Published in final edited form as:

J Opt Soc Am A Opt Image Sci Vis. 2009 November ; 26(11): B94–109.

Pattern recognition in correlated and uncorrelated noise

Brianna Conrey and Jason M. Gold*

Department of Psychological and Brain Sciences, Indiana University Bloomington, 1101 E. 10th Street, Bloomington, Indiana 47405 USA

Abstract

This study examined how correlated, or filtered, noise affected efficiency for recognizing two types of signal patterns, Gabor patches and three-dimensional objects. In general, compared with the ideal observer, human observers were most efficient at performing tasks in low-pass noise, followed by white noise; they were least efficient in high-pass noise. Simulations demonstrated that contrast-dependent internal noise was likely to have limited human performance in the high-pass conditions for both signal types. Classification images showed that observers were likely adopting different strategies in the presence of low-pass versus white noise. However, efficiencies were underpredicted by the linear classification images and asymmetries were present in the classification subimages, indicating the influence of nonlinear processes. Response consistency analyses indicated that lower contrast-dependent internal noise contributed somewhat to higher efficiencies in low-pass noise for Gabor patches but not objects. Taken together, the results of these experiments suggest a complex interaction among signals, external noise spectra, and internal noise in determining efficiency in correlated and uncorrelated noise.

1. INTRODUCTION

Ideal observer analysis is a powerful technique that has led to remarkable insights about human vision, from basic physiology [1] to complex visual cognition [2]. However, the low human efficiencies, or ratios of ideal to human performance, for performing tasks involving complex patterns, such as faces, objects, and letters [3–5], have presented challenges to researchers who study vision using either traditional ideal observer analysis [6] or sequential ideal observer analysis, a modeling approach that gradually builds constraints into an otherwise-ideal model [7]. Another set of modeling approaches, which taken together we might call the “ideal signal search” approach, stems from the observation that human observers should be most efficient at processing signals that are well matched to their inherent perceptual and cognitive constraints [8–11].

The research presented here was motivated by the insight that in addition to the signal, the external noise added to the signal, and specifically the correlation structure of that noise, could also be chosen to match the constraints of human vision. Most psychophysical research uses Gaussian white noise, which is uncorrelated in the spatial and frequency domains, and so relatively little is known about how the choice of an external noise spectrum affects human efficiencies and strategies for performing visual pattern recognition tasks, especially when the patterns to be recognized are complex. In addition, previous research on visual pattern recognition in correlated noise, much of it done in the context of medical imaging, has used only relatively simple patterns. Therefore, in order to begin to examine the interaction between noise spectrum and signal type, this research investigated efficiencies and human strategies for

*Corresponding author: jgold@indiana.edu.

OCIS codes: 330.4060, 330.5000, 330.5510.

recognizing two types of patterns, or signal types. One of the types, oriented Gabor patches, is relatively simple and is commonly used in psychophysical research, and the second type, three-dimensional objects, is more complex and is commonly encountered in everyday settings.

This work drew on two primary areas of the research literature. First, previous work on natural scene statistics informed the choice of the external noise spectra used in this study. Natural scenes are generally accepted to have amplitude spectra that fall off approximately as $1/f$, with f = spatial frequency [12–16]. In addition, certain properties of the visual system, in particular the logarithmic spacing of frequency channels [17,18], make it well suited for perceiving stimuli with $1/f$ structure [19].

This work also drew on previous research on how the structure of noise interacts with visual perception. One key finding from this literature is that, for simple stimuli such as Gaussian bumps, humans are able to adjust their strategy for performing a task based on the spatial frequency content of noise [20–22]. In addition, humans are able to partially compensate for redundancies introduced by the correlations in low-pass noise, meaning they can perform at least certain tasks in low-pass noise in a quasi-ideal way [23–27]. However, they are not able to compensate for the correlations in high-pass noise [23,24,28,29], perhaps because this type of noise correlation is not commonly encountered in the visual environment [25].

The research presented here examined how efficient human observers are at performing psychophysical tasks with different types of signals in white, low-pass, and high-pass noise. To address this question, thresholds were estimated for human and ideal observers performing two different tasks in white noise, low-pass noise, and high-pass noise. One of the tasks studied was Gabor patch orientation discrimination, which involves a signal that is localized in frequency space. The other task, three-dimensional object recognition, was chosen to be more naturalistic and also because it involves signals that have broadband spectra. Human and ideal observer thresholds in each were compared to obtain a measurement of human efficiency, which allows the comparison of human performance across tasks and noise conditions.

2. METHOD

A. Observers

Four observers (two female, two male) participated in the experiment. Two (MD and PM) were naïve to the purpose of the experiment; BC and JG were authors. All had normal or corrected-to-normal vision. The observers' ages ranged from 18 to 35, with a mean age of 24.5. Each observer completed the experiment within approximately four 1-hour sessions. As laboratory personnel, BC, JG, and PM volunteered their time; MD was compensated at the rate of \$10 per hour.

B. Apparatus

Stimuli were displayed on a Sony Trinitron Multiscan G520 monitor controlled by an Apple G4 computer running Mac OS 9.2.2. The monitor had a resolution of 1024×768 pixels, subtending 16.4°×12.4° of visual angle at the viewing distance of 130 cm. The frame rate was set to 85 Hz. The experiment was conducted in the MATLAB programming environment (version 5.2.1 for MacIntosh) using the Psychophysics Toolbox extensions [30,31].

A Minolta Luminance Meter LS-100 photometer was used to calibrate the monitor, and a 1792-element look-up table was built from the calibration data in order to linearize the display, as described by Tyler *et al.* [32]. Luminance ranged between 0.7 and 103.6 cd/m², with an average luminance of 34.1 cd/m². During the experiment, the software selected appropriate luminance values for each trial from the calibrated look-up tables and constructed an 8-bit look-up table for the display using these luminance values.

C. Signals

Two signal types were utilized in this experiment: Gabor patches and two-dimensional projections of simple three-dimensional geometric objects. Each signal was 200×200 pixels, measuring $7.4 \text{ cm} \times 7.4 \text{ cm}$ and subtending $3.3^\circ \times 3.3^\circ$ of visual angle at the viewing distance of 130 cm.

Gabor patches—Two Gabor patches that were identical except for orientation were used in the orientation discrimination task (leftmost column of Fig. 1). Both Gabor patches had a spatial frequency of 1.5 cycles per degree (5 cycles per image at ± 2 standard deviations) and a sine phase of 0. One had an orientation of 45° and the other had an orientation of -45° .

Objects—Six different geometric objects were chosen for the object-identification task: a cone, a cube, a cylinder, a pyramid with a square base, a sphere, and a square pyramidal frustum (a square pyramid truncated by a plane parallel to its base; see Fig. 1). So as to minimize reliance on size cues, the objects were constructed to have similar heights and widths. For instance, the lengths of the sides of the square bases were equal to the diameter of the circular bases. One two-dimensional projection was used for each object throughout the experiment. All of the objects had the same angle of rotation away from the observer. Rotating the objects enhanced the illusion of three-dimensional depth.

All images were represented in contrast values such that the contrast (c_{xy}) at pixel location (x, y) in an image was given by

$$c_{xy} = \frac{l_{xy} - L}{L}, \quad (1)$$

with L equal to the background luminance and l_{xy} equal to the pixel luminance. The integrated contrast of a given image was computed using root-mean-squared (RMS) contrast C_{RMS} , defined as $n m$

$$C_{RMS} = \sqrt{\frac{1}{nm} \sum_{x=1}^n \sum_{y=1}^m c_{xy}^2}, \quad (2)$$

where n is equal to the width of the image in pixels, m is equal to the image's height in pixels, and nm is the total number of pixels in the image. C_{RMS} was then converted to contrast energy E , which is expressed in degrees of visual angle squared (deg^2) and is defined as

$$E = (C_{RMS})^2 n m a_{\text{pixel}}, \quad (3)$$

where (a_{pixel}) is the area of a pixel in degrees of visual angle squared.

D. Noise Conditions

In addition to a white-noise condition, two low-pass noise conditions and two high-pass noise conditions were used in this experiment. A condition with no external noise was also included so that the noise-masking technique of extrapolating between white-noise and no-noise thresholds could be used to estimate the contrast-invariant component of internal noise present

in the visual system [33–38]. Thus, including the condition with no external noise, six different noise conditions were tested for each type of signal.

In all of the conditions that utilized external noise, a unique sample of noise with the same dimensions as the signals was generated on each trial and added to the trial's signal. The noise fields with filtered spectra were generated to have the same average power spectral density (contrast power per unit bandwidth) as the white noise.

White noise—In addition to being used in the white-noise condition, white-noise fields served as the base for all filtered-noise fields, as noted above. The value for each pixel in a white-noise field was obtained from a Gaussian pseudo-random number generator with a mean of 0 and a contrast variance of 0.05 (power spectral density = $1.3 \times 10^{-5} \text{ deg}^2$). Each value in the noise matrix was treated as a contrast value. The chosen variance for the noise ensured that at least 95% of the values in the distribution would fall within the linear contrast range of the noise display when added to either the Gabor patch or the object signals. Values that exceeded ± 2 standard deviations from the mean were resampled to fall within the range of displayable contrast values.

Low-pass, $1/f$ noise—The amplitude of the low-pass $1/f$ filter fell off as the inverse of spatial frequency. Each low-pass $1/f$ noise field was created by filtering all the frequencies in the amplitude spectrum of a white-noise field by a filter of the form $g(f_i) = 1/f_i$, where f_i is equal to an individual spatial frequency.

Low-pass, flipped f noise—The amplitude of the low-pass *flipped f* noise filter fell off linearly with spatial frequency. Taking $\max(f)$ and $\min(f)$ as the maximum and minimum frequencies represented in the white-noise field and f_i as an individual frequency, the low-pass *flipped f* noise filter was given by $g(f_i) = \max(f) + \min(f) - f_i$.

High-pass, flipped $1/f$ noise—The amplitude of the high-pass *flipped $1/f$* filter increased with the inverse of spatial frequency. Taking $\max(1/f)$ and $\min(1/f)$ as the maximum and minimum values, respectively, in the low-pass $1/f$ noise filter and f_i as an individual frequency, the equation for the high-pass flipped $1/f$ filter was $g(f_i) = \max(1/f_{\text{filt}}) + \min(1/f_{\text{filt}}) - 1/f_i$.

High-pass, f noise—Each high-pass f noise field was created by filtering each frequency f_i in the amplitude spectrum of a white-noise field by a filter of the form $g(f_i) = f_i$, whose amplitude increased linearly with spatial frequency.

Examples of each type of noise and the amplitude spectrum of the associated filter are shown in Fig. 2. Each of the sample filtered noise fields shown in Fig. 2 was created using the sample pictured white-noise field as a base.

E. Viewing Conditions

Viewing was binocular through natural pupils, and a combination forehead and chin rest stabilized the observer's head. The monitor supplied the only source of illumination during the experiment.

F. Procedure

The experiment consisted of two tasks: 1-of-2 Gabor patch orientation identification and 1-of-6 object recognition. Each of these tasks was presented with the five different conditions of external noise and one condition with no noise. Trials were blocked by task and by noise condition. The order of the tasks was counterbalanced across observers, and the order of the noise conditions within each task was randomized for each observer.

On each trial, observers were presented with a signal plus an external noise field mask on a background of average luminance (or a signal alone in the no-noise condition). Signal duration was approximately 360 ms in both tasks. After signal presentation, the display was reset to average luminance and a selection screen with images of the possible signals was presented. Observers used the mouse to select the signal they thought had been presented. After a selection was made, auditory feedback indicated whether the response was correct, and the display was reset to average luminance prior to the beginning of the next trial.

In each condition, an adaptive staircase procedure manipulated the contrast energy of the signals across trials according to the observer's responses. For the Gabor patch conditions, which had only two alternatives and thus a chance performance rate of 50%, contrast energy was manipulated by a staircase tracking the 71% correct point on the psychometric function with a 2-down, 1-up rule. For the shape conditions, which had six alternatives and thus a guessing rate of 16.7%, the staircase tracked the 50% correct point using a 1-down, 1-up rule. The threshold for each condition was estimated by fitting Weibull functions to the data, which consisted of a total of 150 trials. Threshold was defined as the contrast energy yielding 71% correct responses for the Gabor patch conditions and 55% correct responses for the shape conditions.

G. Ideal Observer

The ideal observer for the 1-of- N and 1-of-2 identification tasks in white noise has been derived elsewhere [5,6]. On each trial, the ideal observer maximizes the likelihood of the observed noisy stimulus given the possible noise-free templates that could have been presented, choosing the template with the highest posterior probability given the stimulus [39] or equivalently in our case the template that maximizes the dot product (i.e., $\sum_{i,j} X_{ij} Y_{ij}$) with the stimulus [5]. Derivations for the spatial-domain ideal observer in correlated noise have also been published elsewhere [25,29] (also cf. [40] for a derivation of a nearly ideal frequency domain observer in correlated noise), but a detailed derivation of the spatial-domain ideal observer in correlated noise that reflects our procedures more precisely than previous versions is presented in Appendix A. The ideal observer in correlated noise uses a process called “prewhitening” prior to applying the same decision rule used by the ideal observer in white noise. Prewhitening allows the ideal observer to account for the noise correlations in the presented stimulus and thus assumes that the ideal observer knows the form of the filter used to create the noise. This is also a hidden assumption for ideal observers in white noise, a special case in which the filter is all-pass and does not introduce correlations into the noise [24,41]. Monte Carlo simulations of 5,000 trials per condition were run to estimate ideal performance for each task in each of the external noise conditions.

3. RESULTS AND DISCUSSION

The results of the experiment are summarized in Fig. 3. The top two panels plot contrast energy thresholds for four human observers, the mean across the human observers, and the ideal observer in the Gabor (left panel) and object (right panel) tasks. The bottom two panels plot the corresponding efficiencies for each task.

There are two very notable aspects to these data. First, efficiencies for both kinds of stimuli were extremely low in high-pass noise—far lower than in either low-pass or white noise. Inspection of the thresholds shows that the ideal observer's thresholds were exceptionally low in the high-pass conditions relative to the low-pass and white conditions. This implies that more information was carried by low spatial frequencies for both kinds of stimuli, since the ideal observer's threshold is a direct reflection of the intrinsic difficulty of a task. This result is self-evident in the case of the Gabor stimuli, where the signals were localized in the lower part of the frequency spectrum by design. However, it is a less obvious result in the case of the

objects, where energy (and information) is spread across the frequency spectrum. Apparently, information in the object recognition task is also highly concentrated in the lower spatial frequencies.

Second, efficiency was highest in the presence of a low-pass noise for both kinds of stimuli. For Gabors, efficiency was higher in both kinds of low-pass noise than in white noise; for objects, efficiency was higher in flipped f noise than in white noise. Although these may not appear to be large differences in Fig. 3, note that the y axis spans a total of 5 log units. The mean efficiency in the Gabor task in flipped f noise was ~13%, compared to ~5.5% in white noise. For the objects, mean efficiency in flipped f noise was ~8.4%, compared with ~4.1% in white noise. Thus, low-pass filtering the noise produced as much as a doubling or more in efficiency relative to white-noise conditions.

A. Low Efficiencies in High-Pass Noise

Although the large difference between ideal and human thresholds accounts for why efficiency was so low in the high-pass conditions, it does not explain why human observers were so inefficient in these conditions. One major factor that always places an upper bound on human efficiency is internal noise. Human thresholds in the high-pass noise conditions were in most cases quite similar to thresholds measured without external noise. Levels of external noise that are sufficient to limit human performance result in thresholds that are higher than if no external noise is present [38], suggesting that the amount of external noise in the high-pass conditions may not have been sufficient to limit human performance. Instead, human observers may have been limited by their own internal noise in these high-pass conditions.

To test this possibility, we conducted simulations for a pair of model observers that were ideal in all respects except that each was limited by internal noise. The first model observer was limited by *contrast-invariant* internal noise, or noise whose power spectral density is independent of stimulus contrast. The power spectral density of this type of internal noise in the human observers was estimated using the technique of external noise masking [33–38]. In external noise masking, the amount of contrast-invariant internal noise an observer has is quantified by measuring contrast energy thresholds for the same task in varying amounts of external white noise. The observer's thresholds will increase linearly with the contrast of the external noise [43]. The observer's contrast-invariant internal noise, or equivalent noise, is defined to be equal to the amount of external white noise that doubles the observer's no-noise threshold. This relationship can be quantified with the following equation:

$$E = k(N_e + N_i), \quad (4)$$

where E denotes the contrast energy threshold, k the efficiency of the observer's calculation in reaching decisions, N_e the power spectral density of the external noise, and N_i the power spectral density of the contrast-invariant internal noise, whose value is assumed to be independent of the value of N_e . If $N_e = N_i$, Equation (4) indicates that E will be doubled relative to when $N_e = 0$. This linear relationship allows us to extrapolate the power of N_i from as few as two contrast energy thresholds measured in different amounts of external white noise (including no noise). Because we measured contrast energy thresholds in no external noise as well as at one level of white noise, we were able to estimate the amount of contrast-invariant noise for each of our human observers in both tasks. Thus, the first model observer for each task included a contrast-invariant, Gaussian white internal noise source that was equal to the average of these estimates.

The second model observer had the same amount of contrast-invariant internal noise plus an additional source of *contrast-dependent* internal noise that was proportional to stimulus

contrast, which at high external noise levels is largely driven by the power of the external noise rather than the contrast energy of the signal. Previous experiments have shown that the power of this contrast-dependent noise source is often similar to that of the external white noise [44]. This situation is typically described as having an internal-to-external noise ratio (I/E , or σ_i/σ_e) of 1. The contrast-dependent noise source of our model observer was therefore a Gaussian white noise with the same power spectral density as the external Gaussian white noise that was either presented on the trial or used to generate the trial's filtered noise.

Monte Carlo simulations of 5,000 trials each were conducted to measure the contrast energy thresholds of these two model observers in both tasks and all noise conditions for which the ideal observer was measured. The model observers were ideal except for the addition of the Gaussian white internal noise sources described above. The internal noise source or sources, either contrast-invariant only or contrast-invariant and contrast-dependent, were added to the signal at the same time as the external noise source and thus before the prewhitening process. Although the locus of the internal noise sources in perceptual decision making has been a matter of some controversy [34,45,46], adding the internal noise after prewhitening in our simulations could not account for differences in efficiency observed across noise conditions, because this would be equivalent to increasing the signal-to-noise ratio by the same amount at all spatial frequencies for all noise types and thus would merely reproduce the ideal observer's exact pattern of thresholds but at a somewhat higher level. By contrast, adding the internal noise before prewhitening allows for the possibility that internal noise could affect the stimulus differently at different frequencies depending on the external noise condition. The model observers were assumed to know the statistics of the internal noise and to use this information in calculating the probability of the signal templates given the presented stimulus.

We expect thresholds and efficiencies for a good model observer to show the same pattern across noise conditions as the human observers, even if the absolute levels are not quite the same as for the humans. This expectation is due to a critical assumption of ideal observer analysis that any information lost during processing cannot be recovered and so must affect the threshold [7].

The addition of contrast-invariant internal noise did not substantially affect thresholds for the model observer limited by only contrast-invariant noise. Thus, we can conclude that the amount of contrast-invariant internal noise present in the visual systems of the human observers is not sufficient to explain their relatively low efficiencies in high-pass noise. However, the model limited by both contrast-invariant and contrast-dependent internal noise was more successful at predicting human thresholds and efficiencies. The thresholds and efficiencies for this model are plotted along with the human and ideal observer data in each panel of Fig. 3. These data show that the presence of contrast-dependent noise captures some, but not all, of the patterns of thresholds and efficiencies observed for the human observers in each task. That is, efficiency in the white-noise condition is relatively overestimated and efficiency in the high-pass *flipped* $1/f$ condition is relatively underestimated in both tasks. However, the qualitative effect of relatively lower efficiencies in high-pass noise is captured by this model observer, indicating that the presence of contrast-dependent noise may be responsible for the relatively low efficiencies in high-pass noise.

B. Higher Efficiencies in Low-Pass than in White Noise

Effects of Strategy—We found that efficiency reached its highest point in low-pass noise for both Gabors and objects. In the Gabor task, efficiencies were higher for both kinds of low-pass noise than for white noise, whereas for objects efficiencies were highest for *flipped* $1/f$ noise, followed by white noise and then $1/f$ noise. One possible cause for these two patterns of results is that there were differential interactions between the perceptual system and the spectra of the signal and noise involved in each task. Although both classes of signal contained more

information at low than at high spatial frequencies, the objects had relatively broadband spectra, whereas the Gabor patches were much more localized in frequency space. This explanation assumes a fixed strategy for each task that happened to be better suited to some combinations of signal and noise spectra than to others.

Alternatively, observers may have simply altered their strategy based on the type of external noise used, and these changes in strategy were differentially effective across noise types. We explored this second possibility by using the response classification technique to estimate the linear templates used by observers when recognizing Gabor patches and objects in white, *flipped f*, or *1/f* noise. The response classification technique is a noise-based system identification method that uses the correlation between the contrast of noise and an observer's responses across a series of trials in order to infer the properties of the linear template used by an observer when performing a task [47,48]. The result is a *classification image*, which is a map that shows the relative weight given to each pixel in an image by an observer over the course of an experiment. In the case of our Gabor and object recognition tasks, we used this technique to explore whether observers' strategies (indexed by their classification images) varied across noise conditions, and if so, whether the variations in strategy could predict the variations observed in efficiency.

We measured classification images for a new set of observers for both the Gabor and the object tasks in both kinds of low-pass noise and in white noise. The object task was altered to have only two alternatives (the cube and the frustrum) rather than the six used in the original experiment. The main reason for this is that the number of trials required to measure classification images increases with the number of alternatives [49]. A task with two alternatives is far more tractable in the context of classification-image analysis. Each observer participated in a total of 10,000 trials per noise condition spread across a series of 1-h sessions. The order of the noise conditions was randomly assigned for each observer. One observer (BC) participated in both the Gabor and the object tasks. All other aspects were identical to the original experiment.

The method for calculating classification images in visual 1-of-2 identification tasks using uncorrelated (white) noise as well as the relationship between the classification image and the observer's linear template have been described elsewhere [47,48,50]. Briefly, the noise fields from the experiment are binned according to what signal was presented (S1 or S2) and what the observer's response was (R1 or R2), resulting in four signal-response bins (S1R1, S1R2, S2R1, and S2R2). The noise fields in each bin are then averaged and combined according to the equation

$$C = (\overline{S1R1} + \overline{S2R1}) - (\overline{S1R2} + \overline{S2R2}), \quad (5)$$

which results in the classification image C. In addition to this method of calculating classification images in white noise for 1-of-2 identification tasks, Abbey and Eckstein [23, 41] have derived a general method for calculating classification images in white or correlated noise for two-alternative forced-choice tasks. Appendix B describes a similar derivation of classification images in correlated noise, but for 1-of-2 identification tasks. Ideal classification images in correlated noise can be computed directly from the ideal signal templates (the images are equivalent to a difference image of the two templates) that correspond to "prewhitened" versions of the templates for the same task in white noise.

Figure 4 shows the resulting efficiencies for each observer in each task. These data replicate the same patterns found in the original experiment, validating the use of two rather than six alternatives in the object task. Figures 5 (Gabors) and 6 (objects) show the resulting

classification images when the data are combined across all three observers within each task and noise condition. For the Gabor classification images, the responses associated with the -45° signal were subtracted from the responses associated with the $+45^\circ$ signal. Thus, the polarity of the pixels from the $+45^\circ$ responses was preserved, whereas the polarity of the pixels from the -45° responses was reversed. For objects, the responses associated with the frustum were subtracted from the responses associated with the cube, preserving the polarity of the pixels from the cube responses and reversing the polarity of the pixels from the frustum responses. The top row of each figure displays the ideal classification images calculated from the ideal templates. The middle row shows raw classification images calculated from the combination of all the human data (i.e., 30,000 trials per condition). The bottom row displays the same images but smoothed with a 9×9 convolution kernel (the matrix product of the vector [1 2 3 4 5 4 3 2 1] with itself transposed) so as to make it easier to visualize the structure in the images.

Inspection of Figs. 5 and 6 shows that, for both Gabors and objects, the classification images in the low-pass conditions differ markedly from those of the ideal observer. The differences between human and ideal are less pronounced in the white-noise condition. In the case of the Gabors, the human observers seem to be using a smaller area of the available stimulus than the ideal observer in order to make their decisions. The boundaries between dark and light areas in the low-pass conditions do not appear as clear as those in the white-noise condition. In the case of the objects, the human low-pass classification images contain vertical curved continuous edges that seem to correspond to the edges of the objects. The marked variation in the human classification images across different noise conditions is consistent with the idea that the human observers were adjusting their strategy based on the kind of noise that is present in the stimulus.

The classification analysis depends upon the assumption that the observer uses a linear template to perform the task. If not, it is difficult to predict the impact that various nonlinearities will have on the classification image weights. Thus, the variations in the human classification images across noise conditions seen in Figs. 5 and 6 could be due to the presence of one or more nonlinear processes. If an observer does in fact use a linear decision variable to perform a task, then their classification image, together with information about the observer's performance and the power of the external noise, should allow one to accurately estimate the observer's efficiency in the task [50]. Murray *et al.* [50] have derived an equation for predicting an observer's efficiency from their classification image. Any deviations from linearity will result in either underprediction or overprediction of efficiency, depending upon the nature of the nonlinearity. We used this approach to determine whether differences in the observers' linear templates could explain the efficiency differences we observed across noise conditions. Following Murray *et al.*, we calculated the ratio of the actual to the predicted efficiency for each observer in each condition. This ratio provides a convenient index that allows the goodness of efficiency predictions to be compared across conditions and observers. A ratio that is less than 1 indicates overprediction of efficiency, whereas a ratio that is greater than 1 indicates underprediction of efficiency. Any deviation from a ratio of 1 indicates the presence of a nonlinear strategy or strategies. Although Murray *et al.*'s formula for predicting efficiency based on a classification image was derived for the white-noise case, because our filtered-noise images are calculated with prewhitened noise fields, the formula is also applicable to these cases.

The ratios of actual to predicted efficiencies for the Gabor and object tasks are shown in Fig. 7. These data were computed based on the combined classification images and average efficiencies in each condition. These data show a very interesting difference between the white and low-pass noise conditions. That is, efficiency is just slightly underpredicted in the presence of white noise, whereas it is severely underpredicted in the low-pass noise conditions. This

underprediction is most apparent in the Gabor task, where efficiency is underpredicted by several orders of magnitude in both low-pass conditions. These results indicate that nonlinear processes had an effect on the classification images and that the effects of the nonlinearities were markedly more severe in the presence of low-pass noise.

The presence of significant nonlinearities makes the variations in the classification images across noise conditions seen in Figs. 5 and 6 quite difficult to interpret. Nonlinearities in an observer's strategy can take many forms and include nonlinear transformations of the stimulus, such as taking the variance or covariance [6,51] or log luminance [50] of the pixel values; nonlinearities at the decision level, such as response bias or guessing [50]; and instances in which the observer has multiple templates associated with the same stimulus [3,5,52–55]. This last type of nonlinearity is thought to be relatively common and includes uncertainty about certain aspects of the stimulus, such as its spatial position or phase. Therefore, in addition to assessing the linearity of our classification images by comparing predicted and actual efficiencies for each condition, we also examined the classification subimages [55], or averaged images from the signal-response bins used to calculate the classification images, for evidence of nonlinearities involving multiple templates.

For a linear observer, the subimages associated with a particular response will contain a positive image of the signal corresponding to the response and a negative image of the other signal. The S1R1 and S2R1 bins are expected to look the same, because the same noise pixels should influence observers to respond “S1” regardless of whether S1 or S2 was actually presented. Likewise, the S1R2 and S2R2 bins are expected to look the same. If the bins associated with a particular response category exhibit asymmetry, this is a strong indication of a nonlinear strategy that includes uncertainty about some aspect or aspects of the signal [41,56,57]. For example, it has been shown that significant amounts of spatial uncertainty can result in a strong negative image of the presented signal on error trials (bins S1R2 and S2R1). The accompanying positive images generated from error trials may be absent altogether [58]. Alternatively, if the performance level is high enough, the positive images may appear as a faint haze, a situation known as “signal clamping” by analogy with a technique from neuroscience for mapping a cell's receptive field [55]. The absent or hazy response image reflects averaging of the many spatially shifted templates that the observer has chosen over the presented signal.

Subimages for the combined, smoothed human data and for a simulation of the ideal observer in all three noise conditions of the Gabor patch task are displayed in Figs. 8 (Gabors) and 9 (objects). In the case of the Gabors, the human subimages show large asymmetries between the S1R1 and S2R1 bins (here, the Signal= +45°/Response=+45° and Signal=−45°/Response=+45° panels, respectively) and between the S1R2 and S2R2 bins (the Signal=+45°/Response=−45° and Signal= −45°/Response=−45° panels) in all three conditions. Specifically, in the white-noise-condition error bins (S1R2 and S2R1), the negative image associated with the presented signal is dominant, whereas the positive image associated with the response is not visible. The positive images in the error bins of the $1/f$ noise condition are also faint or absent, although the subimages are much fainter overall in this condition despite much higher efficiencies than in the white-noise condition. The *flipped f* condition's subimages are also faint and appear to display a different pattern than the other two conditions, yet they still show asymmetries that are indicative of a nonlinear strategy. In the case of the objects, all three sets of subimages show marked asymmetries between the correct and the error bins associated with each response, suggesting that the humans may have used nonlinear strategies that involve uncertainties about the signal. In addition, all three sets show signs of signal clamping in that the negative image of the presented signal tends to dominate the images in the error bins.

Effects of Internal Noise—The measure used to predict an observer's efficiency from their classification image depends mostly on the squared dot product of the observer's classification

image with the ideal observer's template; this squared dot product term is affected both by internal noise and by the observer's sampling efficiency, or the similarity of the observer's linear template to the ideal template [50]. Larger amounts of internal noise should lower both predicted and actual efficiency, and so it is unlikely that differences in internal noise across conditions are responsible for the large discrepancies observed between predicted and actual efficiencies. However, without measuring internal noise in all three conditions, we cannot rule out the possibility that differences in internal noise contributed to the results we observed across conditions for both the actual efficiencies and the squared dot product term that is the critical factor for predicting efficiencies.

The total amount of an observer's internal noise can be measured using a technique called response consistency [33,34,44,59]. The basic idea is this: If a physically identical stimulus is presented on multiple trials of an experiment, a noise-free observer will always make the same response to that stimulus, whereas observers with internal noise will not be completely consistent in their responses. How consistent the observer is depends on the ratio of the standard deviation of the observer's internal noise to the standard deviation of the external noise, or the σ_i/σ_e ratio [44]. Thus, an observer's consistency in a double-pass experiment—one in which the first and second halves of the experiment contain a sequence of physically identical trials—can be used to estimate the observer's total amount of internal noise. The response consistency technique measures contrast-invariant and contrast-dependent noise together; however, at high levels of external noise it is expected that any response inconsistency measured will almost exclusively be due to contrast-dependent noise [33,34].

We estimated the contribution of internal noise to the variations in efficiency across noise conditions using the double-pass response consistency technique for two observers in the Gabor task and two observers in the object task with each kind of noise (white, *1/f*, and *flipped f*). The tasks were identical to those used when measuring the classification images described above. To measure response consistency, the observers completed one approximately hour-long session of 1,500 trials per noise condition for a total of three sessions per task. The first and second halves of each session consisted of physically identical trials presented in the same sequence. To achieve this, a 2-down, 1-up staircase tracked the 71% point on the observers' psychometric functions and adjusted signal contrast for the first 750 trials of each session only. The sequence of the signals, contrast levels, and random number generator seeds used to generate the noise fields on each trial during the first half of the experiment were saved and used to present a physically identical sequence of signals and noise fields during the second half. During the second 750 trials, the observers still received accuracy feedback on their responses; however, to ensure that these trials were physically identical to the first 750 trials, signal contrast was not adjusted according to the observers' responses. This meant that the feedback observers received during the second half of the experiment did not necessarily correspond to how the signal contrast was adjusted from trial to trial. However, this lack of correspondence was not typically noticeable.

Thresholds and efficiencies (not shown) were quite similar to those reported with respect to the classification images. The amount of consistency displayed by each observer in each two-pass session can be visualized by plotting the percent correct at each signal contrast level tested as a function of the percent of the observer's responses that were the same at that contrast level on the two passes. Plots of percent agreement versus percent correct for the three noise conditions are shown in Figs. 10 (Gabor) and 11 (objects). The line in each panel reflects the best fit of the data to the equation

$$p_c = m \log_{10} \left(\frac{p_a}{100} \right) + 100, \quad (6)$$

where p_c is percent correct, p_a is percent agreement, and m corresponds to the slope of the line [33,34]. m increases as the ratio of internal to external noise standard deviations (σ_i/σ_e) decreases, because a decrease in internal noise results in higher overall percent agreement and thus a steeper line. More specifically, in our tasks, m and σ_i/σ_e are related by the following equation [34]:

$$\sigma_i/\sigma_e = \alpha + \gamma_1 e^{-\beta_1 m} + \gamma_2 e^{-\beta_2 m}, \quad (7)$$

where α , β_1 , β_2 , γ_1 , and γ_2 are fitted parameters and the σ_i/σ_e that corresponded to m was determined by fitting this equation with values of m estimated through bootstrap simulations for each observer in each condition [42]. Eqs. (6) and (7) are purely descriptive functions and have no theoretical significance in this context.

The response consistency plots for the Gabor patch task show a difference in slope and σ_i/σ_e across conditions for both observers. Specifically, the slope is steepest, and σ_i/σ_e lowest, for the $1/f$ condition, less steep for *flipped f* noise, and least steep for white noise. For the object task, the response consistency plots do not seem to reveal such striking differences across conditions, although the $1/f$ and *flipped f* conditions again seem to have slightly steeper slopes than the white noise condition for both observers.

Taken together, the results of the response consistency experiment reveal interesting differences between the two tasks under study. Specifically, internal noise appears to differ across noise conditions in the Gabor patch orientation discrimination task but not in the object recognition task. In the Gabor patch task, given equally good linear templates for all three conditions and the levels of internal noise estimated for the human observers in this experiment, a model observer would predict the observed pattern of human efficiencies: $1/f > \text{flipped } f > \text{white}$. On the other hand, even though internal noise is lower in the low-pass external-noise conditions than in the white-noise condition, efficiency is dramatically underpredicted by the classification images only for the low-pass conditions in this task. Thus, we now have still stronger evidence that the human observers' linear template in the low-pass conditions is not as highly correlated with the ideal template as is the humans' linear template in white noise. Therefore, either human strategies are more pronouncedly nonlinear in the low-pass noise conditions or the specific nonlinearities have a greater effect on the low-pass linear template than the white-noise linear template. Regardless of the sorts of nonlinearities present or their effects, however, the lower levels of internal noise in the $1/f$ and *flipped f* conditions can help to increase actual efficiencies only in those conditions.

As for the objects, the response consistency results enable us to largely rule out internal noise as a cause for differences in efficiency across noise conditions in this task. It is plausible that nonlinearities in the human strategies for this task could account for some of the observed differences in actual efficiency across conditions. A more puzzling question is why internal noise would differ across conditions for Gabors but not for objects. One hypothesis is that internal noise might be mediated by individual spatial-frequency-tuned channels, so that each channel only contributes internal noise if it is being recruited by the spatial frequency content of what is being perceived. So, for instance, the Gabor patch is localized at a low frequency, and in the low-pass conditions, the external noise also contains most of its power at low frequencies and has very little power at higher frequencies. If channels contribute internal noise only when the stimulus contains relevant frequency content, then perhaps high-frequency-

tuned channels do not contribute much noise in the low-pass conditions of the Gabor patch task. This would predict lower amounts of internal noise in the low-pass noise conditions than in the white-noise condition, in which higher frequencies are better represented. The objects, because they are relatively broadband, contain information (such as edges) at high frequencies, meaning that even in low-pass noise the higher-frequency-tuned channels would still be recruited and would still contribute internal noise. This would lead to relatively equal amounts of internal noise across conditions in the object recognition task, as was found in the response consistency experiment. Spatial frequency channels that contribute internal noise individually have previously been modeled by Abbey and Barrett [28], who considered a channelized Hotelling observer with independent internal noise in each channel that was proportional to the variance of the image noise in that channel.

4. SUMMARY AND CONCLUSIONS

The primary goal of this research was to measure human efficiency for the recognition of spatial patterns varying in complexity in correlated and uncorrelated noise. We found that pattern recognition efficiency peaked in the presence of low-pass *flipped f* noise, followed by either low-pass *1/f* noise (for Gabors) or white noise (for objects). Efficiency was far lower in the presence of high-pass noise than in either low-pass or white noise. Simulations suggested that the low efficiencies in high-pass noise are related to limitations imposed by internal contrast-dependent noise. Possible sources of the variations in efficiency across low-pass and white conditions were explored using response classification (to measure variations in strategy) and response consistency (to measure variations in internal noise). Although we did find marked differences across classification images obtained in white and low-pass noise, the variations in efficiency could not be explained by the observed differences in classification images. Instead, the classification images provided strong evidence of nonlinear components to the human observers' strategies for all conditions, especially the low-pass conditions, in both tasks. The results of the response consistency experiments suggest that variations in internal noise may have contributed somewhat to the differences in efficiency across noise conditions in the Gabor but not the object recognition task.

This research presents some important new insights about pattern recognition in uncorrelated noise. First, although others have suggested that internal noise may explain low efficiencies in high-pass noise [24], our results indicate that this is a plausible explanation and that the limiting noise source is probably contrast dependent rather than contrast invariant. Second, our results build on previous literature showing that efficiency for the same task may vary depending on external noise spectrum by demonstrating that correlated noise, at least low-pass noise, seems to also affect observer strategies, possibly causing observers to use more pronouncedly nonlinear strategies than white noise. Third, an observers' levels of internal noise may also vary depending on the external noise spectrum, even when the external noise spectra are controlled to have the same total power. In practical terms, the findings presented here should be taken as cautionary. Although efficiencies were discovered to be higher in one or both types of low-pass noise than in white noise, depending on the task, the serious nonlinearities evident in the classification images make using low-pass noise a questionable choice if one is assuming that the human observer's strategy is linear. Certainly white noise lends itself better to the well-developed techniques of visual psychophysics. In addition, given the current findings, it is unclear whether one can predict *a priori* that human observers will indeed perform more efficiently in a given type of low-pass noise than in white noise. Performance appears to depend on the signal as well as the noise spectrum, and one cannot always expect that, for example, humans are most efficient in *1/f* noise just because it is most similar to the structured "noise" inherent in environmental statistics.

In conclusion, this investigation has demonstrated that efficiencies, the levels of internal noise, and the linearity of the human strategy for performing a task can be influenced by the correlation structure of the external noise. Future work will focus on developing our understanding of how the interactions between signal and noise spectra affect efficiency.

Acknowledgments

This research was supported by a National Science Foundation (NSF) Graduate Research Fellowship and an Indiana University College of Arts and Sciences Dissertation-Year Fellowship to B.C. and by National Institute of Health (NIH) grants EY015787 and EY019265 to J. M. G. Discussions with Richard Murray and Bosco Tjan helped frame the motivation for this work. We thank Craig Abbey and two reviewers for their helpful comments.

References

1. Hecht S, Schlaer S, Pirenne MH. Energy, quanta and vision. *J Gen Physiol* 1942;25:819–840. [PubMed: 19873316]
2. Liu Z, Knill DC, Kersten D. Object classification for human and ideal observers. *Vision Res* 1995;35:549–568. [PubMed: 7900295]
3. Gold J, Bennett PJ, Sekuler AB. Identification of band-pass filtered letters and faces by human and ideal observers. *Vision Res* 1999;39:3537–3560. [PubMed: 10746125]
4. Solomon JA, Pelli DG. The visual filter mediating letter identification. *Nature* 1994;369:395–397. [PubMed: 8196766]
5. Tjan BS, Braje WL, Legge GE, Kersten D. Human efficiency for recognizing 3-D objects in luminance noise. *Vision Res* 1995;35:3053–3069. [PubMed: 8533342]
6. Green, DM.; Swets, JA. *Signal Detection Theory and Psychophysics*. Wiley; 1966.
7. Geisler WS. Sequential ideal-observer analysis of visual discriminations. *Psychol Rev* 1989;96:267–314. [PubMed: 2652171]
8. Burgess AE, Wagner RF, Jennings RJ, Barlow HB. Efficiency of human visual signal discrimination. *Science* 1981;214:93–94. [PubMed: 7280685]
9. Olman C, Kersten D. Classification objects, ideal observers & generative models. *Cogn Sci* 2004;28:227–239.
10. Watson, AB. The Committee on Vision. *Frontiers of Visual Science*. National Academy Press; 1978. The ideal observer concept as a modeling tool; p. 32–37.
11. Watson AB, Barlow HB, Robson JG. What does the eye see best. *Nature* 1983;302:419–422. [PubMed: 6835375]
12. Burton GJ, Moorhead IR. Color and spatial structure in natural scenes. *Appl Opt* 1987;26:157–170.
13. Field DJ. Relations between the statistics of natural images and the response properties of cortical cells. *J Opt Soc Am A* 1987;4:2379–2394. [PubMed: 3430225]
14. Ruderman DL, Bialek W. Statistics of natural images—scaling in the woods. *Phys Rev Lett* 1994;73:814–817. [PubMed: 10057546]
15. Tolhurst DJ, Tadmor Y, Chao T. Amplitude spectra of natural images. *Appl Opt* 1992;12:229–232.
16. vanderSchaaf A, vanHateren JH. Modelling the power spectra of natural images: statistics and information. *Vision Res* 1996;36:2759–2770. [PubMed: 8917763]
17. Blakemore C, Campbell FW. On existence of neurones in human visual system selectively sensitive to orientation and size of retinal images. *J Physiol (London)* 1969;203:237–260. [PubMed: 5821879]
18. Graham, NVS. *Oxford Psychology Series*. Oxford Univ. Press; 1989. *Visual Pattern Analyzers*; p. xvip. 646
19. Field DJ, Brady N. Visual sensitivity, blur and the sources of variability in the amplitude spectra of natural scenes. *Vision Res* 1997;37:3367–3383. [PubMed: 9425550]
20. Kersten D. Statistical efficiency for the detection of visual noise. *Vision Res* 1987;27:1029–1040. [PubMed: 3660646]
21. Levi DM, Klein SA, Chen IN. What is the signal in noise? *Vision Res* 2005;45:1835–1846. [PubMed: 15797773]

22. Solomon JA. Channel selection with non-white-noise masks. *J Opt Soc Am A* 2000;17:986–993.
23. Abbey CK, Eckstein MP. Estimates of human-observer templates for simple detection tasks in correlated noise. *Proc SPIE* 2000;3981:70–77.
24. Abbey CK, Eckstein MP. Classification images for simple a detection and discrimination tasks in correlated noise. *J Opt Soc Am A* 2007;24:B110–B124.
25. Burgess AE. Visual signal detection with two-component noise: low-pass spectrum effects. *J Opt Soc Am A* 1999;16:694–704.
26. Burgess AE, Li X, Abbey CK. Visual signal detectability with two noise components: anomalous masking effects. *J Opt Soc Am A* 1997;14:2420–2442.
27. Rolland JP, Barrett HH. Effect of random background inhomogeneity on observer detection performance. *J Opt Soc Am A* 1992;9:649–658. [PubMed: 1588452]
28. Abbey CK, Barrett HH. Human- and model-observer performance in ramp-spectrum noise: effects of regularization and object variability. *J Opt Soc Am A* 2001;18:473–488.
29. Myers KJ, Barrett HH, Borgstrom MC, Patton DD, Seeley GW. Effect of noise correlation on detectability of disk signals in medical imaging. *J Opt Soc Am A* 1985;2:1752–1759. [PubMed: 4056949]
30. Brainard DH. The psychophysics toolbox. *Spatial Vis* 1997;10:433–436.
31. Pelli DG. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vis* 1997;10:437–442.
32. Tyler CW, Chan H, Liu L, McBride B, Kontsevich L. Bit-stealing: how to get 1786 or more grey levels from an 8-bit color monitor. *Proc SPIE* 1992;1666:351–354.
33. Gold JM, Bennett PJ, Sekuler AB. Signal but not noise changes with perceptual learning. *Nature* 1999;402:176–178. [PubMed: 10647007]
34. Gold JM, Sekuler AB, Bennett PJ. Characterizing perceptual learning with external noise. *Cogn Sci* 2004;28:167–207.
35. Lu ZL, Doshier BA. Characterizing human perceptual inefficiencies with equivalent internal noise. *J Opt Soc Am A* 1999;16:764–778.
36. Pelli, DG. Ph.D. dissertation. University of Cambridge; 1981. Effects of visual noise.
37. Pelli, DG. The quantum efficiency of vision. In: Blakemore, C., editor. *Vision: Coding and Efficiency*. Cambridge Univ. Press; 1990. p. 3-24.
38. Pelli DG, Farell B. Why use noise? *J Opt Soc Am A* 1999;16:647–653.
39. Geisler, WS. Ideal observer analysis. In: Werner, JS.; Chalupa, LM., editors. *The Visual Neurosciences*. MIT Press; 2004. p. 2 v(various pagings)
40. Braje WL, Tjan BS, Legge GE. Human-efficiency for recognizing and detecting low-pass filtered objects. *Vision Res* 1995;35:2955–2966. [PubMed: 8533334]
41. Abbey CK, Eckstein MP. Classification image analysis: estimation and statistical inference for two-alternative forced-choice experiments. *J Vision* 2002;2:66–78.
42. Efron, B.; Tibshirani, R. *Monographs on Statistics and Applied Probability*. Vol. 57. Chapman & Hall; 1993. *An Introduction to the Bootstrap*.
43. Legge G, Kersten D, Burgess AE. Contrast discrimination in noise. *J Opt Soc Am A* 1987;4:391–406. [PubMed: 3559785]
44. Burgess AE, Colborne B. Visual signal detection. IV. Observer inconsistency. *J Opt Soc Am A* 1988;5:617–627. [PubMed: 3404312]
45. Doshier BA, Lu Z. Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proc Natl Acad Sci USA* 1998;95:13988–13993. [PubMed: 9811913]
46. Lu ZL, Doshier BA. External noise distinguishes attention mechanisms. *Vision Res* 1998;38:1183–1198. [PubMed: 9666987]
47. Ahumada AJ Jr. Classification image weights and internal noise level estimation. *J Vision* 2002;2:121–131.
48. Murray RF, Bennett PJ, Sekuler AB. Optimal methods for calculating classification images: weighted sums. *J Vision* 2002;2:79–104.

49. Watson AB. Multi-category classification: template models and classification images. *Invest Ophthalmol Visual Sci* 1998;39:S912.
50. Murray RF, Bennett PJ, Sekuler AB. Classification images predict absolute efficiency. *J Vision* 2005;5:139–149.
51. Neri P, Heeger DJ. Spatiotemporal mechanisms for detecting and identifying image features in human vision. *Nat Neurosci* 2002;5:812–816. [PubMed: 12101403]
52. Conrey B, Gold JM. An ideal observer analysis of variability in visual-only speech. *Vision Res* 2006;46:3243–3258. [PubMed: 16725171]
53. Eckstein MP, Shimozaki SS, Abbey CK. The footprints of visual attention in the Posner cueing paradigm revealed by classification images. *J Vision* 2002;2:25–45.
54. Pelli DG. Uncertainty explains many aspects of visual contrast detection and discrimination. *J Opt Soc Am A* 1985;2:1508–1532. [PubMed: 4045584]
55. Tjan BS, Nandy AS. Classification images with uncertainty. *J Vision* 2006;6:387–413.
56. Barth E, Beard BL, Ahumada AJ Jr. Nonlinear features in vernier acuity. *Proc SPIE* 1999;3644:88–96.
57. Solomon JA. Noise reveals visual mechanisms of detection and discrimination. *J Vision* 2002;2:105–120.
58. Ahumada AJ, Beard BL. Classification images for detection. *Invest Ophthalmol Visual Sci* 1999;40:3015.
59. Green DM. Consistency of auditory detection judgments. *Psychol Rev* 1964;71:392–407. [PubMed: 14208857]

APPENDIX A: PREWHITENING FOR AN IDEAL OBSERVER IN CORRELATED NOISE

We begin with a Gaussian white-noise field, N_W , with mean contrast 0 and variance σ^2 . To create the desired type of correlated noise, N_C , we filter N_W by the desired noise spectrum in the Fourier domain, maintaining the phase of N_W , as follows:

$$N_C = \mathfrak{F}^{-1}[\text{filt} \cdot A(N_W)e^{i\phi(N_W)}], \quad (\text{A1})$$

where $\mathfrak{F}^{-1}$ is the inverse Fourier transform, filt the amplitude spectrum of the filter, and $A(N_W)e^{i\phi(N_W)}$ the Fourier transform of the white noise in polar form.

On each trial, we add a contrast-adjusted signal S to the correlated noise field from Eq. (A1) to obtain the image I that will be displayed. For convenience in this derivation we show this addition in the Fourier domain, recalling that additions in the spatial and Fourier domains are equivalent:

$$I = S + N_C = \mathfrak{F}^{-1}[A(S)e^{i\phi(S)} + \text{filt} \cdot A(N_W)e^{i\phi(N_W)}], \quad (\text{A2})$$

with $A(S)e^{i\phi(S)}$ the Fourier transform of the signal S in polar form.

The ideal observer must take into account not only the possible templates but also the statistics of the correlated noise. This is most easily achieved through a process called “prewhitening,” which removes the noise correlations and also adjusts the possible templates to preserve the signal-to-noise ratio of the presented image I . Note that prewhitening therefore requires that the ideal observer know the form of the filter applied to produce the noise correlations.

Prewhitening is performed on the image I before submitting it to the decision rule. Essentially, I is reformulated from a signal plus a filtered noise field to an inverse-filtered signal plus a Gaussian white-noise field by dividing out the magnitude spectrum of the filter from Eq. (A2):

$$\begin{aligned} I_{C^{-1}} &= \mathfrak{F}^{-1} \left[\frac{A(S)}{filt} e^{i\phi(S)} + filt \cdot \frac{A(N_W)}{filt} e^{i\phi(N_W)} \right] \\ &= \mathfrak{F}^{-1} \left[\frac{A(S)}{filt} e^{i\phi(S)} + A(N_W) e^{i\phi(N_W)} \right]. \end{aligned} \quad (\text{A3})$$

The resulting prewhitened stimulus, $I_{C^{-1}}$, is equivalent to an inverse-filtered signal plus the original Gaussian white-noise field. The phase spectrum of I is preserved in $I_{C^{-1}}$, and the signal-to-noise ratio is equal for the original and the prewhitened stimuli across spatial frequencies and orientations:

$$SNR_{I_{C^{-1}}} = \frac{\frac{A(S)}{filt}}{A(N_W)} = \frac{A(S)}{filt \cdot A(N_W)} = SNR_I. \quad (\text{A4})$$

To ensure that the prewhitened stimulus $I_{C^{-1}}$ preserves the probability structure of the original image I relative to the ideal template T , the ideal template must also be inverse filtered. In other words, the appropriate comparator for $I_{C^{-1}}$ is not T but rather $T_{C^{-1}}$, the inverse-filtered template.

Equation (A3) leaves us with a signal presented in Gaussian white noise, and thus the ideal decision rule is the same as in previously published derivations for 1-of- N identification tasks and their special case, 1-of-2 identification [5,6].

APPENDIX B: CLASSIFICATION IMAGES IN CORRELATED NOISE

For this derivation, we assume that the observer's templates have equal energy and are adjusted in contrast to match the presented signal. Further, for simplicity, we assume that the observer does not have an internal noise source.

After applying the prewhitening procedure derived in Appendix A, we can write our decision variable, r :

$$r = (S_{C^{-1}} + N_W) \cdot T_{C^{-1}}, \quad (\text{B1})$$

where $S_{C^{-1}}$ is the presented signal with the inverse filter applied, N_W the original field of Gaussian white noise that was filtered to make the presented correlated noise, and $T_{C^{-1}}$ the ideal prewhitened template (i.e., the white-noise signal template with the inverse filter applied). Here, the dot indicates the dot product, i.e., $X \cdot Y = \sum_{i,j} X_{ij} Y_{ij}$.

If we make the substitutions $T' = T_{C^{-1}}$ and $S' = S_{C^{-1}}$, it becomes clear that r is equivalent to the formulation of the decision variable in Murray *et al.* [48], except that here r does not include an internal noise source:

$$r = (S' + N_W) \cdot T'. \quad (\text{B2})$$

Thus, once we have completed the prewhitening step, we can use the same methods for calculating classification images as described in Murray *et al.* [48]. In other words, we can calculate the classification images in the usual way based on the original fields of white noise. However, it should be noted that the estimated template $T' = T_C^{-1}$ reflects the prewhitening process and is not equivalent to the white-noise signal template.

Fig. 1. Leftmost column, the two oriented Gabor patch signals; remaining columns, the six three-dimensional object signals used in the experiment.

Fig. 2. Example noise fields and the associated filters for each of the external noise conditions. Amplitude is not to scale but was normalized to have unit power on average for each spatial frequency and orientation.

Fig. 3.

Top row, human, ideal observer, and noise-limited ideal observer thresholds for the Gabor (left) and object (right) recognition tasks. Bottom row, corresponding efficiencies for the Gabor and object recognition tasks. Error bars correspond to ± 1 standard deviation, estimated by bootstrap simulations [42].

Fig. 4. Efficiencies for the Gabor (left panel) and the 1-of-2 object (right panel) recognition tasks, computed from the response classification data.

Fig. 5.

Classification images calculated from the ideal observer templates (top row) and from the combined human observer data (middle and bottom rows) for the Gabor patch task in all three noise conditions. Human images in the middle row were calculated from the raw data. For the images in the bottom row, the raw data were smoothed with a 9×9 convolution kernel (the matrix product of the vector $[1\ 2\ 3\ 4\ 5\ 4\ 3\ 2\ 1]$ with itself transposed) so as to make it easier to visualize the structure in the images.

Fig. 6.

Classification images calculated from the ideal observer templates (top row) and from the combined human observer data (middle and bottom rows) for the object task in all three noise conditions. Human images in the middle row were calculated from the raw data. For the images in the bottom row, the raw data were smoothed as described for Fig. 5.

Fig. 7.
Actual/predicted efficiency ratios measured from the response classification data in the Gabor and object tasks.

Fig. 8. Classification subimages for human and simulated ideal observers in the Gabor patch task in the white (top panels), $1/f$ (middle panels), and *flipped f* (bottom panels) noise conditions.

Fig. 9. Classification subimages for human and simulated ideal observers in the object task in the white (top panels), $1/f$ (middle panels), and *flipped f* (bottom panels) noise conditions.

Fig. 10.

Response consistency plots for observers JL (upper panels) and BC (lower panels) in the Gabor task in the white (left), *1/f* (middle), and *flipped f* (right) conditions.

Fig. 11.

Response consistency plots for observers SK (upper panels) and BC (lower panels) in the object task in the white (left), *1/f* (middle), and *flipped f* (right) conditions.