How attention extracts objects from noise

Michael S. Pratte, Sam Ling, Jascha D. Swisher, and Frank Tong
Psychology Department and Vanderbilt Vision Research Center, Vanderbilt University, Nashville, Tennessee

Pratte MS, Ling S, Swisher JD, Tong F. How attention extracts objects from noise. *J Neurophysiol* 110: 1346–1356, 2013. First published June 26, 2013; doi:10.1152/jn.00127.2013.—The visual system is remarkably proficient at extracting relevant object information from noisy, cluttered environments. Although attention is known to enhance sensory processing, the mechanisms by which attention extracts relevant information from noise are not well understood. According to the perceptual template model, attention may act to amplify responses to all visual input, or it may act as a noise filter, dampening responses to irrelevant visual noise. Amplification allows for improved performance in the absence of visual noise, whereas a noise-filtering mechanism can only improve performance if the target stimulus appears in noise. Here, we used fMRI to investigate how attention modulates cortical responses to objects at multiple levels of the visual pathway. Participants viewed images of faces, houses, chairs, and shoes, presented in various levels of visual noise. We used multivoxel pattern analysis to predict the viewed object category, for attended and unattended stimuli, from cortical activity patterns in individual visual areas. Early visual areas, V1 and V2, exhibited a benefit of attention only at high levels of visual noise, suggesting that attention operates via a noise-filtering mechanism at these early sites. By contrast, attention led to enhanced processing of noise-free images (i.e., amplification) only in higher visual areas, including area V4, fusiform face area, mid-Fusiform area, and the lateral occipital cortex. Together, these results suggest that attention improves people's ability to discriminate objects by de-noising visual input in early visual areas and amplifying this noise-reduced signal at higher stages of visual processing.

selective attention; equivalent noise; decoding; noise reduction; primary visual cortex; functional magnetic resonance imaging

TO IDENTIFY OBJECTS in real world settings, the visual system must extract relevant information in the presence of visual clutter or noise. Attention underlies our ability to selectively process sensory information (for reviews see Carrasco 2011; Desimone and Duncan 1995; Maunsell and Treue 2006; Reynolds and Chelazzi 2004). However, the mechanisms by which attention separates task-relevant information from irrelevant noise remains unclear.

There are two mechanisms by which attention can improve sensory processing: amplification and noise filtering. An amplification mechanism boosts all visual input, simultaneously increasing the sensory gain of both the relevant signal and any irrelevant noise at the attended location (sometimes referred to as signal enhancement). A noise-filtering mechanism selectively reduces responses to the noise component of the visual input while preserving the relevant signal (sometimes referred to as external noise reduction). These attentional processes have been formally specified within the perceptual template model, proposed by Lu and Dosher (1998), to characterize

human psychophysical performance (see also Dosher and Lu 2000b). According to this model, perceptual performance may be limited by the level of external noise (e.g., white noise) present in the stimulus, or by the level of internal noise present in the visual system. For conditions in which the relevant signal is embedded in levels of external noise that considerably exceed the visual system's level of internal noise, only a noise-filtering mechanism will be capable of improving the processing of that signal. Only noise filtering is effective in this regime, since an amplification mechanism would increase responses to both the signal and the external noise, yielding no net gain in the signal-to-noise ratio (SNR). By contrast, for situations in which the level of external noise is lower than the visual system's level of internal noise, a nonselective amplification of the input will effectively increase the SNR of processing in the visual system. Such amplification will increase the visual system's response to both signal and noise, but as long as the amplified noise level remains below the level of internal noise, then the benefits of amplifying the relevant signal will outweigh the costs of amplifying the external noise component. Thus, an amplification mechanism can enhance performance when external noise falls to low or negligible levels, whereas a noise-filtering mechanism can only enhance performance when the relevant signal to be processed is embedded in substantial external noise.

Noise-filtering and amplification mechanisms can be assessed behaviorally by measuring whether attention enhances an observer's ability to perceive a stimulus in the presence of varying levels of external noise. If attention improves behavioral performance at low noise levels, such effects must arise from amplification (see Fig. 1). Conversely, attentional benefits that occur only at high noise levels would imply a noise-filtering mechanism. By manipulating both external noise and attention, studies of behavioral performance have found that attention sometimes acts as a noise filter (Dosher and Lu 2000b; Lu and Dosher 2000; Lu et al. 2002; Smith et al. 2004), sometimes as an amplifier (Carrasco et al. 2000; Lu and Dosher 1998, 2000; Lu et al. 2000; Morrone et al. 2002), and sometimes as both (Ling et al. 2009).

The goal of the present study was to determine whether attention leads to amplification or noise filtering at various stages of processing in the visual hierarchy, using fMRI and multivoxel pattern classification (Tong and Pratte 2012). Participants viewed line-drawn images of objects belonging to one of four categories (faces, houses, chairs, and shoes, see Fig. 2A). The objects were presented with varying levels of visual noise and were either attended to or ignored. Cortical activity patterns within individual visual areas were analyzed by training a multivoxel pattern classifier to discriminate object category, separately for each noise level and attention condition. Classification performance served as a measure of the amount of
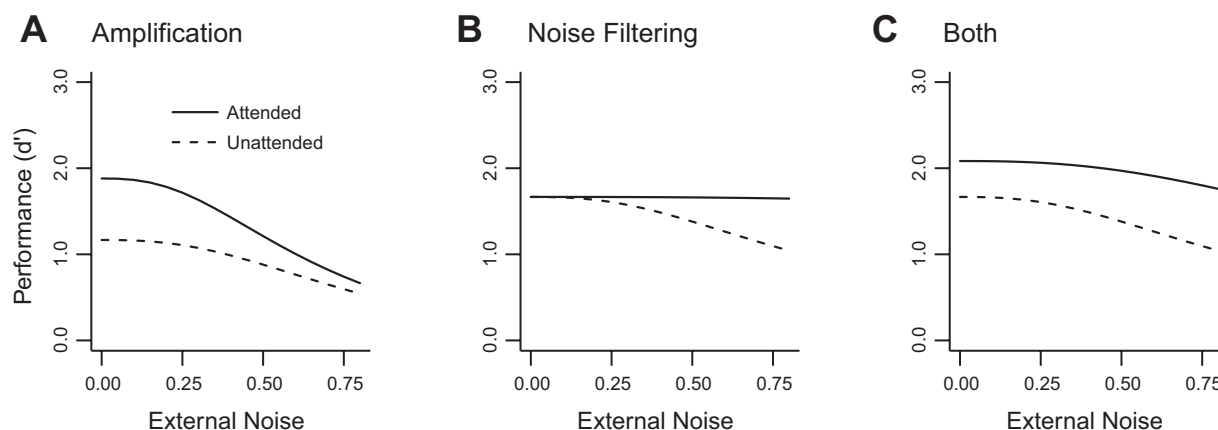
Fig. 1. Predictions of the perceptual template model. Performance is plotted as a function of external noise for attended (solid curves) and unattended (dashed curves) conditions. Performance could be a behavioral measure or the accuracy of fMRI pattern classification. *A*: pure amplification is evidenced by an attentional benefit at low noise levels, but no effect at high noise. *B*: pure noise filtering is evidenced by an attentional benefit at high levels of noise, but no effect at low noise levels. *C*: if both amplification and noise-filtering effects are present, then attention is expected to increase performance accuracy across all noise levels.

object-discriminating information present within the activity patterns for each visual area and condition. Such discriminating information could arise from selectivity to the local low-level features present in the images, or to higher-level object properties (Kietzmann et al. 2012; Kriegeskorte et al. 2008; Naselaris et al. 2009). Regardless of whether a visual area is sensitive to low-level features or complex object properties, this classification approach allowed us to compare performance across attended and unattended conditions within a visual area, and thereby to evaluate whether the effects of amplification or noise filtering are evident in a given cortical visual area. Whereas an amplification mechanism should lead to improved discrimination of objects presented in the absence of visual noise, a noise-filtering mechanism can only lead to improved performance for the target objects embedded in substantial levels of noise. The perceptual template model provides
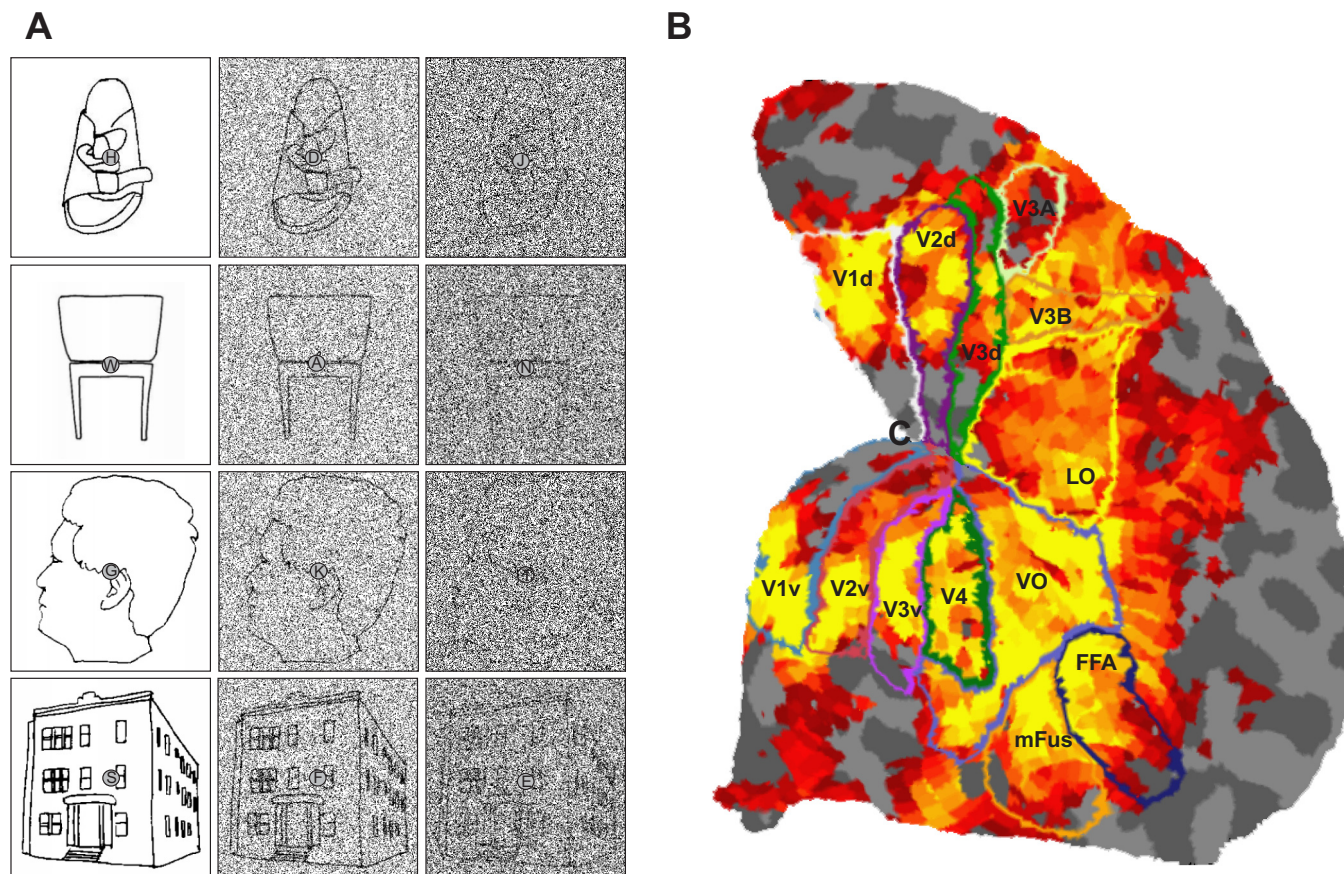

Fig. 2. Experimental stimuli and cortical regions of interest. *A*: stimuli consisted of images of shoes, chairs, faces, and houses. The *left* column shows images with zero noise, the *middle* and *right* columns show the same images with 50% and 75% noise, respectively. The objects were still identifiable in the highest noise condition. *B*: a flattened representation of the right occipital lobe for 1 participant. The activity map shows significant BOLD activity in the visual localizer experiment. VO, ventral occipital cortex; LO, lateral occipital cortex; FFA, fusiform face area; mFus, mid-Fusiform area.

straightforward quantitative predictions regarding how amplification and noise filtering should affect classification performance, based on fundamental principles of signal detection theory. Although this systems-level model does not specify which of several possible neural mechanisms might underlie a complex process such as noise filtering, its application to neural data can prove useful by allowing one to distinguish whether a given brain area exhibits the predicted effects of amplification, noise filtering, or both processes.

We hypothesized that the visual system acts as a signal-processing stream in which the initial sensory input undergoes active noise filtering at an early processing stage, and this noise-filtered signal is subsequently amplified at higher stages of processing. Such an organization would ensure that the system is capable of increasing the SNR of visual processing under both low- and high-noise conditions. Specifically, we predicted that attention would enhance processing in early visual areas only when objects were presented in high levels of visual noise, implying an early-stage noise-filtering mechanism, whereas attention would enhance object processing in high-level visual areas even in the absence of noise, as a result of amplification.

By applying the perceptual template model to the classification performance of individual visual areas, we can quantify the extent to which attention leads to effects of noise filtering and amplification at successive levels of the human visual pathway. An important consideration is that the response properties of any visual area will reflect processing that has occurred at prior stages as well as the processes that occur in that area proper. Here, we focused on how the effects of amplification and noise filtering change across the visual hierarchy to gain insight into the stages at which these processes likely emerge.

MATERIALS AND METHODS

Participants. Six healthy adult volunteers (4 males) aged 22–39, with normal or corrected-to-normal vision, participated in the experiment. All participants gave informed written consent. The study was approved by the Vanderbilt University Institutional Review Board.

Apparatus. The stimuli were generated using Matlab and the Psychophysics Toolbox (Brainard 1997; Pelli 1997). Stimuli were displayed on a rear-projection screen using a luminance-calibrated Eiki LC-X60 LCD projector with a Navitar zoom lens. Participants used a custom bite-bar system to minimize head movements.

Scanning was performed using a Philips 3-Tesla Intera Achieva MRI scanner with an eight-channel head coil, located at the Vanderbilt University Institute for Imaging Science. We used standard gradient-echo echoplanar T2*-weighted imaging to acquire whole brain functional images (TR 2,000 ms; TE 30 ms; 35 slices with no gap; 64×64 in-plane resolution; 3-mm isotropic voxels; $1.5 \times$ SENSE acceleration).

Stimuli and design. The main experiment consisted of 18 functional runs. Observers were instructed to maintain fixation on a central fixation bull's eye (1° diameter) where sequences of letters appeared (0.8° diameter), while line drawings of faces, houses, shoes, or chairs (spanning 12° of visual angle) were concurrently presented. Each object set consisted of 16 different line drawings of objects from that category. The images within each category were more similar to one another than to images in other categories (e.g., image category could be classified based on the pixel values of the images), and these category differences can thereby allow for high classification performance in early visual areas. The images were embedded in noise by randomly changing 0%, 50%, or 75% of the pixels to black or white

(with equal probability) in the 0, 0.5, and 0.75 noise conditions, respectively (Fig. 2*A*). In behavioral pilot studies, we confirmed that the line-drawn objects could still be perceived at the highest level of visual noise.

Each experimental run consisted of 16 consecutive stimulus blocks that lasted 14 s each, with a 16-s fixation-rest period at the beginning and end of the run. For each stimulus block, the objects were either attended or unattended. Each block was preceded by 1 s of fixation, followed by a 1-s cue indicating whether participants should attend to the objects or to the central letters. Following this instruction, 14 randomly selected images from a given category and noise level were presented for 600 ms each, separated by a 400-ms interstimulus interval, while central letters were presented at a rate of 5 items per second. In the unattended blocks, participants performed a letter discrimination task and had to report whenever a "J" or "K" appeared in the temporal sequence of letters by pressing a corresponding key on a button box (target frequency, 10%; average hit rate, 90%). In the attended blocks, participants performed a one-back object-matching task, and had to make a button-press whenever the same image was shown consecutively (20% of all image presentations). Average hit rate for this task was 91% and did not vary as a function of external noise level [$F(2, 5) = 0.92, P = 0.43$].

Stimulus category and task were counterbalanced within each run such that each combination of task and category occurred twice per run. The attention condition alternated between the one-back task and the letter discrimination task across successive blocks, with each run beginning with the one-back task. Critically, stimulus parameters were identical across attention conditions; only the instructions changed. Noise level was varied across runs, and each session began with zero noise followed by 0.50 and 0.75.

The experimental session also included two functional localizer runs used to identify voxels that responded to the region occupied by the stimuli. For the localizer runs, images were randomly selected from all possible categories and noise levels and were presented in stimulus blocks of 12 s, separated by 12 s of fixation (no letters were presented at fixation). Standard linear modeling procedures were used to quantify the extent to which each voxel responded to the images during the localizer, providing a measure of the extent to which each voxel responded to the stimulus region regardless of image category.

fMRI preprocessing. Functional images were simultaneously motion-corrected and aligned to the mean of the eighth run using FSL's MCFLIRT (Jenkinson et al. 2002) and high-pass filtered (50-s cutoff) to remove slow drifts in the MRI signal. Images were then aligned to the surface using Freesurfer's boundary-based registration (Greve and Fischl 2009).

Regions of interest. Each participant took part in a separate scanning session for functional mapping of each region of interest. Retinotopic visual areas V1, V2, V3, V3A, V3B, and V4 were defined on the flattened cortical surface using standard polar retinotopic mapping procedures (Engel et al. 1997; Sereno et al. 1995). To define category-selective object areas, in two additional runs, images of faces, houses, objects, and visually scrambled objects were presented in 12-s blocks separated by 12 s of fixation. Voxels that responded more to faces than objects were visualized on the cortical surface to define the fusiform face area (Kanwisher et al. 1997), and voxels responding more to objects than faces were used to define the mid-fusiform gyrus (Grill-Spector 2003). Voxels on the ventral and lateral cortical surfaces that responded to the experimental localizer, but did not fall within any of these retinotopic or category-selective regions, were labeled as ventral occipital cortex and lateral occipital cortex, respectively (Fig. 2*B*).

Pattern classification analysis. Linear support vector machines (SVM) (Boser et al. 1992; Haynes and Rees 2006; Kamitani and Tong 2005; Norman et al. 2006; Tong and Pratte 2012; Vapnik 1998) were used to classify the viewed object category. The classifiers were trained and tested separately for each visual area, noise level, attentional condition, and participant. The 120 voxels that responded most

strongly to the experimental localizer within each visual area were used as features for classification; results were highly similar across a range of choices for number of voxels included. Responses of each voxel for each block were obtained by temporally averaging the BOLD signal from 6 s after the presentation of the first image (to account for hemodynamic lag) up to and including the fMRI volume that was acquired at the end of the 14-s stimulus block. A leave-one-run-out cross-validation procedure was used to evaluate the amount of object-discriminating information that could be decoded from activity patterns in each visual area. Specifically, a four-class SVM was trained on blocks from all but one run and used to predict the object category of blocks in the left-out test run. Classification accuracy was obtained by averaging performance accuracy across all test runs. Accuracy from the four-choice classifier was converted to d′ units (Ingleby 1973) so that the formal perceptual template model (presented below) could be applied. Results were qualitatively similar when multiclass classification accuracy was considered rather than d′.

Perceptual template model. By measuring the effects of attention on classification performance across noise conditions, the influence of amplification and noise filtering can be formally assessed within a cortical region. The signature of amplification is an attentional benefit for processing stimuli in the absence of external noise, and the signature of external noise filtering is an attentional benefit for processing stimuli embedded in high levels of external noise. In principle, one can simply evaluate performance in the presence of no noise and in high noise to determine the role of these attentional mechanisms (Dosher and Lu 2000a). In our study, we also included an intermediate noise-level condition, which allowed us to fit the fMRI classification data from each visual area with a variant of the perceptual template model (Lu and Dosher 1998). Fitting this model to the fMRI data provided formal estimates of amplification and noise reduction in each visual area, and allowed us to measure how well the formal model is able to characterize the effects of attention on fMRI classification performance.

Our application of this model quantifies how classification performance changes as a function of external noise (N_E) and attention (k), within a visual area. For the unattended condition ($k = 0$), the model included parameters for internal noise (N_I), a scale term (β), and a non-linearity term (γ). The attended condition ($k = 1$) was defined by these same values, modified by amplification (a_{AMP}) and noise filtering (a_{NF}) parameters:

$$d'(N_E, k) = \frac{\beta \times (a_{AMP})^{k\gamma}}{\sqrt{\left(\frac{N_E \times (a_{AMP})^k}{(a_{NF})^k}\right)^{2\gamma} + N_I^2}}$$

Noise filtering improves discriminability by attenuating the impact of the external noise by a scale factor. Amplification improves discrim-

inability by increasing the gain of the stimulus, but it also increases the gain of the noise. Figure 1 shows predicted d′ values plotted as a function of external noise level for a model with reasonable parameter values. In this study, the model was separately fit to the data from each visual area and participant by minimizing squared error. For each visual area, the non-linearity parameter was fixed across participants to aid in parameter identification; all other parameters were free to vary across participants.

RESULTS

Analysis of mean BOLD amplitude. We first examined the effects of external noise and attention on the overall amplitude of the BOLD signal across visual areas (Fig. 3). The statistical significance of these effects was assessed using a repeated-measures ANOVA with visual area, noise condition, and attention condition serving as separate factors. In early visual areas V1 and V2, the addition of external noise increased the BOLD response, as was expected given the higher contrast energy of the noisy images. Higher visual areas, however, exhibited the opposite pattern of results. The addition of noise tended to reduce the mean BOLD response in higher areas, presumably because these areas are comparatively insensitive to the amount of stimulus energy but are more sensitive to the loss of coherent contours and other object-related information that accompanies the addition of noise. This crossover inter-action effect between noise level and visual area was statistically significant [F(16, 80) = 8.8, P < 0.001], suggesting that embedding objects in visual noise modulates the overall BOLD response in qualitatively different ways in low- and high-level visual areas.

Directing attention to the objects led to a marked increase in mean BOLD activity across all visual areas [main effect of attention: F(1, 5) = 47.1, P < 0.001], as has previously been shown (Beauchamp et al. 1997; Chawla et al. 1999; Gandhi et al. 1999; Somers et al. 1999). This attentional effect was greater in retinotopic visual areas than in higher visual areas [attention by area interaction: F(8, 40) = 10.0, P < 0.001]. This pattern of results may simply reflect the fact that the visual localizer excluded the fovea in retinotopic areas, but may not have done so in higher-level areas, resulting in smaller than expected effects of attentional modulation in higher areas.

Critically, the magnitude of the attentional effect did not depend on noise level [attention by noise interaction: F(2, 10) = 1.46, P =
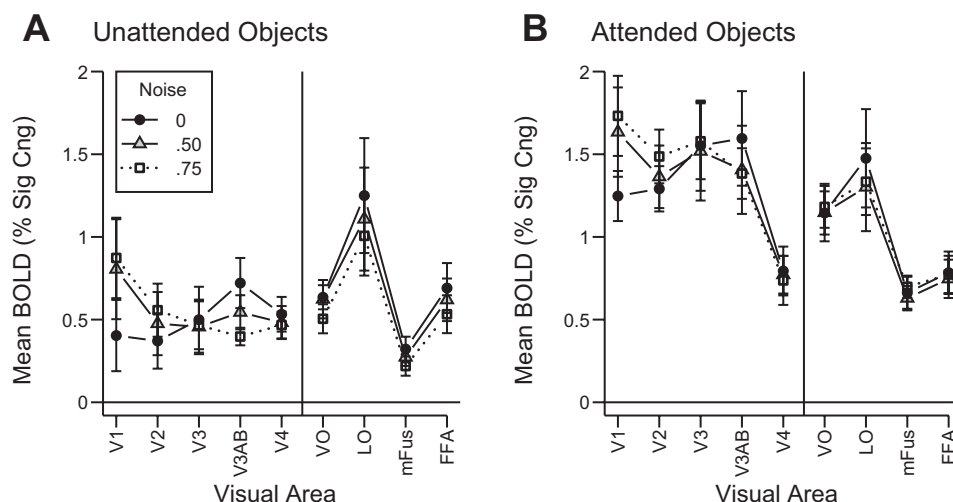

A Unattended Objects

B Attended Objects

Fig. 3. Mean BOLD activity for each noise condition, plotted by visual area separately for unattended (*A*) and attended (*B*) conditions. Error bars denote ±1 SE. Raw BOLD amplitudes were converted to percent signal change units by dividing the mean time course of each voxel by the mean signal for that voxel over the first 16 s of fixation that began each run.

0.28]. Thus, mean BOLD activity did not prove useful for discerning how attention might act to overcome the effects of external noise, as the magnitude of attentional modulation was comparable across noise levels. These findings can be contrasted with the much more informative results we obtained through pattern classification analysis, described below.

Multivoxel pattern classification. We applied multivoxel pattern classification to the activity patterns in each region of interest, to identify the object category that was viewed on each 14-s stimulus block. This analysis allowed us to assess the reliability of cortical activity patterns across attended and unattended conditions, and to evaluate the impact of external visual noise on the reliability of the activity patterns evoked by the distinct sets of object stimuli (i.e., faces, houses, chairs, and shoes). It should be emphasized that the object-selective information found in these activity patterns could arise from sensitivity to local low-level features, as would predominate in early visual areas, or from sensitivity to higher-order shape and object properties, as has been found in high-level visual areas (Kietzmann et al. 2012; Kriegeskorte et al. 2008; Naselaris et al. 2009). Our analyses focused on changes in classification performance within each visual area, to determine the impact of attention on the reliability of visual processing in that region across different levels of external visual noise.

Figure 4, *A* and *B*, shows pattern classification performance for the unattended and attended conditions, respectively. We first highlight the effects of external noise on classification performance. Whereas the addition of external noise led to increases in mean BOLD amplitude in early visual areas, increasing the amount of external noise led to clear decreases in classification performance across all visual areas for both attended and unattended objects [main effect of noise: $F_{(2, 10)} = 121.8$, $P < 0.001$]. This dissociation in the pattern of results highlights the fact that, unlike mean BOLD amplitude, pattern classification accuracy provides a measure of the amount of object-discriminating information available in the activity patterns of a visual area for a given experimental condition. The disruptive effect of external noise was substantial in early areas and more modest in high-level areas [noise by area interaction: $F_{(16, 80)} = 15.4$, $P < 0.001$]. This trend in the data was most clear when comparing classification accuracy across the no-noise and medium-noise conditions. In early visual areas V1, V2, and V3, moderate external noise led to a large drop in
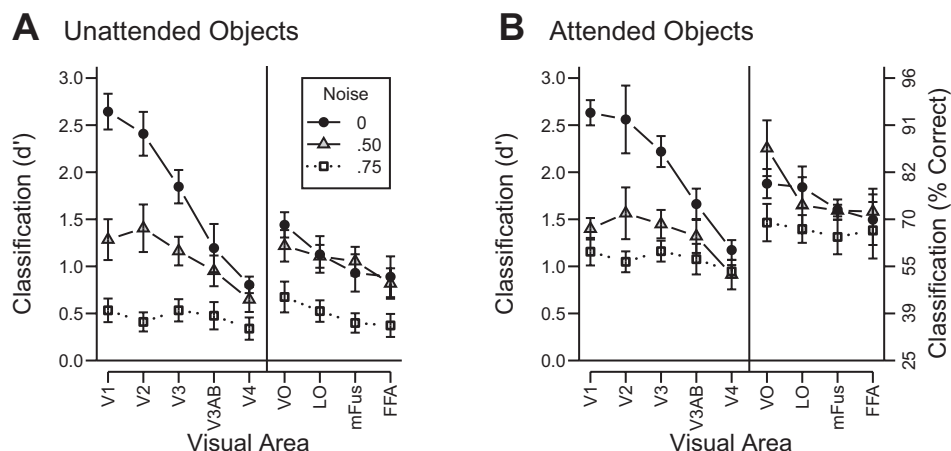
classification accuracy, relative to the noise-free condition, by ~1.0 d′ unit. In contrast, higher visual areas exhibited virtually no difference between the no-noise and medium-noise conditions. Only at the highest level of noise did we observe impaired processing of object information in these high-level areas. The greater sensitivity of early visual areas to external noise likely reflects the fact that the visual representations in these areas are based on low-level visual properties, such as local contrast and orientation, which may be more easily disrupted by the addition of visual noise. The weaker effects of noise in higher-level visual areas is consistent with their sensitivity to complex object properties, which would allow for integration of visual information over larger portions of the visual field.

What impact does attention have on the neural representation of objects in noise? A comparison of Fig. 4, *A* and *B*, reveals that attending to the objects generally increased classification performance [main effect of attention: $F_{(1, 5)} = 178.2$, $P < 0.001$], and that these attentional benefits were greater in high-level areas than in low-level areas [attention by area interaction: $F_{(8, 40)} = 3.6$, $P < 0.05$]. Attention also attenuated the deleterious effects of noise across all visual areas [noise by attention interaction: $F_{(2, 10)} = 4.2$, $P < 0.05$], leading to greater improvements in classification performance for noisy than for noise-free conditions. These results indicate that attention can help to overcome external noise and enhance the representation of object-specific feature information in both low- and high-level visual areas.

Of critical interest was whether the attentional effects observed in each visual area were due to amplification, noise filtering, or a combination of both mechanisms. Recall that an amplification mechanism will produce attentional benefits at low noise levels, whereas a noise-filtering mechanism will produce benefits only at high noise levels. To highlight the differential impact of these two mechanisms, we calculated the extent to which attention enhanced object-discrimination performance for each visual area, separately for high-noise images and no-noise images. In the presence of high external noise, attention improved classification performance in all visual areas of interest (Fig. 5*A*). Thus, attentional benefits of noise filtering were evident throughout the cortical visual pathway, with modulatory effects emerging as early as V1.

In contrast, we found no effect of attentional enhancement in early visual areas (V1 and V2) during the viewing of noise-free

Fig. 4. fMRI pattern classification accuracy plotted across visual areas for each noise level in the unattended (*A*) and attended (*B*) conditions. Error bars denote standard errors. Data are in d′ units (*left* axis); corresponding 4-choice accuracy is shown on the rightmost axis.
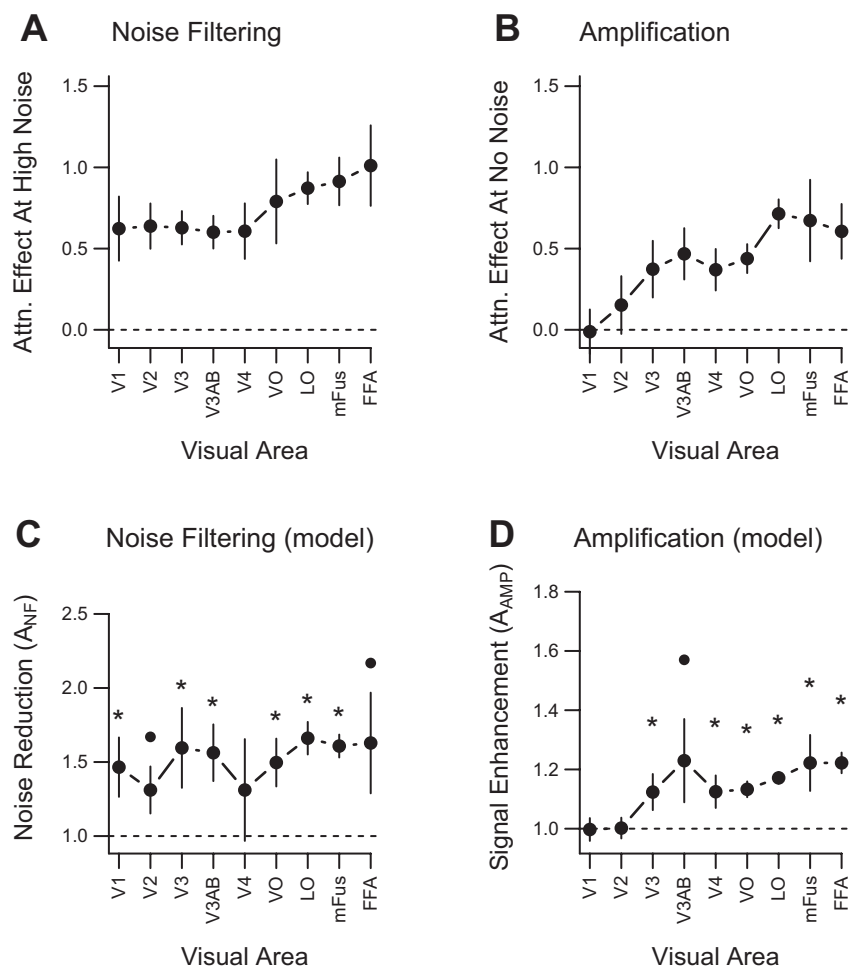
Fig. 5. Measures of noise filtering and amplification for individual visual areas. *A*: attention effect (attended-unattended classification performance) for each visual area in the high-noise condition. *B*: attention effect in the no-noise condition. *C*: parameter estimates of noise filtering from perceptual template model (PTM) fits to each visual area. *D*: parameter estimates of amplification from PTM fits. Parameter estimates that significantly exceed 1.0 indicate a reliable multiplicative effect; asterisks and black circles denote significant ($P < 0.05$) and marginally significant ($P < 0.1$) effects, respectively. Error bars denote standard errors.

objects, or even for objects in moderate levels of noise. As can be seen in Fig. 5*B*, attention enhanced classification performance for noise-free images only in higher visual areas. These results are consistent with the proposal that amplification is absent in early visual areas, and only emerges at higher stages of visual processing. However, the claim of a lack of attentional effects in early areas warrants careful inspection, as the conclusions are based on a failure to reject the null hypothesis. To determine whether the null hypothesis is likely to be true, we used a Bayesian technique akin to the *t*-test (Rouder et al. 2009) to evaluate the likelihood that there is no attentional effect in early areas for noiseless images. This method provides the Bayes factor, which reflects how much more likely the null hypothesis is to be true than the alternative hypothesis of the existence of an effect. Bayes factors of 3.0 or greater are typically interpreted as "substantial" evidence for the null hypothesis (Jeffreys 1961). For area V1, the Bayes factor was 3.5, implying that the null hypothesis of no amplification was 3.5 times more likely than the alternative of amplification being present. For V2, the null was 3.0 times more likely than the alternative, again suggesting a lack of amplification in early visual areas.

Although this Bayesian analysis mitigates statistical problems associated with accepting the null, it might be argued that our failure to observe a benefit of attention for noise-free images was caused by a ceiling effect. For example, classification accuracy of noise-free images for V1 was ~92%, for both attended and unattended conditions. However, a ceiling effect is unlikely to account for our results, as we also found no evidence for attentional enhancement of classification performance in V1 or V2 for objects presented in medium levels of noise. Again, Bayes factors for V1 and V2 were 2.5 and 3.0 in favor of the null hypothesis, respectively, even though classification accuracies for these medium-noise conditions were well below ceiling.

To further quantify the extent to which these attentional mechanisms could account for classification performance in each visual area, we fitted the perceptual template model to the fMRI classification data (see MATERIALS AND METHODS for details). This model provides an estimate of amplification and noise filtering parameters to account for the effects of attention on discrimination performance across a range of noise levels.

Figure 6 shows classification performance as a function of noise level for each visual area. Overall, the model fits the data well, accounting for 89–95% of the within-subject variance. Figure 5*C* shows the noise filtering parameter estimates for each visual area, and Fig. 5*D* shows the amplification parameter estimates. This parametric analysis is highly concordant with the conclusions drawn from assessing the no-noise and high-noise conditions: noise filtering emerges at the level of V1 and is evident in all downstream visual areas, whereas amplification emerges later in the visual hierarchy.

Is it possible that attention led to enhanced classification performance in early visual areas, not by selectively enhancing
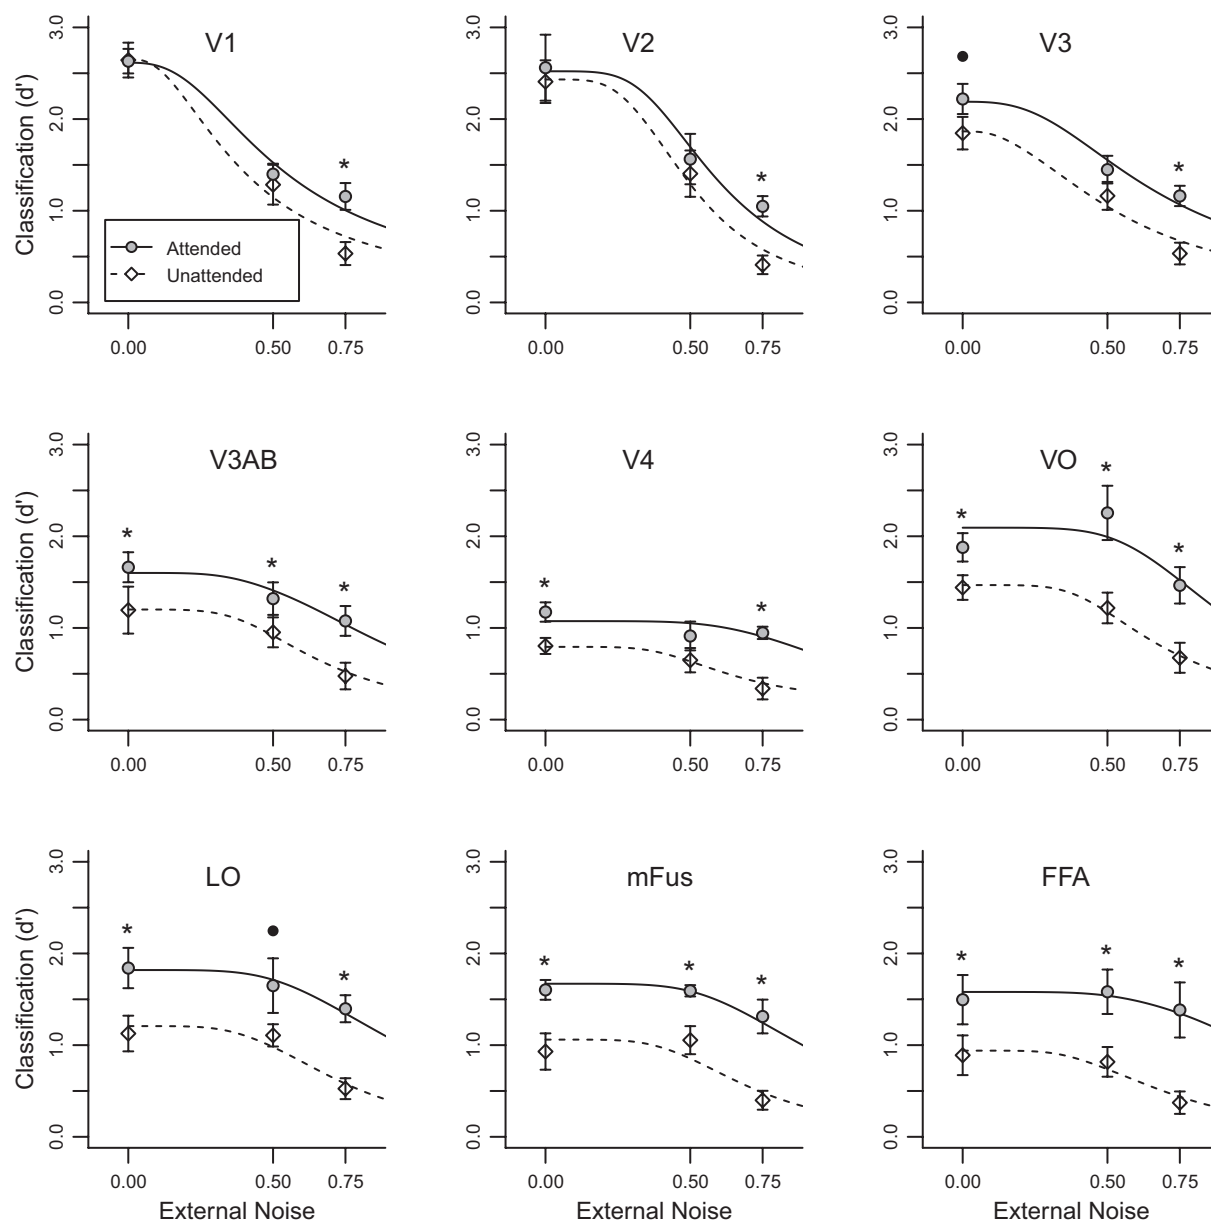
Fig. 6. Classification accuracy in each visual area plotted as a function of external noise level for attended and unattended conditions. Points denote the data with standard errors; asterisks and black circles denote statistically significant ($P < 0.05$) and marginally significant ($P < 0.1$) attentional effects, respectively, for each area and noise level. Curves indicate PTM fits to the fMRI classification data, obtained by generating predictions for each participant and averaging over participants. Whereas the pattern of results in areas V1 and V2 are consistent with pure noise reduction, higher-level areas exhibit both noise reduction and amplification.

responses to the low-level features contained in the attended objects, but rather, by the feedback of a more abstract attentional signal that would be distinct from the pattern of feedforward stimulus-driven responses? A recent study reported that attending to the shape of peripheral objects can lead to object-specific patterns of activity in the foveal representation of early visual areas (Williams et al. 2008). Interestingly, these activity patterns remained stable across changes in object position, implying that a more abstract object-selective signal was being fed back to early visual areas in a manner independent of the low-level visual input.

We consider it unlikely that the feedback effects found in the Williams et al. (2008) study can account for the present findings, as such an account fails to explain why we observed

attentional enhancement in V1 and V2 only in the highest noise condition, with no enhancement under conditions of moderate or no noise. Moreover, our visual localizer produced regions of interest that largely excluded the foveal representation (see Fig. 2B). Nonetheless, to further address this possibility we conducted an analysis to determine whether stimulus-driven responses to unattended objects were effectively enhanced by top-down attention. To do this, we first trained a classifier to discriminate object category using activity patterns in the zero-noise, unattended condition. This trained classifier was then used to predict the viewed object category in all other conditions. Since the classifier used for training was based on the stimulus-driven activity patterns evoked by unattended objects (i.e., unattended signals in the absence of noise), boosts in classification perfor-

mance with attention can be attributed to top-down enhancement of stimulus-driven responses.

The results of this control analysis are shown in Fig. 7. Performance in all conditions was very similar to our original analysis when classifiers were trained separately on each condition (cf. Fig. 4). Of critical interest, we found that attentional effects for the high-noise condition were reliable in both V1 [$t(5) = 3.29$, $P = 0.019$] and V2 [$t(5) = 3.17$, $P = 0.025$], implying that top-down attention serves to enhance the pattern of stimulus-driven responses in early visual areas under conditions of high noise. These findings are consistent with an active noise-filtering mechanism in which feedforward responses to the relevant signal (i.e., the object) are selectively enhanced, relative to responses to the noise.

DISCUSSION

Many recent studies have utilized fMRI and multivoxel pattern analysis to explore how attention can enhance visual processing (Cukur et al. 2013; Jehee et al. 2011; Kamitani and Tong 2005; Peelen et al. 2009; Reddy et al. 2009; Scolari et al. 2012; Serences and Boynton 2007). Here, we expanded this approach by examining the role of attention in extracting information about complex objects in the presence of visual noise. We found that activity patterns in higher visual areas could better distinguish between attended than unattended objects, but that these benefits of attentional processing were equally prominent for noisy and noise-free viewing conditions. In contrast, early visual areas V1 and V2 exhibited improved processing with attention only when objects were presented in high levels of external visual noise. The noise dependency of these attentional effects in V1/V2 was all the more striking given that the mean BOLD response was similarly boosted by attention, regardless of noise level. Our results suggest that an important role of top-down feedback to early visual areas is to extract relevant signals from irrelevant noise. Moreover, these findings raise the possibility that certain effects of attentional feedback in early visual areas might fail to be detected if stimulus conditions are insufficiently noisy or perceptually challenging.

By examining these fMRI data within the signal-processing framework of the perceptual template model (Lu and Dosher 1998), we were able to separately measure the attentional contributions of noise filtering and amplification. These quantitative analyses provided compelling evidence of differences in processing across the visual hierarchy: whereas noise filtering was clearly evident at the earliest stages of visual processing, effects of amplification were found only in higher-level visual areas. These findings support our proposal that top-down attentional mechanisms operate along the visual hierarchy by applying noise filtering at early sites of the visual processing stream, followed by an amplification stage. From a signal processing perspective, such a design would allow the visual system to be well suited at extracting relevant information from both noisy and noiseless inputs. If amplification were instead to precede noise filtering, there would be no benefit of the amplification stage during the processing of noisy visual inputs.

Although we found signatures of noise filtering throughout the visual pathway, this result does not necessarily imply that active noise filtering occurs in every one of these visual areas. Given that cortical areas downstream would inherit any modulation of visual processing from early visual areas, a simpler account is that the reduction of external noise occurred primarily in early visual cortices, with the resulting noise-reduced signal propagated to higher visual areas. Our model-based estimates of attentional amplification suggest a further distinction between low- and high-level visual areas, with amplification emerging only at higher stages of the visual pathway, beyond V1 and V2. The combined effects of noise filtering and amplification provide a good account of the pattern of results in high-level visual areas, in which attentional enhancement occurred across all noise levels (see Fig. 6). In theory, enhancement across all noise levels could also result from a reduction in (multiplicative) internal noise, but psychophysical studies have provided compelling evidence to rule out this possibility (Dosher and Lu 2000a). Moreover, the pattern of results we observe across the visual hierarchy is better explained in terms of sequential stages of noise filtering and amplification.

Although areas V1 and V2 showed no attentional enhancement for objects in the absence of noise, some previous fMRI studies have found attentional enhancement of pattern-specific responses in early visual areas under noise-free conditions. However, these studies typically involved simple stimuli, such as a single orientation grating or multiple stimuli that allowed for competition between distinct neuronal populations. For
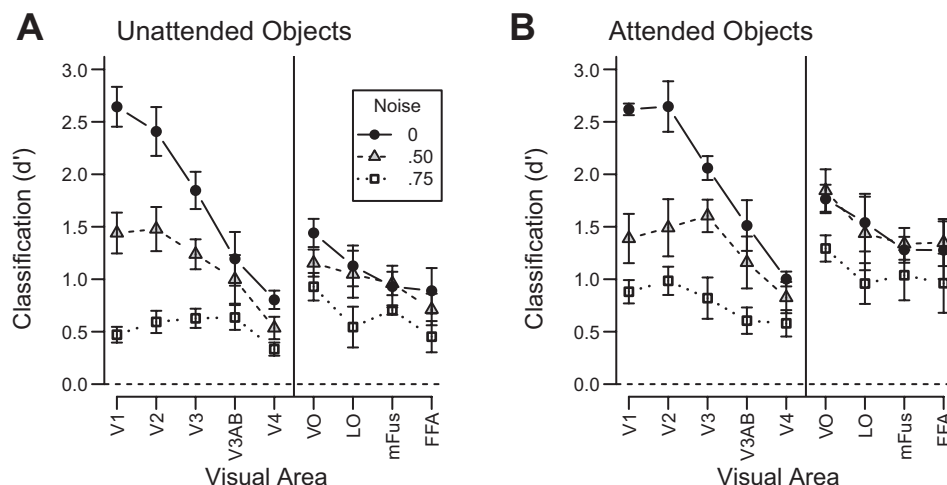

Fig. 7. A pattern classifier was trained on the zero-noise unattended condition and used to classify object category in all other conditions. The results are similar to the above analyses, in which classifiers were separately trained and tested within each condition, suggesting that the nature of object-category information is qualitatively similar across attention and noise conditions, but that the strengths of these patterns are modified by attention and noise.

example, Kamitani and Tong (2005) found that selective attention directed to one of two overlapping orientated gratings could modify orientation-selective activity patterns in early visual areas. Although no visual noise was present in the compound stimulus, this stimulus design allowed for competitive interactions to take place between the attended grating and the unattended grating. Similar competitive interactions are likely to take place when an observer views objects in noise, between neuronal populations responding to the contours of the object and those responding to the noise elements. In another fMRI study, Jehee et al. (2011) found that attending to the orientation of a peripheral (noise-free) grating led to enhanced orientation-selective responses in early visual areas, but that attending to the contrast of that grating did not. The modulatory effects found in their study were attributed to feature-based attention directed at a specific orientation. In the present study, participants had to attend to objects composed of many different local features, and the top-down mechanisms involved in this more complex form of attentional selection may well differ from basic mechanisms of feature-based attention. It will be interesting for future studies to determine the extent to which mechanisms of amplification and noise reduction can modulate activity in early visual areas, and how such factors may depend on the complexity of the stimulus being attended.

The present study provides compelling evidence for attention-based noise filtering in early visual areas. These visual areas might also have passive noise-filtering properties, but the attention-dependent effects found here indicate that a target object is better distinguished from surrounding noise when it becomes the focus of attention. What types of top-down feedback to early areas would allow for noise filtering (i.e., improved SNR), such that the responses to local object features are enhanced or the responses to the noise elements are dampened? It is possible that simple forms of attentional feedback could have contributed to the noise-filtering effects found in early visual areas. For example, given that our line-drawn stimuli had somewhat greater power at lower spatial frequencies, a feature-based attentional mechanism (Martinez-Trujillo and Treue 2004; Maunsell and Treue 2006; Treue and Martinez Trujillo 1999) that enhanced responses to lower spatial frequencies could improve the SNR to some extent. Top-down feedback might also involve more complex filtering operations, such as facilitating the processing of fragmented contours by modulating the strength of lateral connections between neurons with collinear receptive fields. Attentional modulation of collinear interactions has been reported in single-unit recording studies of V1, and both perceptual learning and task-related voluntary attention appear to be important for realizing these effects (Crist et al. 2001; Li et al. 2008; Roelfsema et al. 1998).

However, recent fMRI studies have also found that object-related knowledge can lead to top-down effects in early visual areas. For example, following successful recognition of an ambiguous image of an object, activity patterns in early visual areas are more stable and reliable (Gorlin et al. 2012; Hsieh et al. 2010). In addition, attending to the shapes of peripheral objects can lead to object-specific patterns of activity in the foveal representation of early visual areas, presumably due to top-down feedback of object-related information (Williams et al. 2008). In separate ongoing studies, our lab has found that covertly attending to one of two overlapping objects leads to

systematic biases in the patterns of activity in early visual areas, which involve shifts toward the pattern that would be evoked if the attended stimulus were presented in isolation (see also Hou and Liu 2012). The above studies suggest that high-level representations of shape and object properties could potentially inform top-down feedback signals to early visual areas to facilitate the attentional filtering of objects in noise. Following the initial sweep of feedforward responses to an object in visual noise, the activation of shape-selective neurons in higher areas might lead to top-down feedback to early visual areas, to enhance local responses to relevant object contours and to suppress responses to unaccounted for noise. Such top-down attentional filtering would lead to a refinement of the visual signals that reach high-level areas. Such dynamic interactions between high- and low-level visual areas would be consistent with the proposed role of recurrent interactions in visual perception and attentional selection (Hochstein and Ahissar 2002; Lamme and Roelfsema 2000; Tong 2003), and with the notion that feedback can act as a predictive coding mechanism in early visual areas (Rao and Ballard 1999; Wacongne et al. 2011). It will be interesting if future studies investigate whether attention-based noise filtering relies on top-down feedback that modulates visual responses based on more simple properties, such as spatial frequency or collinearity, or more complex properties.

How might the attentional mechanisms of noise filtering and amplification operate at the neuronal level? We would anticipate that these attentional mechanisms lead to changes in the gain of neuronal responses, as has been demonstrated in many neurophysiological studies of visual attention (McAdams and Maunsell 1999; Reynolds et al. 2000; Treue and Martinez Trujillo 1999). For example, selective top-down enhancement of the fragmented contours corresponding to the target object might serve to enhance the gain of responses to the relevant signal. Such enhancement of relevant contours could also lead to local suppression of responses to noise, as would be predicted by divisive normalization models with local inhibition (Heeger 1992; Reynolds and Heeger 2009). In addition, attention has recently been found to modulate the variability of neuronal firing (Mitchell et al. 2009) as well as the strength of correlated noise responses between neurons (Cohen and Maunsell 2009), and such factors could potentially contribute to the effects of attention found here. Because of the coarse nature of the BOLD response, fMRI studies cannot readily distinguish between these possible types of neuronal modulation, as either an increase in response gain or a reduction in response variability at the neuronal level would lead to similar improvements in the reliability of large-scale activity patterns. However, the present fMRI study does allow us to distinguish between amplification and noise filtering: the strength or reliability of responses is similarly boosted for signal and noise in the former case, whereas in the latter case the signal is enhanced relative to the noise.

Our results further demonstrate how fMRI pattern classification can provide a different window into functional or visual processes than that provided by simple measures of mean BOLD activity. For example, the addition of visual noise led to stronger mean BOLD responses in early visual areas, whereas the presence of such noise severely impaired the reliability of the activity patterns evoked by different objects. The discrepancy between these two measures highlights the fact that

classification accuracy reflected the distinctiveness of the visual signal (i.e., object category) in the presence of uninformative, disruptive noise, whereas mean BOLD activity reflected the overall strength of visual responses to both signal and noise. These results indicate potential limitations of previous reports suggesting that the overall BOLD activity in early visual areas is not disrupted by the addition of visual noise or image scrambling (Grill-Spector et al. 1998; Tjan et al. 2006).

We also observed differences in how attending to the objects affected the mean BOLD signal and classification performance. In particular, we found no attentional effects in classification accuracy for noiseless or medium-noise images in V1 or V2, despite the fact that attending to the objects led to increases in mean BOLD responses. These findings parallel other recent work showing dissociable effects of spatial attention on mean BOLD responses and the reliability of orientation-selective activity patterns (Jehee et al. 2011). It is therefore important to consider how our pattern classification results compare with previous studies that have investigated the effects of attention and noise on mean BOLD activity. For example, a recent fMRI study found that attending to a low-contrast grating presented in high-contrast noise reduced mean BOLD responses in V1 (Lu et al. 2011). This result was interpreted to suggest a noise-filtering mechanism, in which attending to the grating suppressed responses to the irrelevant external noise, thereby leading to an overall decrease in mean BOLD response. This finding concurs with our pattern-classification evidence for noise filtering in V1. However, we also found effects of noise filtering in higher visual areas beyond V1, whereas Lu et al. failed to observe an attentional decrease in mean BOLD response in these areas. It would be interesting for future studies to investigate whether noise filtering and amplification are similarly engaged by simple stimuli such as gratings and complex object stimuli such as we used here, and how these mechanisms can be measured using pattern classification and mean BOLD activity.

In conclusion, we found that attending to objects embedded in noise led to dramatic increases in pattern classification performance, implying that noise filtering is an important function of covert visual attention. The differential effects of attention across the visual hierarchy suggest that the visual system is designed to operate optimally under both low- and high-noise situations. Early visual areas exhibited strong effects of noise filtering but negligible evidence of amplification, whereas high-level visual areas showed effects of both mechanisms, due to the emergence of amplification at higher stages of processing. These results reveal how top-down attention can serve multiple functions, and how its effects on visual processing are not unitary. To date, most neuroimaging and neurophysiological studies of attention have emphasized the specific role of attention in modulating the "gain" or response level of sensory neurons (Buracas and Boynton 2007; Reynolds and Chelazzi 2004). However, such an attentional amplification mechanism would not be sufficient for improving sensory encoding in the presence of external noise, particularly if both signal and noise are spatially overlapping. Instead, multiple mechanisms are needed to explain how attention is able to dynamically enhance visual processing across a variety of viewing conditions, especially those that involve substantial visual clutter or noise.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the author(s).

AUTHOR CONTRIBUTIONS

M.S.P., J.D.S., and F.T. conception and design of research; M.S.P. performed experiments; M.S.P. analyzed data; M.S.P., S.L., J.D.S., and F.T. interpreted results of experiments; M.S.P. prepared figures; M.S.P., S.L., and F.T. drafted manuscript; M.S.P., S.L., J.D.S., and F.T. edited and revised manuscript; M.S.P., S.L., and F.T. approved final version of manuscript.

REFERENCES

Beauchamp MS, Cox RW, DeYoe EA. Graded effects of spatial and featural attention on human area MT and associated motion processing areas. *J Neurophysiol* 78: 516–520, 1997.

Boser BE, Guyon IM, Vapnik VN. A training algorithm for optimal margin classifiers. In: *Proceedings of the Fifth Annual Workshop on Computational Learning Theory,* Pittsburgh, PA, 1992, p. 144–152.

Brainard DH. The psychophysics toolbox. *Spatial Vision* 10: 433–436, 1997.

Buracas GT, Boynton GM. The effect of spatial attention on contrast response functions in human visual cortex. *J Neurosci* 27: 93–97, 2007.

Carrasco M. Visual attention: the past 25 years. *Vision Res* 51: 1484–1525, 2011.

Carrasco M, Penpeci-Talgar C, Eckstein M. Spatial covert attention increases contrast sensitivity across the CSF: support for signal enhancement. *Vision Res* 40: 1203–1215, 2000.

Chawla D, Rees G, Friston KJ. The physiological basis of attentional modulation in extrastriate visual areas. *Nat Neurosci* 2: 671–676, 1999.

Cohen MR, Maunsell JH. Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* 12: 1594–1600, 2009.

Crist RE, Li W, Gilbert CD. Learning to see: experience and attention in primary visual cortex. *Nat Neurosci* 4: 519–525, 2001.

Cukur T, Nishimoto S, Huth AG, Gallant JL. Attention during natural vision warps semantic representation across the human brain. *Nat Neurosci* 16: 763–770, 2013.

Desimone R, Duncan J. Neural mechanisms of selective visual attention. *Ann Rev Neurosci* 18: 193–222, 1995.

Dosher BA, Lu ZL. Mechanisms of perceptual attention in precuing of location. *Vision Res* 40: 1269–1292, 2000a.

Dosher BA, Lu ZL. Noise exclusion in spatial attention. *Psychol Sci* 11: 139–146, 2000b.

Engel SA, Glover GH, Wandell BA. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex* 7: 181–192, 1997.

Gandhi SP, Heeger DJ, Boynton GM. Spatial attention affects brain activity in human primary visual cortex. *Proc Natl Acad Sci* 96: 3314–3319, 1999.

Gorlin S, Meng M, Sharma J, Sugihara H, Sur M, Sinha P. Imaging prior information in the brain. *Proc Natl Acad Sci* 109: 7935–7940, 2012.

Greve DN, Fischl B. Accurate and robust brain image alignment using boundary-based registration. *Neuroimage* 48: 63–72, 2009.

Grill-Spector K. The neural basis of object perception. *Curr Opin Neurobiol* 13: 159–166, 2003.

Grill-Spector K, Kushnir T, Hendler T, Edelman S, Itzchak Y, Malach R. A sequence of object-processing stages revealed by fMRI in the human occipital lobe. *Hum Brain Mapp* 6: 316–328, 1998.

Haynes JD, Rees G. Decoding mental states from brain activity in humans. *Nat Rev Neurosci* 7: 523–534, 2006.

Heeger DJ. Normalization of cell responses in cat striate cortex. *Vis Neurosci* 9: 181–197, 1992.

Hochstein S, Ahissar M. View from the top: hierarchies and reverse hierarchies in the visual system. *Neuron* 36: 791–804, 2002.

Hou Y, Liu T. Neural correlates of object-based attentional selection in human cortex. *Neuropsychologia* 50: 2916–2925, 2012.

Hsieh PJ, Vul E, Kanwisher N. Recognition alters the spatial pattern of FMRI activation in early retinotopic cortex. *J Neurophysiol* 103: 1501–1507, 2010.

Ingleby JD. The separation of bias and sensitivity in multiple-choice tasks. *Perception* 2: 295–305, 1973.

Jeffreys H. *The Theory of Probability*. Oxford: Oxford University Press, 1961.

Jehee JF, Brady DK, Tong F. Attention improves encoding of task-relevant features in the human visual cortex. *J Neurosci* 31: 8210–8219, 2011.

Jenkinson M, Bannister P, Brady M, Smith S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* 17: 825–841, 2002.

Kamitani Y, Tong F. Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8: 679–685, 2005.

Kanwisher N, McDermott J, Chun MM. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J Neurosci* 17: 4302–4311, 1997.

Kietzmann TC, Swisher JD, Konig P, Tong F. Prevalence of selectivity for mirror-symmetric views of faces in the ventral and dorsal visual pathways. *J Neurosci* 32: 11763–11772, 2012.

Kriegeskorte N, Mur M, Ruff DA, Kiani R, Bodurka J, Esteky H, Tanaka K, Bandettini PA. Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* 60: 1126–1141, 2008.

Lamme VA, Roelfsema PR. The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci* 23: 571–579, 2000.

Li W, Piech V, Gilbert CD. Learning to link visual contours. *Neuron* 57: 442–451, 2008.

Ling S, Liu T, Carrasco M. How spatial and feature-based attention affect the gain and tuning of population responses. *Vision Res* 51: 1484–1525, 2009.

Lu ZL, Dosher BA. External noise distinguishes attention mechanisms. *Vision Res* 38: 1183–1198, 1998.

Lu ZL, Dosher BA. Spatial attention: different mechanisms for central and peripheral temporal precues? *J Exp Psychol Hum Percept Perform* 26: 1534–1548, 2000.

Lu ZL, Lesmes LA, Dosher BA. Spatial attention excludes external noise at the target location. *J Vis* 2: 312–323, 2002.

Lu ZL, Li X, Tjan BS, Dosher BA, Chu W. Attention extracts signal in external noise: a BOLD fMRI study. *J Cogn Neurosci* 23: 1148–1159, 2011.

Lu ZL, Liu CQ, Dosher BA. Attention mechanisms for multi-location first- and second-order motion perception. *Vision Res* 40: 173–186, 2000.

Martinez-Trujillo JC, Treue S. Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr Biol* 14: 744–751, 2004.

Maunsell JH, Treue S. Feature-based attention in visual cortex. *Trends Neurosci* 29: 317–322, 2006.

McAdams CJ, Maunsell JH. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J Neurosci* 19: 431–441, 1999.

Mitchell JF, Sundberg KA, Reynolds JH. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63: 879–888, 2009.

Morrone MC, Denti V, Spinelli D. Color and luminance contrasts attract independent attention. *Curr Biol* 12: 1134–1137, 2002.

Naselaris T, Prenger RJ, Kay KN, Oliver M, Gallant JL. Bayesian reconstruction of natural images from human brain activity. *Neuron* 63: 902–915, 2009.

Norman KA, Polyn SM, Detre GJ, Haxby JV. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10: 424–430, 2006.

Peelen MV, Fei-Fei L, Kastner S. Neural mechanisms of rapid natural scene categorization in human visual cortex. *Nature* 460: 94–97, 2009.

Pelli DG. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision* 10: 437–442, 1997.

Rao RP, Ballard DH. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci* 2: 79–87, 1999.

Reddy L, Kanwisher NG, VanRullen R. Attention and biased competition in multi-voxel object representations. *Proc Natl Acad Sci* 106: 21447–21452, 2009.

Reynolds JH, Chelazzi L. Attentional modulation of visual processing. *Ann Rev Neurosci* 27: 611–647, 2004.

Reynolds JH, Heeger DJ. The normalization model of attention. *Neuron* 61: 168–185, 2009.

Reynolds JH, Pasternak T, Desimone R. Attention increases sensitivity of V4 neurons. *Neuron* 26: 703–714, 2000.

Roelfsema PR, Lamme VA, Spekreijse H. Object-based attention in the primary visual cortex of the macaque monkey. *Nature* 395: 376–381, 1998.

Rouder JN, Speckman PL, Sun D, Morey RD, Iverson G. Bayesian t tests for accepting and rejecting the null hypothesis. *Psychon Bull Rev* 16: 225–237, 2009.

Scolari M, Byers A, Serences JT. Optimal deployment of attentional gain during fine discriminations. *J Neurosci* 32: 7723–7733, 2012.

Serences JT, Boynton GM. Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron* 55: 301–312, 2007.

Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268: 889–893, 1995.

Smith PL, Wolfgang BJ, Sinclair AJ. Mask-dependent attentional cuing effects in visual signal detection: the psychometric function for contrast. *Percept Psychophys* 66: 1056–1075, 2004.

Somers DC, Dale AM, Seiffert AE, Tootell RB. Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proc Natl Acad Sci* 96: 1663–1668, 1999.

Tjan BS, Lestou V, Kourtzi Z. Uncertainty and invariance in the human visual cortex. *J Neurophysiol* 96: 1556–1568, 2006.

Tong F. Primary visual cortex and visual awareness. *Nat Rev Neurosci* 4: 219–229, 2003.

Tong F, Pratte MS. Decoding patterns of human brain activity. *Ann Rev Psychol* 63: 483–509, 2012.

Treue S, Martinez Trujillo JC. Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399: 575–579, 1999.

Vapnik VN. *Statistical Learning Theory*. New York: Wiley, 1998.

Wacongne C, Labyt E, van Wassenhove V, Bekinschtein T, Naccache L, Dehaene S. Evidence for a hierarchy of predictions and prediction errors in human cortex. *Proc Natl Acad Sci* 108: 20754–20759, 2011.

Williams MA, Baker CI, Op de Beeck HP, Shim WM, Dang S, Triantafyllou C, Kanwisher N. Feedback of visual object information to foveal retinotopic cortex. *Nat Neurosci* 11: 1439–1445, 2008.