

Why use noise?

Denis G. Pelli

Psychology and Neural Science, New York University, New York, New York 10003

Bart Farell

Institute for Sensory Research, Syracuse University, Syracuse, New York 13244-5290

Received September 10, 1998; revised manuscript received November 3, 1998; accepted November 4, 1998

Measuring the dependence of visual sensitivity on parameters of the visual stimulus is a mainstay of vision science. However, it is not widely appreciated that visual sensitivity is a product of two factors that are each invariant with respect to many properties of the stimulus and task. By estimating these two factors, one can isolate visual processes more easily than by using sensitivity measures alone. The underlying idea is that noise limits all forms of communication, including vision. As an empirical matter, it is often useful to measure the human observer's threshold with and without a noise background added to the display, to disentangle the observer's ability from the observer's intrinsic noise. And when we know how much noise there is, it is often useful to calculate ideal performance of the task at hand, as a benchmark for human performance. This strips away the intrinsic difficulty of the task to reveal a pure measure of human ability. Here we show how to do the factoring of sensitivity into efficiency and equivalent noise, and we document the invariances of the two factors. © 1999 Optical Society of America [S0740-3232(99)01703-2]

OCIS codes: 330.1800, 330.7310, 330.4060, 330.5510.

1. INTRODUCTION

Traditionally, objective studies of perception have measured and explained threshold contrast of a target on a blank background: the contrast required for an arbitrarily selected level of performance. Indeed, most of what we know best about visual processing has come from such studies.¹ However, it is not widely appreciated that visual sensitivity is a product of two factors.² By measuring an additional threshold, on a background of visual noise, one can partition visual sensitivity into two components representing the observer's efficiency and equivalent noise. Although they require an extra threshold measurement, these factors turn out to be invariant with respect to many visual parameters and are thus more easily characterized and understood than their product, the traditional contrast threshold.

Previous authors have presented compelling theoretical reasons for isolating these two quantities in order to understand particular aspects of visual function.³⁻¹³ Here we ignore theory, to focus on the empirical properties of the two factors, especially their remarkable invariances, which make them more useful than sensitivity. When one factor is invariant with respect to a parameter that affects sensitivity, then, of course, the other factor varies proportionally with sensitivity. The surprise is that one or the other of the two factors is invariant with respect to many of the parameters that affect sensitivity. This complementarity is what makes the factoring so useful, confining the explanation to one or the other of the two factors.

A. Energy

For reasons that will become clear shortly, after measuring the threshold contrast we generally convert to energy

units. Contrast energy E is the square of the contrast function summed over the dimensions along which the stimulus varies. For static two-dimensional stimuli (as in Fig. 1 below), signal energy is integrated over space:

$$E = \iint c^2(x, y) dx dy. \quad (1)$$

The contrast function is the normalized deviation of the luminance function from the background level,

$$c(x, y) = [L(x, y) - L_b]/L_b.$$

For letters, energy is the product of "ink" area and squared contrast.

B. Using Noise

Studies of sensitivity measure threshold on a blank field. Here we represent the threshold contrast energy on a blank background by E_0 . We advocate also measuring a second threshold E , on a background of white noise,¹⁴ the stronger the better. Further, we advocate calculating the threshold E_{ideal} of the ideal observer for the same task on the same noise background. Knowing E_0 , E , and E_{ideal} , instead of just E_0 , the experimenter is much better informed and able, as we will see below, to distinguish (and reject) whole classes of explanation. From E_0 , E , and E_{ideal} we particularly recommend calculating high-noise efficiency η^* and equivalent input noise N_{eq} . Efficiency and equivalent input noise have interesting histories^{8,16} and useful theoretical and empirical properties: They capture the two degrees of freedom in the data (E_0 and E) and are the factors of sensitivity.

C. Signal-to-Noise Ratio

We introduce two minor extensions to the usual energy and noise notation. First, lacking any widely accepted symbol for signal-to-noise ratio E/N , we introduce D :

$$D = E/N. \quad (2)$$

Second, in separating the effects of stimulus noise from those of the observer's equivalent noise, the threshold E is less important than the threshold elevation E^* produced by the stimulus noise:

$$E^* = E - E_0. \quad (3)$$

We similarly apply the same asterisk to derived quantities when we substitute threshold elevation E^* for threshold E , e.g., $D^* = E^*/N$.

D. Measuring Threshold: E_0

The value of measuring the ordinary threshold contrast on a blank background remains. After converting it to energy units, we call it E_0 .

Visual thresholds are typically plotted as log threshold versus log X , where X is a parameter of the stimulus, e.g., spatial frequency or background luminance. There is a long tradition of looking for straight sections of such plots and devising an explanation (a model) for that region, e.g., the Rose-DeVries law and Weber law regimes of the threshold-versus-luminance plot, and Ricco's and Bloch's laws for threshold versus size and duration.¹⁷ By measuring threshold again, in noise as explained next, we can streamline this find-the-line approach, arriving at such explanations more readily.

E. Measuring Threshold in Noise: E

The approach is based on the long-standing observation that visual thresholds are elevated in the presence of white noise added to the display. Threshold contrast energy E is linearly related to the displayed noise power spectral density N :

$$E = D^*(N + N_{\text{eq}}), \quad (4)$$

where D^* and N_{eq} are fitted constants. We are not aware of any exceptions to this result for white noise (see Ref. 8 for review; narrow-band masks can produce other results^{18,19}). A linear relationship, a line, is determined by two points, so we can fix the line by measuring two thresholds. Usually the experimenter will measure the traditional threshold E_0 on a blank background and a second threshold E at a high level of noise N —the higher the better, but at least high enough to double threshold, $E > 2E_0$. In Eq. (4), the slope

$$D^* = \frac{E}{N + N_{\text{eq}}} = \frac{E - E_0}{N} = \frac{E^*}{N} \quad (5)$$

and the offset

$$N_{\text{eq}} = \frac{E_0}{E - E_0} N \quad (6)$$

specify the line in a handy way. N_{eq} is the observer's equivalent input noise, which contributes to the effective stimulus, adding to the displayed noise N . D^* is the effective signal-to-noise ratio (i.e., the signal-to-noise ratio

of the effective stimulus) at threshold. For performance on a blank background, the observer is the only noise source, and the effective signal-to-noise ratio is just E_0/N_{eq} .

F. Calculating Ideal Threshold in Noise: E_{ideal}

For any task in which the displayed noise N poses a theoretical limit to attainable performance, it is relevant to ask what threshold the (mathematically defined) ideal observer has. (The ideal makes decisions based on the available information, so as to maximize the probability of a correct response.) When we can calculate the ideal observer's threshold E_{ideal} for the task, we prefer to normalize the human observer's threshold by the ideal's. This gives us the human observer's high-noise efficiency $\eta^* = E_{\text{ideal}}/E^*$.

A word about notation: Tanner and Birdsall⁴ introduced the quantities d' and η . The square of d' is the signal-to-noise ratio $d'^2 = D_{\text{ideal}} = E_{\text{ideal}}/N$ required by an ideal observer to perform as well as the observer under study.²⁰ Efficiency η is the ratio of threshold energies of ideal and human observers $\eta = E_{\text{ideal}}/E$. Using our asterisk again, we introduce high-noise efficiency

$$\eta^* = \frac{E_{\text{ideal}}}{E^*} = \frac{D_{\text{ideal}}}{D^*} = \eta \frac{N + N_{\text{eq}}}{N}, \quad (7)$$

which is the asymptotic value of η when measured in high noise, $N \gg N_{\text{eq}}$. In practice, the nontrivial rms error of threshold estimation makes the two flavors of efficiency, η and η^* , indistinguishable once the noise is high enough to raise threshold at least fourfold, $E > 4E_0$. Like D^* , η^* discounts the no-noise threshold, considering only the threshold elevation $E^* = E - E_0$ produced by the added noise. η^* has previously been called "central," "sampling," and "calculation" efficiency.^{6-8,21} We introduce the new name, high-noise efficiency, because the old names have a much more restricted domain of validity than the measure itself; i.e., the old names are misleadingly limited.

For some very simple but popular tasks, like two-interval forced-choice detection of a known signal in white noise, the ideal signal-to-noise ratio at threshold $D_{\text{ideal}} = d'^2$ can be looked up in a table,²² and the ideal threshold $E_{\text{ideal}} = D_{\text{ideal}}N$ can then be calculated for the actual noise level. For more complex tasks, like letter identification, the analytic solution is nontrivial, and it is easier to implement the ideal as a computer program whose threshold E_{ideal} can be measured by the same procedure used to test the human observer.²³ When E_{ideal} cannot be calculated, the effective signal-to-noise ratio D^* can be used instead of efficiency, with some of the same virtues. We touch on this again in Appendix A. This review considers only luminance noise (static and dynamic), which has proved very useful in the study of the spatiotemporal mechanisms mediating pattern perception, but the same arguments apply to noise in other dimensions, e.g., chromaticity,²⁴ numerosity,²⁵ and disparity.²⁶

Having D_{ideal} , we use the identity $\eta^* = D_{\text{ideal}}/D^*$ to rewrite the empirical finding of linearity [Eq. (4)] as

$$E = \frac{D_{\text{ideal}}}{\eta^*} (N + N_{\text{eq}}). \quad (8)$$

2. EMPIRICAL PROPERTIES OF THE TWO FACTORS

In a nutshell, we measure E_0 and E , find E_{ideal} , and then calculate the equivalent input noise N_{eq} [see Eq. (6)] and the (high-noise) efficiency

$$\eta^* = \frac{E_{ideal}}{E - E_0}. \tag{9}$$

(We drop the prefix “high-noise” from here on.) Efficiency η^* rates the computation underlying our perceptual decisions on the absolute performance scale defined by the ideal observer, while equivalent noise N_{eq} specifies how much noise the observer’s visual system (including transduction) adds to the display. As noted above, there are compelling theoretical reasons to isolate these quantities, but here we invite the reader to be an agnostic empiricist, ignore the theory, and simply consider the merits of the new quantities N_{eq} and η^* as a convenient coordinate frame in which to analyze and report experimental results.

To illustrate this approach, let us examine the effects of size on letter visibility. Figure 1(a) shows letters of variable size and contrast on a blank background. The faintest visible letters trace out the reader’s threshold (or sen-

sitivity) as a function of letter size. Thresholds are plotted as open symbols in Fig. 1(d). There are two limbs. For letters smaller than 1 deg, the slope approaches -1 : Threshold is inversely proportional to size. For letters larger than 1 deg, the threshold curve is very shallow, approaching a slope of 0.18.

Figure 1(b) shows the effect of adding white noise. The noise consists of independently generated square checks; the power spectral density N of the noise equals the product of contrast power c_{rms}^2 and check area. Although the power spectral density is constant only up to frequencies of approximately half a cycle per check, we call the noise “white” because that frequency is sufficiently high that the noise power spectral density is constant over all frequencies to which the mechanism under study is sensitive.^{15,27} The noise elevates threshold more for smaller letters, so the smallest letters are invisible (threshold contrast higher than 1), while the largest letter is hardly affected [compare with Fig. 1(a)]. Noise can be powerful stuff: Adding too much raises threshold unmeasurably high, but adding too little will produce a uselessly small threshold elevation. To be most informative, the noise-masked threshold must be several times the unmasked threshold, yet still measurable.

An easy way to achieve this is to scale the noise pattern

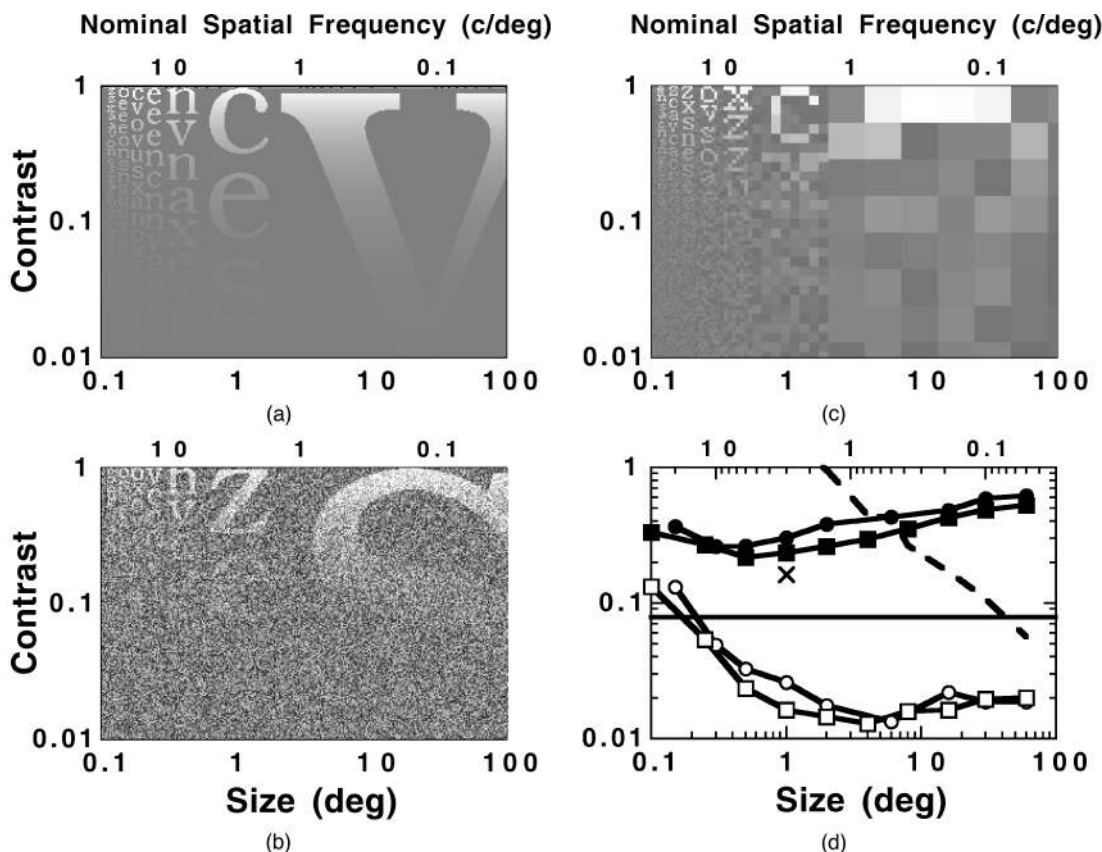


Fig. 1. (a) Letters. The vertical scale indicates the contrast of the letter. The bottom horizontal scale indicates the letter size, assuming a viewing distance of 48 cm. The top horizontal scale indicates nominal spatial frequency, in cycles per degree (c/deg), assuming 3 cycles per letter.²⁷ The actual experiment was similar to this demonstration but showed only one letter at a time, randomly chosen from the entire alphabet. (b) Letters in noise. The noise power spectral density N is $2 \times 10^{-4} \text{ deg}^2$. (c) Letters in noise. The noise is scaled with the signal. The largest letter is an “a.” (d) Threshold contrast for 64%-correct identification, as a function of letter size, with no noise (open symbols), unscaled white noise (dashed curve), and scaled white noise (filled symbols). Circles, data for observer WT; squares, data for observer DM. The horizontal line is the ideal observer’s threshold with scaled noise. The \times is DM’s threshold for 82%-correct letter detection of 1° letters in noise.

with the signal, as shown in Fig. 1(c). (In laboratory testing we display one letter at a time, at a fixed physical size, and merely adjust viewing distance to change the visual size of signal and noise together.) Since the noise power spectral density is proportional to check area,⁸ halving the viewing distance quadruples the power spectral density.

Figure 1(d) shows 64%-correct thresholds for identifying static letters, with and without noise, as a function of letter size. (See Ref. 23 for methods.) The upper horizontal axis is the nominal spatial frequency of the letter, based on the finding that the observer uses the same channel to identify a 1-deg letter and detect a 3-c/deg grating,²⁷ assuming that the channel frequency is inversely proportional to letter size. (This assumption is not quite right, but that's another story.²⁸) The no-noise curve (open symbols) corresponds to Fig. 1(a) and is very much like a traditional contrast sensitivity curve for gratings, rotated 180°. It is rotated because we are plotting threshold versus size instead of 1/threshold (sensitivity) versus 1/size (spatial frequency). The dashed curve shows the effect of adding noise with a fixed power spectral density [as in Fig. 1(b)]; most of the thresholds are outside the usefully measurable range. As shown by the filled symbols, adding scaled noise [as in Fig. 1(c)] yields useful thresholds; it raises threshold but keeps the task doable.

The horizontal line in Fig. 1(d) represents threshold in noise for the ideal observer, measured by the same procedure as for the human observers. The ideal observer makes maximum-likelihood choices based on the same stimulus information as provided to the human observer. It is mathematically defined and is typically implemented as a computer program. Its threshold is the lowest possible for the task.

Figure 2(a) takes the threshold contrasts of Fig. 1(d) and replots them as contrast energy. As before, the open symbols represent threshold measured on a blank background, and the filled symbols represent threshold measured in scaled noise. The remaining graphs of Fig. 2, which plot N_{eq} and η^* , are derived from the energy thresholds plotted here in Fig. 2(a).

Figure 2(b) plots the observers' equivalent noise N_{eq} , computed by Eq. (6) from the data in Fig. 2(a). This shows that the observers' N_{eq} increases with size. Recall that, without noise, threshold contrast changes only slightly, and nonmonotonically, over this range [Fig. 1(d)]. Extensive parametric equivalent noise studies have shown what this N_{eq} curve is made of.²⁹ The curve has two limbs. Along the horizontal limb, N_{eq} is photon noise, which is inversely proportional to luminance, and, along the rising limb, N_{eq} is neural noise arising in the visual cortex, i.e., after the two eyes' signals have been combined. Note that the equivalent noise is independent of task. Back in Fig. 2(a) the thresholds are very different for the binary-decision task of detecting 1-deg letters (×) and the 26-way decision task of identifying 1-deg letters (squares and circles), but here we can see that both tasks yield similar estimates of the equivalent noise level. Raghavan²⁹ found the equivalent noise estimates based on identifying 2 gratings, 2 or 26 letters, or 2000 words to be identical.

Figure 2(c) plots efficiency η^* , computed by Eq. (9) from Fig. 2(a). These data also appear in Pelli *et al.*²³ As noted there, for letters of 0.5 deg and larger (as large as 60 deg), there is a shallow negative log-log slope of

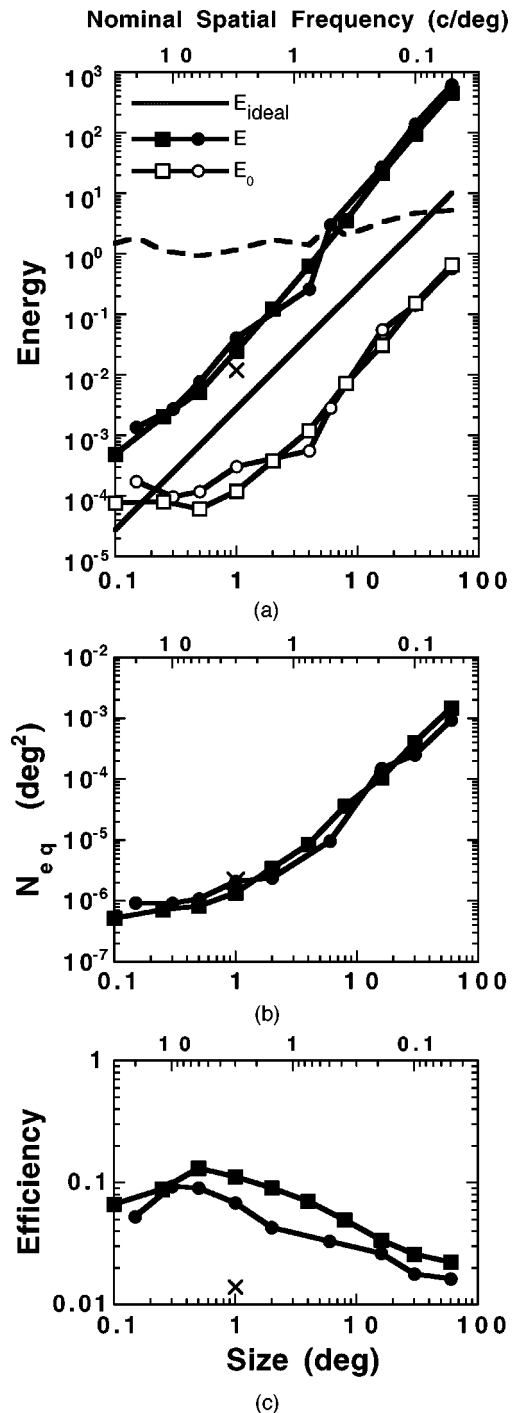


Fig. 2. (a) Threshold energy for identification as a function of letter size, replotted from Fig. 1(d). Energy equals squared contrast times letter area. The × is DM's threshold for detection. (b) Equivalent input noise of the observer, as a function of letter size, computed by Eq. (6) from the data shown in (a). Note that the equivalent noise for detection (×) is about the same as that for identification (filled symbols). (c) Efficiency η^* as a function of letter size, computed by Eq. (9) from the data shown in (a). The efficiency for detection (×) is one tenth that for identification (filled symbols).

-0.35 , indicating that we see small (0.5-deg) letters best. Below 0.5 deg, efficiency drops as the observers approach their acuity limit of 0.1 deg. The mere fivefold variation of efficiency in Fig. 2(c) is very small, considering the 600:1 size range and the $600^2 = 360,000:1$ area range of the letters.³⁰

Not shown here are experimental results revealing the invariance of letter identification efficiency with respect to most of the other visual parameters studied. Efficiency is independent of duration (60–4000 ms), contrast, and eccentricity of the letter.²³ That viewing conditions have so little effect on efficiency suggests that the observer's decision process is similar across wide variations in stimulus presentation. We were surprised that efficiency should be independent of eccentricity, since it is well known that threshold without noise grows with eccentricity unless the signal is scaled to compensate for the falling cortical magnification.³¹ The efficiency results indicate that eccentricity and the cortical magnification factor affect only the equivalent input noise, not efficiency, and that, judging by efficiency, tasks at different eccentricities are performed most similarly when the signals are the same size. Note that efficiency *is* sensitive to changes in the nature of the computations demanded by the task, varying tenfold across alphabets²² and between detection and identification.

In returning to our goal of explaining sensitivity, i.e., the two limbs of the no-noise threshold curve in Fig. 1(d), how can we explain the steep rise in threshold for small letters and the shallow rise for large letters? As we promised at the outset, we can now partition sensitivity into two empirical factors. At zero noise, Eq. (4) reduces to the product of signal-to-noise ratio and equivalent noise:

$$E_0 = D^* N_{\text{eq}}, \quad (10)$$

and, similarly, Eq. (8) reduces to the product of reciprocal efficiency and equivalent noise:

$$E_0 = \frac{D_{\text{ideal}}}{\eta^*} N_{\text{eq}}. \quad (11)$$

The extra scalar in Eq. (11), the ideal signal-to-noise ratio at threshold D_{ideal} , is independent of most experimental parameters (e.g., letter size), except for the kind of task (e.g., detection versus identification) and the threshold criterion (e.g., 64% versus 82% correct). Of the two empirical factors, efficiency η^* is too flat a function of letter size [Fig. 2(c)] to explain much of the 10^4 effect of size on threshold E_0 in Fig. 2(a). However, N_{eq} [Fig. 2(b)], like threshold, has two distinct limbs, with a break at 1 deg. For smaller letters the log-log slope of N_{eq} versus size approaches 0 (i.e., signal independent), and for larger letters it approaches 2 (i.e., scales with the signal area). As noted above, Raghavan²⁹ showed that the zero-slope section of the equivalent noise curve is dominated by photon noise and that the +2-slope section is dominated by cortical noise (i.e., noise arising after binocular combination of the two eyes). Because efficiency is nearly constant, the zero-slope limb of N_{eq} (small letters) yields zero slope of threshold energy [Fig. 2(a)] and a -1 slope of threshold contrast [Fig. 1(d)]. [Recall that energy is the product of area and squared contrast [Eq. (1)].] If efficiency were

perfectly constant, the +2-slope limb of N_{eq} in Fig. 2(b) (large letters) would yield a +2-slope limb of threshold energy and a zero-slope limb of threshold contrast. In fact, for larger letters the efficiency curve has a slope of -0.35 , resulting in a +2.35 slope of threshold energy and a $0.35/2 = 0.18$ slope of threshold contrast.

Traditionally, one would measure only the no-noise threshold [open symbols in Fig. 1(d) or 2(a)]. Previous investigators have derived valuable insights from similar sensitivity curves. For example, Banks *et al.*¹⁰ noted that the small-signal limb has the same -1 slope as the ideal observer limited by photon noise, suggesting that our ability to see small signals is limited by photon noise and showing that we have constant overall efficiency η for small signals. However, the large-signal limb remained unexplained: Why is contrast sensitivity so nearly constant as letters grow from 1 to 60 deg? Measuring a threshold in noise helps answer the question.

For large letters we uncover two effects. The more profound effect is that the observer's equivalent noise scales with the signal [Fig. 2(b)], suggesting that it arises in a part (or parts) of the visual pathway that is selective for the signal scale. The smaller, but no less intriguing, effect is that efficiency is highest for 0.5-deg letters and gradually drops as letter size is increased [Fig. 2(c)]; we had expected scale invariance. (Majaj *et al.*²⁸ follow up this result, finding other, more dramatic, ways in which letter identification depends on size.)

The small-letter results confirm the conclusion of Banks *et al.*¹⁰ However, Banks *et al.* noted that they could get their result (same slope of human and ideal) only if all their signals were geometrically similar. By adding noise and factoring sensitivity into efficiency and equivalent noise, Raghavan²⁹ found that equivalent noise is independent of signal extent and kind (e.g., same for 1-deg letters and 3-c/deg gratings), which makes its assignment to photon noise easier and more secure.

The constant overall efficiency for grating detection (at high spatial frequencies) found by Banks *et al.*¹⁰ is a tiny 0.2%, whereas for letters Fig. 2(c) reports a respectable high-noise efficiency of close to 10%. Most of the inefficiency in the overall result is due to equivalent noise, which Banks *et al.* could not separately estimate because they did not add noise.

Measuring a threshold in noise also helped in characterizing the effect of eccentricity. The evidence that eccentricity affects equivalent noise, not efficiency (except near the acuity limit) rejects the whole class of explanations that are based on scaling everything by the cortical magnification factor. [To be fair to the proponents of that theory, the equivalent noise does increase with eccentricity in a way that is predicted by ganglion cell density (or cortical magnification factor),²⁹ but this is a different kind of theory.]

3. CONCLUSION

Even when the experimenter's immediate goal is to explain sensitivity, it may be easier, and thus quicker, to achieve this by adding a measurement with noise, to partition the problem into two factors that can be explained separately. Equivalent noise, being independent of task,

invites explanation in terms of known properties of visual neurons: their density, gain, variance, and physiological thresholds. Efficiency, being largely independent of viewing conditions, invites explanation in terms of the computation that combines the distributed stimulus and prior information to yield a decision. The classical paradigm of explaining the straight segments in plots of log threshold versus log X remains fruitful, but the slopes are now immediately interpretable. Zero slope (at peak efficiency) indicates a computation that scales with the stimulus parameter, while rising and falling slopes of ± 1 indicate fixed-size computations.

APPENDIX A: FACTORING SENSITIVITY BY ADDING NOISE

In practice, the ideas presented here can play out in four different ways, depending on whether the experiment measures thresholds with and without noise and whether the task is one for which it is practical to compute an ideal threshold. For many questions it is possible to define tasks for which we can measure human performance and calculate ideal performance, but there are some questions for which a subjective criterion is essential (e.g., a unique hue) so that we cannot meaningfully specify what ideal performance would be. And there are some important practical tasks, such as diagnostic radiology, for which it is difficult to calculate ideal performance.³²

1. One Threshold

If the experimenter measures only one threshold, then the result will have only one degree of freedom. Cases *a* and *b* distinguish whether we can calculate the ideal threshold:

Case a. Measure E_0 . Plot E_0 . Most of the classic vision literature follows this paradigm, e.g., the effects of area (Ricco), duration (Bloch), or luminance (Weber and Rose–DeVries); the effects of all three¹⁶; or the effects of spatiotemporal frequency.³³

Case b. Measure E , and calculate E_{ideal} . Plot overall efficiency η , where

$$\eta = \frac{E_{\text{ideal}}}{E}. \quad (\text{A1})$$

This case arises when the noise is known but not directly under experimental control, so there is no zero-noise threshold. Examples include Barlow's^{34,35} quantum efficiency measurements, the studies by Tanner and Swets³⁶ and by Banks *et al.*,¹⁰ and other studies reviewed by Geisler.¹¹

2. Two Thresholds

If the experimenter measures threshold with and without noise, then the results will have two degrees of freedom. Again, cases *a* and *b* distinguish whether we can calculate E_{ideal} :

Case a. Measure E_0 and E . Factor sensitivity, $E_0 = D^* N_{\text{eq}}$, and plot the factors D^* and N_{eq} , where

$$D^* = \frac{E - E_0}{N}, \quad (\text{A2})$$

$$N_{\text{eq}} = \frac{E_0}{E - E_0} N. \quad (\text{A3})$$

Examples in the literature of this approach include the work of Rose,^{3,37} Nagaraja,³⁸ van Meeteren and Boogaard,³⁹ and Engstrom.⁴⁰ As it happens, most of these studies used objective tasks, for which they could have computed the ideal threshold, but they did not.

Case b. Measure E_0 and E , and calculate E_{ideal} . Factor sensitivity, $E_0 = (D_{\text{ideal}}/\eta^*)N_{\text{eq}}$, where $D_{\text{ideal}} = E_{\text{ideal}}/N$, and plot the two factors η^* and N_{eq} :

$$\eta^* = \frac{D_{\text{ideal}}}{D^*} = \frac{E_{\text{ideal}}}{E - E_0}, \quad (\text{A4})$$

$$N_{\text{eq}} = \frac{E_0}{E - E_0} N. \quad (\text{A5})$$

It seems that Burgess *et al.*⁶ and Pelli^{7,8} were the first to determine all three thresholds for the same task, though Barlow²¹ had already made it clear that this would allow a partitioning of sensitivity into two factors. This powerful technique has been applied to contrast discrimination⁹ and to detecting and identifying patterns,⁴¹ letters,²³ and solid objects.⁴²

ACKNOWLEDGMENTS

We thank Bill Geisler for helpful comments on an early draft, which led us to refine the terms and symbols to allow description of all four cases in Appendix A without conflicting with existing notation. We also thank Debbie Moore for running some of these experiments and for preparing the graphs. Some of the results shown in Fig. 2 also appear in Pelli *et al.*²³ This work was supported by National Institutes of Health grant EY04432 (to D. Pelli).

The authors can be reached by e-mail at denis@psych.nyu.edu and at bart_farell@isr.syr.edu, respectively.

REFERENCES AND NOTES

1. For examples, see T. N. Cornsweet, *Visual Perception* (Academic, New York, 1970).
2. We are using the terms "factor" and "product" loosely, referring to both multipliers and divisors.
3. A. Rose, "The sensitivity performance of the human eye on an absolute scale," *J. Opt. Soc. Am.* **38**, 196–208 (1948).
4. W. P. Tanner, Jr. and T. G. Birdsall, "Definitions of d' and η as psychophysical measures," *J. Acoust. Soc. Am.* **30**, 922–928 (1958).
5. A. van Meeteren and H. B. Barlow, "The statistical efficiency for detecting sinusoidal modulation of average dot density in random figures," *Vision Res.* **21**, 765–777 (1981).
6. A. E. Burgess, R. F. Wagner, R. J. Jennings, and H. B. Barlow, "Efficiency of human visual signal discrimination," *Science* **214**, 93–94 (1981).
7. D. G. Pelli, "Effects of visual noise," Ph.D. dissertation (Cambridge University, Cambridge, UK, 1981).
8. D. G. Pelli, "The quantum efficiency of vision," in *Vision: Coding and Efficiency*, C. Blakemore, ed. (Cambridge U. Press, Cambridge, UK, 1990), pp. 3–24.
9. G. E. Legge, D. Kersten, and A. E. Burgess, "Contrast discrimination in noise," *J. Opt. Soc. Am. A* **4**, 391–404 (1987).
10. M. S. Banks, W. S. Geisler, and P. J. Bennett, "The physical limits of grating visibility," *Vision Res.* **27**, 1915–1924 (1987).

11. W. S. Geisler, "Sequential ideal-observer analysis of visual discriminations," *Psychol. Rev.* **96**, 267–314 (1989).
12. A. J. Ahumada, Jr., "Putting the visual system noise back in the picture," *J. Opt. Soc. Am. A* **4**, 2372–2378 (1987).
13. H. B. Barlow and D. G. Pelli, eds., special issue on the statistical efficiency of natural and artificial vision, *J. Opt. Soc. Am. A* **4**(12) (1987) and **5**(4) (1988), Pts. 1 and 2, respectively.
14. White noise is indistinguishable, by the system under study, from noise whose samples are all stochastically independent. In practice, the noise samples (checks) are usually created independently, and it is enough to make sure the checks are too small to be resolved. Thus the power spectral density is constant over the range of frequencies that affect the system under study. Typically one achieves this by displaying a random checkerboard, each cell randomly black or white, or sampled from a truncated Gaussian distribution, with checks no bigger than one quarter of the period of the grating to be detected (Ref. 15), since detection of the grating is mediated by an octave-wide channel. When the mediating mechanism is unknown, the relevant band is still restricted by the visual optics. Checks finer than 2 per cycle of the optical cutoff frequency will produce white noise.
15. H. Kukkonen, J. Rovamo, and R. Näsänen, "Masking potency and whiteness of noise at various noise check sizes," *Invest. Ophthalmol. Visual Sci.* **36**, 513–518 (1995).
16. A. E. Burgess, "The Rose model, revisited," *J. Opt. Soc. Am. A* **16**, 633–646 (1999).
17. H. B. Barlow, "Temporal and spatial summation in human vision at different background intensities," *J. Physiol. (London)* **141**, 337–350 (1958).
18. A. E. Burgess, X. Li, and C. K. Abbey, "Visual signal detectability with two noise components: anomalous masking effects," *J. Opt. Soc. Am. A* **14**, 2420–2442 (1997).
19. A. E. Burgess, "Prewhitening revisited," in *Medical Imaging, 1998: Visual Perception*, H. L. Kundel, ed., Proc. SPIE **3340**, 55–64 (1998).
20. Note a subtle difference in notation. Tanner and Birdsall (Ref. 4) used one-sided power spectral density N_0 , whereas we use the two-sided $N = 2^{-k}N_0$, where k is the dimensionality of the noise (e.g., $k = 2$ for two-dimensional space), which simplifies the equations (see Ref. 8).
21. H. B. Barlow, "Retinal and central factors in human vision limited by noise, in *Vertebrate Photoreception*, B. Barlow and P. Fatt, eds. (Academic, New York, 1977).
22. P. B. Elliott, "Appendix 1—Tables of d' ," in *Signal Detection and Recognition by Human Observers*, J. S. Swets, ed. (Wiley, New York, 1964), pp. 651–684.
23. D. G. Pelli, C. W. Burns, B. Farell, and D. C. Moore, "Identifying letters," *Vision Res.* (to be published).
24. K. R. Gegenfurtner and D. C. Kiper, "Contrast detection in luminance and chromatic noise," *J. Opt. Soc. Am. A* **9**, 1880–1888 (1992).
25. A. Burgess and H. B. Barlow, "The precision of numerosity discrimination in arrays of random dots," *Vision Res.* **23**, 811–820 (1983).
26. J. M. Harris and A. J. Parker, "Efficiency of stereopsis in random-dot stereograms," *J. Opt. Soc. Am. A* **9**, 14–24 (1992).
27. J. A. Solomon and D. G. Pelli, "The visual filter mediating letter identification," *Nature (London)* **369**, 395–397 (1994).
28. N. Majaj, D. G. Pelli, P. Kurshan, and M. Palomares, "The role of spatial frequency channels in letter identification," *Vision Res.* (to be published).
29. M. Raghavan, "Sources of visual noise," Ph.D. dissertation (Syracuse University, Syracuse, New York, 1995).
30. The fivefold deviation from constant efficiency is only $\sqrt{5}$ -fold in contrast because efficiency, like energy, is proportional to squared contrast. Parish and Sperling [D. H. Parish and G. Sperling, "Object spatial frequencies, retinal spatial frequencies, noise, and the efficiencies of letter discrimination," *Vision Res.* **31**, 1399–1416 (1991)] found a hint of this nonzero slope over the 32:1 range of size that they tested.
31. J. Rovamo, V. Virsu, and R. Näsänen, "Cortical magnification factor predicts the photopic contrast sensitivity of peripheral vision," *Nature (London)* **271**, 54–56 (1978).
32. A. Burgess, "Image quality, the ideal observer, and human performance of radiologic decision tasks," *Acad. Radiol.* **2**, 522–526 (1995).
33. J. G. Robson, "Spatial and temporal contrast-sensitivity functions of the visual system," *J. Opt. Soc. Am.* **56**, 1141–1142 (1966).
34. H. B. Barlow, "A method of determining the overall quantum efficiency of visual discriminations," *J. Physiol. (London)* **160**, 155–168 (1962).
35. H. B. Barlow, "Measurements of the quantum efficiency of discrimination in human scotopic vision," *J. Physiol. (London)* **160**, 169–188 (1962).
36. W. P. Tanner, Jr. and J. A. Swets, "A decision-making theory of visual detection," *Psychol. Rev.* **61**, 401–409 (1954); "The human use of information. I. Signal detection for the case of the signal known exactly," *Trans IRE PGIT-4*, 213–221 (1954).
37. A. Rose, "Quantum effects in human vision," *Adv. Biol. Med. Phys.* **5**, 211–242 (1957).
38. N. S. Nagaraja, "Effect of luminance noise on contrast thresholds," *J. Opt. Soc. Am.* **54**, 950–955 (1964).
39. A. van Meeteren and J. Boogaard, "Visual contrast sensitivity with ideal image intensifiers," *Optik (Stuttgart)* **37**, 179–191 (1973).
40. R. W. Engstrom, "Quantum efficiency of the eye determined by comparison with a TV camera," *J. Opt. Soc. Am.* **64**, 1706–1710 (1974).
41. A. Burgess, "Visual signal detection. III. On Bayesian use of prior knowledge and cross correlation," *J. Opt. Soc. Am. A* **2**, 1498–1507 (1985).
42. B. S. Tjan, W. L. Braje, G. E. Legge, and D. Kersten, "Human efficiency for recognizing 3-D objects in luminance noise," *Vision Res.* **35**, 3053–3069 (1995).