Behavioral/Cognitive

Feature-Based Attention Multiplicatively Scales the fMRI-BOLD Contrast-Response Function

Joshua J. Foster[1,2] and Sam Ling[1,2]

[1]Department of Psychological and Brain Sciences, Boston University, Boston, Massachusetts 02215, and [2]Center for Systems Neuroscience, Boston University, Boston, Massachusetts 02215

fMRI plays a key role in the study of attention. However, there remains a puzzling discrepancy between attention effects measured with fMRI and with electrophysiological methods. While electrophysiological studies find that attention increases sensory gain, amplifying stimulus-evoked neural responses by multiplicatively scaling the contrast-response function (CRF), fMRI appears to be insensitive to these multiplicative effects. Instead, fMRI studies typically find that attention produces an additive baseline shift in the BOLD signal. These findings suggest that attentional effects measured with fMRI reflect top-down inputs to visual cortex, rather than the modulation of sensory gain. If true, this drastically limits what fMRI can tell us about how attention improves sensory coding. Here, we examined whether fMRI is sensitive to multiplicative effects of attention using a feature-based attention paradigm designed to preclude any possible additive effects. We measured BOLD activity evoked by a probe stimulus in one visual hemifield while participants (6 male, 6 female) attended to the probe orientation (attended condition), or to an orthogonal orientation (unattended condition), in the other hemifield. To measure CRFs in visual areas V1-V3, we parametrically varied the contrast of the probe stimulus. In all three areas, feature-based attention increased contrast gain, improving sensitivity by shifting CRFs toward lower contrasts. In V2 and V3, we also found an increase in response gain, an increase in the responsivity of the CRF, that was greatest at inner eccentricities. These results provide clear evidence that the fMRI-BOLD signal is sensitive to multiplicative effects of attention.

Key words: attention; BOLD; contrast-response functions; feature-based attention; fMRI

Significance Statement

fMRI plays a central role in the study of attention because it allows researchers to precisely and noninvasively characterize the effects of attention throughout the brain. Electrophysiological studies have shown that attention increases sensory gain, amplifying stimulus-evoked neural responses. However, a growing body of work suggests that the BOLD signal that is measured with fMRI is not sensitive to these multiplicative effects of attention, calling into question what we can learn from fMRI about how attention improves sensory codes. Here, using a feature-based attention paradigm, we provide evidence that the BOLD signal can pick up multiplicative effects of attention.

Introduction

Attention plays a central role in perception, prioritizing processing of relevant information. One way that attention improves neural coding is by increasing sensory gain, selectively amplifying stimulus-evoked responses in sensory areas (Hillyard et al., 1998; Carrasco, 2011). For instance, attention increases the spike rate of individual neurons in nonhuman primates (McAdams and Maunsell, 1999; Treue and Martínez-Trujillo, 1999) and the amplitude of visually evoked potentials measured with EEG in humans (van Voorhis and Hillyard, 1977; Morgan et al., 1996). More recent work has characterized how the effect of attention scales with stimulus drive by varying stimulus contrast to measure contrast-response functions (CRFs). Unit-recording and EEG studies that have taken this approach have found that attention multiplicatively scales CRFs, either vertically scaling the CRF, increasing responsivity (response gain, Fig. 1a), or shifting
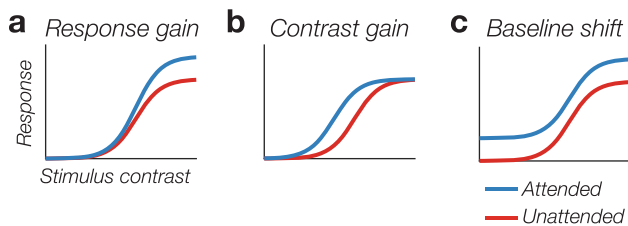
Figure 1. Attentional modulations of CRFs. Attention can modulate CRFs in three ways. **a**, Response gain: attention increases the amplitude of the CRF, increasing responsivity. **b**, Contrast gain: attention shifts the CRF leftward, increasing sensitivity. **c**, Baseline shift: attention produces an additive increase in the CRF.

the CRF toward lower contrasts, increasing sensitivity (contrast gain, Fig. 1b) (Reynolds et al., 2000; Martínez-Trujillo and Treue, 2002; Williford and Maunsell, 2006; Y. J. Kim et al., 2007; Itthipuripat et al., 2014a,b, 2019; Foster et al., 2021). These sensory gain effects have served as the empirical backbone for computational models of attention (Reynolds and Heeger, 2009).

fMRI has long played a central role in attention research, allowing researchers to precisely and noninvasively measure the effects of attention throughout the human brain. Like electrophysiological studies, fMRI studies report larger responses to attended stimuli than to unattended stimuli (Carrasco, 2011). It is often assumed that these effects reflect an amplification of stimulus-driven activity. However, fMRI studies of attention that have measured CRFs reveal a qualitatively different pattern than electrophysiological studies. Unlike the multiplicative effects seen in electrophysiology, fMRI studies have found that attention produces an additive baseline shift in the BOLD signal (Fig. 1c) (Somers and McMains, 2005; Buracas and Boynton, 2007; Murray, 2008; Pestilli et al., 2011; Sprague et al., 2018; Itthipuripat et al., 2019; but see X. Li et al., 2008). Because a baseline shift is independent of stimulus strength, and is seen in the absence of visual input, this pattern suggests that attentional effects measured with fMRI reflect top-down input to visual cortex, rather than a multiplicative modulation of stimulus-driven activity (Murray, 2008; Itthipuripat et al., 2014a). Understanding whether the BOLD signal is sensitive to multiplicative attention effects is critical for understanding the degree to which fMRI can be used to study attention: if the BOLD signal is not sensitive to modulations of stimulus-driven activity, then this drastically limits what fMRI can tell us about how attention improves sensory codes.

Although studies of spatial attention suggest that fMRI is insensitive to multiplicative effects of attention, a seminal study of feature-based attention suggests otherwise. Saenz et al. (2002) measured the effect of feature-based attention on the BOLD response to a spatially unattended probe stimulus. They found that attention to a specific feature (e.g., red or upward motion) increased the BOLD response to the probe stimulus when it matched the attended feature. This effect appears to reflect a multiplicative amplification of the probe-evoked response rather than a baseline shift. Although feature-based attention does produce baseline shifts in neural activity (Chelazzi et al., 1993; Serences and Boynton, 2007; Stokes et al., 2009), these baseline shifts are feature-specific, selectively increasing the baseline activity of neurons tuned for the attended feature. Thus, the baseline activity of the entire neuronal population, when aggregating across neurons with different feature preferences, is not expected to vary with the attended feature. Although the results of Saenz et al. (2002) hint at a multiplicative effect of

feature-based attention, no study has assessed the effect of feature-based attention on the full CRF, which is essential to distinguish between different kinds of attentional modulation because multiple mechanisms can account for an increase in the response at a single contrast level (Fig. 1).

Here, we provide that definitive test. We manipulated feature-based attention in one hemifield, and measured responses to a probe stimulus in the other hemifield, exploiting the well-established spatial spread of feature-based attention (Treue and Martínez-Trujillo, 1999; Saenz et al., 2002). Critically, however, we parametrically manipulated the contrast of the probe stimulus to measure the CRF in visual cortex, using a recently developed paradigm that captures compressive nonlinearities in the BOLD CRF (Vinke et al., 2022). In doing so, we found robust multiplicative effects of feature-based attention on the CRF throughout early visual cortex. Throughout early visual cortex, feature-based attention increased contrast gain, shifting CRFs leftward. For V3 and a subset of eccentricities in V2, we also found an increase in response gain. These results provide clear, positive evidence that BOLD responses are sensitive to multiplicative effects of feature-based attention on stimulus-driven responses.

Materials and Methods
Participants
Thirteen healthy volunteers (7 male and 6 female, mean age = 27.8 years, SD = 6.3 years) participated in the study. Participants were between 19 and 42 years old, and reported normal or corrected-to-normal visual acuity. Our target sample size was 12 participants. We replaced one participant because of very weak activation of visual cortex in the independent visual localizer scan (see Visual localizer). Thus, we were not confident that the signal-to-noise ratio would be adequate to measure CRFs for this subject. The decision to exclude this subject was made solely on the basis of their visual localizer data, and we did not analyze their data in the feature-based attention task. Thus, our final sample comprised 12 participants (6 male and 6 female, mean age = 27.6 years, SD = 6.6 years). All participants provided written informed consent, and completed a screening form to ensure that they had no MRI-related contraindications. Participants received monetary compensation for their time ($15/h for behavioral sessions and $40/h for MRI sessions), except for two participants who were the authors of the study. All procedures were approved by the Boston University Institutional Review Board.

Apparatus and stimuli
Stimuli were generated using MATLAB (The MathWorks) and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). In the MRI scanner, participants viewed stimuli on a γ-corrected, rear-projection screen (ProPixx DLP LED, VPixx Technologies; refresh rate: 60 Hz; resolution: 1024 × 768 pixels; viewing distance: ~99 cm) that was mounted inside the scanner bore. In behavioral sessions, participants were seated at a desk with their chin on a padded chin rest and viewed stimuli on a γ-corrected LCD monitor (Display++, Cambridge Research Systems; refresh rate: 100 Hz; 1440 × 1080 pixels; viewing distance: ~135 cm). In all sessions, the display was the only light source in the testing room. Participants responded using a two-button response box in the scanner, and using the "1" and "2" keys on the number pad of a standard keyboard in behavioral sessions.

Stimuli were presented on a mean-luminance background (~150 cd/m²). A light-gray fixation dot (0.2°) was presented in the center of the display. The stimuli were three wedge-shaped gratings (2.2 cycles/°, inner edge: 1.0°, outer edge: 8.5°, SD of Gaussian roll-off: 0.1°; Fig. 2a). The wedges in each hemifield spanned 70° of polar angle above and below the horizontal meridian. The display was divided into a relevant hemifield that contained task-relevant stimuli, and a probe hemifield that contained an irrelevant probe stimulus to evoke visual responses (Fig. 2a). We counterbalanced the relevant hemifield (left or right) across observers. In the relevant hemifield, the wedge was divided into two gratings, one in the upper visual field and one in the lower visual
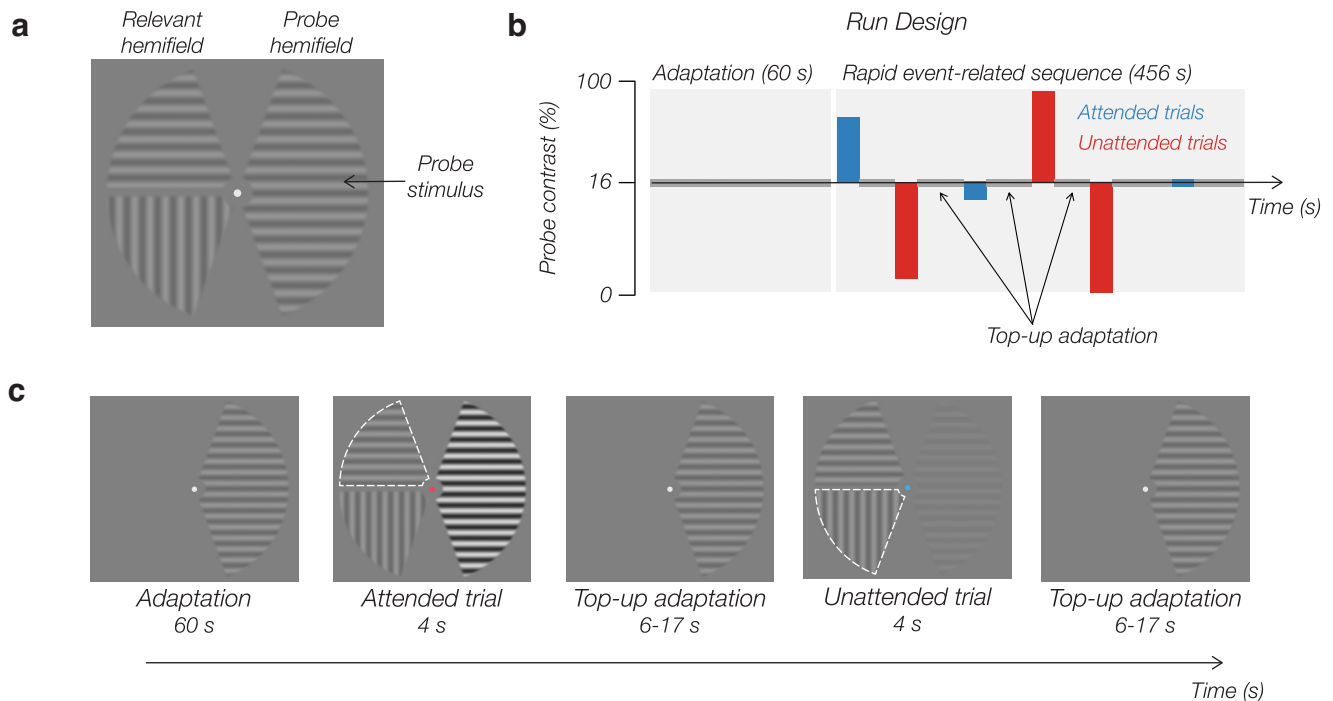
Figure 2. Feature-based attention task. **a**, Stimulus configuration. On each trial of the task, observers were cued to attend to one of two gratings in the relevant hemifield, while we measured responses evoked by a probe stimulus in the other hemifield. **b**, Each fMRI run began with a 60 s adaptation period during which the probe grating was presented at 16% contrast to adapt visual neurons to this contrast level. Following the initial adaptation period, 4 s trials were interleaved with periods of top-up adaptation. The contrast of the probe stimulus was parametrically varied across trials (0%-96%). **c**, Illustration of the task sequence. On each trial, the color of the fixation dot cued observers to attend one of the two gratings in the relevant hemifield (marked by the white dotted outline that was not visible to observers). We manipulated feature-based attention by varying whether observers attended the horizontal or vertical grating in the relevant hemifield. On attended trials, observers attended to the horizontal grating, which matched the orientation of the probe. On unattended trials, observers attended the vertical grating, which did not match the orientation of the probe. The spatial frequency of the gratings in this figure has been reduced for clarity.

field, separated by a 0.5° gap. On each trial of the task, one of these gratings was oriented horizontally, and the other oriented vertically, with each orientation appearing in the upper and lower position equally often. In the probe hemifield, a horizontal probe grating spanned both the upper and lower visual fields. The spatial phase of each grating randomly updated at 10 Hz, with the two horizontal gratings moving in phase together. The two gratings in the relevant hemifield were always presented at 16% Michelson contrast. The probe stimulus was presented at 10 different contrast levels (0%, 2.67%, 4.0%, 5.33%, 8.0%, 16%, 32%, 48%, 64%, 96%), which allowed us to measure probe-evoked CRFs.

Experimental design

Feature-based attention task. Participants completed a rapid event-related task (Fig. 2b,c). We used a contrast-adaptation procedure to promote compressive nonlinearities in CRFs (Vinke et al., 2022) (see Task and behavior). Each run began with a 60 s adaptation period, during which the probe stimulus was presented at 16% contrast (the adapting contrast). The initial adaptation period was followed by a 456 s rapid event-related trial sequence, with 4 s trials interleaved with top-up adaptation periods (6-17 s) to maintain the adaptation state (Fig. 2b). The timing of the trials and order of conditions (20 total: 2 attention conditions × 10 probe contrasts) were sequenced using the Optseq2 schedule optimization tool (Dale, 1999). Participants completed 9 or 10 runs of the task in the scanner. Each run included 40 trials (2 trials per condition). Therefore, we obtained 18-20 trials per condition for each participant.

During each trial, two gratings were presented in the relevant hemifield, and the contrast of the probe stimulus changed (Fig. 2c). The color of the fixation dot (pink or blue) cued the observer to attend the upper or lower grating in the relevant hemifield (with the color-location mapping counterbalanced across observers). The color cue onset 0.5 s before the 4 s trial period and remained on-screen until the end of the trial. On half of trials, the cued grating was horizontal, such that the attended orientation matched the orientation of the probe stimulus (attended trials). On the other half of trials, the cued grating was vertical, such that the

attended orientation did not match the probe stimulus (unattended trials). The observer's task was to monitor the cued grating for targets, which were occasional, brief (100 ms) reductions in the grating's spatial frequency. On each trial, 0, 1, or 2 target reductions could occur. Targets never appeared during the first 0.7 s or the final 0.2 s of the trial. On trials with two targets, the targets were always separated by at least 0.6 s to avoid poor detection of the second target because of the attentional blink (Raymond et al., 1992; Duncan et al., 1994), and the second target always occurred during the final second of the trial to ensure that observers needed to remain engaged with the task for the full trial duration. Similar spatial frequency perturbations occurred in the two unattended gratings to encourage observers to selectively attend to the cued grating. For the probe grating, these perturbations were an increase in spatial frequency rather than a decrease, which ensured that, on attended trials, observers could not perform the task by simply monitoring for a relative decrease in spatial frequency between the cued horizontal grating and the collinear probe grating in the other hemifield. Following the trial period, observers reported how many targets they detected with a button press (no response = zero targets; left button = one target; right button = two targets). Responses needed to be made within 2 s of the end of the trial. At the end of the response window, the fixation dot turned green or red for 0.5 s to indicate whether the response was correct or incorrect, respectively.

Staircase procedure. Before the MRI session, participants completed 1 or 2 behavioral practice sessions (~1.5 h each) of the feature-based attention task to learn the task and to perform initial staircasing to equate difficulty between the attended and unattended conditions. For each condition, we ran separate staircases for trials in which the attended wedge was in the upper or lower visual field to account for potential differences in performance between locations. We titrated difficulty by adjusting the size of the spatial frequency decrement with a weighted up/down procedure: after a correct response, we decreased the size of the spatial frequency decrement by 5%; after an incorrect response, we increased the size of the spatial frequency decrement by 17.6%. The staircases operated continuously throughout the behavioral session(s) and

the MRI session. During the behavioral sessions, participants completed staircasing runs of the task until performance stabilized in each staircase (7-15 runs, mean = 11.6, SD = 2.0). This procedure successfully equated accuracy across conditions during the MRI session: mean accuracy was 76.6% (SD = 2.4) in the attended condition and 77.0% (SD = 2.0) in the unattended condition.

Visual localizer. Participants performed one run of a visual localizer task to identify voxels that were responsive to the wedge stimuli used in the feature-based attention task. The wedge stimuli presented in the localizer task were identical to those in the feature-based attention task (see Apparatus and stimuli), except that all three gratings were horizontally oriented and were presented at full contrast. We alternated 16 s periods of blank and stimulation 8 times, followed by a final 16 s blank period (272 s total). To encourage fixation and maintain alertness, participants monitored the fixation dot for 300 ms decrements in luminance that occurred every 3-5 s, reporting these target events with a button-press. The hit rate was >90% for all participants (mean = 98.2%, SD = 2.4%). We analyzed the visual localizer data using a standard general linear model (GLM) analysis in Freesurfer (Fischl, 2012).

Eye tracking

We monitored eye position in the scanner using an MR-compatible EyeLink 1000 Plus infrared eye tracker (SR Research). Eye position was sampled at 500 Hz, and we calibrated the eye tracker as needed at the start of each run of the task. After removing samples contaminated by blinks, we calculated the mean eye position (relative to the fixation dot) during the 4 s trial period for each condition. The eye tracking data confirmed that observers did a good job of maintaining fixation. Because the relevant hemifield (left or right) was counterbalanced across observers, we flipped the sign of horizontal coordinates for some observers so that positive values indicate eye position toward the relevant hemifield. For vertical eye coordinates, positive values indicate eye positions above the fixation dot, and negative values indicate eye positions below the fixation dot. We observed a very small bias (<0.1°) in vertical eye position toward the location of the cued grating: mean vertical eye position was 0.069° (SEmean = 0.097) when the upper grating was cued and –0.046° (SEmean = 0.116) when the lower grating was cued, $t_{(11)} = 3.29$, $p = 0.007$, but horizontal eye position did not vary as a function of location of the cued location (upper cued: mean = 0.157°, SEmean = 0.055, lower cued: mean = 0.166°, SEmean = 0.048, $t_{(11)} = -0.71$, $p = 0.491$). More importantly, however, eye position did not differ between the attended and unattended conditions: mean vertical eye position (attended: mean = 0.014°, SEmean = 0.106; unattended: mean = 0.010°, SEmean = 0.106, $t_{(11)} = 0.49$, $p = 0.631$) and horizontal eye position (attended: mean = 0.160°, SE = 0.052; unattended: mean = 0.163°, SEmean = 0.050, $t_{(11)} = -0.28$, $p = 0.785$) were near-identical across conditions. Thus, any effects of feature-based attention cannot be attributed to differences in eye position.

MRI data acquisition

All neuroimaging data were acquired at the Boston University Cognitive Neuroimaging Center, on a 3T Siemens Prisma scanner (Siemens Healthcare) using the vendor's 64-channel headcoil. A whole-brain anatomic scan was acquired using a T1-weighted magnetization-prepared rapid gradient multiecho MPRAGE sequence (van der Kouwe et al., 2008), with the following parameters: 1.0 mm^3 voxels, 176 sagittal slices, TR = 2530 ms, TEs = 1.67, 3.55, 5.41, and 7.27 ms, TI = 1100 ms, flip angle = 7°, FOV = 256 mm, GRAPPA (Griswold et al., 2002) acceleration = 2. BOLD data were collected via a T2*-weighted EPI pulse sequence that used multiband RF pulses and simultaneous multislice (SMS) acquisition (Feinberg et al., 2010; Moeller et al., 2010; Setsompop et al., 2012; J. Xu et al., 2013; Cauley et al., 2014). For the task runs, the scan parameters were: 2.0 mm^3 voxels, 70 interleaved axial-oblique slices (25 degrees toward coronal from ACPC alignment), TR = 1000 ms, TE = 30 ms, FA = 64°, FOV = 208 mm, SMS factor = 5, GRAPPA acceleration = 2. The SMS-EPI acquisition used the CMRR-MB pulse sequence from the University of Minnesota.

Anatomical analysis

We analyzed whole-brain T1-weighted anatomic data using the standard 'recon-all' pipeline provided by FreeSurfer (Fischl, 2012) to generate cortical-surface reconstructions, which enabled surface-based registration of functional data to structural data. This allowed us to align population receptive field (pRF) data (see Population receptive fields) to the functional volume space for the feature-based attention task.

fMRI preprocessing

Functional BOLD time-series were corrected for EPI distortions using a reverse phase-encoded method (Andersson et al., 2003) implemented in FSL (Smith et al., 2004). The fieldmap-corrected data were preprocessed with FS-FAST (Fischl, 2012) to apply standard motion-correction procedures, Siemens slice timing correction, and boundary-based registration between functional and anatomic spaces (Greve and Fischl, 2009). No spatial smoothing was applied. To ensure good voxel-wise correspondence of functional data across runs, we applied robust rigid registration (Reuter et al., 2010) using the middle time point of each run. Additionally, for runs of the feature-based attention task, we excluded the first 60 s of data that corresponded to the initial adaptation period (Fig. 2b,c; see Feature-based attention task), before the time-series for each voxel was linearly detrended, high-pass filtered (low cutoff = 0.01 Hz), and converted to z scores. We concatenated the data for all runs of the feature-based attention task. For one participant, we excluded one run of the feature-based attention task because of excessive head motion.

Population receptive fields

pRF mapping was conducted for each participant in a separate scan session. Observers completed 3-5 scans of each of two types of mapping runs: (1) expanding/contracting ring and bar sweep stimuli; and (2) rotating wedge stimuli. The stimuli were composed of a pink noise background with colored objects and faces of varying spatial scale (refreshed at 15 Hz) on mean luminance background (Kay et al., 2013), as used in the Human Connectome Project 7T Retinotopy dataset (Benson et al., 2018). The data were analyzed with the analyzePRF toolbox for MATLAB, which implements the compressive spatial summation pRF model (Kay et al., 2013). Only voxels within the cortical ribbon of the occipital lobe were included in pRF analysis, which were identified using a visual area network label generated using an intrinsic functional connectivity atlas (Yeo et al., 2011). The results of the pRF analysis were used to manually draw ROI labels for early visual areas V1, V2, and V3.

Voxel selection

We used the estimated pRFs in conjunction with the visual localizer (see Visual localizer) data to select voxels for analysis in the feature-based attention task. For each ROI, we excluded voxels whose pRF was centered outside the apertures of the wedge stimuli used in the FBA task, whose pRF model fit was poor ($R^2 < 10\%$), or whose size was unreasonably small (RF < 0.1°). Of the remaining voxels, we identified the 50% most visually responsive voxels based on the localizer data for analysis. On average, this procedure left 160 ± 42 (SD) voxels in V1, 147 ± 25 in V2, and 129 ± 11 voxels in V3 for each hemisphere.

Statistical analysis

We used a GLM to estimate the deconvolved fMRI hemodynamic response function (HRF) in each condition. Specifically, we modeled time points from 0 to 24 s (25 time points) relative to trial onset using a finite impulse response approach, and included run-specific terms to account for any baseline differences between runs. Given each voxel's time-series, we solved a GLM using standard linear regression to obtain 25 weights for each condition, providing an estimate of the voxel's HRF in each condition.

We measured CRFs by taking the average of the deconvolved HRF over a fixed window 3-9 s after stimulus onset for each probe contrast level, per attention condition. We fit the resulting CRFs in each condition with a Naka–Rushton of the following form:

$$R(c) = G_r \frac{c^n}{c^n + G_c^n} + b$$

Where R is the measured BOLD response and c is stimulus contrast. The function had four free parameters: baseline (b), which determines

the offset of the function from zero; response gain (G_r), which determines the amplitude of the function above baseline; contrast gain (G_c), which determines the semisaturation point of the function; and an exponent (n), which determines the slope of the function. We used the "scipy. optimize.curve_fit" function from SciPy (Virtanen et al., 2020) to fit the Naka–Rushton function to measured CRFs. We restricted the b parameter to be between −10 and 10, G_r to be between 0 and 10 (with 10 being a value that far exceeds the observed amplitudes of the CRFs), G_c to be between 0% and 100% contrast, and n to be between 0.1 and 10.

As Itthipuripat et al. (2019) have pointed out, in the absence of a fully saturating CRF, it is possible to obtain unrealistically large estimates of G_r when the best-fitting function saturates outside the range of possible contrasts (i.e., 0%-100%). If the best-fit function saturates above 100% contrast, the maximum value of the function can well exceed the largest measured response. Thus, following Itthipuripat et al. (2019), rather than reporting G_r and G_c, we instead obtained a measure of response gain, R_{max}, by calculating the amplitude of the best-fit Naka–Rushton function at 100% contrast and subtracting the baseline (i.e., $R_{max} = R(100) - b$), and a measure of contrast gain, C_{50}, by calculating the contrast at which the function is halfway between baseline and the response at 100% contrast.

We used a subject-level resampling procedure to test for differences in the parameters of the fitted Naka–Rushton function across conditions. We drew 10,000 bootstrap samples, each containing $N - 1$ many subjects sampled with replacement, where N is the sample size. We drew samples of size $N - 1$ to correct for bias toward narrow confidence intervals in small samples (Hesterberg, 2011). For each bootstrap sample, we fitted Naka–Rushton functions to the mean CRFs across subjects in the bootstrap sample for each condition. For each Naka-Rushton parameter, we calculated the difference for the parameter between the attended and unattended conditions for each bootstrap sample, which yielded a distribution of 10,000 values. We tested whether these difference distributions significantly differed from zero in either direction by calculating the proportion of values >0 or < 0, and doubling the smaller value to obtain a two-sided p value.

Data and code availability
All data and code are available on Open Science Framework at https://osf.io/gaxd2/.

Results

Task and behavior
Observers performed a feature-based attention task (Fig. 2; see Feature-based attention task). We manipulated feature-based attention in one visual hemifield (the relevant hemifield), and measured fMRI-BOLD responses to a spatially unattended probe stimulus in the other hemifield (the probe hemifield). On each trial, observers viewed three flickering gratings (with spatial phase randomly updating at 10 Hz): a pair of gratings (one vertical, one horizontal) in the relevant hemifield, and a horizontal probe grating in the opposite hemifield (Fig. 2a). The color of the fixation dot cued observers to covertly attend to one of the two gratings in the relevant hemifield (while maintaining fixation, see Eye tracking). On half of the trials, observers were cued to attend the horizontal grating in the relevant hemifield that matched the orientation of the probe grating (attended trials). On the other half of trials, observers were cued to attend the vertical grating (unattended trials). Thus, we could measure responses evoked by the irrelevant probe stimulus, while feature-based attention was manipulated in the relevant hemifield. During each trial of the task, observers monitored the cued grating for decrements in spatial frequency, reporting whether they detected 0, 1, or 2
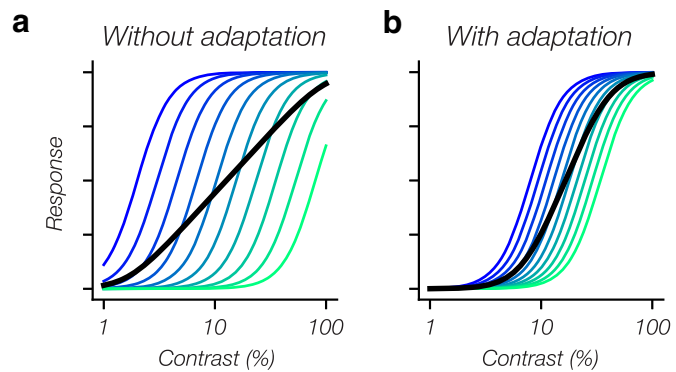

Figure 3. Illustration of how contrast adaptation may promote nonlinear CRFs measured with fMRI. *a*, fMRI measurements aggregate across many neurons that may vary considerably in their semisaturation point without adaptation (colored curves), which yield a near-linear average CRF (black curve). *b*, Unit recordings have shown that contrast adaptation shifts the semisaturation point CRFs of neurons toward the adapting contrast. Therefore, adaptation may bring the CRFs of individual neurons into closer alignment, yielding a nonlinear CRF when aggregating across the population.

decrements. Observers practiced this task outside the scanner for several hours before the fMRI session. We used a staircasing procedure to adjust task difficulty (see Staircase procedure), which ensured that accuracy was equated between the attended (mean = 76.6%, SD = 2.4) and unattended (77.0%, SD = 2.0) conditions during the fMRI session.

To measure the full CRF in each attentional condition, we parametrically varied the contrast of the probe stimulus (0%-96%). One challenge for measuring CRFs with fMRI is that the BOLD response often appears to scale linearly with stimulus contrast (e.g., Boynton et al., 1999; Murray, 2008). This is rather unlike CRFs that are typical of individual neurons, which are characterized by a compressive nonlinearity, saturating at higher contrasts (Albrecht and Hamilton, 1982; Williford and Maunsell, 2006). To overcome this challenge, we took advantage of a contrast-adaptation procedure, which recent work from our laboratory has shown yields BOLD CRFs that show a compressive nonlinear response, more closely mirroring the known properties of CRFs measured at the level of individual neurons (Vinke et al., 2022). Vinke et al. (2022) proposed that linear BOLD CRFs measured without contrast adaptation may reflect the fact that the fMRI measurements aggregate across populations of neurons that vary widely in their semisaturation points (Fig. 3a). Unit recording studies have provided clear evidence that contrast adaptation modulates contrast gain, shifting the semisaturation point of neurons toward the adapting contrast (Ohzawa et al., 1982; Sclar et al., 1989). Thus, contrast adaptation may yield nonlinear CRFs measured with fMRI by bringing the CRFs of the neuronal population into closer alignment (Fig. 3b). Indeed, Vinke et al. (2022) found that, when observers were adapted to 16% contrast stimulus, BOLD CRFs in early visual cortex showed a compressive nonlinearity, saturating at higher contrasts. Following Vinke et al. (2022), we began each run of the task with a 60 s adaptation period, during which observers were adapted to the probe stimulus at 16% contrast (the adapting contrast), and the trials were interleaved with top-up adaptation periods to maintain the adaptation state (Fig. 2b, c) (compare Gardner et al., 2005; Vinke et al., 2022).

Spatial attention within the relevant hemifield modulates cortical responses
Spatial attention increases baseline BOLD activity in visual cortex (e.g., Tootell et al., 1998; Buracas and Boynton, 2007; Murray, 2008). Therefore, before examining the effects of feature-based
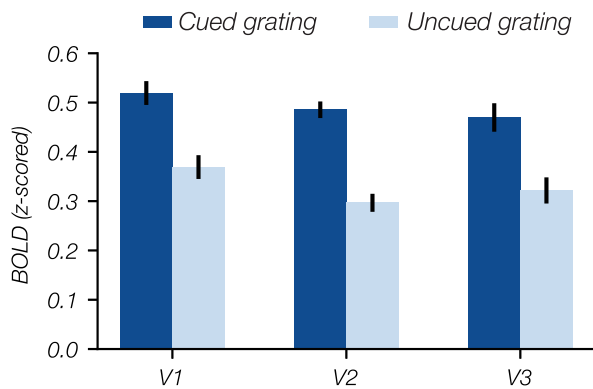
Figure 4. Spatial attention modulated cortical responses to stimuli in the relevant hemifield. Using pRF estimates, we identified voxels in V1-V3 that responded to each of the gratings (top and bottom) in the relevant hemifield. Bar plot represents the response measured 3-9 s after trial onset for voxels responsive to the cued grating (dark bars) and for voxels responsive to the uncued grating (light bars) for visual areas V1, V2, and V3. In all three areas, we observed greater activation of voxels selective for the cued grating than of voxels selective for the uncued grating. Error bars indicate ± 1 SEM across subjects.

attention, we first sought to verify that observers spatially attended to the cued grating in the relevant hemifield. To this end, we examined whether the BOLD response in the hemisphere contralateral to the relevant hemifield varied depending on which grating (upper or lower) was cued. Of the visually responsive voxels (see Voxel selection), we identified voxels whose pRF center fell within the upper or lower grating in the relevant hemifield, so that we could measure voxel responses when the grating they respond to was cued or uncued (Fig. 4). We found that voxel responses were larger when their preferred grating was cued than when it was uncued (two-sided paired t tests, V1: $t_{(11)} = 17.15$, $p < 0.001$; V2: $t_{(11)} = 12.83$, $p < 0.001$, V3: $t_{(11)} = 11.84$, $p < 0.001$), confirming that observers deployed spatial attention to the cued grating.

Feature-based attention multiplicatively scales probe-evoked responses
Having verified that observers attended the cued grating, we turned to our main question: how does feature-based attention modulate the BOLD CRF evoked by the probe stimulus? To measure CRFs in visual areas V1-V3, we deconvolved HRFs for stimulus-responsive voxels contralateral to the probe grating as a function of probe contrast and attention condition (Fig. 5a). Figure 5b shows CRFs for the attended and unattended conditions for each of the three visual areas. To characterize the effect of feature-based attention, we fit a Naka–Rushton function to the CRFs measured in each condition (Fig. 5b; see Statistical analysis). For each attention condition, we estimated four parameters of the Naka–Rushton function: a baseline parameter (b), which determines the additive offset of the function from zero; a response gain parameter (R_{max}), which determines how much the function rises above baseline; a contrast gain parameter (C_{50}), which measures the semisaturation point of the function, with changes in C_{50} producing horizontal shifts of the CRF; and a slope parameter (n), which determines how steeply the slope rises (compare Itthipuripat et al., 2019; Foster et al., 2021). Table 1 shows the Naka–Rushton parameter estimates for the attended and unattended conditions. We found no evidence for an additive baseline shift (a change in b). Instead, in all three visual areas, feature-based attention increased contrast gain, shifting the CRFs to the left (a decrease in C_{50}). In V3, we also found that

feature-based attention increased response gain (an increase in R_{max}). These results provide clear positive evidence that the BOLD signal measured with fMRI is sensitive to multiplicative changes in sensory gain because of feature-based attention.

It is important to note that the measured CRFs did not clearly saturate, which can make it difficult to differentiate between a change in contrast gain and a change in response gain. However, we found that CRFs in each attention condition converged at high contrasts, particularly in V1 and V2. This pattern cannot be explained by a pure change in response gain, which gives us confidence that feature-based attention truly modulates contrast gain. Importantly, despite any uncertainty about whether feature-based attention modulates contrast gain or responses gain, our data provide clear evidence for a multiplicative effect of attention on BOLD CRFs.

The effect of feature-based attention varies with voxel eccentricity preference
Recent work has found that CRF saturation varies with the eccentricity of voxel receptive fields, with CRFs saturating at lower contrasts (i.e., lower C_{50} values) for voxels with pRFs centered at inner eccentricities (i.e., closer to fovea) than for voxels that respond to the visual periphery (Vinke et al., 2022). Therefore, in an exploratory analysis, we examined whether the effect of feature-based attention varies with voxel eccentricity. We divided voxels in each visual area into three eccentricity bins (spanning 1.0°-2.5°, 2.5°-4.8°, and 4.8°-8.5°) that roughly equated the number of voxels in each bin for each visual area. Figure 6 shows CRFs for each visual area as a function of voxel eccentricity preference (for Naka–Rushton parameter estimates, see Table 2). Like Vinke et al. (2022), we found that C_{50} values gradually increased with eccentricity. The contrast-gain attentional effect that we observed at the whole-ROI level was consistent across eccentricity bins (Table 2). Moreover, we found that feature-based attention increased both contrast gain and response gain in V2 and V3 for the inner eccentricities, suggesting that the nature of attentional modulation varies with the eccentricity preference of neural populations. However, given the exploratory nature of this analysis, further work is necessary to test the robustness of this pattern.

Reconciling a contrast-gain effect with the normalization model of attention
At first glance, our finding that feature-based attention primarily increases contrast gain appears to be at odds with predictions from the normalization model of attention (NMA) (Reynolds and Heeger, 2009), a prominent computational model of attention. Whereas we found that feature-based attention primarily increases contrast gain, Herrmann et al. (2012) showed that the NMA predicts that feature-based attention selectively increases response gain, leaving contrast gain unchanged, and reported psychophysical data that support this prediction. To try and shed light on this discrepancy, we turned to simulations using the NMA to investigate whether our contrast gain effect can be reconciled with the model (Fig. 7a). For both of the simulations reported below, we used Reynolds and Heeger's (2009) MATLAB implementation of the NMA, which is available online (http://www.cns.nyu.edu/heegerlab/index.php?page=software&id=attentionModel). Table 3 summarizes the model parameters used in our simulations.

In our first simulation, we reproduced Herrmann et al.'s (2012) finding that feature-based attention selectively increases response gain (Fig. 7b). Following Herrmann et al. (2012), the stimulus input to the model was two orthogonal orientations, as
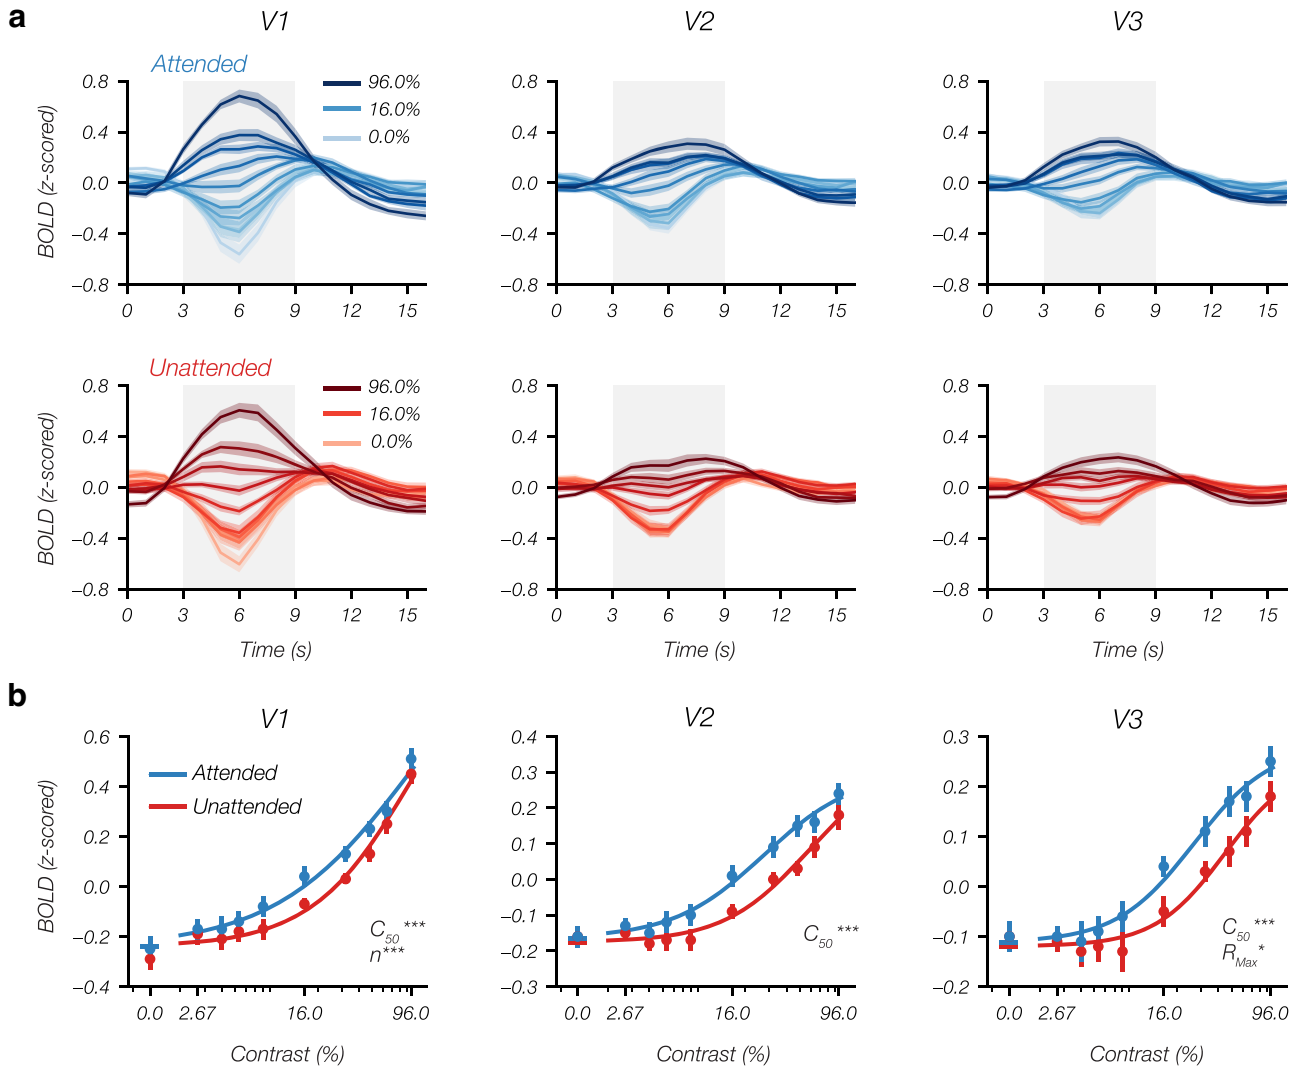
Figure 5. The effect of feature-based attention on probe-evoked CRFs. ***a***, Deconvolved HRF from stimulus-activated voxels contralateral to the probe stimulus in V1-V3 as a function of probe contrast in the attended (top) and unattended (bottom) conditions. To obtain CRFs, we measured the average response in each condition in a fixed window 3-9 s after stimulus onset (shaded gray area). ***b***, CRFs for the attended and unattended conditions for each ROI. The parameters that differed between conditions are shown in the bottom right corner of each plot. *$p < 0.05$. ***$p < 0.001$. See also Table 1. Error bars indicate $\pm$1 bootstrapped SEM across subjects.

Table 1. Naka–Rushton parameter estimates and resampling statistics results corresponding to the whole-ROI analysis shown in Figure 5b

ROI	R_{Max}	C_{50}	b	n
	Attended (mean $\pm$ bootstrapped SEM)			
	Unattended (mean $\pm$ bootstrapped SEM)			
	p value			
V1	0.72 ± 0.05	28.19 ± 3.04	-0.24 ± 0.05	0.87 ± 0.08
	0.68 ± 0.05	39.27 ± 4.12	-0.24 ± 0.04	1.18 ± 0.13
	$p = 0.240$	***$p < 0.001$	$p = 0.986$	***$p < 0.001$
V2	0.39 ± 0.04	21.53 ± 3.70	-0.16 ± 0.03	1.32 ± 0.29
	0.35 ± 0.04	36.29 ± 5.44	-0.18 ± 0.02	1.42 ± 0.27
	$p = 0.112$	***$p < 0.001$	$p = 0.465$	$p = 0.730$
V3	0.35 ± 0.03	22.80 ± 3.38	-0.11 ± 0.03	1.54 ± 0.32
	0.30 ± 0.03	35.01 ± 4.73	-0.12 ± 0.02	1.71 ± 0.31
	*$p = 0.025$	***$p < 0.001$	$p = 0.477$	$p = 0.566$

*$p < 0.05$. ***$p < 0.001$.

was the case in their psychophysical experiment, and we measured the response for a simulated neuron that was tuned for one of the two presented orientations (marked by a green cross) when that orientation was attended or unattended (for simulation parameters, see Table 3). Consistent with Herrmann et al.'s (2012) results, we found that feature-based attention selectively increased response gain under these conditions.

In our second simulation, we modeled the details of our own fMRI experiment (Fig. 7c). Our experiment differed from that of Herrmann et al. (2012) in two important ways. First, whereas Herrmann et al. (2012) cued observers to attend one of two orthogonal orientations, in our study we presented a single orientation in the probe hemifield and manipulated the attended orientation in the other hemifield. To simulate our paradigm, we provided a single orientation as stimulus input, and we varied whether attentional gain was centered on that orientation or on the orthogonal orientation. Second, Herrmann et al. (2012) measured the response of neurons tuned for one of the two presented orientations when that orientation was attended or unattended. This decision is reasonable when modeling behavioral performance, which may largely depend on the gain of neurons tuned for the target feature. However, responses measured with fMRI aggregate across many neurons, regardless of their orientation selectivity. Thus, in our second simulation, we measured the average activity across the entire orientation-selective
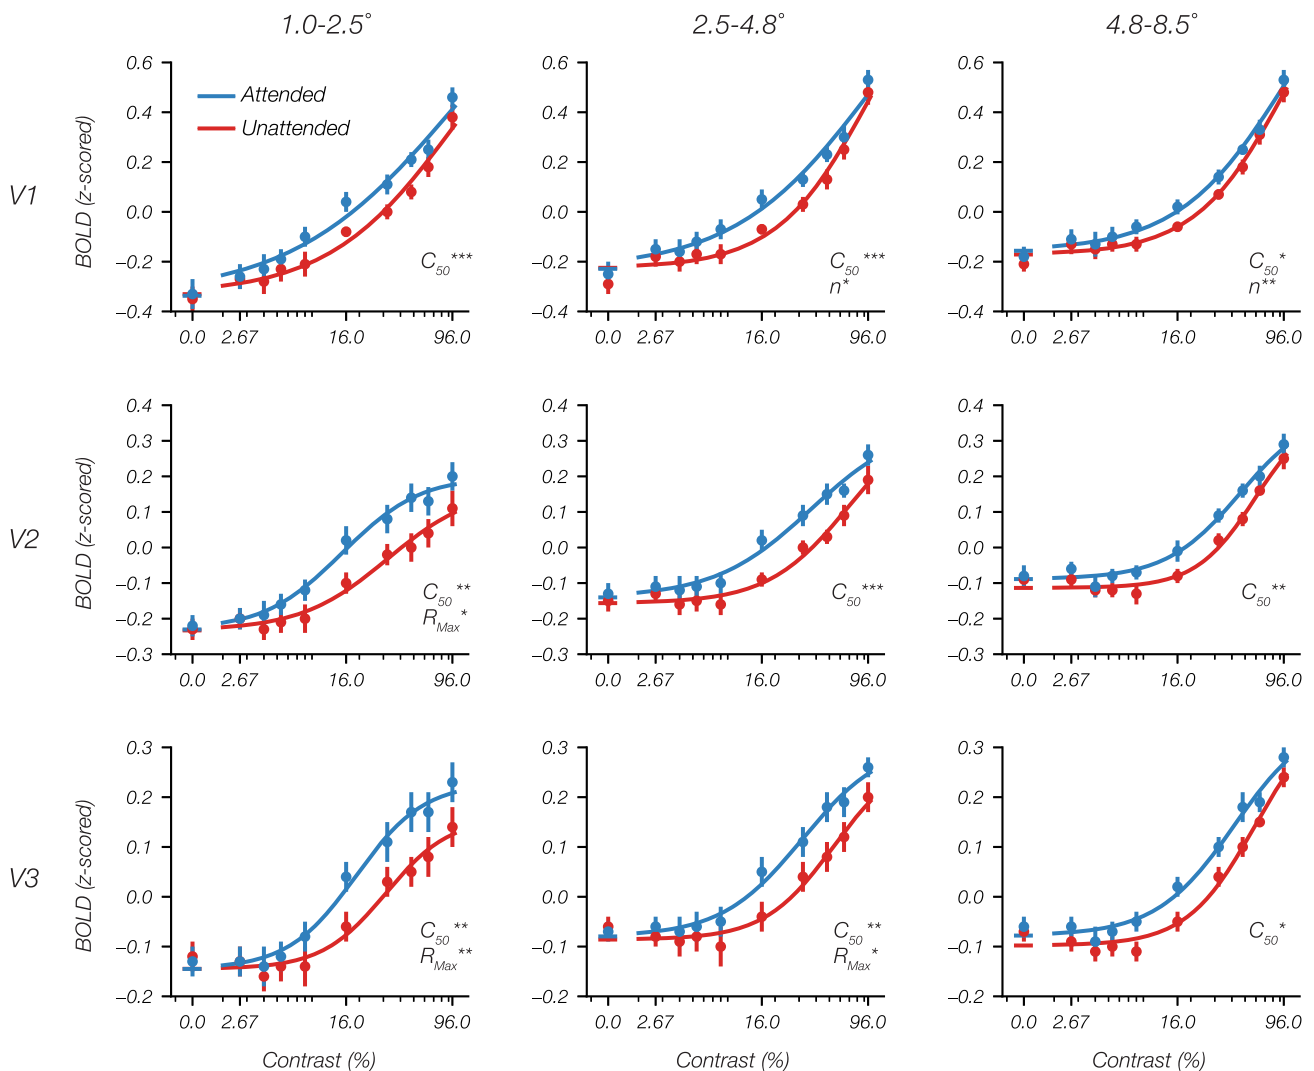
Figure 6. The effect of feature-based attention on probe-evoked CRFs as a function of voxel eccentricity preference. For each visual area (rows), we split voxels into bins based on eccentricity derived from pRF estimates (columns). The parameters that differed between conditions are shown in the bottom right corner of each plot. $*p < 0.05$. $**p < 0.01$. $***p < 0.001$. See also Table 2. Error bars indicate ± 1 bootstrapped SEM across subjects.

population. Apart from these two changes, the parameters of this simulation were otherwise identical to those used in the previous simulation (see Table 3). Under these conditions, we found that feature-based attention purely increased contrast gain (Fig. 7c).

It is striking that feature-based attention did not increase response gain of the population response in our second simulation. Past work has found that the effect of attention predicted by the NMA can vary substantially across neurons with different tuning preferences (e.g., Hara et al., 2014). Thus, we examined the effect of feature-based attention across neurons with different orientation preferences (Fig. 7d). We found that feature-based attention uniformly increased contrast gain across neurons with different orientation preferences, but only increased response gain for neurons whose preferred orientation was close to the probe orientation. For neurons tuned for other orientations, feature-based attention decreased response gain. This decrease in response gain reflects the fact that these neurons receive little stimulus drive, but are subject to greater suppressive drive (which was uniform across orientation) when attention is directed to the probe orientation.

Our simulations show that the contrast-gain effect that we observed is consistent with the NMA. However, we also found an increase in response gain in V2 and V3 that was most pronounced at inner eccentricities (Fig. 6). Can this finding be reconciled with the NMA? Our simulation finding that the effect of feature-based attention on response gain varies with the tuning preferences of simulated neurons (Fig. 7d) provides one possible explanation. Given this simulation result, an increase in response gain would be expected in in the population response if the probe orientation (horizontal) was overrepresented. In this case, neurons that show an increase in response gain will contribute more to the population average. It is well established that the cardinal orientations (vertical and horizontal) are indeed overrepresented in visual cortex, with a greater proportion of cells tuned for those orientations than for oblique orientations (Mansfield, 1974; Coppola et al., 1998; B. Li et al., 2003; Fang et al., 2022). Furthermore, neurophysiological and psychophysical evidence has shown that this cardinal bias is strongest near the fovea, declining with eccentricity (Mansfield, 1974; Berkley et al., 1975; Rovamo et al., 1982; X. Xu et al., 2007; but see Fang et al., 2022) and, relatedly, is stronger in neurons

Table 2. Naka–Rushton parameter estimates and resampling statistics results as a function of voxel eccentricity (corresponds to Fig. 6)

ROI	Eccentricity	Attended (mean ± bootstrapped SEM) Unattended (mean ± bootstrapped SEM) p value			
		R_{Max}	C_{50}	b	n
V1	1.0°-2.5°	0.76 ± 0.06	22.10 ± 2.81	−0.34 ± 0.06	0.73 ± 0.10
		0.68 ± 0.06	30.96 ± 4.18	−0.33 ± 0.05	0.94 ± 0.11
		p = 0.081	***p < 0.001	p = 0.759	p = 0.127
	2.5°-4.8°	0.71 ± 0.06	28.19 ± 3.46	−0.23 ± 0.05	0.87 ± 0.09
		0.69 ± 0.06	42.05 ± 5.28	−0.22 ± 0.04	1.27 ± 0.19
		p = 0.435	***p < 0.001	p = 0.913	***p < 0.001
	4.8°-8.5°	0.68 ± 0.06	36.87 ± 3.07	−0.16 ± 0.04	1.10 ± 0.09
		0.67 ± 0.05	43.18 ± 3.29	−0.17 ± 0.03	1.31 ± 0.15
		p = 0.611	*p = 0.011	p = 0.359	**p = 0.009
V2	1.0°-2.5°	0.41 ± 0.05	13.96 ± 2.46	−0.23 ± 0.03	1.53 ± 0.43
		0.33 ± 0.05	24.82 ± 6.02	−0.23 ± 0.03	1.36 ± 0.50
		*p = 0.037	**p = 0.001	p = 0.912	p = 0.722
	2.5°-4.8°	0.38 ± 0.04	25.24 ± 4.26	−0.14 ± 0.03	1.23 ± 0.41
		0.34 ± 0.04	41.07 ± 6.23	−0.16 ± 0.03	1.48 ± 0.37
		p = 0.146	***p < 0.001	p = 0.278	p = 0.419
	4.8°-8.5°	0.37 ± 0.04	34.58 ± 4.27	−0.09 ± 0.02	1.60 ± 0.27
		0.37 ± 0.04	45.50 ± 3.31	−0.11 ± 0.02	2.01 ± 0.35
		p = 0.985	**p = 0.001	p = 0.071	p = 0.268
V3	1.0°-2.5°	0.35 ± 0.04	17.80 ± 2.56	−0.14 ± 0.03	1.74 ± 0.42
		0.27 ± 0.04	27.88 ± 4.84	−0.14 ± 0.03	1.86 ± 0.51
		**p = 0. 006	**p = 0.002	p = 0.991	p = 0.731
	2.5°-4.8°	0.33 ± 0.03	25.49 ± 4.44	−0.08 ± 0.02	1.47 ± 0.29
		0.28 ± 0.03	39.48 ± 6.05	−0.09 ± 0.02	1.73 ± 0.43
		*p = 0.019	**p = 0.001	p = 0.559	p = 0.495
	4.8°-8.5°	0.35 ± 0.03	31.49 ± 4.48	−0.08 ± 0.02	1.52 ± 0.23
		0.34 ± 0.03	42.07 ± 4.92	−0.10 ± 0.02	1.76 ± 0.32
		p = 0.640	*p = 0.012	p = 0.070	p = 0.409

*p < 0.05. **p < 0.01. ***p < 0.001.

tuned for high spatial frequencies (B. Li et al., 2003; Shen et al., 2014). Thus, the response gain effect that we observed in our fMRI data, which was most pronounced at inner eccentricities, might reflect the well-established overrepresentation of cell tuned for the (horizontal) probe orientation that is known to be greater at inner eccentricities. Although the cardinal bias can parsimoniously explain the response gain effect that we observed, more work is needed to test this account. For example, this account would predict that the increase in response gain that we observed should only be seen when the probe orientation is one of the cardinal orientations. If the probe were an oblique orientation instead, then the probe orientation should be underrepresented, and no increase in response gain, or perhaps a decrease in response gain, should be observed.

Discussion

One primary way that attention improves neural coding is by increasing sensory gain, selectively amplifying stimulus-evoked neural responses (Hillyard et al., 1998; Carrasco, 2011). However, substantial evidence suggests that fMRI is not sensitive to these multiplicative effects of attention. fMRI studies that have examined the effect of spatial attention on CRFs have not observed the multiplicative effects that are seen in electrophysiological measures. Instead, attention produces an additive baseline shift in the BOLD signal (e.g., Buracas and Boynton, 2007; Murray, 2008; Sprague et al., 2018), calling into question what fMRI can tell us about how attention modulates sensory codes. In contrast with studies of

spatial attention, a seminal study found that feature-based attention amplified the BOLD response to a probe stimulus when it matched the attended feature (Saenz et al., 2002), an apparent multiplicative attention effect that is difficult to explain if fMRI is not sensitive to modulation of sensory gain. However, no previous study has examined the effect of feature-based attention on the full CRF to definitively characterize the nature of this effect. Here, we provided that critical test, measuring the effect of feature-based attention on CRFs in early visual areas V1-V3. We found robust multiplicative effects of attention in early visual cortex. In all three areas, feature-based attention increased contrast gain, improving neural sensitivity. In V2 and V3, we also observed an increase in response gain, primarily in voxels that respond to inner eccentricities, increasing neural responsivity. Together, our results show that BOLD CRFs are multiplicatively modulated by feature-based attention.

The role of contrast adaptation in characterizing attentional modulation of CRFs

Computational models of attention make specific predictions about how attention modulates CRFs (e.g., Reynolds and Heeger, 2009). A major obstacle to testing these predictions in humans is that CRFs measured with fMRI are typically linear (e.g., Boynton et al., 1999), unlike the CRFs of individual neurons, which are characterized by a compressive nonlinearity, saturating at higher contrasts (e.g., Albrecht and Hamilton, 1982). Linear CRFs pose a problem for testing computational models of attention for two reasons. First, models that are grounded in divisive normalization inherently predict nonlinear
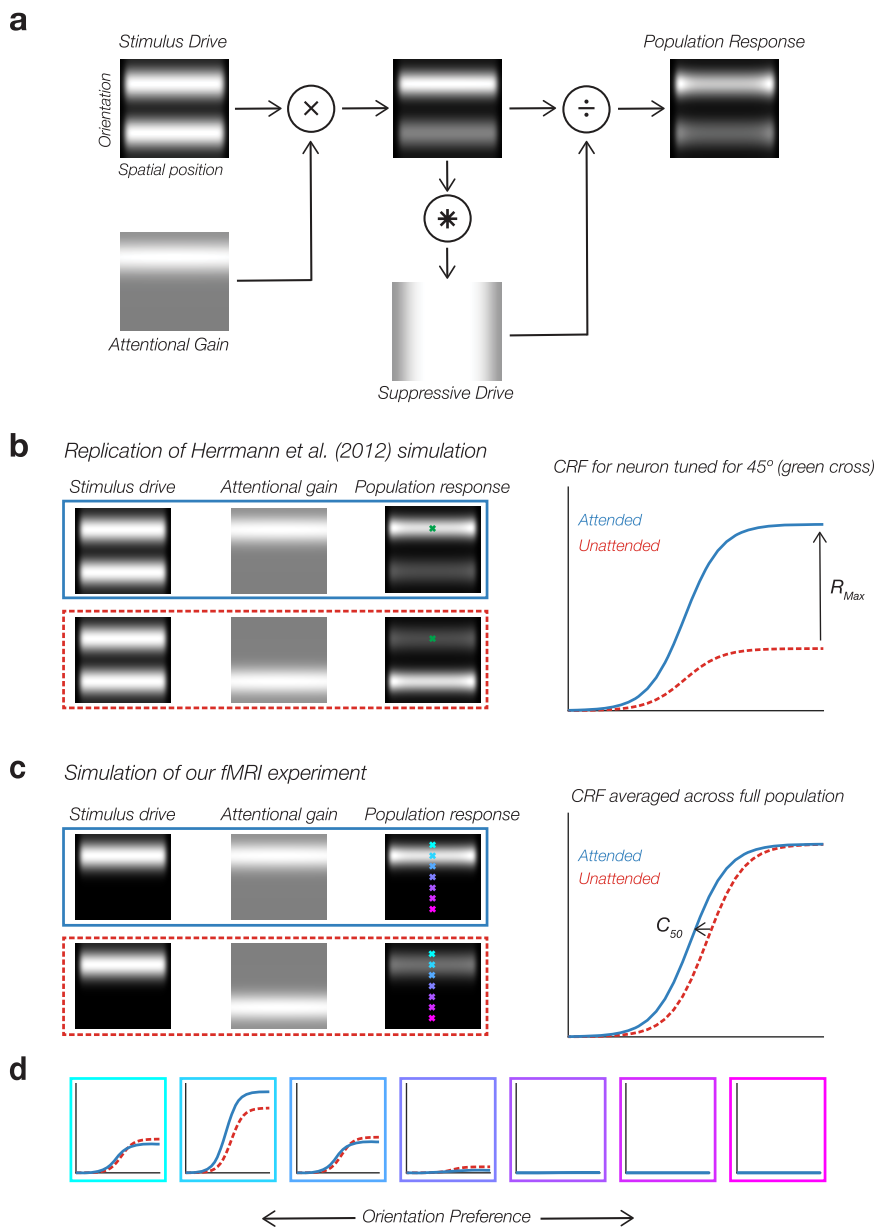
Figure 7. NMA and simulation results. ***a***, Schematic of the NMA (Reynolds and Heeger, 2009). Top left, The stimulus drive (brightness) as a function of the neuron's receptive field center (horizontal axis) and orientation preference (vertical axis). In this example, two spatially overlapping stimuli of orthogonal orientations were presented to the model. The stimulus drive is multiplied by an attentional gain (bottom left). The resulting product is convolved with a suppressive field to obtain suppressive drive. The final population response is calculated by dividing the product of stimulus drive and attentional gain by the suppressive drive. In our simulations, the suppressive field was Gaussian across space and uniform across orientation (for simulation parameters, see Table 3). ***b***, Replication of the Herrmann et al. (2012) simulation. We presented two orthogonal orientations to the model (stimulus drive), and varied which orientation was attended (attentional gain), measuring the response of a neuron (green cross) when its preferred orientation was attended (blue line) or unattended (red dashed line). Right, The predicted CRFs. Like Herrmann et al. (2012), we found that feature-based attention selectively increased response gain. ***c***, To model our own fMRI experiment, we presented a single probe orientation, and varied whether that orientation was attended or unattended. Because fMRI aggregates activity across neurons regardless of their orientation preference, we measured the average response across the entire orientation-selective population. Under these conditions, we found that the NMA predicts that feature-based attention selectively increases contrast gain. ***d***, Probing further, we found that the effect of feature-based attention in our second simulation (***c***) varied as a function of the orientation preference of individual neurons. All neurons showed an increase in contrast gain. However, only neurons tuned for the probe orientation showed increased response gain. Neurons tuned for other orientations showed a decrease in response gain.

CRFs (Reynolds and Heeger, 2009), and linear CRFs are at odds with this prediction. Second, differentiating between attentional mechanisms, such as response gain and contrast gain, only makes sense in the context of nonlinear CRFs. Fortunately, recent work suggests that the linear CRFs that are seen with fMRI are an artifact of the BOLD signal aggregating across neurons that vary widely in their semisaturation points (Fig. 3a). Studies have shown that contrast adaptation modulates contrast gain, shifting CRFs toward the adapting contrast (Ohzawa et al., 1982; Sclar et al., 1989; Gardner et al., 2005). Thus, contrast adaptation may enable measurement of nonlinear CRFs with fMRI by bringing the semisaturation points of individual neurons into closer alignment (Fig. 3b). Indeed, Vinke et al. (2022) found that contrast adaptation revealed compressive nonlinearities in population CRFs measured with fMRI.

Our study is the first to use this contrast-adaptation approach to examine how attention modulates CRFs. Like Vinke et al. (2022), our CRFs showed compressive nonlinearities, with the degree of saturation increasing across the visual hierarchy (greatest in V3), and greater for voxel tuned for inner eccentricities (Fig. 6). Although we did not see clear saturation in all visual areas or across all eccentricities, this is a big step forward given that saturation was altogether absent in previous studies (e.g., Boynton et al., 1999; Murray, 2008; Pestilli et al., 2011). There is still room to optimize this approach, and more work is needed to understand why contrast adaptation does not produce clear saturation in the periphery. One possibility is that longer adaptation periods are needed to see more uniform saturation across eccentricities and visual areas. Another possibility is that the stimulus must more precisely match the preferred spatiotemporal tuning (e.g., spatial frequency and temporal frequency tuning) at each eccentricity. Nevertheless, our study highlights the opportunity that the contrast-adaptation approach provides to better understand how attention modulates CRFs in human visual cortex.

One potential concern with the contrast-adaptation approach is that adaptation might change how attention modulates CRFs. Indeed, past work has shown that attention and contrast adaptation can interact. For instance, attending a high-contrast stimulus potentiates adaptation to that stimulus (Ling and Carrasco, 2006). We designed our study to

Table 3. NMA simulation model parameters[a]

	Simulation 1 (Fig. 7b)	Simulation 2 (Fig. 7c)
Orientation input(s)	−45° + 45°	−45° only
Attended orientation	−45° vs 45°	−45° vs 45°
Attentional field size		
Space	Uniform	Uniform
Orientation	30°	30°
Stimulation field size		
Space	5 au	5 au
Orientation	20°	20°
Suppressive field size		
Space	20 au	20 au
Orientation	Uniform	Uniform
Measured neuron(s)	Tuned for −45°	Population average

[a]The attentional field, stimulation field, and suppressive field were modeled as Gaussians, with size specified as their SD; au, arbitrary units.

prevent such interactions between attention and contrast adaptation. Most importantly, we interleaved our attention conditions within each run of the task to ensure that the adaptation state was matched between two conditions. The visual system constantly adapts to current visual input. Thus, in our task (Fig. 2), adaptation not only occurs during the initial 60 s adaptation period and the top-up adaptation periods between trials, but also during the trials themselves. The purpose of the top-up adaptation periods was to readapt observes to the adapting stimulus, but if attention conditions were blocked, then attention-induced amplification of neural responses during trials could lead to accumulated differences in the adaptation state between blocks. By interleaving and randomizing the order of conditions within a block, we ensured that the adaptation state was matched across conditions. Thus, we can be confident that differences between conditions solely reflect attentional modulation.

Why have fMRI studies of spatial attention not found multiplicative effects?
Given the central role of fMRI in studying the neural mechanisms of attention, it is essential that we understand whether fMRI-BOLD is sensitive to multiplicative changes in sensory gain. Our results provide clear, positive evidence that fMRI is sensitive to these effects. However, our results stand in stark contrast with studies of spatial attention that have repeatedly found additive baseline shifts rather than multiplicative effects (e.g., Buracas and Boynton, 2007; Murray, 2008; Sprague et al., 2018; Itthipuripat et al., 2019). If the BOLD signal is sensitive to multiplicative changes in sensory gain, then why are these effects not seen in studies of spatial attention?

In a recent study, Liu et al. (2021) suggested that additive effects because of spatial attention may reflect nonlinearities of the intravascular response of large draining veins, which dominate the BOLD response measured with commonly used EPI sequences at 3T (Zhong et al., 1998). In an initial experiment using such a sequence, Liu et al. (2021) found that spatial attention produced an additive effect on the BOLD response, consistent with previous work. However, in a second experiment, Liu et al. (2021), used a balanced steady-state free precision sequence that provided greater spatial resolution (see Lee et al., 2008). In this experiment, they found that attention had an additive effect for voxels that fell within large draining veins, but a multiplicative effect for other voxels. Indeed, Liu et al. showed in subsequent modeling that known nonlinearities in the intravascular response of large vessels can produce additive effects given a multiplicative effect in neuronal activity. Thus, Liu et al. argued that nonlinearities

in the intravascular response that are unrelated to neuronal activity may explain the additive effect that has been reported in studies of spatial attention. Our results, however, suggest that this explanation cannot be the whole story. We found that feature-based attention multiplicatively scaled CRFs even though we used an EPI sequence at 3T that was comparable to past studies of spatial attention that have reported additive effects (e.g., Buracas and Boynton, 2007; Murray, 2008). Thus, while CRFs measured with commonly used EPI sequences may be distorted to some degree by nonlinearities in intravascular responses, these nonlinearities do not appear to explain the discrepancy between the multiplicative effects that we observed because of feature-based attention and the additive effects that have been reported in studies of spatial attention.

Work by Hara et al. (2014) provides another possible explanation for this discrepancy. Using the NMA, Hara et al. (2014) simulated the effect of spatial attention on population responses and found that the NMA predicts that the effect of spatial attention on the CRFs of simulated neurons depends on their tuning properties: spatial attention increased response gain for neurons that were well tuned for the feature value of the stimulus, but instead increased contrast gain for neurons tuned for other feature values. Hara et al. (2014) found that the net effect of these modulations on the population CRF was a near-additive effect. Thus, the additive effects of spatial attention reported in past fMRI studies may reflect the fact that the responses measured in these studies aggregated across neurons with different tuning preferences, rather than fMRI being inherently insensitive to multiplicative effects of attention. By contrast, in our simulations of feature-based attention, we found that the NMA predicts that feature-based attention, unlike spatial attention, will produce a multiplicative contrast-gain effect in the population response. Thus, feature-based attention may provide a clearer testbed than spatial attention for examining whether fMRI is sensitive to modulation of sensory gain.

Although Hara et al.'s (2014) simulations offer a plausible explanation for why fMRI studies of spatial attention have not detected multiplicative effects of attention, this account is yet to be empirically tested. Fortunately, recent advances in computational neuroimaging techniques provide an opportunity to test this account. Inverted encoding models (IEMs), for instance, allow researchers to reconstruct the activity of feature-selective channels (or neuronal populations) from voxelwise patterns of activity in visual cortex (Brouwer and Heeger, 2009, 2011). The IEM approach has proven to be a powerful tool for characterizing how attention modulates feature-selective population responses (e.g., Garcia et al., 2013; Sprague and Serences, 2013). Furthermore, this approach has been used to measure the effect of sensory manipulations (e.g., cross-orientation suppression) on the CRF of different orientation-selective channels (Brouwer and Heeger, 2011). Thus, the IEM approach provides an opportunity to test Hara et al.'s (2014) account that spatial attention increases response gain in the channel tuned for the stimulus value and increases contrast gain for channels tuned for other feature values.

Testing Hara et al.'s (2014) account is especially important considering a recent study by Itthipuripat et al. (2019), who examined the effect of spatial attention on CRFs measured with both fMRI and EEG in the same human observers performing the same attention task. Consistent with previous work, they found that spatial attention produced a near-additive baseline shift in the fMRI-BOLD response, but increased response gain of EEG responses (e.g., the P1 component and steady-state visually evoked potentials), providing clear evidence that fMRI and EEG

measurements are sensitive to different neural signals that are modulated by attention in different ways. Although Hara et al.'s (2014) account can explain the additive effect of attention measured with fMRI, Itthipuripat et al. (2019) pointed out that this account is difficult to reconcile with the EEG results because EEG, like fMRI, measures population-level activity, and would also be expected to show an additive shift. That being said, it is also possible that EEG may be particularly sensitive to neural signals that are primarily sensitive to changes in response gain, and much more work is needed to further understand why EEG and fMRI exhibit qualitatively distinct attentional effects.

Conclusions

In conclusion, attention improves sensory processing by increasing sensory gain, amplifying sensory responses evoked by attended stimuli. However, available evidence has suggested that the fMRI-BOLD signal is not sensitive to these multiplicative effects of attention. In the current study, we found that feature-based attention multiplicative scales CRFs in early visual cortex measured with fMRI, primarily increasing contrast gain of BOLD CRFs, an effect that is broadly consistent with predictions of how feature-based attention should modulate population neural responses derived from the NMA. Our results provide clear, positive evidence that fMRI is sensitive to multiplicative modulations of sensory gain. Given the central role of fMRI in the study of visual attention, an important direction for future work will be understanding the boundary conditions under which such multiplicative effects are observed.

References

Albrecht DG, Hamilton DB (1982) Striate cortex of monkey and cat: contrast response function. J Neurophysiol 48:217–237.

Andersson JL, Skare S, Ashburner J (2003) How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. Neuroimage 20:870–888.

Benson NC, Jamison KW, Arcaro MJ, Vu AT, Glasser MF, Coalson TS, Van Essen DC, Yacoub E, Ugurbil K, Winawer J, Kay K (2018) The Human Connectome Project 7 Tesla retinotopy dataset: description and population receptive field analysis. J Vis 18:23.

Berkley MA, Kitterle F, Watkins DW (1975) Grating visibility as a function of orientation and retinal eccentricity. Vision Res 15:239–244.

Boynton GM, Demb JB, Glover GH, Heeger DJ (1999) Neuronal basis of contrast discrimination. Vision Res 39:257–269.

Brainard DH (1997) The Psychophysics Toolbox. Spat Vis 10:433–436.

Brouwer GJ, Heeger DJ (2009) Decoding and reconstructing color from responses in human visual cortex. J Neurosci 29:13992–14003.

Brouwer GJ, Heeger DJ (2011) Cross-orientation suppression in human visual cortex. J Neurophysiol 106:2108–2119.

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

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

Cauley SF, Polimeni JR, Bhat H, Wald LL, Setsompop K (2014) Interslice leakage artifact reduction technique for simultaneous multislice acquisitions. Magn Reson Med 72:93–102.

Chelazzi L, Miller EK, Duncan J, Desimone R (1993) A neural basis for visual search in inferior temporal cortex. Nature 363:345–347.

Coppola DM, White LE, Fitzpatrick D, Purves D (1998) Unequal representation of cardinal and oblique contours in ferret visual cortex. Proc Natl Acad Sci USA 95:2621–2623.

Dale AM (1999) Optimal experimental design for event-related fMRI. Hum Brain Mapp 8:109–114.

Duncan J, Ward R, Shapiro K (1994) Direct measurement of attentional dwell time in human vision. Nature 369:313–315.

Fang C, Cai X, Lu HD (2022) Orientation anisotropies in macaque visual areas. Proc Natl Acad Sci USA 119:e2113407119.

Feinberg DA, Moeller S, Smith SM, Auerbach E, Ramanna S, Gunther M, Glasser MF, Miller KL, Ugurbil K, Yacoub E (2010) Multiplexed echo planar imaging for sub-second whole brain fMRI and fast diffusion imaging. PLoS One 5:e15710.

Fischl B (2012) FreeSurfer. Neuroimage 62:774–781.

Foster JJ, Thyer W, Wennberg JW, Awh E (2021) Covert attention increases the gain of stimulus-evoked population codes. J Neurosci 41:1802–1815.

Garcia JO, Srinivasan R, Serences JT (2013) Near-real-time feature-selective modulations in human cortex. Curr Biol 23:515–522.

Gardner JL, Sun P, Waggoner RA, Ueno K, Tanaka K, Cheng K (2005) Contrast adaptation and representation in human early visual cortex. Neuron 47:607–620.

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

Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, Haase A (2002) Generalized autocalibrating partially parallel acquisitions (GRAPPA). Magn Reson Med 47:1202–1210.

Hara Y, Pestilli F, Gardner JL (2014) Differing effects of attention in single-units and populations are well predicted by heterogeneous tuning and the normalization model of attention. Front Comput Neurosci 8:12.

Herrmann K, Heeger DJ, Carrasco M (2012) Feature-based attention enhances performance by increasing response gain. Vision Res 74:10–20.

Hesterberg T (2011) Bootstrap. WIREs Comp Stat 3:497–526.

Hillyard SA, Vogel EK, Luck SJ (1998) Sensory gain control (amplification) as a mechanism of selective attention: electrophysiological and neuroimaging evidence. Philos Trans R Soc Lond B Biol Sci 353:1257–1270.

Itthipuripat S, Ester EF, Deering S, Serences JT (2014a) Sensory gain outperforms efficient readout mechanisms in predicting attention-related improvements in behavior. J Neurosci 34:13384–13398.

Itthipuripat S, Garcia JO, Rungratsameetaweemana N, Sprague TC, Serences JT (2014b) Changing the spatial scope of attention alters patterns of neural gain in human cortex. J Neurosci 34:112–123.

Itthipuripat S, Sprague TC, Serences JT (2019) Functional MRI and EEG index complementary attentional modulations. J Neurosci 39:6162–6179.

Kay KN, Winawer J, Mezer A, Wandell BA (2013) Compressive spatial summation in human visual cortex. J Neurophysiol 110:481–494.

Kim YJ, Grabowecky M, Paller KA, Muthu K, Suzuki S (2007) Attention induces synchronization-based response gain in steady-state visual evoked potentials. Nat Neurosci 10:117–125.

Lee JH, Dumoulin SO, Saritas EU, Glover GH, Wandell BA, Nishimura DG, Pauly JM (2008) Full-brain coverage and high-resolution imaging capabilities of passband b-SSFP fMRI at 3T. Magn Reson Med 59:1099–1110.

Li B, Peterson MR, Freeman RD (2003) Oblique effect: a neural basis in the visual cortex. J Neurophysiol 90:204–217.

Li X, Lu ZL, Tjan BS, Dosher BA, Chu W (2008) Blood oxygenation level-dependent contrast response functions identify mechanisms of covert attention in early visual areas. Proc Natl Acad Sci USA 105:6202–6207.

Ling S, Carrasco M (2006) When sustained attention impairs perception. Nat Neurosci 9:1243–1245.

Liu C, Guo F, Qian C, Zhang Z, Sun K, Wang DJ, He S, Zhang P (2021) Layer-dependent multiplicative effects of spatial attention on contrast responses in human early visual cortex. Prog Neurobiol 207:101897.

Mansfield RJ (1974) Neural basis of orientation perception in primate vision. Science 186:1133–1135.

Martínez-Trujillo J, Treue S (2002) Attentional modulation strength in cortical area MT depends on stimulus contrast. Neuron 35:365–370.

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.

Moeller S, Yacoub E, Olman CA, Auerbach E, Strupp J, Harel N, Uğurbil K (2010) Multiband multislice GE-EPI at 7 Tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. Magn Reson Med 63:1144–1153.

Morgan ST, Hansen JC, Hillyard SA (1996) Selective attention to stimulus location modulates the steady-state visual evoked potential. Proc Natl Acad Sci USA 93:4770–4774.

Murray SO (2008) The effects of spatial attention in early human visual cortex are stimulus independent. J Vis 8:2.

Ohzawa I, Sclar G, Freeman RD (1982) Contrast gain control in the cat visual cortex. Nature 298:266–268.

Pelli DG (1997) The VideoToolbox software for psychophysics: transforming numbers into movies. Spat Vis 10:437–442.

Pestilli F, Carrasco M, Heeger DJ, Gardner JL (2011) Attentional enhancement via selection and pooling of early sensory responses in human visual cortex. Neuron 72:832–846.

Raymond JE, Shapiro KL, Arnell KM (1992) Temporary suppression of visual processing in an RSVP task: an attentional blink? J Exp Psychol Hum Percept Perform 18:849–860.

Reuter M, Rosas HD, Fischl B (2010) Highly accurate inverse consistent registration: a robust approach. Neuroimage 53:1181–1196.

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

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

Rovamo J, Virsu V, Laurinen P, Hyvärinen L (1982) Resolution of gratings oriented along and across meridians in peripheral vision. Invest Ophthalmol Vis Sci 23:666–670.

Saenz M, Buracas GT, Boynton GM (2002) Global effects of feature-based attention in human visual cortex. Nat Neurosci 5:631–632.

Sclar G, Lennie P, DePriest DD (1989) Contrast adaptation in striate cortex of macaque. Vision Res 29:747–755.

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

Setsompop K, Gagoski BA, Polimeni JR, Witzel T, Wedeen VJ, Wald LL (2012) Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. Magn Reson Med 67:1210–1224.

Shen G, Tao X, Zhang B, Smith EL 3rd, Chino YM (2014) Oblique effect in visual area 2 of macaque monkeys. J Vis 14:3.

Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TE, Johansen-Berg H, Bannister PR, De Luca M, Drobnjak I, Flitney DE, Niazy RK, Saunders J, Vickers J, Zhang Y, De Stefano N, Brady JM, Matthews PM (2004) Advances in functional and structural MR image analysis and implementation as FSL. Neuroimage 23:S208–S219.

Somers DC, McMains SA (2005) Spatially-specific attentional modulation revealed by fMRI. In: Neurobiology of attention, pp 377–382. San Diego: Academic.

Sprague TC, Itthipuripat S, Vo V, Serences JT (2018) Dissociable signatures of visual salience and behavioral relevance across attentional priority maps in human cortex. J Neurophysiol 119:2153–2126.

Sprague TC, Serences JT (2013) Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. Nat Neurosci 16:1879–1887.

Stokes M, Thompson R, Nobre AC, Duncan J (2009) Shape-specific preparatory activity mediates attention to targets in human visual cortex. Proc Natl Acad Sci USA 106:19569–19574.

Tootell RB, Hadjikhani N, Hall EK, Marrett S, Vanduffel W, Vaughan JT, Dale AM (1998) The retinotopy of visual spatial attention. Neuron 21:1409–1422.

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

van der Kouwe AJ, Benner T, Salat DH, Fischl B (2008) Brain morphometry with multiecho MPRAGE. Neuroimage 40:559–569.

van Voorhis S, Hillyard SA (1977) Visual evoked potentials and selective attention to points in space. Percept Psychophys 22:54–62.

Vinke LN, Bloem IM, Ling S (2022) Saturating nonlinearities of contrast response in human visual cortex. J Neurosci 42:1292–1302.

Virtanen P, et al. SciPy 1.0 Contributors (2020) SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat Methods 17:261–272.

Williford T, Maunsell JH (2006) Effects of spatial attention on contrast response functions in Macaque area V4. J Neurophysiol 96:40–54.

Xu J, Moeller S, Auerbach EJ, Strupp J, Smith SM, Feinberg DA, Yacoub E, Uğurbil K (2013) Evaluation of slice accelerations using multiband echo planar imaging at 3T. Neuroimage 83:991–1001.

Xu X, Anderson TJ, Casagrande VA (2007) How do functional maps in primary visual cortex vary with eccentricity? J Comp Neurol 501:741–755.

Yeo BT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR, Fischl B, Liu H, Buckner RL (2011) The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J Neurophysiol 106:1125–1165.

Zhong J, Kennan RP, Fulbright RK, Gore JC (1998) Quantification of intravascular and extravascular contributions to BOLD effects induced by alteration in oxygenation or intravascular contrast agents. Magn Reson Med 40:526–536.