SPIE Proceedings, vol. 3016, paper 1 (1997).

Image quality and entropy masking

Andrew B. Watson[a], Robert Borthwick[b], & Mathias Taylor[b]

[a]NASA Ames Research Center, Moffett Field, CA 94035-1000[-]
[b]Foothill College, Los Altos Hills, CA

ABSTRACT

Image quality models usually include a mechanism whereby artifacts are masked by the image acting as a background. Scientific study of visual masking has followed two traditions: *contrast masking* and *noise masking*, depending primarily on whether the mask is deterministic or random. In the former tradition, masking is explained by a decrease in the effective gain of the early visual system. In the latter tradition, masking is explained by an increased variance in some internal decision variable. The masking process in image quality models is usually of the gain-control variety, derived from the contrast masking tradition. In this paper we describe a third type of masking, which I call *entropy masking*, that arises when the mask is deterministic but unfamiliar. Some properties and implication of entropy masking are discussed. we argue that image quality models should incorporate entropy masking, as well as contrast masking.

Keywords: noise, contrast gain control, normalization, pattern recognition, signal detection theory, spatial vision, background, entropy masking, image compression, DCTune

1. IMAGE QUALITY

Over the past three decades, in concert with the growth in digital imaging technology, there have been many attempts to develop models or metrics for image quality that incorporate elements of human visual sensitivity. Many of these efforts have been concisely reviewed by Ahumada[1], and many of the general issues are reviewed in Watson[2]. While details vary, most of the models incorporate the modules or processes indicated in Figure 1.

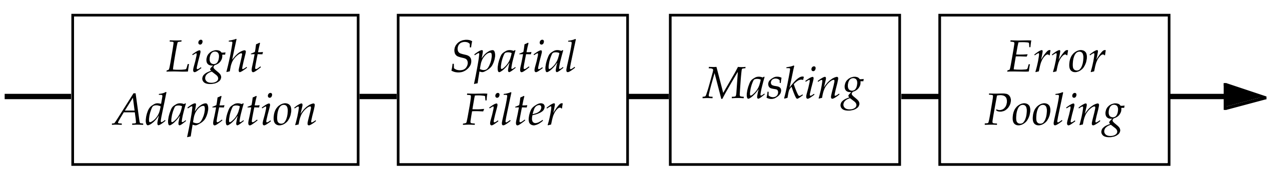

Figure 1. Components of image quality models.

The luminance image is first subjected to a light adaptation process, which typically involves a conversion to contrast, and possibly further adjustments. The spatial filter will incorporate the contrast sensitivity function, and may also include spatial channels tuned for narrow bands of spatial frequency and orientation. At the masking stage, the filter responses are altered in some way to reflect the reduced visibility of signals presented on a contrasty background. At the final stage, some computation is performed to express the error of the test image relative to a reference image, and these errors are combined in some way to yield a scalar quality metric.

2. AN EXAMPLE: DCTUNE

As a concrete example of this general outline, I briefly describe the DCTune quality metric. To achieve practical efficiencies in the optimization of DCT-based image compression, this metric uses the blocked Discrete Cosine Transform (DCT) as the filter bank. The image error is computed as the arithmetic difference between DCTs of test and reference images. The error DCT coefficients are weighted by their absolute visibility, as a function of DCT frequency, the given display mean

[-] Further author information: email: abwatson@mail.arc.nasa.gov ; WWW: http://vision.arc.nasa.gov/

luminance and display visual resolution, so that they may be expressed in jnd units. Light adaptation is achieved by using the DC coefficient from each block (the average graylevel in the block) to adjust the visibilities in each block appropriately.

The errors are further adjusted by a contrast masking process, in which each error is attenuated by a power function of the corresponding coefficient in the reference image, also converted to jnds. The resulting masked jnds are then pooled over frequency and over space using a Minkowski metric, to yield a total perceptual error. The inverse of this error is the DCTune Quality Metric. The metric is so calibrated that a value of 1 indicates a visually lossless image.

To illustrate the success of this metric, Figure 2 shows the results from an experiment in which two observers judged whether briefly presented images appeared to be compressed[3]. Five different images, and five bit-rates (0.25, 0.5, 0.75, 1.0, and 8.0 bits/pixel) were used. By fitting a psychometric function to the proportion of “not compressed” judgments, we were able to estimate the approximate bit-rate at which the image first appeared uncompressed. This bit-rate was then converted to a DCTune Quality for that image.

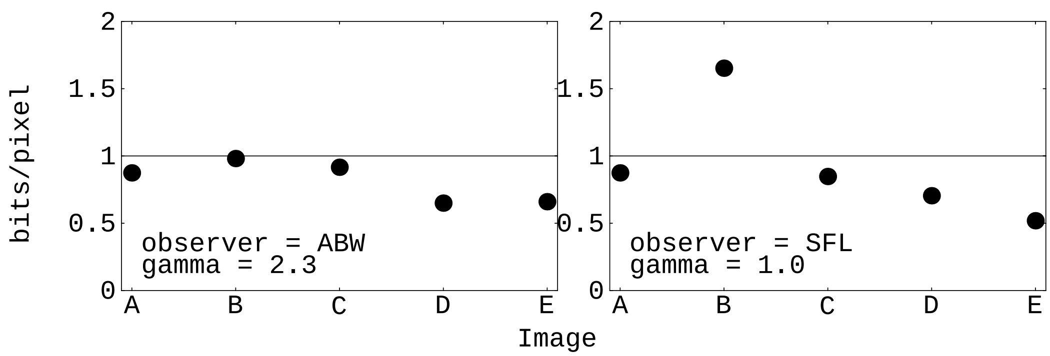

Figure 2. Thresholds for perceptually lossless compression, in units of DCTune Quality, for five images.

The main point of this figure is that, with a few exceptions, the values are close to 1, as the model predicts. The ability of image quality metrics to predict absolute quality levels is extremely important in many applications, from graphic arts, to electronic broadcast media, to medical imaging; consequently the success of this metric is heartening.

However, the images in this experiment were small, the observers were experienced, and the images quite familiar. In more recent experiments there are suggestions that for unfamiliar, larger images, or less experienced observers, the metric may overestimate the ability of observers to detect artifacts. In other words, the metric *underestimates* the quality of images. There are many possible reasons for this systematic shortfall; here we consider whether the model of masking incorporated in this metric, and most existing metrics, fully captures the essence of masking by complex patterned backgrounds.

3. MASKING

In the DCTune metric, as in most image quality metrics, masking occurs due to adjustment of gain by the background contrast. If we consider the scientific literature on visual masking, we can identify two distinct, nearly non-communicating traditions, which I will loosely call *contrast masking* and *noise masking*. Consider the four masks in Figure 3.

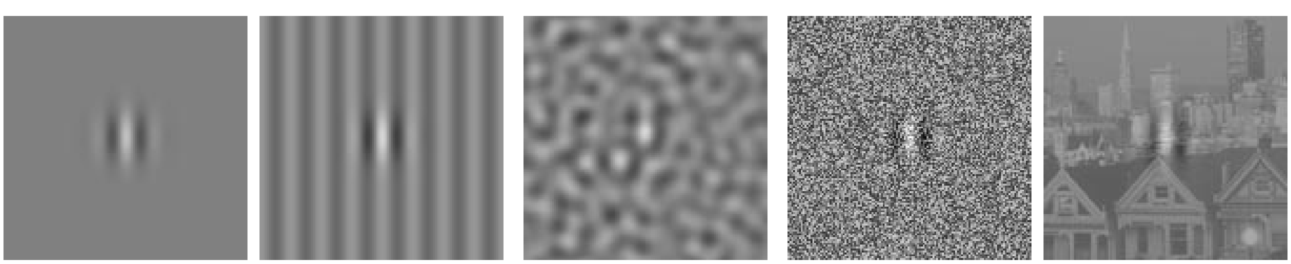

a b c d e

Figure 3. A Gabor target (a) added to four different backgrounds (b-e). Background contrasts are approximately those used in the experiment; the target contrast is 0.5.

In (d) is a uniform white noise background. Although it cannot be pictured, we may consider an additional case when this noise is dynamic in time. Traditionally, the threshold elevation produced by either the static or dynamic white noise has been interpreted as a form of *noise masking*. In this tradition the mask exerts its effect by increasing the variance of the decision variable, which might be the response of a single neuron or the result of cross-correlation with a particular template. In the case of the cosine background (b), the traditional explanation is a change in the differential gain of the relevant visual mechanisms, either because of a saturating transducer function or because of a contrast gain control process.

These two masks differ in two respects. The first is that one is random, and the other deterministic. As we shall show later, this difference may be less important than it seems. The second difference is that these two masks lie at either ends of a continuum in which the bandwidth of the mask is progressively narrowed. The puzzle is this: with which tradition shall we interpret the masking effect produced by the bandpass noise? Certainly current models of contrast gain control would predict that the bandpass noise should raise threshold, but should there not also be a contribution from "noise" masking? And what about the background in (e)? With which tradition, or combination of traditions, should we interpret its threshold elevating effect?

4. EXPERIMENTS

To clarify some of these issues we conducted a small series of experiments. They resemble in various respects experiments undertaken by others[4, 5], but were designed to focus in on a few key issues. All experiments consisted of measurement of contrast thresholds for a target added to a mask, using a two-interval forced choice procedure. Additional methodological details are provided in an appendix.

4.1. Target and Masks

The test stimulus in all cases is a Gabor function with a spatial frequency of 4 cycles/degree, an orientation of 0, and a frequency bandwidth at half height of 1 octave (Figure 3a). The experimental series employed five backgrounds, as illustrated in Figure 3. From left to right they are: a uniform field (no mask); a vertical cosine grating of 4 cycles/degree, contrast 0.172, and zero phase (a bright bar is centered on fixation); a sample of isotropic bandpass filtered noise with center frequency of 4 cycles/deg, one octave bandwidth, and contrast of 0.5; a sample of uniformly distributed white noise with a contrast of 0.595; and a natural image with a contrast of 0.309.

One goal of this series of experiments is to disentangle the effects of "noise" masking and contrast gain control. While details of contrast gain control models differ, most adjust the gain by a quantity that is close to the contrast energy of the mask. Thus equating contrast energy of various masks is a rough way of equating their effects upon the gain control system. To this end, contrasts for the cosine, bandpass noise, and image masks were selected (0.172, 0.5, and 0.309) that would equate their contrast energies (5,546 10^{-6} deg^2 seconds). For the white noise, a contrast that would yield the same contrast energy would not produce sufficient masking, so a contrast of 0.595 was used, which yields a contrast energy of 8 x 5,546 = 44,368 10^{-6} deg^2 seconds.

At this point we digress to introduce a new unit that we believe will prove useful in vision science. This unit, the Barlow, is the contrast energy of a visual target, in units of $degree^2$ seconds, multiplied by 10^6. Discussion of the many virtues of this unit is beyond the scope of this paper, but we note that it is a unit that incorporates all the dimensions of a contrast target (unlike contrast), and it has the attractive feature that the threshold for "what the eye sees best" is about 1 Barlow[6]. The Barlow has a decibel equivalent, dBB = 10 Log_{10}(Barlow), where "best" threshold is 0 dBB. Another useful attribute of dBB is that for a given target, it differs from dB by an additive constant. With the aid of these units, we can now describe our masks as having contrast energies of 37.44 dBB (cosine, bandpass, and image) or 46.47 dBB (white).

4.2. Masking Conditions

The experiments contained eight conditions, each characterized by a particular type of mask and presentation. The conditions and associated mnemonics were as follows.

None: The threshold contrast for the target was measured with no mask.

Cosine: The cosine mask was presented in both intervals of the forced-choice trial.

Random: A new sample of bandpass noise was used in each interval of each trial.

Twin: A new sample of bandpass noise was used on each trial, but the same sample was used for the two intervals.

Fixed: The same sample of noise was used in each interval of every trial.

White: A random sample of uniform white noise was used in each interval of every trial.

FixedWhite: A fixed sample of uniform white noise was used in each interval of every trial.

Image: The fixed "natural" image was used in each interval of every trial.

Where new samples of noise (bandpass or white) were used, they were obtained by random circular shifting of a single sample in horizontal and vertical dimensions. Because the bandpass noise was constructed by filtering in the discrete Fourier transform domain, it is periodic and circular shifts reveal no discontinuities.

4.3. Results

Mean thresholds for the three observers in the eight conditions, as well as group means, are shown in Figure 4 and Table 1. The results for each masking condition are considered in the following sections.

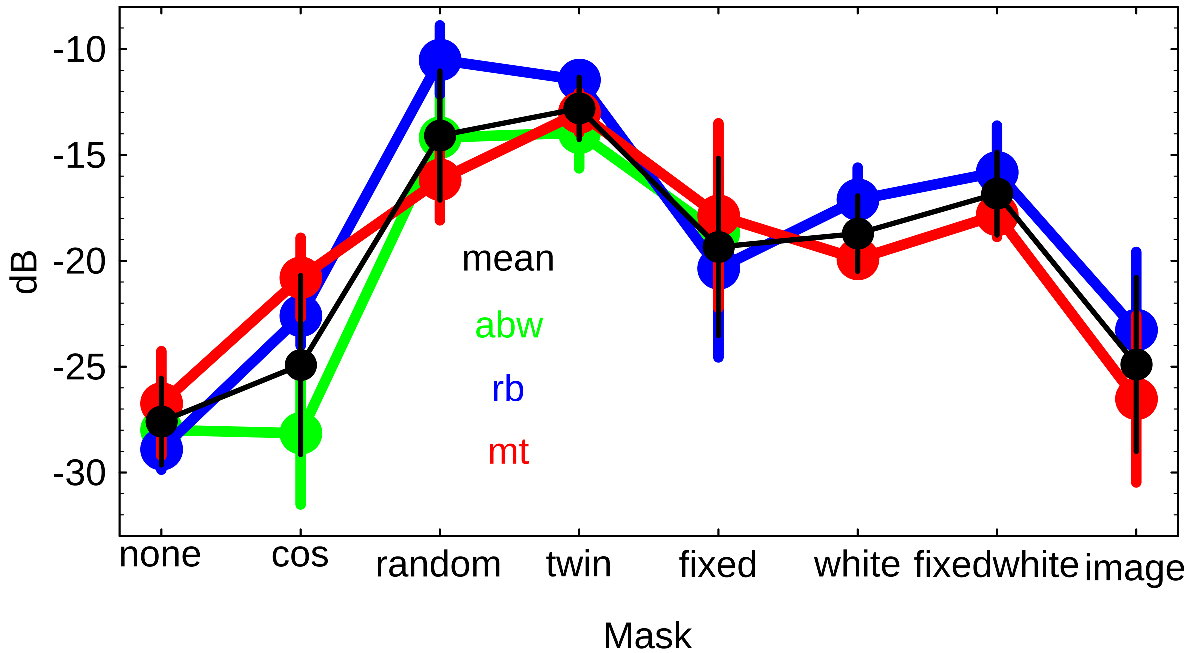

Figure 4. Contrast thresholds for detection of a Gabor target in the presence of eight different backgrounds. Data for three observers and their mean are shown.

	Condition							
	none	cos	random	twin	fixed	white	fixedwhite	image
threshold(dB)	-27.59	-24.92	-14.08	-12.80	-19.34	-18.71	-16.82	-24.89
elevation(dB)	0.	2.67	13.51	14.79	8.24	8.88	10.77	2.70

Table 1. Mean thresholds and threshold elevations for the eight experimental conditions. Elevations are relative to the no mask ("none") condition.

4.3.1. None

The threshold (mean = -27.59 dB, 7.07 dBB) is comparable to similar measurements in the literature[7]. The three observers are in good agreement.

4.3.2. Cosine

The cosine mask elevates threshold for two observers by about 6 dB, while for the third observer a surprising reduction in threshold was observed. The elevation is about that expected from prior studies of masking of a Gabor by a cosine[7]. As discussed above, the theoretical consensus is that this elevation is due to a contrast gain control process [7, 8, 9].

4.3.3. Random

Despite being equated for their gain control effect, the random noise yields a much larger threshold elevation (13.51 dB). This is of course not surprising. The randomness of the mask, quite apart from its contrast gain effects, would be expected to elevate threshold.

4.3.4. Twin

In the ideal observer model, a template matched to the signal, and adjusted for the spectral characteristics of the noise, is cross-correlated with the image in both intervals. That interval yielding the largest response is selected. Many other sub-optimal models share this template idea.

If the template model were correct, then adding the same sample of noise to both intervals would have no effect. It would merely add a constant to the response in both intervals. Yet the threshold elevation is essentially identical to that produced by a completely random noise. Similar results have been produced by other authors[4, 5, 10]. A tentative conclusion is that the masking in this case is due at least in part (possibly entirely) to what we will for the moment call "noise" masking. It cannot be due entirely to contrast gain control, because a) it is matched for gain control effect with the cosine background, which produces much less masking, and b) because whatever gain control effect it does induce, it should be the same as that induced by the random mask, which must in addition exhibit a noise masking effect.

In the context of modeling the observer's detection strategy, the presumption must be that the same strategy is used in both *random* and *twin* conditions. This excludes any model in which the observer subtracts the images in the two intervals, or cross-correlates with a template that is matched to the signal.

4.3.5. Fixed

Some further insight is provided by the fixed noise condition. Here performance is considerably (5.9 dB) better than in the *random* or *twin* conditions. However, this performance is only obtained after some experience with the fixed noise sample. Figure 5 shows how threshold is gradually reduced (with some reversals) as the observers gradually learn what to look for. Evidently, over the course of many trials with the same background, observers are able to reduce their threshold almost to the level of the no-mask condition, or at least to the level of the cosine mask of the same contrast energy. It appears that the learning process has ameliorated whatever masking capacity the fixed bandpass noise had, above its contrast gain control effect.

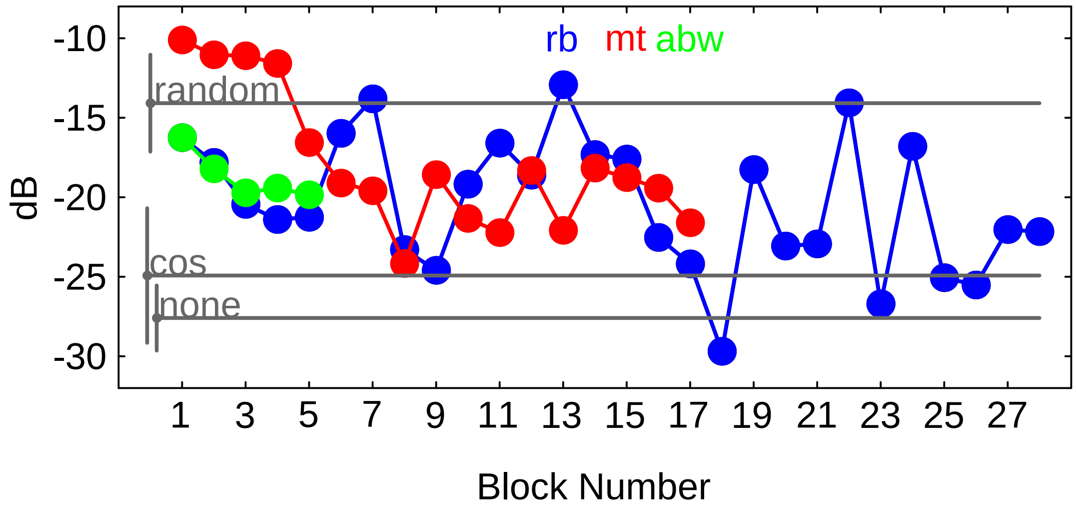

Figure 5. Contrast threshold in the *fixed* condition as a function of block number for three observers. The gray data points and horizontal lines are the means for the three observers for the *none, random,* and *cos* conditions.

4.3.6. White

Recall that the white noise mask had a contrast energy eight times that of the cosine or bandpass masks, in order to produce a substantial masking effect. The larger energy is presumably required because much of it is squandered in spectral regions not considered by the observers detection strategy. But consequently absolute comparisons between white and bandpass noises are not made here. Rather, we compare this threshold with that for the *fixedwhite* condition considered next.

4.3.7. FixedWhite

As shown in Figure 4 and Table 1, thresholds for white and fixedwhite conditions are essentially the same. Freezing the white noise does not appear to reduce its masking effect, unlike the case for the bandpass noise. Figure 6 shows the progression in performance over the course of seven blocks for two observers. There is no evidence of learning, though we cannot discount the possibility that further training might help. Others have found some learning of white noise after lengthy training[4]. We can at least conclude that learning of this mask is considerably more difficult than for the fixed bandpass noise.

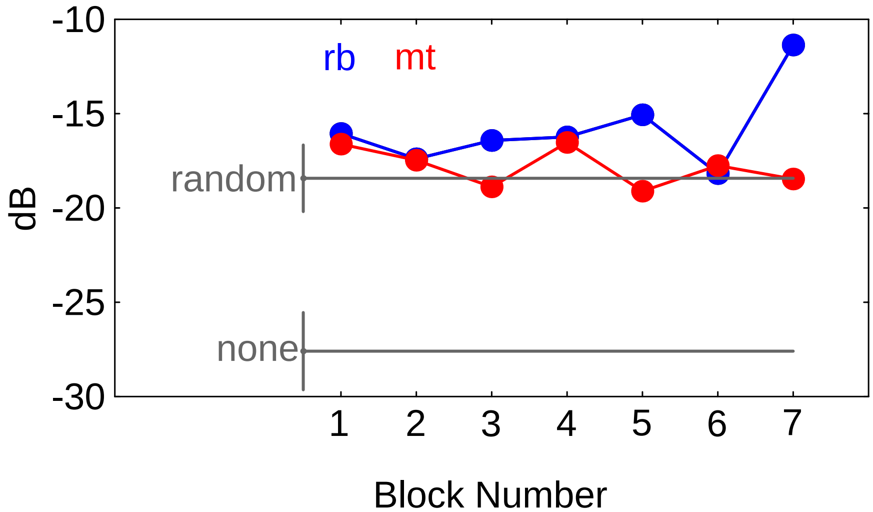

Figure 6. Contrast threshold as a function of block number for two observers. The gray data points and horizontal lines are the mean thresholds for the three observers for the *none* and *random* conditions.

4.3.8. Image

The final mask we used was a “natural” image, whose contrast energy was matched to the cosine and bandpass noises. The mean threshold here is not significantly different from the *cos* or *none* conditions. But here again, it is instructive to look at the trend over blocks. Starting from a threshold comparable to that for the random condition, thresholds decline rapidly, asymptoting in the region of the *cos* or *none* conditions. Here it appears that learning again obliterates the masking effect, and does so with great rapidity. We venture the observation that images differ in some property, akin to simplicity, that determines the ease with which they are learned, and with which this form of masking is obliterated.

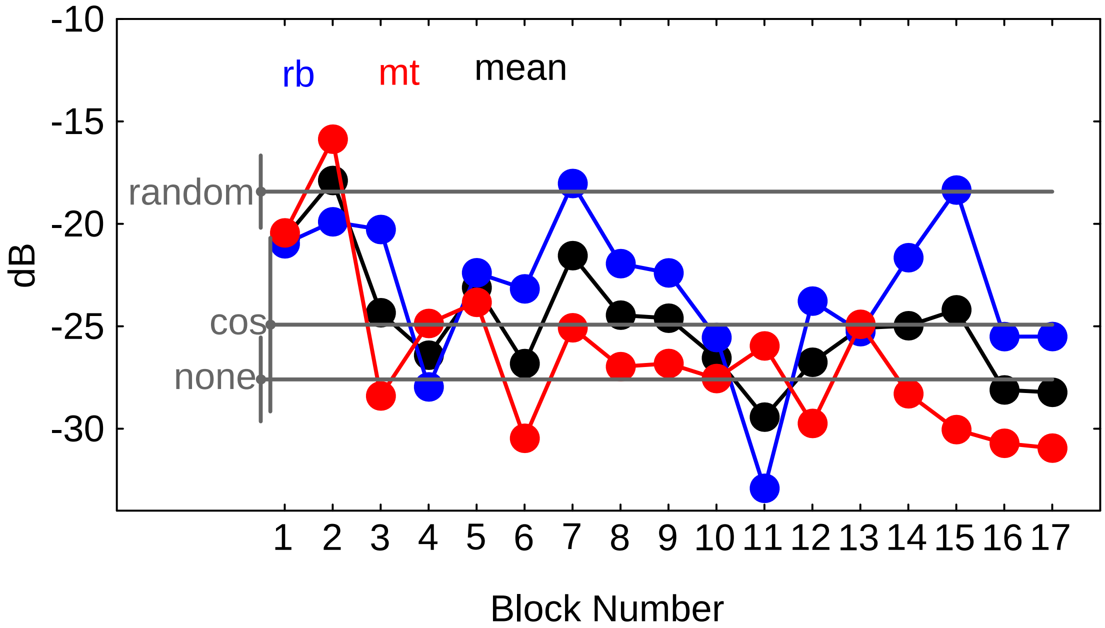

Figure 7. Contrast threshold for a Gabor target with an image mask, as a function of block number for two observers. The gray data points and horizontal lines are the mean thresholds for three observers for the *none, cos,* and *random* conditions.

5. DISCUSSION

5.1. Not Just Gain Control

The first observation is that fixed backgrounds do not elevate threshold exclusively through a gain-control process. Otherwise the cosine mask would produce as much masking as the bandpass noise. Likewise, if only a passive, deterministic gain control were at work, there would be no learning; but there is.

5.2. Not Just Randomness

If it is not gain control, by what process does the bandpass noise elevate threshold? A second observation is that it is not the randomness *per se* that is effective. The *random*, *twin*, and initial blocks of the *fixed* and *image* conditions all yield similar thresholds, though only the first would have an effect in conventional models of noise masking, that is, would increase the variance in the response of a template model.

5.3. No Discounting the Background

If it is neither traditional “noise masking nor gain control, what is the source of the threshold elevations for twin, fixed, and image conditions? The twin condition in particular suggests that the observer cannot “discount the background,” as supposed by traditional models of noise masking. In these models, the observer cross-correlates a template with the image in the two intervals, subtracts the first from the second, and chooses the second interval if the result is greater than zero. The subtraction has the effect of discounting any fixed noise in the two intervals.

If the observer is not behaving in this way, what are they doing? A precise theory is beyond the scope of this paper, but we can suggest some vague outlines. One is that as the observer learns a sample of fixed noise, they are developing something like a new template. But it is not a template matched to the signal; rather it is a template matched to the signal + noise. To emphasize, the observer cannot discount the background, so they must integrate the background into their template.

5.4. *Learnability = Simplicity*

The trends in threshold versus block number reveal the following: 1) the cosine shows no learning, and the elevation it produces is consistent with gain control; 2) the image is learned very rapidly; 3) the fixed bandpass noise is learned, but slowly, 4) the fixed white noise is either not learned, or learned very slowly. These results suggests a further principle: that the learnability and speed of learning of the background is a function of its simplicity. This would be a more powerful assertion if we had ready methods for computing the simplicity or complexity of an image. In the present case, we may point to the degrees of freedom of the image, as reflected in the number of significant Fourier coefficients, which is 2 for the cosine, about 99 for the image, about 476 for the bandpass noise, and 16384 for the fixed white noise.

5.5. *Entropy Masking*

Since we have seen that many backgrounds reduce visibility through a process distinct from noise masking or contrast masking, it may be useful to adopt a new term to describe this process. I propose the term “entropy masking”, to reflect the notion that the masking is a function of the degree to which the mask is unknown. Entropy is a measure of the information in a signal. Information is by definition that which we do not know. While formal measures of entropy may or may not prove useful in this context, the term seems to capture rather well the phenomenon at issue. As the observer learns a mask, its entropy declines, as does the amount of entropy masking. For the curious, the zero order entropies of the masks are: cos 3.125 bits; bandpass 7.02 bits; image 7.3 bits; white 7.97 bits. Of course, these do not take into account the higher order correlations in the image, or the knowledge of the observer.

It should be understood that we do not propose “entropy masking” as a mechanism, any more than “noise masking” is a mechanism. Rather it is intended to describe a property of the mask that, in concert with particular detection mechanisms and strategies, results in threshold elevation. These mechanisms and strategies remain to be described in detail, though we have already suggested that a key concept is the inability of the observer to discount the background.

It is possible that the concept of entropy masking should be generalized to incorporate the entropy of the signal as well as the background. In its parametric form, this traditionally goes by the name of signal uncertainty. If an observer is to develop a template, or more generally an algorithm, for discriminating target+background from mask alone, then ignorance of the target is as critical as ignorance of the background.

5.6. *Similarity Model*

To this point we have not directly addressed the question of what strategy the observer employs to identify the target interval in a forced-choice trial. From the twin condition we have concluded that it cannot be the ideal strategy of subtracting the template cross-correlations from each interval. We also concluded that the observer employs the same strategy for both random and twin conditions. What is this strategy?

We propose a *similarity* model as a candidate for further investigation. In this scheme, the observer computes a measure of similarity between a template or model and the stimulus received in each interval. The interval with the larger similarity is selected. Although this model shares some properties with the ideal, it predicts rather different behavior, especially in the twin condition. This is illustrated in Figure 8. In the present context, the critical feature of the similarity model is that, like human observers, it does not discount the background. And furthermore, if the mask is fixed and feedback is provided, we may suppose that the observers template migrates from $\mathbf{s}$ to $\mathbf{x}_1$, that is, a new template is learned that incorporates both signal and mask. We hope to study the further implications of this model, and how it accommodates learning phenomena, in future work.

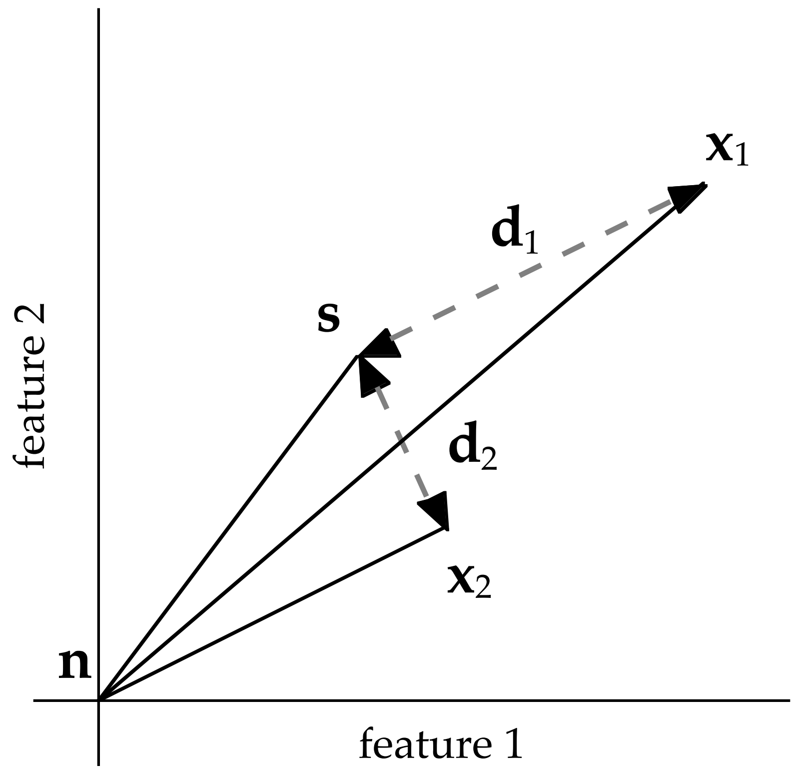

Figure 8. Illustration of behavior of the similarity model in the twin condition. The stimuli received in the two intervals are $\mathbf{x}_1$ and $\mathbf{x}_2$, which in turn consist of the signal (**s**) and null (**n**) plus the noise sample (equivalent to $\mathbf{x}_2$). Here the template is assumed to be equal to the signal **s**, and similarity is represented by Euclidean distance in a two-feature space. On this trial, the distance d_2 is smaller than d_1, so the observer selects the wrong interval. The ideal observer would compute the quantity $(\mathbf{x}_1 - \mathbf{x}_2) \cdot \mathbf{s}$, which would discount the twin noise sample.

5.7. An Ecological and a Historical Explanation

To conclude this portion of the discussion, we allow ourselves a brief speculation concerning why the observer cannot discount the background. Cross-correlation and subtraction are not complicated operations, and would seem to be well within the capacity of known neural mechanisms. One possible explanation is that, in the “natural” visual world, backgrounds are rarely combined additively with the target. More often, the combination rule is occlusion[11]. Under such circumstances, use of the similarity strategy may be most effective. The cross-correlation model for the ideal observer, which depends upon and assumption of additive noise, no doubt gained its dominance because in the realm of audition, in which signal detection theory made its first conquests, signals and noise are combined additively.

5.8. Implications for Image Quality

One impetus for this work was the desire to improve our models of image quality through an improved understanding of spatial masking. What are the implications of entropy masking for the design of image quality metrics? Should these metrics incorporate entropy masking, and if so, how should the masking effect be computed? It appears that the apparent quality of an image, in the sense of the invisibility of any artifacts, will be a function of the observers internal model of what the image *should* look like, which in turn will depend upon such things as their previous experience with the same or similar images. For applications where this internal model is faint, entropy masking should be incorporated. Where the model is highly developed, entropy masking should have little effect. On the second question, further research is clearly required, but simple measures such as the zero order entropy of the image provide a starting point.

At a minimum, however, those involved in the design of image quality metrics should be aware of the phenomenon of entropy masking. If a model is calibrated with simple, low entropy masks such as cosines, and then applied to complex high entropy masks such as a novel "natural" image, a failure of prediction can be expected to result.

6. APPENDIX: RELATION TO PREVIOUS WORK

Swift and Smith[5] measured thresholds for a vertical sinusoid in the presence of a fixed background consisting of the sum of eight vertical sinusoids, and found considerable learning effects. At the learning asymptote, they found that the function relating threshold to mask contrast had a slope (around 0.65) consistent with what we have called a gain-control process. This agrees with our notion that the entropy masking is removed by learning, leaving only the gain-control effect. They also conclude that the masking effect is due to ignorance of the background, though they do not acknowledge that this ignorance would pose no problem for the ideal observer.

Daly has considered the effect of learning on masking with reference to the results of Swift and Smith[12]. In the design of his image quality model, his solution has been to alter the exponent of the within-channel masking power function: 1 for no learning and 0.65 for highly learned patterns. He fixes this exponent at different values for different frequency bands. This solution clearly cannot take into account the image-specific knowledge that the observer may have, including learning of a specific mask.

7. APPENDIX: GENERAL METHODS

All experiments consisted of measuring contrast threshold for a target superimposed upon a background. A two-interval forced-choice procedure was used: the background appeared in both intervals; the target in only a random one; and the observer selected the interval that appeared to contain the target. Within a block of trials, the contrast of the background was fixed. From trial to trial the contrast of the target was varied using the QUEST staircase procedure[13]. The experiment was conducted using the Psychophysica software package[14]. Each experiment was divided into blocks of 32 trials, and a single threshold was estimated from each block by fitting a Weibull psychometric function[14, 15].

The target was always a Gabor function with a spatial frequency of 4 cycles/degree, an orientation of 0, and a bandwidth of one octave, corresponding to a scale of 0.352 degrees or 1.41 cycles. The background was either a zero contrast uniform field, a cosine with a frequency of 4 cycles/degree and an orientation of 0, a sample of isotropic bandpass filtered noise, a sample of uniformly distributed white noise, or a digital photographic image. The bandpass noise was created by filtering, in the DFT domain, a sample of uniform noise and a filter consisting of the convolution of a Gaussian and a ring impulse. In the *random* and *twin* conditions, the noise was shifted (with wrap-around) by random number of pixels horizontally and vertically. Because the noise was filtered in the DFT domain, it has a toroidal boundary and hence circular shifting does not reveal any discontinuities.

Stimuli were displayed using the Cinematica software[16], on an Apple Monochrome monitor. The display was linearized. The mean luminance was 24 cd m^{-2} A fixation point was present at all times except during the stimulus presentation. Stimuli were viewed binocularly with natural pupils from a distance of 118 cm, yielding a display visual resolution of 64 pixels/degree. Apart from the display, the room was dark. The contrast of both target and background varied over time as a Gaussian function with a scale of 8 frames and a total duration of 16 frames. The display frame rate was 60 Hz.

The three observers were abw, mt, and rb, all males with respective ages of 45, 18, and 29 years. All three were correctedmyopes.

8. ACKNOWLEDGMENTS

This work was supported by NASA Grant 199-06-12-39. I am grateful to Josh Solomon, Al Ahumada, and Miguel Eckstein for useful discussions, and to Cynthia Null for support and encouragement of the NASA Vision Group.

9. REFERENCES

1. A. Ahumada, Jr., "Computational Image Quality Metrics: A Review," Society for Information Display International Symposium, Digest of Technical Papers, 24, 305-308 (1993).
2. A.B. Watson, Digital images and human vision, MIT Press, Cambridge MA (1993).

SPIE Proceedings, vol. 3016, paper 1 (1997).

3. A.B. Watson, "DCTune perceptual optimization of compressed dental X-Rays," Medical Imaging, SPIE Proceedings, 3031, (1997).
4. A.J. Ahumada, Jr. and B.L. Beard, "Image Discrimination Models: Detection in Fixed and Random Noise," Human Vision, Visual Processing, and Digital Display VIII, 3016, (1997).
5. D.J. Swift and R.A. Smith, "Spatial frequency masking and Weber's Law," Vision Res., 23(495-506), (1983).
6. A.B. Watson, H.B. Barlow and J.G. Robson, "What does the eye see best?," Nature, 302(5907), 419-422 (1983).
7. J.M. Foley, "Human luminance pattern mechanisms: masking experiments require a new model," Journal of the Optical Society of America A, 11(6), 1710-1719 (1994).
8. A.B. Watson and J.A. Solomon, "A model of visual contrast gain control and pattern masking," Journal of the Optical Society A, in press, (1997).
9. P.C. Teo and D.J. Heeger, "Perceptual image distortion," ICIP-94, II, 982-986 (1994).
10. M.P. Eckstein, A.J. Ahumada, Jr. and A.B. Watson, "Image discrimination models predict human detection in medical images," Human Vision, Visual Processing, and Digital Display VIII, 3016, (1997).
11. J.B. Mulligan, "Nonlinear combination rules and the perception of motion transparency," Vision Research, 33(14), 2021-2030 (1992).
12. S. Daly, "The visible differences predictor: an algorithm for the assessment of image fidelity quality," Digital images and human vision, A. B. Watson, , MIT Press, Cambridge, MA (1993).
13. A.B. Watson and D.G. Pelli, "QUEST: A Bayesian adaptive psychometric method," Perception and Psychophysics, 33(2), 113-120 (1983).
14. A.B. Watson and J.A. Solomon, "Psychophysica: Mathematica notebooks for psychophysical experiments," Spatial Vision, in press, (1997).
15. A.B. Watson, "Probability summation over time," Vision Research, 19, 515-522 (1979).
16. J.A. Solomon and A.B. Watson, "Cinematica: A system for calibrated, Macintosh-driven displays from within Mathematica," Behavior Research Methods, Instruments, & Computers, 28(4), 607-610 (1996).