta(10/B)$ ($p < .001$ for B.B., $p < .01$ for S.B.). That is, $0 \approx \Delta(10/5) < \Delta(10/2.5) < \Delta(10/B)$. The curves in Figures 11 and 12 are not parallel. Furthermore, since for the 30 msec exposure 10/2.5 sensitivity is strictly less than 10/B sensitivity for both observers ($p < .001$ for both observers), we have the additional result that none of the curves in these two figures is parallel to the abscissa.

Two important conclusions follow from the nature of the curves in Figures 11 and 12. The observation that these curves are not parallel to the abscissa indicates lack of independence between a 10 cycle/degree grating and suprathreshold background gratings separated by more than an octave in frequency. In addition, the observation that the 300 msec and 30 msec curves are not parallel

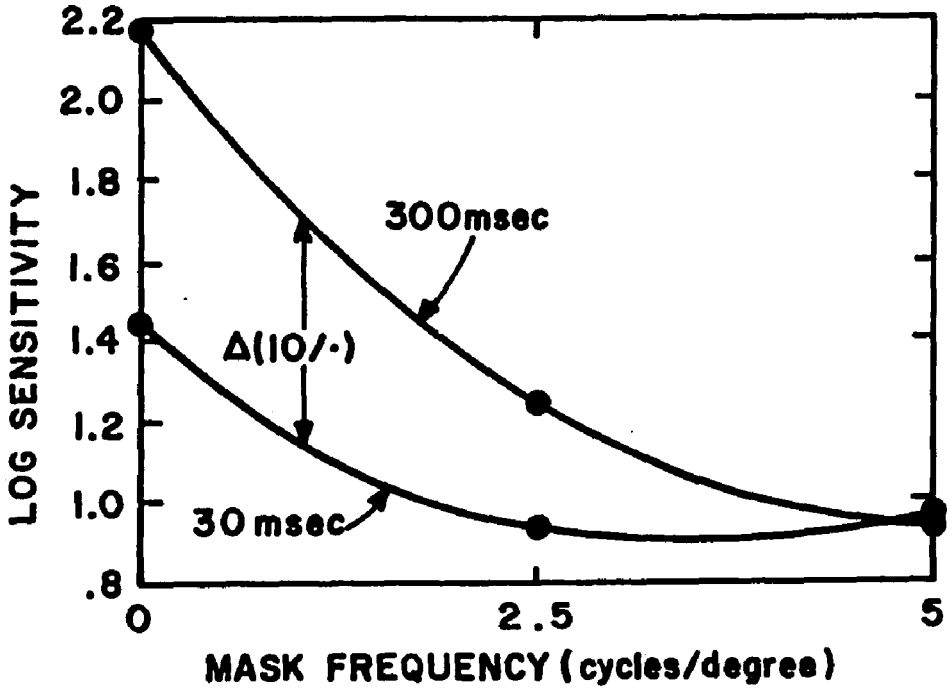


FIGURE 11. Log contrast-sensitivity for a 10 cycle/deg grating as a function of background, mask frequency, for 30 msec and 300 msec exposure duration. Results for observer B.B. The separation between these curves at a particular mask frequency is Δ(10/.). See text for details.

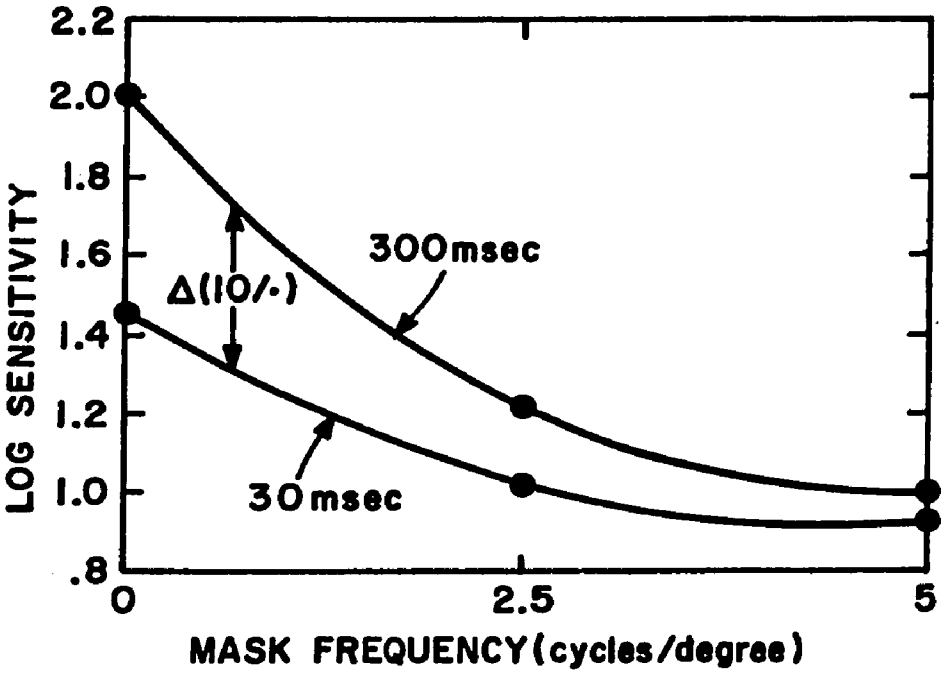


FIGURE 12. Log contrast-sensitivity for a 10 cycle/deg grating as a function of background, mask frequency, for 30 msec and 300 msec exposure durations. Results for observer S.B. The separation between these curves at a particular mask frequency is Δ(10/.). See text for details.

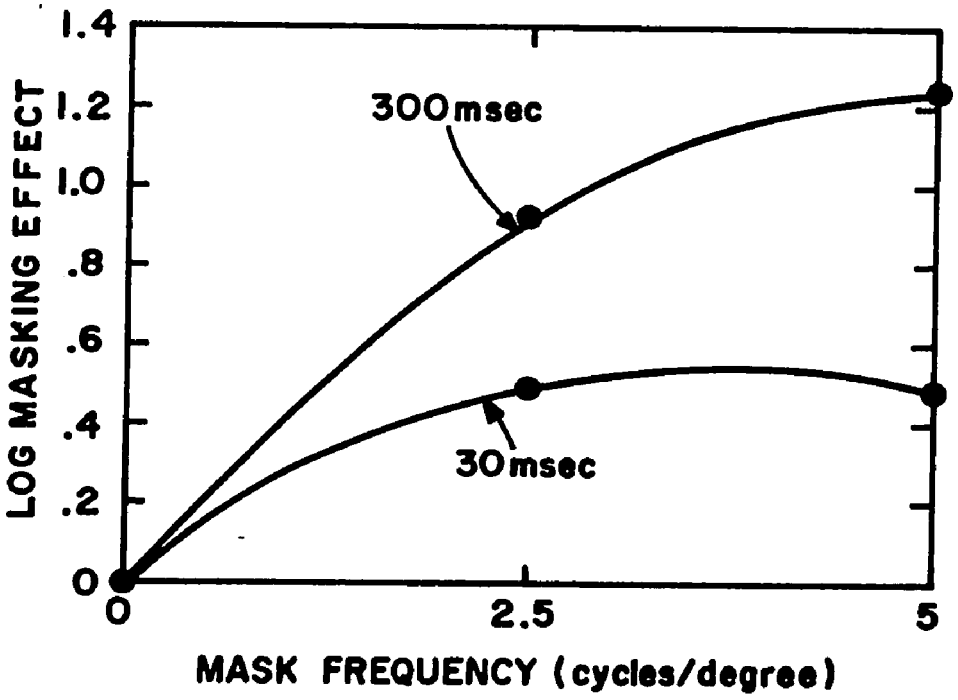


FIGURE 13. Log masking effect of a background frequency on 10 cycle/deg sensitivity for two exposure durations. Results for observer B.B. The masking effect of a background frequency is defined as the induced sensitivity reduction when compared to the blank background situation. This curve is derived directly from Fig. 11.

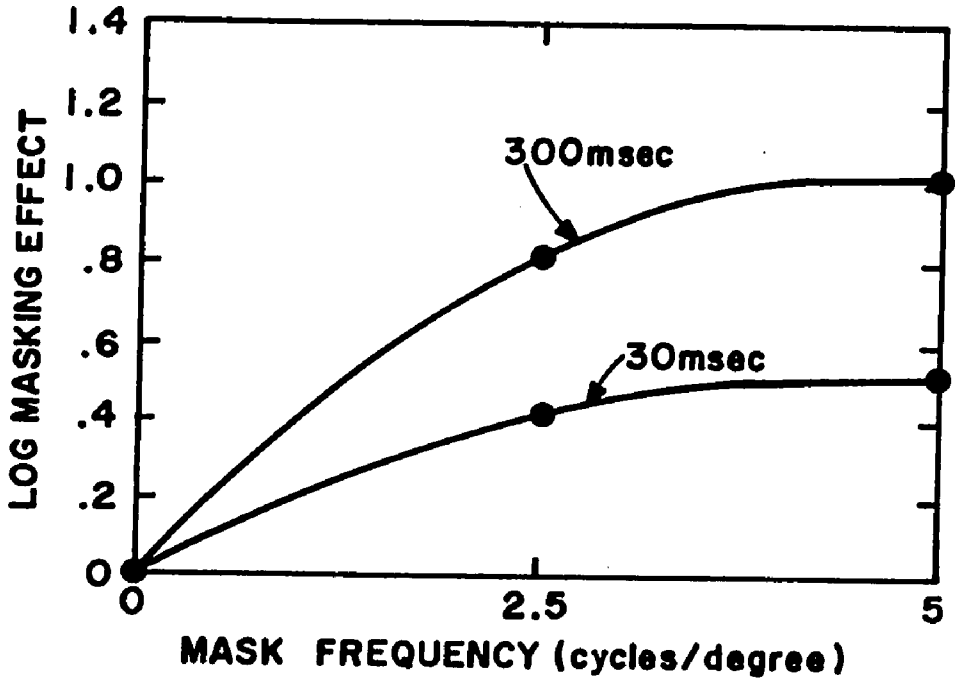


FIGURE 14. Log masking effect of a background frequency on 10 cycle/deg sensitivity for two exposure durations. Results for observer S.B. The masking effect of a background frequency is defined as the induced sensitivity reduction when compared to the blank background situation. This curve is derived directly from Fig. 12.

for either observer indicates that the time course of interaction between a 10 cycle/degree grating and its background is different for different frequency background gratings. Violation of the independent channels hypothesis is unequivocal; time dependent cross-talk for well separated spatial frequencies is very evident. Linearity is not a tenable proposition for suprathreshold stimuli.

Support for this assertion is also found in the adaptation studies of Tolhurst (1972), and of Nachmias, et.al. (1973) referred to earlier. Our finding of non-linearity in spatial frequency processing is not new. However, by our use of what can be called a simultaneous masking technique we have gained some additional information about the nature of spatial frequency interaction that is not easily obtained using adaptation procedures. That is, some temporal aspects of these spatial interactions are also revealed. We will elaborate on this shortly.

The effect of changing background frequency for 10 cycle/deg detection is further emphasized by the curves in Figures 13 and 14. These curves depict a simple transformation on the curves in the preceeding pair of figures. By inverting the curves in Figures 11 and 12, and assigning a value of zero to the effect of masking by a blank field we have a measure of the degree of masking induced by a particular background. Figures 13 and 14 display the result of this transformation.

The pair of curves for each observer should coincide if the time course of interaction is independent of background, and the

curves should coincide with the abscissa if interaction is absent altogether.

The most important point that emerges from consideration of the results displayed in Figures 11 - 14 is that the detectability of a grating against the background of a different frequency grating is sufficiently non-linear to warrant caution against drawing any simple general conclusions. With this caveat in mind, aspects of the data that bear on the apparent conflict between the results of experiments 1 and 2 will be pointed out.

The differences between 10 cycle/degree grating detectability at 300 msec and 30 msec exposures for various backgrounds can be used to expose different processing rates for the different backgrounds. We have gathered two types of information with our simultaneous masking technique; the degree to which a suprathreshold noise grating masks a test grating, and the time course of this masking. The degree of masking and the time course of masking are distinct issues. There are reasons from prior considerations to suspect that a suprathreshold 5 cycle/deg grating would be more effective than a 2.5 cycle/deg grating for decreasing sensitivity to a 10 cycle/deg grating. One would expect that the likelihood of cross-talk increases as separation along the stimulating dimension decreases. An entirely independent question is whether the 5 cycle/deg masking is more immediate than that of the 2.5 cycle/deg masking. This question addresses itself to the relative processing rates of suprathreshold stimuli.

It has been pointed out that for both observers: $0 \approx \Delta$ (10/5) $< \Delta$ (10/2.5). The observation that $0 \approx \Delta$ (10/5) leads one to conclude that the masking effect of a suprathreshold 5 cycle/deg grating is complete in less than 30 msec. This implies that visual processing of such a 5 cycle/degree grating is completed in less than 30 msec. (Recall that 30 msec is in the range of exposure durations used in the picture recognition studies of experiment 2). Next we note that Δ (10/2.5) > 0, and furthermore Δ (10/2.5) $> \Delta$ (10/5). The masking effect of a suprathreshold 2.5 cycle/deg grating is not complete in 30 msec, its maximal effect takes longer to achieve than for 5 cycle/deg masking. This implies that complete visual processing of such a 2.5 cycle/deg grating takes longer than 30 msec. Thus, 5 cycle/deg processing is more rapid than 2.5 cycle/deg processing in suprathreshold situations.

It is important to note that this type of inference about relative processing rates from simultaneous masking is valid only if the frequencies considered are used to mask the same test grating. For example, it is not meaningful to compare Δ (10/5) and Δ (2.5/10) to determine the relative rates of 5 cycle/deg and 10 cycle/deg processing. This is because the test signals, 10 cycles/deg and 2.5 cycles/deg, are different and are, themselves, processed at different rates.

The relative processing rates we arrived at for threshold stimuli in experiment 1 indicated that 2.5 cycle/deg processing is more rapid than 5 cycle/deg processing. We find, however, that this ordering of processing rates is not maintained for suprathreshold

gratings. In fact, the ordering becomes reversed. Of course, it is pure speculation to argue from this result that in suprathreshold situations processing becomes more rapid as spatial frequency increases. This result, however, is not inconsistent with such a premise, and such a premise is not inconsistent with the picture plus noise masking results of experiment 2.

It is revealing to reconsider the relation $0 \approx \Delta(10/5) < \Delta(10/2.5) < \Delta(10/B)$ from the point of view of processing time of the 10 cycle/deg target. One can say that the processing rate of the 10 cycle/deg target depends critically on its background. Processing of target information is complete in less than 30 msec if viewed against a 5 cycle/deg background, and becomes progressively less complete as the frequency of the background becomes separated from that of the test grating (recall that a blank background is equivalent to a zero frequency background). This observation reemphasizes the highly non-linear nature of suprathreshold processing.

Non-linear processing is a complicated business. We see that the simplest generalization of a spatial frequency visibility experiment leaves us squarely in a situation of non-linear processing. For this reason alone we can regard the apparent inconsistency between the results of experiments 1 and 2 as not being particularly troublesome. As we have noted, the connection between these two experiments rests heavily on the assumption of linearity. From the evidence in this latest experiment, it seems that this assumption is inaccurate. The blank background and non-blank background grating

detection results are inconsistent in a number of respects that indicate conflict between our first two experiments is to be expected.

CHAPTER 5

SUMMARY AND COMMENT

The question that was posed at the outset of this study was whether it is useful to formulate theories of visual recognition in terms of statements they make about the processing of spatial frequencies. Virtually all models of recognition propose an early feature extraction stage. The visual features being extracted in such models are typically thought of as local elements, such as line segments, curves, corners, and the like. This is somewhat reasonable, since these models deal with the recognition of alpha-numeric material, which is composed of line segments, curves, corners, and the like. However, when one begins to consider visual stimuli such as pictures, features of this sort become less compelling. A more global type of feature for such situations would be more intuitively satisfying. The psychophysical evidence indicating the existence of narrowly tuned spatial frequency analyzers in the visual system suggests spatial frequency content in a narrow frequency band as a potential alternative to the local element feature concept. An experimental investigation of the applicability of the concept of perceptual feature to spatial frequency, therefore, was deemed worthwhile. The common theme of the experiments we performed was an examination of the relative benefit that increased viewing time of a stimulus produces for the processing of spatial frequency information.

Our first experiment was essentially a replication of a study by Nachmias (1967). Using a low frequency and a high frequency square wave grating, he found that the low frequency grating reached a sensitivity plateau at an exposure duration for which the high frequency sensitivity was still increasing. This leads one to conclude that visibility processing of the low frequency square wave grating is more rapid than that of the high frequency grating. A number of potential complications are introduced by the use of square wave gratings, so we re-examined this situation using sinusoidal luminance modulation.

Five sine wave gratings, at one octave frequency spacing, were selected for study. Our results were generally consistent with those of Nachmias; that is, low frequency visibility processing did tend to be more rapid than high frequency processing. Using our results we isolated a low frequency neighborhood and a high frequency neighborhood, that were clearly processed at different rates, for our picture recognition experiment. These neighborhoods covered 2.5 cycles/deg and 10 cycles/deg, respectively.

For the picture recognition experiment we prepared two-dimensional, band-limited, equal bandwidth, noise patterns, one with the noise band centered near 2.5 cycles/deg and another with the noise band centered near 10 cycles/deg. Observers viewed briefly exposed pictures with the noise superimposed for the first half of the exposure, the second half of the exposure, or not introduced at all. The noise insertion was accomplished without changing either the picture contrast or the mean luminance from

that fraction of the exposure when the noise was absent. We found that the low frequency noise had no observed effect on picture recognition for brief exposures, whereas the high frequency noise produced a masking effect when presented in the first half of the picture viewing interval. High frequency noise did not significantly degrade recognition performance for insertion in the latter part of the exposure, hence this type of noise pattern yields some of the characteristics of forward masking.

The low frequency noise covers spatial frequencies that have lower thresholds and are more rapidly processed than spatial frequencies covered by the high frequency noise, according to the results of our first experiment. The effects of these two types of noise patterns were, therfore, unexpected.

A possible explanation for this apparent conflict is that masking effects of the noise patterns are indeed consistent with the characteristics of single frequency processing observed in our first experiment, however the recognition process for briefly exposed pictures makes use of only high spatial frequency information. Selective attention to the higher spatial frequencies could be a perceptual strategy employed in this type of task. Evidence for spatial frequency selective attention has, in fact, been reported by Carpenter and Ganz (1972).

However, a simpler explanation for the apparent conflict between the results of our first two experiments is that single frequency visibility processing does not accurately simulate the type of suprathreshold, multiple frequency processing involved in picture recognition. An assumption of linearity is strongly relied

on in considering these two situations comparable. An experiment was devised to test this assumption.

In the first experiment we examined single spatial frequency processing against a blank background. In our third experiment we examined frequency processing against a non-blank background. In particular, we studied exposures of the luminance distribution:

$$L = k[1 + \alpha M_a \sin 2\pi f_a x + (1 - \alpha) M_b \sin 2\pi f_b x]$$

$$\alpha, M_a, M_b \in [0,1].$$

Exposures were preceeded and followed by a blank adaptation field.

The threshold value of α was determined for detection of the presence of the test grating of frequency f_a in the presence of the noise grating of frequency f_b. From threshold α we determined the test grating contrast threshold, as well as a more general signal-to-noise measure. The validity of a linearity assumption would require indifference of these threshold measures to changes in background frequency, for backgrounds at least an octave separated from the test frequency.

Besides being a test of linearity, this approach can also be said to approximate a figure-ground separation task that is involved in a picture recognition situation. In this experiment the figure is the test grating and the ground is the noise grating. In picture perception some informative complex of frequencies could be identified as figure, and unimportant spatial frequencies could be identified

as ground. The validity of this analogue is not critical to our experiment.

The non-blank background experiment revealed a number of facts which indicated absence of linearity. Hence close agreement between the picture with inserted noise mask results and characteristics of visibility processing of a single spatial frequency against a blank background would be the surprising outcome, not the other way around. The situations are simply not comparable.

We were also able to use the results of this last experiment, which may be considered a simultaneous masking situation, to reveal some aspects of suprathreshold spatial frequency processing. In particular, we compared the masking effects of the suprathreshold background gratings and concluded that 5 cycle/deg processing is more rapid than 2.5 cycle/deg processing; 5 cycle/deg processing is complete for the exposure durations used in the picture recognition experiment whereas 2.5 cycle/deg processing is not.

The primary outcome for this last experiment is not a potential explanation of the picture plus noise findings, but the resolution of the apparent conflict between processes observed for single frequency, blank background processing and processes observed for noise masked picture recognition. In light of the evidence we have found for spatial frequency interaction, it would be somewhat ambitious to attempt to accurately simulate the multiple frequency processing involved in picture recognition using sine wave gratings. It is not likely we have done this in our two spatial frequencies situation.

It seems fair to conclude that spatial frequency analysis is not a promising vehicle for modeling the general process of visual recognition. There is, however, one aspect of this process that might be profitably considered in a spatial frequency framework

It was pointed out that the non-blank background grating detection task we have examined could be viewed as a task involving figure-ground separation. Figures 11-14 above illustrate the fact that 10 cycle/deg grating detection becomes more difficult as the background frequency gets closer to 10 cycle/deg. It is true that the closer in frequency two gratings are to each other, the more they resemble each other. For the more typical figure-ground situation one would expect that figure-ground separation becomes more difficult as the figure and ground become more similar. So it seems to be with sine wave gratings.

Neisser's sequential stages of preattentive processing and focal attention were restated as spatial frequency operations, in the introduction of this study, as an example of how one might go about reformulating a recognition theory in terms of spatial frequency analysis. From the figure-ground separation as spatial frequency analysis point of view it might have been more accurate to isolate only the preattentive stage for consideration, since Neisser asserts that the passive perceptual segregation of figure from ground is the function of the preattentive processes. This is the only strictly passive stage in his constructive framework and, therefore, might be the only aspect of this theory that can be approached from

a study of passive sine wave grating processing. And this approach would be meaningful only if background gratings are included. It might, indeed, turn out that a clearer understanding of the perceptual feat of figure-ground separation could emerge from such considerations.

APPENDIX 1

SUFFICIENT VISUAL ANGLE TO AVOID FREQUENCY DISTORTION IN FINITE EXTENT GRATINGS

The intended luminance distribution for each of the gratings in our experiments is:

$$h(x) = k[1 + m \sin 2\pi f_o x]$$

where:

f_o = grating frequency

x = degrees visual angle

m = grating modulation

k = average grating luminance

However, the gratings subtend 3.7^o, so the true luminance distribution is:

$$h_t(x) = k[1 + m \sin 2\pi f_o x] \cdot [u(x + x_o) - u(x - x_o)]$$

where $u(\cdot)$ is the unit step function:

$$u(x-a) = \begin{cases} 0, & x < a \\ 1, & x > a \end{cases}$$

and

$$x_o = 1/2\,(3.7^o) = 1.85^o.$$

The amplitude of the Fourier transform of $h_t(x)$ is depicted in Figure A1. In this figure $f_1 = \frac{1}{x_o} = \frac{1}{1.85^o} = .54 \text{ deg}^{-1}$. If the signal band, centered at f_o, is kept from overlapping with the aperture band, centered at zero frequency, distortion of the intended signal is avoided. This is effectively accomplished if we choose f_1 sufficiently small (or x_o sufficiently large), so

that for all f_o, $f_o - f_1 > f_1$, or $f_o > 2f_1$. If this is satisfied for the smallest chosen f_o, it is satisfied for all f_o. We have $f_1 = .54 \text{ deg}^{-1}$, and the smallest f_o is 1.25 deg^{-1}. We therefore have $f_o = 1.25 > 2f_1 = 1.08$. Signal distortion due to the finite grating extent should be absent for the range of frequencies we have selected for study.

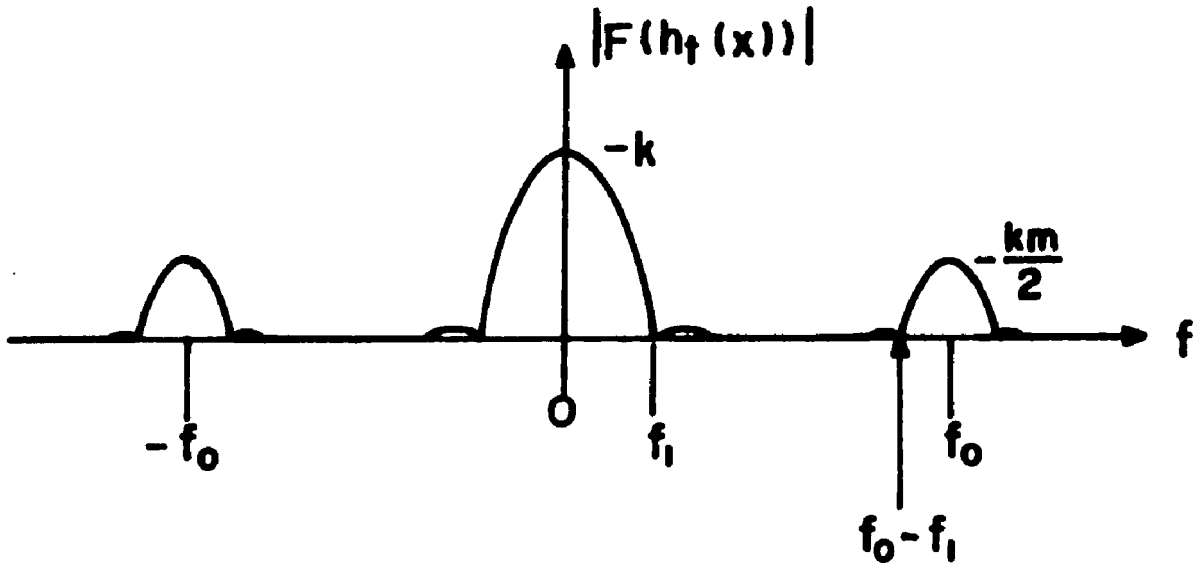


FIGURE A1. Amplitude of the frequency spectrum of a sine wave grating of frequency f_o subtending a finite visual angle. See text for details.

APPENDIX 2

PICTURE-STIMULI FOR THE RECOGNITION EXPERIMENT

The picture-stimulus slides used in our second experiment were kindly provided by Professor Robert G. Pachella. These slides were prepared from photographs of common objects found in magazines. Each object-picture was cut out of a magazine, placed on a blank background, photographed, and converted to slide format. The resulting slide had the object-picture in its center and was transparent elsewhere. The eight pictures used in our experiment were of the following objects:

fire extinguisher	coffee pot
roller skate	combination lock
bicycle	door knob
flag	wheel chair .

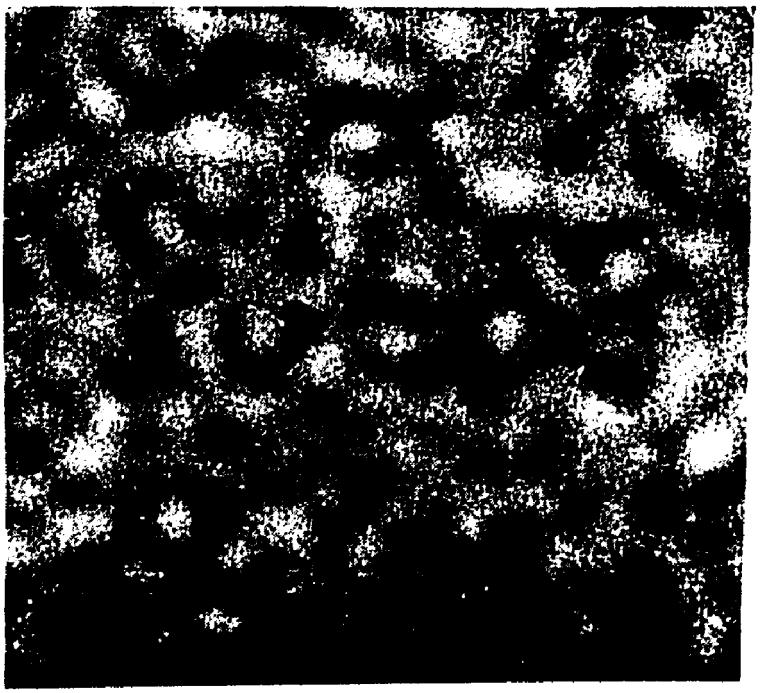
APPENDIX 3

PICTURE DISTORTION BY FILTERED NOISE SUPERPOSITION: A SOURCE OF CONTRAST HOLES AND PLUGS

Reproductions of the high and low frequency filtered noise patterns used in Experiment 2 are shown in Figure A2. The reproduction process has introduced a fair amount of distortion, but the renditions here are adequate for the current discussion.

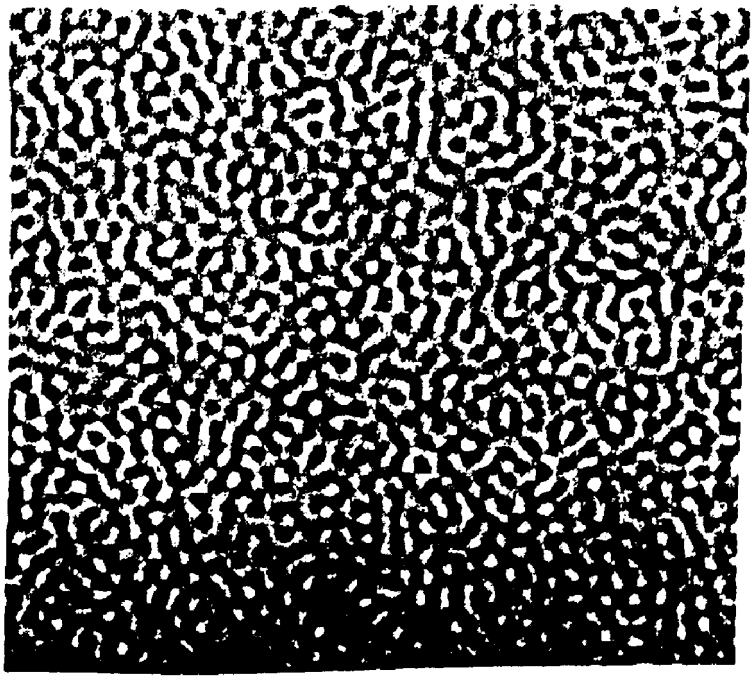
The light and dark regions in these patterns correspond to relatively opaque and relatively transparent regions of the noise slides. Light passing through a noise slide is appropriately modulated by these local variations in transmissivity, and then optically superimposed on the image of a picture slide. The result of this optical superposition is the intended addition of noise to picture information. The effect of this noise superposition is simply to introduce local picture contrast perturbations, thereby distorting the picture.

For example, in a local region corresponding to relatively high noise slide transmissivity the local picture contrast will be lower than would be the case if the noise slide transmissivity were at its mean value, and in a local region corresponding to relatively low noise slide transmissivity the local picture contrast will be higher than would be the case if the noise slide were at its mean value. To make this point more apparent, it may be helpful to recall that contrast is defined as $(L_{max} - L_{min})/(L_{max} + L_{min})$, so luminance increase by optical superposition yields a decrease in contrast. When the blank-fixation slide is used, as in the no-noise trials,



LOW FREQUENCY NOISE PATTERN

FIGURE A2



HIGH FREQUENCY NOISE PATTERN

picture contrast is, of course, not subject to local perturbations of this sort, since all local regions of the fixation slide are at the mean transmissivity value.

The fact that the picture contrast perturbations caused by the noise superposition can be at higher contrast than these regions would be at in the no-noise condition has interesting implications.

It is descriptive to refer to the perturbed regions with enhanced contrast as contrast holes. That is, if one considers the diffuse fixation slide luminance, in the no-noise condition, as being a dilution field with appropriate contrast reducing properties, then its replacement by a noise slide can be said to introduce a dilution field with holes that allow picture contrast to pass with relatively little reduction. In addition, a noise dilution field also produces contrast reducing regions which are the complements of the contrast holes. For want of a better term, such regions will be referred to as contrast plugs.

For practical reasons, only one noise pattern was prepared for each range of spatial frequencies to be masked. Also, for practical reasons, one of the noise patterns was used exclusively in an experimental session. The consistent use of a particular noise pattern for a session places the contrast holes so that they are effectively congruent with the same picture regions for every noise trial. It is, therefore, possible that a subject could become able to take advantage of information enhancement at the never changing spatial locations of the contrast holes for the picture identification task.

To an extent, this likelihood was minimized by some of our procedures. The practice and experimental sessions were kept short, with a reasonably small number of trials, so that the subjects did not become overly practiced. Response feedback was not given after each experimental trial, so awareness of contrast hole locations would be slow to emerge. Subjects were not informed about noise trials prior to a session. When interviewed later, subjects reported that they were not aware of the occurrence of noise trials.

One would expect that the possible identification enhancement afforded by the existence of contrast holes would be related to the size of these holes in a straightforward way; the larger the hole, the greater the benefit to identification. Inspection of the filtered noise patterns shown in Figure A2 reveals that the general size of the contrast holes is different for the two patterns. The contrast holes associated with the low frequency pattern are considerably larger than the holes associated with the high frequency pattern. Therefore, one might expect that if the anomalous effect of noise trials yielding higher recognition accuracy than no-noise trials were to occur, we would find that the low frequency noise trials are capable of producing greater recognition enhancement than are the high frequency noise trials.

Do the experimental results, summarized in Table 3 above show any evidence of such an outcome? We do, in fact, find that low frequency noise for late noise trials yields 3.3% higher accuracy than for the corresponding no-noise trials (early noise trials produced a similar difference), whereas high frequency noise

for late noise trials yields 1.4% greater accuracy than for the corresponding no-noise trials. Statistical analysis attributed these differences to random variability, they were not significant. However, one can not prove the null hypothesis. The outcome, for the presumed non-masked late noise trials, of recognition enhancement for both types of noise patterns with greater enhancement for the low frequency noise pattern would be caused randomly with probability 1/6 = .167, suggesting that contrast hole effects are perhaps hidden in the noise of the data. One could argue that the lack of significance of these hole effects in our experiment indicates we have adequately guarded against them, and there is no reason to be concerned about our consistent use of the same noise patterns.

It is important to point out that even if the contrast hole effects were significant we would still be able to make use of the early and late filtered noise masking results. The questions as to whether the contrast plugs laterally mask the contrast holes, and the size effects of these plugs and holes address themselves to the global aspects of filtered noise masking which we have studied, and are not really confounded by keeping the locations of the plugs and holes constant.

One way that contrast hole effects might become more evident with consistently used noise patterns would be to use subjects who were highly practiced. We have gathered some data that points in this direction.

After a break of a few days, three of the subjects that had participated in the experiment, and had been informed about the use

of noise trials, returned for additional sessions. These subjects were selected from those who had been tested with the middle frequency noise patterns, had no-noise trials included in all sessions, and had two practice sessions. Thus, they had been through five separate sessions before they were recruited for the extra sessions. Each of these extra sessions contained twice the number of trials as had been used in the experimental sessions. These subjects were, therefore, highly practiced, and become even more practiced in the post-experiment sessions. (A fourth subject was involved in one of these extra sessions, but did not participate in the other. Her results are, therefore, not included here, but her data were consistent with that of the others). Results for these extra sessions are summarized in Table A1.

Table A1

	Early Noise	Late Noise	No Noise	$\bar{\tau}$ (Average Exposure msec)
Low Frequency Noise	86.3%	82.9%	70.9%	29
High Frequency Noise	70.4%	76.3%	65.9%	26

Statistical analysis for results with three subjects would be pointless. Nevertheless, it is interesting to observe that we obtain what seems to be an exaggerated version of the earlier, possibly hole induced, recognition enhancement on noise trials. For late noise, low frequency noise trials yield 12.0% higher recognition accuracy than no-noise trials, while the corresponding difference for high frequency noise is 10.4%. Similarly, for early noise, low

frequency noise trials yield 15.4% higher recognition accuracy than no-noise trials, whereas the corresponding difference for high frequency noise trials is 4.5%. Once again, the possibly noise induced enhancement is coupled with superior low frequency noise enhancement. What looks like noise unmasking in this data occurs in a manner consistent with what at first glance look like random effects in the results reported in Table 3.

The purpose of this Appendix is to explain how superimposed filtered noise distorts a picture stimulus, and how this particular type of distortion creates an opportunity for recognition enhancement over the corresponding noise-free situation. This potential relative recognition enhancement would be an effect of what we have called contrast holes. Although such an effect was not found to be significant in our noise masked picture recognition results, there was a not too subtle hint in these data that hole effects could emerge with further experimental investigation. Preliminary results with three highly practiced subjects suggests that hole effects would probably be significant with observers of this sort.

The intent here, however, is not to investigate contrast hole effects in a situation where they probably would be found. We only want to point out that in a situation where they could be found, as in the second experiment of this study, they are minimal.

APPENDIX 4

MIDDLE FREQUENCY SPATIAL NOISE MASKING IN PICTURE RECOGNITION

In experiment 2, all 18 subjects were tested in both low frequency, 0-3.7 cycle/deg, and high frequency, 7.4-11.1 cycle/deg, noise masking sessions, as described above. The first twelve subjects were also tested in an additional session, where a middle frequency, 3.7-7.4 cycle/deg, spatial noise pattern was employed for picture masking. For these twelve subjects, the counterbalancing of the type of noise used in a particular session was for the three types of noise.

The results for the 3.7-7.4 cycle/deg noise masking sessions are summarized in Table A2.

Table A2

Middle Frequency Noise Session Results

Early Noise	Late Noise	No Noise	($\bar{\tau}$ Average Exposure msec)
72.9%	77.5%	75.8%	32

The table entries indicate average response accuracy for late noise, early noise, and no-noise trials. Since the first six subjects were not given no-noise trials, a least-squares estimate of the missing data was obtained using the procedure of Cochran and Cox (1957). Analysis of variance on the arcsine transformed results revealed that treatment effects were not significant.

The data for middle frequency noise masking does suggest a mild version of the effects obtained with high frequency noise

masking (cf. Table 3). However, given the fact that the point to be made by this experiment would be adequately demonstrated by a comparison of the high frequency and low frequency noise masking results, it was decided not to invest the time necessary to run more sessions with the hope that the suggestion of a weak effect with middle frequency noise masking could be rendered significant.

REFERENCES

Attneave, F. and Arnoult, M.D. (1956). The quantitative study of shape and pattern perception. Psychological Bulletin, 53 (6), 452-471.

Baron, J. (1971). The threshold for successiveness. Perception & Psychophysics, 10(4A), 201-207.

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

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

Bracewell, R. (1965). The Fourier transform and its applications. New York: McGraw-Hill.

Carpenter, P.A. and Ganz, L. (1972). An attentional mechanism in analysis of spatial frequency. Perception & Psychophysics, 12, 57-60.

Cicerone, C.M. (1974). Theory of luminance detection and measurement of the impulse responses and frequency responses of the color mechanisms of the eye. Michigan Mathematical Psychology Program, 74-6.

Cochran, W.G. and Cox, G.M. (1957). Experimental Designs, New York: Wiley.

Festinger, L., Allyn, M.R. and White, C.W. (1971). The perception of color with achromatic stimulation. Vision Research, 11, 591-612.

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

Green, D.G. (1968). The contrast sensitivity of the colour mechanisms of the human eye. Journal of Physiology, 196, 415-429.

Green, D.G. (1969). Sinusoidal flicker characteristics of the color-sensitive mechanisms of the eye. Vision Research, 9, 591-601.

Harmon, L.D. and Julesz, B. (1973). Masking in visual recognition: effects of two-dimensional filtered noise. Science, 180, 1194-1197.

Hebb, D.O. (1949). The organization of behavior. New York: Wiley.

Hoekstra, J., van der Groot, D.P.A., van den Brink, G., and Bilsen, F.A. (1974). The influence of the number of cycles upon the visual contrast threshold for spatial sine wave patterns. Vision Research, 14, 365-368.

Hubel, D.H. and Wiesel, T.N. (1959). Receptive fields of single neurones in the cat's striate cortex. Journal of Physiology, 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.

Kelley, D.H. (1974). Spatio-Temporal characteristics of color-vision mechanisms. Journal of the Optical Society of America, 64, 983-990.

Nachmias, J. (1967). Effect of exposure duration on visual contrast sensitivity with square-wave gratings. Journal of the Optical Society of America, 57, 421-427.

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.

Narasimhan, R. (1970). Picture languages, in Picture Language Machines, S. Kaneff, ed., New York: Academic Press.

Neisser, U. (1967). Cognitive Psychology. New York: Appleton-Century-Crofts.

Pollehn, H., and Roehrig, H. (1970). Effect of noise on the modulation transfer function of the visual channel. Journal of the Optical Society of America, 60, 842-848.

Ratliff, F. (1966). Mach bands: quantitative studies on neural networks in the retina. San Francisco: Holden Day.

Rosenfeld, A. (1969). Picture processing by computer, New York: Academic Press.

Rumelhart, D.E. (1970). A multicomponent theory of the perception of briefly exposed visual displays. Journal of Mathematical Psychology, 7, 191-218.

Rumelhart, D.E. and Siple, P. (1974). Process of recognizing tachistoscopically presented words. Psychological Review, 81, 99-118.

Sachs, M., Nachmias, J., and Robson, J.G. (1971). Spatial frequency channels in human vision. Journal of the Optical Society of America, 61, 1176-1186.

Schade, O. (1956). Optical and photoelectric analog of the eye. Journal of the Optical Society of America, 46, 721-739.

Schober, H.A.W. and Hiltz, R. (1965). Contrast sensitivity of the human eye for square-wave gratings. Journal of the Optical Society of America, 55, 1086-1090.

Selfridge, O. (1959). Pandemonium: a paradigm for learning, in The mechanisation of thought processes, London: H.M. Stationery Office.

Snedecor, G.W. and Cochran, W.G. (1967). Statistical methods. Ames, Iowa: Iowa State University Press.

Tolhurst, D.J. (1972). Adaptation to square-wave gratings: inhibition between spatial frequency channels in the human visual systems. Journal of Physiology, 226, 231-248.

Wetherill, G.B. and Levitt, H. (1965). Sequential estimation of points on a psychometric function. British Journal of Statistic Psych., 18.

Winer, B.J. (1971). Statistical principles in experimental design. New York: McGraw-Hill.