DEGRADATION OF PICTURE QUALITY BY SPECKLE IN COHERENT MAPPING SYSTEMS

Thesis by

Vijaya Narayan Korwar

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

1980

(Submitted April 11, 1980)

to my parents

ACKNOWLEDGEMENT

I would like to thank the various persons and organizations whose help I have received during the course of this work. In particular, I would like to express my deep gratitude to my thesis advisor, Professor John Pierce, for his understanding, encouragement and guidance throughout my stay at Caltech. I would also like to thank Dr. Edward Posner very much for all his help over the past few years and for his helpful discussions regarding my thesis.

I am very grateful for the help of various people at the Jet Propulsion Laboratory; especially Dr. Chialin Wu for his many helpful discussions, and Ray Piereson and Chialin Wu for allowing me to use the computing facilities in their group, with which all the simulated pictures required in this work were produced.

The generous financial support given by Caltech, the Jet Propulsion Laboratory and Bell Laboratories is appreciated.

And last but not least, I would like to thank Tim, Ashok and Mado for all their help in the production of the thesis.

ABSTRACT

When a coherent imaging system is used to map a rough surface, for instance when a synthetic aperture radar (SAR) system is used to map the earth, the resulting picture of the surface is degraded by random intensity fluctuations called speckle. Speckle obscures the intensity variations caused by the inherent reflectivity differences that identify various features in the picture. The problem of specifying the extent of the degradation caused by speckle in pictures meant to be examined by a human observer was investigated. In particular, the problems considered were those in which an observer has to (a) detect a small feature immersed in a somewhat darker background; (b)detect a grating consisting of alternating bright and somewhat darker lines; (c) distinguish between two or four specified geometrical forms. In each case the picture was corrupted by speckle.

In investigating each of these problems, a plausible theoretical model was developed for the decision process used by the observer in his detection or discrimination task. This model was used to relate the probability of his making a correct decision to the relevant picture parameters such as contrast ratio between the reflectivities of various parts of the picture, number of looks per pixel, picture size, and dimension of the features or lines. These calculations were verified by experiments in which the decisions

made by an observer examining computer simulations of speckle- corrupted pictures were noted.

Results of calculations for the reasonable SAR parameters of 1 dB contrast ratio and 12 looks per pixel showed that, in order to achieve a probability of correct decision of 0.95, (a) a small square in a 100 by 100 pixel background needs to be about 7 pixels on a side; (b) a 100 by 100 pixel grating of line pairs needs to have lines about 2 pixels wide; (c) a simple geometrical form (a specific one) needs to be at least 12 pixels on a side to be distinguished from another (specific) form of the same size, when these two forms are the only possible alternatives. These results illustrate that detectability of line pairs is a poor criterion for characterizing picture quality, while form discrimination imposes the most stringent requirements on the imaging system.

CONTENTS

1. INTRODUCTION

1.1 Nature Of The Speckle Problem

Consider a surface whose rms (root-mean-square) roughness is of the order of magnitude of the wavelength of visible light. Assume that the surface has a uniform optical reflectivity, so that if viewed in ordinary incoherent light, all parts of the surface appear equally bright. However, if the same surface is viewed in coherent light such as laser light, the surface appears to have a random intensity distribution which is called a speckle pattern. Such patterns are easily seen even by a casual observer when highly coherent laser light is shone on a rough surface.

Although speckle was first discovered and so named in connection with optical systems [1], the same effect occurs in other regions of the electromagnetic spectrum, for example, in acoustical imaging systems, in radar astronomical systems, and in radar systems [2]-[8]. In radar, the "clutter" or intensity distribution resulting from microwaves reflected by a diffuse scattering surface is actually a speckle pattern. Speckle patterns formed in partially coherent light are very useful in obtaining information about the scattering surface. However, if a coherent imaging system is used for mapping a surface, then

speckle is an undesirable effect, because it degrades the quality of the photograph obtained. Synthetic Aperture Radar (SAR) is a high resolution, coherent radar technique used for mapping surfaces like the earth from an aircraft or spacecraft [9]-[10]. We will concern ourselves only with the picture-degrading effects of speckle, with particular reference to SAR systems, although the results are applicable to pictures obtained by any other coherent mapping system.

Microwave and optical images are obtained by translating the variations in reflectivity of different portions of the surface being imaged into optical intensity variations on a photographic film. As in any other imaging system, there is, of course, a smallest resolvable dimension, which defines the size of the resolution cell in the surface being mapped. Each such resolution cell corresponds to one pixel in the final picture. In imagery produced by incoherent systems, the resolution cell size has a physical meaning; for instance, a feature of the size of a resolution cell, which has a somewhat different reflectivity from its immediate surroundings, is actually visible in the picture. However, this is not necessarily the case for coherent imagery, because of the random intensity fluctuations caused by speckle; all that one pixel defines is the size of the grains or "speckles" in the picture.

1.2 Brief theory of speckle

In order to explain how speckle arises, we take the specific

case of a SAR system. Consider the composition of the signal return from a single resolution cell, i.e. the signal reflected by the cell. If the surface being imaged has a roughness or height variation comparable to the wavelength of the electromagnetic radiation used for the imaging, the phases of the contribution from various scattering centers differ by amounts of the order of magnitude of 2π and can be considered to be independent random variables uniformly distributed between 0 and 2π. If the imaging process is such that the signal return from a single resolution cell consists of a coherent superposition of the contributions of a large number of these elementary scattering centers, then the intensity of the return from each resolution cell can be considered a random variable with an exponential distribution. Since two adjacent resolution cells represent the contributions of two different sets of elementary scatterers, the intensities on the film corresponding to these adjacent resolution cells are independent random variables (under ideal processing conditions), and the picture consists of a random collection of bright and dark spots which obscure the inherent reflectivity variations of the surface. This high spatial frequency variation of intensity in the picture causes speckle.

The theoretical conditions under which the exponential intensity distribution (or equivalently, the Rayleigh amplitude distribution) for the signal return holds are well-known [7]. Marcum and Swerling [11] state that the

Rayleigh distribution for amplitude holds in most diffuse scattering cases. Experimental evidence for the Rayleigh amplitude distribution has been found by analyzing lunar radar returns [7], radar returns from the ocean [12] and light scattered from diffuse surfaces [13].

1.3 Speckle in SAR systems

The question of visual examination of speckled pictures with the intention of finding features in them has only recently arisen with the development of SAR systems for mapping surfaces like the ocean surface and that of Venus. In some applications, conventional (as opposed to synthetic aperture) radar is used in mapping, for instance, a runway as seen from a (stationary) high tower [14]. But the interest here is not in the diffuse reflection or clutter, but in some smooth target like an aircraft on a runway, which is relatively free from the effects of speckle. The only interest in clutter in these cases is in ways to separate the desired signal from the clutter.

Figure (1.1), courtesy Chialin Wu of JPL, shows a SAR picture of a portion of Imperial Valley, California, generated by digitally processing SAR data taken by the experimental mission SEASAT. The large, grainy patch at the top right corner of the picture is a patch of water, which has a uniform radar reflectivity, but appears grainy because of speckle. Similarly, the interior of each of the several

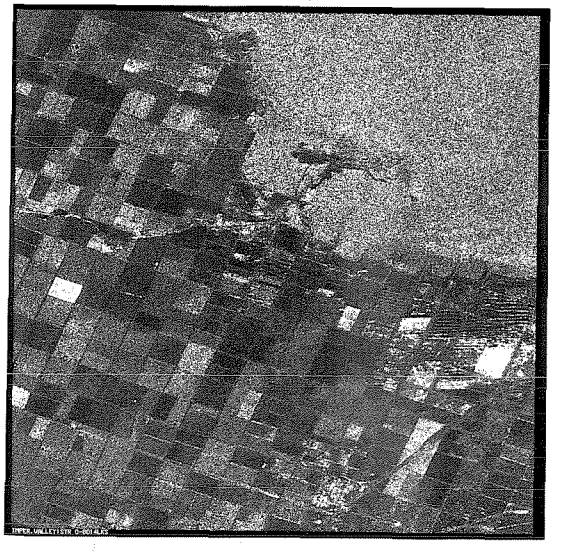

Figure (1.1): Digitally processed SAR image of a part of Imperial Valley, California.

(Courtesy Chialin Wu, JPL)

small rectangular patches in the rest of the picture ought to have appeared uniform, because each of them represents a single vegetable or crop field, a few miles on a side, which again has a uniform radar reflectivity.

1.4 Reduction of speckle

(a) Looks: It is well known that speckle effects are reduced by averaging incoherently over several independent returns from a single resolution cell, i.e. by averaging over several independent estimates of the power return from a single resolution cell. The standard deviation of the average of L independent, identically distributed random variables is $L^{-1/2}$ times the standard deviation of each, while the mean of their average is equal to the mean of each. Therefore, the intensity fluctuation ratio, defined as the ratio of rms variation in intensity to mean value of intensity should decrease by a factor of $1/L^{1/2}$ after this averaging and present a truer reproduction of the original scene [8].

The picture in Figure (1.1) has 4 looks per pixel, but we see that this is not a large enough number of looks to suppress the speckle in the picture.

If the averaging is done coherently, i.e. if independent estimates of amplitude (rather than power) of the return are added coherently, there is no reduction in speckle [15].

These considerations indicate that feature detection and recognition in a given speckled picture ought to improve by taking more "looks" (from now on, unless otherwise specified, "looks" means independent estimates) at each resolution cell, and averaging incoherently over these looks. Examining SAR pictures of a terrain obtained by processing actual SAR data,we can see qualitatively that the pictures appear better with increasing number of looks [8]. Butman and Lipes [16] simulated pictures starting with an unspeckled ERTS picture and adding speckle to it. This picture was then processed to include the effects of varying numbers of looks and the resulting pictures, which are reproduced in Figure (1.2) here, show the same qualitative improvement with increasing number of looks.

(b)Contrast:In the above paragraphs, we considered how a picture with given contrasts between the various features and the background could be improved. But this intensity contrast is, to some extent, a function of the imaging system design and we would like to consider the effect of varying this contrast.

Consider the problem of detecting a small feature of just-resolvable cell size, whose intensity in the absence of speckle would be (1+b) times that of its immediate surroundings. The actual intensity of the SAR return from this bright cell is a random variable with a mean value equal to (1+b) times that of the return from the background.

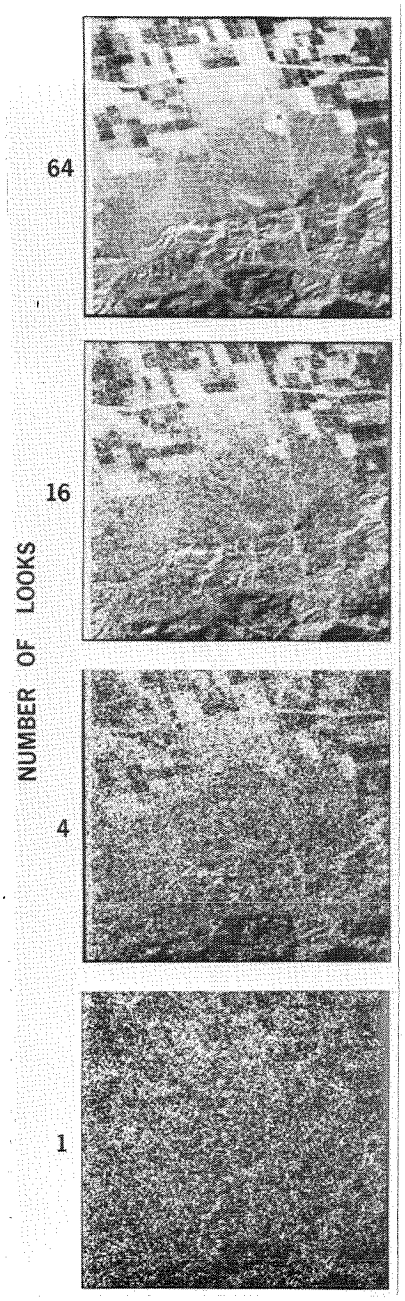

Figure (1.2): Improvement in speckled picture quality with increasing number of looks

(Courtesy Richard Lipes, JPL)

There is a non-zero probability, p, that the return from the bright cell is smaller than that from the background cells. As the number of looks per resolution cell increases, the probability density functions of these random variables become more peaked in the vicinity of their mean values and this probability, p, decreases (which is another way of explaining the effect described in (a) above). But it is also true that, for a given form of probability density functions, this probability decreases as the contrast ratio (1+b) increases. We will be interested in determining how much this probability decreases as contrast ratio increases.

1.5 Relating speckle and system parameters

In the previous section, we described the importance of looks and contrast in speckle reduction. These parameters, in turn, relate to the design parameters of the imaging system, as we explain below.

(a) Looks: In SAR systems, independent estimates of a pixel intensity are obtainable by various means [8], [17]:

(i) by using different carriers sufficiently separated in frequency so that the returns from a single resolution cell at these frequencies are independent;

(ii) by looking at a single resolution cell from various aspect angles.

Airborne or spacecraft-borne SAR systems achieve high azimuthal resolution by a technique of processing the radar

return so as to synthesize a large antenna aperture from a small physical one. Often the technique for processing involves some kind of sub-aperture processing which accomplishes (ii) above [18]. The method consists of separately processing L fractions of the synthetic aperture so as to get an azimuth resolution that is 1/L times the best one obtainable with that aperture. Then the L estimates so obtained are combined, so that we sacrifice resolution for looks. Or, for a given resolution, the aperture for L looks needs to be L times as large as the aperture required for 1 look. This makes the processing requirements for L looks about L times as complex as those for 1 look, if a specified resolution is to be attained.

Similarly, if the frequency diversity technique of (i) is used, there are L times as many carriers, and therefore about L times as much processing, for L looks as there is for 1 look.

(b) Contrast: In SAR processing by 2-dimensional matched filtering of the signal return [9], the autocorrelation function of the filter impulse response should, ideally, be a delta-function. In practice, it is a peaked function with sidelobes, and the integrated side-lobe power appears as noise in the processed return. An increase in ISLR (Integrated Sidelobe Ratio), i.e. the ratio of integrated sidelobe power to the power between 3-dB points of the central peak, has the effect of decreasing the contrast between a feature and its background (in addition to making

adjacent resolution cell returns somewhat correlated). Reducing the ISLR while meeting the other SAR specifications requires careful system design.

1.6 Scope of work

From the discussion in section (1.4), it can be seen that looks and contrast are parameters that determine the quality of the picture. From the discussion in section (1.5), it can be seen that the system design determines these two parameters. In order to be able to design a SAR imaging system that produces some tolerable level of picture quality, we need, first, to specify more precisely what we mean by picture quality, and then, to determine how the parameters L and b relate to this picture quality. This is what we attempt to do in this work.

The aspects of picture quality in which we are interested are those of detection and recognition of small features. We consider these problems in idealized cases where the pictures of interest contain well-defined geometrical shapes. We consider these pictures rather than pictures with real scenes because realistic pictures have too many intensity variations to permit any quantitative evaluation of picture quality. Moreover, it is usual, in testing the performance of an imaging system, to do so by observing how it images artificial forms, which are often geometrical. For instance, in testing television system

performance, line-pair gratings are normally used as the test pattern [19].

We begin by demonstrating the picture-degrading effects of speckle and the improvement with looks and contrast using simulated pictures of eye-charts like those used in vision testing. Because of the complexity of these pictures, we can make only a few quantitative statements about these pictures.

We then investigate, more quantitatively, the degradation caused by speckle and the effects of looks and contrast, in three specific cases. These are:

(a) the detection of small features in speckled backgrounds, where, in the absence of speckle, the small features would be somewhat brighter than the background;

(b) the detection of line-pair gratings, when the picture is degraded by speckle;

(c) discrimination between 2 or 4 specific geometrical forms, when the forms are degraded by speckle.

2. DEMONSTRATION OF SPECKLE EFFECTS

2.1 Overview

In later chapters, we present a mathematical analysis and experimental verification of geometrical feature and line-pair detection and geometrical form discrimination. Real scenes have complicated intensity distributions so that it is difficult to apply our results to them, in general. In this chapter, we show the effects of speckle, looks and contrast on some pictures of an intermediate level of complexity. We chose eyecharts because they consist of shapes that are well-defined and familiar to everyone. It is theoretically possible, but not practicable, to extend all the analysis of geometrical shapes of later chapters to the situation in this chapter; however, in Chapter 5, we show how to apply some of the approximate formulas derived there to the present case. The details of the simulations used in generating the eyecharts are the same as in the later chapters and are described in Chapter 4. However in the case of the eyecharts, we did not perform experiments as we did for the cases of later chapters. This was because the eyecharts are merely intended as a demonstration of speckle effects on picture quality.

2.2 Eyecharts

In Figures 2.1(a)-(p), we show simulated pictures of an eyechart with varying numbers of looks L and varying

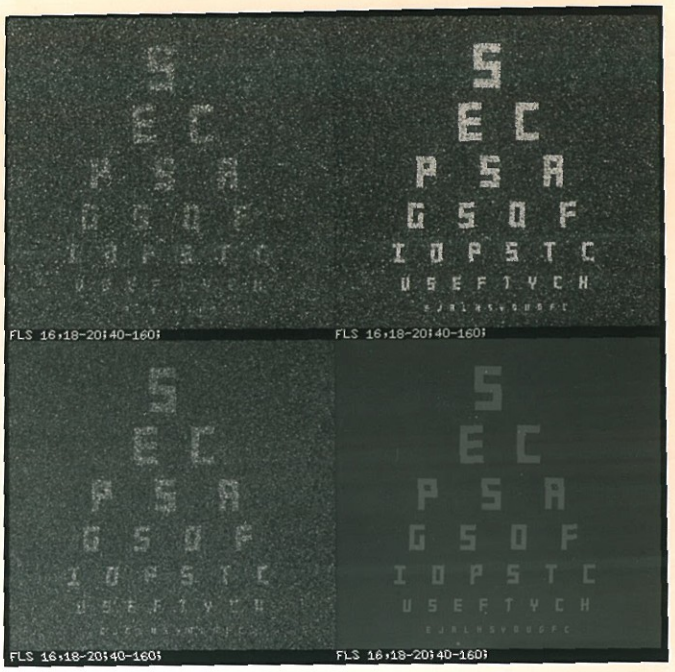

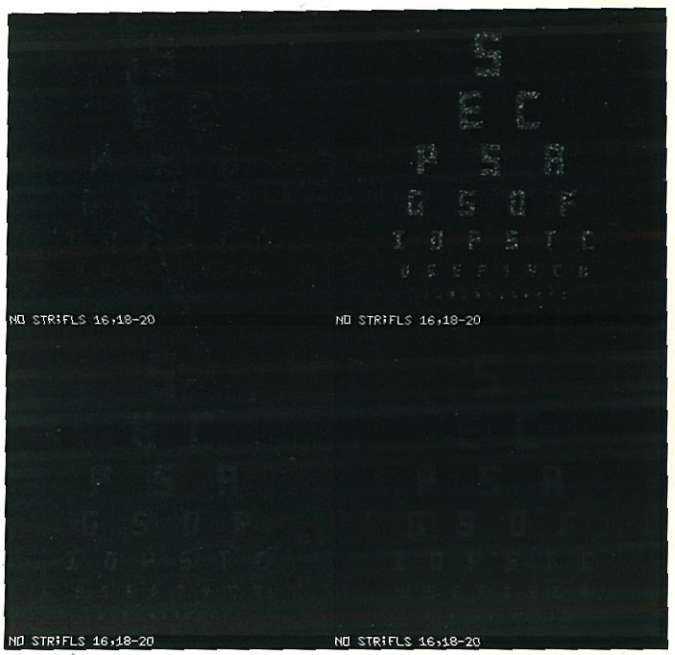

Figure (2.1): Eyecharts (a)-(h) reading from left to right and top to bottom

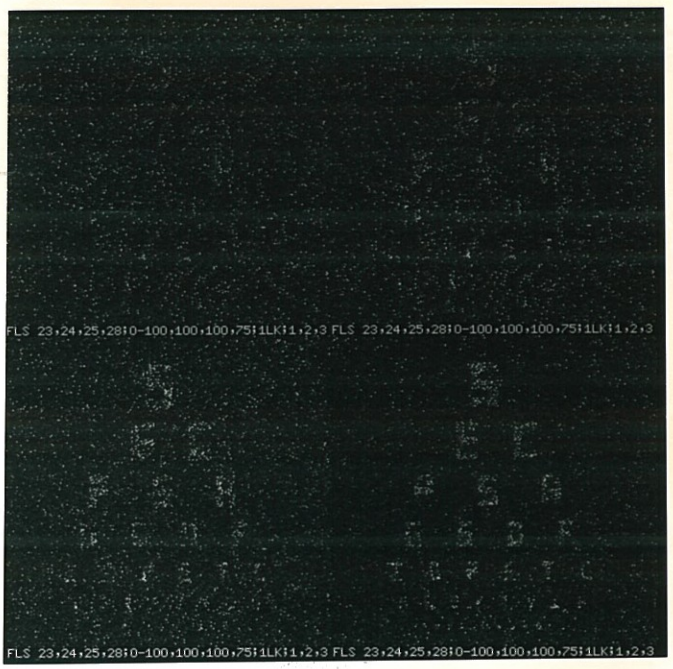

Figure (2.1): Eyecharts (i)-(p) reading from left to right and top to bottom

Table (2.1): Eyechart parameters

Figure	Number of looks per pixel, L	Contrast ratio (1+b)	(1+b) in dB	Stretch
2.1(a)	25	1.26	1.0	40-160
2.1(b)	25	1.80	2.5	40-160
2.1(c)	50	1.26	1.0	40-160
2.1(e)	25	1.26	1.0	0-255
2.1(f)	25	1.80	2.5	0-255
2.1(g)	50	1.26	1.0	0-255
2.1(i)	1	1.26	1.0	0-100
2.1(j)	1	1.60	2.0	0-100
2.1(k)	1	2.00	3.0	0-100
2.1(l)	1	2.57	4.0	0-75
2.1(m)	1	3.16	5.0	0-75
2.1(n)	1	4.00	6.0	0-75
2.1(o)	1	3.16	5.0	0-255
2.1(p)	1	4.00	6.0	0-255

Note: Figures 2.1(d) and (h) have no speckle;

(d) has a stretch 40-160, (h) is unstretched.

contrasts (1+b). All the pictures have 256*256 pixels; they are all the same set of letters in the same positions. The unspeckled set of letters with 1 dB contrast with respect to the background is shown in Figures 2.1(d) and (h). The letters were designed so as to have the minimum redundancy, given that the thickness and maximum height and width of all the letters were to be the same. It can easily be seen that a set of letters each constructed within a rectangle of height 5 and width 3 pixels has this property. Only 15 letters of the alphabet can be unambiguously defined using this rectangle. These are A (or R), C, D (or O), E, F, G, H, I, J, L, P, S, T, U, Y, which is a set large enough for our purposes.

The lowest line in all the pictures in Figure (2.1) is of size 5*3 pixels. The nth line from the bottom has a size 5n*3n pixels. The parameters used in the various pictures are summarized in Table (2.1).

Various stretch parameters have been used in the pictures. A stretch of say 40-160 means that the intensities of the pixels in the range 40-160 were linearly rescaled to cover the range 0-255. Those below 40 were set to 0, and those above 160 were set to 255. The stretched pictures in Figures 2.1(a), (b) and (c), and the pictures (e), (f) and (g) which are unstretched, are otherwise identical sets. (However, even in the unstretched pictures, values above 255 are set to 255, so that, to some extent,

there is still some saturation of intensities). Stretching seems to make no difference to recognizability in the 2.5 dB pictures (b) and (f), while in the low contrast (1 dB) pictures, the stretched versions look better.

We see by examining the pictures that the pictures with 25 looks and 2.5 dB appear better than the ones with 50 looks and 1 dB. Thus, as a preliminary result, it can be seen that increasing the contrast from 1.26 to 1.8 is better than increasing the number of looks from 25 to 50.

We note here that the process of stretching increases the variance of the intensity distribution for each of the pixels in the surround and in the feature itself, as well as the difference between the means of the distributions, all by the same factor. We come back to this point later on. Thus the probabilities of detection, as calculated in this work, should be unaffected by a histogram stretch. However, the process of "increasing the contrast", referred to in the previous paragraph, is that of increasing the value of (1+b), so that the difference between the means of the background and feature intensity distributions increases, as does the variance of the intensities of the pixels in the bright features, while the variance of the intensities of the pixels in the surround is left unaltered. Therefore this increase of contrast increases the probability of correct detection. But, as stated in Chapter 1, this increase in contrast can be done only by a change in design of the imaging system (as opposed to stretching, which can

be done at the final stage of image enhancement).

The improvement in readability of the letters as contrast is increased is most pronounced at the lower contrasts, i.e., in Figures 2.1(i), (j) and (k). If we compare the various 1-look pictures and the 25-look pictures for detectability of the lowest detectable line, we find that the 25-look, 1 dB pictures are about equivalent to the 1 look, 4 dB or 5 dB pictures. That is, changing the contrast from 1.26 to 2.57 or 3.16 is roughly equivalent to increasing the number of looks from 1 to 25. The theory described in Chapter 5 predicts approximately the same result. The same sort of conclusions can also be drawn by considering that an increase in number of pixels in a letter acts exactly the same way as increase in number of looks, and then comparing appropriate lines from different pictures. We do not elaborate on this method of comparison here, but we describe it in Chapter 5 and the results support the ones mentioned in this paragraph.

3. MATHEMATICS OF SPECKLE

3.1 Overview

Speckle is such a ubiquitous phenomenon that a considerable amount of work has been done on the mathematics of speckle [20]. Most of this work has been in connection with optical systems, and most of the mathematical treatments have taken into account the properties of a diffracting aperture and of light propagating through free space before reaching the image plane. For our purposes, we choose one of the simplest representations, since the representation of speckle is the starting point for further mathematical analysis needed in this work.

We adopt the method used by Butman and Lipes [16], which we repeat in sub-section 3.2.1 below. We will assume, throughout this chapter, that in addition to speckle, we also have additive white Gaussian noise (AWGN) [21] corrupting the signal. This inclusion of AWGN does not change the analysis very much, and is a way of taking into account, approximately, various sources of noise, like receiver noise and antenna noise. We shall see that this AWGN merely leads to a redefinition of the contrast ratio (1+b), and, with this newly defined "effective contrast ratio", all other results are just the same as if no AWGN were present. However, none of our simulated pictures had AWGN in them, so the validity of this effective contrast

ratio in accounting for AWGN has not been verified.

Throughout this work, we will be using the terms (pixel) intensity and (pixel) power interchangeably, unless otherwise stated, since we are not interested in varying the area of each pixel; in effect, we assume a system of units in which the pixel area is 1, so that power equals intensity.

3.2 Pixel SNR (PSNR)

3.2.1 Definition of PSNR

In this sub-section, we deal with the first-order statistics, i.e. mean and variance, of speckle. We also define a signal-to-noise ratio (SNR) that accounts for both speckle and AWGN, which we call [16] the "Pixel SNR" or PSNR. This PSNR provides one way of explaining the improvement in picture quality with increase in number of looks; we do this in sub-section 3.2.2.

Each estimate of a pixel intensity in a SAR system is derived by envelope- detecting an in-phase (I) and a quadrature-phase (Q) signal (so called because they are, respectively, in phase and 90 degrees out of phase, with a carrier frequency reference). The I and Q signals are themselves contaminated by system noise which is AWGN, so that the two detected voltages v_I and v_Q are given by

$$v_I = a_I + n_I \tag{3.1}$$

$$v_Q = a_Q + n_Q \tag{3.2}$$

where a_I and a_Q are the in-phase and quadrature signal voltages in the pixel in the absence of AWGN; and n_I and n_Q are the corresponding contaminating voltages due to this noise.

We assume here an idealized SAR processing scheme in which the voltages v_I and v_Q for each pixel are independent of those for all other pixels, so that we have 2K independent random variables associated with a picture of K pixels.

The signal power s, noise power n, and total power v in the processed return leading to any given pixel are given by

$$s = a_I^2 + a_Q^2 \tag{3.3a}$$

$$n = n_I^2 + n_Q^2 \tag{3.3b}$$

$$v = (a_I + n_I)^2 + (a_Q + n_Q)^2 \tag{3.3c}$$

With an ideal receiver and processing scheme, the a_I, a_Q, n_I and n_Q are all independent and the a_I and a_Q are identically distributed with a variance which we denote by P/2, while the n_I and n_Q are identically distributed with a variance denoted by N /2. All four variables have a mean value equal to 0. Then the variables s, n, v are all

exponentially distributed, because they are the sums of squares of two identically distributed, independent Gaussian random variables with zero means. The means of s, n, v are given by

$$\bar{s}=P \tag{3.4a}$$

$$\bar{n}=N \tag{3.4b}$$

$$\bar{v}=P+N=V \tag{3.4c}$$

It has been shown [22] that the image of a uniformly diffusing object illuminated with coherent light and imaged through a lens system has a mean intensity (equal to the standard deviation of the intensity) proportional to the intensity that would be observed if the same object were imaged in incoherent light. For SAR images, whether optically or digitally processed, the same results will hold, so that the P_0 appearing above is proportional to what we will call the "actual" power reflected by the resolution cell, i.e. the power reflected if speckle were absent. The proportionality constant can be taken as 1 because we are only interested in relative intensities and not in absolute intensities in the picture.

We have, so far, been considering the power return in one look. We now consider what happens when each pixel is formed by averaging incoherently over L estimates or looks.

Using v_L to represent the total (signal plus noise) power in the L-look estimate of pixel intensity, we get

$$v_L = \frac{1}{L}\sum_{i=1}^{L} v_i \tag{3.5}$$

where each v_i is now equivalent to the v in Equation (3.3).

The conventional SNR that accounts for the receiver noise is defined by

$$SNR = P / N \tag{3.6}$$

However, if we want to define a signal-to-noise ratio that accounts for both speckle and receiver AWGN, we may do so by defining the "pixel signal- to-noise ratio" or PSNR as

$$\begin{aligned} PSNR &= (\text{mean signal power}) / \\ &\quad (\text{standard deviation of the variable v}) \\ &= P / \sqrt{\text{variance of v}} \end{aligned} \tag{3.7}$$

Since the L estimates v_i in Equation (3.5) are independent, their variances add and so the variance of v_L is $1/L^2 * L$ times the variance of v_i ,i.e., it is equal to V^2/L; however, the mean of the signal component of v_L is obtained by adding the means of the L estimates and dividing by L, so

that it is equal to P_s. Therefore, we have

$$PSNR = P_s / (v/\sqrt{L}) = \sqrt{L} / (1+(SNR)^{-1}) \qquad (3.8)$$

where SNR is the quantity defined in Equation (3.5).

3.2.2 Advantage of looks

We can now explain why taking independent looks and averaging them incoherently improves picture quality. From Equation (3.8), we see that the L-look mean/(standard deviation) ratio is $\sqrt{L}$ times the corresponding 1-look quantity. This means that the random intensity fluctuation that causes the picture to appear granular is reduced and so the true intensity variations of the scene are less obscured. This should improve feature detectability in the picture.

3.3 Probability density functions

There are two techniques for analyzing the effect of random noise on the detectability of a signal: the SNR approach and the probability of detection approach. The SNR approach starts with a suitable definition of SNR as we did above. When the noise corrupting a signal is not Gaussian, the performance of the system predicted by the two approaches is not identical [23]. In our case, we are interested in speckle noise and we are concerned about the probability of

detection of features. Therefore, we need to develop the probability of detection approach. In order to do so, we first need the probability density functions of speckled signals.

So far, our notation has been general. We will now change it slightly so that it is applicable to that used in some of the later chapters. We present the probability density function for the intensity of a pixel formed from M effective looks, which we define more precisely below. Consider a small feature, B, of size D*D pixels, whose intensity B in the absence of speckle is uniform throughout the feature. We will use B to represent both the feature itself and the quantity L*(the total intensity of the feature,T). We assume that each of the D*D pixels representing the feature is formed by averaging L independent estimates of the return. Therefore, L*{the pixel intensity} for each of the pixels is formed by summing these L independent estimates. In Chapter 5, we will be interested in the total intensity T reflected by these D*D pixels in the feature, which is simply the sum of D*D pixel intensities. Thus, B=L*T is obtained by summing L*D*D intensities which are all independent estimates of the same intensity, i.e., by summing M=L*D*D "effective looks". Since we will only be interested in relative intensities, L*T is just as good for our purposes as T, and, for simplicity, we shall deal with the probability density

function of L*T. If we need to obtain the probability density function for a single pixel intensity, all we need to do is to consider the case where D=1.

The total signal plus noise power in the sum of M returns, and therefore in B has a mean value of

$$P_B=M*(P_0+N_0) \tag{3.9}$$

where we now use the symbols P_0 and N_0 for the mean signal and noise powers in any one of the M returns.

If we now have a new feature, G, whose reflectivity is (1+b) times that of feature B, then the total signal plus noise power in M returns from this feature G has a mean value of

$$P_G=M*(1+b)P_0+M*N_0 \tag{3.10}$$

If we now define an effective contrast, c, by

$$c = \frac{bP_0}{P_0 + N_0} = \frac{b}{1 + (SNR)^{-1}} \tag{3.11}$$

then the total signal plus noise power received from feature G is seen to be expressible as (1+c) times the total power received from feature B. Thus the effective contrast factor (1+c) accounts for both AWGN and speckle. Therefore, if we

replace (1+b) wherever it appears in this work by (1+c), that ought to take care of both AWGN and speckle. Since none of our simulated pictures have AWGN in them, we have used (1+b) in all our analysis, rather than (1+c).

We now consider the probability density functions (pdf) of the variables

$$B = \sum_{i=1}^{M} v_{Bi} \tag{3.12}$$

$$G = \sum_{i=1}^{M} v_{Gi} \tag{3.13}$$

where G/L represents the randomly varying intensity reflected by a feature of size D*D pixels, whose actual intensity is effectively greater than that of its background by a factor of (1+b). B/L represents the intensity reflected by a region of the background of size D*D pixels. Thus B and G have the same meaning that they had in the equations so far. Here,

$$v_{B_i} = s_{B_i} + n_{B_i} \tag{3.14}$$

$$v_{G_i} = s_{G_i} + n_{G_i} \tag{3.15}$$

are the total signal plus noise powers in the processed returns from the ith of M independent estimates of intensity of the features B and G respectively. Thus each v_{Bi} has a

mean value of MP_0 and each v_{Gi} has a mean value of $MP_0(1+b)$.

Then the variables, defined by $X=B/(P_0/2)$ and $Y=G/(P_0(1+b)/2)$ are both chi-squared distributed with 2M degrees of freedom [24]. Equivalently, $X/2=B/P_0$ and $Y/2=G/(P_0(1+b))$ are gamma distributed with parameter M. Specifically, the pdfs of B and G are

$$p_G(y) = \frac{1}{[P_0(1+b)]^M} y^{(M-1)} \frac{e^{-y/P_0(1+b)}}{\Gamma(M)} \qquad (3.16)$$

$$p_B(x) = \frac{1}{(P_0)^M} x^{(M-1)} \frac{e^{-x/P_0}}{\Gamma(M)} \qquad (3.17)$$

These are the desired probability density functions. Since B and G are multiples of a gamma distributed or of a chi-squared distributed random variable, we will refer to B and G themselves as being gamma or chi-squared distributed, meaning that they follow the density functions given in Equations (3.16)-(3.17) above.

For large M, we can approximate the distribution functions of B and G by Gaussian distribution functions having the proper means and variances, i.e., the mean values are MP_0 and $MP_0(1+b)$ respectively, and the standard deviations are $1/\sqrt{M}$ times the respective means. This approximation results from the Central Limit Theorem [21] which is applicable because G and B are sums of a large

number of independent, identically distributed random variables. These approximate Gaussian probability density functions are

$$p_B(y) \approx \frac{1}{\sqrt{(2\pi M)}\, P_0} \exp[-(y-MP_0)^2/2P_0^2 M] \qquad (3.18)$$

$$p_G(x) = \frac{1}{\sqrt{(2\pi M)}\, P_0(1+b)} \exp\{-[x-MP_0(1+b)]^2/2P_0^2(1+b)^2/M\} \qquad (3.19)$$

However, a better approximation [2.8] for a chi-squared distributed variable, such as Y, is to represent $\sqrt{(2Y)}$, rather than Y itself, by a Gaussian distributed variable with a mean value of $\sqrt{(4M-1)}$. This approximation gives good results for $2M>100$, for most purposes. We will use this approximation in some cases, and the expressions for B and G using this and some other approximations are given where they are needed.

4. EXPERIMENTAL METHOD: GENERAL FEATURES

4.1 Overview

In order to verify the detectability calculations we will make in the coming chapters, we perform psychophysical experiments using computer-simulated pictures. The details of the simulations and experiments used depend on the specific calculations being verified, and these are described in the later chapters. However, all the simulations and experiments have certain features in common and we will describe them in this chapter.

4.2 The simulations

We consider pictures generated by the computer, which simulate the output of an ideal processing system in the case where there is no additive system noise. Sometimes, speckle is included under the term "noise", but in our work, we use the term "noise" to mean only system noise, excluding speckle.

To generate the various speckled pictures we use in this work, we started by generating a matrix of intensity values representing the "actual" or unspeckled picture of interest. Then we degraded the picture by speckle as follows. Each of these intensity values which we may call P_I was replaced by a random intensity v_L corresponding to the v_L in Chapter 3. This was made to have the desired gamma

density function by averaging L independent random variables, v_i, each having an exponential probability density function, and a mean value equal to the actual or unspeckled intensity P_I. Thus in the case of the pictures of Chapter 5, each of the random intensities v_{Gi} for the small bright features had a mean value of $P_0(1+b)$, while each of the random intensities v_{Bi} for the background had a mean value of P_0.

The pictures were all generated on a Dichomed image generator connected on-line to the computer (SEL-32) used for generating the random variables. The Dichomed converts intensities of resolution cells represented by 8- bit pixels (range 0-255) into optical intensities exposing a photographic film. For initial examination, black and white Polaroid type 52, Polapan 4*5 Land film was used. Higher quality negatives could also be obtained and printed and the pictures shown in this dissertation are obtained by this method.

4.3 Histogram stretching

We also studied, for each picture, the effect of uniformly stretching the picture histogram to make the full dynamic range of 0-255 available for pixel values initially lying in some range R. This range R was chosen by generating the picture histogram and examining it to find the range in which more than 99% of all pixels lay. This is a popular

technique used in picture enhancement. One effect of this stretching is to change the appearance of the picture, making the background look much more grainy. Theoretically, as noted in Chapter 2, the stretching process should not affect the probability of detection of features or line pairs or the probability of correctly discriminating forms, if only the probability density functions are considered. However, in cases where the contrast between a feature and its background is very low, stretching should help in detecting features, because the film is unable to resolve fine variations in gray shades, and because the viewer's eye can more easily spot well-defined contrasts. The same sort of considerations should apply to line pair detection and form discrimination. In all our experiments, we used stretched or unstretched pictures, whichever gave better feature, line or form detectability for the particular set of parameters used in each picture.

4.4 Picture and pixel size

In the calculations that follow, we assumed each pixel to be square and non-overlapping with its neighbors. To ensure that this was so in the pictures generated, each pixel was expanded to fill a square of size at least 2*2 pixels centered on the original pixel. Thus even though the Dichomed output beam does not have a square profile, the final pixels are of uniform intensity in a square region and almost independent of adjacent pixels. In some pictures, a

particular pixel was expanded to fill a square of 4*4 pixels or more, but the appearance of the final picture is about the same as when each pixel was expanded to fill a 2*2 square. Throughout this work, when we refer to "pixel", we mean this composite pixel of square shape. Each such composite pixel corresponds to exactly one resolution cell of the real or imaginary scene being imaged. The results would, however, be approximately valid even if this "squaring" of pixels were not done.

Most of the pictures generated have 200*200 pixels (meaning the composite pixels referred to above) and there are up to 2048*2048 "actual pixels" in each picture. Some of the pictures have 100*100, 256*256, 300*300 or 512*512 pixels. The Polaroid Land film used in the experiments has a size of 3.5*4.5 inches, and an actual image area of 3.5*3.5 inches and a resolution of 35 to 40 lines per mm. However, once the negatives are made, there are many possible print sizes.

4.5 Imaging system characteristics

All the pictures we show, unless otherwise stated, were generated using the "log select" switch on the Dichomed. This causes the input exposure codes, which are proportional to the pixel intensities generated on the computer, to be linearly related to the photographic density, that is, the logarithm of the intensity transmittance of the film [26].

It is reasonable to assume [26] that the reflectance of a developed film is proportional to its transmittance, and [27] that the observer's eye responds to the logarithm of the intensity of light entering it. After the logarithm extractor in the viewer's eye, we assume that there comes a probabilistic evaluation by his brain. Thus, with the log select, it is reasonable to assume that the response evoked (by a pixel of intensity P), at the final stage where its probability is evaluated is proportional to the gamma-distributed pixel intensity generated by the computer. This sequence of assumptions makes it reasonable that the probability densities that are finally required in decision making are the gamma-density functions described in Chapter 3.

In many of the cases in Chapter 5, we also generated pictures using the linear select, which causes the exposure codes to be linearly related to the film intensity transmittance. After preliminary experiments, in which we found that there was no clear-cut improvement or worsening in the detectability of features, we decided to use only the log select.

This is the extent to which we consider the imaging system and film characteristics. In our analysis, we neglect other complexities and imperfections in the imaging process.

4.6 Absolute and relative brightness levels

In all our pictures, we are concerned with two pixel intensity levels, which we denote by P_0 and $P_0(1+b)$. In all the pictures, the "effective contrast ratio" (1+b) is an important parameter in the detection problem. However, the absolute intensities P_0 and $P_0(1+b)$ are not important in any of our problems, because the overall brightness level of the picture can be very easily changed in the developement of the photograph and in all the cases we consider, the picture is assumed to have an overall brightness level that is clearly above the threshold of detectability. There may be some small dependence of detectability of features (or line-pairs or form or whatever) on absolute brightness levels, because of the properties of the visual detection mechanism, but we neglect this effect in our calculations, and experiments that we performed for the purpose seemed to justify this assumption.

We also assume, since we are investigating the degrading effects of speckle, that the differences in brightness that the observer is looking for would be definitely detected in the absence of speckle. Thus, we do not consider the effect of (1+b) being so small as to be near the observer's just-noticeable-difference level of relative brightness (which is normally about 1% of the average brightness of the two levels involved [27]) or near the intensity-resolving limit of the photographic system.

By looking at the unspeckled versions of the pictures simulated, we found that the contrast ratios of 1 dB and above, which we used, satisfy this requirement.

4.7 The experimental procedure

Psychophysical experiments were conducted using the computer-simulated pictures to verify the theoretical results that we derive in the succeeding chapters. In all of the experiments, several pictures with a fixed set of parameters were shown to an observer, who was asked to make some decision. The proportion of his responses that were correct, $\hat{p}_c$, was then used as an estimate of the probability of correct decision p_c for the particular case. The theoretically predicted p_c and experimentally estimated $\hat{p}_c$ were compared to determine the correctness of the theory.

In all experiments, the choice of exactly what directions to give the observer was made following the principles that

(a) the observer ought to be told all the possible details he could know without biasing him; and that

(b) the experimental situation should, as far as possible, resemble the real situation which might occur when a scientist, such as an oceanographer or geologist, is studying the pictures. Thus, we tell the observer in our experiments such things as one might expect the scientist to know in

some real case which our experiment parallels.

The first of the above principles follows the procedure given by Green and Swets [28] (in the Appendix on Experimental Procedures). The details of the second principle are given later.

In our experiments, some pictures were initially shown to an observer (TJG) with a given set of directions. The observer then made a decision for each sample or picture, and his responses were compared with the correct ones. The probability of error was estimated by the percentage of errors made by the observer in his decision task.

The same procedure was repeated with these same pictures and with the author (VNK) as observer and the results agreed with those obtained with TJG above to within a 5% significance level as measured by the t-test [29]. It was possible to use VNK as observer, because of the way in which the pictures were designed. For instance, in the line pair detection experiments, each picture consisted of a set of patterns that randomly had or did not have a line pair grating in them, so that VNK had no a priori knowledge of whether her decision would be correct. There is a similar justification for using VNK in the other sets of experiments. Making observations on the many pictures that were needed in the experiments was too time-consuming to make the use of any other observer feasible.

5. DETECTION OF SMALL FEATURES

5.1 Overview

By detectability of a feature, we mean the ability to see that the feature is present without necessarily being able to distinguish its shape. A circle (which might represent a crater) may be a good feature to use in our calculations and simulations. But we use squares because the mathematics is simpler. We consider the probability, p_c, of an observer's (correctly) detecting a small square of D*D pixels of uniform intensity, in a uniform background of W*W pixels, and brighter than this background by a factor (1+b).

In this chapter, we make certain assumptions about the visual process involved in the detection problem stated above and derive a mathematical relationship between the probability of detection p_c and various parameters of the picture. We simulated several pictures of the type that we made the calculations for, using various sets of parameters. We then used psychophysical experiments to check our calculations. The nature of the pictures considered in this chapter is such that it was not feasible to perform a large enough number of experiments so that we can say that the theory developed here is the only correct one. However, the experiments are extensive enough to show that the predicted probabilities are very close to the experimental ones and that the theory developed here at least provides a way of

predicting the probability of detection given the various picture parameters. We feel this is useful, because the ability to resolve small features like rocks and craters is an important characteristic of an imaging system, and it is desirable to be able to specify the level of confidence with which such features are detectable in the imagery produced by the system.

We also derive some simple approximate equations which, we feel, will help in choosing the system parameters (1+b), i.e. contrast ratio, and L, i.e. number of looks per pixel, which will lead to desired probabilities of correct detection of small features.

5.2 The parameters

In this section, we define the specific problem for which we made our calculations and simulations used for our experiments. The picture we are interested in consists of a square background of W by W pixels, in which a single bright feature is located at some place that is unknown to the observer. The feature is a square of size D by D pixels, where D<<W. In the absence of speckle, the background has a uniform intensity, P_0, and the feature has a uniform intensity of $P_0(1+b)$. However, we are interested in speckled pictures, so that neither the background nor the feature appear uniformly intense any more, and the location of the feature becomes difficult to determine. We assume

that the picture corresponds to something that might be obtained by an L-look SAR system, so that each pixel intensity is obtained by summing L independent estimates of the unspeckled pixel intensity.

It is reasonable to assume that the detectability of the feature depends on the values of W, D, L, and the contrast ratio (1+b) between the feature and background intensities. The absolute intensities P_0 and $P_0(1+b)$ are not important in our problem, but their ratio is, as explained in Chapter 4.

5.3 Theoretical calculations

5.3.1 The detection model

We make our calculations concerning the detectability of features according to a detection model that uses two basic assumptions. Although we make our calculations for the case of a square of size D*D pixels, the general method is applicable to the problem of detecting small features of any shape.

(1) The first assumption we make is that the viewer's visual system, by which we mean the entire eye-brain system involved in visual perception, responds to the total optical power reflected by a stimulus, provided the stimulus subtends a small enough angle at the eye (i.e., a small enough visual angle). This has, in fact, been demonstrated, by various workers, to be true under many different

conditions. To quote from Budrikis [30], "it is clear that before detection (of light) takes place (in the visual system) the input distribution is subjected to spatial and temporal filtering...Both the spatial and temporal filter characteristics show bandpass resonances. However, a multiplicity of data on detection of stimuli on plain backgrounds can be explained by a filter with only low-pass characteristics." We assume that this result can be extended to the speckled backgrounds that are of interest to us. We are not going to be concerned with the temporal filtering, since the viewer in all cases of interest to us has unlimited time for observation. Visual transfer function curves, such as the one given by Budrikis [30], suggest that we can assume that spatial (intensity) summation occurs within about 10' arc of visual angle subtended by the stimulus. Cornsweet [27] also gives this angle as the one in which spatial summation occurs. The impulse response of the transfer function of the visual system will not be exactly equuivalent to a summation, but it will be a weighted summation of intensity within some visual angle. We assume that it is a simple summation, so that the exact details of the visual transfer function of each observer are not considered. Appendix (5.1) summarizes more details of the visual system of interest here.

The pictures we used in our experiments were of a small enough size so that, in most of them, the bright features to

be detected subtend at most 10' arc at the eye at normal viewing distances. Even if this was not the case, the observer was told to adjust his viewing distance so that he could best make a decision. This is a reasonable thing to do, because our calculations are intended to correspond to situations where a scientist is looking at a picture and is free to view it under the most advantageous circumstances.

Furthermore, we assume that, in these real situations, the scientist knows what size of object he is looking for in the map. Therefore, in our experiments, we told the observer the size of square he should look for. Under these circumstances, we can assume that the observer's visual system responds to the total intensity of areas of size D*D pixels. Since we assume that each pixel is formed by averaging L independent estimates of the pixel intensity, the total intensity of each of these D*D sized squares is formed by summing $M=LD^2$ independent estimates of intensity of the (uniformly reflecting) square, and dividing this sum by L. Since all the areas have the same L, we will neglect this factor of L, and think in terms of the sum of M intensities (for each of the areas), in which case we can use the pdf's of Chapter 3.

We will say more about this assumption of visual summation in the next subsection (5.3.2).

(2) The second assumption we make is about the decision

technique used by the observer. In this connection, we first consider how an ideal detector or receiver which responds to the total intensity of an area of size D*D pixels would behave. One way in which such a detector might detect a target of the type considered here would be by looking at the total power returned from every square of size D*D pixels and determining which one was the brightest. We assume that the human observer acts as an ideal detector in detecting visual signals in noise. This assumption has often been made before in work on vision and shown to be reasonable [28], [31], [32].

5.3.2 Justifying the detection model

In the previous subsection, we made the assumption that the observer's visual system responds to the sum total power from all pixels in D*D pixel areas. In order to justify the assumption that his visual system responds to D*D sized areas and not areas of some other size, we argue that the observer is looking for targets of that size and therefore mentally evaluates areas of that size. Even if this is not a sufficiently strong reason to support our assumed detection model, there are various other arguments that we can make to justify using this model.

(1) First of all, the idea of using the total power from all the pixels in the feature as the important parameter determining its detectability is certainly reasonable if the picture is held far enough away from the eye so that the

angle which the feature subtends at the eye is just equal to the angular resolution of the eye, which is about 1' arc [33]. A scientist studying a SAR map to detect bright features is free to hold it far away from him so that he sees several adjacent pixel intensities blurred or merged together. As long as the feature is not too large in extent, the viewing distance required to make most or all its pixels appear to merge is small enough so that the amount of light from the picture entering the eye provides a detectable stimulus to the eye.

(2) An alternative to the detection model described in sub-section 5.3.1 goes as follows. The observer decides that an area of D*D pixels is the feature he is looking for if more than half (or some other fraction, f, greater than half) of the pixels comprising it are brighter than some threshold level. This assumes that the observer looks at the individual pixels rather than the total power from some area of size D*D, and would thus apply to large targets as well as small ones. We did a few sample calculations based on this assumption, where the threshold level with which each pixel intensity is compared is set by a Neyman-Pearson strategy [34]. This method is more complicated than the one described in sub-section 5.3.1 because there are several parameters to be chosen (similar to the one parameter p_t that we will describe for the model of sub-section 5.3.1). However, for a choice of these parameters leading to a

highly reliable detection of the target, the required values of M for given b (or b for given M) are very close to the ones predicted by the model of sub-section 5.3.1. Therefore, we decided to use the latter model.

(3) Examination of preliminary simulations performed to verify the results predicted in sub-section 5.3.1 give us a clue to the validity of the model. These preliminary simulations are not sufficient in number to be used in detailed experiments like the ones finally made. However, they do indicate that, for a wide range of parameters b and M, the probability of the observer's making a correct decision as a function of b and M follows the same trend, whether the pixel intensities are directly displayed on the film or their square roots, and whether the "linear" or "log" select option (Chapter 4) is used. Since the method of sub-section 5.3.1 calculates the probabilities on the basis of a summation of pixel intensities by the observer, the probability of correct decision ought to depend on the specific function of intensity displayed. However, we may regard the sum of the pixel powers as a measure of detectability of the feature. This measure is related to some other possible measures we could define, such as the sum of the logarithms of the pixel powers or the sum of the square roots of the powers, or the number of pixels in the feature with powers greater than a certain threshold value. All these measures have the property that they increase if

the power of some of the pixels in the feature increases. Because of the correlation between the various measures, a probability of correct decision based on any of these measures ought to yield a fairly similar trend of results. This explains why the method of choosing a square, half of whose pixels exceed some given level in power, as a feature gives about the same results as the method of sub-section 5.3.1. It also explains why the particular (monotonically increasing) function of intensity displayed is not too important. The simulations and experimental results are not accurate enough to allow us to choose between the various models. Our aim here is merely to provide an approximate method of specifying the relation between the system parameters b and M so that features of given sizes can be detected. Since the model and the method of sub-section 5.3.1 give results that agree quite well with the experiments on the simulated pictures, the model is acceptable.

5.3.3 Theoretical analysis

5.3.3.1 The probability of correct detection

In this section, we outline the analysis used in our theoretical calculations of the probability of correct detection using the model of sub-section 5.3.1.

Since we are considering small square features, the length or width W of the picture is much greater than the

side D of the square. Then there are about $V=W^2/D^2$ independent squares of size D^2 pixels in the picture. (This number may or may not be an integer, but for large values of V, we may replace V by the integer closest to V, for our purposes.) There are actually about $N=W^2$ possible squares of size D^2 in the picture, because each of the W^2 pixels can be considered to be, for instance, the central pixel of a square of size D^2 pixels (neglecting the effects at the four edges of the picture). Only V of these N squares are independent in the sense of being non-overlapping.

The question arises : should we use N or V as the number of squares in the background that the feature should exceed in total power, in order to be correctly detected to be the desired feature? We will refer to these background squares as comparison squares. If there were a square grid of size D by D pixels superposed on the picture, we could ask the observer to confine his attention to D*D pixel squares defined by the grid. In this case, there are definitely only $V=W^2/D^2$ comparison squares in the background. However, this is not a very realistic situation, and so we chose not to have such a grid in the picture. In the actual, gridless case, we could take $N=W^2$ as the number of comparison squares, provided that in our calculations we took into account the correlation between the W^2 comparison squares. Where $N=W^2$ is large, as it is in all our cases of interest, the lack of independence of all N

squares is not very important, provided we are interested in values of p_c greater than about 1/2, as we shall clarify later.

According to the model we are using, a correct decision occurs whenever the feature has a total power greater than that of each of the N comparison squares. If we represent the total power reflected by the i^{th} one of these N comparison squares as B_i and the power reflected by the feature as G, then a correct decision occurs whenever G exceeds each of the B_i. Thus we can express the probability of correct decision p_c as

$$p_c = \Pr \{G>B_1,\ G>B_2,\ldots.G>B_N\} \qquad (5.1)$$

where Pr{x} is the probability of the event described by x. We can also write this as (5.2)

$$p_c = \Pr \{G> \text{ the largest of the N variables } B_i\ ,\ i = 1,2,\ldots N\}$$

Or, defining the probability distribution function of the largest of the N variables B_i by $A_N(x)$, i.e.,

$$A_N\ (\ x\ ) = \Pr\{\text{largest of the N variables } B_i \leq x\ \} \qquad (5.3)$$

we can write

$$p_c = \int_0^\infty p_G\ (x)\ \ A_N\ (x)\ \ dx \qquad (5.4)$$

If the N powers B_i are independent, then $A_N(x)$ can be written as $F_N(x)$ where

$$F_N(x) = (F_B(x))^N \qquad (5.5)$$

Here, $F_B(x)$ is the probability distribution function of the B_i. Therefore,

$$F_N(x) = \left[\int_0^x p_B(y) \, dy \right]^N \qquad (5.6)$$

If the variables B_i are not independent, then we must express $A_N(x)$ as the joint probability distribution function of the B_i, given by

$$F_{1,2,....N}(x,x,....x) = Pr\{B_1 \leq x, B_2 \leq x, ... B_N \leq x\} \qquad (5.7)$$

In the case of independent powers B_i, we can write p_c as

$$p_c = \int_0^\infty p_G(x) \; F_N(x) \, dx \qquad (5.8a)$$

$$p_c = \int_0^\infty p_G(x) \left[\int_0^x p_B(v) \, dv \right]^N dx \qquad (5.8b)$$

In the case of the B_i not being all independent, we can write

$$p_c = \int_0^\infty p_G(x) \; F_{1,2,....N}(x,x,....x) \, dx \qquad (5.9)$$

5.3.3.2 Extreme value theory

The expression $A_N(x)$ obtained above cannot be expressed in a simple closed form for a general probability density function (pdf) of the B_i, and certainly not for a chi-squared or Gaussian distribution of B_i. However, for large N and large x, an asymptotic expression can be found for $A_N(x)$ [35], which has the same form for a large class of distributions of the variable B_i (which we shall refer to as

the initial distribution). This is obtained from "extreme value theory." For initial distributions of the exponential type, to which the chi-squared and the Gaussian belong, the asymptotic distribution is referred to by Gumbel [35] as the first asymptote, and is defined by

$$A_{Na}(x) = \exp[-\exp\{a_N (x - x_N)\}] \qquad (5.10)$$

where x_N is the so-called characteristic largest value of the initial distribution, and a_N is the so-called intensity function evaluated at this point. More precisely, for a population with pdf equal to f(x) and a distribution function equal to F(x), we define the characteristic largest value x_N in a sample of size N from this population by

$$F_N(x_N) = 1-1/N \qquad (5.11)$$

The "intensity function" m(x) for the distribution is defined by

$$m(x) = \frac{f(x)}{1-F(x)} \qquad (5.12)$$

and, if we use Equation (5.11), the intensity function evaluated at $x=x_N$ becomes

$$a_N = N\, f(x_N) \qquad (5.13)$$

5.3.3.3 Interdependence of the N background squares

As we have pointed out in the previous section, the $N=W^2$ comparison squares in the background are not all independent, so that the correct expression for $A_N(x)$ ought

to be given by Equation (5.7) and not Equation (5.6), which is a much simpler expression. However, if we can use the first asymptote instead of the exact expression for $A_N(x)$, (and we show, in the next sub-section, that we can) then the interdependence of the N values of B_i causes no special problems. This is because, for a large number of identically distributed random variables, N, the probability distribution of the largest value is given by the same asymptotic expression as that given in Equation (5.10), whether or not the variables are interdependent (see Gumbel [35]). A justification for this, Gumbel notes, is that the asymptotic distribution of the extremes depends on the properties of the initial distribution only for large values of the variate, where the influence of the interdependence may vanish.* This idea was first put in mathematical form by G.S.Watson [36]. Thus, in our problem, where N is indeed large, we will use $N=W^2$ together with the first asymptote in the circumstances where the latter is applicable, which are described in the next sub-section. In any case, when we are making a comparison between the first asymptote and the exact expression for $A_N(x)$, we use Equation (5.6) instead of the more correct Equation (5.7), because the explicit form of Equation (5.7) is difficult to obtain. For this reason, the "exact expression (value) for p_c" that we use in the following paragraphs and sections is not really exact since it refers to Equation (5.8) rather than to Equation (5.9).

* See Appendix (5.2) for details.

5.3.3.4 Using the first asymptote

In our application of the first asymptote, we want to use it as part of an integrand where the integration runs from 0 to infinity (Equation5.4). For this purpose, we need to ascertain the range of its validity and also to check how close it is to the correct expression in this range. We note here that even if the first asymptote can be used in Equation (5.4) for the entire range 0 to infinity, a closed-form evaluation of p_c is still not possible, and it has still to be evaluated by computer. However the computer evaluation in this case is much easier than a computer evaluation of Equation (5.8). This is because we are interested in p_c of about 0.5 or more. Hence the value of the integral is of the order of magnitude of 0.5, while the integrand of Equation (5.8) is extremely small over most of its range, because N is very large. We do not have this problem if we use the first asymptote instead of $F_N(x)$ in Equation (5.8), because there is no N^{th} power of any integral to be determined. Also if we use the first asymptote, we need to evaluate a single integral, as opposed to the double integral of Equation (5.8), where, in addition the inner integral $F_B(x)$ has to be evaluated with a high degree of precision because it is then raised to a large power. These problems did, in fact arise when we evaluated the exact integral of Equation (5.8) in order to compare the results with those obtained using the first asymptote, Equation 5.10.

The first asymptote is a good approximation for the function $A_N(x)$ only for large enough x. Therefore, we do not use it to evaluate the entire integral from 0 to infinity. Instead, we use the technique described in the next 3 paragraphs, which is valid for p_c not too small; it is certainly valid for p_c greater than 0.5.

For large values of N, the distribution function $A_N(x)$ has a step-function-like appearance, with the transition occurring near the characteristic largest value x_N. This can be seen to be true from the form of the first asymptote, Equation (5.10), or by reflecting on the form of the exact expression for $A_N(x)$, Equation (5.6) for the case of independent B_i. For instance, for $N=10^4$, in order for $A_N(x)$ to be equal to 0.05, $F_B(x)$ has to be 0.99970047, while for $A_N(x)$ to be 0.95, $F_B(x)$ has to be 0.99999487. The two values of x at which $A_N(x)$ is 0.05 and 0.95 respectively, are therefore very close, which results in a step-function-like curve for $A_N(x)$. Figure (5.1) shows the nature of the first asymptote $A_{Na}(x)$, along with the initial probability density functions $p_B(x)$ and $p_G(x)$ for the background and feature powers, all for the case where $N=10^4$, $M=LD^2=35$ and $b=1.0$ respectively. The cases we are interested in all have N greater than 10^4, so that the $A_{Na}(x)$ curve in Fig. (5.1) is quite representative.

The integrand for the calculation of p_c in Equation (5.4) is a product of the two factors $A_N(x)$ and $p_G(x)$. The

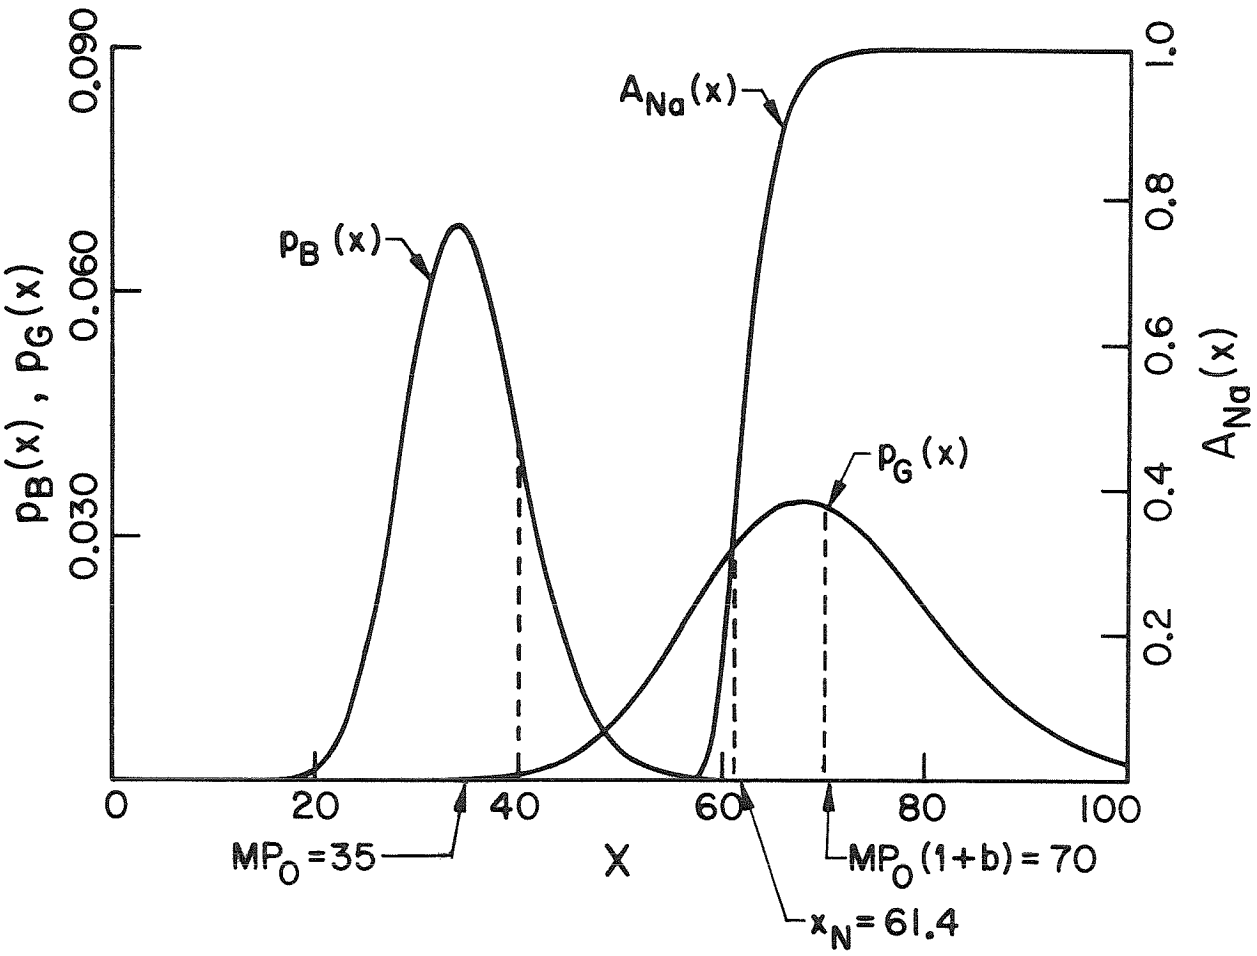

Figure (5.1): Probability density functions $p_B(x)$, $p_G(x)$ with means MP_0, $MP_0(1+b)$ and first asymptote $A_{Na}(x)$ with characteristic largest value x_N; all have $N=10^4$, $M=35$, $b=1$, $P_0=1$

first of these factors becomes very small for low values of x as we go away from $x=x_N$, and the second one is very small for both low and high values of x, far from the peak which occurs near $(1+b)P_0$. If the integral for p_c is to have any appreciable value at all, it must be because the peak of the density function $p_G(x)$ occurs either at a value of x greater w than x_N, or at a value of x that is less than x_N but close to it, so that both the factors of the integrand are reasonably large in an overlapping region within the range

of integration. This requires that the mean $(1+b)P_0$ of the pdf $p_G(x)$ be large if N is large, or in other words, that the contrast must be high to ensure a high value of p_c.

We find by actual calculation that the first asymptote is a good approximation for $A_N(x)$ at least in the region $x>x_0$ where $A_N(x)$ is greater than about 0.05. Over the range of b and M investigated (which is the range covered in Table 5.1) the value of the integral of Equation (5.4) is negligible for $x<x_0$. By negligible, we mean that these contributions to p_c are no more than 0.05. Thus, for p_c greater than, say, 0.5, we may neglect the contribution to the integral of Equation (5.4) from this region, while the contribution from the remaining region may be evaluated using the first asymptote in place of $A_N(x)$ in Equation (5.4).

5.3.3.5 Relation between b, M and N given p_c

In the previous sub-sections, we considered the problem of finding p_c for given values of b,L,D,W. We will refer to this as the analysis problem. We are also interested in the converse problem, which we will refer to as the design problem, of determining some of these parameters when the others are given, so that a desired value of p_c is obtained. Since we cannot obtain a closed form expression for p_c even using the first asymptote as in the previous sub-section, this converse problem cannot be easily solved by the method used for the analysis problem, unless we are willing to use

an iterative approach. In this sub-section, we will make further approximations so as to obtain a simple relation between b, $M=LD^2$ and $N=W^2$ that will result in a specified $p_c = p_t$ which we will call the threshold b, M, N and p_c respectively. The method will give good results provided the specified value of p_c is not too small. It is reasonable that, if we want to design a system that will enable reliable detection of features of a certain size, we would like to design it for a fairly high probability of correct detection, for instance 0.95. The method we describe will be applicable to this case; if greater accuracy is desired in specifying the p_c obtainable with the system, we can always refine the design using the more exact approach of the previous sub-section.

The approximate relation between b, M and N is obtained in several steps.

(1) We first note that, as shown in Gumbel [35], for N>100, the agreement between $A_N(x)$ and $A_{Na}(x)$ is very good over the range

$$x_0 \leq x \leq x_1 \qquad (5.14)$$

defined as the interval in which

$$0.05 \leq A_N(x) \leq 0.95 \qquad (5.15)$$

This statement is for a Gaussian initial distribution. As N increases, the agreement gets better. Also, for the actual chi-squared initial distribution, the agreement is even

better, as we verified in some cases. The Gaussian is not a very good distribution from this point of view (of convergence of $A_N(x)$ to the first asymptote), whereas the exponential distribution is, and, in this application, the chi-squared resembles the exponential for large values of the argument. We use this property to divide the range of integration of Equation (5.4) into 3 parts, defined by

$$\begin{aligned} 0 \leq x < x_0 \qquad & \text{(Range 1)} \\ x_0 \leq x \leq x_1 \qquad & \text{(Range 2)} \\ x_1 < x < \infty \qquad & \text{(Range 3)} \end{aligned} \tag{5.16}$$

Since $A_N(x)<0.05$ over Range 1, we see from Equation (5.4) that we make an error of much less than 0.05 if we replace the integral, over this range, in Equation (5.4) by 0. In range 3, we may, again with an error of much less than 0.05, set the value of $A_N(x)$ equal to 1. The middle range is the one over which the first asymptote is valid.

(2) We find our approximate relation between b, M, N and p_c by requiring that the integral over Range 3 be equal to p_c, and substituting the value of 1 for the factor $A_N(x)$. We find that doing this leads to a simple relation of the type we want. Obviously, this relation will be reasonably correct only if the integral over the other two ranges is negligible in comparison to that over Range 3. We have already stated that the integral over the first range is

less than 0.05. For the case where p_c is equal to 0.5, we made a very rough estimate of the integral over Range 2, by using the value of the integrand at the midpoint of the range. We calculated this mid-range value of the integrand by using the first asymptote in place of $A_N(x)$, since this is a valid approximation in this range. (This is the only use we made of the first asymptote in this calculation, and this is the only reason for choosing the three ranges instead of combining the first two ranges into one.) We found that over the range of M and b and N that we are interested in (as given in Table (5.1) at almost the end of this chapter) the integral over Range 2 is between 0.2 and 0.3, when the overall p_c is about 0.5. However, if $p_c>0.9$, we find that the error in p_c caused by neglecting the integral over the first two ranges is less than 10%. This dependence of the integral over the first two ranges on the value of p_c can be explained by the same arguments as those given at the end of the previous sub-section. That is, if b increases, the peak of the curve of $p_G(x)$ in Figure (5.1) shifts to the right, the non-zero region of the integrand shifts to the right, and p_c increases.

(3) We are now ready to obtain the desired relationship. To do this, we use the fact [25] that, for a nmber of degrees of freedom 2M greater than 100, the chi-squared distribution of a random variable denoted by χ^2 can be approximated by assuming that $(2\chi^2)^{1/2}$ has a Gaussian distribution. This is

more accurate than directly using the Gaussian distribution for χ^2 (as we could according to the Central Limit Theorem [21]). In deriving our approximate relation between b, M, N and p_c, we assume that this Gaussian approximation to $(2\chi^2)^{1/2}$ holds even for $2M<100$, recognizing, however, that our answers will be in error to some extent for small M.

Let the chi-squared distribution function with n degrees of freedom be denoted by $P_{\chi^2}(z \mid n)$, and the area under the tail of the Gaussian distribution function by the so-called Q function Q(.), and let z_1 and z_{p1} be the $100p_1\%$ point of these two distributions respectively. We are interested in the case where the value of p_1 is $(0.95)^{1/N}$, but for any general value of p_1, we have the following approximate relationship between z_1 and z_{p1}:

$$z_1 = 1/2 \ (z_{p1} + (2n-1)^{1/2} \ )^2 \qquad (5.17)$$

where

$$P_{\chi^2}(z_1 \mid n) = p_1 \qquad (5.18)$$

and

$$Q(z_{p1}) = p_1 \qquad (5.19)$$

Here the Q(.) function is defined as usual by

$$Q(z) = \int_z^\infty \frac{1}{(2\pi)^{1/2}} \exp(-v^2/2) \ dy \qquad (5.20)$$

Thus, the x_1 of Equation (5.16) is given approximately by

$$x_1 \simeq (P_0/4)[\, z_{p1} + (4M-1)^{1/2} \,]^2 \qquad (5.21)$$

because the $X_i = B_i/(P_0/2)$ are chi-squared distributed (see Chapter 3) and Equation (5.17) then applies to the X_i.

The integral of Equation (5.4) then has a contribution from Range 3 of Equation (5.16) that is given by

$$p_{c3} = \int_{x_1}^{\infty} p_G(x)\, dx \qquad (5.22)$$

or, in terms of the pdf of the $Y_i = \dfrac{G_i}{P_0(1+b)/2}$, which are chi-squared distributed, we have

$$p_{c3} = \int_{y_1}^{\infty} p_Y(y)\, dy = Q_{\chi^2}(\, y_1 \mid 2M \,) \qquad (5.23)$$

where

$$y_1 = \frac{x_1}{P_0(1+b)/2} \qquad (5.24)$$

and where

$$Q_{\chi^2}(z \mid n) = 1 - P_{\chi^2}(z \mid n) \qquad (5.25)$$

is the integral, between z and infinity, of the chi-squared pdf with n degrees of freedom. Now, using the Gaussian approximation for $(2Y_i)^{1/2}$ [25], we have

$$Q_{\chi^2}(z \mid n) \simeq Q(z_a) \qquad (5.26)$$

where

$$z_a = (2z)^{1/2} - (2n-1)^{1/2} \qquad (5.27)$$

Thus, with 'z' replaced by y_1 of Equation (5.23) and using Equations (5.21), (5.24) and (5.26), we have

$$p_{c3} \simeq Q\left[\frac{z_{p1} + (4M-1)^{1/2}}{(1+b)^{1/2}} - (4M-1)^{1/2}\right] \qquad (5.28)$$

If we now decide to make p_{c3} equal to the given p_c, which we have denoted by p_t, then the argument in square brackets in Equation (5.28) must be equal to the $100p_t\%$ point of the normal distribution, which we denote by z_{pt}. That is,

$$\frac{z_{p1}(N) + (4M-1)^{1/2}}{(1+b)^{1/2}} - (4M-1)^{1/2} = z_{pt} \qquad (5.29)$$

where we have denoted z_{p1} by $z_{p1}(N)$ to emphasize that $p_1=(0.95)^{1/N}$ is a function of N. Equation (5.29) is the desired relationship between b, M and N for a given value, p_t, of p_c. For $b<<1$, we can get a very simple relation between b and M

$$bM^{1/2} = \text{constant}. \qquad (5.30)$$

This shows that, for low contrasts, increasing b is more effective than increasing M. This agrees with what we found from the eyecharts in Chapter 2.

We could choose p_1 to be $(0.99)^{1/N}$ instead of $(0.95)^{1/N}$. The reason for choosing p_1 to be $(0.95)^{1/N}$ is simply that, for our purposes, a 5% error (which is about what this choice leads to in cases of our interest) in our calculations is tolerable.

5.3.3.6 Choice of p_t

In order to be able to make predictions about the quality of pictures obtainable with a given set of parameters, we need to choose a suitable value for the acceptable probability of correct decision, p_t in the previous sub-section. We now show how to make this choice of p_t. We use $p_t=0.95$ in all our numerical calculations.

(1) We know that, if M=0, the probability of correct decision is 0, while for M=infinity, it should be 1. It would be desirable to be able to determine a threshold value of M, denoted by M_t, such that, for a given b and N, an increase in M above $M=M_t$ makes little difference to the probability of correct decision. To see what value of p_t gives an M_t satisfying this requirement, we show a typical plot of p_c versus M, for fixed b=2.16 and $N=10^4$. The calculations for this graph were done by computer, using the method we have discussed in this section. For values of p_c greater than 0.5, the extreme-value method was used to evaluate the integral of Equation (5.4), and for lower values of p_c, the exact expression of Equation (5.6) was used. The results are plotted in Figure (5.2), which shows $p_e=(1-p_c)$ versus M. Also shown in this figure, are the experimentally obtained estimates, $\hat{p}_e=(1-\hat{p}_c)$, which were determined as explained in the next section.

We see from this figure that if p_t is chosen to be about 0.95, increasing M above the corresponding M_t value would cause a change in p_c that would be hard to detect experimentally. That is, it would not lead to a noticeable improvement in picture quality. Thus, for example, in Figure (5.2), increasing M from 2 to 20, i.e., by a factor of 10, reduces the error probability $p_e=1-p_c$ from 0.94 to 0.05. On the other hand, increasing M from 20 to any higher value causes p_e to reduce from 0.05 to at most 0, and in the kind of application that we have here, this reduction would not be easily noticed, whereas the reduction from 0.94 to 0.05 would be, as can be verified from the simulated pictures shown in the next section. (In an area like coding theory, however the region of interest would be that in which p_e decreases from 0.05 to 0, rather than from 0.94 to 0.05.)

(2) The use of the first asymptote in place of $A_N(x)$ in Equation (5.4) becomes more accurate as p_c increases. In order to be able to use this simpler calculation for p_t it is desirable to have p_t fairly large. If p_t is chosen to be 0.95, as we have done, the approximate Equation (5.29) becomes applicable and can be used for solving the design problem of finding one of the parameters b, L, D and N, given the others and a specified p_t.

(3) Finally, if we want to specify a p_t and obtain the value of feature size D at which the feature is detectable with

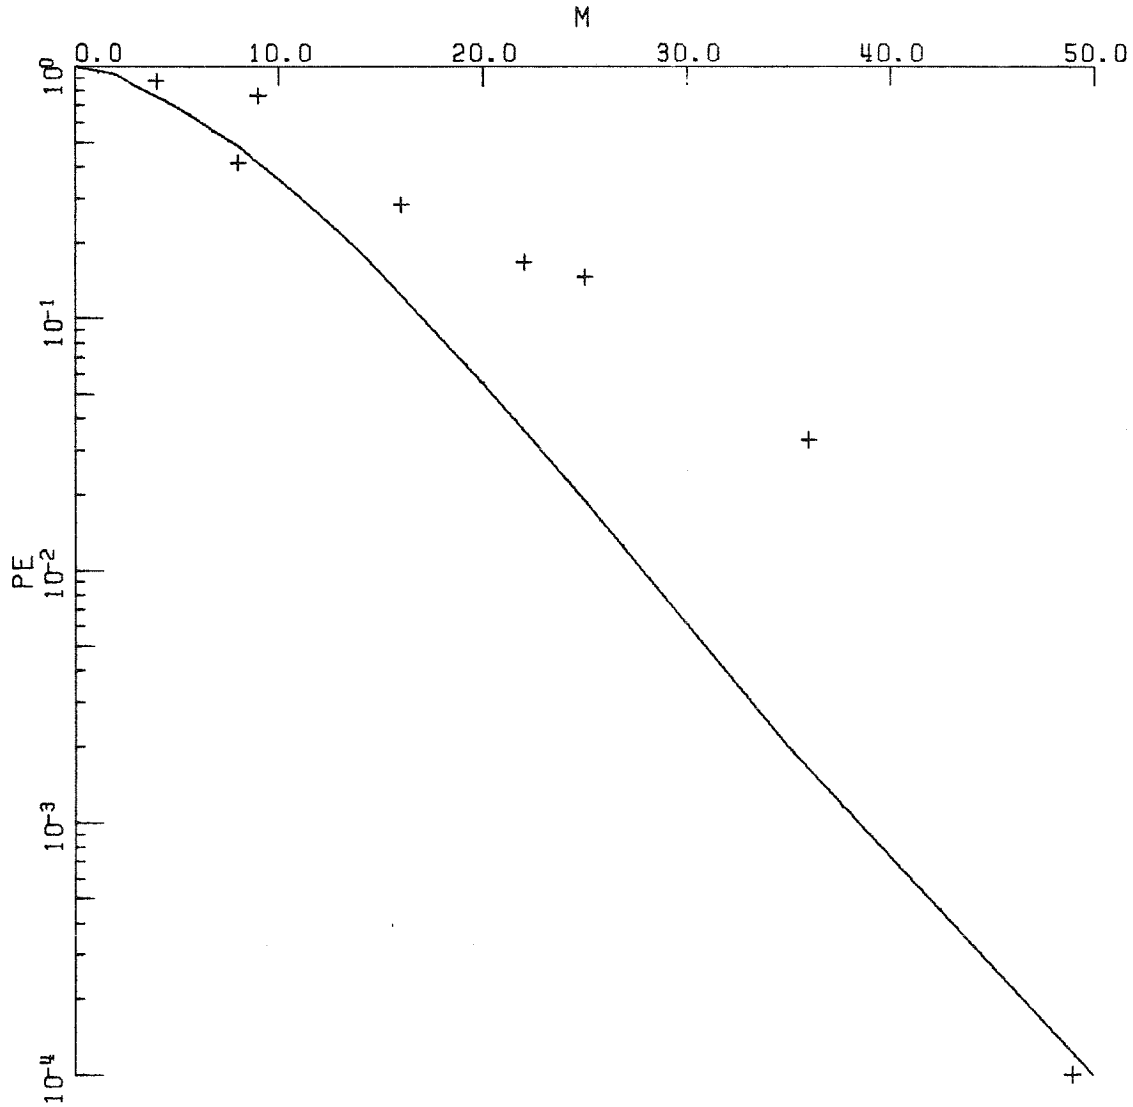

Figure (5.2): Theoretical and experimental probability of error, p_e versus M for $N=10^4$, b=2.16

Table (5.1): Detection of D*D square in 100*100 picture

(1+b)	(1+b) dB	L	D	M	# of observations n	Theoretical p_c	Experimental $\hat{p}_c$ lower bound	$\hat{p}_c$	$\hat{p}_c$ upper bound
1.26	1	7	6	252	12	0.51	0.05	0.25	0.55
		14	6	504	48	0.65	0.37	0.52	0.67
		8	8	512	48	>0.65	0.63	0.77	0.88
		10 *40	8 4	640	60	0.95	0.71	0.82	0.93
		18 *162	6 2	648	84	>0.95	0.85	0.93	0.98
2.0	3	2	6	24	24	--	0.04	0.17	0.37
		1	5	25	60	0.63	0.20	0.32	0.45
		1	6	36	48	0.72	0.52	0.67	0.80
		9	2	36	24	0.72	0.33	0.54	0.74
		1	7	49	96	--	0.59	0.71	0.82
		7	3	63	12	0.95	0.64	0.92	1.0
		2	6	72	12	>0.95	0.80	0.92	1.0
		1	10	100	48	>0.95	0.93	1.0	1.0

Table (5.1) (continued)

(1+b)	(1+b) dB	L	D	M	# of observations n	Theoretical p_c	Experimental $\hat{p}_c$ lower bound	$\hat{p}_c$	$\hat{p}_c$ upper bound
3.16	5	1	2	4	48	0.24	0.05	0.13	0.25
		2	2	8	24	0.52	0.37	0.58	0.78
		*8	1						
		1	3	9	60	0.59	0.14	0.24	0.36
		1	4	16	60	0.85	0.59	0.72	0.82
		22	1	22	24	0.96	0.63	0.83	0.96
		1	5	25	48	0.98	0.72	0.85	0.94
		1	6	36	60	0.999	0.88	0.97	1.0
		1	7	49	48	0.9999	0.93	1.0	1.0
		1	8	64	12	>0.9999	0.76	1.0	1.0
6.32	8	1	2	4	12	—	0.45	0.75	0.95
		8	1	8	24	0.95	0.63	0.83	0.96
		1	3	9	12	0.95	0.77	1.0	1.0

* Pictures with these L and D values were considered equivalent to those in the preceding row (which have the same M), in experiments as well as in theory.

Table (5.2): Predicted threshold parameters for $p_t = 0.95, N = 10^4$

(1+b) dB	M	D for L=12
1	641	7.3
3	65	2.3
5	22	1.4
8	8	0.8

near certainty, we might need to have p_t even higher than 0.95.

5.3.4 Results of theoretical analysis

The results of the theoretical calculations of p_c for various sets of parameters b, L, D, all for $N=10^4$, are shown in Table (5.1) along with the experimental results, obtained as described in the coming section. In the theoretical calculations, as well as in the experiments, we concentrated on the sets of parameters with b=2.16; theoretical calculations for various sets of such parameters were done in order to obtain Figure (5.2).

The calculations for the other contrast ratios considered are merely illustrative, and are not done for the complete range of parameter sets simulated, since the calculations are time-consuming, especially for p_c lower than 0.5, where we used the exact expression for $A_N(x)$ instead of the asymptotic one.

Table (5.2) shows the calculated threshold parameters M, given 4 different values of b, required for $p_c=p_t=0.95$. The calculations are all for the case of $N=10^4$, and were all made using the design relationship Equation (5.29). From the values of $M=LD^2$, in this table, we can calculate D for given values of L. The values of D for a typical case, L=12, are also shown in Table (5.2).

5.4 Simulations and experiments

5.4.1 The experimental procedure

Ideally, we can verify the theoretical calculations of probability, p_c, of correctly detecting a square in a darker surround by using a large number of simulated pictures, which we will refer to as "samples". The samples were generated and the experiments conducted with a viewer as described in Chapter 4, and the percentage of his correct decisions used as an estimate of p_c.

Since the theory is applicable only to cases where the object to be detected is small compared to the size of the picture, it would take a very long time to simulate a number of pictures large enough to verify the calculated probabilities of correct detection. This problem is not serious in simulations like the ones in Chapters 6 and 7 of this thesis, where the size of each pattern used in the experiments is small. In the present case, the need to simulate a 100*100 pixel picture for each experimental observation limits the number of samples of pictures that can be simulated.

For these reasons, it is not possible, with our experiments, to determine whether the theory used is the correct one, as opposed to other possible theories that can be used to predict p_c, and we do not attempt such a verification. We merely hope to verify that the model we

have used gives acceptably good results and that our predicted parameters b, L, D, N for p_c equal to p_t, where p_t is about 0.95 are approximately correct.

In our experiments, the observer knew the size of feature to be detected and the contrast ratio and was asked to designate one square of size D*D in the picture as the feature. The observer was aware that, because of speckle, the features might not look square, and that the apparent shape of the feature was not relevant.

To analyze the results, we compared, for each sample, the observer's choice of feature location with the actual one. If the actual location coincided with the observer's choice, then the decision on that particular sample was noted as correct; from a number, n, of such samples the probability of correct decision was estimated by the fraction of the total number of decisions that were correctly made. We denote this experimental estimate of p_c by $\hat{p}_c$.

5.4.2 The simulations

Several samples corresponding to various sets of parameters (L, b, D), but all with the same picture size (100*100 pixels) were simulated and shown to the observer. To reduce the number of pictures we actually simulated 200*200 pixel pictures and had 4 bright features in each picture, one in each quadrant of the picture, but the location of each

feature in each quadrant was random. In the experiments, 3 quadrants were masked at all times by a black card-paper mask, so that the net result was that each picture resulted in 4 samples of size 100*100 pixels. The sample pictures shown in Figure (5.3) are these 200*200 pictures. All four samples in each of the pictures simulated had the same set of parameters b, L, D. A detailed discussion of the pictures of Figure (5.3) is given in the sub-section on "Examples of simulated pictures".

The number of samples with a given set of parameters ranged from 12 to 72. In order to demonstrate the variation of p_c with total equivalent number of looks, $M=LD^2$, for a fixed b, we simulated at least 48 samples of various sets of parameters that had a fixed b=2.16 (5 dB contrast ratio), but L and D varying from set to set. The results of experiments using these sets of parameters are shown in Figure (5.2), which shows $\hat{p}_e=(1-\hat{p}_c)$, i.e., the experimental estimates of probability of error, as a function of M.

Other values of b simulated were b=0.26 (1 dB contrast), 1.0 (3 dB contrast) and 5.318 (8 dB contrast). The range of L and D corresponding to the 3 dB contrast case was also as extensive as that for the 5 dB case, but some of these sets had only 12 or 16 samples. The range of values of L and D for the other values of b (3 dB and 8 dB cases) was not so extensive. The sets of parameter values in these cases were chosen to lie in a small range around the

theoretically expected value of M_t required for p_c equal to p_t=0.95, and are merely expected to demonstrate approximately the M value required to obtain this threshold probability of correct decision p_t.

5.4.3 Examples of simulated pictures

5.4.3.1 Overview

In our experiments, we used a large number of simulated pictures. It is neither feasible nor very revealing to include all of them here. However, we show, in Figure (5.3), some examples of these pictures. An examination of these pictures will show the plausibility of the theoretical model and some of the results derived in section (5.3), in particular, Equation (5.29). Each picture shows 4 samples with different sets of parameters. In the experiments, these simulated pictures were viewed with 3 quadrants covered or masked at a time. Each picture simulated had 200 * 200 pixels, but the last 6 lines at the bottom were replaced by the caption before obtaining the photograph. The caption under each picture gives the parameters used in producing that picture, and also a code, such as B26 or S10 #1, which identifies the picture. All the pictures shown here have D, L and (1+b) in dB specified in the caption. In addition, we have also specified the amount of histogram stretching. This is written as "no str" for the case of no stretch or as "str a-b" for the case where the pixel intensities lying in the range a-b were uniformly

rescaled to extend over the available range 0-255 of pixel intensities.

5.4.3.2 The samples

(1) We first discuss the photograph which shows the 4 samples designated as B6, S30, B9 and B22. B6 and S30 are unspeckled pictures. The 4 features, one in each quadrant, are clearly seen in these pictures. B9 and B22 are samples with speckle, where the parameters are chosen to be at (or near) the predicted threshold parameters for 95% probability of correct detection. See Table (5.2) which predicts $M=LD^2=8$ for the 8 dB case of B9, and $M=LD^2=63$ for the 3 dB case of B22. We can see the high detectability of features in each of these samples.

(2) The second photograph shows samples called B6, B11, B28 and B29. The B6 shown here is the speckled 8-look version of the (unspeckled) one discussed above. B6 and B11 are 5 dB pictures, with M=8 and M=22 respectively, which correspond to values below and at the threshold given in Table (5.2). The 4 features in B6 are almost invisible, whereas in B11 they are quite clear. Thus a comparison of B6 and B11 illustrates the validity of the predicted threshold for the 5 dB case. B28 and B29 similarly illustrate the 8 dB threshold parameters.

(3) The third photograph shows B26, B27, S10 #1 and S9 #3. B26 and B27 are 5 dB samples with M=16 and 25 respectively,

(1)

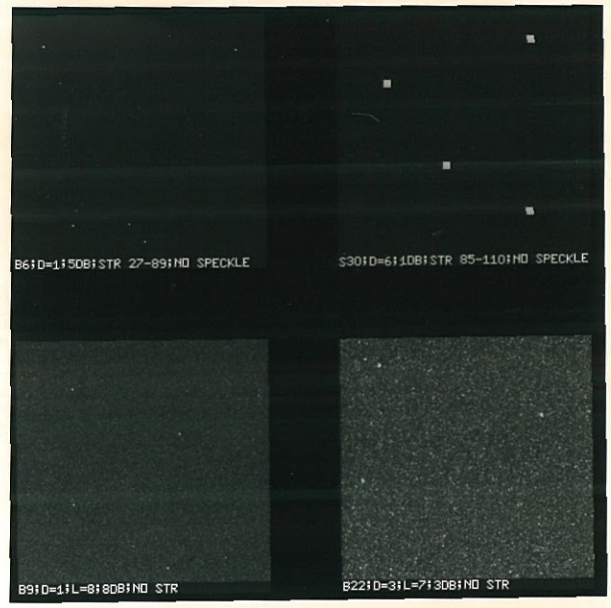

(2)

Figure (5.3): Detection of squares--examples of simulated pictures; photographs 1 & 2

(3)

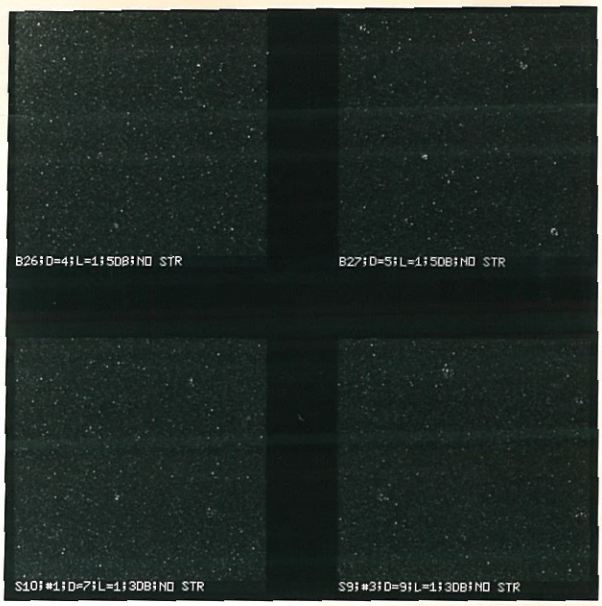

(4)

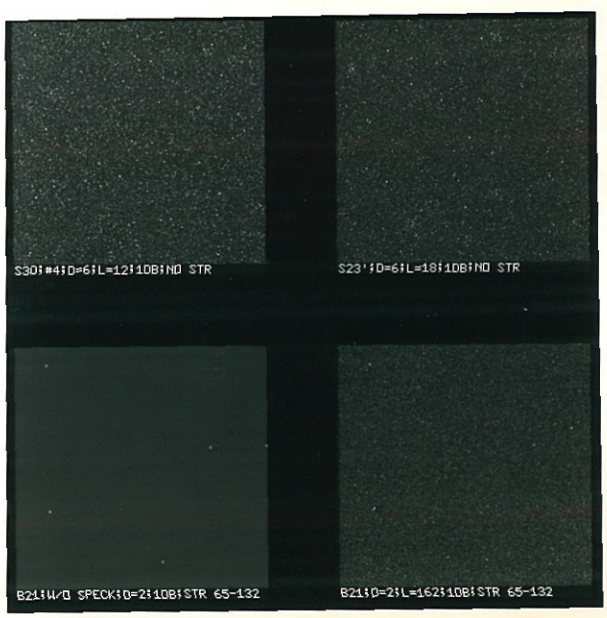

Figure (5.3) continued: Photographs 3 & 4

and again illustrate the validity of the threshold parameters, just as B6 and B11 did. The only difference is that B6 and B11 differ in their D values, while B26 and B27 differ in L. This shows the validity of using $M=LD^2$ as an "effective number of looks per feature". Similar conclusions follow by comparing B29 and B22 (of photograph 1).

S10 #1 and S9 #3 are 3 dB pictures with M=49 (below threshold) and M=81 (above threshold). Again, the predicted threshold for the 3 dB case appears reasonable.

(4) The fourth photograph shows 1 dB samples. S30 #4 and S23 have M=432 and 648 respectively, and show the validity of the 1 dB threshold parameters predicted. S30 #4 shows that in a 1 dB, 12 look picture, a D=6 feature is not easily detectable. We will refer to this fact in Chapter 6. The other two samples are B21 without speckle and the speckled version of the same sample, B21, which has M near threshold. A comparison of B21 and S23 again justifies the use of $M=LD^2$.

5.4.4 Results and conclusions

The experimental and the theoretically expected probabilities of correct decision, $\hat{p}_c$ and p_c respectively, are shown in Table (5.1). The upper and lower bounds on the $\hat{p}_c$ shown in Table (5.1) are the 95% confidence limits calculated from the standard formula for confidence limits on estimates of proportions such as the $\hat{p}_c$ are [29]. As

stated earlier, in both the theoretical and experimental calculations, we concentrated on the sets of parameters with b=2.16. The theoretical results for the sets with other values of contrast ratio are not so extensive.

From Tables (5.1)-(5.2) and Figure (5.2), we see that

(1) the experimental p_c are usually lower than the theoretical ones, but both follow the same sort of increase with b and M;

(2) however, the difference between the experimental and theoretical p_c is small for high p_c (above approximately 0.95) --this supports our use of a threshold probability of correct decision p_t equal to 0.95;

(3) the simulations for the various values of b with M around the predicted threshold values M_t required for $p_c=p_t=0.95$, show that the analytically obtained M_t values are reasonable criteria to use in predicting, approximately, the L and D combination required to make a bright feature of contrast ratio (1+b) detectable with reasonable certainty.

5.5 Application to eye-charts

We now apply the theoretical analysis that we have described in this section to explain some of the results seen in the eye-charts of Chapter 2. Because of the greater complexity of the eye-chart problem, we can analyze only certain aspects of the problem by the methods developed so far in

this chapter.

In the 1-look per pixel eyecharts discussed in Chapter 2, the lowest line in the 1-dB contrast case is not detectable (while all the upper ones are), but as the contrast increases, there is a point at which it becomes visible. This probability of detection of the lowest line differs from the p_c we have been dealing with so far in this chapter because the situation here is slightly different. If we consider the letters in the lowest line as small features, the probability of detection of each of these letters is independent of all the others, and we may consider that any one of these becomes detectable if its total intensity exceeds that of an area of comparable size in the surround. By surround, we mean here the region at the bottom of the picture below the second-to-last line, in which we assume there is some number, N, of "comparison areas". This extends the argument for squares that we have used so far to other shapes and ought to be at least roughly valid, since we are merely looking for a small bright cluster and not trying to recognize any shapes.

Now, in all the pictures, the letters in the upper lines of the eye-chart are large enough to be clearly visible as letters, if not recognizable. Therefore, an observer, when asked to say how many lines of letters are present, will include the last line if one or more of the

The fact that 9 features are present (there are 9 letters in the last line) again increases the probability of detection of at least one of them. But, on the other hand, it is very likely that the observer needs to see more than one feature in the last line before he decides that the line is present. Finally, the presence of the other lines of letters will perhaps influence him to expect the same intensity of a letter in the last line as there is in a bright area of the same size in one of the upper lines.

All these differences from the simple situation assumed in section (5.3) will invalidate the use of Equation (5.29) in the same form in the present situation. However, we may expect that, if $z_{p1}(N)$ and z_{pt} in that equation are suitably redefined (to be, perhaps, a value that is specific to the configuration of letters in the particular eye-chart used), then the equation will still hold. That is, we assume that the general functional relationship between b and M given in Equation (5.29) will be valid for pictures more complicated than the one it was derived for. Thus, we may test the hypothesis that

$$(4M-1)^{1/2}\ [(1+b)^{1/2} - 1] \ = \ K_{pic} \qquad (5.31)$$

at the detection threshold, where K_{pic} is a constant that depends on the spatial intensity variation in the picture in the absence of speckle, and on the desired threshold

probability, p_t, of correct detection of the last line in the picture.

We find reason to support Equation (5.31) if we examine the speckled eye-charts of Chapter 2. Thus, for the 25 look, 1 dB pictures, the lowest line is just detectable and we assume that it represents the threshold of detectability; if we now look for the same level of detectability (subjectively judged) of the lowest line in the 1 look pictures, we find that it occurs in the 4 dB or 5 dB pictures. This is the experimental result. To see what Equation (5.31) predicts, we first note that in all the pictures, the lowest line has 5*3=15 pixel letters, so that in the 25-look, 1 dB pictures, the lowest line has M=25*15=375, b=0.26, so we find, by substituting in Equation (5.31), that K_{pic} has the value of

$$K_{pic} = 4.74 \qquad (5.32)$$

For the 1-look pictures, where M=5*3*1=15 for the last line, in order to obtain the same level of detectability, we need

$$(1+b)^{1/2} - 1 = \frac{4.74}{(4M-1)^{1/2}} = 0.617 \qquad (5.33)$$

or

$$10 \log (1+b) = 4.2 \qquad (5.34)$$

Thus, the contrast needs to be 4.2 dB, which agrees very well with the observed value, stated in Chapter 2.

APPENDIX (5.1)

Relevant properties of the visual system

A-5.1.1 The visual system

A great deal of research has been done on understanding the details of human visual perception, both by physiological methods and psychophysical methods. Cornsweet [27] gives many of the details and several references. Our interest here is in the spatial filtering properties of the human eye. There is considerable evidence for this spatial filtering and many visual phenomena, such as Ricco's law (which states that the threshold intensity of a small bright object in a dark field is, within limits, inversely proportional to its size) and the Mach band phenomenon (light lines appear to darken adjacent dark lines) can be explained on the basis that adjacent receptors in the human retina excite or inhibit each other. Cornsweet [27] shows how these phenomena agree with the retinal interaction theory. It has been shown that this interaction does occur and can explain both spatial and temporal effects (time delays in transient response of the eye are an example of the latter) [37], [38]. Helson [38] concludes that

(1) the area of an object being detected in a darker background acts as a luminance, i.e, increase in area is equivalent to increase in luminance in this respect;

(2) in the retina, relatively weak neural interactions

summate, while intense stimulation in one area inhibits weaker stimulations in neighboring ones.

These interactions explain the modulation transfer function of the eye obtained experimentally [39],[40], as shown by Cornsweet [27].

Applying the above results to our problem is complicated by the fact that our pictures are viewed directly by the observer in normal conditions; we do not have the controlled experimental conditions (e.g. dark adaptation, uniform background) under which the visual transfer function was determined [39], [40], or those under which Cornsweet [27] shows that the visual system integrates intensities lying within 10' arc of each other. However, the visual transfer function has been successfully applied to image enhancement problems [41] where the experimental conditions referred to above are not met either, but are similar to the normal viewing conditions used in our experiments and by scientists studying SAR maps. As another example of the effects of spatial filtering under ordinary viewing conditions, we have the work of Harmon et and Julesz [42], whose results show that looking at a picture corrupted by high-frequency noise from a viewing distance that is greater than normal helps improve its recognizability. The blurring of the picture caused by viewing it from a distance is equivalent to low-pass filtering the picture.

We assume, therefore, that our observer's eyes integrate intensities within an area subtending 10' arc at the eye, and that he is able to resolve targets subtending 1' arc at the eye, as is generally accepted. (All the visual angles referred to here are those subtended by any one dimension of the object at the observer's retina.) A-A-5.1.2 Visual angles in our pictures

The simulated pictures used in our experiments have various picture sizes (in inches) and numbers of lines and pixels per picture. We give here typical calculations that will enable a quick conversion of feature size to visual angle. In those pictures where 200 pixels fit into 3.5 cm, and are viewed at a distance of 40 cm (16"), which is the normal viewing distance, each pixel subtends a visual angle of about 3' arc. Thus, a feature of about 2 or 3 pixels on a side would subtend less than 10' arc at the retina at the normal viewing distance, and it is reasonable to apply the summation approximation to it. For larger features, the viewing distance needs to be larger, for instance for D=6 to 7, the viewing distance needs to be doubled. Since, at a distance of a few times the normal viewing distance, the picture intensity variations are still visible, it is reasonable to try to locate features of the kind of sizes we have used in our experiments by looking at them from afar, thus averaging out the speckle. This is what a scientist looking for small features might well do.

APPENDIX (5.2)

Interdependence of background powers

Watson [36] shows that if the variables in a strictly stationary sequence are m-dependent and unbounded above, the asymptotic distribution of the largest in a sample of size N is the same as in the case of independence, provided Equation (A-5.1) below is satisfied. A sequence of random variables $\{X_i\}$ is called m-dependent if $|i-j|>m$ implies that X_i and X_j are independent. For our 2-dimensional sequence of $N=W^2$ background powers, this is true in both dimensions for m=D. In order to justify the use of Equation (5.10) with $N=W^2$, all we need to do is to show that Equation (A-5.1) holds.

$$\lim_{c\to\infty} \frac{1}{Pr(x_i>c)} \cdot \max_{|i-j|<m} Pr(x_i>c,\ x_j>c) = 0 \qquad (A\text{-}5.1)$$

This has been shown to be true [36] for X_i normally distributed, so we can use this result for large M, where our variables approach Gaussian variables. For general M, it suffices to show that for any two of the dependent variables X_1 and X_2 in our problem,

$$\frac{Pr(X_1>c,\ X_2>c)}{Pr(X_1>c)} \to 0 \quad \text{as } c\to\infty \qquad (A\text{-}5.2)$$

Now (setting $P_0=2$ for simplicity) X_1 and X_2 are each formed by summing (M_0+M_1) exponential random variables, of which M_0

are common. If we let the common part be denoted by Y, we have

$$X_1 = Y+Z_1$$
$$X_2 = Y+Z_2 \qquad \text{(A-5.3)}$$

where Y, Z_1 and Z_2 are independent random variables with $2M_0$, $2M_1$, $2M_1$ degrees of freedom (dof). Then we can write Equation (A-5.2) in terms of the pdfs of Y and Z_1 (or Z_2) as

$$\frac{\int_0^\infty p_Y(y)dy \left(\int_c^\infty p_{Z_1}(x_1-y)dx_1\right)^2}{\int_0^\infty p_Y(y)dy \int_c^\infty p_{Z_1}(x_1-y)dx_1} \to 0 \quad \text{as } c\to\infty \qquad \text{(A-5.4)}$$

By direct calculation, this can be verified to be true for cases where M_0 and M_1 are small. We can also give a heuristic argument for the validity of Equation (A-5.2) in our problem. Equation (A-5.2) implies that the joint probability of X_1 and X_2 being very large is small compared to the probability that X_1 is very large. One would expect that if this is not true, it is because the common part, Y is large. However, using the asymptotic expansion for the distribution function, we get

$$\frac{Pr(Y>c)}{Pr(X_1>c)} = \frac{Q_{\chi^2}(c|2M_0)}{Q_{\chi^2}(c|2M_0+2M_1)} =$$

$$\frac{\frac{c^{(M_0-1)}e^{-c/2}}{\Gamma(M_0)2^{(M_0)}}\left(2+2^2\frac{(M_0-1)}{c}+\ldots\right)}{\frac{c^{(M_0+M_1-1)}e^{-c/2}}{\Gamma(M_0+M_1)2^{(M_0+M_1)}}\left(2+2^2\frac{(M_0+M_1-1)}{c}+\ldots\right)} \sim c^{-M_1}\to 0 \text{ as } c\to\infty \qquad \text{(A-5.5)}$$

This makes it plausible that Equation (A-5.2) should hold in our problem, and therefore that the asymptotic distribution

of the largest of the B_i should be the same as in the case of independence. However, we note that, from a practical point of view, p_c is not strongly dependent on the exact value of N for large N, at least for $p_c > 0.9$, so that an order of magnitude estimate of N is sufficient at least for predicting threshold parameters as we do in sub-section 5.3.3.5, which is one of our main concerns.

6. LINE PAIR DETECTION

6.1 Overview

A common criterion used in television picture characterization is line pair resolution. A picture consisting entirely of alternating bright and dark lines is used as a test picture. As the width of the lines is made smaller, there comes a stage at which the average viewer can see only a uniform gray picture instead of the line pattern. The inverse of the width of the lines, d, or some quantity related to it, like the spatial (fundamental) frequency of the lines at this stage, is a figure of merit of the TV system. The lines used in the test pattern usually have a square-wave or sinusoidal spatial intensity distribution.

People sometimes attempt to use the same criterion, that is, line pair resolution, to characterize SAR systems and others that produce speckled pictures. In the absence of speckle, if the system bandwidth(BW) was high enough and the signal-to-noise ratio (SNR) was high enough, we would have a situation similar to that with a good TV system. However, large system BW and high receiver SNR are not sufficient to produce a good picture in a coherent imaging system, because of speckle effects. Thus a line pair of a spatial frequency resolvable with a speckle-free system may or may not be resolvable in one with speckle. We might still attempt to characterize the overall system by the

spatial frequency of resolvable lines, hoping to account for the degradation due to speckle by a decrease in the system performance.

In this chapter, we investigate the relationship between the various parameters characterizing a line pair grating and the probability, p_c, of correctly detecting it. We find that for a typical system, and a test pattern of a size that might typically be used, for example, one of 100*100 pixels, line pair detection is a useless way to characterize resolution in the presence of speckle.

A more stringent criterion is the detectability of small features, which we dealt with in a previous chapter. Thus, for a contrast ratio and number of looks per pixel at which line pairs in a picture consisting entirely of line pairs are recognizable, a small object may not be detectable.

In this chapter, as before, we assume that the degradation in the picture is caused by speckle alone, i.e. we consider a system that is speckle- limited rather than BW- or SNR- limited, so as to isolate the effects of speckle.

6.2 The method used

6.2.1 Overview

Our approach to the problem consists of a combination of theory and experiment that will enable us to relate p_c to the parameters of the grating. Because current understanding of the visual mechanism is incomplete, and because of the difficulty of making exact calculations in the theory that we use, our analytical calculations are merely plausible. We do not attempt to predict analytically the functional dependence of p_c on the various parameters of the grating. Instead, we assume (and justify) a relationship of the form

$$p_c = p_c(SSNR) \qquad (6.1)$$

where SSNR is a plausible signal-to-speckle-noise ratio expression that involves all the grating parameters.

We derive a reasonable expression for SSNR and then rely on experiment to give the desired functional relationship of Equation (6.1). We then use this functional relationship to predict p_c given a set of grating parameters and vice versa.

6.2.2 The line pair gratings

There are various parameters that could influence the detectability of a line-pair grating. All the gratings we consider are rectangular in shape and consist of lines running vertically, since we felt that in a real SAR

picture, if the user were looking for some feature resembling a grating, he would be free to turn the picture around so that the lines, if present, would be vertical. Two of the picture parameters are the height, V, expressed in number of pixels in the (vertical) length of each line, and width, D, expressed in number of pixels in the (horizontal) width of each line. Another parameter possibly affecting the detectability is the number of pixels, H, in the width of the grating in the horizontal direction. Thus there are HV pixels in the grating, and H/2D pairs of lines or cycles. H, V, D may be called the geometrical parameters of the grating.

We consider line pairs with a square-wave distribution of intensity, where the average intensity of the bright lines is $(1+b)*P_0$ and that of the dark lines is P_0. Thus, in the absence of speckle we would have the bright lines uniform in intensity and (1+b) times brighter than the dark lines, which are also uniform in intensity across their widths. We assume that each pixel in the grating is obtained by averaging L independent estimates of the pixel intensity, that is, there are L looks per pixel in the picture. L and b are what may be called the system parameters of the grating.

Thus our grating parameters are b, L, H, V, D.

6.2.3 Visual detection mechanism

There are two distinct mechanisms which the observer might use to make his decisions about whether or not the signal is present. In one method, he makes his decision by looking at where he thinks the bright lines in the picture ought to be and determining whether all or most of them appear brighter than all or most of the regions where the dark lines ought to be. This may be called a space-domain approach.

The other mechanism, which is the one we assume in this chapter, involves a comprehensive look at the whole picture, and may be thought of in terms of spatial Fourier transforms. We assume that the grating is passed through a filter in the visual system of the observer and the amplitude of the output envelope observed. If it lies above a certain threshold, then the observer decides that the signal is present; otherwise he decides it is absent.

On considering the way in which one examines a high-frequency grating, it seems reasonable that the first mechanism, if it operates at all, does so when only a few lines are present, or being looked for, in the visual field. This is because, when there are a large number of lines of reasonably high contrast in the grating, one is aware of the presence of the signal even without looking at every line. It could also be that, whatever the mechanism actually used, either of the above two methods would be useful in predicting the approximate probability of correct decision

for a given case, and our experiments would be unable to choose between the two mechanisms. Because of this uncertainty of approach, we decided, at least as a beginning, just to try to obtain a definition of a signal-to-speckle-noise ratio that would correlate well with the experimental probabilities of correct decision, p_c. In other words, we have used the SNR approach rather than the probability approach of Chapter 3.

6.2.4 Evidence for assumptions

(a)There is evidence, neurological as well as psychophysical, that there exist, in the visual nervous system, several visual channels, each tuned to a different limited range of spatial frequencies. Campbell and Robson, in an investigation of gratings of different waveforms, were the first to propose this [43]. Other findings since then support their hypothesis. According to Stromeyer and Julesz [44], "the idea that the visual system decomposes a visual scene into a set of sinusoidal gratings of specific spatial frequency, orientation, phase positions and contrast is contrary to our introspection.....Nevertheless, this idea is consistent with several studies using simple gratings as stimuli." It has been demonstrated by them that vertical sinusoidal gratings are obscured or "masked" by noise whose spatial frequencies are close to the grating frequency. Their results show that only the noise within +1 or -1 octave of the grating frequency affects the detection

of the grating. They also lend support to the theory that there exist channels in the visual system that are selectively tuned to different spatial frequencies. In our experiments, we asked the observer to look at a picture which consisted of a speckled version of either a uniform background or a rectangular grating and decide between the two alternatives. Thus the observer saw either plain speckle noise or a noisy rectangular grating.

Because of the results in masking just stated, we can assume that the observer's visual system acts as a series of filters of various center frequencies f_0 and BW equal to $(2f_0-0.5f_0)=1.5f_0$. We assume that if the signal (the grating) is present, the observer's visual channel that comes into play is the one that is tuned to a frequency equal to the fundamental of the grating frequency f_g, so that the filter BW is $1.5f_g$. Even when there is no signal present in the picture, we assume that since the observer is told what frequency of grating he is looking for, the visual channel he uses is the one tuned to f_g. Thus, in either case, the light reflected by the picture is passed through a filter of central frequency f_g and BW $1.5f_g$.

(b)To justify the assumption that a filter output envelope is examined, we note first that this is what an ideal detector would do if detecting a sinusoid of known frequency and unknown phase in additive noise [21] and extend the decision procedure, though not the resulting curves of p_c

versus SNR, to the present case of speckle. (We cannot extend these curves to the present case because speckle is not additive noise and because the visual filter is not narrowband, as it is for the ideal detector.) Then we note that it is common practice in interpreting psychophysical results in audio or visual experiments to assume that the human observer is only somewhat worse than an ideal detector. This assumption is often correct (see Green and Swets [28]). Here we merely assume the same detection process as an ideal detector, and not the same performance.

6.3 Signal-To-Speckle-Noise Ratio (SSNR)

6.3.1 Overview

In this section, we show how to obtain one plausible definition of SSNR. Going by this SSNR, we are led to the conclusion that line-pair resolution is not a very good criterion for characterizing speckle. We shall elaborate upon this later in the chapter. Because of this conclusion, as well as because of the difficulties of analysis mentioned in section (6.2), we did not pursue the attempt to predict the value of p_c given the grating parameters.

The SSNR expression obtained contains one parameter, which we call 'a', that lies somewhere between 0 and 1. We decided to investigate two possible values of 'a', and the two corresponding expressions for SSNR are referred to by subscripts 1 and 2.

We first consider the case where the grating has a width H equal to the correlation length of the filter output (which is approximately 1/BW =1/1.5f_g, where f_g is the center frequency of the grating that the observer is looking for) and a height V equal to 1 pixel. We denote the corresponding SSNR expression by $SSNR_p$. We then extend the results to the case of larger V and H.

6.3.2 $SSNR_p$

Using the speckle statistics described in Chapter 3,we can represent each pixel intensity in the grating by a Gaussian random variable with mean values P_0 and $P_0(1+b)$ for the dark and bright pixels respectively. We will also define two new quantities related to P_0 and b by

A_b = the 'bias' intensity

$$= P_0 *(1+b/2) \qquad (6.2)$$

A_d =the 'signal' intensity

$$= P_0 *b/2 \qquad (6.3)$$

In this analysis, each pixel intensity is written as the sum of a signal component of intensity equal to the intensity it would have in the absence of speckle, plus a noise component that accounts for the departure of the actual observed value from this ideal noiseless value. We then derive the SSNR expression by considering that the input intensity spatial

distribution goes through the observer's visual filter of bandwidth $1.5f_g$ extending from $0.5f_g$ to $2f_g$. To obtain this expression, we first determine the signal and noise powers at the filter output.

(a) Noise power output: The total intensity of each pixel (which includes both noise component and signal component, if any) can be written as

$$i(x)=g*s(x) \tag{6.4}$$

where s(x) is the spatial distribution of intensity in the absence of speckle and g is a Grv with mean 1 and variance 1/L.

If g is rewritten as

$$g=1+r \tag{6.5}$$

then r is a Grv with mean 0 and variance 1/L. Also s(x) in the present case is

$$s(x)=[A_b+A_d\,\mathrm{squ}(x/x_0)] \tag{6.6}$$

where we define

$$\mathrm{squ}(a)=\begin{cases} 1 & \text{if the largest integer less than a is odd;} \\ -1 & \text{if the largest integer less than a is even} \end{cases} \tag{6.7}$$

$$x_0 = 1/2f_g \qquad (6.8)$$

Thus, we get

$$i(x) = (1+r)[A_b + A_d \, squ(x/x_0)] \qquad (6.9)$$

The amplitude V_i of the input to the filter is proportional to i(x); for simplicity, we take the proportionality constant to be 1. Therefore

$$V_i = (1+r)[A_b + A_d * squ(x/x_0)] \qquad (6.10a)$$

This can be written as the sum of a signal component V_s and a noise component V_n, i.e.

$$V_n(x) = (1+r)A_b + r * A_d * squ(x/x_0) \qquad (6.10b)$$

$$V_s(x) = A_d * squ(x/x_0) \qquad (6.10c)$$

Because of the spatially varying term $r * A_d * squ(x/x_0)$ in $V_n(x)$, the noise has a mean square value which is a function of position x, and hence is not stationary. However, a spectral density function may be defined even for a non-stationary process, as the Fourier transform of the time- (or, in this case, space-) autocorrelation function (acf) of the process [45]. We may further take the the ensemble-average of the time- (space-) acf, if that is necessary to make the result meaningful [46]. Alternatively, if the noise falls into a category of random process called "broad-sense stationary" [47], the space-acf may be averaged

over a spatial cycle to get a meaningful acf. $V_n(x)$ does fall into this category if we can approximate the square-wave by one of infinite extent.

Once we have thus defined the spectral density of the noise, the results on linear filtering of noise also become applicable here. Thus, we can show in the same manner as in the conventional stationary noise case [21], that the power spectral density $S_0(f)$ at the filter output is given by

$$S_0(f)=S_i(f)*|H(f)|^2 \qquad (6.11)$$

where H(f) is the filter transfer function.

We now apply these results to the case at hand. In Appendix (6.1) we show that the acf of $V_n(x)$ is given by

$$C(y)=\begin{cases} A_b^2+((d-y)/d)*(A_b^2+A_d^2)/L & \text{for } y<d \\ A_b^2 & \text{otherwise} \end{cases} \qquad (6.12)$$

where d is the thickness of each pixel in the horizontal or x direction.

The power spectral density of $V_n(x)$ is then

$$P_1(f) = A_b^2\,\delta(f) + \frac{(A_b^2 + A_d^2)}{L}\, d\left[\frac{\sin(\pi fd)}{(\pi fd)}\right]^2 \qquad (6.13)$$

Now if this noise is passed through the visual filter, we

can use Equation (6.11) to obtain the output noise spectral density. The visual filter has a center frequency f_g and a flat passband from $0.5f_g$ to $2f_g$. We can see that with D pixels per width of each line and a thickness d for each pixel,

$$f_g = 1/(2Dd) \qquad (6.14)$$

The filter passes unchanged that portion of the spectrum $P_i(f)$ that lies between 0.5/(2Dd) and 1/(Dd), so that the total output noise power is

$$P_{On} = (A_b^2 + A_d^2)/L \int_{0.5/(2Dd)}^{1/(Dd)} d \left[\frac{\sin(\pi fd)}{(\pi fd)}\right]^2 df \qquad (6.15)$$

To get a simple approximate estimate, we replace the $(\sin(.)/(.))^2$ function in Equation (6.13) by a rectangular function extending up to the first zero of the original function, i.e., up to f1=1/d, and of the same area as the$(\sin(.)/(.))^2$ function. We find that the height of this rectangle should be 1/2 the peak height of the$(\sin(.)/(.))^2$ function. Using this approximation, we get, for P_{On},

$$P_{On} = \frac{(A_b^2 + A_d^2)}{2L} * (1.5/2D) \qquad (6.16)$$

This is the desired expression.

(b) Signal power output: The input signal is

$$V_s(x) = A_d * \mathrm{squ}(x/x_0) \qquad (6.10c)$$

We can see that, irrespective of the value of D, only the fundamental frequency component of the square-wave grating lies in the bandwidth $0.5f_g$-$2f_g$ that is passed by the visual filter. Since the amplitude of the fundamental is $4/\pi * A_d$, the power output from the filter, after averaging over a few spatial periods, is

$$P_s = 1/2 * (4/\pi A_d)^2 \qquad (6.17)$$

where the factor 1/2 comes from averaging the sinusoidal fundamental. (c) $SSNR_p$: We can now write an expression for the signal-to-speckle-noise ratio $(SSNR)_p$ for the case where V=1 and there is only one horizontal correlation interval available for observing the output. We get

$$SSNR_p = \frac{P_s}{P_{0n}} = 1/2 \quad (4/\pi)^2 \frac{A_d^2}{(A_b^2 + A_d^2)} \; 2L \; \left(\frac{2D}{1.5}\right) \qquad (6.18)$$

Note that the $SSNR_p$ is proportional to the product D*L of the number of pixels per line and the number of looks per pixel. This is in accord with the averaging or blurring property of the eye-brain combination which we assumed in Chapter 5, which enables the eye to consider the total intensity of pixels in the line width as one unit. Thus, there are, equivalently, M=L*D looks per line width. In Chapter 5, this results in an improvement of signal-to-noise ratio by a factor of $\sqrt{M}$. Here, the eye has been considered to respond to the square of the input light intensity since we assume a square-law detector, so that the SSNR

improvement is by a factor of M.

6.3.3 SSNR for general H*V

In order to extend the result to case where V>1 and more than one correlation interval is present in the horizontal direction, i.e. H>2D/1.5, we assume that the observer's eye-brain system considers a portion of the output envelope of length equal to one correlation interval (i.e. 2D/1.5 pixels) in the horizontal direction and one pixel length in the vertical direction (V=1) as the basic unit, and, from each such unit available, forms an independent estimate of the output amplitude.

To relate the horizontal extent, i.e. H pixels, of the grating to the number of independent estimates available from the horizontal direction, we consider the filtering operation in the Fourier transform domain. Then we can see that the output extends for one correlation interval if H<2D/1.5 (we will ignore this unrealistic case in which there is not even one whole cycle in the grating); else it extends for a length equal to the length of H pixels. Thus there are approximately H/(2D/1.5) correlation intervals in the filter output and hence this many independent estimates of output envelope amplitude available from the horizontal direction. To account for the approximation, we insert a factor a_h and get a_h*H/(2D/1.5) for the number of estimates from the horizontal direction. Further, if there are V vertical pixels, then there are a_v*V independent estimates

possible from the vertical direction for each horizontal correlation length. The factor a_v accounts for the possibility that the eye looks at more than one vertical pixel in making one estimate. Thus there are $a_v*a_h*HV(1.5/2D)$ independent estimates obtained.

The precise value of the factor a_v*a_h is not important provided we assume that it is the same for all H and V. This is equivalent to assuming that the observer keeps his decision criteria constant from picture to picture, which is also implicit in the assumption, made in section 6.2, that the observer is only somewhat worse than an ideal detector. Then the factor a_v*a_h merely changes the SSNR expression by a constant factor and since our aim is to find, experimentally, the functional relationship between p_c and SSNR, we may set this factor equal to 1. Specifically, in section (6.4), we find that

$$p_c = 1-1/2*\exp(-k*SSNR) \qquad (6.19)$$

where k is obtained by a least-squares technique. If a_v*a_h were not set equal to 1, k would change accordingly.

We now have that the observer's brain forms $N=V(H/(2D/1.5))$ independent estimates of the presence of the signal. It can combine these estimates in various ways to reach a decision. If an ideal envelope detector is used and post-detection integration performed [21], then the use of N

independent estimates leads to an improvement in SSNR by a factor of $N^{1/2}$ over that achievable with a single estimate. With perfect pre-detection integration, the improvement factor is N. In general, we may assume that the improvement factor will be N^a, where a lies between 1/2 and 1, if some sort of imperfect pre-detection integration is assumed. If some imperfect post-detection integration is assumed, a can be even less than 1/2. We investigated two values of a, namely 1/2 and 3/4, and got two SSNR expressions, which we will call $SSNR_1$ and $SSNR_2$ respectively. We get, for a grating of H*V pixels,

$$SSNR_1 = SSNR_p \left[HV\ (1.5/2D)\right]^{1/2} \qquad (6.20a)$$

so that

$$SSNR_1 = \frac{16}{\pi^2}\frac{A_d^2}{A_d^2+A_b^2} L \left[\frac{2D}{1.5}\right]\left[HV\frac{1.5}{2D}\right]^{1/2} \qquad (6.20b)$$

or, in other words,

$$SSNR_1 = \frac{16}{\pi^2}\frac{(b/2)^2}{(1+b/2)^2+(b/2)^2} L \left[HV\frac{2D}{1.5}\right]^{1/2} \qquad (6.20c)$$

and, for $SSNR_2$, we get

$$SSNR_2 = \frac{16}{\pi^2}\frac{A_d^2}{A_d^2+A_b^2} L \left[\frac{2D}{1.5}\right]\left[HV\frac{1.5}{2D}\right]^{3/4} \qquad (6.21a)$$

or,

$$SSNR_2 = \frac{16}{\pi^2} \frac{(b/2)^2}{(1+b/2)^2+(b/2)^2} L(HV)^{3/4} \left[\frac{2D}{1.5}\right]^{1/4} \qquad (6.21b)$$

Thus, we have found two plausible expressions for SSNR. Although these expressions are logical only for the case of large number of looks per pixel, L, we have assumed them to be the correct expressions even for small L, in analyzing our experimental results.

6.4 Simulations and experiments

6.4.1 The procedure

In order to test the usefulness of the SSNR expressions we obtained, we performed computer simulations of gratings with various sets of parameters. Each picture obtained, which we will also refer to as a sample, had a size of 200*200, 256*256 or 300*300 pixels and consisted of a matrix of smaller rectangular-shaped figures which we shall refer to as "patterns". Each of these patterns was of size H*V pixels and represented a speckle-corrupted version of either "noise" or "signal-plus-noise". The patterns representing "noise" were obtained by taking a V*H pattern of uniform intensity $P = P_0[(1+b/2)^2+(b/2)^2]^{1/2}$, and replacing the intensity of these V*H pixels by V*H independent, gamma-distributed random variables of mean value P and variance P^2/L that represented L-look, speckle-corrupted pixels of average intensity P. The patterns representing signal-plus-noise were similarly obtained by starting with a figure of size V*H pixels that had vertical line-pairs, with each

line having a thickness of D pixels and with each pixel having an intensity of either P_0 or $P_0(1+b)$ depending on whether it belonged to a dark or a bright line. Each picture shown to the observers consisted of a rectangular matrix of k*p patterns. Each pattern in the picture was randomly chosen to be, with equal probability, either a noise or a signal-plus-noise pattern. The parameters characterizing each pattern, namely the contrast ratio (1+b), the number of looks per pixel, L and the values of V, H and D, were the same for each pattern in a given picture, so that the only non-random difference between any two patterns was the presence or absence of a signal, i.e. a line-pair grating. The number of patterns, k*p, varied from picture to picture depending on V and H and the total number of pixels in the picture.

Whenever a line-pair grating was present in a pattern, it had a random phase, i.e. the phase of the square-wave cycle represented in the leftmost pixel of the pattern was randomly chosen for each pattern in the matrix. This phase-randomizing ensures that an envelope detector is indeed the ideal detector [21].

The experimental procedure was to show pictures obtained with various sets of the 5 parameters, b, L, H, V, D, to an observer, who was given certain directions and asked to decide whether or not a line-pair grating was present in each of the patterns. The information to be made

known to the observer was chosen as explained in Chapter 4; specifically, (a) he knew that the signal and no-signal patterns occurred with equal probability; he was also told: (b) the number of cycles to expect in each pattern if a signal were to be present; (c) the contrast ratio (1+b); (d) the number of looks per pixel, L.

The effect of revealing fact (b) above is that the observer knows what center frequency of signal to look for. Thus, the assumption of known center frequency of signal expected, that we used in subsection 6.2.4, is satisfied.

The other 3 pieces of information given to the observer let him know what contrast and clarity to look for in making his decision. His knowledge of these pieces of information makes it reasonable to assume that he can calculate the probability density functions of signal-plus-noise and noise and use some decision strategy such as the Bayes strategy [21] to make his decision. This kind of assumption, that the observer's eye-brain combination seems to make such probability calculations, has often been made and shown feasible in psychophysical experiments [28].

Figure (6.1) shows the geometry of patterns in a sample, and Figure (6.2) some examples of the speckled pictures used, as well as how some of them would look in the absence of speckle. We will now describe these pictures.

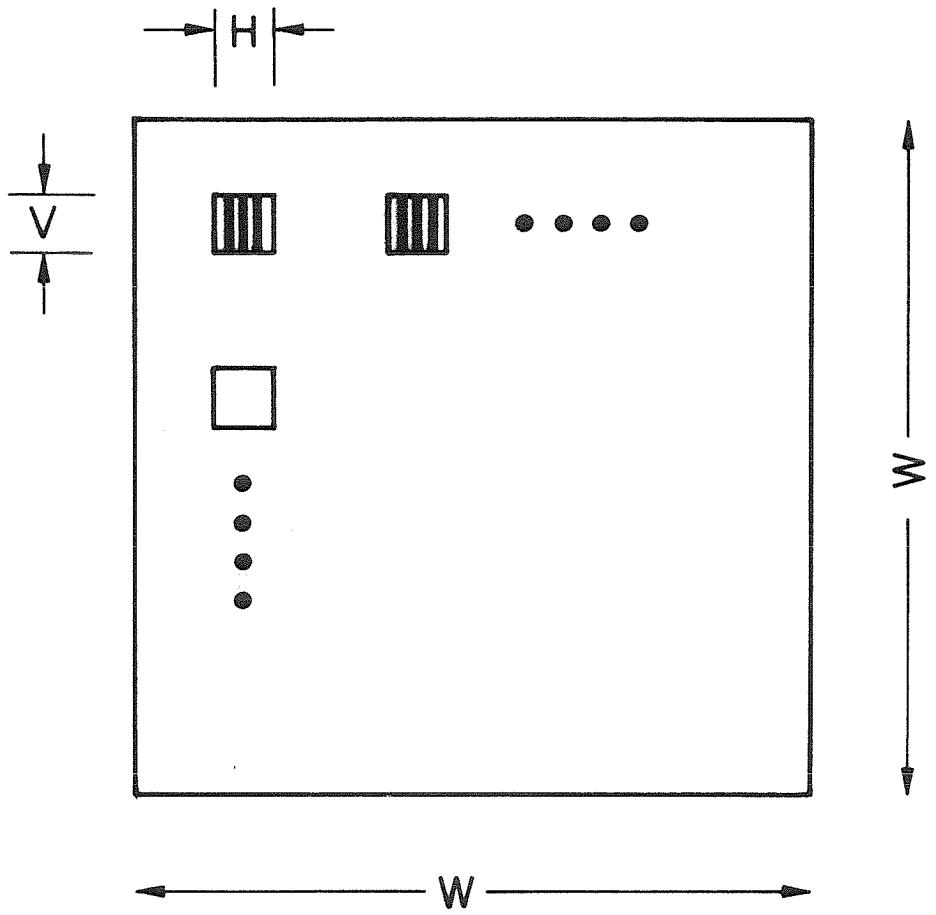

Figure (6.1): Geometry of patterns in a picture
in line-pair detection experiments

(1)

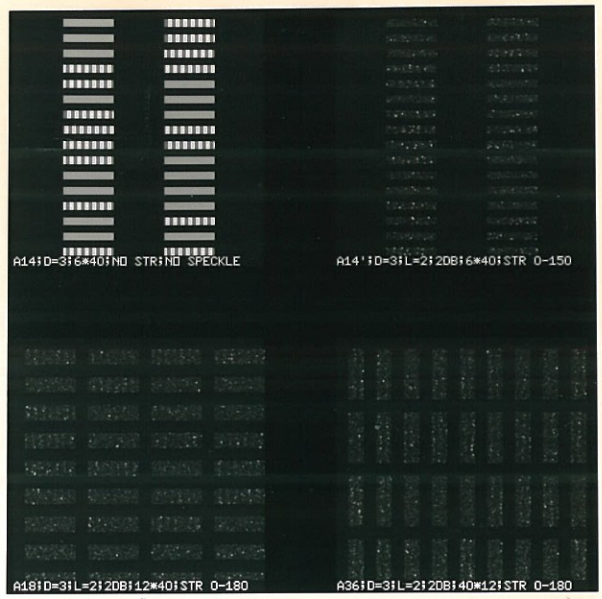

(2)

Figure (6.2): Line-pair detection; examples of simulated pictures

6.4.2 Examples of simulated pictures

6.4.2.1 Overview

As in Chapter 5, we show here only a few of the simulated pictures used in our experiments. Here, we show only 8 samples as opposed to the 16 shown in Chapter 5. However, each of these 8 samples contained many patterns, whereas each sample in Chapter 5 could show only 4 features because of the need there to show a large background for each feature. All the samples except L1 are 200*200 pixels, except that the last 6 lines are replaced by the caption. (We specify the size of the samples, for completeness, although it is irrelevant to the detection problem.) The sample called L1 has 100*100 pixels. The captions specify the line thickness D in pixels, the contrast ratio in dB, the number of looks per pixel L and the size of each pattern V*H in pixels. However, in A16 and L1, there isn't enough room to specify all we need to, and we will supply the missing information in the text.

6.4.2.2 The samples

(1) The first photograph shows A14 without speckle, A14 with speckle, A18 and A36. All these are 2 dB samples with D=3 and A18 and A36 have L=2. The first sample shows that the lines are clearly detectable in the absence of speckle. The second one shows the degradation caused by speckle.

A comparison of A14 and A18 shows the improvement in

detectability of lines as V is increased. A36 has V*H=40*12, while A18 has V*H=12*40. It is not really possible to see the effect of interchanging V and H by examining one example of such a pair of samples, unless many other samples are included. But, at any rate, comparing A18 and A36 gives an idea of what this interchange does. In addition, it is clear that A36 has lines that are more detectable than A14, which shows that increasing V from 6 to 40 more than compensates for reducing H from 40 to 12.

(2) The second photograph shows L1, which is a 100*100 sample, and A47, A45, A44, each of which is 200*200. L1 has parameters D=2, L=12, 1 dB contrast ratio, V*H=100*100, and is intended to show how clearly the lines can be detected in such a picture. We contrast this with the poor detectability of features in S30 #4 of Chap. 5, which has the same L and contrast ratio, and a much larger D. This illustrates that the detectability of small features is a much more stringent criterion than detectability of line pairs.

A47, A45 and A44 are all 1 look, 3 dB pictures with D=2. They show that detectability of lines increases as V*H increases from 8*8 through 16*16 to 64*64. In the sample A44, the probability of correct decision is practically 1.

6.4.3 Results of experiments

Table (6.1) shows the results of the experiments for the various values of parameters used. The table also shows, for each set of parameters, the values of SSNR as given by the two expressions of Equations (6.20)-(6.21). We attempt to see if one of these two expressions is much better than the other in predicting the observed experimental results. The various columns in Table (6.1) are

(a) the parameters H, V, D, b, L

(b) the probability of correct decision p_c obtained experimentally as the percentage of patterns correctly decided upon by the observer

(c) $SSNR_1$ calculated from Equation (6.20)

(d) $SSNR_2$ calculated from Equation (6.21)

(e) the number of patterns, n, on which the experimental p_c is based.

6.4.4 Analysis of the results

We wish to determine whether one or both of the SSNR expressions defined in section 6.3 is a meaningful expression that helps us obtain the probability of correct decision p_c for line gratings with a given set of parameters. In order to do this, we plotted, corresponding to each of the SSNR expressions, the value of p_c obtained experimentally versus the SSNR, fitted a regression curve

Table (6.1): Line pair gratings

#	(1+b)	(1+b) dB	L	D	V	H	p_c	$SSNR_1$	$SSNR_2$
1	1.26	1.0	1	2	64	64	0.80	2.213	13.855
2	1.26	1.0	1	2	96	96	0.75	3.320	25.453
3	1.26	1.0	1	1	120	120	0.63	2.934	29.912
4	1.26	1.0	1	3	120	120	0.92	5.082	39.366
5	1.26	1.0	1	3	100	100	0.89	4.235	29.947
6	1.26	1.0	1	3	85	85	0.94	3.600	23.468
7	1.26	1.0	1	3	64	64	0.64	2.711	15.333
8	1.26	1.0	1	4	12	40	0.69	1.071	3.300
9	1.26	1.0	1	4	24	40	0.75	1.515	5.550
10	1.26	1.0	1	4	24	80	0.83	2.143	9.334
11	1.26	1.0	1	4	80	24	0.81	2.143	9.334
12	1.26	1.0	1	4	60	60	0.61	2.934	14.956
13	1.26	1.0	12	2	4	4	0.66	1.660	1.396
14	1.26	1.0	12	2	4	8	0.60	2.347	2.347
15	1.26	1.0	12	2	8	4	0.77	2.347	2.347
16	1.26	1.0	12	2	8	8	0.76	3.320	3.948
17	1.60	2.0	3	2	12	8	0.72	3.934	7.323
18	1.60	2.0	3	2	8	8	0.68	3.212	5.403
19	1.60	2.0	3	2	6	8	0.70	2.782	4.354
20	1.60	2.0	3	2	8	6	0.99	2.782	4.354
21	1.60	2.0	2	3	40	12	0.93	7.183	19.992
22	1.60	2.0	2	3	12	40	0.81	7.183	19.992
23	1.60	2.0	2	3	120	6	0.73	8.798	27.098

Table (6.1) (continued)

#	(1+b)	(1+b) dB	L	D	V	H	ρ_c	$SSNR_1$	$SSNR_2$
24	1.60	2.0	2	3	40	6	0.78	5.079	11.887
25	1.60	2.0	2	3	6	40	0.74	5.079	11.887
26	1.60	2.0	2	3	25	6	0.81	4.016	8.356
27	1.60	2.0	2	3	6	24	0.59	3.934	8.104
28	2.00	3.0	1	1	8	4	0.70	1.059	2.344
29	2.00	3.0	1	1	8	8	0.86	1.498	3.942
30	2.00	3.0	1	1	16	16	0.64	2.995	11.149
31	2.00	3.0	1	2	8	8	0.58	2.118	4.688
32	2.00	3.0	1	2	4	4	0.86	1.059	1.657
33	2.00	3.0	1	2	16	16	0.99	4.236	13.258
34	2.00	3.0	1	2	64	64	1.00	16.943	106.068
35	3.16	5.0	2	3	40	6	1.00	21.332	49.925
36	3.16	5.0	2	3	40	12	1.00	30.000	83.964

and calculated the index of correlation corresponding to this curve. The details of the methods for these statistical calculations are standard in statistical theory [29], [48], [49].

The experimental data are given in Figures 6.3(a) and 6.3(b). Initially, we tried fitting a straight line to the observation points, but an examination of the points and the line showed that a straight line is not the right curve to fit the data. For large values of SSNR, p_c should approach 1 asymptotically and not exceed it, as a straight line would predict. Many of our data points have p_c close to 1, so that even for the range of values of the data points, a straight line is inadequate. A curve that has the proper characteristics is one that relates p_c linearly to the negative of the exponential of SSNR.

For the exponential regression curve, we specified that the point SSNR=0, p_c=1/2 lie on the curve, because, for SSNR=0, no signal is present, but the observer is expecting a signal half the time, so that the percentage correct should be 50% for an unbiased observer. We assumed a curve given by

$$\hat{p}_c = 1 - 1/2 \exp[(-\lambda (SSNR)] + e \qquad (6.22)$$

and determined λ so as to minimize the sum of squares of the errors represented by e in the above equation. In this chapter, $\hat{p}_c$ is used to denote the value predicted by the

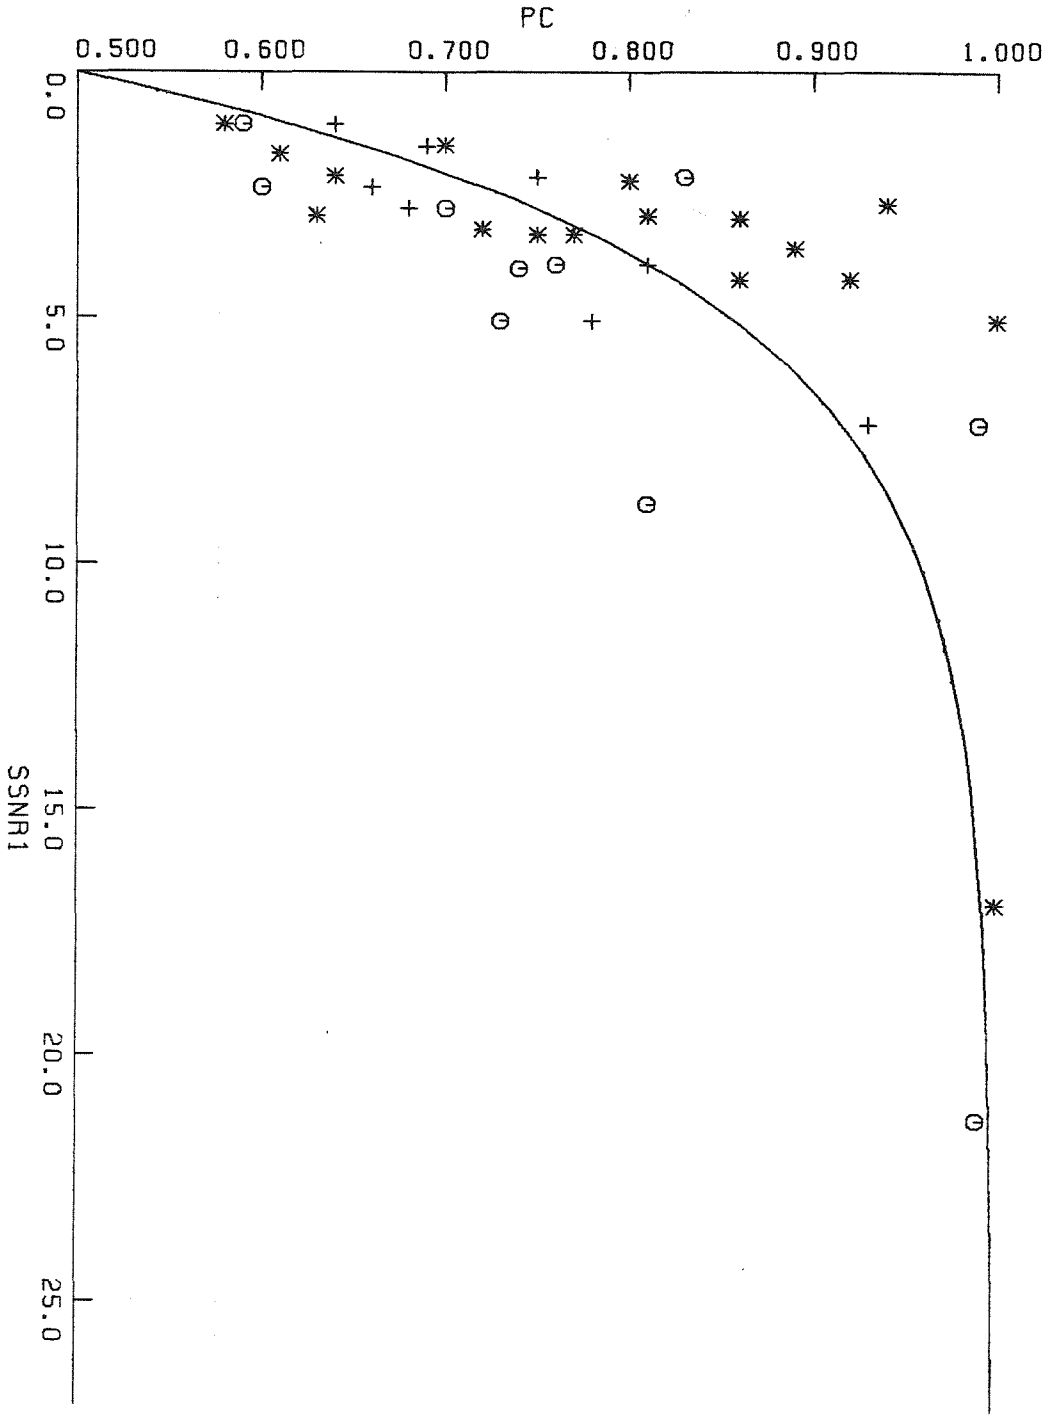

Figure 6.3(a): Data points and regression curves; experimental p_c versus $SSNR_1$
o data points with V>H
+ data points with V<H
* data points with V=H

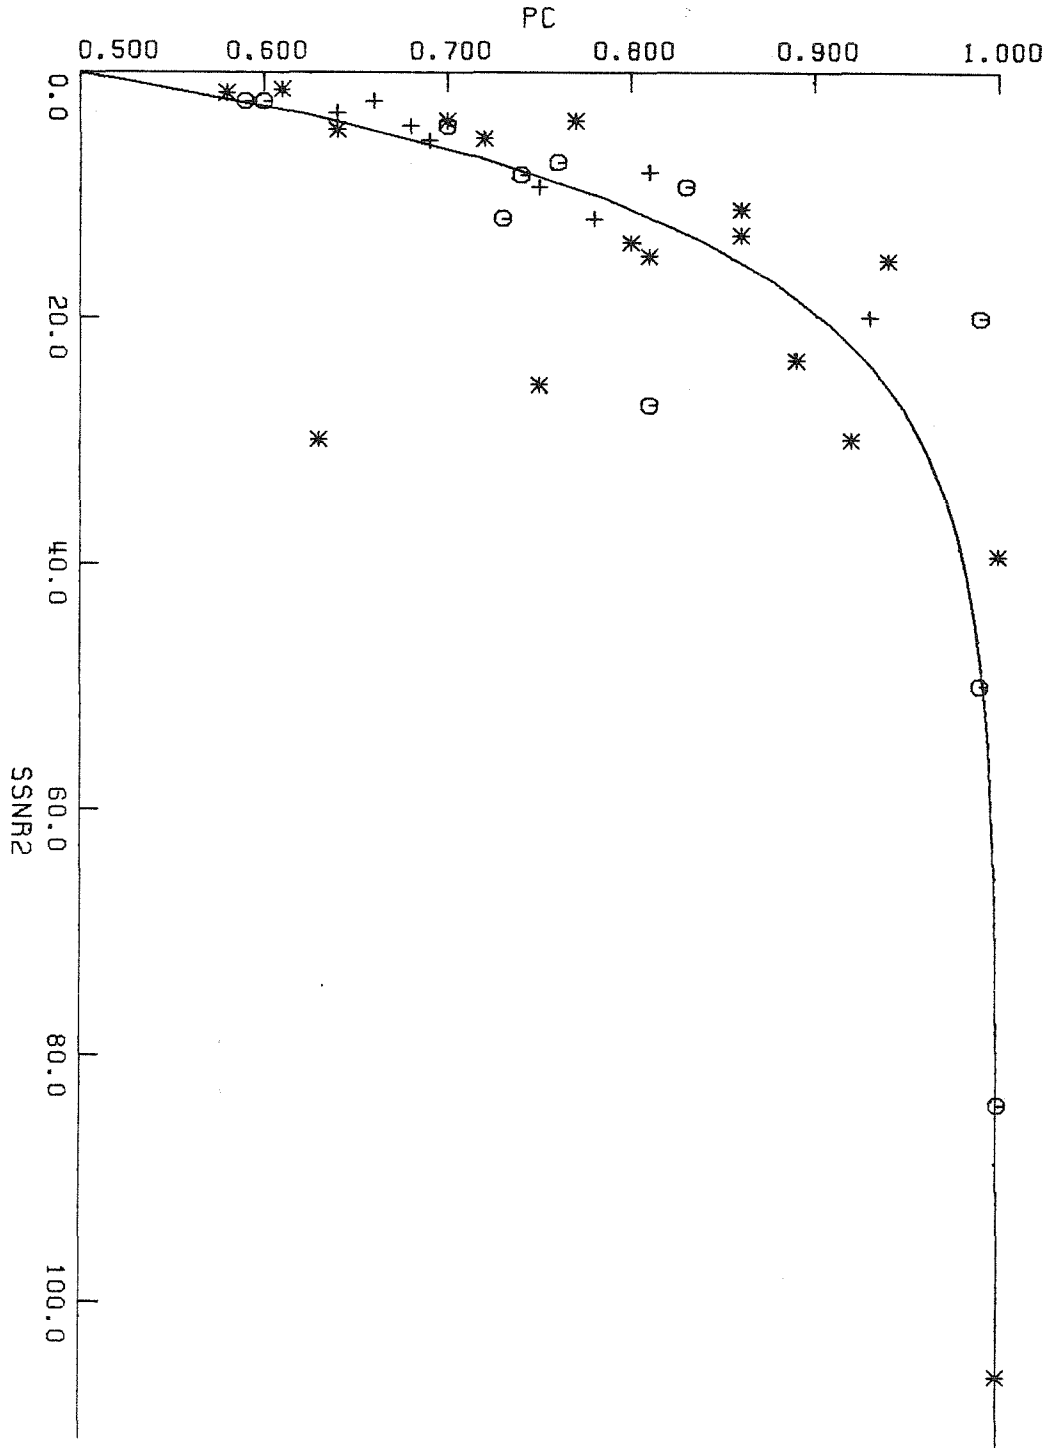

Figure 6.3(b): Data points and regression curves; experimental p_c versus $SSNR_2$
o data points with V>H
+ data points with V<H
* data points with V=H

regression curve. The equation for λ that results has to be solved iteratively, and this was done by computer. The resulting values of λ and the two regressions curves for the two assumed expressions for SSNR of Equations (6.20)-(6.21) are given below with subscripts 1 and 2 corresponding to the two SSNR expressions.

$$\lambda_1 = 0.24818 \quad (6.23a)$$

$$\hat{p_c} = 1 - 1/2 \exp[-0.24818 \ (SSNR)_1] \quad (6.23b)$$

$$\lambda_2 = 0.08264 \quad (6.24a)$$

$$\hat{p_c} = 1 - 1/2 \exp[-0.08264 \ (SSNR)_2] \quad (6.24b)$$

The resulting regression curves are shown in Figures 6.3(a) and 6.3(b) along with the data points.

6.4.4.1 Comparing $(SSNR)_1$ and $(SSNR)_2$

Since Equations (6.23b) and (6.24b) are nonlinear in the parameter λ they fall in the category of nonlinear regression equations. Hence, standard goodness-of-fit tests that can be used for linear regression are not strictly valid here. However, following [29], we may still use these standard tests to get some idea of goodness of fit of the two curves. The test we used is applicable only if there are some repeated observations for the same value of the independent variable and if the errors corresponding to e in Equation (6.22) are normally distributed. This normality

assumption was tested using the Kolmogorov-Smirnov test [48] and found to be valid. We do have repeat observations in our experimental results, but only if the (SSNR) expressions of Equations (6.20) and (6.21) are assumed to be correct in fitting Equations (6.23b) and (6.24b) respectively, to the corresponding data. That is, we have observations for 2 or more combinations of b, D, L, H, V that give the same SSNR value.

We find that, to the extent that the test can be considered applicable to the present nonlinear case, the fit for the regression curve of p_c on $(SSNR)_2$ is better than that for the regression curve of p_c on $(SSNR)_1$. However this does not tell us that the $(SSNR)_2$ expression is better than the $(SSNR)_1$ expression, because the assumed form of the regression curves may not be the correct one. There are, however, two possible comparisons that we might make about the validity of the two SSNR expressions, which we give below.

(a) The index of correlation [49] between two random variables X, Y fitted by a regression curve described by

$$\hat{Y} = f(X) \qquad (6.25)$$

is given by i_c, which is defined as

$$i_c = (1 - \frac{\sigma_z^2}{\sigma_y^2})^{1/2} \qquad (6.26a)$$

where σ_z^2 is the variance in the deviations z_i, which are defined by

$$z_i = Y_i - f(X_i) \qquad (6.27)$$

and σ_y^2 is the variance in the observed Y. The Y_i and X_i in the above equations are the data points, for i=1,2,..N, where N is the number of observations made. We find that this index of correlation, i_c, has the values

$$i_{c1} = 0.781 \qquad (6.28a)$$

$$i_{c2} = 0.767 \qquad (6.28b)$$

where the subscripts 1 and 2 refer to the cases of the two SSNR expressions considered. Since the values of i_c for the two SSNR expressions considered are so close, we cannot infer that either one is more correct than the other. However, we note that both indices of correlation are fairly high, so that either of the SSNR expressions, together with the corresponding regression equation, is useful in predicting the observed results.

(b) We may compare the values of $\sigma_e^2 = SS_e/(df)$ obtained from the "error-sum-of-squares", SS_e, and its number of degrees of freedom (df). Here,

$$SS_e = \sum_{i=1}^{k} \sum_{j=1}^{n_i} (y_{ij} - \bar{y}_i)^2 \qquad (6.29a)$$

$$(df) = \sum_{i=1}^{k} n_i - k \qquad (6.29b)$$

where there are k groups of repeat observations with n_i

observations in the i^{th} group, Y_{ij} is the j^{th} observation in the i^{th} group and $\overline{Y_i}$ is the mean of the i^{th} group. This variance, σ_e^2, arises as a result of pure experimental errors in the observations and any possible errors in the assumed SSNR expression. It is reasonable to assume that the experimental errors are the same whether $SSNR_1$ or $SSNR_2$ is used, since the same basic set of observations is used in both cases. Hence, any difference in SS_e in the two cases must arise out of the difference between the two SSNR expressions.

The values of $SS_e/(df)$ are given, for the cases of the two SSNR expressions, by

$$\sigma_{e_1}^2 = \left[SS_e/(df)\right]_1 = 0.002775 \qquad (6.30a)$$

$$\sigma_{e_2}^2 = \left[SS_e/(df)\right]_2 = 0.006877 \qquad (6.30b)$$

Thus, $SSNR_1$ appears to give less error than $SSNR_2$. But if these variances are compared using the F-test, their ratio is found to be insignificant at the 5% level of significance.

We find, from our two comparisons, that we cannot definitely choose between the two SSNR expressions. However, either of them will predict the experimental results reasonably well. It is also possible to make various other comparisons to try to choose between the two expressions [49], but we did not feel they were necessary in this case, since there are too many other expressions

possible, for instance, just by varying the value of 'a' used in deriving Equation (6.20a).

6.4.4.2 Regression of (SSNR) on p_c

It is also possible to obtain a regression equation of SSNR on p_c, e.g., to estimate a parameter γ, to minimise the variance of e', where

$$\hat{(SSNR)} = -\gamma [\ln 2(1-p_c)] + e' \qquad (6.31)$$

This expression would enable us to calculate (SSNR), given a value of p_c.

In a case where there was perfect correlation between p_c and SSNR, we would find that $\gamma = 1/\lambda$. But we did not perform this regression, because it requires another iterative solution, and it is sufficient for our purposes to be able to find p_c, given SSNR.

6.4.4.3 Interchangeability of H and V

We also used Table 6.1 and Figure 6.3 to check, in two ways, the assumption that a pattern with H_0 horizontal and V_0 vertical pixels is just as correctly recognizable as one with V_0 horizontal and H_0 vertical pixels, other things being equal. This assumption is implicit in the fact that the SSNR expression has H*V occurring in it only as a product H*V.

(a) We calculated the statistical significance of the difference in p_c for all the pairs of observations having

this property (that the values of H and V were interchanged for the two pictures) and found that none of these differences were significant at the 5% level.

(b) We made a visual confirmation of the hypothesis as follows. In Figures 6.3(a) and (b) we have drawn all points corresponding to observations with H>V with one sort of symbol and all points corresponding to V>H with another sort of symbol. Scanning Figures 6.3(a) and (b), we find that points drawn with both kinds of symbols are similarly scattered, which makes it plausible that a long, narrow grating is neither more nor less recognizable than a short, wide one, if the product H*V is the same for both.

6.4.5 Conclusions from the experiments

From the fairly high index of correlation between the p_c and $(SSNR)_1$, and between p_c and $(SSNR)_2$, we conclude that either of the two SSNR expressions given in section (6.3) is a useful figure-of-merit expression in characterizing the detectability of a speckled line-pair grating. The regression curves we have obtained in Equations (6.23)-(6.24) define the relationship between p_c and each of the SSNR expressions within the range of p_c and $(SSNR)_1$ and $(SSNR)_2$ covered in Figure (6.3) and Table (6.1), which is a fairly wide range. They can be used to predict the p_c that can be expected for any given set of parameters within the given range.

6.5 General Conclusions

From the previous sections, it appears reasonable that the parameters b, L, H, V, D characterizing a grating affect the probability p_c, of correctly deciding on whether it is present or not, via the expression

$$(SSNR) = \frac{16}{\pi^2} \frac{(b/2)^2}{(1+b/2)^2+(b/2)^2} \left[\frac{L}{1.5}\right]^{2D\left[\frac{H\ V\ 1.5}{2\ D}\right]^a} \qquad (6.32)$$

where the value of a may be chosen as 1/2 or 3/4 to get pretty good results. We cannot, on the basis of our experiments, choose between these two values of a; perhaps, some other value of a between 1/2 and 1 is even more appropriate.

We can also obtain two other general conclusions:

(i) We see that, at least within the range of parameter values investigated, increasing the height V or the width H of a grating increases the SSNR, and hence the probability of correct decision. It follows, then, that merely specifying the frequency of the grating (which, for a given width of pixel, d, specifies the number of pixels per line, D) will not reveal whether a certain combination of system parameters, b and L, is sufficient to make the grating correctly detectable; the height and width (V and H) of the grating also need to be given. Or, conversely, within limits, any combination of b, L and D will make a grating detectable if present or recognized as noise otherwise,

provided the height and width of the grating are made sufficiently large. This is not the case with the usual pictures that are corrupted by additive white Gaussian noise (AWGN) as in TV systems, where the just-detectable line-pair frequency for a standard-sized picture is sufficient to characterize the system performance.

(ii) One reason for considering line-pair gratings was to determine whether line-pair gratings could be used as test patterns for characterizing the quality of pictures obtained by the imaging system. First of all, as we noted above, the height and width, as well as the frequency of the test pattern, would have to be specified if a line-pair grating were chosen as the test pattern. Assume that these specifications are chosen to have the values that they might typically be expected to have for a SAR picture, considering the kind of image-processor that is normally used for SAR pictures. Thus, a picture of 256*256 pixels might be used as the test picture. We now show that, for typical SAR parameters, a line-pair grating of such a size would be almost certainly detected by a trained observer used to looking for line pair gratings, as in our experiments, and as might be expected of a scientist examining a SAR map.

We consider a picture of 100*100 pixels; a more realistic 256*256 picture ought to be even more detectable, according to th discussion in (i) above. We also choose a grating with 2-pixel wide lines, D=2. From the SSNR

expressions derived, any larger value of D should make the lines even more detectable. We consider that we have L=12, b=0.26 (corresponding to a 1 dB contrast), D=2, H=V=100. For these parameter values, we obtain

$$(SSNR)_1 = 41.5 \qquad (6.33)$$

From the regression curve of Equation (6.23), or by interpolation in Figure 6.3(a), we see that the p_c for this value of $(SSNR)_1$ is >0.999.

If we use $(SSNR)_2$ instead, we find that

$$(SSNR)_2 = 386 \qquad (6.34)$$

and using Equation 6.24(a), we find that p_c is again predicted to be >0.999.

Thus a practised observer should almost certainly be able to see lines with D=2 in a 100*100 picture of 1dB contrast and 12 looks per pixel. This was, in fact, found to be the case, with a simulated picture having these parameters. The sample L1 in Figure (6.2) illustrates this point.

Actually, even a D=1 line grating is visible with p_c>0.95, as we find from both regression curves and experiments. However, it requires a practised observer to demonstrate that p_c>0.95 in this case. We use D=2 so as to be able to demonstrate our point by presenting a sample (L1) in which a casual observer can see the lines.

We now consider the detectability of a small feature in a speckled picture with comparable parameters. From Chapter 5, we find that, for a 100*100 speckled picture with 12 looks per pixel, a square whose actual intensity (in the absence of speckle) is 1 dB above that of its background, of size D*D pixels, would have to have D>7 in order to be detectable with 95% probability of correct decision. An otherwise similar square with D=2 would be practically undetectable. As an illustration, we see that the D=6 features in the sample S30#4 in Figure (5.3) are hard to detect.

We see from the comparison above, that a speckled imaging system that is capable of making a test pattern consisting of a grating with 2-pixel lines clearly resolvable (by which we mean detectable) might be totally unable to resolve a square 2 pixel on a side and having the same contrast ratio as the grating. Since it is more likely that the imaging system will be used for detecting features rather than lines, resolvability of line-pair gratings is a useless way of characterizing the quality of the image. Resolvability of squares, as in Chapter 5, is a much more stringent criterion for testing image quality.

APPENDIX (6.1)

POWER SPECTRAL DENSITY OF NOISE IN CHAPTER 6

Here, we calculate the power spectral density of the random noise process $V_n(x)$, given by

$$V_n(x) = A_b + r[A_b + A_d \, squ(x/x_0)] \qquad (A\text{-}6.1)$$

where

$$squ(a) = \begin{cases} 1 & \text{if the largest integer less than a is odd;} \\ -1 & \text{if the largest integer less than a is even} \end{cases} \qquad (A\text{-}6.2)$$

We assume that the function $V_n(x)$ extends infinitely long in space, although this is not actually true for our problem. We will assume the results derived here to be approximately applicable to our actual finite-extent problem.

We obtain the power spectral density by first calculating the autocorrelation function (acf) of the process. We note that the process considered here is not stationary so that the ensemble-acf obtained by the conventional method of ensemble averaging the product of the function values at a pair of points is not a function of their separation alone, as we require in order to take the Fourier transform. However, for the process considered here, we can use one of two techniques to obtain an acf,

both of which give the same results. The first technique [46] consists of obtaining the spatial acf $C_1(y)$ for a sample function $V_{n1}(x)$, for a specified shift, y, of the whole function, and then ensemble averaging this $C_1(y)$ to get the acf for the whole process. In the second technique [47], we obtain the product of the values of a sample function $V_{n1}(x)$ at a pair of points, x and x+y, find the ensemble average of this product, and then average over all x in the cycle. The cycle here has a width 2d where d is the width of one pixel. We use the first technique here because it is simpler, since, for the process considered here, the spatial acf $C_1(y)$ is the same for all sample functions so there is no need for ensemble averaging. A-

A-6.1.1 The case where $|y|<d$

The expression for acf that we derive in this section is valid only for the case where the extent of the shift, y, of the sample function $V_{n1}(x)$ is less than the width of each line, d.

We show, in Figure A-6.1, two adjacent pulses in the function $V_{n1}(x)$ with amplitudes A_1 and A_2, where

$$A_1 = A_b + r_1 (A_b + A_d) \quad \text{(A-6.2a)}$$

$$A_2 = A_b + r_2 (A_b - A_d) \quad \text{(A-6.2b)}$$

The r_1 and r_2 in the above equations are independent Gaussian random variables, each having a mean 0 and a variance 1/L.

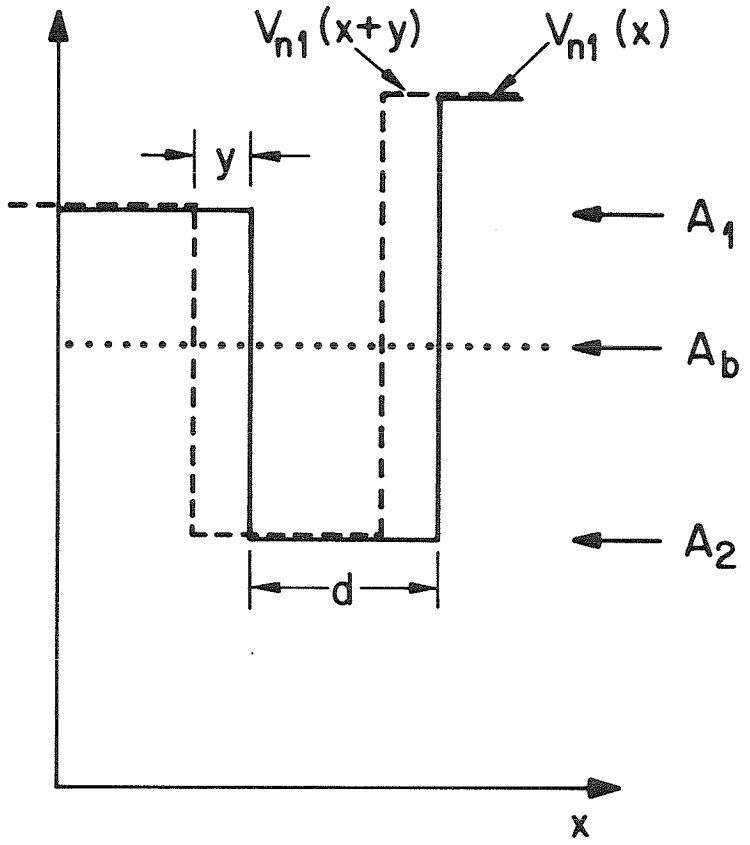

Figure (A-6.1): Two typical pulses in a noise sample function $V_{n1}(x)$ and the shifted function $V_{n1}(x+y)$

We also show the same two pulses shifted by an amount y (in dotted lines). From looking at this figure and remembering that calculating the spatial acf involves multiplying $V_{n1}(x)$ by $V_{n1}(x+y)$ and obtaining the integral of this product, we see that

$$C_1(y) = \lim_{X \to \infty} \frac{1}{2X} \int_{-X}^{X} V_{n1}(x)\ V_{n1}(x+y)\ dx \qquad \text{(A-6.3)}$$

For $|y|<d$, this becomes

$$C_1(y) = \frac{1}{d} E \left[(d-y) \frac{A_1^2}{2} + (d-y) \frac{A_2^2}{2} + yA_1A_2 \right] \qquad \text{(A-6.4)}$$

where E[x] is the expected value of x obtained by ensemble averaging. This ensemble average occurs here, although only one sample function $V_{n1}(x)$ is being considered, because the 'r's in the different pulses of width d forming $V_{n1}(x)$ are all independent. Therefore, in the infinitely long wave $V_{n1}(x)$, there are so many independent values of r, that averaging over all the contributions to each of the three terms in Equation (A-6.4) as is required, is equivalent to finding the ensemble average value for each term. To perform the ensemble average indicated in Equation (A-6.4), we substitute for A_1 and A_2 from Equations (A-6.2a)-(A-6.2b) and use the fact that the 'r's are independent and have mean values of 0 and variances of 1/L and end up with

$$C_1(y) = A_b^2 + \frac{(d-y)}{Ld}(A_b^2 + A_d^2) \qquad \text{for } |y|<d \qquad \text{(A-6.5)}$$

A-6.1.2 The case where $|y|>d$

In the case where the shift y has any magnitude greater than d, we can see that Equation (A-6.3) for $C_1(y)$ reduces to

$$C_1(y) = E[yA_1A_2] \qquad \text{(A-6.6)}$$

This leads to

$$C_1(y) = A_b^2 \qquad \text{for } |y|>d \qquad \text{(A-6.7)}$$

A-6.1.3 The power spectrum

Taking the Fourier transform of $C_1(y)$ as expressed by Equations (A-6.5) and (A-6.7), we get, for the power spectrum of $V_n(x)$,

$$P_1(f) = A_b^2\,\delta(f) + \frac{(A_b^2 + A_d^2)}{L}\, d \left[\frac{\sin(\pi fd)}{(\pi fd)}\right]^2 \qquad \text{(A-6.8)}$$

7. FORM DISCRIMINATION

7.1 Overview

In Chapter 5, i.e., the part on detection of small features, we ignored the internal structure of the features to be detected. While the detectability of small features provides a good measure of one aspect of picture quality, the aim of most pictures is to provide good enough quality so that we can identify the form of a feature. Usually, the identification problem can be thought of as one in which the observer has various possible alternatives in mind - perhaps a large number of them - and he believes the given feature to be a speckled version of one of these forms. In this chapter, we consider some simple cases of this idealized form identification problem. Specifically, we consider three cases of discrimination, i.e., discriminating among two possible forms, two possible orientations of a given form and four possible orientations of a given form. Our aim here is to try to relate the parameters of the possible alternatives to their discriminability.

In order to be able to make any theoretical predictions, we have chosen some simple geometrical forms for which calculations can be made reasonably easily. The forms chosen are such as to have very little complexity of structure, and they are easily identified in the absence of speckle.

Not much is known about the exact perceptual mechanism involved in form discrimination. Therefore, the calculations we make are based on an assumed nearly optimal discriminator. There is some evidence to support this assumption. In any case, we checked our calculations by computer simulation of the pictures for various sets of parameters and psychophysical experiments in which these pictures were shown to observers, and found reasonable agreement with the theory.

7.2 The approach used

7.2.1 Overview

To focus attention on the internal structure of the features of interest, we isolate the features from their speckled backgrounds, i.e., the features are assumed to be set in an otherwise uniform completely dark background, rather than the more reasonable speckled background. This makes possible a simpler mathematical calculation, since the background speckle need not be accounted for. It also means that a large number of small features can be simulated in each picture used in the experiments, since we need only show the observer a small feature through a hole in a dark mask. This allows us to use a large number of samples of each set of parameters in the experiments, contrary to the case for the detection of features discussed in Chapter 5.

We note that, with a speckled background,

discriminability would be worse, as we verified by means of a few simulated pictures. This has, in general, been found to be true in visual form determination experiments in the past [50].

The general approach we use here is the same as in the problem of detection of small features in speckled backgrounds. That is, we make some assumptions about the discrimination process used by the observer, choose a plausible theory to calculate the probability, p_c, of correctly distinguishing between the allowed shapes, then run experiments in which an observer looks at simulated pictures with various parameters and compare the probability of correct decision $\hat{p}_c$ obtained experimentally with the theoretical probability. In the experiments, the observer is told what the possible shapes are and then asked to determine which of the possible shapes the speckled patterns seem closest to.

7.2.2 The forms used

Throughout this chapter, whenever we refer to "form" or "shape", we mean the speckle-free feature from which the speckled "pattern" is obtained.

We considered first the problem of discrimination when there are only two allowed shapes, differing slightly from one another. The problem of discriminating between the two shapes shown in Figure 7.1(a) was considered first. The

theory, which we describe in the next section, predicts that the two shapes in Figure 7.1(b) are just as differentiable from each other as the two shapes in Figure 7.1(a). In order to verify this experimentally, we simulated some pictures that had the two shapes in Figure 7.1(b) as the allowed shapes, although the majority of the pictures simulated here are for the problem of distinguishing between the two shapes in Figure 7.1(a). We will refer to the simulated pictures having the two shapes of Figure 7.1(a) as the Set 1 pictures, and the ones corresponding to Figure 7.1(b) as the Set 2 pictures.

We also considered the problem of distinguishing between two possible orientations of a given shape. The shape and the two orientations considered are shown in Figure 7.2(a). We then extended the theory to the case where four alternative shapes are present and made a few simulations for the four orientations shown in Figure 7.2(b). We will refer to the set of pictures corresponding to the shapes in Figure 7.2(a) as Set 3 and the ones corresponding to Figure 7.2(b) as Set 4.

We will refer to the shapes shown in Figure 7.1(a) as U and O, those in Figure 7.1(b) as O and S, those in Figure 7.2(a) as U and R, and those in Figure 7.2(b) as U,D,L,R.

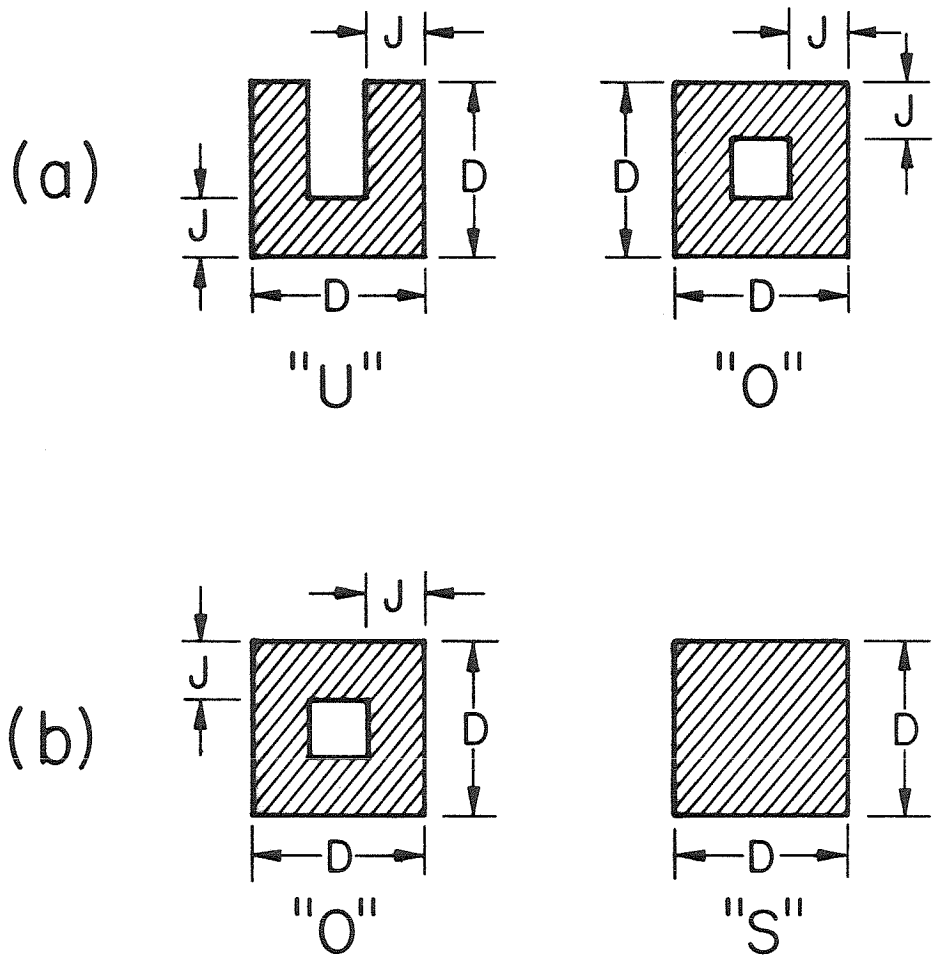

Figure (7.1): Forms used in "2 alternative forms" experiments
(a) Set 1
(b) Set 2

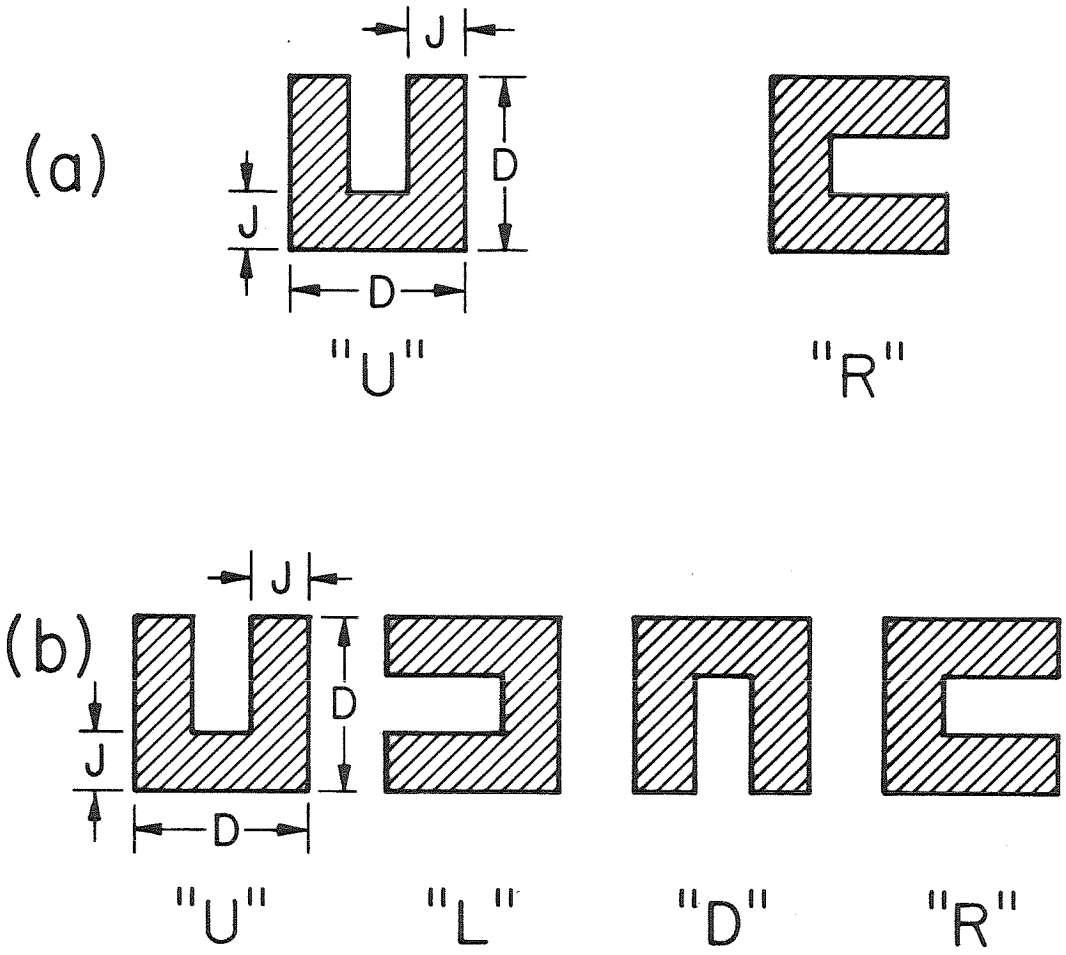

Figure (7.2): Forms used in
(a) 2 orientations experiments (Set 3)
(b) 4 orientations experiments (Set 4)

7.2.3 The parameters

In the absence of speckle, all the shapes described above are obtained from squares of size D pixels * D pixels, by making some portions brighter than the others. In the absence of speckle, all pixels in these brighter portions have intensities P_0 (1+b) while the remaining pixels in the squares have intensities P_0. The brighter portions in all the shapes occur in strips of width J pixels along the edge of the squareof side D and are shown shaded in Figures (7.1) and (7.2). In all cases, D=3J so that the darker portions (if any exist in the shape) are also of thickness J pixels.

All pixels in all the speckled pictures have L looks, i.e., are obtained by summing L independent estimates of the speckle-corrupted pixel intensity. Thus, our parameters are L, b, and J (or D).

7.3 Theoretical calculations

7.3.1 General method

Currently there isn't enough known about the perceptual processes involved in discrimination of forms, so that the use of any visual model and calculations based on it, as in the previous chapter, is not feasible. However, some guidelines do exist.

Previous work on discrimination has shown [51], as one would expect, that two patterns become more easily discriminable from one another as the "difference" between

them increases. This difference may be any kind that is visually perceptible. One reasonable way of objectively characterizing the difference between the various forms we consider is used here. We start by representing the two forms in terms of a few basic spatial components, which we can choose to be orthogonal. One simplified choice of orthogonal basis vectors is the set of 9 squares of size J pixels * J pixels that characterizes each figure. Or, we could increase the complexity greatly but get a more accurate representation by choosing the D^2 pixels in each pattern as D^2 orthogonal basis vectors. We chose the former, simpler approach. The various shapes are then characterized by 9 components, each one representing the total power reflected by the J^2 pixels of one of the 9 squares of size J pixels * J pixels. There is some sacrifice in accuracy in using this representation, because the intensity variation within the individual squares of size J^2 pixels is ignored and only the sum retained. This representation would be acceptable if, for example, we were considering machine recognition and the input to the machine was obtained from a sensor that summed the intensities of the J^2 pixels in each square somewhat like the sensor considered by Dainty [23].

It may or may not be reasonable to consider the human observer's perceptive process as performing the sort of summation described above. But we may try to justify the

use of this model by the following arguments:

(a) If the observer is told, as he is in our experiments, exactly what the forms are, his brain will be examining the 9 squares of size J^2 pixels in the pattern, so it is reasonable that he will mentally divide the pattern into 9 squares of size J^2 pixels and fix his attention on the 9 squares, examining each as one unit. Of course, he is limited in his judgement of how large the squares should be, but since the observer in all our experiments initially saw a high signal-to-noise version of the pattern he would examine, this limitation is not too great. This method of showing the observer a high contrast version of whatever he is looking for, has the same effect as a training period for the observer (Green and Swets [28], Appendix on Experimental Procedures) that makes the assumption of his ability to judge relative sizes reasonable.

(b) Once we assume, as in (a) above, that the observer's attention is focused on a particular square of size J^2 pixels, then the same arguments as in Chapter 5 apply to justify the use of the sum of intensities in the square as a parameter, i.e., either we may assume that his visual system actually performs the spatial summation, or that it responds to some function of the individual pixel intensities in the square of size J^2 pixels that may be approximated by their sum. We also considered another function of the J^2 pixel intensities which assumed that the J^2 pixel intensities are

independent estimates that are combined suboptimally by the visual system, so that they effectively act as if only $(J^2)^{1/2}=J$ such estimates were available. The use of the square root in this suboptimal combination of estimates comes from an analogy with incoherent summation in the case of line pair detection. Because the same detection mechanism is not used here, the assumption may be somewhat unrealistic. However, the calculated probabilities of correct decision p_c based on this suboptimal combination of the estimates was found to be a useful lower bound for the experimental $\hat{p}_c$, while the ones based on the optimal use of the estimates formed an upper bound to the experimental $\hat{p}_c$.

Using these assumptions, we calculated the probability of correct decision given one of the shapes, as the probability of identifying the given speckled pattern as that shape, and then the overall probability of correct decision for equally likely shapes. We now describe the decision mechanisms assumed:

(c) In the case of distinguishing between two shapes of the kind shown in Figure 7.1(a) or Figure 7.1(b), the observer has to look at a particular one of the 9 components which we call C and decide whether, in the given pattern, this component is "bright" or "dark", i.e., is a speckled version of intensity $P_0(1+b)$ or P_0. In Figure 7.1(b), this component C is the central square of the shape; in Figure 7.1(a), it is the central one in the top row. The observer

needs a standard which tells him what "bright" and "dark" should look like. We assume that he derives this standard from the squares that he knows to be dark or bright and from his knowledge of b. In the case of the Figure 7.1(b) shapes, he has some direct idea of what "bright" squares should look like because the 8 outer squares are "bright", and then, from his knowledge of b, he is assumed to be able to judge what "dark" looks like. In the case of the Figure 7.1(a) shapes, he knows what "dark" looks like from the central square and what "bright" looks like from the 7 surrounding squares that are known to be bright. This difference between the cases of Figure 7.1(a) and Figure 7.1(b) might make a small difference to the experimental $\hat{b}_c$ for the two cases. Because our experiments are not sensitive enough to detect this difference in our calculations, we assume its affect to be negligible. Thus the same calculations should apply for both the cases, Figure 7.1(a) and Figure 7.1(b). We now describe the method used for these calculations. We assume that the observer's perceptual process is able to use his knowledge of L and b and "bright" and "dark" standards to evaluate the threshold intensity at which the two conditional probability densities involved (i.e. the one given C "bright" and the one given C "dark") are equal. Then he makes a decision about component C being "bright" or "dark" according as the observed intensity of C falls above or below the threshold. Thus he acts as a maximum a posteriori probability (MAP) detector

[21]. Since we assume both possibilities are a priori equally likely, the MAP detector is equivalent to a Maximum Likelihood (ML) detector [21]. This assumption of MAP or ML detection by the observer is at least plausible, and it has been verified in some cases in the past [28].

(d) The case of distinguishing between two or four orientations of a shape, as in Figure 7.2(a) and Figure 7.2(b) is somewhat different from the cases discussed in paragraph (c) above. In the case of the two orientations of Figure 7.2(a), it is reasonable to assume that the observer makes his decision by determining whether the central square in the top edge (C1) or the right edge (C2) is the brighter one. There he does not need to make any judgement about "bright" or "dark" levels as in the cases of Figure 7.1(a) and Figure 7.1(b). This process is equivalent to the one used in the detection of orthogonal signals in communications theory [21]. The idea of using this orthogonal signal theory in the determination of orientation of a noisy shape, for the case of additive white Gaussian noise, has been used before [31]. In our case, however, the noise is not AWGN.

7.3.2 Remark on method used

Finally, we note that even if the observer does not actually behave according to these assumptions, the probabilities resulting from his visual and perceptual mechanism, whatever it actually is, can at least be compared experimentally with

the ones calculated for the idealized observer which we assumed. This is what we hope to do with our simulations and experiments. We find that as the parameters b, L, J are varied, the experimental and calculated probabilities follow the same trends, and that the experimental probabilities mostly fall between upper and lower bounds represented by the calculations based on the assumed optimal and suboptimal combinations of the J^2 intensity estimates within each basis square. These results at least lend support to the arguments used in our calculations, although they do not justify them, since other alternative theories might give a similar agreement with theory. In any case, the close agreement of the experimental results with the theory shows that the theory is one reasonable way to make predictions of probabilities for the kind of problems considered here, within the range of parameters investigated.

7.3.3 Details of method of calculation

7.3.3.1 Case of two alternatives of Figure 1

We first consider the case where there are two alternative forms as in Figures 7.1(a) and 7.1(b). Figure (7.3) illustrates schematically the calculations for this case. Although the two shapes may be characterized by nine orthogonal components, they differ in only one such component, so that, to calculate p_c, only this component, C, need be considered.

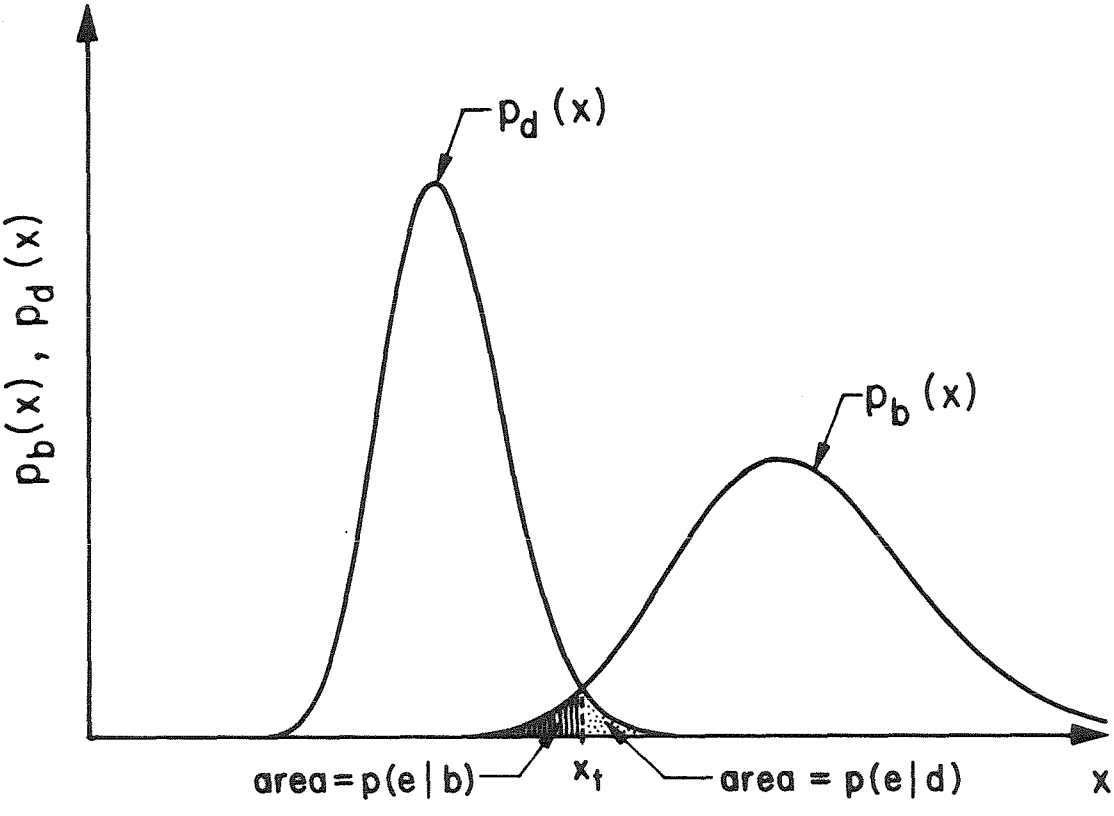

Figure (7.3): Schematic for ML detection mechanism assumed in "2 alternative forms" calculations

For any one of the nine basis vectors, the corresponding component can be either MP_0 or $MP_0(1+b)$. Given either of these values, a priori, we can calculate the probability density function for the corresponding total intensity following Chapter 3. This density function is a gamma-distributed random variable with M degrees of freedom; where $M = LJ^2$ in the optimal case, or $M = LJ$ in the suboptimal case [see subsection (7.3.1)]; thus

$$p_d(x) = \frac{1}{P_0^M} x^{(M-1)} \frac{e^{-x/P_0}}{\Gamma(M)} \quad \text{for } x>0 \tag{7.1}$$

$$p_b(x) = \frac{1}{[P_0(1+b)]^M} x^{(M-1)} \frac{e^{-x/P_0(1+b)}}{\Gamma(M)} \quad \text{for } x>0 \tag{7.2}$$

with the probability densities being 0 for $x < 0$. Here, $p_d(x)$, $p_b(x)$ are, respectively, the probability density functions for the cases where the relevant component C of the shape is the darker or brighter alternative, i.e., has a mean value of P_0M or $P_0(1 + b)M$. Thus, for the case of Set 1, shown in Figure 7.1(a), $p_b(x)$ is the probability density for component C of the form 0, while $p_d(x)$ is the one for the form U. In the case of Set 2, shown in Figure 7.1(b), $p_b(x)$ is the probability density for the relevant component C in the shape of S, while $p_d(x)$ is that for shape 0.

The calculation of p_c for the ML detector (in this case equivalent to the MAP detector) now proceeds as follows:

The intersection of the two density curves represented

by Equations (7.1) and (7.2) gives us a threshold intensity of x_t. If the observed intensity falls above this, then the probability of occurrence of x is higher if the brighter alternative is assumed to be present than if the darker is assumed present. So the ML detector chooses the brighter alternative. If $x < x_t$, it chooses the darker alternative. By equating the density functions of Equations (7.1) and (7.2), we get, for the threshold intensity x_t,

$$x_t = P_0\{ M \ln(1+b)\} \left(\frac{1+b}{b}\right) \qquad (7.3)$$

The probability of error given that the brighter alternative is actually present is

$$p(e|b) = \int_0^{x_t} p_b(x)dx \qquad (7.4)$$

The probability of error given that the darker alternative is actually present is

$$p(e|d) = \int_{x_t}^{\infty} p_d(x)dx \qquad (7.5)$$

The overall probability of error, for equally likely bright and dark alternatives is

$$p_e = \frac{1}{2}[p(e|b) + p(e/d)] \qquad (7.6)$$

and the probability of correct decision is

$$p_c = 1 - p_e \qquad (7.7)$$

The probability of correct decision can be calculated from the above equations, using either numerical integration or tables of the gamma distribution. Since these tables are not available for all the desired values of M in our simulation and for all the x_t values that occur, we used numerical integration.

For large values of M, we can use the Gaussian approximations for the pdf's involved (see Chapter 3). Since extensive tables of the Gaussian distribution are readily available, this method is useful. We checked the results obtained by using this approximation with the numerical integration results using the exact Equations (7.1)-(7.7) for values of M equal to 18, 24 and 25, and with b = 1.0. We found that the values of $p(e|b)$, $p(e|d)$ and p_c calculated by these two methods agreed to within $\pm$ 0.02 or better, which is a close enough agreement for our purposes. Hence, for all values of $M \geq 25$, we used the Gaussian approximation. For all values of $M < 25$, we used the exact method with numerical integration to calculate p_c. These calculated probabilities p_c, and the experimental ones, $\hat{p_c}$ are shown in Tables 7.1(a) and 7.1(b). See the subsection on 'Tables of Results' for further details. The calculations for all cases are done for both, $M = LJ^2$ and $M = LJ$.

The use of the Gaussian distribution along with a further approximation for $b \ll 1$, leads to the following useful quick approximation to the value of the threshold x_t and p_c:

$$x_t \simeq P_0 M\left(1+ \frac{b}{2}\right) \tag{7.8}$$

and

$$p_c \simeq 1 - \frac{1}{2}\left\{ Q\left(\frac{b\sqrt{M}}{2(1+b)}\right)+ Q\left(\frac{b\sqrt{M}}{2}\right)\right\} \tag{7.9}$$

where $$Q(z) = \int_z^\infty \frac{1}{(2\pi)^{1/2}} \exp(-y^2/2)\ dy . \tag{7.10}$$

7.3.3.2 Case of 2 orientations of Figure 7.2(a)

Here we assume that the observer looks at 2 vector components C_1 and C_2 of the pattern he is shown, and decides which of them is brighter. He does not have to try to decide, as in section (7.3.3.1), whether the absolute level of a given component of the pattern is closer to P_0M or to $P_0(1 + b)M$.

We can see that the probability of correct decision is the probability that the component which is actually brighter is seen as brighter, and is therefore given by

$$p_c = \int_0^\infty p_G(x)\,dx \int_0^x p_B(y)\,dy \tag{7.11}$$

where $p_G(x)$, $p_B(x)$ are given by the same expressions as $p_b(x)$ and $p_d(x)$ in Equations (7.1) and (7.2) respectively.

We use the notation $p_G(x)$ and $p_B(x)$ to emphasize the similarity with the calculations in Chapter 5. By substituting these gamma density functions for $p_G(x)$ and $p_B(x)$ from Equations (7.1) and (7.2) in Equations (7.11), we can obtain p_c using numerical integration.

Alternatively, for M > 25, we can use the Gaussian approximation for $p_G(x)$ and $p_B(x)$, and obtain the approximation

$$p_c \simeq 1 - Q\left(\frac{b\sqrt{M}}{\sqrt{2[1+(1+b)^2]}} \right) \qquad (7.12)$$

The agreement between the p_c calculated by this equation and the exact ones is very good for all the cases investigated, since all of them had M > 51 and b = 0.26.

These calculated results are shown in Table 7.2(a), along with the experimental ones. The calculated probabilities for all cases are shown assuming $M = LJ^2$ as well as M = LJ.

7.3.3.3 The case of four alternative orientations as in Figure 2(b)

In this case, the observer decides which of four basis squares (top center, bottom center, left center or right center) is brightest. The probability of correct decision is then

$$p_c = \int_0^\infty p_G(x)\,dx \left[\int_0^x p_B(y)\,dy\right]^3 \qquad (7.13)$$

These p_c were calculated by numerical integration for a few cases and are shown, along with the experimental values, in Table 7.2(b). Calculations are shown for both $M = LJ^2$ and $M = LJ$.

7.3.4 Analogy with optimum detector theory

The decision mechanisms in all the discrimination problems discussed above are analogous to some case or other of optimal detection of signals in noise. To see the analogy more clearly, we make the Gaussian approximation to the probability density functions of Equations (7.1) and (7.2), and assume a low contrast, $b \ll 1$. Then the two probability density functions have approximately equal variances MP_0^2, and look like those for signals in additive white Gaussian noise (AWGN) in problems of optimal detection [21]. The mean values of the intensities here then correspond to the signal amplitudes in the optimal detector case and the speckle-induced variability of intensity corresponds to the variability of the received signal-plus-noise vector caused by additive noise in the receiver. By working out the details, we can obtain the approximate Equations (7.10) and (7.12).

7.4 The simulations and experiments

7.4.1 The procedure

In order to test the validity of our calculation of p_c, we performed computer simulations of four different sets of forms for many different sets of parameters. The most extensive simulations were for the case of Set 1 in Figure 7.1(a), i.e., for distinguishing between U and O .

The procedure used was to generate pictures of size $(300)^2$ pixels, each picture corresponding to a fixed value of b, L, J. Each picture, which we will also refer to as a sample, consisted of a matrix of patterns each of size D * D pixels as in Chapter 6. By 'pattern', we mean a speckle-corrupted version of any one of the allowed alternative forms. That is, for the sets of Figure (7.1) and Figure 7.2(a), a random matrix of equally likely 0's and 1's was generated by computer. Then, corresponding to the 0's, one of the two alternative forms, say S0, was generated, and, corresponding to the 1's, the other form S1 was generated. For Set 4, a random matrix of equally likely 0's, 1's, 2's and 3's was similarly used. The number of patterns per picture, N, was as many as could reasonably be fitted in the picture. N varied from picture to picture depending on the value of D, but was in the range of 64 to 256. The background in which the patterns were set was, as stated before, completely dark. The arrangement of each picture is shown schematically in Figure (7.4). Some examples of the

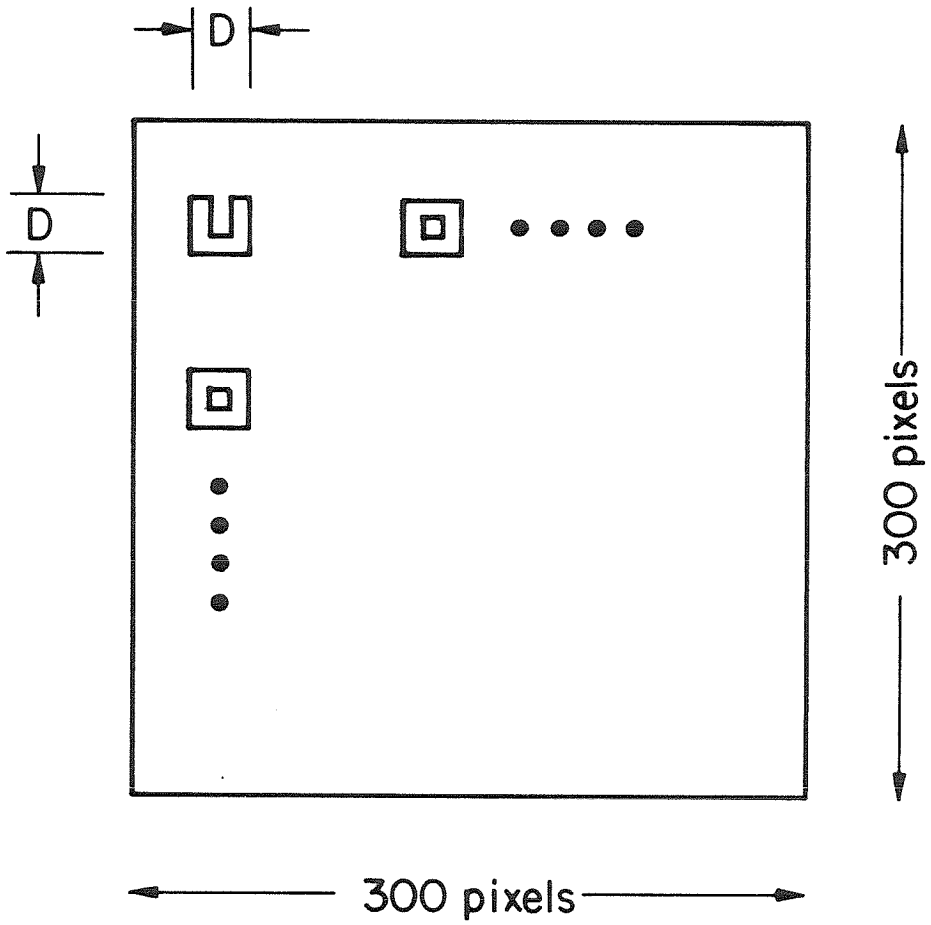

Figure (7.4): Arrangement of patterns in a picture in form discrimination experiments

pictures simulated are shown in Figure (7.5) and are discussed in the next subsection.

The experiments consisted of showing speckled pictures obtained with various sets of parameters, b, J (or D), and L, to an observer, who was given certain directions and asked to decide which of the specified possible forms each pattern in the picture represented. The facts revealed to him, which were chosen following the procedure given in Green and Swets [28] were the following. He knew

(a) what exactly the various possible forms were

(b) that they were all equally probable

(c) the contrast ratio (1+b)

(d) the number of looks per pixel L

(e) the value of J.

The responses of the observer were then compared to the original (correct) responses. From the proportion of errors made by the observer for each shape, we can estimate the probability of error given each shape and the overall probability of correct decision. These experimental results were then compared with the analytical ones.

7.4.2 Examples of simulated pictures

7.4.2.1 Overview

As in Chapter 6, the size of each sample is irrelevant to the discrimination problem, but we specify it for completeness; all samples shown are 200*200 portions of an original 300*300 picture, which is why some of the patterns at the edges are incomplete. The last 6 lines are the caption, which specifies, among other things, J, L, (1+b) in dB, and the stretch parameters.

7.4.2.2 The samples

(1) The first photograph contains samples M2, R16, R10 and M5. All these are samples consisting of patterns that are either O's or U's.

M2 is a 3 dB, 34 look, J=3 picture which is almost like an unspeckled picture in the sense that all the patterns are clearly recognizable.

R16 also has J=3 and 3 dB contrast but L=8, and the patterns are much harder to distinguish than in M2. Thus, a comparison of R16 and M2 shows the improvement in discriminability with increase in L.

R10 and M5 both have 1 dB contrast and L=20, but J=5 in R10 and 3 in M5. The patterns in R10 are much more easily identified than those in M5, which shows the improvement in discriminability with increase in pattern size. We will refer to the poor discriminability in M5 in the Conclusions section in this chapter.

(1)

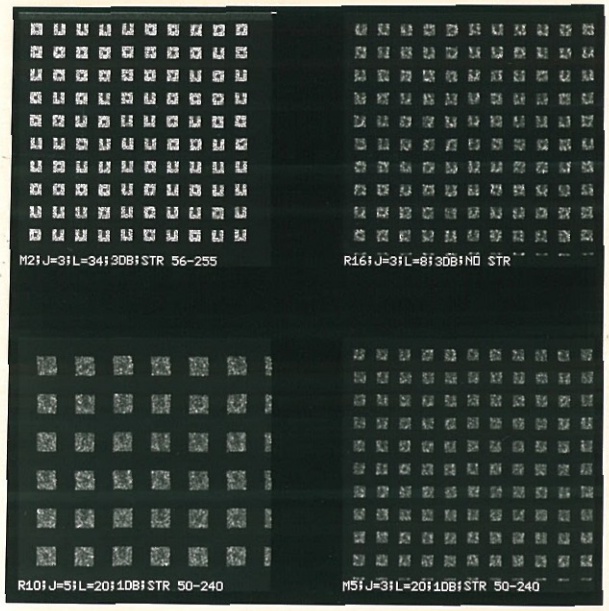

(2)

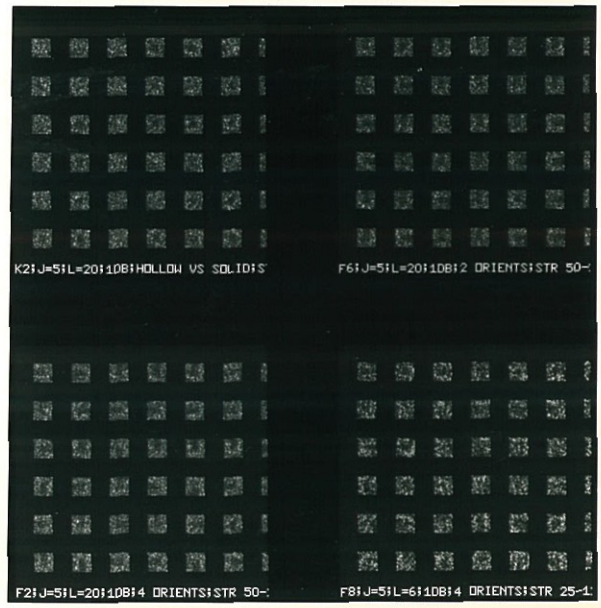

Figure (7.5): Form discrimination experiments; examples of simulated pictures

(2) The second photograph shows K2, F6, F2 and F8. The captions show what the patterns (see Figure (7.1)) in each of these are. K2 has O's and S's while R10 has U's and O's. Although a comparison of K2 and R10 is not sufficient proof, it does illustrate the fact that for the same parameters, the forms U and O are just as discriminable as S and O. (See photograph 1 for R10)

F6 has patterns with one of two orientations, U, D, R, L. F2 and F6 have the same parameters, and a comparison indicates that it is harder to distinguish between 4 alternatives (as in F2) than between 2 (as in F6). A comparison of F8 and F2 shows the improvement in discriminability as L increases.

7.4.3 Tables of results

Tables 7.1(a), 7.1(b), 7.2(a), 7.2(b) show the experimental and calculated probabilities for the 4 sets in Figures (7.1) and (7.2). The tables give, for each picture, at least the following:

(a) the parameters b, L, D

(b) the probabilities of correct decision $\hat{p}_c$ obtained experimentally

(c) the probabilities of detection p_c obtained analytically

(d) In addition, for the cases in Sets 1 and 2, where only two alternatives are present, and the probabilities of error given each alternative are expected, theoretically, to be

Table 7.1(a): Two alternative forms; Set 1 (U and O)

(1+b) dB	L	J	M= (a) LJ_2 (b) LJ^2	# of samples	Probabilities Theoretical p(e|O)	p(e|U)	p_c	Experimental p(e|O)	p(e|U)	$\hat{p}_c$	Bounds on $\hat{p}_c$ (cf. Ch.5) lower, upper
1	4	7	28 196	49	0.28	0.26	0.73 0.95	0.24	0.18	0.80	0.66, 0.90
	6	5	30 150	81	0.27	0.26	0.75 0.92	0.15	0.29	0.78	0.67, 0.86
	5	6	30 180	64	0.27	0.26	0.75 0.94	0.21	0.11	0.84	0.73, 0.92
	15	3	45 135	256	0.26	0.18	0.78 0.91	0.09	0.44	0.69	0.63, 0.75
	8	6	48 288	64	0.22	0.19	0.79 0.98	0.12	0.06	0.91	0.81, 0.97
	7	7	49 243	49	0.22	0.19	0.79 0.96	0.22	0.08	0.86	0.73, 0.94
	10	5	50 250	81	0.22	0.19	0.79 0.97	0.22	0.07	0.85	0.75, 0.92
	17	3	51 153	256	0.24	0.17	0.79 0.92	0.21	0.06	0.86	0.81, 0.90
	20	3	60 180	256	0.22	0.15	0.82 0.94	0.13	0.27	0.78	0.72, 0.83
	13	4	52 208	144	0.22	0.19	0.79 0.95	0.15	0.20	0.82	0.75, 0.88
	20	5	100 500	81	0.14	0.11	0.87 0.99	0.09	0.09	0.91	0.82, 0.96
	34	3	102 306	356	0.14	0.11	0.88 0.98	0.04	0.24	0.83	0.79, 0.87

Table 7.1(a) (continued)

(1+b) dB	L	J	M= (a) LJ (b) LJ^2	# of samples	Probabilities Theoretical p(e|0)	p(e|U)	p_c	Experimental p(e|0)	p(e|U)	$\hat{p}_c$	Bounds on $\hat{p}_c$ lower, upper
3	1	5	5 25	100	0.27	0.18	0.78 0.96	0.18	0.12	0.85	0.76, 0.91
	3	2	6 12	---	0.24	0.16	0.80 0.88	----	----	----	----, ----
	2	3	6 18	256	0.24	0.16	0.80 0.93	0.10	0.14	0.88	0.83, 0.92
	8	3	24 72	256	0.07	0.02	0.95 0.998	0.014	0.017	0.98	0.96, 1.0
	6	4	24 96	144	0.07	0.02	0.95 0.999	0.05	0.0	0.98	0.94, 1.0
	5	5	25 125	81	0.07	0.02	0.95 0.999	0.0	0.0	1.0	0.96, 1.0
	34	3	102 306	100	0.0005	0.0003	0.999 $1-10^{-6}$	0.0	0.0	1.0	0.96, 1.0
5	1	1	1 1	160	0.41	0.19	0.70 0.70	0.36	0.31	0.67	0.59, 0.74
	1	2	2 4	192	0.29	0.15	0.78 0.87	0.21	0.22	0.78	0.71, 0.84
	4	2	8 16	192	0.09	0.04	0.93 ——	0.05	0.03	0.96	0.92, 0.98
	3	3	9 27	256	0.06	0.04	0.95 ——	0.05	0.0	0.98	0.95, 0.99

Table 7.1(b): Two alternative forms; Set 2 (O and S)

(1+b) dB	L	J	M= (a) LJ_2 (b) LJ^2	# of samples	Probabilities Theoretical p(e|S)	p(e|O)	p_c	Experimental p(e|S)	p(e|O)	$\hat{p}_c$	Bounds on $\hat{p}_c$ lower, upper
1	6	5	30	100	0.27	0.26	0.75	0.23	0.10	0.84	0.75, 0.91
			150				0.92				
	10	5	50	100	0.22	0.19	0.79	0.26	0.14	0.80	0.71, 0.87
			250				0.97				
	17	3	51	256	0.24	0.17	0.79	0.09	0.28	0.82	0.77, 0.86
			153				0.92				
	20	5	100	100	0.14	0.11	0.87	0.14	0.02	0.92	0.85, 0.97
			500				0.99				

Table 7.2(a): Two orientations; Set 3 (U and R)

(1+b) dB	L	J	M= (a) LJ_2 (b) LJ^2	# of samples	Probabilities Theoretical p_c	Experimental $\hat{p}_c$	Bounds on $\hat{p}_c$ lower, upper
1	6	5	30	100	0.81	0.85	0.76, 0.91
			150		0.977		
	10	5	50	100	0.87	0.96	0.90, 0.99
			250		0.995		
	17	3	51	256	0.88	0.86	0.81, 0.90
			153		0.98		
	13	4	52	144	0.88	0.87	0.80, 0.92
			208		0.99		
	20	5	100	100	0.95	1.0	0.96, 1.0
			500		0.9998		

Table 7.2(b): Four orientations; Set 4 (U, D, L, R)

(1+b) dB	L	J	M= (a) LJ (b) LJ^2	# of samples	Probabilities Theoretical p_c	Experimental $\hat{p}_c$	Bounds on $\hat{p}_c$ lower, upper
1	6	5	30	100	0.64	0.70	0.60, 0.79
			153		0.95		
	10	5	50	100	0.75	0.82	0.73, 0.89
			250		0.99		
	17	3	51	256	0.75	0.82	0.77, 0.86
			153		0.95		
	13	4	52	144	0.75	0.81	0.74, 0.87
			208		0.973		
	20	5	100	100	0.88	0.95	0.89, 0.99
			500		0.999		

different, these probabilities are also given.

For cases where M<25 in Tables 7.1(a) and (b) and in Table 7.2(a), the values listed as `Calculated Probabilities' were obtained as described in section 7.3 by numerical integration of the appropriate exact expressions. For cases where M>25, we used the expressions based on the Gaussian approximation, which gave results within +-0.02 or better of the exact ones in these cases, as we verified in a few cases. For all cases of Table 7.2(b), the exact expression was used.

The results on calculated probabilities include both those made with the assumption of $M=LJ^2$ (rows labeled (b)) and those with the assumption M=LJ (rows labeled (a)).

7.4.4 Analysis of results and conclusions from experiments

From our theoretical and experimental calculations, we arrived at the following conclusions:

(a) We find that the p_c values obtained experimentally almost always lie between those calculated analytically for the two cases $M=LJ^2$ and M=LJ.

(b) The experimental estimates of probability of error given the individual shapes p(e|b) and p(e|d) in cases where they are different for the different shapes don't show the same order relation as the corresponding calculated probabilities, i.e., don't agree with the theory about which

of the two is greater. However, this may be due to experimental error in measuring small quantities, since the probabilities of error in all the cases dealt with are fairly low. Or, it could be due to an observer bias which changes from picture to picture.

(c) The use of the Gaussian approximation to the gamma distribution leads to simple approximate expressions for p_c that are, for large M, very close to the ones obtained using the exact gamma distributions and numerical integration.

7.5 Conclusions

In this chapter, we made certain assumptions in order to calculate the probabilities of correct decision in the problem of discriminating between certain geometrical forms. We checked our calculations experimentally using computer simulated pictures. From the results of our experiments, we arrive at the following conclusions:

(a) The conclusions of paragraph (a) of section 7.4.3 are valid for the fairly wide range of parameters covered experimentally. This fact lends credence to the theories about the identification process used in making the analytical calculations, although it does not actually prove that these theories are correct. There is an insufficient number of patterns for each set of parameters to validate completely the theories used, since other plausible theories might equally well produce the same degree of agreement with

the experimental probabilities. However, whatever the mechanism used, the results indicate that the present method of calculation can be used to make predictions about p_c. The simple approximate expressions for the low- contrast, high M cases should be especially useful in designing a system to provide a desired p_c for any of the cases considered.

(b) We now wish to compare the discriminability of forms in a given picture with their detectability. For instance, consider the case where we are trying to identify a feature in a speckled background, and if we assume b=.26, L=13, we see from Table 7.1(a) that, for p_c=0.95, we need J=4 or D=12, even by the optimistic estimate which uses $M=LJ^2$ rather than M=LJ, and even if we know that there is only one alternative form allowable that differs in one orthogonal component. This is much larger than the D=7.3 required for the detection of that feature in a 12 look 1 dB 100*100 picture. Note that the form discrimination results are based on a completely dark background. If the background in which the patterns are set were also speckled, the probability of correct decision would decrease because there is, in general a worsening of form recognizability as the background is made non-uniform.

As another example, consider sample M5 in the first photograph of Figure (7.5). This sample has 1 dB contrast and D=9, L=20 and only two possible patterns in the picture.

The patterns are seen to be hard to identify in this sample. If we merely had to detect one of these patterns located in a speckled background, it would be clearly detected, as shown in Chapter 5. We can see from the photograph in Chapter 5 that the 1 dB, M=648 cases of S23 and B21 give good detectability, and an M=1620 picture like the one we are discussing would be even better.

Thus, by combining the results of Chapters 5 and 7, we are able to specify how much harder it is to discriminate between two or more forms than it is to detect them.

8. SUMMARY AND CONCLUSIONS

8.1 General

We have investigated the problem of characterizing the degradation caused by speckle in pictures generated by SAR or other coherent mapping systems. For this purpose, we have considered pictures in which, in the absence of speckle, the intensity at any point in the picture can be only one of two values P_0 or $P_0(1+b)$. The contrast ratio $(1+b)$ and the number of looks per pixel, L are picture parameters that are important in determining the quality of the picture. We have considered, both theoretically and experimentally, the extent of picture quality improvement as either of these parameters increases. The eye-charts of Chapter 2 provide a good visual demonstration of speckle effects, and, in Chapter 5, we showed how the theoretical model used for detection of features predicts some of the quantitative relations between b, L and picture quality in the case of the eye-charts.

We have considered the problem of determining whether or not some intensity variation seen in a speckled picture is inherent in the picture. That is, would the variation be present in the picture if speckle were eliminated, or is it an artifact produced by speckle? The answer, for a picture with given parameters b and L and picture size W*W, depends on what kind of inherent intensity variation we are looking

for. We have considered problems that fall into one of three categories of intensity variation that we might be looking for.

In each category, we made some theoretical calculations that served as a guideline for our experiments. In the experiments, we showed simulated pictures to an observer who was given a specific decision task, and we noted the percentage of correct responses made by the observer. The results of these experiments agree with the calculations based on the particular theoretical model used to predict the outcome of the experiments. However, the number of experiments made is not large enough to rule out other plausible theoretical models. But the experiments do show that the theoretical models we have used give reasonably good results. We now discuss these three categories in more detail.

8.2 The three categories

(a) The first category deals with detection of a small square of D*D pixels in a picture of W*W pixels. We related theoretically, the probability p_c, of the observer's correctly detecting this feature, to the parameters b, L, D, W, and verified this relation experimentally. We also obtained a simple formula that can be used to design the system parameters b and L required for a specified theoretical probability, $p_c = p_t = 0.95$, of correctly detecting

a square of D*D pixels in a picture of N=W*W pixels. We first obtain M from the formula, repeated below, and then use $L=M/D^2$. In all our equations, Q(x) is the integral of the normal probability density function from x to ∞.

$$\frac{z_{p1}(N) + (4M-1)^{1/2}}{(1+b)^{1/2}} - (4M-1)^{1/2} = z_{pt} \qquad (5.29)$$

where $Q(z_{pt}) = p_t$; $\quad Q[z_{p1}(N)] = p_1 = (0.95)^{1/N}$ (5.30)

(b) The second category deals with the detection of a V*H pixel grating consisting of alternating bright and dark lines, each one D pixels wide. We defined a speckle-signal-to-noise ratio (SSNR) expression that involves all the grating parameters b, D, L, V, H and fitted a regression curve to the data of experimentally obtained probability of correct decision p_c. The SSNR expression and the regression equation are repeated below. Together, they can be used to predict p_c for a grating of given parameters. We also derived another very similar expression for SSNR, which can be used equally well to predict p_c.

$$SSNR_1 = \frac{16}{\pi^2} \frac{(b/2)^2}{(1+b/2)^2+(b/2)^2} L \left[HV \frac{2D}{1.5}\right]^{1/2} \qquad (6.20c)$$

$$\hat{p_c} = 1 - 1/2 \exp[-0.08264\ (SSNR)_2] \qquad (6.24b)$$

(c) The third category deals with distinguishing between 2 or 4 alternative geometrical forms. Each of these forms is derived from a D*D pixel square, consisting of 9 smaller

squares, each of size J*J pixels. One or two of these nine smaller squares is "dark" while the others are brighter by a factor (1+b).

(1) We first considered the problem of discriminating between the two forms which we have designated as S and O. The results here also apply to the case of distinguishing between U and O. One of the simpler formulas derived, which is valid for L>>1, b<<1, is quoted below. It relates the probability of correct decision p_c, to the parameters b, L, J.

$$p_c \simeq 1 - \frac{1}{2}\left\{ Q\left(\frac{b\sqrt{M}}{2(1+b)}\right) + Q\left(\frac{b\sqrt{M}}{2}\right)\right\} \qquad (7.9)$$

where $M=LJ$ and $M=LJ^2$ lead to lower and upper bounds on p_c and $Q(x)$ is defined as in the case of the equation quoted above. (2) We also considered the problem of distinguishing between 2 orientations, U and R, and the problem of distinguishing between 4 orientations U, D, R, L. We quote a simple formula for the case of large number of looks per pixel for the 2-orientation problem.

$$p_c \simeq 1 - Q\left(\frac{b\sqrt{M}}{\sqrt{2[1+(1+b)^2]}}\right) \qquad (7.12)$$

where $M=LJ$ and $M=LJ^2$ lead to lower and upper bounds on p_c and $Q(x)$ is defined as in the case of the two equation above.

8.3 Threshold parameters

For each of the three categories of problems considered

above, we found the required threshold value of D for a probability of correct decision $p_c = 0.95$, for a picture with parameters that are typical for a SAR system.

For a 1 dB contrast ratio and a 12 look picture, we found the following results:

(1) To detect a small bright speckled square in a 100*100 pixel picture, we need D equal to 7 or 8.

(2) To detect a speckled line pair grating of 100*100 pixels, a line thickness D equal to 2 will more than suffice.

(3) To determine the form of a speckled feature, even in the favorable circumstances in which the background is unspeckled, and only one alternative form is available, we need D>12.

These results demonstrate that it is not possible to characterize the quality of the imagery produced by a coherent system by one single parameter, such as resolution cell size or highest spatial frequency of grating resolvable by the system. Thus, if for a given set of parameters b and L, a line pair grating of width D is detectable with near certainty, a square of size D may not be detectable. Again, if a D*D square can be located with near certainty, its form cannot necessarily be determined. The detectability of line pair gratings is a poor criterion for characterizing the

quality of the picture. Discriminability of form imposes the most stringent conditions on the system, in terms of the b and L required to make a D*D form identifiable.

9. REFERENCES

1. J. D. Rigden and E. I. Gordon, Proc. IRE 50, 2367 (1962).
2. A. Labeyrie, Astron. Astrophys. 6, 85 (1970).
3. W. T. Welford, J. Opt. Soc. Am., 66, 1172 (1976).
4. N. George, C. R. Christensen, J. S. Bennett, and B. D. Guenther, J. Opt. Soc. Am., 66, 1282 (1976).
5. P. E. Green, Radar Astronomy, ed. J. V. Evans and T. Hagfors (McGraw-Hill, New York, 1968).
6. Acoustical Holography, ed. P. S. Green, (Plenum, New York, 1974).
7. P. Beckmann and A. Spizzichino, The Scattering of Electromagnetic Waves from Rough Surfaces (MacMillan, New York, 1963).
8. L. J. Porcello, N. O. Massey, R. B. Innes, and J. M. Marks, J. Opt. Soc. Am., 6, 1305 (1976).
9. R. O. Harger, Synthetic Aperture Radar Systems: Theory and Design (Academic, New York, 1970).
10. J. J. Kovaly, Ann. N.Y. Acad. Sci. 187, 154 (1972).
11. J. I. Marcum and P. Swerling, IRE Trans. IT-6, 59 (1960).
12. D. E. Kerr, ed., MIT Radiation Laboratory Series (McGraw Hill, New York, 1951).
13. G. Gould, S. F. Jacobs, J. T. La Tourrene, M. Newstone, and P. Rabinowitz, Appl. Optics 3, 648 (1964).

14. L. N. Ridenour, ed., MIT Radiation Laboratory Series (McGraw Hill, New York, 1947).

15. J. W. Googman, J. Opt. Soc. Am. 66, 1145 (1976).

16. S. A. Butman and R. G. Lipes, Deep Space Network Progress Report 42-29, Jet Propulsion Lab., CA, 46 (1975)

17. A. Jain, Ph. D. Thesis, Calif. Inst. of Tech., Pasadena, Calif. (1973).

18. J. S. Zelenka, J. Opt. Soc. Am. 66, 1295 (1976).

19. O. H. Schade, Image Quality (RCA, Princeton, 1975).

20. Laser Speckle and Related Phenomena, ed. J. C. Dainty, (Springer, Berlin, 1975).

21. J. M. Wozencraft and I. M. Jacobs, Principles of Communications Engineering (Wiley, New York, 1965).

22. L. H. Enloe, Bell Sys. Tech. J. 46, 1479 (1967).

23. J. C. Dainty, Optica Acta 18, 327 (1971).

24. K. S. Miller, Multi-dimensional Gaussian Distributions (Wiley, New York, 1964).

25. Handbook of Mathematical Functions, ed. M. Abramowitz and I. A. Stegun, (NBS, 1970).

26. C. E. K. Mees, The Theory of the Photographic Process (MacMillan, New York, 1954).

27. T. N. Cornsweet, Visual Perception (Academic, New York, 1970).

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

29. W. J. Dixon and F. J. Massey, Introduction to Statistical Analysis (McGraw Hill, New York, 1969).

30. Z. L. Budrikis, IEEE proc. 60, 771 (1972).

31. T. W. Bernard, A Symposium on Sampled Images, Perkin-Elmer, Norwalk, 1971.

32. J. M. Findlay, Vision Res. 9, 157 (1969).

33. S. H. Bartley, Vision (Hafner, New York, 1963).

34. C. W. Helstrom, Statistical Theory of Signal Detection (Pergamon, New York, 1960).

35. E. J. Gumbel, Statistics of Extremes (Columbia Univ. Press, New York, 1958).

36. G. S. Watson, Ann. Math. Stat. 25, 798 (1954).

37. F. Radcliff, M. K. Hartline, and W. H. Miller, J. Opt. Soc. Am. 53, 110 (1963).

38. H. Helson, J. Opt. Soc. Am. 53, 179 (1963).

39. M. L. Davidson, J. Opt. Soc. Am. 58, 1300 (1968).

40. F. W. Campbell and G. Green, J. Physiol. London 181, 576 (1965).

41. T. G. Stockam, Jr., Proc. IEEE 60, 828 (1972).

42. L. D. Harmon and B. Julesz, Science 180, 1194 (1973).

43. F. W. Campbell and J. G. Robson, J. Physiol. London 197, 551 (1968).

44. C. F. Stromeyer and B. Julesz, J. Opt. Soc. Am. 62, 1221 (1972).

45. A. D. Whalen, Detection of Signals in Noise (Academic, New York, 1971).

46. W. M. Brown, IRE Trans. IT-4, 137 (1958).

47. J. R. Pierce and E. C. Posner, Introduction to Communication Science and Systems (Plenum, New York, 1980).

48. B. Ostle, Statistics in Research (Iowa State Univ. Press, 1963).

49. M. Ezekiel and K. A. Fox, Methods of Correlation and Regression Analysis (Wiley, New York, 1959).

50. The Eye, ed. H. Davson, (Academic, New York, 1962).

51. Form Discrimination as Related to Military Problems, ed. J. W. Wulfech and J. H. Taylor (Publication 561, NAS-NRC, Washington D. C., 1957)