

Adaptive Allocation of Attentional Gain

Miranda Scolari and John T. Serences

Perception and Cognition Laboratory, Department of Psychology, University of California, San Diego, La Jolla, California 92093-0109

Humans are adept at distinguishing between stimuli that are very similar, an ability that is particularly crucial when the outcome is of serious consequence (e.g., for a surgeon or air-traffic controller). Traditionally, selective attention was thought to facilitate perception by increasing the gain of sensory neurons tuned to the defining features of a behaviorally relevant object (e.g., color, orientation, etc.). In contrast, recent mathematical models counterintuitively suggest that, in many cases, attentional gain should be applied to neurons that are tuned away from relevant features, especially when discriminating highly similar stimuli. Here we used psychophysical methods to critically evaluate these “ideal observer” models. The data demonstrate that attention enhances the gain of the most informative sensory neurons, even when these neurons are tuned away from the behaviorally relevant target feature. Moreover, the degree to which an individual adopted optimal attentional gain settings by the end of testing predicted success rates on a difficult visual discrimination task, as well as the amount of task improvement that occurred across repeated testing sessions (learning). Contrary to most traditional accounts, these observations suggest that the primary function of attentional gain is not to enhance the representation of target features per se, but instead to optimize performance on the current perceptual task. Additionally, individual differences in gain suggest that the operating characteristics of low-level attentional phenomena are not stable trait-like attributes and that variability in how attention is deployed may play an important role in determining perceptual abilities.

Introduction

Perceptual expertise requires selectively attending to relevant aspects of a scene (Fahle, 2004). For example, when examining an x-ray, a radiologist must attend to subtle variations in shape and color to determine whether or not an anomaly is malignant. Typically, voluntary shifts of attention are thought to enhance perceptual acuity by selectively modulating the gain—and thus the signal-to-noise ratio—of those sensory neurons that respond maximally to the defining features of the relevant target (i.e., those neurons that are tuned to the target feature) (Desimone and Duncan, 1995; McAdams and Maunsell, 1999; Martinez-Trujillo and Treue, 2004; Boynton, 2005; Maunsell and Treue, 2006). Applying attentional gain to these sensory neurons is optimal when detecting the presence (vs absence) of a stimulus or discriminating one stimulus from a very dissimilar set of distractors (e.g., when searching for a high-contrast tumor in an otherwise unremarkable image) (Fig. 1A) (McAdams and Maunsell, 1999; Hol and Treue, 2001; Martinez-Trujillo and Treue, 2004). In contrast, recent theoretical and psychophysical work indicates that enhancing the response of neurons tuned to the target feature is suboptimal when performing a difficult discrimination between two very similar stimuli (e.g., subtle variations in the color of two adjacent tissue types) (Regan and Beverley, 1985;

Purushothaman and Bradley, 2005; Butts and Goldman, 2006; Jazayeri and Movshon, 2006, 2007; Navalpakkam and Itti, 2007). In this case, gain should be applied to neurons tuned slightly away from the target because they are more sensitive to small changes in the relevant feature value (Fig. 1B) (Navalpakkam and Itti, 2007). However, the shape of this “off-channel” attentional gain profile, and the extent to which human observers are capable of adaptively engaging such a computationally optimal strategy, remains largely unexplored.

Here, two simple models were used to predict how attentional gain should be deployed while performing a difficult discrimination between stimuli that are very similar [the “optimal gain hypothesis” by Navalpakkam and Itti and a model based on Fisher information (FI)]. Next, psychophysical studies examined how closely human observers emulate the ideal attentional gain functions predicted by each model, as well as the perceptual consequences of how people differentially deploy attentional gain. On average, across all subjects, attentional gain was applied to the sensory neurons that best discriminated targets from distractors, even in situations in which these neurons were not tuned to the target-defining feature. Moreover, after several testing sessions, the amount of attentional gain deployed to informative sensory neurons came to predict both individual differences in visual search efficiency as well as the amount of improvement that occurred across all experimental sessions (i.e., learning). Most traditional theories posit that attention facilitates perception by increasing the firing rate of neurons that are tuned to relevant features, which is assumed to enhance the cortical representation of the target (for review, see Boynton, 2005). However, the present results suggest that attention adaptively modulates the gain of the most informative sensory neurons to maximize the probability of successfully performing a specific perceptual task.

Received Nov. 25, 2008; revised Aug. 12, 2009; accepted Aug. 19, 2009.

This work was supported by National Institutes of Health Grant R21-MH083902 (J.T.S.). We thank Harold Pashler, Steven Yantis, and Edward Awh for helpful comments and Nicole Panzer and Lily Wu for assistance with data collection.

Correspondence should be addressed to either Miranda Scolari or John Serences, Department of Psychology, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0109, E-mail: mscolari@ucsd.edu or jserences@ucsd.edu.

DOI:10.1523/JNEUROSCI.5642-08.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/2911933-10\$15.00/0




Figure 1. Understanding how attentional gain should be applied to sensory neurons in the context of orientation detection and discrimination tasks. Although various features such as shape and color are discussed in Introduction, we focus here on orientation for simplicity and because it is used in the subsequent psychophysical studies. **A**, Modulating the gain of the most responsive neurons is optimal during target detection (or coarse discrimination) because they respond maximally to the relevant feature and minimally to the irrelevant feature(s), thus resulting in a high ratio of spiking evoked by the targets compared with distractors (a high SNR). **B**, When performing a difficult fine discrimination, a neuron tuned to the target feature (solid black line) does not discriminate targets and distractors very well (SNR 1). However, a neuron tuned to an exaggerated target feature (dashed line) undergoes a large change in firing rate because its tuning function has a steeper slope at the to-be-discriminated orientation (that is, SNR2 > SNR1). Vertical dashed lines indicate the target (90°) and the distractor(s) (85°). **B** was adapted with permission from Navalpakkam and Itti (2007), their Figure 4.

Materials and Methods

Subjects

Thirty-one subjects from the University of California, Irvine (UCI) community were recruited to participate in experiments 1 and 2, each of whom gave written informed consent according to Institutional Review Board requirements at UCI. Experiment 1 was actually run after experiment 2, and the presentation order was reversed here for clarity. Fourteen of the subjects who participated in experiment 1 also participated in experiment 2, along with an additional 17 subjects (M.S. was a subject in both experiments). Most subjects completed eight blocks (108 trials per block) during the single experimental session in experiment 1 (M.S. completed five blocks). In experiment 2, subjects completed five 1.5 h sessions, each consisting of eight blocks (108 trials per block), and each was held on separate days. Data from three observers were discarded from experiment 2 because of difficulty setting contrast detection thresholds (i.e., floor effects). Subjects were compensated \$10/h for their participation.

Seventeen subjects from the University of California, San Diego (UCSD) community completed experiment 3 (M.S. and one other subject also participated in experiments 1 and 2), each of whom gave written informed consent according to Institutional Review Board requirements at UCSD. One subject was discarded from analysis because of difficulty setting contrast detection thresholds (floor effects). Subjects completed eight blocks (108 trials per block) during each of three 1.5-h sessions, each held on different days. Subjects were compensated \$8/h for their participation.

Materials

All stimuli were generated using the Matlab programming language (version 7.6; MathWorks) with the Psychophysics Toolbox (version 3) (Brainard, 1997; Pelli, 1997). For experiments 1 and 2, stimuli were displayed on a 17 inch cathode ray tube (CRT) monitor running at 100 Hz. For experiment 3, stimuli were displayed on either a 19 inch CRT monitor running at 100 Hz or a 17 inch CRT monitor running at 85 Hz (any given subject used the same monitor for all sessions). The luminance output of the monitor was measured using either a Minolta LS110 (experiments 1 and 2) or a United Detector Technology S380 (experiment 3) photometer and linearized in the stimulus presentation software.

General description of experimental approach

In each of the three experiments, two basic trial types were intermixed: one dominant (two-thirds of all trials) and one nondominant (one-third of all trials). The idea behind this general design scheme was to use the dominant trial type to induce an attentional set and then to probe the consequences of this attentional set using the nondominant trial type.

Because all experiments shared this common structure, we first describe the dominant trial types for each of the three experiments and then describe the nondominant trial types.

Experiments 1–3: dominant “attentional set” trials

Experiments 1 and 2: fine discrimination task. Fine discrimination trials made up two-thirds of the total trials in both experiments 1 and 2 (see below for details on the remaining one-third of the trials in each experiment). Targets and distractors were Gabor patches (Gaussian windowed sinusoidal gratings) with a radius of 3° visual angle and spatial frequency of 2 cycles/°. The computer screen was divided into four quadrants, and each quadrant contained a single Gabor. The center of each Gabor was vertically and horizontally offset from the fixation point by 3° visual angle. Three of the Gabors were rendered in the same orientation (distractors), and the fourth differed from the others by 5° (target). There were 36 equally represented potential target and distractor orientations (5° steps over 180°) to equate sensory stimulation at each possible stimulus orientation; this was done to ensure that subsequent estimates of contrast detection thresholds were not biased by passive sensory adaptation (see section on nondominant trial types below). A central precue was presented at the start of each trial for 1.25 s, which indicated both the orientation of the distractors (via the orientation of the cue line) and the rotational offset of the target from the distractors (via the color of the cue) (Fig. 2). For example, a red cue indicated that the target would be rotated 5° clockwise from the cue line, and a green cue indicated that the target would be rotated 5° counterclockwise from the cue line (color assignments were counterbalanced across subjects). The search array was presented for 500, 1000, 1500, or 2000 ms and was immediately followed by four pattern masks comprising truncated Gaussian noise presented for 250 ms (mean luminance of each mask was middle gray, maximum was white, and minimum was black). Subjects were instructed to emphasize accuracy and indicated the quadrant of the target with an unsped button-press response (numbers 1, 2, 4, and 5 on the number pad of a standard QWERTY keyboard). The target appeared in each quadrant an equal number of times over the course of an experimental session. To help maintain motivation and to encourage preparation for the fine discrimination task, correct responses were rewarded with 10 points, and incorrect responses were penalized with −2.5 points, although the points had no monetary value. Feedback was presented at the end of each trial for 500 ms and indicated whether the response was correct or incorrect, how many points the subject had earned for the previous trial, and how many points the subject had earned in total.

Experiment 3: coarse discrimination task. Coarse discrimination trials made up two-thirds of the total trials. The details of the coarse discrimination trials were nearly identical to those of the fine discrimination trials used in experiments 1 and 2, with the following exceptions. First, the target differed from the distractors by 90°. Because a rotational offset of 90° clockwise or counterclockwise from the cue would result in the same orientation, the direction of the rotational offset of the target from the distractors was no longer relevant, and the central cue was therefore always rendered in red (instead of in red or green to indicate a clockwise or counterclockwise rotational offset, as in experiments 1 and 2). Second, to equate the difficulty of the coarse discrimination task in experiment 3 with the difficulty of the fine discrimination task used in experiments 1 and 2, we presented the search array in experiment 3 for a very brief temporal interval; the exposure duration was fixed for each subject and ranged from 30 to 71 ms (presentation rate was adjusted on a subject-by-subject basis in units of either 10 or 11.7 ms to guard against ceiling and floor effects, and the mean \pm SEM exposure across subjects was 58 ± 2.93 ms). Note also that this resulted in one exposure duration per session used for each subject in experiment 3, as opposed to the four exposure durations (500–2000 ms) used in experiments 1 and 2. Finally, the Gabor patches were rendered at 40% contrast, as opposed to 100% contrast as in experiments 1 and 2, to further increase the difficulty of the coarse discrimination. Feedback was presented at the end of each trial just as in experiments 1 and 2.

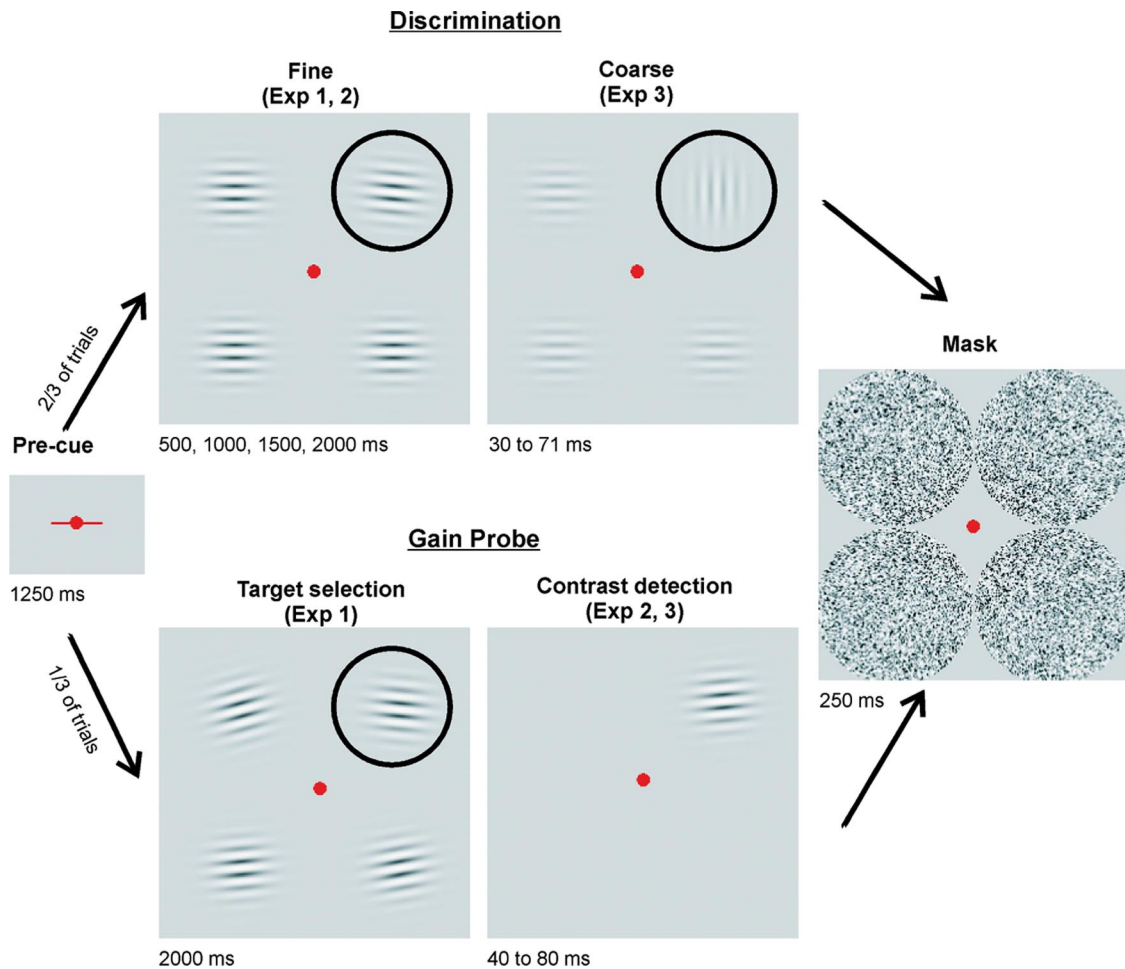


Figure 2. Schematic of the experimental paradigms (for details, see Materials and Methods). The black circle surrounding some of the Gabors indicates that they are the target of search and was not presented in the actual study.

Experiments 1–3: nondominant “attentional probe” trials (one-third of all trials)

Experiment 1: target selection task. This task was modeled after the psychophysical procedure developed by Navalpakkam and Itti (2007). The same precue used on the dominant fine discrimination trials (see above) was presented at the start of each target selection trial, so subjects did not know which task they would perform until the stimulus display appeared. After the cue, one Gabor patch was presented in each quadrant for 2000 ms, and each Gabor was rendered at a different orientation. The target was rendered at the orientation indicated by the colored cue (i.e., the target was the same orientation as would be expected on a fine discrimination trial, either 5° clockwise or counterclockwise from the cue line depending on cue color, and the target orientation is henceforth referred to as 0°). The three distractors were all rendered in different orientations, rotated by $\pm 5^\circ$, $\pm 10^\circ$, $\pm 15^\circ$, or $\pm 20^\circ$ from the target, in which positive values refer to rotation in the direction indicated by the color of the cue, and negative values refer to rotation in the direction opposite of that indicated by the cue. For example, if a red cue indicated that the target was rotated clockwise with respect to the cue line, then by convention, all distractors rotated clockwise from the target would be denoted with a positive value and all distractors rotated counterclockwise would be denoted with a negative value. Thus, positive rotational offsets denote “exaggerated targets,” and negative offsets denote the distractor orientation (-5° from the target) and “exaggerated distractors.” A post-stimulus prompt was used to inform the subjects that they were to perform either the dominant fine discrimination task (described above) or the target selection task, and then they indicated the location of the target with an unspeeded button-press response (although the postcue was not strictly necessary given that the subjects’ task was always to find the

target). Subjects were not given feedback about the accuracy of their response on target selection trials so that this information could not be used to adjust subsequent responses. Although this task was conceptually similar to the paradigm of Navalpakkam and Itti, there were some potentially important differences: we used a smaller search array (4 vs 25 stimuli) and varied the cued target orientation from trial to trial, whereas they used a single fixed target and distractor orientation. For an expanded explanation of these differences, see Discussion.

Experiments 2 and 3: contrast detection task. The same precue used on the dominant fine (experiment 2) or coarse (experiment 3) discrimination trials was presented at the start of each trial, but only a single Gabor patch was flashed briefly in a randomly selected quadrant (for experiment 2, the exposure duration of the single Gabor was either 50 or 70 ms, set on a subject-by-subject basis to ensure that each participant could see the stimulus when rendered at full contrast; for experiment 3, the exposure duration of the single Gabor ranged between 40 and 80 ms). The Gabor stimulus was immediately followed by four pattern masks, one presented in each quadrant (same type of Gaussian noise mask described above). The subjects’ task was simply to make a button-press response indicating which of the four quadrants contained the single oriented stimulus. The orientation of the Gabor was selected from a set of nine possible rotational offsets with respect to the expected target orientation: 0° , $\pm 5^\circ$, $\pm 10^\circ$, $\pm 20^\circ$, or $\pm 40^\circ$ (in experiment 2) or 0° , $\pm 5^\circ$, $\pm 10^\circ$, $\pm 20^\circ$, or $\pm 90^\circ$ (in experiment 3). In experiment 2, the sign of the rotational offset of the Gabor depended on the central cue, with positive values indicating a rotation in the cued direction. For example, if the cue indicated that the target would be rotated 5° clockwise from the cue line, then, by convention, positive values were assigned to stimuli rotated clockwise and negative values were assigned to stimuli rotated counter-

clockwise with respect to the expected target orientation. In experiment 3, positive rotational offsets were always rotated counterclockwise from the expected target orientation, and negative offsets were always clockwise; the sign of the offset was absolute because the expected target orientation and the cue line were orthogonal (and hence a target rotated +90° from the cue was identical to a target rotated −90°).

The dependent measure of interest in both experiments 2 and 3 was the contrast level required at each orientation offset to achieve 75% correct performance at reporting the quadrant that contained the target (determined using the QUEST algorithm) (Watson and Pelli, 1983). A separate staircase was run for each orientation offset (in experiment 3, contrast levels were independently adjusted for distractors offset by ±90° from the target to equate the number of trials in each staircase, although these two offsets in fact referred to physically identical stimuli). Critically, these contrast detection probe trials were intended to assess how subjects were applying gain to different populations of orientation-selective neurons in preparation for the dominant task, thus producing an estimated “attentional gain function” for each subject. Because sensitivity in a detection task depends primarily on the gain level of neurons tuned to the to-be-detected feature (Regan and Beverley, 1985), we used the estimated contrast detection thresholds as a proxy for the level of neural gain applied to different subsets of orientation-selective neurons in early visual cortex while subjects prepared for the dominant task. Thus, an observation of lower contrast detection thresholds for a particular orientation implies that neurons tuned to that orientation underwent more attentional gain. Because each of the 36 possible stimulus orientations was cued with equal probability (see above), neurons tuned to each orientation received an equal amount of sensory stimulation over the course of the experiment. This ensured that estimates of contrast detection thresholds were safeguarded against biases that might have been introduced by differential sensory adaptation effects. Subjects received 2.5 points for correct responses and were not penalized for incorrect responses (feedback was provided at the end of each trial as described above).

The interpretation of the contrast detection thresholds from experiments 2 and 3 depends on the assumption that subjects were preparing for dominant fine or coarse discrimination trials, as opposed to preparing for the contrast detection task. To encourage the desired attentional set, contrast detection trials were presented infrequently (one-third of the total trials) and were deemphasized in both the task instructions and point system, so there was little motivation for the subjects to place priority on this secondary task. More importantly, the orientation of the contrast detection probe was randomly selected from a range of possible offsets with respect to the cued target orientation (±40° in experiment 2 and ±90° in experiment 3). Therefore, subjects could not have effectively increased the gain of neurons tuned to the contrast detection probe because its orientation was unknown in advance.

Predicting attentional gain: optimal gain hypothesis

Navalpakkam and Itti (2007) recently proposed an elegant model, which we term here the optimal gain hypothesis, to describe how attention should be deployed when performing a visual search task that requires discriminating a target from a uniform field of distractors (for more details and for a description of how their model predicts attentional deployment under a wide array of other search conditions as well, see their Appendix A):

$$g_i = \frac{f_i(t)}{f_i(d)} \left/ \left[\frac{1}{N} * \sum_{i=1}^N \frac{f_i(t)}{f_i(d)} \right] \right. \quad (1)$$

When both target and distractor features are known in advance, the optimal attentional gain (g_i) that should be applied to neuron i can be estimated using Equation 1, where $f_i(t)$ is the response of a neuron to the target, $f_i(d)$ is the response of the same neuron to the distractor(s), and N is the total number of neurons in the population (equation adapted from Navalpakkam and Itti, 2007) (see also Regan and Beverley, 1985; Butts and Goldman, 2006; Jazayeri and Movshon, 2006). The first term in Equation 1 captures the ratio of the response evoked by the target to the response evoked by the distractor(s); the second term is a normalizing

factor that reflects the average response ratio across all neurons and can be effectively ignored for the present purposes. According to Equation 1, attentional gain should primarily be applied to the neurons that undergo the largest positive firing rate change in response to targets compared with distractors. When performing a coarse discrimination task—say when searching for a horizontal target among vertical distractors—the optimal gain model predicts that gain should be applied to the neurons that are tuned to the target feature, an observation that has been demonstrated empirically by single-unit recording studies (Martinez-Trujillo and Treue, 2004) (for review, see Maunsell and Treue, 2006). Conversely, in a difficult fine discrimination task, gain should be applied to those neurons tuned to an orientation rotated slightly beyond the target (termed an exaggerated target feature) (see Fig. 3A). For example, when searching for a 90° target among 85° distractors, positive attentional gain should be applied to neurons tuned to ~95–120°, assuming an average tuning bandwidth of ~40–50°, which is an appropriate estimate for orientation-selective neurons in primary visual cortex (V1) (Snowden, 1992; Ringach et al., 2002). In contrast, no gain should be applied to neurons that respond nearly equally well to targets and distractors, and neurons that respond more to the distractor(s) than to the target should be suppressed (i.e., neurons tuned to ~50–85°, henceforth termed exaggerated distractors).

Predicting attentional gain: Fisher information

Fisher information is a related but alternative metric that can also be used to predict how attention might influence the gain of sensory neurons to facilitate visual search. However, instead of directly generating an estimate of attentional gain, FI measures how well each neuron distinguishes between the target and distractor stimuli (Seung and Sompolinsky, 1993; Pouget et al., 2001). Formally, FI for a neuron is defined as the derivative of the firing rate with respect to the relevant stimulus parameters (orientation, contrast, spatial frequency, etc.), weighted by the amount of noise in the neural response (for a full mathematical treatment of FI for various target/distractor configurations and for a detection task, see Itti et al., 2000, particularly their Appendix A and B). For the simple case of fine discrimination between two adjacent orientations, FI is given by

$$FI_i = \frac{f_i'(\theta)^2}{n_i(\theta)}, \quad (2)$$

where $f_i'(\theta)$ is the differential firing rate of the neuron to the target and distractor orientations, and $n_i(\theta)$ is the variance of the firing rate, which, under the assumption of Poisson noise, is equal to $f_i(\theta)$. In this context, information conveyed by a single neuron is high along regions of the tuning function that undergo the largest firing rate modulation in response to target and distractor orientations. Note that the FI metric defines the “informativeness” of a neuron simply as the differential firing rate evoked by targets and distractors, without regard for the sign of the difference. FI falls to zero at regions of the tuning function in which similar stimuli evoke approximately the same response and the slope of the tuning function goes to zero.

Using Equation 2, FI can be computed for each neuron in a population with respect to discriminating a 90° target from 85° distractors (see Fig. 3B). The most informative neurons are those tuned slightly away from the to-be-discriminated stimulus features (i.e., those neurons whose tuning functions have high slopes around the target and distractor orientations). Given this prediction of how much information each neuron contributes to discriminating the target from distractors, we can infer how attentional gain should be most effectively applied; contrary to the optimal gain hypothesis, the FI metric suggests that attentional gain should be applied to neurons tuned to orientations on either side of the target and distractor features, because these populations of neurons are equally informative.

Although the discussion above focuses on fine discrimination, FI can also be used to compute the effectiveness of a single neuron in discriminating between any two arbitrary stimulus values (Itti et al., 2000). For example, when faced with a detection task (stimulus present/absent), the most informative neurons are those tuned to the target because they undergo the largest change in firing rate. When performing a coarse

discrimination (as in experiment 3), the FI metric also predicts gain profiles that diverge from the optimal gain model. FI holds that attentional gain should be applied to sensory neurons that prefer the target and/or to those that prefer the distractor because both of these neuronal populations will undergo a similarly large change in firing rates when stimulated by either stimulus (however, the changes in firing rate will have opposite signs).

Summary of model predictions

In the context of difficult fine discriminations, which is the main focus of the present study, the optimal gain hypothesis predicts that attentional gain should be applied to neurons that respond more to targets compared with distractors (i.e., neurons tuned to an exaggerated target feature). In contrast, the FI metric suggests that attentional gain should be applied to neurons that undergo a large differential response to the target and distractor, regardless of the sign of the difference. In the context of a coarse discrimination (experiment 3), both models predict high gain for neurons tuned to the target, with FI additionally predicting gain deployed to neurons tuned to the distractors.

It is important to note that the exact shape of the gain functions shown in Figure 3, *A* and *B*, depend on the assumed bandwidth of orientation-selective neurons, and, although we have assumed a single average bandwidth across subjects and orientations, these could in reality differ. However, changes in bandwidth should only affect the dispersion of the gain functions that are actually measured during the experiments and are not likely to account for any systematic biases that are observed in the attentional gain profiles. For example, a subject with high-bandwidth tuning functions might boost the gain of neurons tuned $>20^\circ$ from the target during a fine discrimination because those neurons have the maximal derivative at the target/distractor orientations, whereas a subject with low-bandwidth tuning functions might boost the gain of neurons tuned only 10° from the target during a fine discrimination. That said, we can think of one highly implausible case in which the bandwidth might influence the pattern of attentional gain during the fine discrimination task (experiment 2). If the sensory neurons that support orientation perception have such narrow bandwidths that a 5° offset is effectively a coarse discrimination, then enhancing the gain of neurons tuned to the target orientation would be optimal even during a fine discrimination. However, this would require tuning bandwidths of much less than 5° , which is far smaller than the estimated neuronal bandwidth size of $\sim 40\text{--}50^\circ$ in either monkey V1 or human V1. Furthermore, accuracy in the fine discrimination task was low for all subjects, even at long exposure durations (see Fig. 4), indicating that the task was not easy for any of our subjects and arguing against the notion that a 5° offset was treated as a coarse discrimination.

Correlating changes in attentional gain with visual search performance

Correlation analyses were used to examine how the relative gain that subjects applied to different orientation-selective neurons—as estimated using contrast detection thresholds—affected their performance on the fine discrimination search task in experiment 2. The goal of these analyses was to determine whether the attentional gain profiles predicted overall fine discrimination performance and the amount of improvement on the fine discrimination task that occurred across repeated testing sessions. First, two indices were computed: (1) a difference score between the contrast detection threshold at the target orientation and the mean thresholds for the distractor and exaggerated distractor orientations (-5° , -10° , -20° , and -40°), and (2) a difference score between the contrast detection threshold at the target orientation and the mean thresholds for the exaggerated target orientations ($+5^\circ$, $+10^\circ$, $+20^\circ$, and $+40^\circ$). Because both indices share a common data point (contrast detection threshold at the target orientation), they are not completely independent; therefore, no direct comparisons between the indices were performed. According to the optimal gain model, those subjects who most successfully enhanced the gain of neurons tuned to exaggerated target orientations should fare best on the fine discrimination visual search task. However, the FI metric predicts that gain applied to neurons flanking the target in either direction should be equally predictive of




Figure 3. Predicting attentional gain using the optimal gain hypothesis and Fisher information. *A*, Hypothetical attentional gain function for each neuron in a population based on the optimal gain model proposed by Navalpakkam and Itti (2007) when discriminating a 90° target from an 85° distractor. The model predicts that the highest degree of gain should be applied to neurons tuned to exaggerated target features because these neurons undergo the largest positive firing rate change in response to targets compared with distractors (see Eq. 1 in Materials and Methods and Fig. 1*B*). *B*, FI for each neuron in a population for discriminating a 90° target from an 85° distractor. Information is high on both sides of the target orientation because these neurons undergo a large differential response to targets and distractors (regardless of the sign of the difference). Note that the exact shape of the optimal gain function in *A* and the FI function in *B* depends on the bandwidth of the underlying sensory neurons (which was 45° in these simulations, in line with estimates of both primate and human bandwidths in V1).

visual search performance. Note that all correlation coefficients reported in Results and their associated p values were computed using ordinary least-squares linear regression; however, p values based on robust regression are also reported to evaluate the possibility that the effects were unduly influenced by outliers.

Results

Experiment 1: inferring attentional gain using a target selection task

The optimal gain and FI metrics depicted in Figure 3 predict different ways in which attentional gain might be deployed to support difficult fine discriminations. To evaluate how individuals actually deploy attentional gain when faced with a fine discrimination, and specifically to determine whether attentional gain can be flexibly deployed rather than simply applied to neurons tuned to the target, we first used a modified version of a task developed by Navalpakkam and Itti (2007) (Fig. 2). Recall that the optimal gain hypothesis (Fig. 3*A*) predicts that subjects will enhance the gain of neurons tuned to exaggerated target orientations, whereas the FI metric (Fig. 3*B*) predicts that subjects will modulate the gain of neurons tuned just away from the target in either direction.

On two-thirds of the trials (termed “fine discrimination” trials), subjects had to identify the spatial position of the target stimulus with an unspeeded button-press response. Accuracy improved as the exposure duration of the search array increased (one-way repeated measures ANOVA, $F_{(3,39)} = 8.235$, $p < 0.001$) (Fig. 4*A*, dotted line). However, performance was well below ceiling for all exposure durations, indicating that, although the target was unique from the uniform field of distractors, the orientation offset was small enough so that the target did not “pop out” from the distractors.

The remaining one-third of the trials consisted of a secondary target selection task that was used to infer how subjects were deploying attentional gain in preparation for the expected fine discrimination task (Fig. 2) (see Materials and Methods). Navalpakkam and Itti (2007) reasoned that they could estimate how attentional gain was being applied to different populations of neurons based on the frequency with which each presented orientation was selected in place of the actual target. For example, if subjects were boosting the gain of neurons tuned to an exaggerated target feature in accord with the optimal gain hypothesis,

then Navalpakkam and Itti (2007) reasoned that distractors oriented $+5^\circ$ or more beyond the target should be selected with higher frequency than the actual target orientation.

Figure 4B shows the percentage of times a particular orientation was reported as the target (of the total number of times that each orientation was presented, because only three of the possible nontarget orientations were displayed on a given trial). Of all available orientations, subjects most often selected the Gabor patch rotated $+10^\circ$ from the target; the Gabor rotated -20° from the target was selected the least. A one-way repeated-measures ANOVA revealed a significant bias in the distribution of responses ($F_{(8,104)} = 25.75, p < 0.001$). These data are consistent with the behavioral results of Navalpakkam and Itti (2007) that were interpreted to indicate a bias in the distribution of attentional gain toward the exaggerated target feature. However, because subjects were searching for a target embedded in distractors that shared a common orientation on two-thirds of the trials, subjects may have developed an internal representation of the target as being more different from the distractors than it actually was (because of a so-called “repulsion effect”) (Gibson and Radber, 1937; Coltheart, 1971; Pouget and Bavelier, 2007). Therefore, instead of indexing changes in attentional gain, subjects may have disproportionately reported an exaggerated feature in place of the target because their internal representation of the target orientation was skewed by contextual factors. Pouget and Bavelier (2007) raised the possibility that this repulsion effect might actually be related to a biased distribution of attentional gain.

Experiment 2: inferring attentional gain using a contrast detection probe

Here, we designed an alternate approach to estimating attentional gain functions that did not rely on assumptions about the relationship between repulsion effects and attention. As in experiment 1, a fine discrimination task was performed on two-thirds of all trials. However, on the remaining one-third of the trials, the amount of attentional gain applied to neurons tuned to various orientations ($0^\circ, \pm 5^\circ, \pm 10^\circ, \pm 20^\circ, \text{ or } \pm 40^\circ$ from the target) was estimated using a contrast detection task (Fig. 2) (see Materials and Methods). Because detection sensitivity depends primarily on the gain level of neurons tuned to the to-be-detected feature (Regan and Beverley, 1985; Itti et al., 2000), we reasoned that subjects should be most sensitive to detect stimuli rendered in orientations corresponding to the neuronal populations receiving attentional gain in preparation for the expected fine discrimination task (and this increase in sensitivity should manifest as lower contrast detection thresholds). Because the reported attribute of the display—the location of the single Gabor—was orthogonal to stimulus orientation, the presence of any repulsion effects induced by the dominant fine discrimination task should not have biased responses about spatial position during the contrast detection trials (although the precise manner in which attentional gain was applied to orientation-selective neurons may depend to some extent on the subject’s internal representation of the target orientation).

As in experiment 1 described above, average accuracy on the main fine discrimination task improved as the exposure duration of the search array increased (one-way repeated measures ANOVA, $F_{(3,81)}$




Figure 4. Psychophysical data revealing how subjects deploy attentional gain in experiments 1 and 2. **A**, Accuracy on the main fine discrimination (FD) task as a function of search array exposure duration (experiment 1, dotted line; experiment 2, solid line). **B**, Proportion of trials that stimuli rendered in each possible orientation were selected in place of the target (experiment 1). Positive values along the x-axis refer to rotation in the direction indicated by the color of the cue, and negative values refer to rotation in the direction opposite of that indicated by the cue. For example, if a red cue indicated that targets were rotated clockwise with respect to distractors, then by convention all distractors rotated clockwise from the target would be denoted with a positive value and all distractors rotated counterclockwise would be denoted with a negative value. **C**, Normalized contrast detection thresholds for the entire group of 28 subjects in experiment 2. The x-axis labels refer to orientation offset of the to-be-detected Gabor from the target orientation, following the same sign convention used in **B**. Note that since there is only one distractor orientation in experiment 2, positive rotational offsets denote exaggerated target features and negative offsets denote the distractor feature (-5° from the target) and exaggerated distractor features. All error bars are ± 1 SEM.

$= 90, p < 0.001$) (Fig. 4A, solid line), and performance was well below ceiling for all exposure durations.

Figure 4C shows the normalized thresholds estimated for each orientation offset on the relatively rare contrast detection trials. Normalization was performed by subtracting the mean contrast level across all orientation offsets for each subject. This was done to remove between-subject variability because orientation offset was a within-subject manipulation (thus the normalization had no impact on the shape of the gain function or the repeated-measures statistics) (for the non-normalized data, see supplemental Fig. 1, available at www.jneurosci.org as supplemental material). Keep in mind that increases in attentional gain applied to neurons tuned to a particular orientation should give rise to lower contrast detection thresholds associated with that orientation.

As shown in Figure 4C, contrast detection thresholds were high near the target orientation ($0^\circ, +5^\circ$) and low for more exaggerated target orientations ($+10^\circ, +20^\circ$), in rough accord with the optimal gain hypothesis depicted in Figure 3A (one-way repeated measures ANOVA, $F_{(8,216)} = 2.2, p < 0.05$). However, contrast detection thresholds were similarly low for the distractor and exaggerated distractor orientations ($-5^\circ, -10^\circ$ from the target orientation). The shape of the gain function averaged across all days is slightly asymmetric $\sim 0^\circ$, as contrast detection thresholds were lowest at -5° on the one side, and at $+10^\circ$ to $+20^\circ$ on the other; however, by day 5, this asymmetry disappeared (see Fig. 6 and surrounding text). The observed pattern of results directly conflicts with the predictions of the optimal gain model because neurons tuned to the distractor and exaggerated distractor orientations respond relatively less to the target, and therefore attentional suppression should be evident via higher contrast detection thresholds at these orientations (Fig. 3A). The complete contrast detection threshold function is instead more consistent with predictions based on the FI metric (Fig. 3B). One caveat is that contrast detection thresholds were slightly lower at the distractor orientation than at the exaggerated distractor orientation, although neurons tuned to the exaggerated distractor orientation are probably more informative given the average bandwidth of neurons in V1 ($\sim 45^\circ$) (Fig. 3B).

Perceptual consequences of differences in attentional gain

We next used a correlation analysis to evaluate the relationship between how individual subjects applied attentional gain to different orientations and their success on the main fine discrimination task (see Materials and Methods). The optimal gain model (Fig. 3A) predicts that subjects who most strongly enhanced the

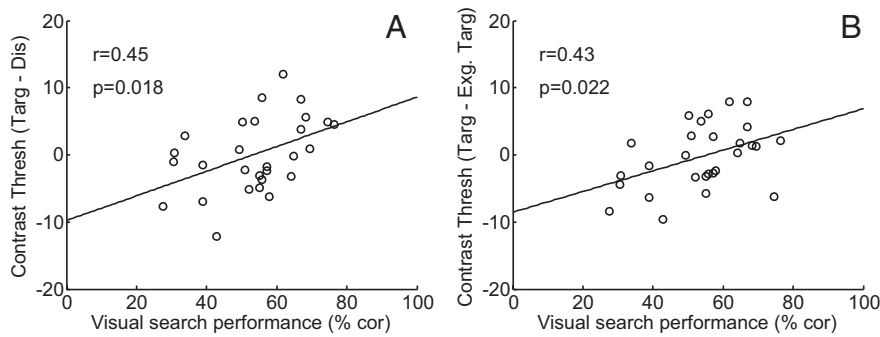


Figure 5. Correlation between attentional gain and fine discrimination accuracy in experiment 2. Correlation on the last day of behavioral testing between the TD index and fine discrimination accuracy (**A**) and between the TET index and fine discrimination accuracy (**B**).

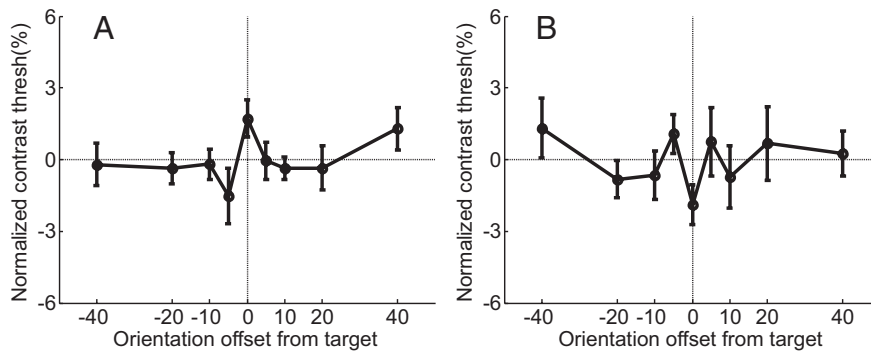


Figure 6. Attentional gain functions for subjects with the highest (**A**) and lowest (**B**) fine discrimination accuracy on the last day of behavioral testing. The most successful subjects had relatively low thresholds for orientations flanking the target on both sides (shown in **A**), whereas the less successful subjects had the lowest threshold at the target orientation (shown in **B**). All error bars are ± 1 SEM.




Figure 7. Correlation between attentional gain and learning in experiment 2. Correlation between TD index (**A**) and the TET index (**B**) and the amount of improvement across testing sessions (accuracy on day 1 subtracted from accuracy on day 5). **C**, Accuracy across testing sessions for the upper and lower half of subjects, sorted based on fine discrimination accuracy on day 5. Not only did the subjects who had the highest accuracy on day 5 also have higher accuracy across all testing sessions, they also showed more improvement with practice. Note that, although the rightmost data points in **C** are predetermined to be different because of the nature of the grouping, all relevant statistics reported in Results were computed only on data from days 1–4 to avoid a non-independence error (hence the break in the line connecting data points on day 4 and day 5). All error bars are ± 1 SEM.

gain of neurons tuned to exaggerated targets should be better at the fine discrimination task. Conversely, the FI metric predicts that visual search performance should be related to the relative gain applied to neurons tuned to either side of the stimuli. We tested these predictions by computing difference scores between contrast detection thresholds at the target orientation and the mean detection thresholds for stimuli rotated in the direction of either the distractor or beyond $[-5^\circ, -10^\circ, -20^\circ, \text{ and } -40^\circ$, forming the target – distractor (TD) index] or beyond the target $[+5^\circ, +10^\circ, +20^\circ, \text{ and } +40^\circ$, forming the target – exaggerated target (TET) index]. Although there was no significant predictive relationship between either of these gain indices and visual search performance across the first four testing sessions (all p values >0.15 ; mean $p = 0.42$), a robust predictive relationship with both

indices emerged on the last day of testing (Fig. 5*A, B*) ($r_{(26)} = 0.45, p < 0.025, p_{\text{robust}} < 0.05$; $r_{(26)} = 0.43, p < 0.025, p_{\text{robust}} < 0.05$, respectively). Recall that high contrast detection thresholds should be associated with less neural gain, hence the positive slope of the regression line. These correlations suggest that subjects who more effectively enhanced the gain of neurons tuned to orientations on either side of the target stimulus performed better on the fine discrimination task compared with those who applied gain to neurons tuned to the target orientation. To further investigate whether the best performing subjects were simultaneously applying attentional gain to both exaggerated targets and distractors, we next compared the gain functions from the best performing half of subjects ($n = 14$) (Fig. 6*A*) and the worst performing subjects ($n = 14$) (Fig. 6*B*) based on their overall fine discrimination accuracy on the last testing session. The functions depicted in Figure 6 are qualitatively different: the best performing subjects had a higher average threshold at the target orientation and relatively low thresholds elsewhere, whereas the subjects who were not as successful on the visual search task had a low threshold at the target orientation and relatively high thresholds elsewhere (between subjects t test revealed a significant difference in the thresholds at the target orientation, $t_{(26)} = 2.37, p < 0.025$). Collectively, these data demonstrate that, by the end of testing, the most successful subjects were those who tended to enhance the gain of neurons tuned just away from the target orientation in either direction; this pattern is most consistent with predictions based on the FI metric.

We next examined the predictive relationship between contrast detection thresholds during the last testing session and the amount of improvement that occurred on the visual search task over all five testing sessions (in which learning is defined as accuracy on day 1 subtracted from accuracy on day 5, collapsed across search array exposure durations).

Both the TD and TET indices predicted the amount of learning that occurred across testing sessions (Fig. 7*A, B*) ($r_{(26)} = 0.39, p < 0.05, p_{\text{robust}} < 0.05$; $r_{(26)} = 0.41, p < 0.05, p_{\text{robust}} < 0.05$, respectively). To more directly convey the nature of the learning effects, we again divided the best performing subjects and the worst performing subjects based on visual search data from day 5 ($n = 14$ per group as described above); Figure 7*C* depicts search accuracy during every testing session for each group. Because group membership was determined based on data from day 5, we avoided a non-independence error by running a two-way mixed-factor ANOVA using only data from days 1–4 to examine how performance between the two groups differed across testing sessions [between-subject factor: accuracy group (two levels, good/poor);

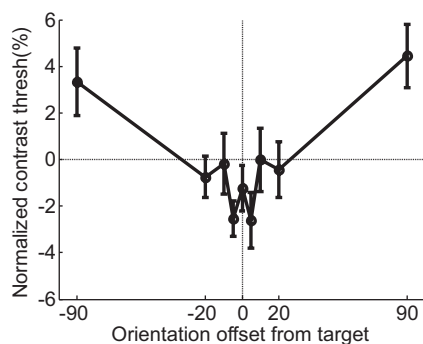


Figure 8. Normalized contrast detection thresholds for all 16 subjects when engaged in a coarse discrimination task (experiment 3). All error bars are ± 1 SEM.

within-subject factor: testing session (four levels, days 1–4)]. Performance generally improved with practice (main effect of testing day, $F_{(3,78)} = 25.1$, $p < 0.001$) and the best performing subjects on day 5 were also better on all other days, including day 1 (main effect of group, $F_{(1,26)} = 22.4$, $p < 0.001$; t test comparing accuracy on just day 1, $t_{(26)} = 2.4$, $p < 0.025$). Most interestingly, however, was the observation that, although the subjects who performed best on day 5 started the experiment with higher accuracy on day 1, they showed more improvement over testing sessions (interaction between testing day and group, $F_{(3,78)} = 10.2$, $p < 0.001$). Interpreted in the context of the contrast detection threshold functions shown in Figure 6, this interaction suggests that those subjects who came to apply attentional gain to orientations flanking the target were also the subjects who improved the most on the task with practice.

Experiment 3: coarse discrimination

To determine whether the gain patterns reported above were specifically related to fine discriminations or whether they were idiosyncratically related to some other aspect of our experimental design, a new group of subjects performed a conceptually similar task involving a coarse discrimination [i.e., find a target rotated 90° from the distractors (Fig. 2) (see Materials and Methods)]. In this situation, subjects should apply attentional gain to neurons tuned to the target orientation (based on the optimal gain account) (Martinez-Trujillo and Treue, 2004; Navalpakkam and Itti, 2007) or to both the target and the distractor orientations (based on the FI account).

Average accuracy in the coarse discrimination task was 55%, indicating that the task was not trivial and that difficulty was approximately equated with the fine discrimination task used in experiment 2. The normalized attentional gain function averaged across all subjects is displayed in Figure 8; the target orientation is indicated by 0° , and the distractor orientation is indicated by $\pm 90^\circ$. A one-way repeated-measures ANOVA revealed that sensory gain was differentially modulated across stimulus orientations ($F_{(8,120)} = 3.73$, $p < 0.001$), and contrast detection thresholds were relatively low around 0° and highest at $\pm 90^\circ$ (comparison of target threshold with distractor thresholds, $t_{(15)} = 2.77$, $p = 0.014$, averaged across $+90^\circ$ and -90° , which were collapsed for simplicity because they were not significantly different, $t_{(15)} = -0.71$, $p = 0.46$, and because they were physically identical stimuli; see Materials and Methods). Interestingly, contrast detection thresholds were also low across a range surrounding the target orientation (e.g., -5° and $+5^\circ$). We speculate that low thresholds in the neighborhood of the target were driven by two factors. First, any imprecision in a subject's ability

to infer the exact orientation of a target rotated 90° from the cue would result in enhanced gain for all neurons tuned to the general vicinity of the target. Second, monkey physiology research suggests that a broader pooling of neurons may be beneficial when making coarse discriminations, whereas pooling across only the most sensitive neurons may be most beneficial when making a fine discrimination (Purushothaman and Bradley, 2005). Nevertheless, the observed gain function is strikingly different from that observed in experiment 2 (compare Fig. 8 with Fig. 4C). Furthermore, the observation of enhanced gain only around the target orientation (and not the distractor orientation) is most consistent with optimal gain theory (as opposed to FI), as well as with existing single-unit recording data (Martinez-Trujillo and Treue, 2004).

Discussion

Here, we used a psychophysical procedure to show that, on average, contrast detection thresholds were lower for flanking orientations around the target when subjects were faced with a very difficult fine discrimination (Fig. 4C). Based on the hypothesized relationship between contrast detection thresholds and neural gain, we propose that neurons tuned to these orientations underwent a larger attentional modulation. Thus, attention maximizes the differential response associated with targets and distractors during a difficult perceptual discrimination, regardless of the sign of this difference (a notion formally captured by the FI metric; see Eq. 2) (Seung and Sompolinsky, 1993; Pouget et al., 2001). More generally, this demonstration of off-channel attentional gain reveals that attention does not simply operate to enhance the activity of neurons tuned to the target but instead maximizes the amount of information available for performing a specific perceptual task. This distinction is important because it is inconsistent with the common intuition that attention primarily increases the perceptual quality of the target (Carrasco et al., 2004; Liu et al., 2009). Instead, attention can bias neural activity away from a veridical representation of the target and toward a more abstract pattern that is specifically tailored to maximize perceptual acuity.

Within the larger group of 28 subjects in experiment 2, we observed individual differences in attentional gain that predicted accuracy on the fine discrimination task (Figs. 5, 6) and the amount of improvement that occurred across repeated testing sessions (Fig. 7). This latter observation is consistent with reports from single-unit physiology suggesting that perceptual learning enhances the firing rates of the most informative sensory neurons or those that undergo the largest firing rate change in response to targets and nontargets (Schoups et al., 2001; Yang and Maunsell, 2004; Raiguel et al., 2006). However, the type of learning we report here is conceptually different from most previous investigations of perceptual learning because the target and distractor orientations were not fixed across the entire experiment. Instead, the orientations were fixed with respect to the cue, so subjects were learning to more efficiently deploy attentional gain based on the advance information provided by the cue as opposed to learning to discriminate a specific visual feature per se.

We speculate that deploying attentional gain to neurons tuned to both sides of the target may be advantageous because the location of the stimulus could then be inferred based on the output of two decision rules (a "max" and a "min" rule) (Zhaoping and May, 2007). For example, consider the response of four distinct populations of neurons that all prefer an exaggerated target feature but that only receive input from one stimulus in the display (that is, the spatial receptive field of each population is restricted

to a single quadrant). The response of these neurons will be relatively weak when stimulated by a distractor and stronger when stimulated by a target. Target discrimination might then be based on the location associated with the neural population that produces the largest response (application of a max rule). Conversely, target discrimination might be based on neurons that respond more to the distractors than to the target (i.e., neurons tuned to an exaggerated distractor), and therefore the target could be found by applying a min rule. Moreover, if attentional gain is simultaneously applied to neurons tuned to both exaggerated target and exaggerated distractor orientations, then the response might be based on the outcome of both of these decision rules, thereby improving the probability of success. However, caution is warranted because this explanation is completely *post hoc*; additional investigation is required to precisely specify how simultaneously enhancing the gain of neurons tuned to orientations flanking the target leads to a more efficient “readout” of activity in early visual areas during perceptual decision making.

In experiment 3, we confirmed that the pattern of results shown in Figure 4C was unique to fine discriminations; when subjects were engaged in a coarse discrimination task, contrast detection thresholds were lowest around the target orientation and highest at the distractor orientation (Fig. 8). Although both the optimal gain model and the FI metric predict a low threshold for the target orientation, the FI metric incorrectly predicts low thresholds for the distractor orientation as well, because neurons tuned to the distractor should be equally discriminating (albeit by responding more to a distractor than to a target). The reason why the FI metric correctly predicted gain patterns for fine discriminations but not coarse discriminations is not entirely clear. However, when the target and distractors are orthogonal to each other (as in a coarse discrimination), the signal-to-noise ratio (SNR) is quite high for neurons tuned to the target orientation and is probably not the limiting factor in search performance. Therefore, taking into account contributions from neurons tuned to the distractor orientation may be unnecessary in most situations. In contrast, in a fine discrimination task, the overall SNR is relatively low, and thus it may typically be advantageous to apply gain to all neurons that undergo a large differential firing rate, regardless of the sign of the difference, to maximize the probability of discriminating the target.

Navalpakkam and Itti (2007) also performed behavioral experiments to examine how attentional gain is applied during a difficult visual search task (e.g., discriminate a 55° target from 60° distractors). Although we replicated their observation of a selective response bias toward exaggerated target features in experiment 1, the results from experiment 2 are more consistent with the FI metric. Given the disparate conclusions, it is important to consider how the experimental designs varied. First, we designed experiment 2 specifically to avoid any undue influence of response bias induced by a repulsion effect (see section of Results related to experiment 1). Second, we used a trial-by-trial cueing design, as opposed to a block design, to equate sensory stimulation at every possible orientation to rule out differential sensory adaptation as a confounding factor when estimating contrast detection thresholds in experiments 2 and 3 (see Materials and Methods). This trial-by-trial fluctuation in the cued orientation, combined with the use of color to indicate the rotational offset of the target, raises the possibility that subjects did not adopt a robust attentional set for the relevant target feature. However, we used an identical cueing procedure in experiment 1 and the systematic and robust target selection bias we observed confirms that subjects were capable of updating the cued orientation on a

trial-by-trial basis. Third, the orientation of our attention cue indicated the distractor feature as opposed to the target, which may have encouraged a strategy of boosting the gain of neurons tuned to the distractor and to exaggerated distractors. However, the response bias toward exaggerated targets observed in experiment 1 demonstrates that subjects were able to use the cue to correctly infer the rotational offset of the target from the distractors (Fig. 4B). Likewise, the distractor orientation was cued in experiment 3, yet the data revealed lower contrast detection thresholds around the target orientation, despite the fact that the target was rotated 90° with respect to the distractors. The overall pattern of results across all three experiments therefore demonstrates that subjects were capable of using the orientation and color of the cue to accurately infer the relevant target feature on a trial-by-trial basis.

Although the present report focuses on understanding how attentional gain operates in the context of a difficult orientation discrimination task, we expect that similar principles will apply to other types of perceptual judgments as well. This is particularly true given that the orientation discrimination task used here likely relies to a large degree on gain modulations in V1, in which attention effects are thought to be relatively small compared with extrastriate visual areas such as V4 or the middle temporal area (MT) (Kastner et al., 1998; Saenz et al., 2002). Thus, in other situations—say when discriminating between two similar directions of motion—we predict that the influence of attention on the most informative sensory neurons should be even larger. In addition, an intriguing possibility is that the shape of attentional gain functions might be qualitatively distinct at different points along the visual hierarchy within the context of the same perceptual task. In the present experiment, for example, we purposefully used stimuli that were defined by a single critical attribute (orientation) so that we could assess attentional gain functions using relatively straightforward psychophysical procedures. However, consider a conceptually similar task that required discriminating between two stimuli that were more complex (e.g., two letters such as R and A). In this case, off-channel gain in V1 might help to distinguish the orientation of each component line, whereas neurons in higher-order visual areas that are sensitive to constellations of features might benefit from gain applied to neurons that are maximally responsive to each letter. This type of mixed strategy might be especially advantageous when dealing with complex natural images that engender simultaneous analysis at many levels of detail. We therefore predict that recording neural activity at multiple levels of the visual hierarchy will reveal that attention optimizes cortical representations of relevant stimuli in a far more complex manner than has been appreciated to date.

The present observation of off-channel gain also complements recent data that highlights the flexible and adaptive nature of attentional modulations. For example, David et al. (2008) recently demonstrated that the orientation and spatial frequency tuning preferences of neurons in V4 shift toward behaviorally relevant features contained in natural scenes; analogous shifts have also been observed in auditory cortex (David et al., 2008; Mesgarani et al., 2008). Spatial receptive fields in V4 and MT also shift toward attended stimuli, leading to an increase in the overall number of neurons that encode sensory information (Connor et al., 1997; Tolias et al., 2001; Womelsdorf et al., 2008). Although these previous reports did not explicitly determine whether attention-mediated changes in tuning characteristics are optimized in an information-theoretic sense, some interesting predictions follow from the present results. For example, future experiments might require discriminating between two natural

images that differed by varying degrees in terms of orientation and spatial frequency composition. Using single-unit recording and the spectral receptive field estimation techniques used by David et al. (2008) (see also Theunissen et al., 2001; Wu et al., 2006), experimenters could determine whether orientation and spatial frequency tuning functions were shifted in a manner that maximized the amount of information available for performing the specified perceptual task. Although many such questions remain to be addressed, the emerging view is that attention does not simply amplify the response of sensory neurons that are tuned to the target of search. Instead, attention optimizes the gain of sensory neurons in a highly flexible and adaptive manner to facilitate whatever perceptual task is currently relevant to the observer.

References

- Boynton GM (2005) Attention and visual perception. *Curr Opin Neurobiol* 15:465–469.
- Brainard DH (1997) The Psychophysics Toolbox. *Spat Vis* 10:433–436.
- Butts DA, Goldman MS (2006) Tuning curves, neuronal variability, and sensory coding. *PLoS Biol* 4:e292.
- Carrasco M, Ling S, Read S (2004) Attention alters appearance. *Nat Neurosci* 7:308–313.
- Coltheart M (1971) Visual feature-analysers and aftereffects of tilt and curvature. *Psychol Rev* 78:114–121.
- Connor CE, Preddie DC, Gallant JL, Van Essen DC (1997) Spatial attention effects in macaque area V4. *J Neurosci* 17:3201–3214.
- David SV, Hayden BY, Mazer JA, Gallant JL (2008) Attention to stimulus features shifts spectral tuning of V4 neurons during natural vision. *Neuron* 59:509–521.
- Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18:193–222.
- Fahle M (2004) Perceptual learning: a case for early selection. *J Vis* 4:879–890.
- Gibson JJ, Radner M (1937) Adaptation, aftereffect and contrast in the perception of tilted lines. *J Exp Psychol* 20:453–467.
- Hol K, Treue S (2001) Different populations of neurons contribute to the detection and discrimination of visual motion. *Vision Res* 41:685–689.
- Itti L, Koch C, Braun J (2000) Revisiting spatial vision: toward a unifying model. *J Opt Soc Am A Opt Image Sci Vis* 17:1899–1917.
- Jazayeri M, Movshon JA (2006) Optimal representation of sensory information by neural populations. *Nat Neurosci* 9:690–696.
- Jazayeri M, Movshon JA (2007) A new perceptual illusion reveals mechanisms of sensory decoding. *Nature* 446:912–915.
- Kahneman D (1973) *Attention and effort*. Englewood Cliffs, NJ: Prentice-Hall.
- Kastner S, De Weerd P, Desimone R, Ungerleider LG (1998) Mechanisms of directed attention in the human extrastriate cortex as revealed by functional MRI. *Science* 282:108–111.
- Liu T, Abrams J, Carrasco M (2009) Voluntary attention enhances contrast appearance. *Psychol Sci* 20:354–362.
- Martinez-Trujillo JC, Treue S (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr Biol* 14:744–751.
- Maunsell JH, Treue S (2006) Feature-based attention in visual cortex. *Trends Neurosci* 29:317–322.
- McAdams CJ, Maunsell JH (1999) Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J Neurosci* 19:431–441.
- Mesgarani N, David SV, Fritz JB, Shamma SA (2008) Phoneme representation and classification in primary auditory cortex. *J Acoust Soc Am* 123:899–909.
- Mruczek RE, Sheinberg DL (2007) Context familiarity enhances target processing by inferior temporal cortex neurons. *J Neurosci* 27:8533–8545.
- Navalpakkam V, Itti L (2007) Search goal tunes visual features optimally. *Neuron* 53:605–617.
- Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis* 10:437–442.
- Pouget A, Bavelier D (2007) Paying attention to neurons with discriminating taste. *Neuron* 53:473–475.
- Pouget A, Deneve S, Latham PE (2001) The relevance of Fisher Information for theories of cortical computation and attention. In: *Visual attention and cortical circuits* (Braun J, Koch C, Davis JL, eds). Cambridge, MA: MIT.
- Purushothaman G, Bradley DC (2005) Neural population code for fine perceptual decisions in area MT. *Nat Neurosci* 8:99–106.
- Raiguel S, Vogels R, Mysore SG, Orban GA (2006) Learning to see the difference specifically alters the most informative V4 neurons. *J Neurosci* 26:6589–6602.
- Regan D, Beverley KI (1985) Postadaptation orientation discrimination. *J Opt Soc Am A* 2:147–155.
- Ringach DL, Bredfeldt CE, Shapley RM, Hawken MJ (2002) Suppression of neural responses to nonoptimal stimuli correlates with tuning selectivity in macaque V1. *J Neurophysiol* 87:1018–1027.
- Saenz M, Buracas GT, Boynton GM (2002) Global effects of feature-based attention in human visual cortex. *Nat Neurosci* 5:631–632.
- Schoups A, Vogels R, Qian N, Orban G (2001) Practising orientation identification improves orientation coding in V1 neurons. *Nature* 412:549–553.
- Seung HS, Sompolinsky H (1993) Simple models for reading neuronal population codes. *Proc Natl Acad Sci U S A* 90:10749–10753.
- Snowden RJ (1992) Orientation bandwidth: the effects of spatial and temporal frequency. *Vision Res* 32:1965–1974.
- Theunissen FE, David SV, Singh NC, Hsu A, Vinje WE, Gallant JL (2001) Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli. *Network* 12:289–316.
- Tolias AS, Moore T, Smirnakis SM, Tehovnik EJ, Siapas AG, Schiller PH (2001) Eye movements modulate visual receptive fields of V4 neurons. *Neuron* 29:757–767.
- Watson AB, Pelli DG (1983) QUEST: a Bayesian adaptive psychometric method. *Percept Psychophys* 33:113–120.
- Womelsdorf T, Anton-Erxleben K, Treue S (2008) Receptive field shift and shrinkage in macaque middle temporal area through attentional gain modulation. *J Neurosci* 28:8934–8944.
- Wu MC, David SV, Gallant JL (2006) Complete functional characterization of sensory neurons by system identification. *Annu Rev Neurosci* 29:477–505.
- Yang T, Maunsell JH (2004) The effect of perceptual learning on neuronal responses in monkey visual area V4. *J Neurosci* 24:1617–1626.
- Zhaoping L, May KA (2007) Psychophysical tests of the hypothesis of a bottom-up saliency map in primary visual cortex. *PLoS Comput Biol* 3:e62.