Supplemental Data

Attention narrows position tuning of population responses in V1

Jason Fischer[1,2] & David Whitney[1,2]

[1]Center for Mind and Brain
[2]Department of Psychology
University of California, Davis
Davis, CA 95616

Supplemental Experimental Procedures

Stimuli

The stimuli in the main experiment consisted of four Gabors (sine wave grating within a Gaussian contrast envelope), one presented in each quadrant of the visual field around a central fixation point. The Gabors appeared at five possible eccentricities: 8.430, 8.735, 9.039, 9.343, and 9.647 degrees from fixation. The positions of the Gabors were manipulated by applying a skew to the Gaussian contrast envelopes [1]. A central condition had no skew applied to the contrast envelope; its size (standard deviation) was 1.66 degrees and its centroid was at 9.039 degrees eccentricity. In two more foveal conditions the contrast envelopes were skewed by 0.19 and 0.38 degrees toward fixation, and in two more eccentric conditions the contrast envelopes were skewed by 0.19 and 0.38 degrees away from fixation. Eccentricities were manipulated in this way rather than by shifting the peak contrast because skewing the contrast envelope is a better method of isolating the mechanisms that perform centroid analysis [1]. A sixth condition consisted of a fixation baseline in which only the fixation point was present.

In all conditions the spatial frequency of the Gabors was 0.38 cycles/degree, peak contrast was 87% (Michelson), and each Gabor was flickered in counterphase at 7.5 Hz. The phase of each Gabor was randomized on every trial. During scanning, the six conditions were randomly interleaved in 36 ten-second blocks per run. Subjects maintained fixation at a central point (0.39 deg diameter) throughout the entire experiment.

In a control experiment, we presented windowed noise rather than Gabors. We first created random noise with a Michelson contrast of 87%, with each light or dark element in the noise occupying a 4 x 4 block of pixels. We then rotated the image to a random orientation (to avoid peaks in the power spectrum at any given orientation) and added skewed Gaussian contrast envelopes just as with the Gabors in the main experiment. During presentation, we updated the noise and its orientation at 10Hz, so that subjects saw dynamic random noise rather than flickering Gabors. In all respects besides this difference in stimuli, the experimental design was identical for the control experiment and the main experiment.

Attention task

We presented stimuli in a 10-second blocked design in order to test the influence of sustained attention on the distribution of BOLD response in the cortex. This blocked presentation also allowed for improved SNR in the measured BOLD response versus an

event-related design. In half of the functional runs, referred to here as *attended* runs, we ensured that subjects attended to the location of the Gabors for the entire duration of the trial using a sustained attention task [2]. In the remaining functional runs, which we refer to as *unattended* runs, subjects' attention was held at the fixation point, and they ignored the Gabors in the periphery.

The same task-related stimuli were present in attended and unattended trials. At a randomly chosen time during the first 8 seconds of each ten second block, a small texture pattern (either circular or radial grating, chosen randomly; 1.09 degrees diameter) appeared for 500 ms superimposed on one of the 4 Gabors (chosen randomly), at an eccentricity of 9.04 degrees (center to center) from fixation. The flashed texture appeared during every condition, including the fixation baseline. During the last two seconds of each ten second block a second flashed texture appeared at the center of a randomly chosen Gabor for 500 ms, and a white annulus (0.98 degrees diameter) was presented continuously around the fixation point, indicating that subjects should make a response. The second textured flash matched the first one with a probability of 50%.

In attended runs, we instructed subjects to maintain fixation at all times and make two judgments. First, subjects discriminated the eccentricity of the Gabors (the degree of skew in their envelopes) in a 5 alternative-forced-choice task. Subjects also needed to continuously monitor the 4 Gabors for the texture flashes. At the end of each trial, when the white annulus cue was presented around the fixation point, subjects indicated both the eccentricity of the Gabors and whether or not the two presented textures matched. This task ensured that subjects attended to the surrounding Gabors for the entire duration of each trial.

Additionally, a small grating annulus (either circular or radial grating, chosen randomly; 0.98 degrees diameter) appeared for 500 ms surrounding the fixation point at least 7 times but no more than 11 times per trial, and one texture always appeared once more than the other. These flashes were uncorrelated with the textured flash superimposed on the Gabors and were therefore uninformative. During attended runs, we instructed subjects to ignore the flashes at the fixation point, and to attend entirely at the positions of the Gabors. During runs in which the Gabors were unattended, subjects ignored the eccentricity of the Gabors and the appearance of the textures, and instead counted how many times each grating annulus appeared around the fixation point. When the white annulus cue was presented at the end of each trial, subjects made a 2AFC judgment of which grating had appeared more often. Thus, during attended and unattended runs (referring to whether or not the Gabor stimuli were attended to), demanding tasks held subjects' attention at the location of the Gabors and at the fixation point, respectively. Attended and unattended runs were identical in terms of the stimuli presented; the only difference between the two was in the judgments that subjects were instructed to make.

fMRI data collection and analysis

Seven subjects participated in the main experiment, and two subjects participated in windowed noise control experiment. Scanning protocols were approved by the University of California, Davis, Human Subject Review Board. Imaging was conducted on a 3-Tesla Siemens TRIO scanner located at the UC Davis Imaging Research Center. Each participant's head was rested in a Siemens eight-channel phased-array head coil. Stimuli were back-projected with a Digital Projection

Mercury 5000HD projector (75 Hz) onto a semi-transparent screen from outside the bore. A mirror angled at 45 deg, located 9.5 cm directly above the subject, provided a reflected view of the stimuli. Functional images were acquired with a gradient-recalled echo EPI sequence. Whole-brain structural images were collected with a high resolution (1 mm^3) Turbo Spin Echo scan that was used to align functional images. The acquisition parameters were: TR = 2000 ms, TE = 26 ms, FA = 90 deg, FOV = 22 x 22 cm^2, voxel size = 1.528 x 1.528 x 2.5 mm^3, 20 slices per volume. The imaging volume was parallel to and centered on the calcarine sulcus, covering the occipital lobe.

All fMRI data processing, including linear trend removal, 3D motion correction, and GLM analyses, was conducted with Brain Voyager QX (Brain Innovation B.V., Maastricht, The Netherlands). The images were not spatially smoothed. A correction for serial correlations (removal of first-order autocorrelation) was used prior to all GLM analyses. Each subject's high resolution anatomical image was transformed to Talaraich coordinates [3], and the data for each functional run was individually aligned to the subject's Talairach-transformed anatomical image. Performing individual alignments for each functional run, within each subject, mitigated any effect of subject movement between functional runs.

To generate an individual map of BOLD response for each eccentricity condition, we fit a GLM to the data with five predictors (corresponding to the five Gabor positions). Each predictor consisted of a boxcar function representing the blocked presentations of the condition, convolved with a canonical HRF. We separately contrasted each of these five eccentricity conditions against a sixth predictor (the fixation baseline) to generate a unique map of BOLD response for each stimulus condition. We performed this process separately for all subjects, creating distinct maps for attended and unattended runs.

ROI definition

We defined the regions of interest (ROIs) for areas V1 through V4 by tracing reversals in each subject's retinotopic mapping in the cortex [4]. In localizer runs, which were presented separately from the main experimental runs, we presented flickering bowtie stimuli to identify the borders of visual areas V1 through V4. The bowties consisted of radial sine wave patterns of 11.79 deg radius, subtending an arc of 8.16 deg. The bowties flickered in counterphase at 7.5 Hz. There were three conditions in these retinotopy runs; in two conditions, the bowties were centered on the vertical or horizontal meridians, and the third condition was a fixation baseline in which only the fixation point was present. Conditions were randomized in 36 ten-second blocks. We performed general linear model (GLM) analyses on the data in the retinotopy runs to define visual areas V1 through V4. Bowties along the vertical meridians were contrasted with those on the horizontal meridians, which yielded a striated map of activity across early visual areas; we defined the boundaries of each visual area V1 through V4 using the mirror reversals in the representations of the horizontal and vertical meridians, as in Sereno et al. [4].

Correlation analysis

To obtain estimates of the precision of retinotopic coding within the V1 through V4 regions of interest, we produced a position discrimination plot for each ROI. To produce each position discrimination plot, we began with the five maps of BOLD response, corresponding to the five different stimulus eccentricities. Given any two maps, we found the correlation between them within the ROI

by plotting t values from one map against the corresponding t values from the other map, on a voxel-by-voxel basis. The resulting plot had as many points as there were voxels in the ROI. We computed a Pearson r value for every such voxel plot, and converted the r values to Fisher z scores so that they could be linearly compared with each other. We produced a voxel plot and a corresponding correlation for every pairing of activity maps, resulting in ten z scores. To create a final position discrimination plot, we plotted each of the ten z scores against the spatial separation (in deg. visual angle) of the two stimuli that produced that particular correlation (smaller separations typically implied larger z scores, and vice versa; [2, 5]). Adjacent conditions had centroids separated by .304 degrees; more distant conditions were separated by multiples of this value (.609, .912, and 1.216 degrees). We plotted subjects together (Fig. 2), and fit a regression line to the position discrimination plot in order to provide an index of the precision of retinotopic coding within the specific ROI that was used to create the plot. The linear regression included an additional random effects term to account for plotting multiple subjects together. This model took the form $z_{ijk} = \beta_0 + \tau_i + \beta_1 x_{jk} + \varepsilon_{ijk}$, where *i* indexed the subjects and the pair (*j*, *k*) indexed the ten stimulus pairings; τ_i provided for rigid vertical shifts between subjects that did not affect the slope. The more negative the slope of the regression line, the better the position discrimination. We produced position discrimination plots for attended and unattended conditions separately in order to allow a comparison of the precision of spatial coding with and without attention. To test whether this difference in slopes was significant, we computed Z statistic on the difference of r values between the attention and fixation plots, given by $Z = \frac{r_1' - r_2'}{\sqrt{\frac{1}{N_1 - 3} + \frac{1}{N_2 - 3}}}$, where r_1' and r_2' are the Fisher-transformed Pearson r values from the attended and unattended plots, and N_1 and N_2 are the number of points in each plot respectively.

This correlational process is based on a simple concept: in a cortical area which encodes object position, the activity resulting from two spatially adjacent stimuli will be highly correlated, while the activity resulting from more distant stimuli will be less strongly correlated. Thus a z score plot constructed as described above will have a negative slope if BOLD response strongly discriminates stimulus position, and a slope near zero if position information is not encoded in the voxels within the ROI for which the plot was constructed. The slope of such a plot indexes the precision of position coding because it indicates the minimum distance between stimuli required to produce significantly different BOLD response patterns [5].

Slice-based analysis

To directly examine the profile of the BOLD response, we measured activity in V1 along a slice across each subject's inflated cortex. We created inflated meshes from subjects' high-resolution anatomical scans, and then mapped a GLM contrast of the most eccentric stimulus condition versus the most foveal condition onto the inflated meshes to identify the retinotopic location of BOLD response in V1 and to define the direction of retinotopic progression for each subject (Supplemental Fig. 4A). We then drew a slice on each subject's right hemisphere surface map, through the region of peak activation in V1 and along the calcarine sulcus, parallel to the direction of retinotopic progression of the stimuli. We measured the BOLD response along each slice separately for attended and unattended conditions, and plotted BOLD

amplitude for each as a function of cortical distance (millimeters in Talairach space) along the slices.

To characterize the shape of the BOLD response, we fit a difference-of-Gaussians (DoG) function to each subject's slice data (Supplemental Fig. 4B shows four subjects). The DoG function had the form $y = a_1 e^{-\frac{(x-b_1)^2}{2c_1^2}} + a_2 e^{-\frac{(x-b_2)^2}{2c_2^2}}$, where a_1 was constrained to be positive and a_2 was constrained to be negative, so that the first and second halves of the equation represent the positive and negative components of the curve, respectively. We were then able to compare the height, width, and peak positions of the positive lobes of the attended and unattended curves by comparing the a_1, c_1, and b_1 parameters respectively. We performed such comparisons using three-way ANOVAs (stimulus position x attention condition x subject), carried out in Matlab 7.1 (The Mathworks, Natick, Massachusetts). When comparing activity at the endpoints of the BOLD response profile, we located the negative lobes of a BOLD slice by taking the smallest responses in the foveal and eccentric halves of the response profile. These provided additional "parameter estimates" for the amplitude of the negative dips – we used this approach, rather than using the amplitude parameter for the negative component of the DoG curve, in order to measure the foveal and eccentric dips separately. After obtaining these measurements of the amplitude of the negative dips, we compared them for attended and unattended conditions using three-way ANOVAs, just as with the DoG fit parameters.

Modeling

To model the affects of gain and tuning on the slope of a position discrimination plot, we first created a prototypical BOLD response profile using subjects' BOLD response profiles from the slice-based analysis. We aligned the response curves for the unattended conditions from all subjects peak-to-peak, normalizing by peak amplitude, and plotting them together (Figure 4A). Then we fit a difference-of-Gaussians (DoG) function to this combined data (black curve in Figure 4A), and used a discretely sampled version of this curve with 200 points in subsequent simulations.

To generate the green curve in Figure 4B, representing the position discrimination slope for unattended conditions, we used two copies of the empirically-defined DoG curve to simulate two different Gabor positions in a Montecarlo simulation. In each iteration, random noise was added to each curve independently, to produce an SNR of 15.33. We estimated this SNR from subjects' BOLD response maps using the correlations at the smallest separation for each subject. SNR is related to correlation by $\frac{S}{N} = \frac{\rho}{1-\rho}$ in the case where two copies of the same signal are corrupted by independent noise sources of the same energy. Since the correlations at the smallest separation reflect very similar but non-identical signals, our SNR estimate was necessarily conservative, and indicative of a lower bound. This estimate reflects the SNR of the BOLD response maps (resulting from GLM contrasts), and not the raw fMR timecourse. We computed SNR for the unattended runs for each subject and then averaged across subjects to yield this value for modeling. In each iteration of the Montecarlo simulation, the two curves were offset from each other by a peak-to-peak separation of δ in the range of [0.0, 4.0] mm, corresponding to the x axis in Fig. 4B & C. A correlation was then computed between the two curves to produce a Fisher z. At each separation δ, 5000 iterations were completed in this way, with new random noise added in

each iteration. We averaged the resulting z scores at each separation, and these average z scores were plotted against the separation δ at which they were computed to produce the green curve in Figure 4B & C. This curve simulates the slope of a position discrimination plot for unattended stimuli (Fig. 2), modeled on a continuous scale.

To generate the remaining curves in Figure 4B & C, we ran the same Montecarlo simulation, after scaling the DoG curves. For the gain plot (red), we increased the amplitude of the DoG curves by 10% to simulate a multiplicative amplitude gain, and all other parameters were kept the same. For the tuning plot (blue), we decreased the width of the DoG by 10% to simulate position tuning narrowing, and all other parameters were kept the same. For the combined gain and tuning plot (purple), both a 10% amplitude gain and a 10% narrowing were applied.

To test how the correlation between patterns of activity is affected by other levels of gain and tuning, we used the same pair of DoGs as described above, now fixed at a realistic separation of 1.51 mm, which was established by taking all subjects' average peak-to-peak separation of the BOLD response profiles produced by the most eccentric and most foveal conditions (corresponding to the largest separation on the position discrimination plot). We added noise to each curve as before, and carried out Montecarlo simulations by varying the amount of signal gain or width narrowing applied to the modeled curves (Figure 4D). In the case of signal gain, we modeled the outcome of multiplicatively scaling the response curves around zero (corresponding to response gain, in which baseline activity is not modulated by attention), multiplicatively scaling the response curves after adding back in a baseline constant (corresponding to activity gain, in which baseline activity is also scaled by attention), and simply adding a baseline constant (corresponding to additive gain). We modeled the last two over a range of baseline constants (from 1% - 1000% of the amplitude of the modeled BOLD response profile), but the value of the baseline constant had no impact on the resulting correlations; activity gain produced an identical curve to that of response gain (solid red curve in Fig. 4D), and additive gain had no impact on the correlations at any level of gain (dashed red line in Fig. 4D). We performed all modeling in Matlab 7.1.

Supplemental Figure 1

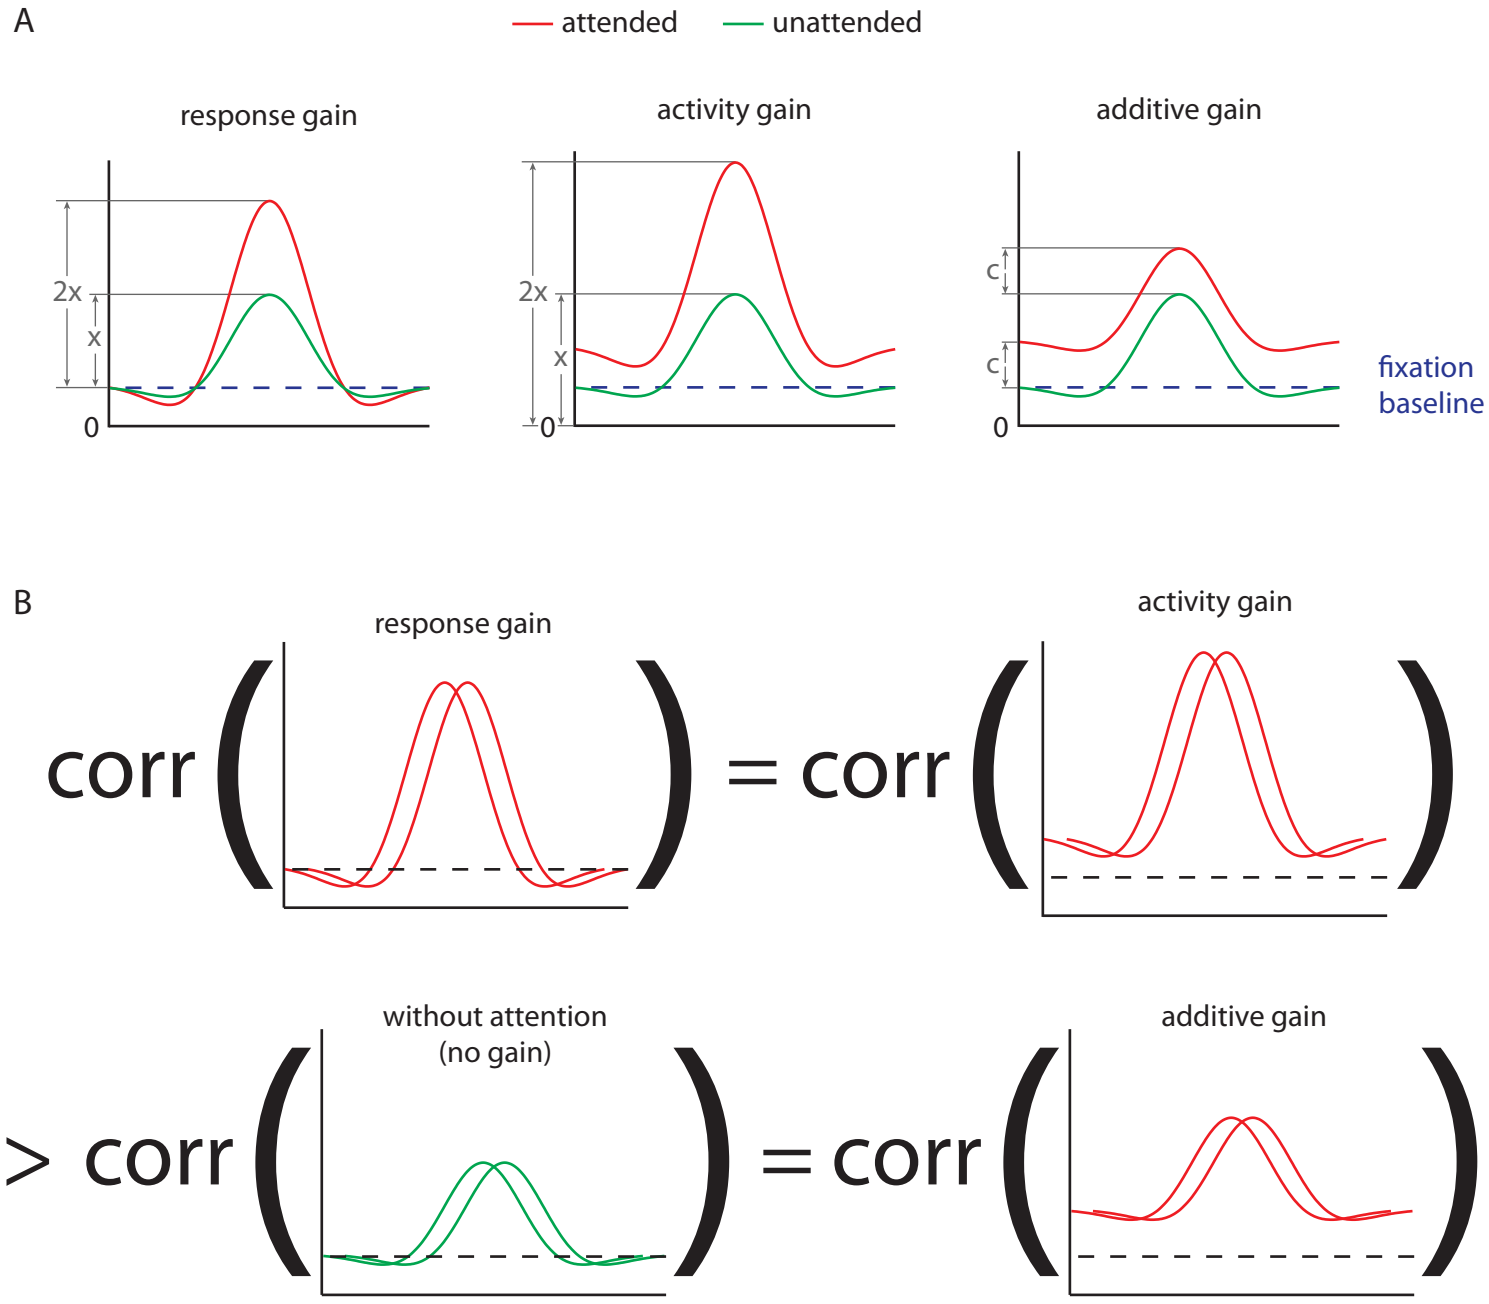

Supplemental Figure 1: The effect of response gain, activity gain, and additive gain on the correlation between BOLD patterns. In attention research, distinctions have been drawn among *response gain,* in which attention has a multiplicative effect on above-baseline activity (stimulus-evoked responses), *activity gain*, in which all activity, including baseline, is multiplicatively increased by attention, and *additive gain*, in which the activity level is increased by a constant, rather than scaled. These three scenarios are shown in panel **A)**. While activity gain yields larger overall responses than response gain, the shapes of the BOLD response curves generated by response gain and activity gain are identical; activity gain yields a vertically-shifted version of the response gain curve. Similarly, additive gain yields a vertically-shifted version of the unattended (green) curve. **B)** Response gain and activity gain produce an equivalent increase in the correlation between two curves, because they yield identically-shaped attended (red) curves. For the same reason, additive gain leaves the correlation between two curves unchanged, because it merely shifts the curves vertically. In this study, we use the term *signal gain* to refer to a vertical scaling of the BOLD response that could result from either response gain or activity gain. Signal gain always increases the correlation between two BOLD response curves, assuming the presence of noise, and positive initial correlation between the unattended curves (see Supplemental Figure 2 for a proof).

Supplemental Figure 2

Let X and Y be random variables with variances σ^2_x and σ^2_y, respectively.

Say that X and Y are corrupted by random noise ε_1 and ε_2, both with mean 0 and variance greater than zero, such that ε_1 and ε_2 are independent of X, Y, and each other. Then

$$\rho(X+\varepsilon_1, Y+\varepsilon_2) = \frac{Cov(X+\varepsilon_1, Y+\varepsilon_2)}{\sigma_{X+\varepsilon_1}\sigma_{Y+\varepsilon_2}} = \frac{Cov(X,Y)+Cov(X,\varepsilon_2)+Cov(\varepsilon_1,Y)+Cov(\varepsilon_1,\varepsilon_2)}{\sigma_{X+\varepsilon_1}\sigma_{Y+\varepsilon_2}}$$

$$= \frac{Cov(X,Y)}{\sigma_{X+\varepsilon_1}\sigma_{Y+\varepsilon_2}}$$ by the independence of X and ε_2, Y and ε_1, and ε_1 and ε_2.

Now, looking at the denominator,

$$\sigma_{X+\varepsilon_1}\sigma_{Y+\varepsilon_2} = \sqrt{\sigma^2{}_{X+\varepsilon_1}}\sqrt{\sigma^2{}_{Y+\varepsilon_2}} = \sqrt{\sigma^2{}_X+\sigma^2{}_{\varepsilon_1}}\sqrt{\sigma^2{}_Y+\sigma^2{}_{\varepsilon_2}};$$

Thus, $$\rho(X+\varepsilon_1, Y+\varepsilon_2) = \frac{Cov(X,Y)}{\sqrt{\sigma^2{}_X+\sigma^2{}_{\varepsilon_1}}\sqrt{\sigma^2{}_Y+\sigma^2{}_{\varepsilon_2}}}.$$

Now, introduce a multiplicative gain of magnitude *a* to X and *b* to Y, where *a, b* > 1.

The new correlation is given by

$$\rho(aX+\varepsilon_1, bY+\varepsilon_2) = \frac{Cov(aX,bY)}{\sqrt{\sigma^2{}_{aX}+\sigma^2{}_{\varepsilon_1}}\sqrt{\sigma^2{}_{bY}+\sigma^2{}_{\varepsilon_2}}} = \frac{abCov(X,Y)}{\sqrt{a^2\sigma^2{}_X+\sigma^2{}_{\varepsilon_1}}\sqrt{b^2\sigma^2{}_Y+\sigma^2{}_{\varepsilon_2}}} =$$

$$\frac{abCov(X,Y)}{a\sqrt{\sigma^2{}_X+\frac{\sigma^2{}_{\varepsilon_1}}{a^2}}b\sqrt{\sigma^2{}_Y+\frac{\sigma^2{}_{\varepsilon_2}}{b^2}}} = \frac{Cov(X,Y)}{\sqrt{\sigma^2{}_X+\frac{\sigma^2{}_{\varepsilon_1}}{a^2}}\sqrt{\sigma^2{}_Y+\frac{\sigma^2{}_{\varepsilon_2}}{b^2}}}.$$

Now, since $\sigma^2{}_{\varepsilon_1}, \sigma^2{}_{\varepsilon_2} > 0$ and $a^2, b^2 > 1$,

$$\sqrt{\sigma^2{}_X+\frac{\sigma^2{}_{\varepsilon_1}}{a^2}} < \sqrt{\sigma^2{}_X+\sigma^2{}_{\varepsilon_1}} \text{ and } \sqrt{\sigma^2{}_Y+\frac{\sigma^2{}_{\varepsilon_2}}{b^2}} < \sqrt{\sigma^2{}_Y+\sigma^2{}_{\varepsilon_2}}.$$

Thus,

$$\rho(aX+\varepsilon_1, bY+\varepsilon_2) = \frac{Cov(X,Y)}{\sqrt{\sigma^2{}_X+\frac{\sigma^2{}_{\varepsilon_1}}{a^2}}\sqrt{\sigma^2{}_Y+\frac{\sigma^2{}_{\varepsilon_2}}{b^2}}} > \frac{Cov(X,Y)}{\sqrt{\sigma^2{}_X+\sigma^2{}_{\varepsilon_1}}\sqrt{\sigma^2{}_Y+\sigma^2{}_{\varepsilon_2}}} = \rho(X+\varepsilon_1, Y+\varepsilon_2),$$

when X and Y covary positively. Therefore, a multiplicative gain applied to one or both signals in the presence of uncorrelated noise increases the correlation between those signals, given that they covary positively.

Supplemental Figure 2: Algebraic proof that multiplicative gain increases the correlation between two positively correlated signals. We assume that two signals of arbitrary shape are positively correlated, and are in the presence of nonzero noise. Applying a multiplicative scaling by a factor greater than one to one or both signals (as is the case with response gain or activity gain) increases the correlation between the two signals by increasing the signal strength while the noise level remains constant. Importantly, if the original signals are negatively correlated, then a multiplicative scaling of one or both signals will increase the strength of the correlation between them, yielding a smaller (more negative) correlation. In this study, all correlations that we measured were positive.

Supplemental Figure 3

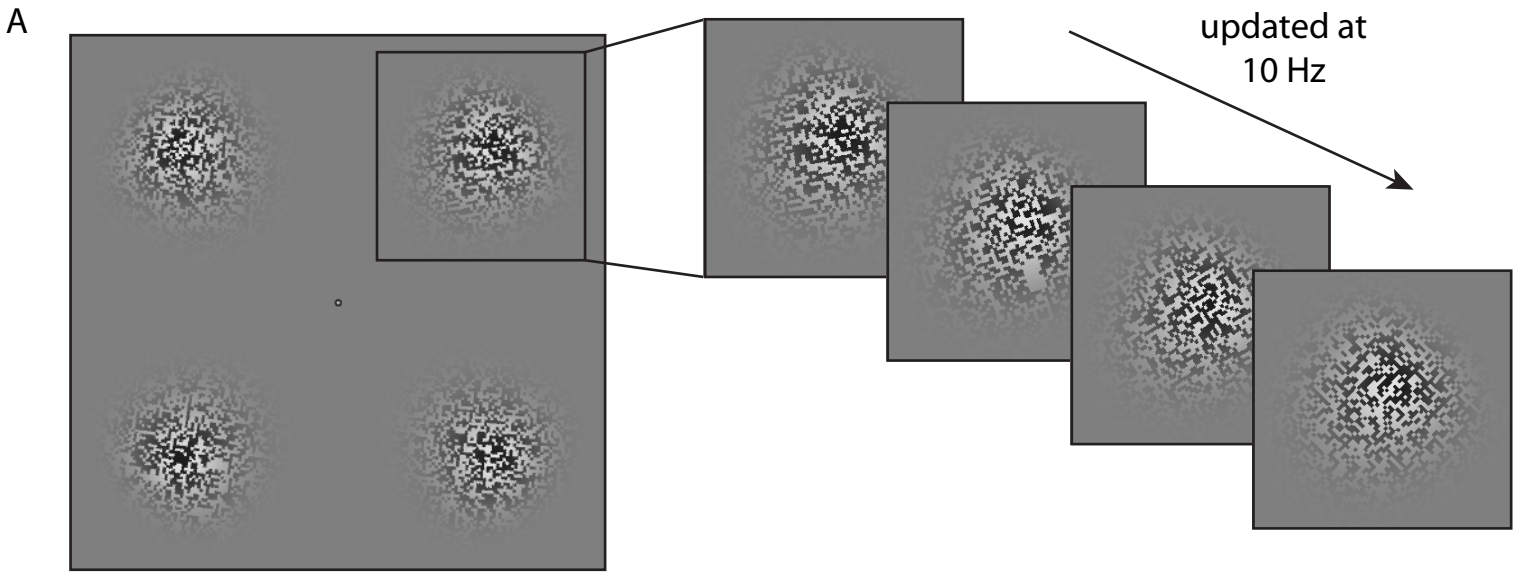

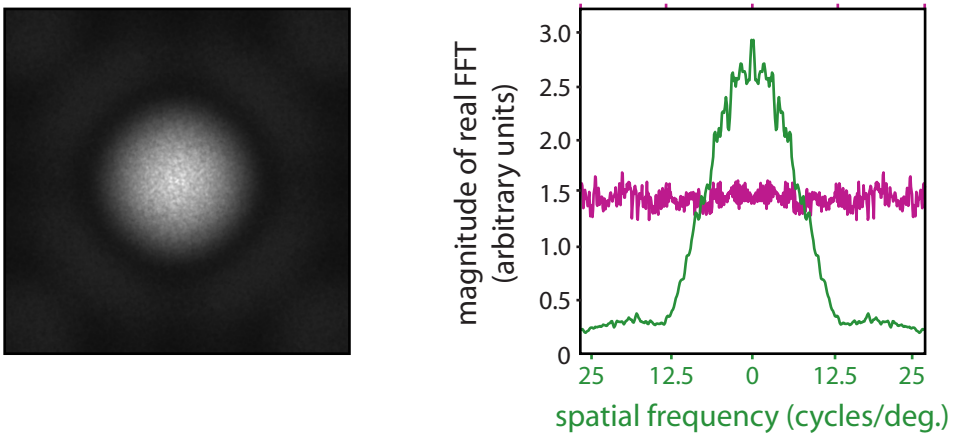

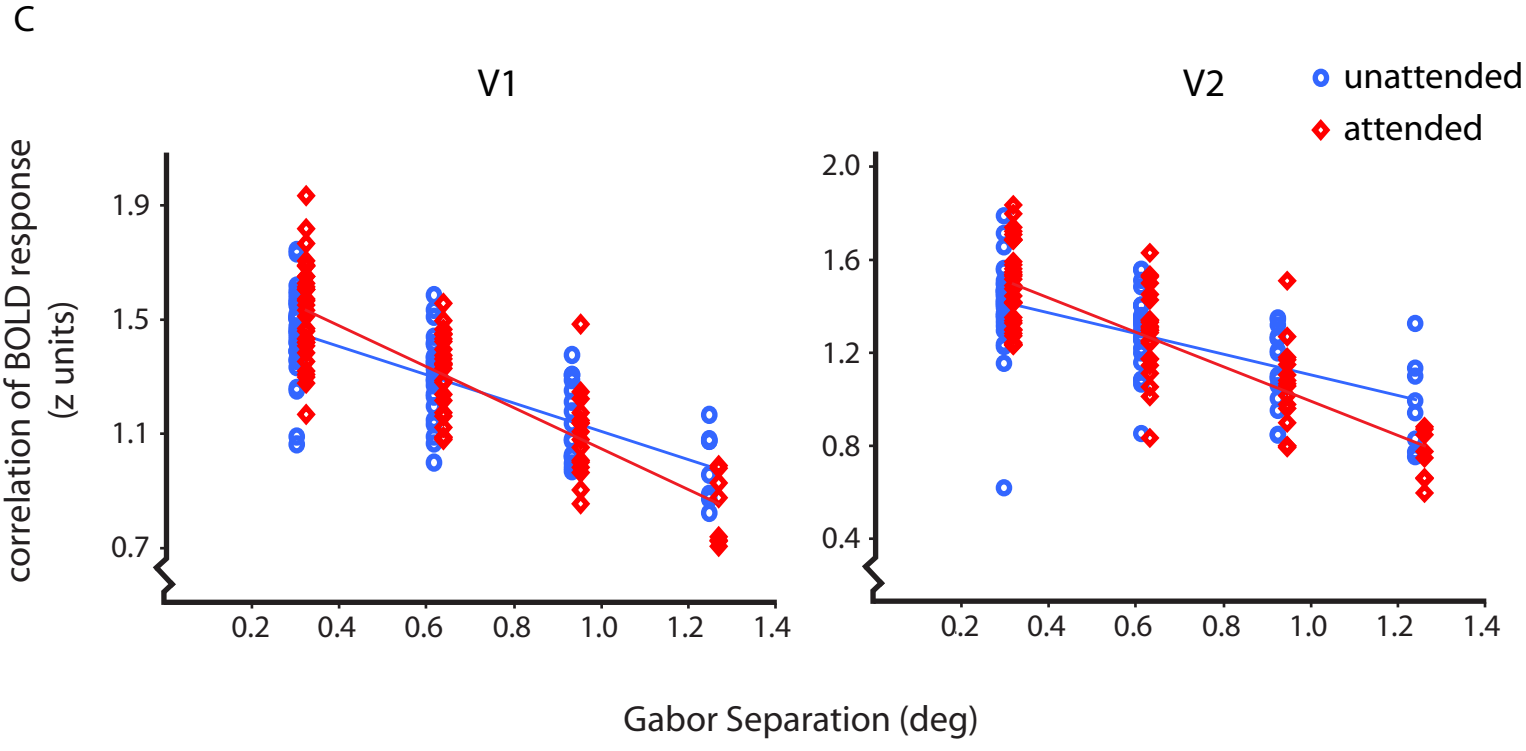

Supplemental Figure 3: Control experiment to rule out feature-based attention. **A)** To ensure that the effects we measured in the main experiment were due to spatially-directed, rather than feature-based attention, we performed a control experiment with two subjects in which we presented Gaussian-windowed noise instead of Gabor patches. Dark and light patches in the noise were 4 pixels by 4 pixels, allowing us to update the orientation of the grid on which the noise was created rapidly (10 Hz). The result was a stimulus that was broadband in spatial frequency and orientation. **B)** On the left is the 2d Fourier transform of a noise stimulus corresponding to the middle eccentricity, and on the right is a visualization of the power across the full range of spatial frequencies at orientation 0 (green), and orientations at 10 cycles/deg (purple; the DC component is omitted for plotting purposes). There are no dramatic spikes in the power spectrum, and power over spatial frequencies drops smoothly in a roughly 1/f fashion. In all other respects besides the stimulus, the control experiment was identical to the main experiment. **C)** We plotted both subjects together as in the main experiment (Figure 2); here, individual runs are plotted separately, and the x positions of the attended and unattended data are offset slightly for visualization purposes. Both V1 and V2 exhibit the crossover between attended and unattended correlations that we found in the main experiment, and there was a significant drop in the correlations with attention at the largest separation ($t = 2.92$, $p = 0.02$ in V1; $t = 3.27$, $p = 0.01$ in V2) indicating reduced spread of the BOLD response with attention. The increase in correlations at the largest separation corresponding to signal gain was not quite significant ($t = 1.87$, $p = 0.07$ in V1; $t = 1.65$, $p = 0.11$ in V2), but by comparing direct measurements of the largest responses in attended and unattended conditions, we verified that attention increased the peak responses in both ROIs (Subj. DH: $t = 4.03$, $p < 0.001$ in V1, $t = 9.42$, $p < 0.0001$ in V2; Subj. NS: $t = 3.87$, $p = 0.001$ in V1, $t = 6.31$, $p < 0.0001$ in V2). Because these noise stimuli contained no dominant spatial frequency or orientation, we can be confident that spatially-directed attention narrows population position tuning.

Supplemental Figure 4

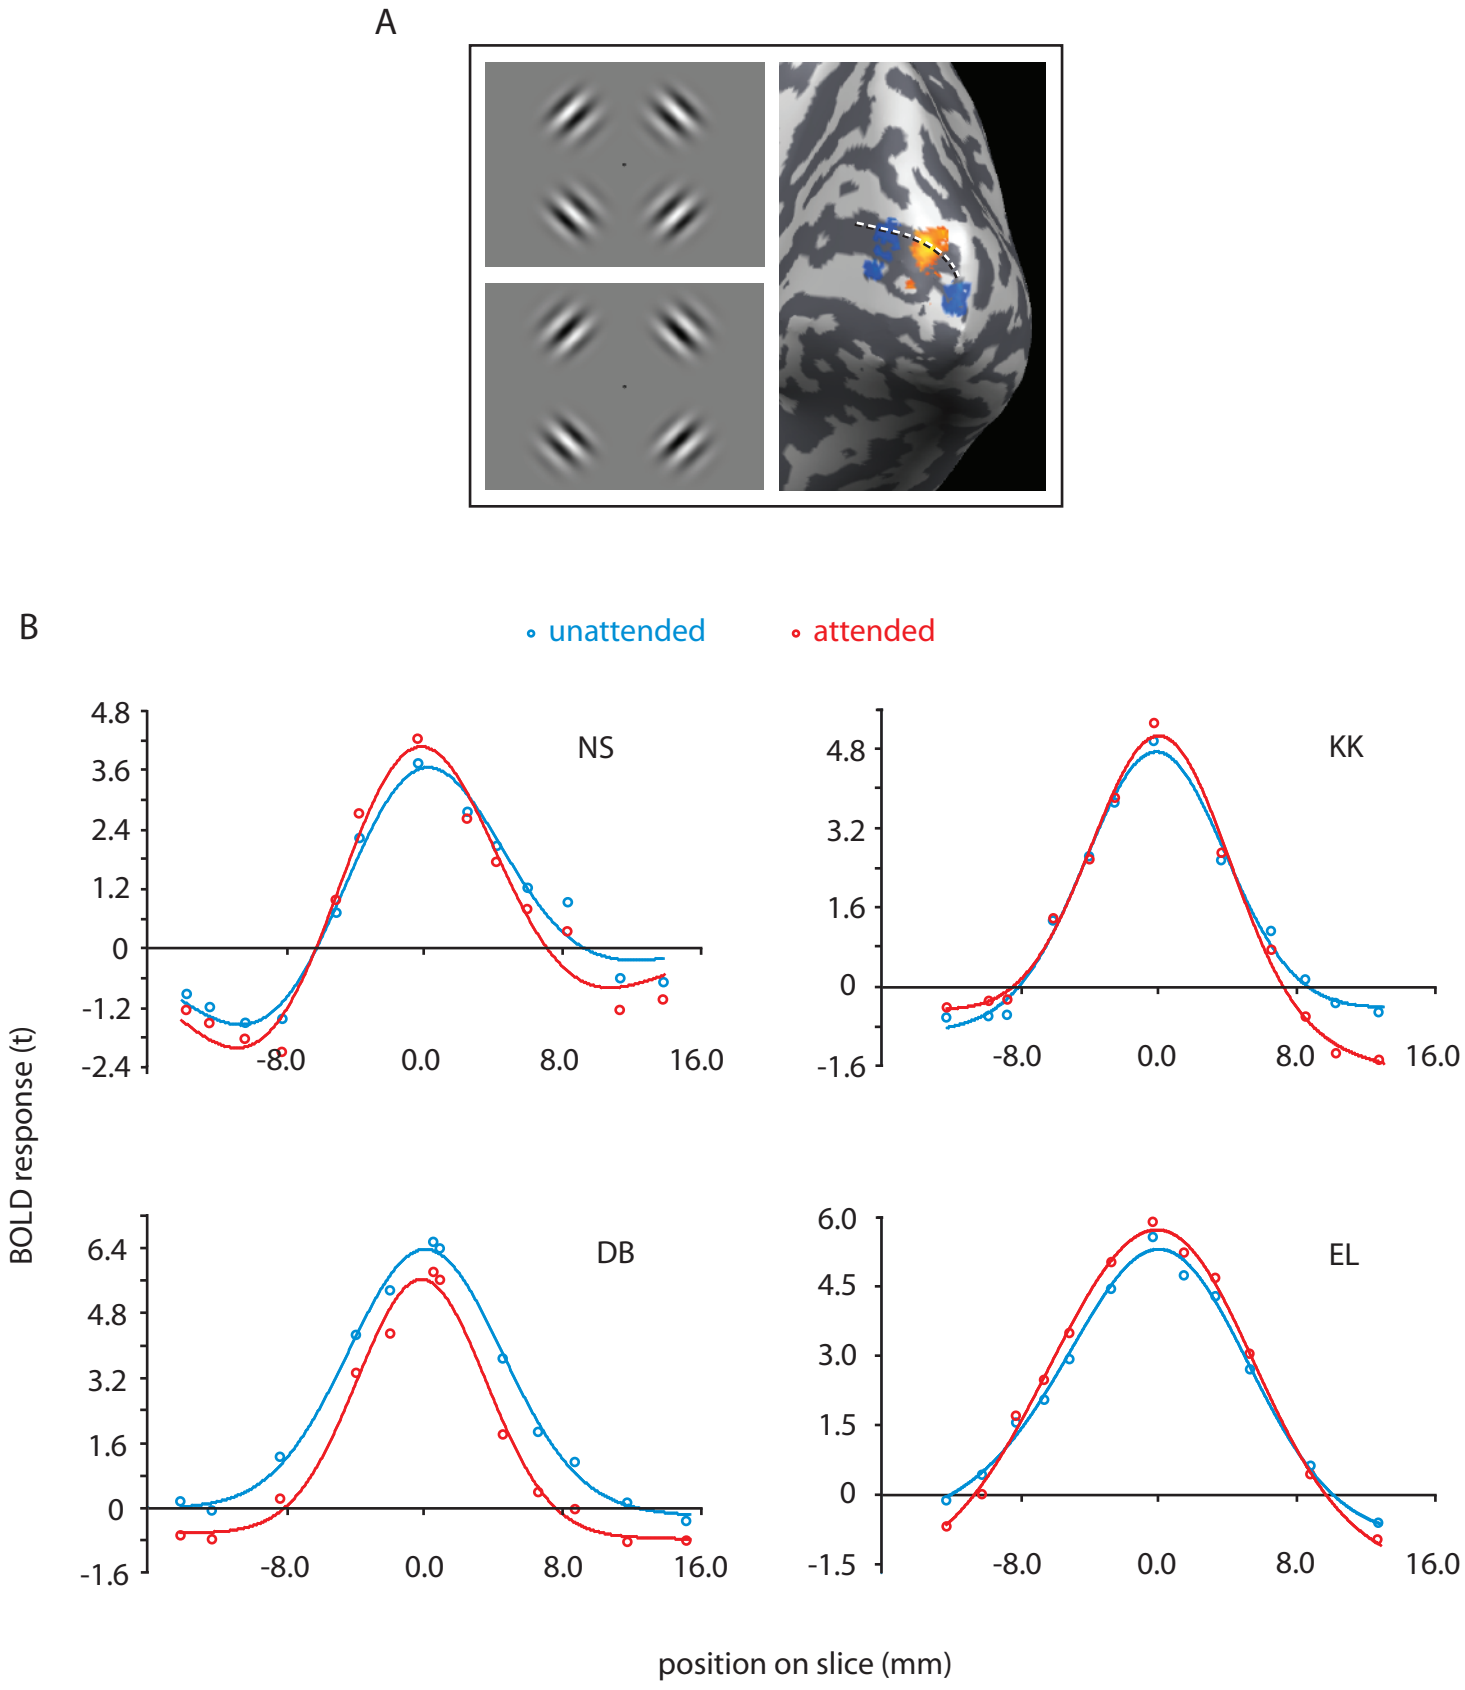

Supplemental Figure 4: Direct measurement of the shape of the BOLD response. **A)** We found the direction of retinotopic progression of the stimuli by contrasting the least and most eccentric stimulus positions in a general linear model. We overlaid each subject's resulting map on his or her inflated cortex, and drew a slice through the peak activation in V1, along the direction of the progression of the stimulus eccentricity manipulation. We measured the BOLD response for each of the five stimulus eccentricities, for attended and unattended trials, and fit difference of Gaussians curves to each set of measurements to characterize the shape of the BOLD response profiles. We compared the amplitude, width, and phase parameters of the curve fits for attended and unattended conditions, and found significant narrowing ($F = 10.24$, $p = 0.002$) and signal gain ($F = 12.57$, $p < 0.001$) with attention, but no change in the positions of the BOLD responses ($F = 0.19$, $p = 0.67$). We also tested for a decrease in the BOLD response with attention at the locations of the two negative lobes by taking the lowest responses at the foveal and eccentric edges of each subject's attended and unattended BOLD response curves. We found that both edges of the stimulus representation were significantly modulated by attention ($F = 8.28$, $p = 0.006$ at the foveal edge; $F = 11.13$, $p = 0.001$ at the eccentric edge). These results rule out a spatially asymmetric effect of attention on the BOLD response profiles, as well as showing that attention acted over the entire stimulus representation. **B)** We created a visualization of the BOLD response profiles for four subjects by aligning the five eccentricity conditions peak-to-peak and averaging across the five conditions. Though subtle, a decrease in the spread of the BOLD response is evident in the fact that the attended (red) curves have a smaller width than the unattended curves at the x-axis.

Supplemental Figure 5

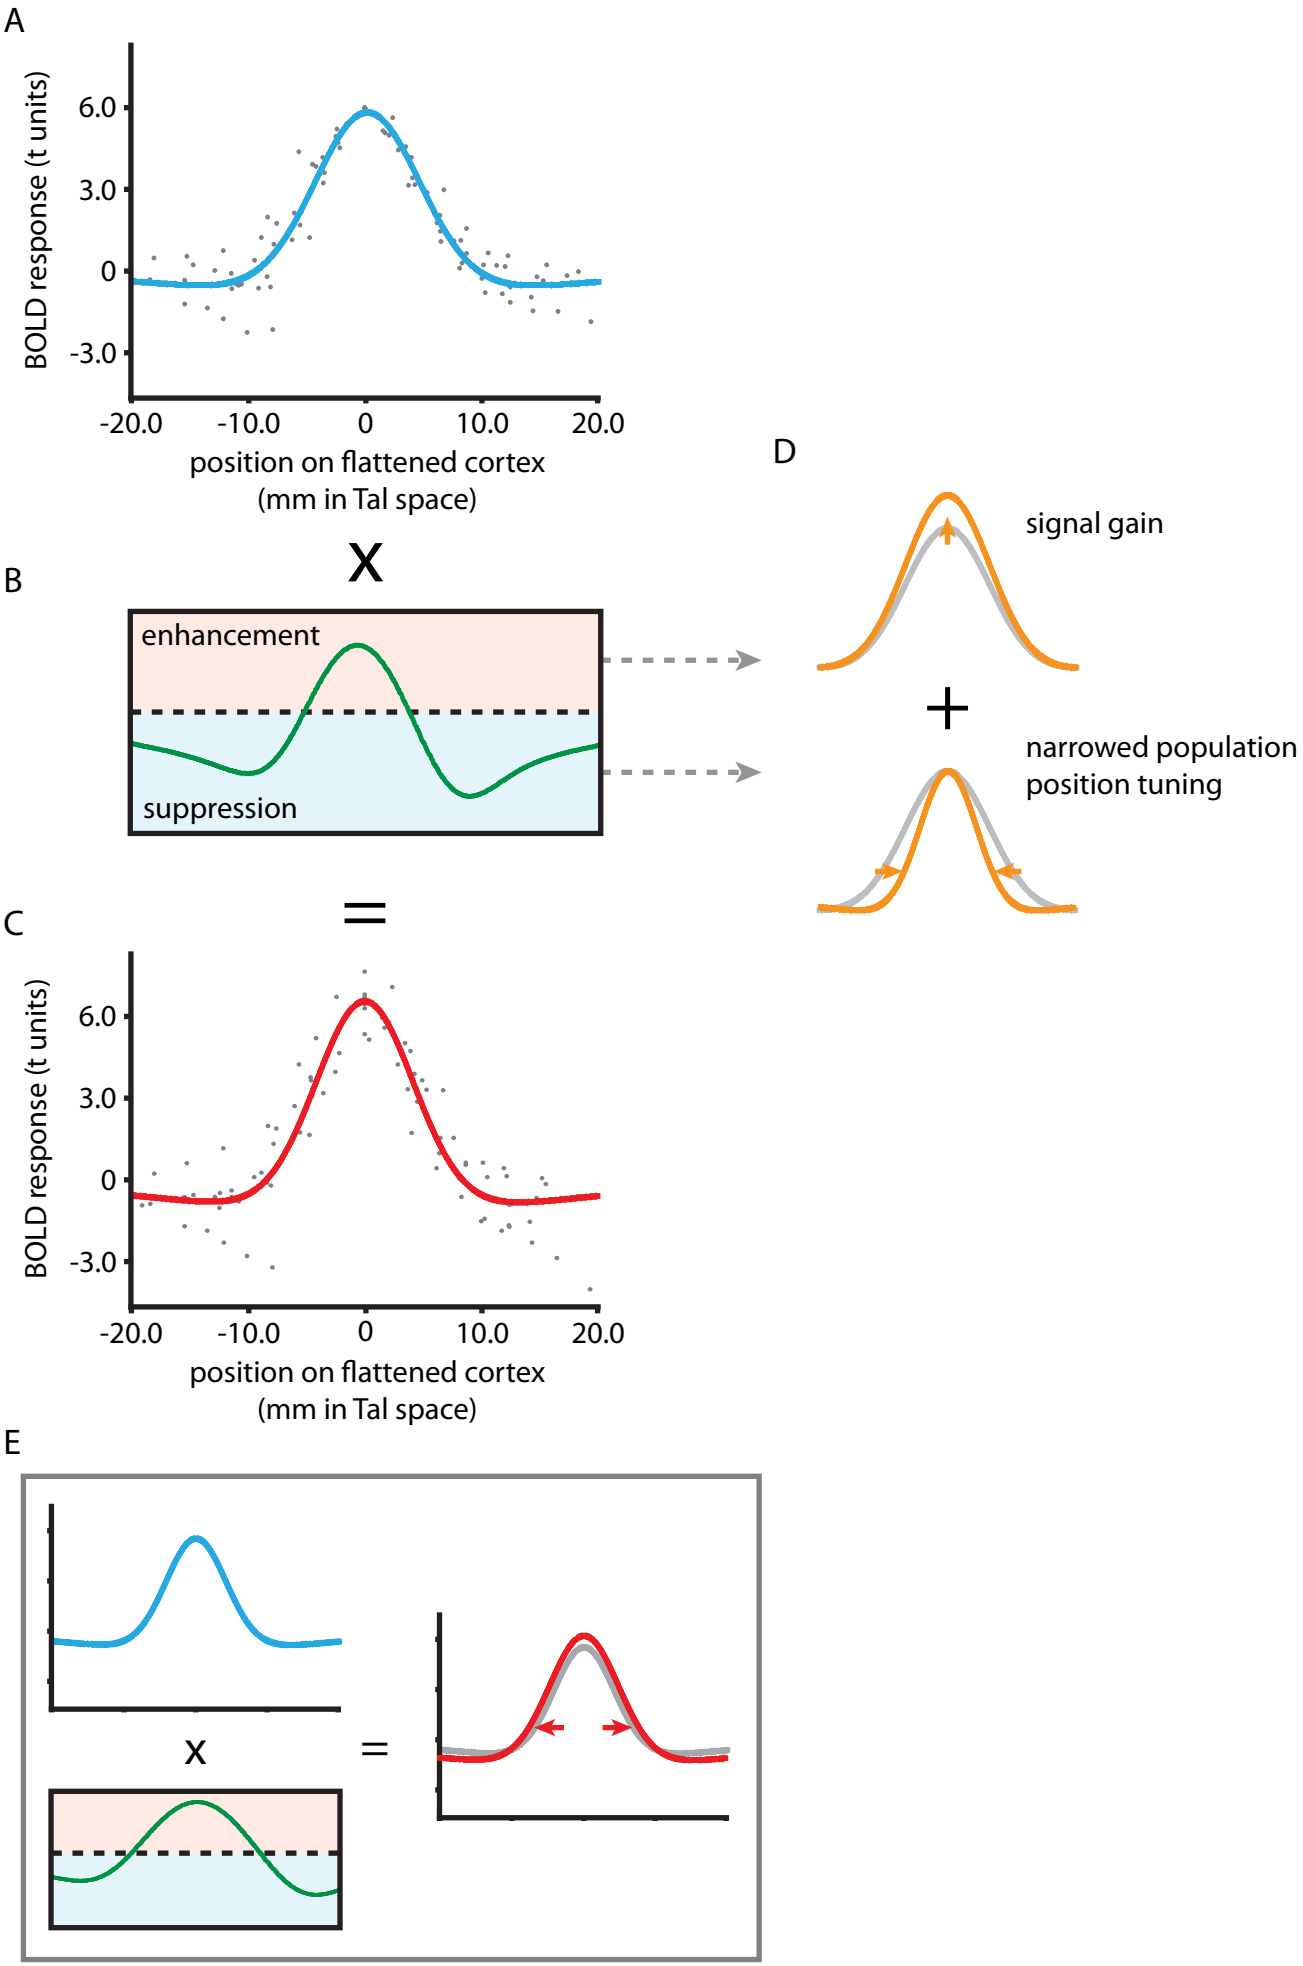

Supplemental Figure 5: Visualizing signal gain and position tuning narrowing as a spatially heterogeneous gain function. **A) - C)** The simultaneous signal gain and narrowing produced by focused attention can be thought of as a spatially distributed field of enhancement and suppression. We generated group BOLD response profiles for unattended (blue; panel **A**) and attended (red; panel **C**) stimuli in order to visualize how attention scaled the BOLD response as a function of position along the slice. We plotted all subjects' slice data together (see Supp. Fig. 4), normalizing to the group mean peak BOLD response for the unattended data. The blue and red profiles are difference of Gaussians curve fits to the attended and unattended slice data, respectively. By deconvolving the attended and unattended curves, we recovered the attentional modulation kernel in panel **B**, showing the factor by which attention scales the BOLD response as a function of position along the slice. The dashed horizontal line is the line of unity; where the green curve intersects this line, attention left the BOLD response unchanged. Above the line of unity, attention scaled the BOLD by a factor greater than 1, and below the line of unity, attention scaled the BOLD by a factor less than 1. Because the regions of suppression encompass much of the positive lobe of the unattended BOLD response profile in **A)** above, scaling the blue curve by the attentional modulation kernel yields a narrowed red curve in panel **C** below. **D)** The attentional modulation kernel, which scales the BOLD independently at each location along the slice, can be parsimoniously summarized by two simple operations acting on the entire distribution: vertical and horizontal scaling of the population response. Each of these operations is directly related to visual resolution: vertical scaling affects the signal-to-noise ratio, and horizontal scaling affects the overlap between neighboring response distributions. **E)** The existence of suppression in unattended regions need not predict narrowed population position tuning at the location of the stimulus. As reported by Tootell et al. [6], Somers et al. [7], Smith et al. [8], Müller & Kleinschmidt [9], and others, focused attention can yield a decrease in the BOLD response relative to baseline at locations remote from the attended region. These studies established that attention can have heterogeneous effects on the BOLD response which extend far beyond the attended location, but, critically, they did not resolve how attention alters the distribution of population responses corresponding to the stimulus itself. As illustrated here, the mere existence of such suppressive regions does not imply that attention will narrow the distribution of BOLD response at the location of the stimulus. In the possibility depicted in panel **E**, attention increases the peak of the BOLD response but also *broadens* the distribution of the BOLD, despite having a suppressive effect in the surrounding regions. In light of this possibility, previous studies cannot reach a conclusion about how attention affects population position tuning because they did not directly measure the distribution of the BOLD response at the location of the stimulus, as we have done here.

Supplemental References

1. Whitaker, D., McGraw, P.V., Pacey, I., and Barrett, B.T. (1996). Centroid analysis predicts visual localization of first- and second-order stimuli. Vision Res *36*, 2957-2970.
2. Bressler, D., Spotswood, N., and Whitney, D. (2007). Negative BOLD fMRI response in the visual cortex carries precise stimulus-specific information. PLoS ONE *2*, e410.
3. Talairach, J., and Tournoux, P. (1988). Co-planar stereotaxic atlas of the human brain : 3-dimensional proportional system: an approach to cerebral imaging, (Stuttgart ; New York, New York: G. Thieme; Thieme Medical Publishers).
4. Sereno, M.I., Dale, A.M., Reppas, J.B., Kwong, K.K., Belliveau, J.W., Brady, T.J., Rosen, B.R., and Tootell, R.B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. Science (New York, N.Y *268*, 889-893.
5. Fischer, J., and Whitney, D. (2007). Precise discrimination of object position in the human pulvinar. Hum Brain Mapp.
6. Tootell, R.B., Hadjikhani, N., Hall, E.K., Marrett, S., Vanduffel, W., Vaughan, J.T., and Dale, A.M. (1998). The retinotopy of visual spatial attention. Neuron *21*, 1409-1422.
7. Somers, D.C., Dale, A.M., Seiffert, A.E., and Tootell, R.B. (1999). Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. Proceedings of the National Academy of Sciences of the United States of America *96*, 1663-1668.
8. Smith, A.T., Singh, K.D., and Greenlee, M.W. (2000). Attentional suppression of activity in the human visual cortex. Neuroreport *11*, 271-277.
9. Muller, N.G., and Kleinschmidt, A. (2004). The attentional 'spotlight's' penumbra: center-surround modulation in striate cortex. Neuroreport *15*, 977-980.