Article

CellPress

Attraction of Position Preference by Spatial Attention throughout Human Visual Cortex

Barrie P. Klein,[1] Ben M. Harvey,[1] and Serge O. Dumoulin[1,*]
[1]Experimental Psychology, Helmholtz Institute, Utrecht University, Heidelberglaan 1, 3584 CS Utrecht, the Netherlands
*Correspondence: s.o.dumoulin@uu.nl
http://dx.doi.org/10.1016/j.neuron.2014.08.047

SUMMARY

Voluntary spatial attention concentrates neural resources at the attended location. Here, we examined the effects of spatial attention on spatial position selectivity in humans. We measured population receptive fields (pRFs) using high-field functional MRI (fMRI) (7T) while subjects performed an attention-demanding task at different locations. We show that spatial attention attracts pRF preferred positions across the entire visual field, not just at the attended location. This global change in pRF preferred positions systematically increases up the visual hierarchy. We model these pRF preferred position changes as an interaction between two components: an attention field and a pRF without the influence of attention. This computational model suggests that increasing effects of attention up the hierarchy result primarily from differences in pRF size and that the attention field is similar across the visual hierarchy. A similar attention field suggests that spatial attention transforms different neural response selectivities throughout the visual hierarchy in a similar manner.

INTRODUCTION

Voluntary attention directed at a visual location, i.e., endogenous spatial attention, represents a major influence of our cognitive state on neural processing and visual perception (Anton-Erxleben and Carrasco, 2013). In order to concentrate processing resources at the attended location, attention affects neural response properties (Martínez-Trujillo and Treue, 2002; Martinez-Trujillo and Treue, 2004; McAdams and Maunsell, 1999; Treue and Maunsell, 1999). These effects are thought to be stronger at higher levels of visual processing (Buffalo et al., 2010; Cook and Maunsell, 2002; Montijn et al., 2012; O'Connor et al., 2002; Posner and Gilbert, 1999).

Models based on Gaussian interactions capture many aspects of these attentional effects. These models incorporate a multiplication of two Gaussian components (Reynolds and Heeger, 2009; Womelsdorf et al., 2008). In the case of spatial attention, the first of these, the stimulus-driven receptive field, represents each neuron's response selectivity in the absence of attention. The second component, the attention field, represents attention's influence and is centered at the attended location, regardless of the receptive fields with which it interacts.

The product of these two components predicts the neural receptive field measured under the influence of attention. Specifically, this Gaussian multiplication predicts a shrinkage of receptive fields and a reweighting of the spatial response selectivity toward the attended location (Figure 1B). Indeed, such reweighting of response selectivity toward attended locations has been observed in individual neurons in macaque V4 (Connor et al., 1996, 1997) and MT (Womelsdorf et al., 2006, 2008), leading to attraction of position preferences toward the attended location. However, neither the entire visual field nor other visual areas were sampled. Gaussian interaction models predict that this reweighting of response selectivity, leading to changes in preferred position, should occur in all receptive fields throughout the visual field representation and visual hierarchy. Here, we ask whether preferred position changes are evident (i) in humans, (ii) throughout the visual field, and (iii) throughout the visual hierarchy. Capturing these effects with Gaussian interaction models allows us to examine (iv) how well these models describe changes in preferred position and (v) how the attention field changes throughout the visual hierarchy.

To characterize these changes in response selectivity, we used high-field functional MRI (fMRI) (7T) to measure population receptive fields (pRFs) (Dumoulin and Wandell, 2008) in five human subjects. Subjects fixated on the center of a display showing a visual field mapping stimulus (Amano et al., 2009; Dumoulin and Wandell, 2008; Harvey and Dumoulin, 2011; Zuiderbaan et al., 2012) while performing an attention-demanding contrast discrimination task to the left or right of the display (Figure 1A). At these task locations lay circular patches of changing contrast-defined pink noise that randomly and independently increased in contrast. Subjects were instructed to report contrast increments of one covertly attended noise pattern, alternating sides between scanning runs (see Experimental Procedures and Movie S1).

The pRF describes the aggregate spatial response selectivity of the neural population within a given cortical location in terms of preferred position (in horizontal and vertical dimensions) and spatial extent of the tuning function (σ) for both attention
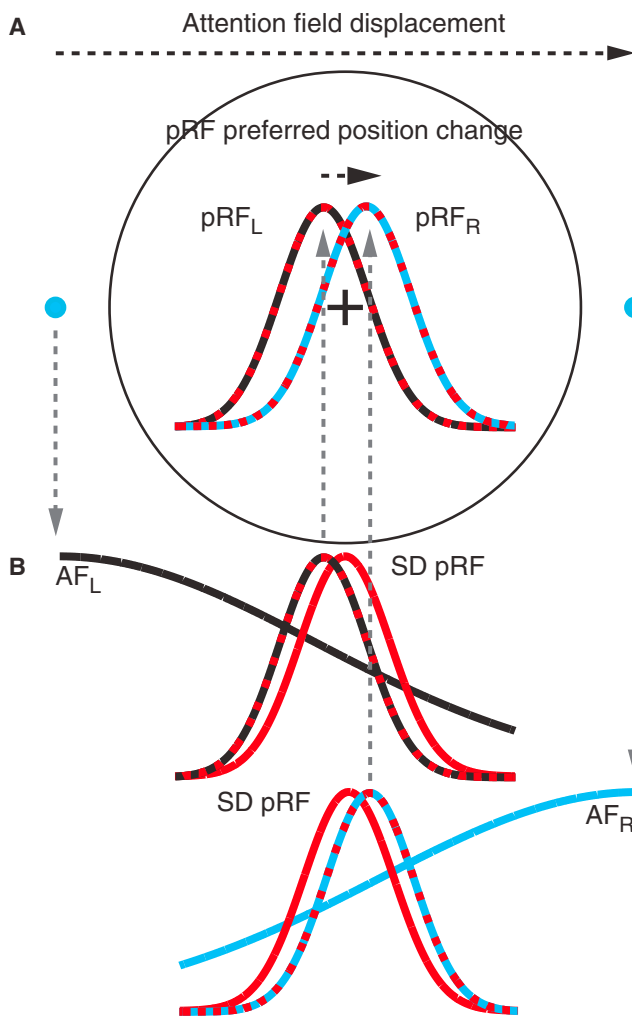
Figure 1. The Stimulus Paradigm and Attention Field Model
(A) Subjects fixated at the cross while performing an attention-demanding contrast discrimination task at either the left or right blue dot. The black circle marks the region where a conventional mapping stimulus was shown to derive the pRF properties (Dumoulin and Wandell, 2008).
(B) We modeled the effect of spatial attention at each attended location as a Gaussian attention field (AF, solid black and blue Gaussians). Multiplying each attention field with the stimulus-driven pRF (SD pRF, solid red Gaussian) results in the measured population receptive fields (pRF_L and pRF_R, black/red and blue/red Gaussian). This model predicts that measuring the pRF preferred position in either attention condition results in a pRF preferred position change between conditions.

sitions between the two conditions that produce equally sized attention fields at different locations is captured by the equation

$$\text{pRF preferred position change} = \frac{(\mu_{AFr} - \mu_{AFl})\,\sigma_{SD}^2}{\sigma_{AF}^2 + \sigma_{SD}^2}, \quad \text{Equation 1}$$

where the σ_{SD} is the stimulus-driven pRF size, σ_{AF} is the attention field size, and $(\mu_{AFr} - \mu_{AFl})$ represents the distance between the two attention field positions (see Experimental Procedures). Comparing between symmetrical attended locations cancels out the effects of distance to the attended location (Equation 10) such that the observed change in position preference does not depend on visual field position. Therefore, this approach highlights the changes in pRF preferred position due to spatial attention.
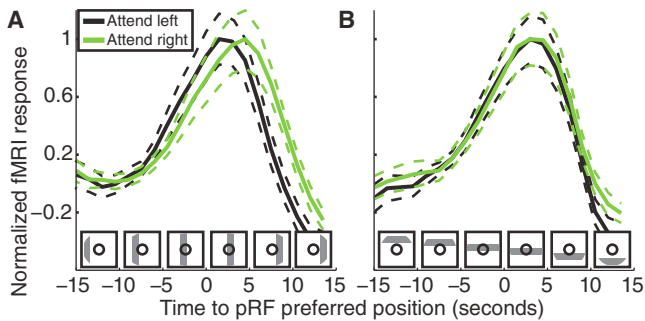
Because we know both attended locations and because the stimulus-driven pRF size can be calculated from the measured pRF size and the attention field size (Equation 5), this Gaussian interaction model only has one parameter to fit: the attention field size (Equation 1). This parameter represents the magnitude of attention's influence on spatial response selectivity, with smaller attention field sizes producing larger changes in preferred position.

We found attention-induced changes in pRF preferred positions across all visual field positions and all visual field maps measured (V1 to IPS4). These preferred position changes increase up the visual hierarchy and are strongly correlated with pRF size. A Gaussian interaction model (Equation 1) captures these preferred position changes well and estimates the extent of the attention field that would predict the changes seen in each visual field map. The extent of the attention field does not show a systematic progression through the visual hierarchy. Indeed, a model constraining the attention field to be common across all visual field maps does similarly well in explaining the observed changes in pRF preferred position. As such, our model suggests that the increase in pRF preferred position changes up the visual hierarchy may be captured by similar transformations of the visual field representation. Thus, variations in the extent of preferred position change arise primarily from variations in the stimulus-driven pRF size. The similarity of this transformation suggests biological constancies in the implementation of spatial attention throughout the visual hierarchy.

RESULTS

fMRI Time Series Differ between Attentional Conditions
To examine the effect of the attentional manipulation on the raw fMRI time series, we computed the average response to each bar sweep relative to the center of each cortical location's pRF (see Experimental Procedures). We averaged the responses to all bar crossings and attend-left and attend-right conditions separately. Figure 2 presents the average time series obtained from voxels in V3a within 2.5° of eccentricity in one subject when the subject attended the target on the left (black dashed line) or right (green dashed line) of fixation. When the bar moved from left to right (Figure 2A), the response rose and peaked earlier when the subject attended to the left compared to when they attended to the right of the display. However, fMRI

conditions separately. We assume that performing the task induces an attention field centered at the task location (Womelsdorf et al., 2008). Despite an identical stimulus and stimulus-driven pRF in both conditions (Figure 1B; SD pRF), Gaussian interaction models predict a change in the measured pRF position preference (Figure 1B; pRFl and pRFr) because of the change in attention field position (Figure 1B; AFl and AFr) between the two conditions. This change reflects the reweighting of the underlying spatial response selectivity toward the attended location. The hypothesized change in pRF preferred po-

Figure 2. Averaged Raw fMRI Responses and the SEM

(A) Averaged fMRI responses to a horizontal bar sweep for V3a aligned to their respective pRF position. Black line and data points represent the attend-left condition; green line and data points represent the attend-right condition.
(B) Identical representation as in (A) for a vertical bar sweep. The insets demonstrate the bar position relative to the mean pRF position. The fMRI responses differ for the horizontal bar sweep, but not the vertical bar sweep. Averaged raw fMRI responses are shown as solid lines, and the SEM is shown as dashed lines.

responses are not affected by attended location when the bar moved from the top to the bottom of the stimulus display (Figure 2B). This indicates that the pRF's response profile changes in the direction of the attentional manipulation, i.e., in the horizontal direction.

PRF Position Estimates Differ between Attentional Conditions

We summarize these fMRI responses using pRF models (Dumoulin and Wandell, 2008). For each cortical location and in each attention condition, these models describe a 2D difference of Gaussians function with four parameters: horizontal position (x), vertical position (y), size (standard deviation) of the positive Gaussian (σ_1), and size (standard deviation) of the negative Gaussian (σ_2) representing the suppressive surround (Zuiderbaan et al., 2012). We used the position and size parameters of the positive Gaussian to compare the spatial response selectivity of each cortical location between the two attended locations.

To summarize the pRF preferred position changes, we divided the cortical locations in each visual field map into 16 bins based on the mean pRF position preference across both conditions. For each section, we compared the mean pRF position when the subject attended the left and right targets (Figure 3A; Figure S1, available online, shows the steps involved in the binning procedure). Preferred positions change with attended location across the entire sampled visual field, with the direction of change centered along the horizontal axis.

To quantify the observed position changes, we describe the relationship between visual field eccentricity and preferred position change or pRF size using a linear function. We extract the middle of the stimulus eccentricity range (2.5°) from this function and use this as a representative value for pRF size or position change in each visual field map. Figures 3B and 3C illustrate the pRF preferred position change in each visual field map for one example subject and the average across all subjects,

respectively. Positive values indicate preferred position changes in the direction that the attended location changed. pRF preferred position changes are significantly larger than zero (dependent samples t test, all p values $< 10^{-6}$, corrected for upsampling and corrected for multiple comparisons using false discovery rates [FDRs]; Benjamini and Hochberg, 1995) for each visual field map when data from all subjects are combined. The pRF preferred position change increases up the visual hierarchy, demonstrating an increasing effect of attention. Similarly, pRF sizes increase across the visual hierarchy (Figure 3D), consistent with previous reports (Amano et al., 2009; Harvey and Dumoulin, 2011). The pRF sizes in each visual field map correlate with the preferred position changes (Figures 3E and S3C) (Pearson's correlation coefficient: r = 0.93, p < 10⁻⁵, n = 15). To examine the dependency of the preferred position changes on attention task and stimulus design, the experiment was repeated in three subjects (two of which participated in the main experiment) with a different stimulus and task (Movie S2). Despite the different task and mapping stimulus, the pRF preferred position changes measured during this experiment are highly correlated with those measured in the main experiment (Figures S3A and S3B; Pearson's correlation coefficient: r = 0.88, p < 10⁻⁵). This demonstrates that pRF preferred position changes are not limited to our specific stimulus setup or task.

Finally, we examined how much the cortical location of the stimulus's neural representation changes between our two conditions due to these changes in pRF preferred positions. To calculate the pRF preferred position changes in millimeters along the cortical surface (Figure 3F), we computed the cortical magnification factor (see Experimental Procedures) of each cortical location on the gray/white matter boundary and multiplied this with its pRF preferred position change. Again, we use a linear function to describe the relation between the cortical shifts and eccentricity and extract a representative value at 2.5° eccentricity. Like the changes in pRF preferred positions, the shift in activation along the cortical surface increases up the visual hierarchy.

Eye Movements toward the Targets Cannot Explain pRF Preferred Position Changes

Prior to fMRI scanning, subjects were trained to perform the task with minimal eye movements. Only if the subjects were able to do the task while keeping eye movements within a standard deviation of 0.25° did they continue to do the fMRI scanning sessions. We did not record the eye movement during the scanning sessions because of technical constraints. Because eye movements toward the attended targets might move pRF preferred position estimates toward the attended location, we simulated the effect of eye movements. We created three sets of simulated data by using a single data set from each participant and incorporated eye movements toward the attended target of 0.2°, 1°, and 2° (Levin et al., 2010). Next, we computed the pRF preferred position changes for each visual field map and plotted these changes as a function of pRF size (Figure 3E, gray dashed lines). Simulated directionally biased eye movements induce changes in pRF preferred positions but without a strong increase up the visual hierarchy or
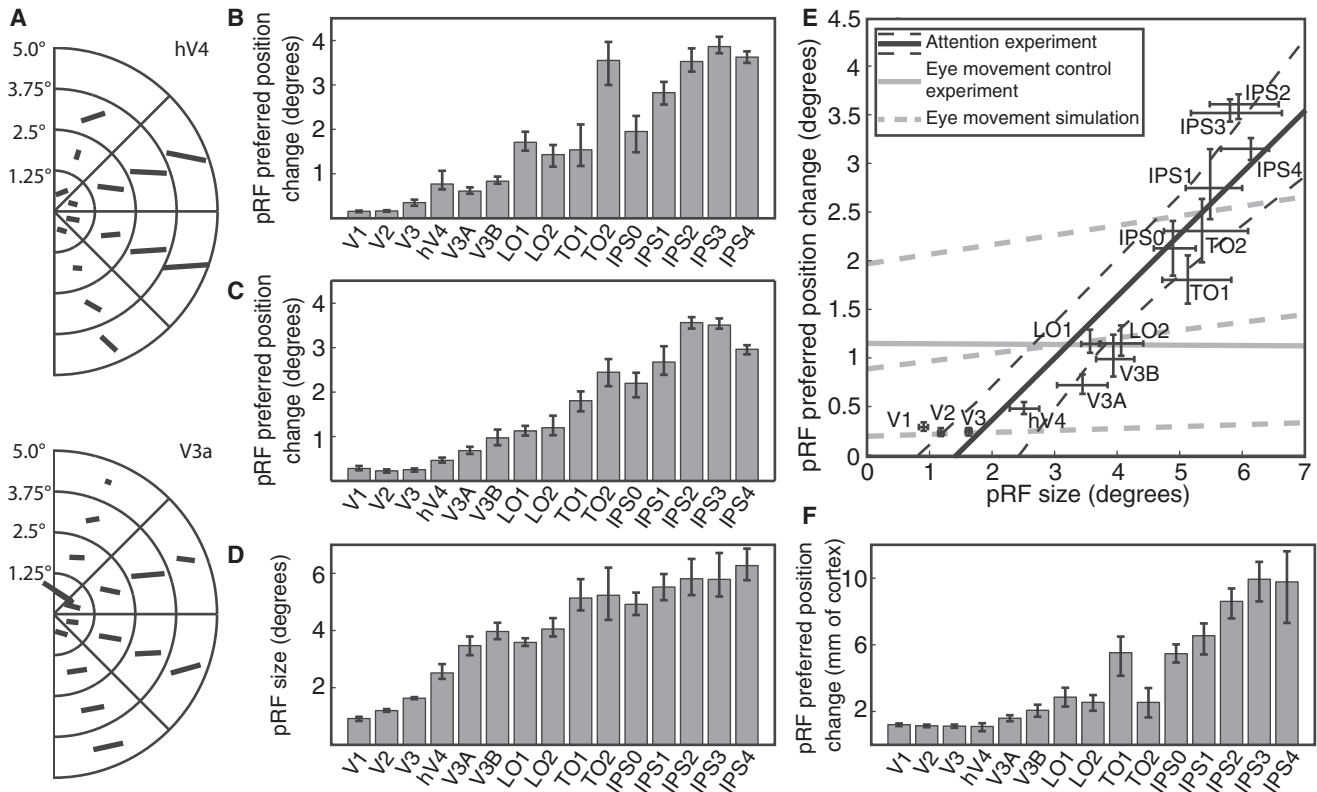
Figure 3. Measured pRF Preferred Position Changes and Their Relation to pRF Size

(A) Average size and direction of pRF preferred position change between the two attention conditions in different sectors of the visual field for one subject. All preferred position changes are in the expected horizontal direction. See Figure S2 for data from all visual field maps for each individual subject and all subjects combined. Positions of the lines are adjusted to pass through the center of each sector.

(B) Horizontal preferred position change for each examined visual field map from a single subject, showing that preferred position changes increase up the visual hierarchy.

(C) Similar to (B) but for all subjects combined.

(D) pRF sizes for each visual field map for all subjects combined.

(E) Differences in preferred position changes across visual field maps are strongly correlated with average pRF sizes (black crosses, from C and D). The solid black line is the best linear fit to these changes. The solid gray line represents the fit when gaze locations are separated by 1° between conditions (see Figure S3C). Gray dashed lines represent linear fits to data from eye movement simulations of 0.2°, 1°, and 2° toward the attended location. Together, these data demonstrate that eye movements produce similar pRF position changes for all pRF sizes and cannot explain the observed pRF position changes for the attention experiment.

(F) Preferred position changes converted to cortical surface distances. All error bars represent 95% confidence intervals determined by bootstrapping.

relationship with pRF size. To confirm these simulation results with experimental data, we collected fMRI data for two of our subjects while they viewed the same pRF mapping stimulus. The fixation point was positioned 0.5° either left or right of the center of the mapping stimulus, resulting in a 1° difference in eye position between the two conditions. Participants responded to color changes of the fixation point. Again, we computed the pRF preferred position changes and determined the linear fit between average pRF sizes and preferred position changes in each visual field map (Figure 3E, solid gray line [median fit]; Figure S3C, black data points). Again, pRF preferred position changes resulting from different gaze positions do not change with pRF size or visual field map. Thus, eye movements toward the attended targets induce a near-constant change in pRF preferred position regardless of pRF size and cannot explain the increase in pRF preferred position change with pRF size.

Furthermore, a relationship between preferred position change and pRF size was also observed within every visual field map, including V1, V2, and V3 (Pearson's correlation coefficient: $r = 0.18, 0.29$, and 0.32, respectively; all p values $< 10^{-6}$, corrected for upsampling), which again is inconsistent with the pRF size-independent effect of eye movements. As eye movements toward the targets effectively add a constant to all changes in preferred position, the maximum eye movement-induced addition to changes in pRF preferred position can be estimated from the change in pRF preferred position in the smallest pRFs in V1. We estimate this as 0.17° (95% confidence interval [c.i.]: 0.16°–0.19°). When measured outside the scanner, subjects' eye positions moved toward the attended target by a mean of 0.06° (95% c.i.: −0.04°–0.16°). Consequently, eye movements cannot explain the observed relationship between pRF size and pRF preferred position change either within or between visual field maps.
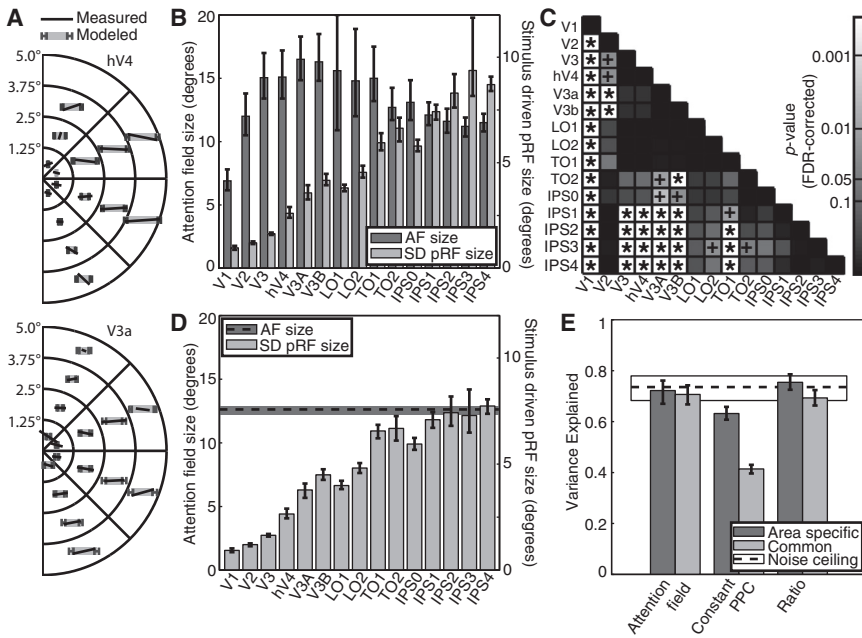
Figure 4. Attention Field Model Results

(A) Polar plots similar to Figure 3A, with the attention field (AF) model's predicted preferred position changes presented as light gray bars.

(B) Estimated AF size (dark gray bars) and stimulus-driven (SD) pRF size (light gray bars, right axis) from fitting an AF separately to each visual field map.

(C) Statistical significance of differences between AF size estimates. Plusses and asterisks mark significant differences between AFs for FDR correction and the more conservative Bonferroni correction, respectively.

(D) Estimated AF size (dashed line) and its 95% confidence intervals (dark gray area) and the estimated SD pRF size when a common AF is fit to all visual field maps.

(E) Variance explained grouped across all visual field maps for different models of the relationship between pRF size and preferred position change, with area-specific relationships between parameters (dark gray bars) or common parameters across all areas (light gray bars). The noise ceiling (dashed line with confidence intervals in white) represents the amount of preferred position change in one hemifield that can be explained by the preferred position changes in the complementary hemifield. All error bars are 95% confidence intervals.

Modulation by Gaussian Attention Fields Captures Changes in pRF Preferred Position

Gaussian interaction models of attention explain the pRF preferred position changes using two components: the stimulus-driven pRF and the attention field. The increasingly large change in pRF preferred position seen up the hierarchy could be explained by an increase in stimulus-driven pRF size and/or a decrease in attention field size up the hierarchy. To examine these two possibilities, we estimated both components from the measured pRF sizes and preferred position changes. For each visual field map and subject (and data grouped across subjects), we determined the attention field size that captured the most variance in the pRF preferred positions between conditions across all cortical locations. The attention field estimation procedure (see Experimental Procedures) used changes in binned pRF preferred positions and the mean pRF size across both attention conditions. To quantify the variance explained by attention field models, we used a 400-fold cross-validation procedure, fitting an attention field to randomly selected halves of the binned data and evaluating the fit on the complementary half of the data. To quantify the maximum explainable variance given the noise in the data, the noise ceiling (Machens et al., 2004; Mante et al., 2005), we measured how much of the binned pRF preferred position changes in one hemifield are captured by the preferred position changes in the complementary hemifield.

Figure 4A shows the measured pRF preferred position changes from Figure 3A overlaid with the change predicted by the best-fitting attention field for each visual field map. Figure 4B shows the estimated attention field size (AF size, dark gray bars) and stimulus-driven pRF size (SD pRF size, light gray bars) for each visual field map. The attention field sizes shown are the median and 95% confidence intervals of all fitting iterations; SD pRF

sizes are a representative value extracted at the middle of the stimulus eccentricity range (2.5°). SD pRF sizes increase systematically up the visual hierarchy, whereas the AF sizes are similar across many visual field maps, with no clear tendency to increase or decrease (perhaps increasing through early visual field maps and decreasing through later maps). As such, SD pRF sizes predict differences in pRF preferred position changes between visual field maps, while AF sizes do not. The expected decrease in AF size up the hierarchy (which would increasingly strongly attract pRF preferred positions) was not found. Yet, differences in AF size between visual field maps were found. To examine whether there are any significant differences between the estimated attention fields, we first determined the overlap between the bootstrap distributions seen in Figure 4B, i.e., the probability that the attention field estimates are different for any two visual field maps. Probabilities were corrected for multiple comparisons using FDR (Benjamini and Hochberg, 1995). This reveals that attention field estimates in V1 and IPS1–IPS4 are significantly smaller than those in other visual field maps (Figure 4C). However, when excluding pRFs whose centers lay outside the stimulus window in either attention condition, the IPS maps do not have significantly different attention field sizes.

At the same time, the lack of systematic changes in AF size up the hierarchy raises the possibility that AF size is similar in all visual field maps. To examine the amount of variance captured by differences in attention field size, we reran our attention field modeling procedure constrained to fit a common attention field to all visual field maps. A common attention field model (Figure 4D) explains 70.7% (95% c.i.: 66.7%–74.3%) of the variance in pRF preferred positions between conditions, which is not significantly different from attention fields for each individual visual field map, which explains 72.2% (95% c.i.: 67.0%–76.1%)
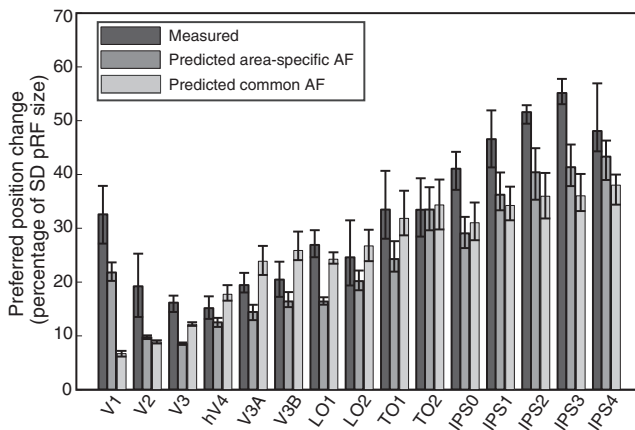
Figure 5. The Measured and Predicted pRF Position Change between the Two Conditions in Percentage of the Standard Deviation of the Stimulus-Driven pRF

Position changes predicted by both a separate (gray bars) and single (light gray bars) attention field model are shown. Data are obtained by linearly fitting eccentricity binned data points. All error bars represent 95% confidence intervals. Measured pRF position change is shown with dark gray bars.

straining the linear ratio model to use the same ratio for all visual field maps significantly reduces the variance explained by the model (75.5% versus 69.3%, p = 0.005). In sum, we prefer the attention field framework because (1) it explains the measured pRF preferred position changes at least as well as other relationships to pRF size and (2) it is an established framework proposed by previous studies.

Preferred Position Changes Are Less than 60% of the SD pRF Size

To examine the magnitudes of preferred position changes relative to the stimulus-driven pRF sizes, we divided pRF preferred position change between conditions by the estimated stimulus-driven pRF size (Figure 5, dark gray bars). This demonstrates that the preferred position changes are always less than 60% of the standard deviation of the stimulus-driven pRF. Furthermore, the ratio between preferred position changes and stimulus-driven pRF sizes increases up the hierarchy with pRF size. This increasing ratio is predicted by Gaussian interaction models, which predict a nonlinear relationship between stimulus-driven pRF size and pRF preferred position change (Equation 1).

DISCUSSION

We found that voluntary spatial attention attracts the preferred position of pRFs throughout the visual field toward the attended location, and this attraction increases up the visual hierarchy. The pRF attraction by attention was well captured by a model that conceptualizes attention's influence as a Gaussian attention field that multiplicatively interacts with stimulus-driven neural response properties. Both area-specific and common attention fields across the visual hierarchy do similarly well in explaining the observed changes in pRF preferred position. However, in V1 and IPS maps, where methodological limitations may affect estimates (see below), attention field estimates are likely smaller than other visual field maps. Nevertheless, this suggests that the increase in pRF preferred position changes up the hierarchy results primarily from the increase in the stimulus-driven pRF size rather than variation in attention field size.

Although we found some differences in attention field size estimates between visual field maps, the increase in variance explained by fitting area-specific attention fields appears negligible, with the possible exception of V1. The attention field size in V1 may be smaller than those in other visual field maps, which would indicate a larger influence of attention than in other visual field maps. However, we believe that the smaller attention field estimated in V1 may reflect methodological issues. Although eye movements cannot explain the pattern observed across the visual hierarchy, we suspect that due to V1's small pRF sizes in particular, small eye movements toward the attended location will lead to a larger underestimation of attention field size in earlier, compared to later, visual field maps. Indeed, this may obscure a decreasing attention field across the visual hierarchy. On the other hand, attention field sizes in IPS maps may also be underestimated. PRFs here are very large and so fall largely outside the stimulus area, which can lead to an underestimation of pRF size, which in turn leads to underestimating the attention

of the variance (p = 0.64, Figure 4E). When examining the variance explained for separate visual field maps (Figure S4), no robustly significant reductions in variance explained are found when fitting a common attention field (all p values > 0.46, FDR corrected for multiple comparisons). However, statistical results from V1 depend on the analysis parameters. Consequently, the data suggest that an area-specific attention field may fit better in V1, but we do not have the statistical power to strongly accept or refute this difference. Thus, the variance explained by area-specific attention fields is not significantly greater than that explained by a common attention field fit to all visual field maps, except perhaps in V1.

Although the Gaussian attention field model explains the observed preferred position changes well, we examined whether two other relationships can explain the observed changes. First, to test whether variance in pRF size affects preferred position changes, we fitted a constant preferred position change (constant PPC) in each visual field map irrespective of the pRF size. Gaussian attention field models, whether area specific or common, explain significantly more variance than constant preferred position change models (p = 0.001) (Figure 4E). Furthermore, the effect of pRF size on preferred position change within a visual field map is also supported by significant correlations between the pRF sizes and preferred position changes of cortical locations within every visual field map (all p values < 10^{-6}).

Second, to test whether the relationship between pRF size and preferred position change is described by a linear function (rather than the slightly nonlinear function predicted by Gaussian models), we tested models with the best-fitting linear relationship in each visual field map. This linear ratio model perform very similarly to Gaussian attention field models, as the Gaussian models closely approximate a linear function over small ranges of pRF sizes (Figure 4E). However, unlike for AF models, con-

field size needed to capture the observed preferred position changes. These methodological issues make it hard to reject the hypothesis that a common attention field affects all visual field maps. In sum, we find little evidence that attention fields change across the visual hierarchy. The main contributor to the increasingly large changes in pRF preferred position up the hierarchy seems to be increases in pRF size. While our results cannot exclude the possibility that the attention field changes up the visual hierarchy, our results raise the possibility that attention transforms the neural representation of the visual field similarly throughout the visual hierarchy.

The measured changes in preferred pRF position are consistent with single-neuron electrophysiological reports of receptive field preferred position changes in macaque V4 (Connor et al., 1996; Connor et al., 1997) and MT (Womelsdorf et al., 2006; Womelsdorf et al., 2008). These studies reported a change of the receptive field's center of mass toward the attended location, akin to a skew of a Gaussian profile. This description suggests not a change in the position of the receptive field, but rather a modification of its shape. Niebergall et al. (2011a) have suggested an expansion of the RF toward the attended location, although these results are also consistent with an attraction of the RF toward the attended location. Our pRF analysis does not include a shape parameter but describes the pRF as an isotropic Gaussian. Similarly, our model assumes a Gaussian attention field and stimulus-driven pRF. As the product of two Gaussians is also Gaussian, our measured pRF should be Gaussian. Consequently, a skew or expansion of the Gaussian profile toward the attended location would manifest itself as a change in the Gaussian envelope location. Furthermore, we do not believe we have the power to distinguish between a skew, an expansion, or a shift of the Gaussian profile. As such, movements of the Gaussian envelope, expansions, or skews of the profile (or a combination of these factors) are complementary interpretations of our pRF preferred position changes.

Gaussian models of spatial attention have much in common with several qualitative models of attention (such as spotlights and zoom lenses). Gaussian models are attractive here because they make quantitative predictions of effects on receptive fields. The idea of a spotlight that selects one location for attention but leaves the rest of the visual field representation unaffected does not describe the results we see. However, Gaussian models are a simplified description of attention's complex effects that provide a starting point for more detailed characterization of the attention field. For example, deviations from Gaussian response selectivity are known, both at single-neuron and population levels (Britten and Heuer, 1999; Heuer and Britten, 2002; Kay et al., 2013). Gaussians are a mathematical abstraction that captures the observed changes and allows straightforward modeling. Taking into account these deviations from Gaussian assumptions may capture more details of the measured response changes, as may more complex models, such as addition of a suppressive surround component surrounding the attention field Gaussian (Hopf et al., 2006) or a subsequent normalization stage (Reynolds and Heeger, 2009).

Attention fields may reflect the spatial distribution of synaptic gain, and a synaptic gain field would be consistent with our modeling framework. However, it is also possible that attention affects the amplitude, rather than the spread, of the gain distribution differently in different visual field maps. As the amplitude of the attention field is not included in the equations for Gaussian multiplication (Equations 3 and 4), our modeling framework suggests that attention field amplitude will not affect preferred position changes. As such, attention may have effects on neural response amplitude that are independent of preferred position changes. By focusing on preferred position changes, our framework cannot assess this factor.

Many models also address the issue of dividing attention between multiple targets (McMains and Somers, 2004; Niebergall et al., 2011b). We focus on one target here, and it is unclear how Gaussian models would extend to multiple targets. We also note that any stimulus used to map a receptive field will draw some exogenous attention itself, in effect making a further focus of attention. Here we focus on the difference between two attention locations with the same mapping stimulus, and so perhaps similar exogenous attention influences, which may reduce this problem.

While our model describes the measured changes in pRF preferred positions well, we do not propose a biological mechanism that may implement the attention field. Previous work based on electrophysiology suggests that modulation of response gain as a function of distance to the attended location underlies the influence of attention (Reynolds and Heeger, 2009). Indeed, an alternative explanation of our pRF preferred position changes could be that they reflect asymmetrical response gain across the neural population rather than a change in the preferred position of the neural RF. We suggest the latter because (1) we do not observe systematic asymmetrical response amplitude changes across different cortical locations, (2) similar preferred position changes have been observed in single neurons, and (3) this alternative explanation is similar to the mechanism behind attention effects proposed here (a normalized Gaussian distributed gain). Specifically, asymmetrical multiplicative gain changes at one level of the visual hierarchy will manifest itself as a change in preferred position of a receptive field that samples from that level (Maunsell and McAdams, 2000). Briefly, if locations within a neural RF's input get amplified more when nearer to the attended location, the location that produces the largest response, i.e., the preferred position, will change and specifically move toward the attended location. In other words, the RF profile will be reweighted toward the attended location.

We measure the aggregated receptive field of a population of neurons, not the receptive field of a single neuron (Dumoulin and Wandell, 2008; Smith et al., 2001). Both are measures of neural spatial selectivity, but at different spatial scales. The pRFs that we measure are affected by several factors that result in pRF size estimates being larger than those of neural RFs (Dumoulin and Wandell, 2008; Harvey and Dumoulin, 2011; Smith et al., 2001). As a consequence, our estimate of the attention field is affected similarly; this estimate will be larger than when derived from single-neuron response selectivities.

Links to neural activity must also take into account that fMRI measures hemodynamic changes reflecting neural activity indirectly. We believe differences in the hemodynamic response function (HRF) between our conditions cannot explain our

results. As we measure from the same neural population twice, we do not expect the HRF between our two conditions to vary, and we have used identical HRF models in analyzing data from both conditions. Last, although eye movements can affect pRF preferred position estimates, we demonstrate that eye movements cannot produce the results seen using eye movement measurement prior to scanning, measurements of the effects of different gaze positions on pRF preferred positions, and simulations of the effects of eye movements.

A Gaussian multiplication predicts attraction of the population's preferred position across the entire visual field representation. As our stimulus range is limited to 11° in diameter, and receptive fields found in higher visual field maps are very large, we cannot characterize response properties of populations whose stimulus-driven pRF does not include the attended locations in these higher visual field maps. However, we find effects of attention on preferred position in V1, V2, and V3 for populations whose stimulus-driven preferred position is located over 4 standard deviations away from either attended location. As such, attention affects neural response properties even when response selectivities are far from the attended location.

Furthermore, we found that attention affects pRF preferred positions across the entire visual hierarchy and that its effects increase up the hierarchy, consistent with previous reports of other modifications of neural responses by attention (Haenny and Schiller, 1988; Motter, 1993). While the effect of attention on preferred position becomes larger, the transformation of neural responses produced by attention (i.e., the attention field) is similar throughout the visual hierarchy. Consequently, as visuospatial response selectivity becomes broader, the effect of attention becomes stronger.

A constant attentional transformation implies that attention acts either uniformly across the visual hierarchy or early in the hierarchy, and the resulting effects on neural responses are inherited by later visual field maps. In line with the latter interpretation, previous studies have shown that attention can affect visual processing as early as the lateral geniculate nucleus (LGN) (McAlonan et al., 2008; O'Connor et al., 2002). Arguing against the hierarchical propagation of attentional effects in early visual processing, Buffalo et al. (2010) show that gain changes induced by attention occur with shorter latency in higher visual field maps.

In sum, our results extend previous neurophysiological findings that neural position preferences are affected by attention to a substantial part of the entire visual hierarchy. As such, visual field representations throughout the visual cortex depend on the top-down attentional state. In this global view of attention's effects, the increasing change in position preference up the hierarchy depends on the sharpness of position tuning, with attention producing similar transformations of visuospatial response selectivity throughout the visual cortex.

EXPERIMENTAL PROCEDURES

Subjects

Five male subjects participated in the main experiment (ages 25–39). Three male subjects participated in a control experiment with an alternative stimulus and task (ages 24–34; Figure S2), two of whom participated in the main experiment. All had normal or corrected to normal acuity. All experiments were cleared by the ethics committee of the University Medical Center Utrecht.

MRI Acquisition

T1-weighted anatomical MRI data were acquired on a Philips Achieva 3T scanner (Philips Medical Systems) at an isotropic resolution of 1 mm^3, with a field of view (FOV) of 288 × 288 × 175 mm. Repetition time (TR) was 9.958 ms, echo time (TE) was 4.59 ms, and flip angle was 8°. Functional T2*-weighted 2D echo planar images were acquired on a Phillips 7T scanner using a 32 channel head coil (Nova Medical) at a resolution of 2.0 × 2.0 × 2.0 mm, with an FOV of 190 × 190 × 50 mm. TR was 1,500 ms, TE was 25 ms, and flip angle was 80°. Functional scans were each 248 time frames (372 s) long, of which the first 8 time frames (12 s) were discarded to ensure a steady signal state. For each subject, 8 scans were acquired for each condition in alternating order, totaling 16 scans taken over two sessions.

Stimulus Presentation

Visual stimuli were back-projected onto a 15.0 × 7.9 cm screen inside the MRI bore. The subject viewed the display through mirrors. The distance from the subject's eyes to the display was 41 cm. Display resolution was 1,024 × 538 pixels. Stimuli were limited to a circular area filling the screen's vertical dimension, with any area outside this circle remaining at constant mean luminance. This stimulus circle had a radius of 5.5° visual angle.

Main Stimulus and Task

Visual stimuli were generated in MATLAB using PsychToolbox (Brainard, 1997; Pelli, 1997). Visual field mapping stimuli consisted of contrast defined bars of cardinal and diagonal orientations, stepping across the display perpendicular to each bar's orientations (Dumoulin and Wandell, 2008). These bars contained 100% contrast checkerboard patterns with alternating rows of checks moving in opposite directions parallel to the bar orientation. Checkerboard motion direction reversed at random intervals (minimum 4 s). The bar width and checkerboard spatial frequency were both 25% of the stimulus radius (1.37°). The bar moved across the stimulus circle in 20 equal steps, each 0.55°. As bar steps were synchronized with functional volume acquisitions, each bar pass took 20 TRs, 30 s. Bars stepped across the display in the four cardinal directions alternating with the four diagonal directions, totaling eight bar pass directions. After each cardinal bar pass, 30 s of mean-luminance (zero contrast) stimulus was displayed.

We added two pink (1/f) noise patterns in circular apertures 0.2° in radius, centered 5.95° to the left and right of the display center, on which the subjects performed a contrast discrimination task. The noise patches randomly changed orientation every 250 ms. The pattern increased contrast on 5% of orientation changes (randomly chosen and different between scan runs). The magnitude of the contrast increase was determined for each subject before scanning so that subjects found these increases difficult to detect but performed above chance. During one entire scanning run, the subjects performed the task on one location only, alternating between runs. Regardless of the task location, both noise patterns were always present and changed contrast independently. Contrast increment detection was reported by a button press within 1 s of the contrast increment. Note that a short interval between left and right contrast increments and a subsequent detection response can be taken as a detection on both noise patterns.

Before scanning sessions, subjects were trained to perform this task while suppressing eye movements. These eye tracking sessions used a highly accurate Eyelink II system (SG Research).

Contrast changes at the attended location were followed by responses more often than those at the unattended location (36.1% versus 12.8% detected, dependent samples t test, t(79) = 15.70, p < 10^{-6}). Task performance did not differ between left and right conditions (34.5% versus 37.8% detected).

Control Stimulus and Task

The above rapid serial visual presentation (RSVP) task design ensures that subjects attend continually, but subjects do not detect most events, complicating interpretation of their performance. This control experiment presented the same noise targets in a two-alternative forced-choice (2AFC) contrast discrimination task. Here, visual field mapping stimuli were only presented during the task intervals, ensuring local attention during visual stimulation. To ensure results generalize to different stimuli, the bar here revealed a stationary 1/f noise pattern. The outer 1° of the stimulus window and the bar aperture

were convolved with a cosine fade toward their edges. Each TR contained a 2AFC contrast discrimination trial. During each trial, apertures revealed a noise pattern for 166 ms followed by an interstimulus interval of 333 ms, then another noise pattern, then an intertrial interval of 833 ms. The second pattern presentation was flipped along the center of the bar's long axis. During bar presentations, the two circular peripheral noise patterns were shown, as well as a central circular noise pattern (0.12° radius). During the second presentation in each trial, these circular patterns were randomly rotated around their central point. To ensure the center pattern was distinguishable from the bar pattern during periods in which the bar moved behind the center pattern, a mean-luminance circle with a radius of 0.19° was drawn between the two patterns producing a separation of 0.07°. Other parameters of this stimulus were identical to the main stimulus.

All patterns increased in contrast on one randomly chosen presentation in each trail. The mean contrast across both presentations in each trial was identical for all patterns, trials, and subjects. During the intertrial interval, subjects reported which presentation at the attended noise pattern was higher in contrast. The contrast increase needed for 75% correct performance was estimated before scanning. Subjects detected more high-contrast patterns on the attended compared to the unattended side (78.3% versus 52.2%, dependent samples t test, t(47) = 20.5, p < 10^{-6}).Task performance did not differ between left and right conditions (78.7% versus 77.8% detected). For subjects that participated in the main and control experiments, task difficulty was similar in both.

Preprocessing
T1-weighted anatomical scans were resampled to 1 mm^3 resolution. The resulting anatomical image was automatically segmented using Freesurfer (Dale et al., 1999) and then hand-edited to minimize segmentation errors (Teo et al., 1997). The cortical surface was reconstructed at the gray/white matter border and rendered as a 3D surface (Wandell et al., 2000). fMRI analysis was performed in the mrVista software package for MATLAB (freely available at http://white.stanford.edu/software). Head movement artifacts between and within functional scans were measured and corrected for (Nestares and Heeger, 2000). Functional data were then averaged across scans, aligned to anatomical scans (Nestares and Heeger, 2000), and interpolated to the anatomical segmentation.

PRF Data Analysis
PRF sizes and positions were estimated from fMRI data and visual stimulus position time course. Each voxel's fMRI response was predicted using a 2D Gaussian pRF model. The pRF is described by four parameters: the preferred position (x and y parameters), the size (σ_1) of the location to which the voxel responds, and the size of a suppressive surround (σ_2) for each voxel (Zuiderbaan et al., 2012). A detailed description is given elsewhere (Dumoulin and Wandell, 2008). Briefly, the fMRI time course was predicted from the modeled pRF taking into account the stimulus time course and a canonical fMRI HRF (Friston et al., 1998; Glover, 1999; Worsley et al., 2002). Next, we estimated HRF parameters for each condition separately (Harvey and Dumoulin, 2011), then averaged these parameters across conditions. Finally, we re-estimated the pRF parameters for each voxel and condition using this averaged HRF.

PRF preferred positions from the two conditions were first averaged to give a single pRF model for definition of visual field maps. PRF polar angle and eccentricity maps were rendered onto an inflated cortical surface (Wandell et al., 2000), and the positions of visual field maps were determined and defined as regions of interest (ROIs). ROI borders were defined following reversals in the polar angle and eccentricity progressions (Sereno et al., 1995; Wandell et al., 2007), and ROIs were identified following published descriptions of their relative locations (Arcaro et al., 2011; DeYoe et al., 1996; Engel et al., 1997; Wandell et al., 2007). Voxels were excluded from further analysis if their pRF models explained less than 30% of response variance in either condition, if their average eccentricity across both conditions exceeded 5°, or if their mean fMRI signal intensity was lower than that in the surrounding cortex, which suggests a large influence of pial draining veins on the fMRI signal (Olman et al., 2012; Winawer et al., 2010).

Alignment and Averaging of fMRI Time Series
For each voxel and bar pass, the mean pRF preferred position across both conditions was used to determine the temporal interval between the bar

crossing the center of the visual field and the mean pRF preferred position (Dumoulin et al., 2014). Within each ROI, fMRI time series from each condition were then offset by this interval, aligning each voxel's maximum neural response to the same time point. All fMRI time series within the ROI were then averaged together, giving the mean fMRI time series and its SE as shown in Figure 2. When averaging within each visual field map, fMRI time series were first detrended, and each voxel's contribution was weighted by the variance explained by its pRF model. To reduce contributions of continuing hemodynamic responses from previous bar passes, we included only voxels with pRF preferred positions at 0.5°–2.5° eccentricity, pRF sizes below 4°, and at least 60% variance explained in both conditions.

PRF Parameter Analyses
To characterize pRF preferred position changes in each visual field map, we first averaged each voxel's preferred position estimates from both conditions. Assuming that effects of attention are symmetrical across both halves of the visual field representation, we then collapsed all measurements into one visual hemifield, giving a comparison between pRF parameters when the attended location was in the same hemifield as the pRF versus the identical position in the opposite hemifield. We then grouped all voxels by their averaged preferred position into 16 bins, each covering an eccentricity range of 1.25° visual angle and 45° polar angle. Finally, we determined the mean pRF preferred position in each bin and compared these between conditions (Figure S1 provides an overview of each of these steps for a single subject for V1–hV4).

We summarize the pRF parameters in each visual field map by assuming linear relationships between them, described by the equation

$$y = ax + b, \qquad\qquad \text{Equation 2}$$

where y represents a pRF parameter of interest and x represents eccentricity. a and b represent the slope and intercept, respectively, and were estimated by minimizing the residual sum of squares (RSS) to the mean of data in eccentricity ranges sampled at 0.5° steps. We obtained 95% confidence intervals of the fit by bootstrapping (1,000 iterations). A representative value for each visual field map was derived from the fit, and its confidence intervals were evaluated at an eccentricity (x) of 2.5° (Equation 2). We used this procedure to summarize each visual field map's pRF size, preferred position change, their ratios, and the preferred position change along the cortical surface (Figures 3B–3F, 5, and S3).

To estimate the preferred position change along the cortical surface, we multiplied each voxel's pRF preferred position change by its cortical magnification factor (Harvey and Dumoulin, 2011).

Attention Field Modeling
Attention effects on pRF position preferences were modeled by the multiplication of two Gaussians (μ_1,σ_1 and μ_2,σ_2), which produces a third Gaussian (μ_3,σ_3). The properties of this third Gaussian are derived from two equations. First, the standard deviation (size) of their product (σ_3) is given by the equation

$$\sigma_3^2 = \frac{\sigma_1^2 \sigma_2^2}{\sigma_1^2 + \sigma_2^2} \qquad\qquad \text{Equation 3}$$

Second, the mean (position, μ_3) is given by the equation

$$\mu_3 = \frac{\mu_1 \sigma_2^2 + \mu_2 \sigma_1^2}{\sigma_2^2 + \sigma_1^2}. \qquad\qquad \text{Equation 4}$$

In our experiment, we measure the pRF preferred position and size twice, each being the product of the same stimulus-driven pRF (SD pRF) and two attention fields (AF). Critically, we assume the attention fields have different but known positions (the target locations) and the same size (as they have the same task and eccentricity). Therefore, we need to estimate three parameters: (i) SD pRF position (μ_{SD}), (ii) SD pRF size (σ_{SD}), and (iii) AF size (σ_{AF}). These parameters depend on each other such that when one is known the others can be derived by Equations 3 and 4.

Here, we estimate these three parameters using a forward model. First, we systematically vary AF size and SD pRF position. For each value of AF size, we derive the SD pRF size, which is given by the equation

$$\sigma_{SD}^2 = \frac{\sigma_{AF}^2}{\frac{\sigma_{AF}^2}{\sigma_{RF}^2} - 1}, \qquad \text{Equation 5}$$

where σ_{RF} is the mean pRF size across both attention conditions. This equation is derived from Equation 3. We use mean pRF size to give a more reliable estimate of pRF size. Next, we predict pRF position (μ_{RFp}) given these specific parameters:

$$\mu_{RFp} = \frac{(\mu_{SD}\sigma_{AF}^2) + (\mu_{AF}\sigma_{SD}^2)}{\sigma_{SD}^2 + \sigma_{AF}^2}. \qquad \text{Equation 6}$$

We do this for each attention field position, yielding two different pRF position predictions (μ_{RFLp} and μ_{RFRp},). We sum the squared differences between these predicted positions and the measured preferred positions across all bins (i), giving the residual sum of squared errors, SSR:

$$SSR = \sum_{i=1}^{n} \left(\left| \mu_{RFLp(i)} - \mu_{RFL(i)} \right| + \left| \mu_{RFRp(i)} - \mu_{RFR(i)} \right| \right)^2, \qquad \text{Equation 7}$$

where μ_{RFL} and μ_{RFR} are measured pRF preferred positions in the left and right conditions, respectively. We also compute the sum across bins of squared differences between measured preferred positions, the total sum of squared (SST) difference between conditions:

$$SST = \sum_{i=1}^{n} \left(\mu_{RFL(i)} + \mu_{RFR(i)} \right)^2. \qquad \text{Equation 8}$$

From these values, we can compute the proportion of the difference between conditions that is explained by this putative attention field, the variance explained:

$$VE = 1 - \frac{SSR}{SST}. \qquad \text{Equation 9}$$

We chose the SD pRF position (μ_{SD}) and AF size (σ_{AF}) that maximize the variance explained (Equation 9). In the case of a single attention field fit across all areas, we maximized the summed variance explained across all visual field maps. We used a 400-fold cross-validation procedure, which estimates the best fitting attention field for randomly selected halves of the binned data and evaluates the fit by determining the variance explained in the complementary half.

We compute SSR using only horizontal positions because the model predicts no vertical position change between μ_{RFLp} and μ_{RFRp} when the two attention fields have the same vertical position (Equation 6). However, the VE reported in Figures 4 and S4 also takes into account variations in vertical positions. All model outcomes are the median and 95% confidence intervals of all model-fitting iterations. Reported p values describe the proportion of values from one distribution that exceed the values in the second distribution.

Finally, Equation 1 is the difference between two versions of Equation 6 when (1) the two attention field centers are the same distance from fixation in opposite direction, like our attended targets ($\mu_{AFL} = -\mu_{AFR}$), (2) both attention fields have the same size (σ_{AF}), and (3) the stimulus-driven pRF (μ_{SD} and σ_{SD}) is the same for both conditions. Our experimental design and some straightforward assumptions ensure that these conditions are met in our experiment. The preferred position change between conditions ($\mu_{RFL} - \mu_{RFR}$) is given by the following equation, which simplifies to give Equation 1:

$$\mu_{RFR} - \mu_{RFL} = \left[\frac{(\mu_{SD}\sigma_{AF}^2) + (\mu_{AFR}\sigma_{SD}^2)}{\sigma_{SD}^2 + \sigma_{AF}^2} \right] - \left[\frac{(\mu_{SD}\sigma_{AF}^2) + (-\mu_{AFL}\sigma_{SD}^2)}{\sigma_{SD}^2 + \sigma_{AF}^2} \right]$$
$$\text{Equation 10}$$

SUPPLEMENTAL INFORMATION

Supplemental Information includes four figures and two movies and can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2014.08.047.

ACKNOWLEDGMENTS

We thank S. te Pas for helpful comments on drafts of this manuscript. This work is supported by the Netherlands Organisation for Scientific Research

REFERENCES

Amano, K., Wandell, B.A., and Dumoulin, S.O. (2009). Visual field maps, population receptive field sizes, and visual field coverage in the human MT+ complex. J. Neurophysiol. *102*, 2704–2718.

Anton-Erxleben, K., and Carrasco, M. (2013). Attentional enhancement of spatial resolution: linking behavioural and neurophysiological evidence. Nat. Rev. Neurosci. *14*, 188–200.

Arcaro, M.J., Pinsk, M.A., Li, X., and Kastner, S. (2011). Visuotopic organization of macaque posterior parietal cortex: a functional magnetic resonance imaging study. J. Neurosci. *31*, 2064–2078.

Benjamini, Y., and Hochberg, Y. (1995). Controlling the False Discovery Rate: a Practical and Powerful Approach to Multiple Testing. J. R. Statist. Soc. B *57*, 289–300.

Brainard, D.H. (1997). The Psychophysics Toolbox. Spat. Vis. *10*, 433–436.

Britten, K.H., and Heuer, H.W. (1999). Spatial summation in the receptive fields of MT neurons. J. Neurosci. *19*, 5074–5084.

Buffalo, E.A., Fries, P., Landman, R., Liang, H., and Desimone, R. (2010). A backward progression of attentional effects in the ventral stream. Proc. Natl. Acad. Sci. USA *107*, 361–365.

Connor, C.E., Gallant, J.L., Preddie, D.C., and Van Essen, D.C. (1996). Responses in area V4 depend on the spatial relationship between stimulus and attention. J. Neurophysiol. *75*, 1306–1308.

Connor, C.E., Preddie, D.C., Gallant, J.L., and Van Essen, D.C. (1997). Spatial attention effects in macaque area V4. J. Neurosci. *17*, 3201–3214.

Cook, E.P., and Maunsell, J.H.R. (2002). Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey. J. Neurosci. *22*, 1994–2004.

Dale, A.M., Fischl, B., and Sereno, M.I. (1999). Cortical surface-based analysis. I. Segmentation and surface reconstruction. Neuroimage *9*, 179–194.

DeYoe, E.A., Carman, G.J., Bandettini, P., Glickman, S., Wieser, J., Cox, R., Miller, D., and Neitz, J. (1996). Mapping striate and extrastriate visual areas in human cerebral cortex. Proc. Natl. Acad. Sci. USA *93*, 2382–2386.

Dumoulin, S.O., and Wandell, B.A. (2008). Population receptive field estimates in human visual cortex. Neuroimage *39*, 647–660.

Dumoulin, S.O., Hess, R.F., May, K.A., Harvey, B.M., Rokers, B., and Barendregt, M. (2014). Contour extracting networks in early extrastriate cortex. J. Vis. *14*, 18.

Engel, S.A., Glover, G.H., and Wandell, B.A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. Cereb. Cortex *7*, 181–192.

Friston, K.J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M.D., and Turner, R. (1998). Event-related fMRI: characterizing differential responses. Neuroimage *7*, 30–40.

Glover, G.H. (1999). Deconvolution of impulse response in event-related BOLD fMRI. Neuroimage *9*, 416–429.

Haenny, P.E., and Schiller, P.H. (1988). State dependent activity in monkey visual cortex. I. Single cell activity in V1 and V4 on visual tasks. Exp. Brain Res. *69*, 225–244.

Harvey, B.M., and Dumoulin, S.O. (2011). The relationship between cortical magnification factor and population receptive field size in human visual cortex: constancies in cortical architecture. J. Neurosci. *31*, 13604–13612.

Heuer, H.W., and Britten, K.H. (2002). Contrast dependence of response normalization in area MT of the rhesus macaque. J. Neurophysiol. *88*, 3398–3408.

Hopf, J.M., Boehler, C.N., Luck, S.J., Tsotsos, J.K., Heinze, H.J., and Schoenfeld, M.A. (2006). Direct neurophysiological evidence for spatial suppression surrounding the focus of attention in vision. Proc. Natl. Acad. Sci. USA *103*, 1053–1058.

Kay, K.N., Winawer, J., Mezer, A., and Wandell, B.A. (2013). Compressive spatial summation in human visual cortex. J. Neurophysiol. *110*, 481–494.

Levin, N., Dumoulin, S.O., Winawer, J., Dougherty, R.F., and Wandell, B.A. (2010). Cortical maps and white matter tracts following long period of visual deprivation and retinal image restoration. Neuron *65*, 21–31.

Machens, C.K., Wehr, M.S., and Zador, A.M. (2004). Linearity of cortical receptive fields measured with natural sounds. J. Neurosci. *24*, 1089–1100.

Mante, V., Frazor, R.A., Bonin, V., Geisler, W.S., and Carandini, M. (2005). Independence of luminance and contrast in natural scenes and in the early visual system. Nat. Neurosci. *8*, 1690–1697.

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

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

Maunsell, J.H.R., and McAdams, C.J. (2000). Effects of attention on neural response properties in visual cerebral cortex. In The New Cognitive Neurosciences, M.S. Gazzaniga, ed. (Cambridge, MA: MIT Press), pp. 315–324.

McAdams, C.J., and Maunsell, J.H.R. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. J. Neurosci. *19*, 431–441.

McAlonan, K., Cavanaugh, J., and Wurtz, R.H. (2008). Guarding the gateway to cortex with attention in visual thalamus. Nature *456*, 391–394.

McMains, S.A., and Somers, D.C. (2004). Multiple spotlights of attentional selection in human visual cortex. Neuron *42*, 677–686.

Montijn, J.S., Klink, P.C., and van Wezel, R.J. (2012). Divisive normalization and neuronal oscillations in a single hierarchical framework of selective visual attention. Front Neural Circuits *6*, 22.

Motter, B.C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. J. Neurophysiol. *70*, 909–919.

Nestares, O., and Heeger, D.J. (2000). Robust multiresolution alignment of MRI brain volumes. Magn. Reson. Med. *43*, 705–715.

Niebergall, R., Khayat, P.S., Treue, S., and Martinez-Trujillo, J.C. (2011a). Expansion of MT neurons excitatory receptive fields during covert attentive tracking. J. Neurosci. *31*, 15499–15510.

Niebergall, R., Khayat, P.S., Treue, S., and Martinez-Trujillo, J.C. (2011b). Multifocal attention filters targets from distracters within and beyond primate MT neurons' receptive field boundaries. Neuron *72*, 1067–1079.

O'Connor, D.H., Fukui, M.M., Pinsk, M.A., and Kastner, S. (2002). Attention modulates responses in the human lateral geniculate nucleus. Nat. Neurosci. *5*, 1203–1209.

Olman, C.A., Harel, N., Feinberg, D.A., He, S., Zhang, P., Ugurbil, K., and Yacoub, E. (2012). Layer-specific fMRI reflects different neuronal computations at different depths in human V1. PLoS ONE *7*, e32536.

Pelli, D.G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. Spat. Vis. *10*, 437–442.

Posner, M.I., and Gilbert, C.D. (1999). Attention and primary visual cortex. Proc. Natl. Acad. Sci. USA *96*, 2585–2587.

Reynolds, J.H., and Heeger, D.J. (2009). The normalization model of attention. Neuron *61*, 168–185.

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.H. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. Science *268*, 889–893.

Smith, A.T., Singh, K.D., Williams, A.L., and Greenlee, M.W. (2001). Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. Cereb. Cortex *11*, 1182–1190.

Teo, P.C., Sapiro, G., and Wandell, B.A. (1997). Creating connected representations of cortical gray matter for functional MRI visualization. IEEE Trans. Med. Imaging *16*, 852–863.

Treue, S., and Maunsell, J.H.R. (1999). Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. J. Neurosci. *19*, 7591–7602.

Wandell, B.A., Chial, S., and Backus, B.T. (2000). Visualization and measurement of the cortical surface. J. Cogn. Neurosci. *12*, 739–752.

Wandell, B.A., Dumoulin, S.O., and Brewer, A.A. (2007). Visual field maps in human cortex. Neuron *56*, 366–383.

Winawer, J., Horiguchi, H., Sayres, R.A., Amano, K., and Wandell, B.A. (2010). Mapping hV4 and ventral occipital cortex: the venous eclipse. J. Vis. *10*, 1–22.

Womelsdorf, T., Anton-Erxleben, K., Pieper, F., and Treue, S. (2006). Dynamic shifts of visual receptive fields in cortical area MT by spatial attention. Nat. Neurosci. *9*, 1156–1160.

Womelsdorf, T., Anton-Erxleben, K., and Treue, S. (2008). Receptive field shift and shrinkage in macaque middle temporal area through attentional gain modulation. J. Neurosci. *28*, 8934–8944.

Worsley, K.J., Liao, C.H., Aston, J., Petre, V., Duncan, G.H., Morales, F., and Evans, A.C. (2002). A general statistical analysis for fMRI data. Neuroimage *15*, 1–15.

Zuiderbaan, W., Harvey, B.M., and Dumoulin, S.O. (2012). Modeling center-surround configurations in population receptive fields using fMRI. J. Vis. *12*, 10.