Zooming in with attention

Neural and perceptual consequences of attention

Barrie P. Klein

Zooming in with attention

Neural and perceptual consequences of attention

Inzoomen met aandacht

Neurale en perceptuele gevolgen van aandacht
(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht op gezag van de rector magnificus, prof.dr. G.J. van der Zwaan, ingevolge het besluit van het college voor promoties in het openbaar te verdedigen op vrijdag 8 juni 2018 des middags te 2.30 uur

door

Barend Pieter Klein

geboren op 4 september 1985 te Goudswaard

Promotoren:	Prof. dr. S.O. Dumoulin
	Prof. dr. S.F. te Pas
Copromotor:	Dr. C.L. E. Paffen

Dit proefschrift werd (mede) mogelijk gemaakt met financiële steun van de Nederlandse Wetenschaps Organisatie subsidie 406-13-141 aan S. O. Dumoulin en B.P. Klein

voor

Neelie & Juliëtte

Contents

Chapter 1

General introduction

The visual processing bottleneck

Vision is the dominant sense in humans. Most aspects of our society require people to have full functioning vision, from interacting with simple technical devices to interacting with other people. In our brain, about 25% of the cerebral cortex is dedicated to processing visual information, which is more than for any other sense (Van Essen, 2003). Nevertheless, the brain's capacity to process visual information is severely limited by several factors. For example, it is estimated that less than 1% of all neurons in our brain can be active at the same time due to the metabolic costs of brain activity (Lennie, 2003). As a result, for efficient visual processing we need to allocate the limited visual processing capacity dynamically to the information that is most relevant to us at any given time. Visual attention is a mechanism through which we can select visual information to prioritize its processing (Treisman, 1960; Johnston and Heinz, 1978; Carrasco, 2006).

Attentional selection

Typically, two ways of attentional selection are distinguished. In the first way, visual information draws our attention to itself. For example, information that is very distinct from the background or that changes unexpectedly tends to attract attention. Because we have little or no control over this process, we refer to this type of attention as **involuntary attention**. In the second way, we *do* have voluntary control over the information at which our attentional focus is directed. As an example, imagine that you are talking with a friend and want to know what time it is. Of course, you could simply look directly towards the clock right next to your friend's head. However, she might think you are bored with the conversation as soon as she notices that you are looking at the clock next to her. As you want to avoid this you try to read the time on the clock, while you keep looking at your friend's face. In this case, you are using voluntary control over your attentional focus, and even separate it from the location you are looking at. Not surprisingly, this type of attention is referred to as **voluntary attention**.

When we select visual information using attention, we can base this selection on different characteristics of visual information. For example, you can attend a specific location in the visual field. In this case, we are referring to **spatial attention**. Alternatively, you could try to pay attention to a specific line orientation, a color or to a numerosity, i.e. the amount of items in a set. We refer to this as **feature attention**.

Attentional selection as a spotlight or magnifying glass

Regardless of how we select visual information, attentional selection enhances its processing leading to improved performance by human observers. For example, attention speeds up reaction times, improves discriminability, and increases spatial resolution and contrast sensitivity (Treisman, 1960; Johnston and Dark, 1986; Carrasco, 2011; Anton-Erxleben and Carrasco, 2013).

Conceptually, the process of attentional selection can be compared to a spotlight that highlights the attended information or as a magnifying glass that magnifies the attended information. (Eriksen and Hoffman, 1972; Posner, 1980; LaBerge, 1983; Eriksen and Yeh, 1985; Eriksen and St. James, 1986). To illustrate these two concepts, consider a typical scene of a busy street like the one in Figure 1.1A. If you were to attend the taxi parked along the street, the spotlight concept is best illustrated by an increase in the brightness of the attended object (Figure 1.1B), whereas the magnifying glass can be best described by enlarging the representation of the attended object (Figure 1.1C).

Figure 1.1 Attention as a spotlight and magnifying glass. **A** A typical visual scene of a busy street. **B** Attention as a spotlight. Within this concept, attending the taxi (blue-red dashed circle) can be illustrated by an increase in the brightness of the attended information. **C** When considered as a magnifying glass, attention now enlarges the representation of the attended information.

The spotlight and magnifying glass of attention are conceptual descriptions of the effects of attention on visual processing. Nevertheless, these concepts are also useful to describe how attention can affect neural response properties. In this case, the spotlight concept would resemble an increases in neural responses to the attended information, whereas the magnifying glass implies an enlarged neural

representation of the attended information (Moran and Desimone, 1985; Spitzer et al., 1988; Corbetta et al., 1990; Kastner et al., 1998; Brefczynski and DeYoe, 1999).

A powerful way to understand how attention affects neural response properties is by using computational models. One such computational model is the **attention field model** (Womelsdorf et al., 2008; Reynolds and Heeger, 2009). The attention field model includes several components. The two most important of these are the attention field, which represents the influence of attention, and a component that represents a neuron's selectivity for a visual characteristic, like spatial position or numerosity. Importantly, the attention field model predicts that attention both increases neural responses and enlarges the neural representation of attended information. As such, according to the attention field model, the spotlight and magnifying glass of attention are two sides of the same coin. However, although the spotlight concept is supported by many empirical studies (Moran and Desimone, 1985; Spitzer et al., 1988; Corbetta et al., 1990; Kastner et al., 1998; Brefczynski and DeYoe, 1999), the empirical evidence supporting the concept of attention as a magnifying glass in humans is currently limited (Connor et al., 1996, 1997; Womelsdorf et al., 2006; David et al., 2008). **In this thesis we will investigate if attention does indeed acts on visual information as a magnifying glass.**

Attention as a magnifying glass

In order to investigate if attention does indeed act like a magnifying glass we build on the attention field model. According to the attention field model, attention can act like a magnifying glass by modulating a neuron's selectivity for visual characteristics. Specifically, the attention field model predicts that attention increases a neuron's sensitivity for the attended information at the expense of it's sensitivity for information that is not attended. When this effect is applied to a large number of neurons, more neurons will respond to the attended information and fewer will respond to the unattended information. This will produce an enlarged representation of the attended information.

Crucially, in order to act like a magnifying glass, attention has to modulate the selectivity of all neurons that are selective for the attended visual characteristic. Thus, if a spatial location is attended, all neurons that are spatially selective should be modulated. Likewise, if a certain numerosity is selected, the sensitivity of all numerosity selective neurons should change. However, studies demonstrating that attention modulates neural selectivity have measured responses in non-human primates from a limited part of visual cortex only (Connor et al., 1996, 1997; Womelsdorf et al., 2006; David et al., 2008). As such, these measurements are insufficient to determine whether attention indeed acts like a magnifying glass in

humans. Furthermore, it is unclear what mechanism would implement attentional magnification and what its consequences are for human perception.

In order to answer these questions, we utilized two major methodological advances in neuroimaging. First, novel neuroimaging analyses incorporate biological models of visual processing. This allowed us to measure the selectivity of neural populations to various visual characteristics across large parts of visual cortex in humans. In Chapter 2, we combined the attention field model with biological models of selectivity for spatial position. In this way, we assessed if spatial attention indeed acts like a magnifying glass on spatial position across visual cortex in humans. The second major advance we utilized is high-resolution fMRI. This technique allowed us to measure brain responses in ultra-high detail. In Chapter 3 we combined high-resolution fMRI with an attention field model to examine the mechanisms that implement attentional magnification in the spatial domain. In Chapter 4, we combined the attention field model with psychophysical measures to relate attentional magnification in visuospatial processing to changes in human perception. Finally, in Chapter 5, we modified the biological models of spatial selectivity used in Chapter 2 to capture selectivity for numerosity. Doing so, we paved the way to assess if attention can act as a magnifying glass on the processing of stimulus characteristics other than spatial position as well. In the following sections, we explain our approach in more detail.

The organization of spatial selectivity in visual processing

The selectivity of a neuron for a visual characteristic is a fundamental building block of visual processing. A very common neural selectivity is selectivity for spatial position. This spatial selectivity is typically referred to as a **receptive field**, the region in the visual field from which the neuron receives input. Within visual cortex, neurons are organized according to the position of their receptive fields. Neurons with neighboring receptive fields are located at neighboring cortical positions. This organization is called **retinotopic organization**, which means that the neural organization follows the physical layout of the retinae. As such, our visual system forms a neatly organized representation of the visual field with neighboring neurons processing
information from neighboring positions in the visual field. Such a representation is called a **visual field map** (Figure 1.2). Visual cortex contains not just one visual field map, but multiple, possibly up to 25 (Wandell et al., 2007). These visual field maps are, to some extent, hierarchically organized, meaning that the output of one visual field map forms the input of the next.

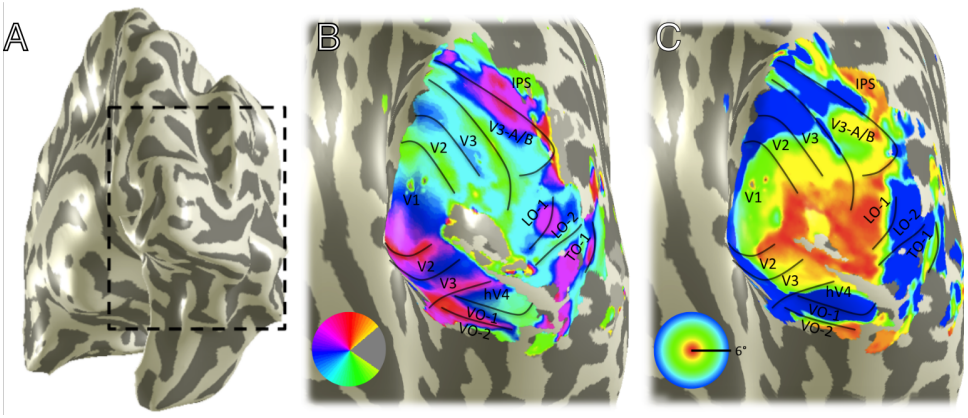

Figure 1.2 Retinotopic organization and visual field maps. **A** 3-Dimensional rendering of the right cerebral cortex. The area within the dashed lines roughly corresponds to the visual cortex. **B** Close-up of the dashed area in A overlaid with a color map. The colors on the rendering correspond to the colors of the circular inset and indicate from which part of the visual field neurons receive input. For example, red colors on the rendering mean that neurons here receive input from the top of visual field. This particular map is called a polar angle map. **C** Similar to B, but now overlaid with an eccentricity map. Colors now indicate the eccentric position (distance from the center of the visual field, which is the point of fixation) from which neurons process information. Notice how both polar angle and eccentricity are organized across the cortical surface and form multiple visual field maps, denoted by V1, V2 etc. Adapted from Wandell et al. (2015).

Attentional magnification in humans

We know from neurophysiological studies in non-human primates that spatial attention modulates the spatial selectivity of a neuron, i.e. it shapes their receptive field. However, neurophysiological studies in non-human primates typically limit their recordings to a few neurons within a single visual field map. As a result, the attentional modulation of receptive fields is only described in two visual field maps (Connor et al., 1996, 1997; Womelsdorf et al., 2006; David et al., 2008). However, in order to act like a magnifying glass, spatial attention has to modulate all receptive fields in every visual field map, as is also predicted by the attention field model. In order to test this, we need a measurement technique that permits simultaneous recordings from multiple visual field maps in humans. Functional Magnetic Resonance Imaging (fMRI) is such a technique. With fMRI we can derive brain

activity from changes in the magnetic properties of blood, enabling us to measure responses from small cubes of brain tissue (voxels), typically no more than 2mm isotropic in size. By exploiting the retinotopic organization of visual field maps, a recently developed analyses of fMRI data allows us to measure the aggregated receptive field of all neurons within a voxel. We call this aggregated receptive field the **population receptive field** (pRF)(Dumoulin and Wandell, 2008). Using fMRI and pRF analysis we can measure the pRF of fMRI voxels across large extents of visual cortex - typically ~100,000 voxels - simultaneously.

When applied to pRFs, the attention field model (Womelsdorf et al., 2008; Reynolds and Heeger, 2009) predicts that pRFs are attracted towards the location of spatial attention (Figure 1.3 A, B and C). With pRF attraction, a neural population's spatial selectivity profile is attracted towards the attended location and the population becomes more sensitive for the attended location and less sensitive for non-attended locations. As a consequence, more neurons will process the attended information, at the expense of information located away from the focus of attention. As such, pRF attraction will enlarge the neural representation of the attended location and thus produces attentional magnification (Figure 1.3 D and E). In Chapter 2, we examine pRF attraction in humans and test several predictions by the attention field model for pRF attraction that are crucial if it is to lead to attentional magnification, namely that pRF attraction should be present both across the entire visual field and across the entire visual cortex. Furthermore, we also examine if pRF attraction, and thus attentional magnification, is related to the size of the pRFs that are attracted. Specifically, larger pRFs are predicted to be attracted more and should lead to more attentional magnification.

Figure 1.3 pRF attraction and attentional magnification. **A** Same image as in figure 1.2, but now overlaid with hypothetical pRFs (white circles). **B** The same image as in A, but now we are attending the taxi on the side of the road. The location of spatial attention is marked by the red dot. **C** Spatial attention attracts the pRFs towards its location, i.e. the taxi. At this point, an increased amount of pRFs cover the attended location resulting in more neurons responding to the attended information. **D** The representation of visual scene in A when attention is not yet directed to the taxi (blue circle). **E** When spatial attention is directed to the taxi, pRFs are attracted towards this location (see panel C) which enlarges the representation of the attended information (blue-red dashed circle). As such spatial attention can act like a magnifying glass by attracting pRFs towards its location.

Mechanisms underlying attentional magnification

Although we found in Chapter 2 that spatial attention indeed acts like a magnifying glass on spatial postion by attracting pRFs towards its location across visual cortex, the biological mechanism that implements this attentional magnification in humans is still unclear. Broadly, two types of mechanisms can implement this magnification, (1) a feed forward and (2) a feedback mechanism (Olshausen et al., 1993; Compte and Wang, 2006; Bobier et al., 2014). In Chapter 3 we assess which of these mechanisms implement the attentional magnification in humans. To this end, we exploit another organizational principle of visual field maps, namely **laminar organization**.

As mentioned above, the visual field maps in visual cortex constitute a hierarchical system. Output from one visual field map forms the input of the next visual field map. We refer to this as feed forward processing. However, visual processing is not 'one way traffic' from lower to higher visual field maps. Higher visual field maps send feedback to lower visual areas which can shape the processing of visual information. Importantly, feed forward and feedback processing remain largely segregated in different layers within visual field maps. Typically, visual cortex is about 2 - 2.5 mm thick and feed forward processing is mainly limited to central cortical layers. In contrast, feedback processing is more located near the edges of the cortex i.e. near the gray/white matter boundary and near the pial surface (Felleman and Van Essen, 1991; Dumoulin et al., 2017). This organizational principle is commonly referred to as laminar organization, as neurons are organized in layers (laminae) running parallel with the cortical surface (Figure 1.4).

The laminar organization is well studied in primary visual cortex (V1), the first stage of cortical visual processing. When we apply the attention field model to the laminar organization of V1, the predicted profile of attentional magnification across cortical thickness is different between a feed forward and feedback implementation. In Chapter 3, we exploit recent advances in fMRI acquisition techniques to measure attentional magnification at an ultra-high resolution (0.7mm isotropic). This way, we can measure profiles of magnification across cortical thickness in V1 and compare them to the predicted profiles.

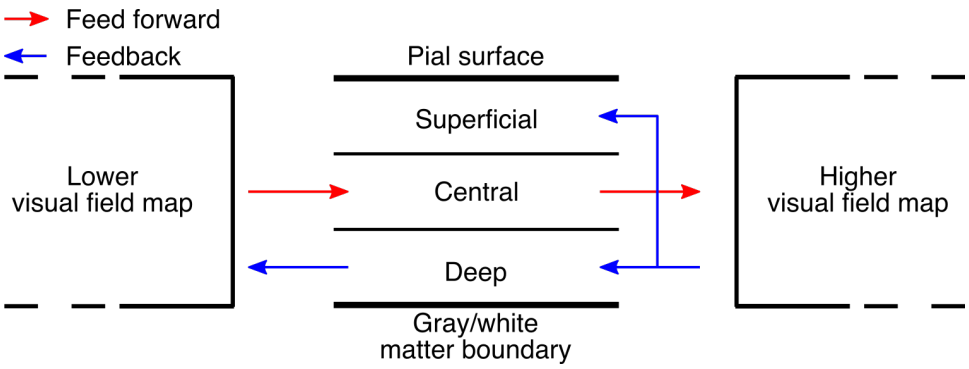

Figure 1.4 Schematic representation of the laminar organization in visual cortex. Feed forward input from lower visual field maps (red arrows) arrives in central portions of the visual cortex. In contrast, feedback from higher visual terminates in superficial and deep cortical portions of visual cortex.

Consequences of attentional magnification for human perception

In Chapters 2 and 3, we describe how spatial attention attracts pRFs in human visual cortex and magnifies the representation of the attended location. We expect that these changes in visual processing have consequences for human perception as well. More specifically, as the representation of the attended location is enlarged by attentional magnification, we expect that human observers perceive the attended location to be larger than it actually is. Consequently, we expect that objects surrounding the attended location are perceived to be further apart than they really are. Moreover, as the amount of magnification is related to the size of the pRFs that underlie the magnification, we expect that the resulting perceptual bias depends on the size of the underlying pRFs as well. In Chapter 4 we test these predictions. We exploit two other principles of visual processing to manipulate the pRF size underlying the perceptual bias. First, the size of pRFs increases with eccentricity, i.e. distance from the center of the visual field (Hubel and Wiesel, 1962; Van Essen et al., 1984; Dumoulin and Wandell, 2008; Harvey and Dumoulin, 2011). Therefore, we expect to measure a larger perceptual bias when objects are placed further away from the center of the visual field. Second, pRF size is related to spatial frequency selectivity. Larger pRFs are selective for lower spatial frequencies (Jones and Palmer, 1987). Consequently, we expect a larger perceptual bias when objects are composed of lower spatial frequencies.

Beyond spatial position

In Chapters 2, 3 and 4, we examine the magnifying glass concept using spatial attention and by measuring it's effects on neural selectivity for spatial position. However, we can direct our attention to much more visual characteristics than just spatial position. For example, we could also attend a certain numerosity.

Numerosity refers to the amount of items in a set. The accuracy with which we can perceive numerosity depends on the amount of items in a set: numerosity perception is most accurate for small amounts and becomes increasingly less accurate for larger amounts (Nieder and Miller, 2003, 2004; Piazza et al., 2004). For example, you can accurately perceive the amount of apples on a table when there are only three of them, but not when there are twenty apples. It is important to realize that the capacity to perceive numerosity is distinct from counting and other mathematical abilities. Animals, infants, and tribes with no numerical language can perceive numerosity (Dehaene et al., 1999, 2008; Brannon and Terrace, 2016). Importantly, some aspects of numerosity processing are very similar to those of primary sensory perception, like vision. Therefore, it has been referred to as a 'number sense' (Burr and Ross, 2008). Because of our innate ability to perceive numerosity and the similarities between numerosity processing and visual processing, we expect that the brain mechanisms underlying numerosity processing are similar to the mechanisms underlying visuospatial processing discussed above. And perhaps, attention may operate like a magnifying glass on the processing of numerosity as well.

In order to examine if attention acts like a magnifying glass on numerosity selectivity using our approach, we need biological models that describe the selectivity for numerosity first. In Chapter 5 we extend the pRF model we have used in earlier chapters to numerosity. We examine numerosity selective neural populations in human visual cortex and we assess how well their organization resembles the organization of spatially selective neural populations. Doing so, we pave the way to apply our approach in Chapter 2, 3 and 4 to numerosity processing and extend the concept of a magnifying glass to visual characteristics beyond spatial position.

Chapter 2

Spatial attention as a magnifying glass

Published as:
Klein, B.P., Harvey, B.M., & Dumoulin, S.O. (2014). Attraction of Position Preference by Spatial Attention throughout Human Visual Cortex. *Neuron* (84), 227-237, doi: 10.1016/j.neuron.2014.08.047

Acknowledgement of author contributions:
BPK, BMH, and SOD designed research. BPK and BMH collected the data. BPK analyzed the data. BPK, BMH, and SOD wrote the paper.

Abstract

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 (McAdams and Maunsell, 1999; Treue and Maunsell, 1999; Martinez-Trujillo and Treue, 2002; Martínez-Trujillo and Treue, 2004). These effects are thought to be stronger at higher levels of visual processing (Posner and Gilbert, 1999; Cook and Maunsell, 2002; O'Connor et al., 2002; Buffalo et al., 2010; Montijn et al., 2012).

Models based on Gaussian interactions capture many aspects of these attentional effects. These models incorporate a multiplication of two Gaussian components (Womelsdorf et al., 2008; Reynolds and Heeger, 2009). 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 2.1B). 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 (Dumoulin and Wandell, 2008; Amano et al., 2009; 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 2.1A). 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 (s) for both attention 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 2.1B; SD pRF), Gaussian interaction models predict a change in the measured pRF position preference (Figure 2.1B; pRFl and pRFr) because of the change in attention field position (Figure 2.1B; 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 positions between the two conditions that produce equally sized attention fields at different locations is captured by the equation:

$$\text{pRF preferred position change} = \frac{(\mu_{AFr} - \mu_{AFl})\sigma_{SD}^2}{\sigma_{AF}^2 + \sigma_{SD}^2}, \qquad \text{Equation 2.1}$$

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 2.10) 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 2.5), this Gaussian interaction model only has one parameter to fit: the attention field size (Equation 2.1). 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 2.1) 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.

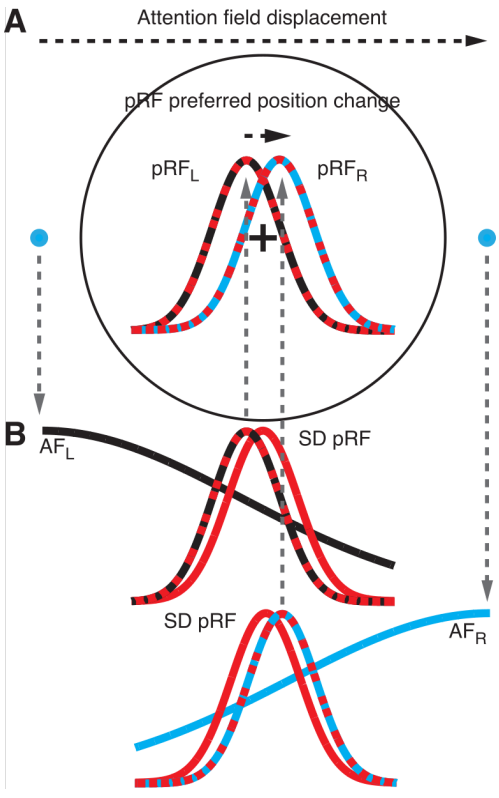

Figure 2.1 The stimulus paradigm and attention field model. **A** Subjects fixated at the cross while performing an attention-demanding contrast discrimination task at either the left or right blue dot. The black circle marks the region where a conventional mapping stimulus was shown to derive the pRF 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 (pRFL and pRFR, 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.

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 2.2 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 2.2A), 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 responses are not affected by attended location when the bar moved from the top to the bottom of the stimulus display (Figure 2.2B). This indicates that the pRF's response profile changes in the direction of the attentional manipulation, i.e., in the horizontal direction.

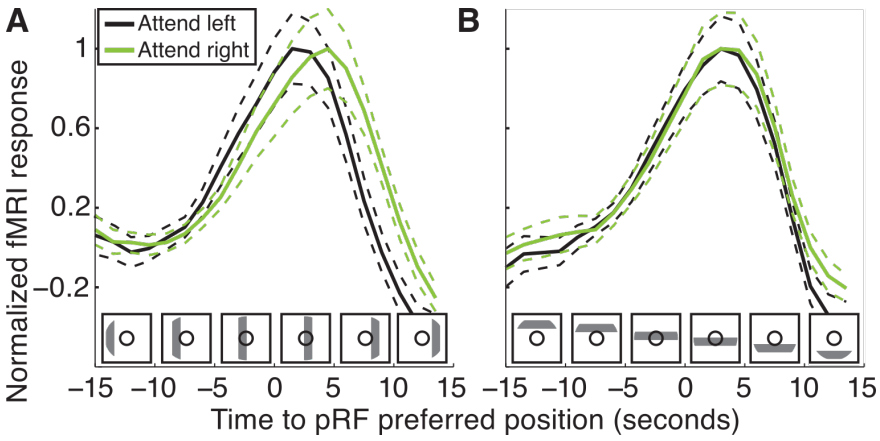

Figure 2.2 Averaged raw fMRI responses and the SEM. **A** Averaged fMRI responses to a horizontal bar sweep for V3a aligned to their respective pRF position. Black line and data points represent the attend-left condition; green line and data points represent the attend-right condition. **B** Identical 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.

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 2.3A; Figure S2.1, 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 2.3B and 2.3C 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 (FDR; 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 (ure 2.3D), 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 2.3E and S2.3C) (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

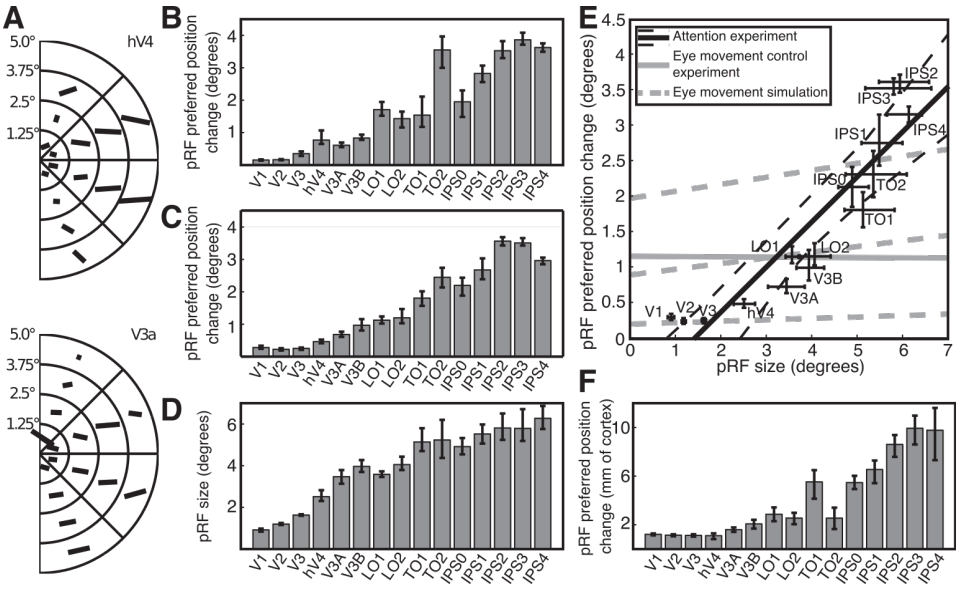

Figure 2.3 Measured pRF preferred position changes and their relation to pRF size. **A** Average size and direction of pRF preferred position change between the two attention conditions in different sectors of the visual field for one subject. All preferred position changes are in the expected horizontal direction. See Figure S2.2 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 S2.3C). 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.

during this experiment are highly correlated with those measured in the main experiment (Figures S2.3A and S2.3B; 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 2.3F), 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 2.3E, gray dashed lines). Simulated directionally biased eye movements induce changes in pRF preferred positions but without a strong increase up the visual hierarchy or 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 2.3E, solid gray line [median fit]; Figure S2.3C, 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.

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 2.4A shows the measured pRF preferred position changes from Figure 2.3A overlaid with the change predicted by the best-fitting attention field for

each visual field map. Figure 2.4B 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 2.4B, 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 2.4C). 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 2.4D) 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%) of the variance ($p = 0.64$, Figure 2.4E). When examining the variance explained for separate visual field maps (Figure S2.4), 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.

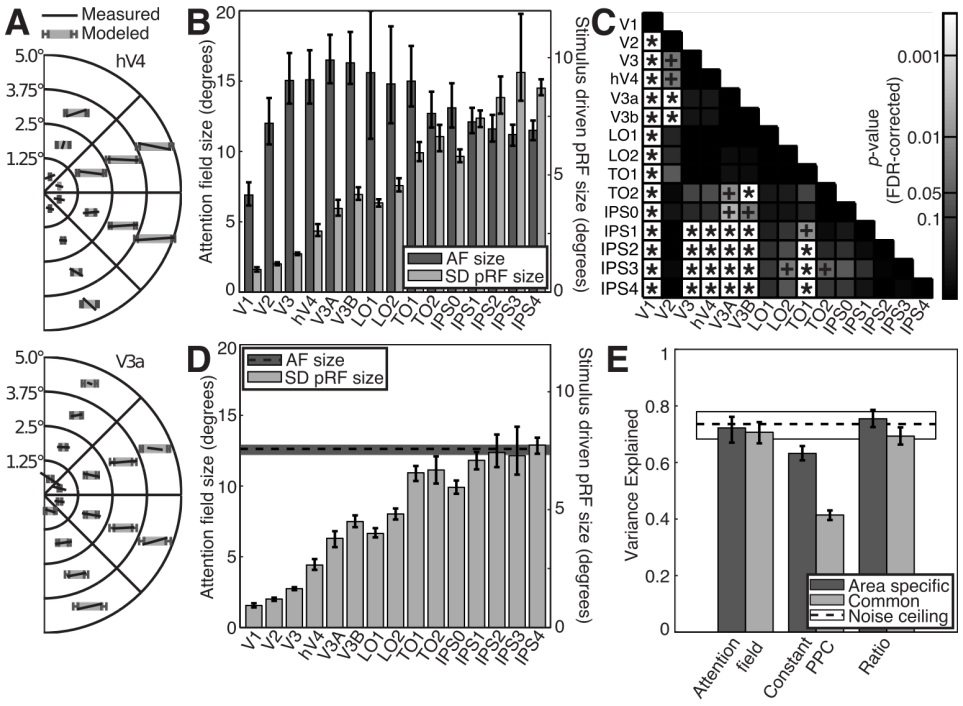

Figure 2.4 Attention field model results. **A** Polar plots similar to Figure 2.3A, with the attention field (AF) model's predicted preferred position changes presented as light gray bars. **B** Estimated AF size (dark gray bars) and stimulus-driven (SD) pRF size (light gray bars, right axis) from fitting 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.

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 2.4E). 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 2.4E). However, unlike for AF models, constraining 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 2.5, 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 2.1).

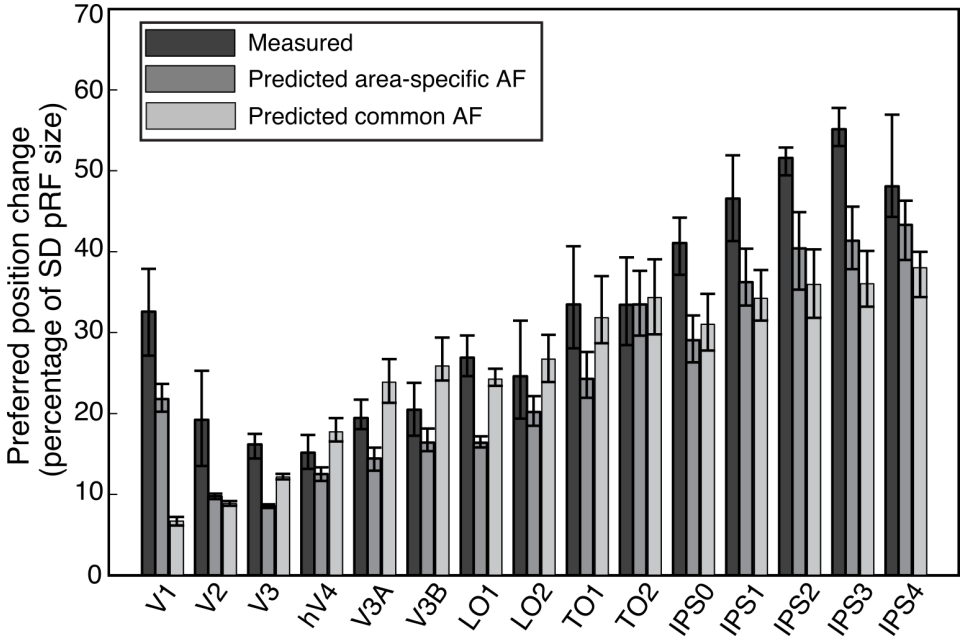

Figure 2.5 The measured and predicted pRF position change between the two conditions in percentage of the standard deviation of the stimulus-driven pRF. Position changes predicted by both a separate (gray bars) and single (light gray bars) attention field model are shown. Data are obtained by linearly fitting eccentricity binned data points. All error bars represent 95% confidence intervals. Measured pRF position change is shown with dark gray bars.

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 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, 1997) and MT (Womelsdorf et al., 2006, 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 (Smith et al., 2001; Dumoulin and Wandell, 2008). 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 (Smith et al., 2001; Dumoulin and Wandell, 2008; Harvey and Dumoulin, 2011). 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) (O'Connor et al., 2002; McAlonan et al., 2008). 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 S2.2), 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 mm3, with a field of view (FOV) of 288 3 288 3 175 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 3 2.0 3 2.0 mm, with an FOV of 190 3 190 3 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 x 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 x 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 (DeYoe et al., 1996; Wandell et al., 2007; Arcaro et al., 2011). 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 (Winawer et al., 2010; Olman et al., 2012).

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 2.2. 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 S2.1 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

$$y = ax + b, \qquad \text{Equation 2.2}$$

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 2.2). 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 2.3B–2.3F, 2.5, and S2.3).

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

$$\sigma_3^2 = \frac{\sigma_1^2 \sigma_2^2}{\sigma_1^2 + \sigma_2^2}. \qquad \text{Equation 2.3}$$

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

$$\mu_3 = \frac{\mu_1\sigma_2^2 + \mu_2\sigma_1^2}{\sigma_2^2 + \sigma_1^2}. \qquad \text{Equation 2.4}$$

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

$$\sigma_{SD}^2 = \frac{\sigma_{AF}^2}{\frac{\sigma_{AF}^2}{\sigma_{RF}^2} - 1}, \qquad \text{Equation 2.5}$$

where σ_{RF} is the mean pRF size across both attention conditions. This equation is derived from Equation 2.3. 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:

$$\mu_{RFp} = \frac{(\mu_{SD}\sigma_{AF}^2) + (\mu_{AF}\sigma_{SD}^2)}{\sigma_{SD}^2 + \sigma_{AF}^2}. \qquad \text{Equation 2.6}$$

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:

$$SSR = \sum_{i=1}^{n} (|\mu_{RFLp(i)} - \mu_{RFL(i)}| + |\mu_{RFRp(i)} - \mu_{RFR(i)}|)^2 , \quad \text{Equation 2.7}$$

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:

$$SST = \sum_{i=1}^{n} (\mu_{RFL(i)} + \mu_{RFR(i)})^2 . \quad \text{Equation 2.8}$$

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

$$VE = 1 - \frac{SSR}{SST} . \quad \text{Equation 2.9}$$

We chose the SD pRF position (μ_{SD}) and AF size σ_{AF} that maximize the variance explained (Equation 2.9). 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 2.6). However, the VE reported in Figures 2.4 and S2.4 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 2.1 is the difference between two versions of Equation 2.6 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 2.1:

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

Equation 2.10

Supplemental movies

Two supplemental movies can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2014.08.047.

Supplementary material

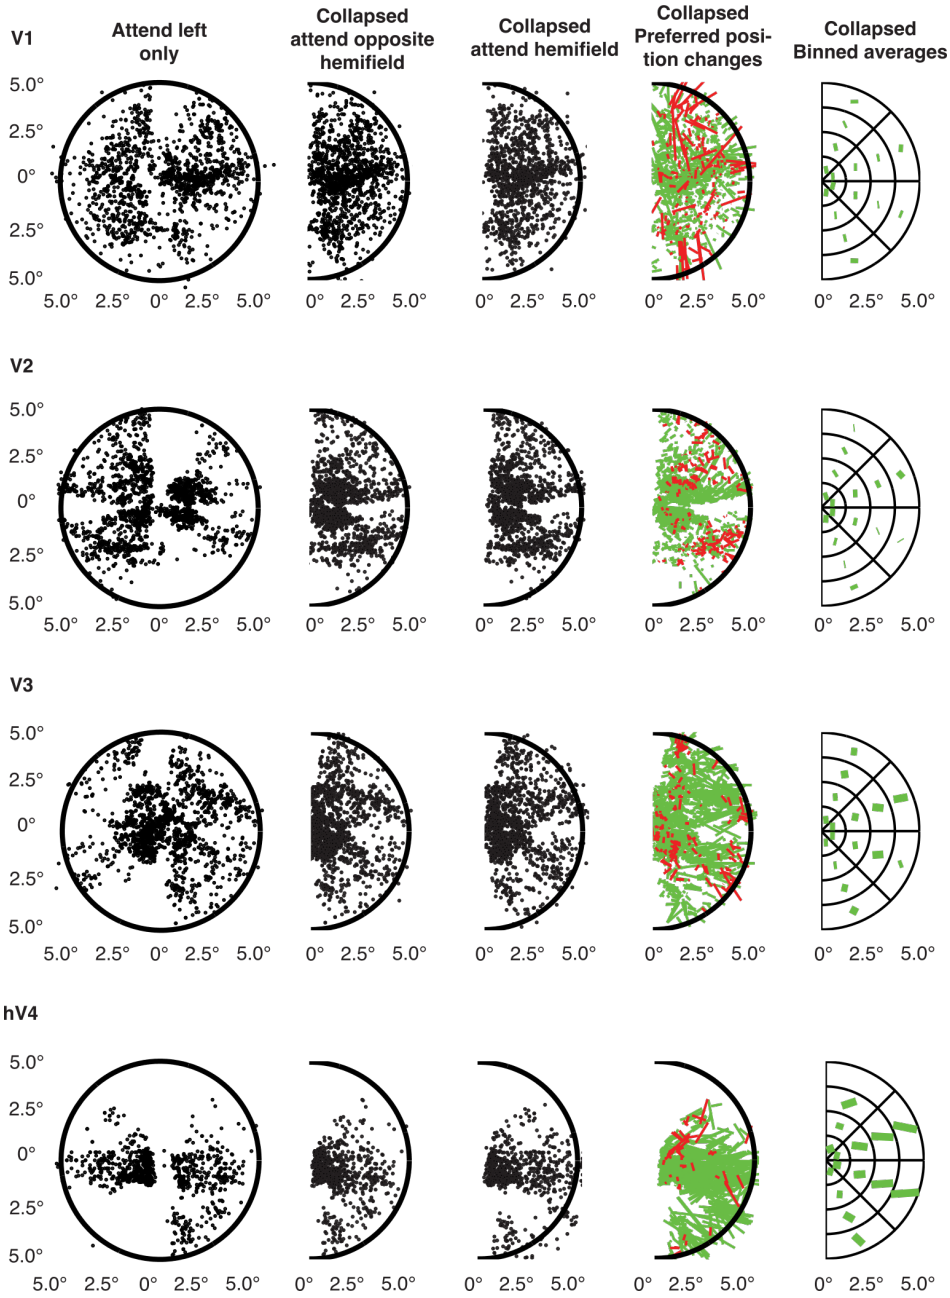

Figure S2.1 related to Figures 2.3A and S2.2. Overview of the steps in binning data from individual cortical locations for display and further analysis, shown for a single subject for areas V1 - hV4. We measure pRF positions when atttending to left or right target (1st column illustrates the pRF positions when attending to the left target). The data is first collapsed into one hemifield based on the average pRF preferred position across both conditions, giving preferred positions when the subject attended in the opposite hemifield (2nd column) and when they attended in the displayed hemifield (3rd column). The preferred position changes are the difference in preferred positions between the two attention conditions for each cortical location (4th column). Finally, preferred position changes are binned based and their average preferred position between the two conditions (5th column).The mean preferred position change is determined for each bin (Collapsed binned averages), together with the mean pRF size in each bin across both conditions. These binned preferred position changes and pRF sizes are used for subsequent analyses.

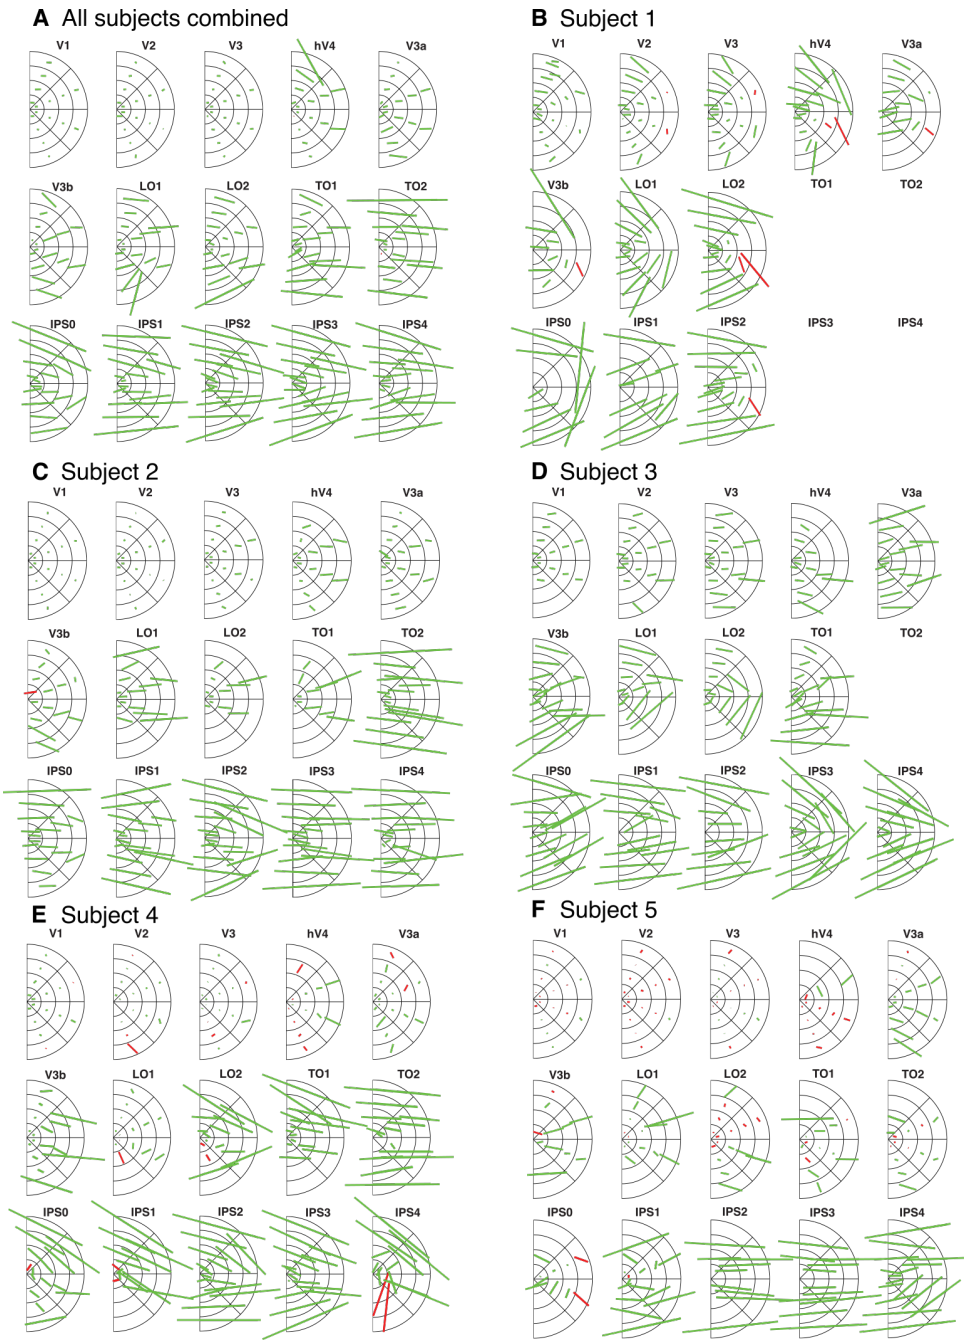
A All subjects combined
V1 V2 V3 hV4 V3a
V3b LO1 LO2 TO1 TO2
IPS0 IPS1 IPS2 IPS3 IPS4
B Subject 1
V1 V2 V3 hV4 V3a
V3b LO1 LO2 TO1 TO2
IPS0 IPS1 IPS2 IPS3 IPS4
C Subject 2
V1 V2 V3 hV4 V3a
V3b LO1 LO2 TO1 TO2
IPS0 IPS1 IPS2 IPS3 IPS4
D Subject 3
V1 V2 V3 hV4 V3a
V3b LO1 LO2 TO1 TO2
IPS0 IPS1 IPS2 IPS3 IPS4
E Subject 4
V1 V2 V3 hV4 V3a
V3b LO1 LO2 TO1 TO2
IPS0 IPS1 IPS2 IPS3 IPS4
F Subject 5
V1 V2 V3 hV4 V3a
V3b LO1 LO2 TO1 TO2
IPS0 IPS1 IPS2 IPS3 IPS4

Figure S2.2 related to Figure 2.3A. Averages size and direction of pRF preferred position change between the two attention conditions in different sectors of the visual field for each individual subject and the average across all subjects. Green and red lines mark preferred position changes whose horizontal component is in the expected or in the opposite direction respectively. Position of the lines are adjusted to cross the center of each bin.

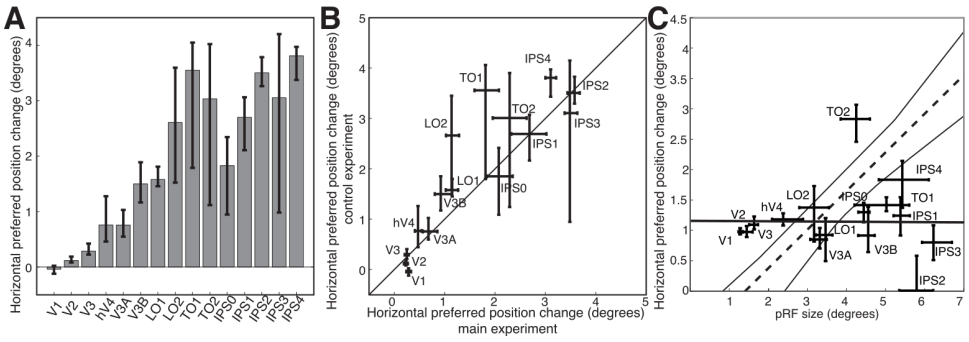

Figure S2.3 related to Figure 2.3B, C and E. Results obtained for the stimulus control experiment (see Movie S2) and eye movement control experiment. **A.** PRF preferred position changes obtained for three subjects combined, following Figure 3B. **B.** PRF preferred position changes found in the control experiment correlated well with those found in the main experiment (Pearson's r = 0.88, $p < 10\text{-}5$). **C.** PRF preferred position changes resulting from a 1° difference in eye position between the two conditions. Points mark average pRF sizes and preferred position changes for each visual area. The solid line is the best linear fit to these points. The dashed lines shows the linear fit to preferred position changes between our two attention conditions. These preferred position changes are unlikely to be caused by eye movements towards the attended targets, as the preferred position changes induced by eye movements do not increase with pRF size. All error bars are 95% confidence intervals determined by bootstrapping.

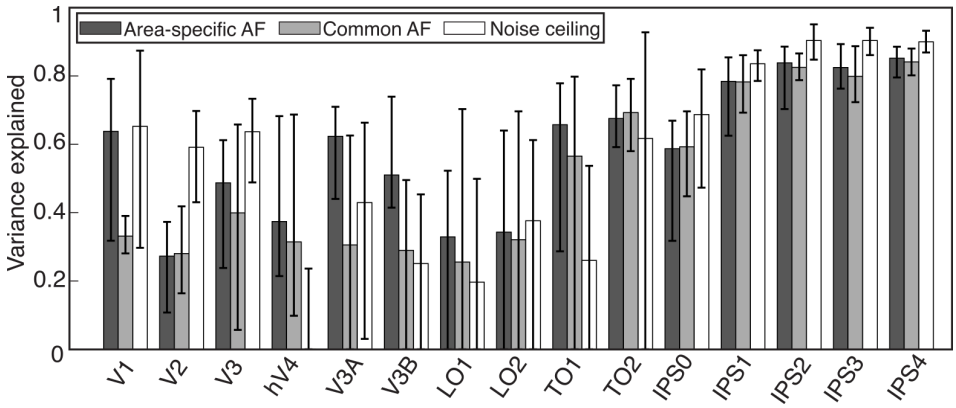

Figure S2.4 related to Figure 2.4D and E**.** Area-specific variance explained by attention field models. Variance in pRF preferred positions explained by area-specific (dark gray bars) and common (light gray bars) attention fields (AF), together with noise ceilings (white bars). We note that where part of the visual field representation is consistently difficult to record from (as in hV4, Fig. S2, Winawer et al., 2010) our noise ceiling estimation procedure cannot effectively compare pRF preferred position changes across hemifields, while AF models can still explain the preferred position changes in the remaining data. Furthermore, any systematic differences in both pRF sizes and preferred position changes across hemifields will be captured by AF models but not the comparisons between hemifields used to estimate noise ceilings. Errorbars represent 95% confidence intervals.

Chapter 3

Mechanisms underlying attentional magnification

Klein, B.P., Fracasso, A., Paffen, C.L.E., te Pas, S.F., & Dumoulin, S.O.
Cortical depth dependent population receptive field attraction by spatial attention in human V1 (under review).

Acknowledgement of author contributions:
BPK, AF, and SOD designed research. BPK and AF collected the data. BPK analyzed the data. BPK, CLEP, SFP, and SOD wrote the paper.

Abstract

Visual spatial attention concentrates neural resources at the attended location. Recently, we demonstrated that voluntary spatial attention attracts population receptive fields (pRFs) toward its location throughout the visual hierarchy. Theoretically, both a feed forward or feedback mechanism can underlie pRF attraction at a given cortical area. Here, we use sub-millimeter ultra-high field functional MRI to measure pRF attraction across cortical depth and assess the contribution of feed forward and feedback signals to pRF attraction. In line with previous findings, we find consistent attraction of pRFs with voluntary spatial attention in V1. When assessed as a function of cortical depth, we found pRF attraction in every cortical portion (deep, center and superficial), although the attraction is strongest in deep cortical portions (near the gray white matter boundary). Following the organization of feed forward and feedback processing across V1, we speculate that a mixture of feed forward and feedback processing underlies pRF attraction in V1. Within this mechanism, the feedback component likely arrives in deep cortical portions.

Introduction

Visual attention is the mechanism through which we concentrate neural resources on relevant visual information. Computationally, the effects of visual attention on both human perception (Herrmann et al., 2010; Klein et al., 2016) and neural responses (Womelsdorf et al., 2008; Reynolds and Heeger, 2009; Klein et al., 2014) can be modeled as an interaction between two components, one representing the influence of attention (attention field) and the other representing a stimulus driven neural response property. Building on this attention field model, we have recently shown that visual attention voluntarily directed at a spatial location attracts the population receptive fields (pRFs) towards the attended location across the visual hierarchy, producing distortions in the perceived location of visual stimuli (Dumoulin and Wandell, 2008; Klein et al., 2014, 2016).

Here, we examine the influence of voluntary spatial attention on pRF position across cortical depth in human V1. Imaging across cortical depth or laminar imaging may reveal unique information about the direction of information flow, specifically whether processes are driven by feed forward or feedback signals (Dumoulin et al., 2017; Lawrence et al., 2017; Self et al., 2017). In line with this notion, we recently showed that pRF size vary across cortical depth. This variation of pRF size across cortical depth closely resembles electrophysiological results and reflects the information flow across cortical depth (Fracasso et al., 2016; Self et al., 2017). Here we extend this approach to examine whether pRF attraction varies across cortical depth in human V1.

We used ultra-high field (7T), sub-millimeter functional MRI (fMRI) to measure pRF position attraction as a function of cortical depth in V1. Using this approach, we first show that voluntary spatial attention attracts pRF preferred positions towards the attended location in V1. Furthermore, we find pRF attraction in every cortical portion, although the attraction is strongest near the gray/white matter boundary and decreases towards superficial cortical depths. Following the organization of feed forward and feedback afferent terminals in V1 (Hubel and Wiesel, 1972; Benevento and Rezak, 1976; Rockland and Pandya, 1979; Blasdel and Lund, 1983; Felleman and Van Essen, 1991), we speculate that a mixture of feed forward and feedback processing underlies pRF attraction in V1, with the feedback component likely arriving in depths close to the white matter.

Materials and Methods

Subjects

Twelve subjects participated in this study (three females, age range 21 - 42, mean age 28.6). All subjects had normal or corrected to normal visual acuity and gave informed consent. Two subjects were excluded from further analysis, one due to imaging artifacts and one to having an attention disorder. All experimental procedures were approved by the ethics committee of University Medical Center Utrecht.

Visual stimuli and experimental design

Visual stimuli were presented by back-projection onto a 15.0x7.9 cm screen inside the MRI bore. Subjects viewed the display through prisms and mirrors, and the total distance from the subject's eyes (in the scanner) to the screen was 35.5 cm. Visible display resolution was 1024x538 pixels.

The stimuli were generated in Matlab using the PsychToolbox (Brainard, 1997; Pelli, 1997). The fixation cross was composed of two diagonal red lines covering the entire display, one pixel wide, intersecting at the center of the screen. Subjects were instructed to fixate the intersection of the two lines. This design facilitates accurate fixation (Schira et al., 2009). Stimuli consisted of a circular dartboard pattern presented behind the diagonal lines and centered on the center of the screen. The pattern consisted of 24 rings each 0.24° visual angle wide. Each ring was divided into 12 black and white segments, each subtending 30°. The rings rotated around the center of the screen, moving at 60° per second. Neighboring rings moved in opposite directions.

The dartboard pattern was viewed through two C-shaped, equi-eccentric apertures. The apertures were centered on the horizontal meridian and subtended 120° each. The apertures were 0.86° visual angle wide and cycled through all eccentricities between 0.58° visual angle and 5.78° visual angle in 12 steps of 0.43° visual angle. On the last step in the stimulus cycle, the apertures were wrapped around the eccentricity range covering the inner most (0.58° visual angle - 1.01° visual angle) and outer most eccentricities (5.345° visual angle - 5.778° visual angle). This creates a cyclical stimulus that is assumed by our functional analysis (see below). The apertures only moved in the outward direction. Each step in the stimulus cycle lasted 4 seconds (1 TR, functional volume acquisition, see below). Consequently, one stimulus cycle lasted 48 seconds (12 TRs). One entire experimental run consisted of 6 stimulus repeats preceded by half a stimulus cycle

to ensure a steady BOLD signal, totaling 78 TRs. We chose this specific stimulus as it allows us to estimate preferred eccentric position from the resulting BOLD signals with using relatively little time points. This is necessary considering the slow repetition time (4 seconds) and limited number of time points per scan (72), which are typical for sub-millimeter functional imaging.

Simultaneous with the dartboard stimuli, we presented two circular $1/f$ noise patterns, 0.24° visual angle in radius. The noise patterns were centered on the horizontal meridian, 6.3° visual angle left and right from fixation. The patterns randomly changed orientation every 250 ms and increased contrast on 5% of orientation changes (randomly chosen and different between functional scans; Figure 3.1).

Subjects were instructed to covertly attend one of the two noise patterns for the duration of one functional scan, and to detect contrast increments on the attended pattern. The attended location alternated between scans. 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 (d' > 0). Regardless of the location that was to be attended, 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 a contrast increment on the left and right noise pattern and a subsequent detection response can be taken as a detection on both noise patterns. Subjects' performance was significantly better for the attended target (d' = 2.05, SD = 0.32), than for the unattended target (d' = 0.81, SD = 0.24) (t = 24.4, $p < 0.001$, two-sided, repeated measures t-test), and we found no differences between performance for the left and right targets ($p = 0.88$, two-sided, repeated measures t-test).

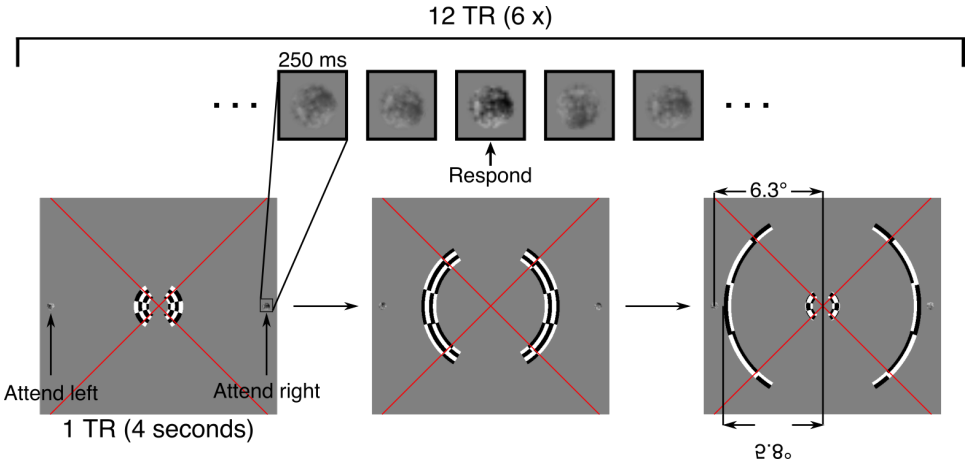

Figure 3.1 Stimulus and task. Subjects fixated the center of the screen, marked by the intersection of two diagonal red lines running across the screen. The stimulus consisted of a rotating dartboard pattern viewed through two C-shaped apertures. The apertures moved from the center of the screen towards the periphery in an expanding fashion. One stimulus cycle lasted 12 TRs (1 TR = 4 seconds) and was repeated 6 times during one functional scan. Concurrently, we presented two 1/f noise patterns left and right of the center of the screen. The noise patterns changed orientation independently every 250 ms. and increased in contrast on 5% of orientation changes. Subjects were instructed to covertly attend either the left or right noise pattern for the duration of one functional scan and report the contrast increments of the attended pattern.

Functional data acquisition

High resolution functional data were acquired using a Philips 7T scanner (Best, Netherlands) and a volume transmit coil for excitation (Nova Medical, MA, USA). Head motion inside the scanner was minimized using a combination of noise-cancelling headphones and foam padding. Functional T2*-weighted 3-dimensional multi-shot EPI (3D-EPI, two shots per slice, 35 slices, 70 shots overall) data were acquired using two custom-built high-density 16-channel surface coils with a of total 32 channels for signal reception (Petridou et al., 2013). The sequence parameters were: TR/TE = 57/28 ms, flip angle: 20°, acceleration factor using SENSE encoding: 3.5 (right-left) × 1.3 (anterior-posterior), echo planar factor: 27, BW (phase-encode): 19.1 Hz/pixel, readout duration ~ 52 ms (with potential blurring in the phase-encode direction estimated at ~ 16%, Haacke et al.,1999), voxel size = 0.70 mm

isotropic, FOV = 131 (right-left) × 120 (feet-head) × 24.5 (anterior-posterior) mm^3, 35 coronal slices, and 28% oversampling in the slice direction. This acquisition sequence produced geometric distortions near the edges of the functional imaging volume. Furthermore, distortions are more severe near the air/tissue interface, for example the edge of cortical gray matter (Truong et al., 2008) and near the basal ganglia due to B0 inhomogeneities resulting from iron storage. By limiting our analyses to primary visual cortex (the calcarine sulcus) we attenuate the effects of geometric distortions in our functional data, as it is away from the edges of the cortical gray matter and basal ganglia. We centered the functional volume on the calcarine sulcus to place it away from the distortions near the edges of the functional volume and minimize their effect on our functional data. Functional volumes were acquired every 4 s and functional scans were each 312 s (78 functional acquisitions) in duration. Each subject completed 6 to 8 functional runs in a single session.

Anatomical data acquisition and processing

For five subjects (S1, S3, S6, S7, S10) anatomical images were acquired using a 3D T1-weighted MPRAGE sequence (TR = 7.48 ms, TE = 3.47 ms, flip angle = 8°, FOV: 250 x 200 x 180 mm, voxel size 0.5 x 0.5 x 0.5 mm).
For two subjects (S2 and S9) anatomical T1-weighted images were acquired using the MP2RAGE sequence (Marques et al., 2010) with the following parameters: TR = 5982 ms, FOV: 220 x 220 x164 mm, voxel size: 0.625 x 0.625 x 0.64 mm, TI1/TI2 = 800/3686 ms, flip angle = 7°/5°. For two subjects (S4 and S8), T1-weighted images were acquired at a resolution of 0.5x0.5x0.8 mm (TR = 7ms, TE = 2.84 ms, flip angle = 8°). All the above anatomical images were acquired on a Philips 7T scanner using a 32-channel head coil. Finally, for one subject (S5), T1-weighted images were acquired on a Philips 3T scanner (TR 10.029 ms TE = 4.6 ms, flip angle = 8°, voxel size 0.75 x 0.75 x 0.8 mm). Anatomical images not acquired at 0.5 mm isotropic resolution were resampled to this resolution. Gray/white matter segmentations were obtained in MIPAV using the TOADS/CRUISE algorithm (Han et al., 2004; Bazin and Pham, 2007) and subsequently manually corrected. We employed the equi-volume model approach to build a coordinate system along cortical depth taking local curvature into account (Waehnert et al., 2014).

V1 ROI definitions

V1 definitions were acquired during separate scanning sessions, or for the purposes of a different experiment. In both cases, we used a regular pRF mapping stimulus, described in detail by Dumoulin and Wandell (2008). Shortly, this stimulus consisted

of a contrast defined, bar shaped checkerboard pattern moving across the visual field in eight different directions (four cardinal, four diagonal). We used a regular pRF modeling procedure (Dumoulin and Wandell, 2008), to estimate each voxel's best fitting pRF as described by its position in the visual field (X and Y) and its extent (standard deviation, sigma). We converted the X and Y positions of every pRF to polar angle and eccentricity estimates, which were rendered on an inflated cortical surface (Wandell et al., 2000). The position of V1 was obtained by following reversals in polar angle and eccentricity progression (Sereno et al., 1995; Wandell et al., 2007). V1 ROI definitions were imported into the subject's high-resolution anatomical space. Finally, we clipped V1 ROI definitions to account for differences in the polar angle and eccentricity coverage between the pRF mapping stimuli used to define V1 and the current experimental stimulus.

Pre-processing of functional data

Functional data was preprocessed using AFNI (Cox, 1996). We corrected for head motion between scans by aligning the first functional volumes for each scan using 3dvolreg. Correction for within-scan motion was done by aligning all the frames of a run to the first frame. We corrected for between and within-scan motion in a single step and averaged the motion corrected images from a single session together. We coregistered the averaged functional image to the motion-corrected and averaged T1 weighted image using an affine transformation. The coregistration was divided into three steps. First, we clipped the T1 weighted anatomy in the anterior-posterior direction, leaving only the occipital lobe. As we used different receive coils for our functional and anatomical data acquisition, we obtained a good starting point for the coregistration by centering the functional image on the clipped anatomy using their respective centers of mass of the reduced FOV volumes, or manually using 3dSlicer (http://www.slicer.org, Fedorov et al., 2012). Second, the averaged functional image was coregistered with the T1 weighted images using an affine transformation via the function 3dAllineate, using the two-pass option. This procedure blurs the functional image and initially allows for a large rotation and shift, and then refines the coregistration using an affine transformation. In the third step the resulting coregistration was further optimized via 3dAllineate, but now using the one-pass option. This does not blur the functional image and thus coregisters the original functional volume with the anatomy. It allows only for a small amount of motion, again using an affine transformation. The obtained transformations were combined in a single affine transformation matrix.

We used local Pearson correlation as the cost function for our coregistration (Saad et al., 2009) but adopted alternative cost functions (such as mutual

information and normalized mutual information) when this initial cost function yielded unsatisfactory results. Our main priority was to obtain an optimal coregistration around the calcarine sulcus. Coregistration output was visually inspected by evaluating the location of anatomical markers as gray matter / white matter (GM/WM) and gray matter / cerebro-spinal fluid (GM/CSF) boundaries in the calcarine sulcus and by the correspondence of the position of large vessels between the T1-w and the averaged functional data.

Functional and statistical analysis

We discarded the first six volumes of every functional run and averaged the functional scans for both conditions (attend left / attend right) separately. We parameterized the fMRI time series using the traveling wave analysis implemented in the mrVista softaware package for Matlab (http://white.stanford.edu/software)(Engel et al., 1994; Engel, Glover, & Wandell, 1997; Sereno et al., 1995). This analysis yields three parameters: phase, amplitude and coherence. The phase gives the temporal delay of the stimulus frequency in the time series in radians. Within our stimulus design, this is a measure of preferred eccentric position. The amplitude gives the BOLD amplitude in percentage signal change at the stimulus frequency. Finally, coherence is the correlation between the harmonic at the stimulus frequency and the fMRI time series. As such, it is a measure of signal quality and reliability of the corresponding phase value. Finally, we interpolated these parameters into the anatomical space, using nearest neighbor interpolation and the transformation computed by the coregistration (see above).

We measured the phase for every voxel in the functional volume twice, once while attention was directed at the hemifield ipsilateral to the voxel (ipsilateral hemisphere) and once while attention was directed at the hemifield contralateral to the voxel (contralateral hemisphere). We computed pRF attraction between conditions, by subtracting the phase estimate measured for a voxel when it was located in the hemisphere ipsilateral to the attended target from the phase estimate for the same voxel when it was located in the hemisphere contralateral to the attended target. These phase differences were wrapped to yield values ranging from -pi to pi, with positive values corresponding to higher preferred eccentric positions in the attended hemifield. Next, we converted the phase differences to degrees of visual angle by dividing by 2pi and multiplied them by the stimulus range (5.2° visual angle), yielding preferred eccentric position changes in degrees of visual angle. We excluded anatomical voxels outside V1 (see above) and those that were located outside the gray matter. Also, we excluded voxels with a coherence value lower than the 25th percentile in either one of the conditions. Additionally, to reject

voxels with pRFs near the stimulus edge, we excluded voxels with an averaged phase less than the 12.5th percentile or more than the 87.5th percentile of the stimulus eccentricities.

We assessed the statistical significance of the preferred eccentric position changes (Figure 3.3D) across V1 using repeated measures t-tests. These t-tests were performed using anatomical voxels as individual data points. As the spatial resolution of the anatomical volumes is higher than the spatial resolution of the functional volumes, the functional volumes were upsampled to match the anatomical resolution. The t-tests reported were corrected for this upsampling. We assessed the variation of the preferred eccentric position changes across cortical depth for both the attention conditions (Figure 3.4E) and eye movement control data and simulation (Figure 3.5B) using linear regression. Similarly, we analyzed the increase in fMRI response amplitude (Figure 3.4D) and change in fMRI response amplitude between conditions (Figure 3.6A) as a function of cortical depth using linear regression. These linear regression analyses used the binned averages for all subjects together as its individual data points. The linear regression weighted the binned averages by the number of voxels each average represents.

Averaged BOLD responses (Figure 3.3C)

To assess differences in BOLD responses, we only included fMRI time series corresponding to voxels included in the phase analysis (see above). We averaged the BOLD responses to all stimulus repeats together, giving the averaged BOLD response to a single stimulus cycle. Next, we used linear interpolation to align the BOLD responses according to their averaged phase across the two conditions. Finally, we averaged the aligned BOLD responses from all voxels together, separately for when attention was directed at the target in the contralateral and ipsilateral hemifield. Conceptually, this analysis yields the averaged BOLD response from both conditions in the hypothetical case that all pRFs in V1 have the same preferred eccentric position when averaged across conditions.

Hypothesized profiles of pRF attraction across cortical depth

To hypothesize how contributions of feed forward and feedback processing to pRF attraction may shape the profile of pRF attraction across cortical depth, we combined an attention field model with the known functional and anatomical organization of laminar connectivity (Hubel and Wiesel, 1974; Felleman and Van Essen, 1991; Fracasso et al., 2016; Dumoulin et al., 2017). In this section, we first apply an attention field model to our experimental design. Then we discuss how the

forward flow of signals across cortical depth affects pRF properties in V1. Finally, we consider how this flow will shape the profile of pRF attraction across cortical depth.

Attention field model. As we summarize the fMRI responses using one parameter, eccentricity, we consider the pRFs underlying the fMRI responses as a one dimensional Gaussian defined along the radial axis (x) (Dumoulin and Wandell, 2008; Fracasso et al., 2016):

$$pRF(x) = e^{(\frac{x-\mu_{pRF}}{2\sigma^2_{pRF}})^2} \quad \text{Equation 3.1}$$

where μ_{pRF} is the preferred eccentric position and σ_{pRF} is the size (standard deviation) of the pRF. We model the effect of attention on preferred eccentric position as a multiplication between two Gaussians (Womelsdorf et al., 2008; Reynolds and Heeger, 2009; Klein et al., 2014). One of these represents the influence of attention, the attention field, whereas the other represents the pRF without the influence of attention, the stimulus driven pRF. This multiplication produces a third Gaussian, representing the pRF under influence of attention. As such, the preferred eccentric position of the pRF under attention ($\mu_{AF\times pRF}$) is given by:

$$\mu_{AF\times pRF} = \frac{\mu_{AF}\sigma^2_{pRF} + \mu_{pRF}\sigma^2_{AF}}{\sigma^2_{AF} + \sigma^2_{pRF}} \quad \text{Equation 3.2}$$

where μ_{AF} and μ_{pRF} represent the positions and σ_{AF} and σ_{pRF} the sizes of the attention field and stimulus driven pRF. Importantly, we compare the preferred eccentric position under two different conditions. Consequently, the preferred eccentric position change between the two conditions is given by:

$$\mu_{AFr\times pRF} - \mu_{AFl\times pRF} = \left(\frac{\mu_{pRF}\sigma^2_{AFr} + \mu_{AFr}\sigma^2_{pRF}}{\sigma^2_{AFr} + \sigma^2_{pRF}}\right) - \left(\frac{\mu_{pRF}\sigma^2_{AFl} + \mu_{AFl}\sigma^2_{pRF}}{\sigma^2_{AFl} + \sigma^2_{pRF}}\right)$$

Equation 3.3

We assume that the stimulus driven pRF sizes are the same for both conditions. Moreover, the voluntary attention task was performed on targets with identical properties and performance was similar for both conditions. Therefore, we also assume the attention field sizes to be similar between the two conditions

($\sigma_{AFr} = \sigma_{AFl}$). Finally, as the attended targets are at the same distance from fixation, the attention fields in both conditions are the same distance from fixation as well ($\mu_{AFr} = -\mu_{AFl}$). Under these assumptions, Equation 3.3 can be simplified to:

$$\Delta\mu_{AF\times pRF} = \frac{\left(\mu_{AFr} - \mu_{AFl}\right)\sigma_{pRF}^2}{\sigma_{AF}^2 + \sigma_{pRF}^2} \qquad \text{Equation 3.4}$$

As such, this model predicts that preferred eccentric position changes ($\Delta\mu_{AF\times pRF}$) are a function of the attention field size, stimulus driven pRF size and the distance between the two attended locations ($\mu_{AFr} - \mu_{AFl}$). Because of our experimental design, the attention field size and distance between attended locations are the same for every pRF: only the pRF size will vary across pRFs. Therefore pRF size will be the major source of variation in preferred eccentric position changes in our design (Klein et al., 2014).

Forward flow across cortical depth. Within the context of V1's neural organization, we assumed that the stimulus driven size and preferred eccentric position of the pRF (without the effect of attention) are the result of feed forward processing. Regarding V1, these feed forward signals predominantly originate in the lateral geniculate nucleus (LGN) and terminate in V1's central cortical depths (Figure 3.2A, pRF; Blasdel & Lund, 1983; Callaway, 1998; Felleman & Van Essen, 1991; Hubel & Wiesel, 1972). Subsequently, neural populations in deep and superficial cortical portions inherit their feed forward, stimulus driven pRFs by sampling from neural populations in central cortical portions (Figure 3.2A, pRF'; Briggs & Callaway, 2001; Callaway, 1998; Fitzpatrick, Lund, & Blasdel, 1985; Fracasso et al., 2016; Maunsell & Gibson, 1992; Self, Kerkoerle, Supèr, & Roelfsema, 2013; Usrey & Fitzpatrick, 1996; Yoshioka, Levitt, & Lund, 1994).

One way to model sampling from one cortical layer to another is as a convolution, where a single neural population in deep and superficial cortical portions receives input from multiple populations in the central cortical portion (Fracasso et al., 2016). This way, the properties of pRFs in deep and superficial cortical portions can be obtained by convolving a function representing the response property at central cortical portions (a pRF Gaussian in this case) and a function representing the sampling function. Between visual field maps, this sampling function is Gaussian shaped (Motter, 2009; Kumano and Uka, 2010; Harvey and Dumoulin, 2011; Haak et al., 2013) and this approach was recently extended to

sampling between layers (Fracasso et al., 2016). As such, pRFs in deep and superficial cortical portions are the product of the convolution between two Gaussian functions, one representing the pRF at central cortical portions and one representing the sampling function from this cortical portion.

Conceptualizing sampling between cortical layers this way highlights two important points: (1) pRF sizes will increase from central to deep and superficial cortical portions (Fracasso et al., 2016) and (2) the pRF position of a neural population in the deep or superficial cortical portion is equal to the Gaussian weighted average of the positions of the pRFs this population samples from in the central cortical portion (Fracasso et al., 2016). In other words sampling between cortical portions does not change pRF positions between cortical portions (Figure 3.2A, $\mu'_{pRF=}\mu_{pRF}$)(Hubel and Wiesel, 1974). So, in sum, the feed forward flow of information produces larger pRF sizes in deep and superficial cortical portions (Fracasso et al., 2016) but no systematic variation of pRF position across cortical depth.

Attention field model applied to information flow across cortical depth. When applied to feed forward processing in V1, the attention field represents an attentional influence that produces pRF attraction by interacting with forward input to V1 via, for example, response modulation at the level of LGN (McAdams and Maunsell, 1999; O'Connor et al., 2002; Compte and Wang, 2006; McAlonan et al., 2008). As forward input to V1 terminates in central cortical portions, we can model pRF attraction in a feed forward process as an interaction between the attention field and the stimulus driven pRF in V1's central cortical portion (Equation 3.4; Figure 3.2B, AF x RF). As discussed above, this attracted pRF position in the central cortical portion will be inherited by the deep and superficial cortical portions (Figure 3.2B, pRF'_{AF}). Thus unlike our speculations in an earlier paper (Klein et al., 2014), sampling between cortical layers as modeled here cannot amplify pRF attraction. In other words, feedfward information flow will inherited the pRF attraction from central layers and is in this case not a fraction of pRF size. Thus, we hypothesize that a feed forward driven pRF attraction yields a uniform pRF attraction across cortical depth (Figure 3.2D).

When applied to feedback processing, the attention field represents an attentional influence that is fed back to V1, where it interacts with feed forward, stimulus driven processing to produce pRF attraction (Figure 3.2C, Compte and Wang, 2006; Bobier et al., 2014). Feedback connections terminate specifically in deep and superficial portions in V1 (Lund et al., 1975; Benevento and Rezak, 1976; Rockland and Pandya, 1979; Felleman and Van Essen, 1991; Yoshioka et al., 1994;

Shipp, 2003). As such, the interaction between the attention field and stimulus driven pRF will specifically occur in deep and/or superficial cortical portions akin to equation 3.4 (Figure 3.2C green and blue arrows respectively). Thus, we hypothesize that a feedback driven pRF attraction yields a non-uniform pRF attraction across cortical depth, specifically with larger attraction at deep and/or superficial depths (Figure 3.2E).

In conclusion, we have highlighted several important concepts. (1) Within our design, pRF attraction will be a function of pRF size. (2) However, inheriting pRF attraction from one location will not increase the pRF attraction. (3) If attention attracts pRFs in V1 via a feed forward process, this attraction will be the same (inherited) across cortical depth. (4) If attention attracts pRFs in V1 via a feedback process, this will happen via feedback afferents in deep and superficial cortical portions yielding (stronger) pRF attraction limited to deep or superficial cortical portion, or both (Muckli et al., 2015; Kok et al., 2016; Hembrook-Short et al., 2017; Nandy et al., 2017).

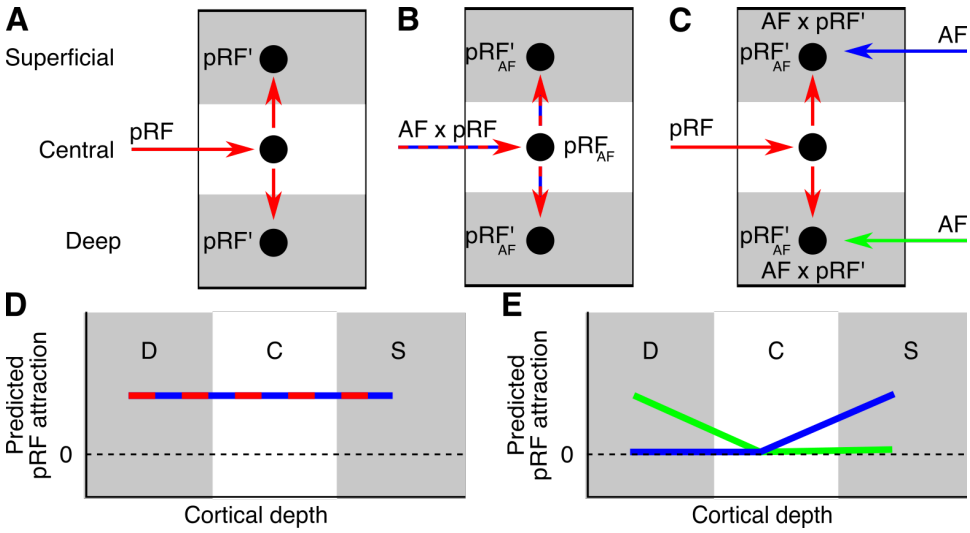

Figure 3.2 Hypothesized profiles of pRF attraction across cortical depth. **A** Neural populations in the central cortical portion obtain their pRFs through forward inputs from the lateral geniculate nucleus (LGN; red horizontal arrow). Following the forward flow of information across cortical depth, neural populations in deep and superficial cortical portions sample from the central cortical portion (pRF', red vertical arrows), resulting in larger pRFs, but with identical preferred eccentric positions (Fracasso et al., 2016). **B** pRF attraction implemented by a feed forward mechanism can be modeled as an interaction between the attention field (AF) and the stimulus driven pRF in the central cortical portion (AF x pRF). This produces the pRF under influence of attention in the central cortical portion (pRF_{AF}), with a preferred eccentric position attracted towards the attended location. Through sampling from central cortical portions, neural populations inherit their pRFs (pRF'_{AF}), together with the attracted preferred eccentric position. **C** pRF attraction implemented by a feedback mechanism can be modeled by an interaction between the attention field and pRFs in deep and superficial cortical portions (AF x pRF', blue and green arrow), producing the pRF under influence of attention in deep and superficial cortical portions (pRF'_{AF}). We speculate that this interaction may occur either in both deep and superficial cortical portions, or can be limited to the deep or superficial portion only. **D** A feed forward implementation of pRF attraction (b) predicts no variation of pRF attraction across cortical depth, as deeper and superficial layers inherit the pRF attraction from central cortical depths. **E** A feedback implementation of pRF attraction (c) predicts that pRF attraction specifically occurs in either deep or superficial cortical depths, or both.

Eye movement controls

Prior to the scanning sessions, we trained subjects on the experimental task outside the scanner while we monitored their eye movements using a highly accurate, head mounted Eyelink II system (SR Research). To estimate the bias in gaze position towards the attended targets, we subtracted median gaze position during the attend left condition from the median gaze position during the attend right condition for every subject separately. Averaged across subjects, the median gaze position difference per condition was 0.046° visual angle, yielding a total bias between conditions of about 0.092° visual angle.

Inside the scanner, we presented the same stimulus as in the main experiment (Figure 3.1), but we shifted the fixation cross 0.1° visual angle to the left or right relative to the center of the stimulus, alternating left and right fixations between scans. This yields a gaze position difference between conditions of 0.2° visual angle, which is twice the size of the averaged bias in gaze position measured

prior to scanning sessions. For this control experiment, the fixation cross's color alternated between red and green and subjects had to report the color changes and ignore the targets left and right of the stimulus. We analyzed the data from this experiment in the same way as the data from the main experiment (see above).

As the averaged bias in gaze position is less than the average main effect on preferred eccentric position change, we also generated a simulated data set with an eye movement bias scaled to match the size of the average preferred eccentric position change in the main experiment. As eye movements towards (or away) from the attended location move pRFs to higher (or lower) eccentricities, they are stimulated later (or earlier) by our stimulus. In order to simulate a larger bias in the BOLD time series measured for the eye movement control experiment, we interpolated the BOLD time series to later time points in the hemispheres contralateral to the direction of the fixation shift (i.e. right (or left) hemisphere when fixation cross is shifted to the left (or right)) and to earlier time points in the hemispheres ipsilateral to the fixation offset.

We determined the amount of interpolation for every TR separately by random sampling from the distribution of gaze positions measured for each subject, adding or subtracting a fixed amount to produce the desired average offset between the two conditions. Doing so, we created 1000 data sets for every subject in the eye movement control condition with an eye movement offset between the two conditions, that, on average, matched the attentional effect observed in the main experiment. As we sampled from the subject's distribution of gaze positions from prescanning sessions, the shift variance was matched to the subject's gaze position variance.

Results

pRFs in V1 are attracted toward the locus of spatial attention

Inside the MRI scanner, subjects fixated the center of the screen while they performed an attention demanding contrast discrimination task 6.3° left or right from fixation, for the duration of one functional scan (Figure 3.1). Following the attention field model (Womelsdorf et al., 2008; Reynolds and Heeger, 2009), we predicted that voluntary attention to either target would result in pRF attraction towards the attended target (Klein et al., 2014). Importantly, this attraction would manifest as higher preferred eccentric positions for pRFs near the horizontal meridian in the hemifield containing the attended location.

To examine preferred eccentric positions near the horizontal meridian during task performance, we measured fMRI responses to two equi-eccentric, C-

shaped, dartboard stimuli. The dartboard stimuli moved in a traveling wave design (Engel et al., 1994; Sereno et al., 1995; Fracasso, Petridou, & Dumoulin, 2016), i.e. from the fixation point towards the attended location, up to 5.8°, in an expanding fashion. By using C-shaped stimuli we limited visual stimulation to the horizontal meridian (Figure 3.1). One functional scan consisted of six stimulus repeats, producing six peaks in each voxel's fMRI response which correspond to the stimulus passing six times through its pRF (Figure 3.3A).

We extracted preferred eccentric positions from the fMRI responses from the attend left and attend right condition separately and assessed the quality of these estimates by overlaying them on a reconstructed cortical surface. As expected, preferred eccentric positions gradually increased along the posterior - anterior axis for both conditions and for all subjects (Figure 3.3B, see Supplemental Figure 3.1 for all subjects). As can be seen from Figure 3.3B, the preferred eccentric positions changed between the two conditions (highlighted by the solid black lines). More specifically, in the left hemisphere, the preferred eccentric positions were higher during the attend right condition (Figure 3.3B, upper panel) than those during the attend left condition for the same voxels (3.3B, lower panel).

We averaged the fMRI responses underlying the preferred eccentric positions from the hemispheres contralateral to the attended target (i.e. left and right hemispheres for the attend right and attend left conditions respectively) and ipsilateral to the attended target (i.e. right and left hemispheres for the attend right and attend left conditions respectively; Figure 3.3C). This revealed that the responses from the contralateral hemispheres were delayed compared to those from the ipsilateral hemispheres. As the stimuli covered higher eccentricities at later time points in the sequence, this delay corresponded to an increase in preferred eccentric position in the hemifield containing the attended target, demonstrating pRF attraction towards the attended target.

We quantified this pRF attraction for every voxel in V1. We measured preferred eccentric positions for every voxel twice, once when it was located in the hemisphere contralateral to the attended target (e.g. attend left right hemisphere voxels) and once when it was located in the hemisphere ipsilateral to the attended target (e.g. attend right for right hemisphere voxels). We subtracted the preferred eccentric positions measured in the ipsilateral hemispheres from those measured in the contralateral hemispheres to give the preferred eccentric position change between the two conditions which measures pRF attraction towards the attended targets for every voxel. This revealed a significant increase in preferred eccentric position in the contralateral hemispheres for every subject separately (Figure 3.3D, two-sided, repeated measures t-test, all p values < 0.001). These preferred eccentric

position changes demonstrate that across V1, voluntary spatial attention attracts pRFs towards its location, as predicted by attention field models.

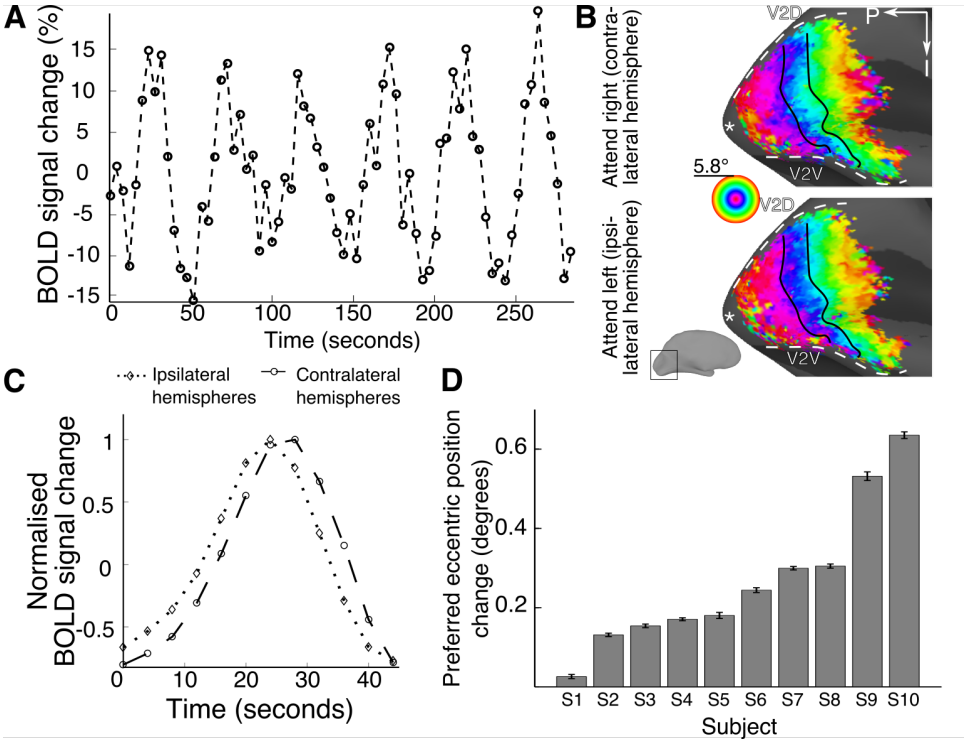

Figure 3.3 Preferred eccentric position changes across V1 **A** fMRI response from one voxel (0.68x0.68x0.70mm) and one condition. We measured six peaks in the fMRI response, corresponding to the stimulus passing through the voxel's pRF six times. **B** Preferred eccentric positions from V1 overlaid on a reconstructed cortical surface of the left hemisphere (inset), for the attend right (upper panel, contralateral hemisphere) and the attend left (lower panel, ipsilateral hemisphere) conditions. Preferred eccentric positions change between the two conditions, as illustrated by the solid black lines). White dashed lines mark the boundary between V1 and V2 ventral (V2v) and V2 dorsal (V2d). The white asterisk marks the foveal representation. The arrows indicate the posterior - anterior (P) and superior - inferior (S) axis. **C** fMRI responses from one subject averaged across stimulus repeats and voxels from the hemispheres ipsilateral to the attended location (dotted line and diamonds) and the hemispheres contralateral to the attended location (dashed line and circles). The fMRI responses from the contralateral hemisphere were delayed relative to the responses from the ipsilateral hemifield. **D** The preferred eccentric positions changed between the two conditions in degrees of visual angle, averaged across V1 entirely, for every subject. Every subject had a significant preferred eccentric position change between the two conditions corresponding to a pRF attraction towards the attended targets. Subjects are sorted by the size of their preferred eccentric position change. Error bars represent the standard errors of the mean.

pRF attraction in V1 is strongest in the deep cortical portion

We hypothesized that feed forward and feedback signals may produce different profiles of preferred eccentric position change across cortical depth in V1 (see Materials and Methods; Figure 3.2). In short, we speculated that if pRF attraction is driven by feed forward signals, this would yield no systematical variation of preferred eccentric position change across cortical depth. If, however, pRF attraction is driven by feedback signals, we would measure stronger preferred eccentric position changes in either deep cortical portions or superficial cortical portions, or both.

We assessed how changes in preferred eccentric position, measuring pRF attraction towards the attended targets, varied across cortical depth in V1. To this end, we acquired high-resolution anatomical images for every subject and computed equi-volume estimates of normalized cortical depth (Waehnert et al., 2014; Figure 3.4A and B and Supplemental Figure 3.1). Next, we imported the estimated preferred eccentric positions (Figure 3.4C and Supplemental Figure 3.1) and fMRI response

amplitude for both conditions into the anatomical space (see Materials and Methods).

We first verified our methods by examining the variation in fMRI response amplitude across cortical depth. As we used a 3D gradient echo (GE) sequence, we should find an increase in fMRI response amplitude with cortical depth (Duvernoy, Delon, & Vannson, 1981; De Martino et al., 2013).

For each subject separately, we divided the depth estimates into ten equally sized bins and computed the averaged fMRI response amplitude for each bin for the contralateral and ipsilateral hemispheres separately and averaged across subjects subsequently. As expected, the fMRI response amplitude increased across cortical depth for both conditions, confirming the validity of our methods and data (Figure 3.4D and Supplemental Figure 3.1). Note that the increase in response amplitude differed between contralateral and ipsilateral hemispheres. We will address these differences in fMRI response amplitude in detail below.

As we found the expected increase in fMRI response amplitude across cortical depth, we assessed how preferred eccentric position changes varied across cortical depth. Again, we divided the depth estimates into ten equally sized bins and computed the averaged preferred eccentric position change per bin for every subject separately (Supplemental Figure 3.1). To average the binned data from all subjects together, we subtracted each subject's averaged preferred eccentric position change (Figure 3.3D) from their binned averages. Next, we averaged the binned averages across all subjects together, weighting each subject's data by the number of voxels they contributed, and added the weighted average preferred position change across all subjects to the averaged binned data (Figure 3.4E). For all subjects combined, we found a significant negative slope of the binned preferred eccentric position changes as across cortical depth (weighted linear regression, slope coefficient = -0.070 degrees of visual angle, $p < 0.001$). Given an averaged preferred eccentric position change across V1 of 0.29 degrees visual angle (Figure 3.3D), this means a decrease of roughly 25% in preferred eccentric position change from deep to superficial cortical portions. This result was the same for subjects whose anatomies were acquired at 0.5 mm resolution (slope coefficient: -0.064 degrees visual angle ($p < 0.001$)) and for subjects whose anatomies were acquired at a lower resolution (slope coefficients: -0.077 degrees visual angle ($p < 0.001$)). In sum, we found that changes in preferred eccentric position induced by spatial attention were larger in the deep cortical portions than in central and superficial cortical portions.

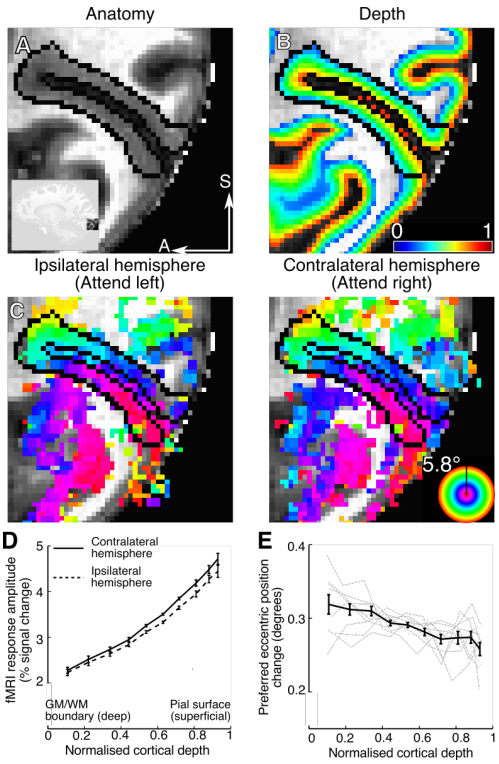

Figure 3.4 Preferred eccentric position changes as a function of cortical depth

A Anatomical image for one of our subjects. The black outline marks the borders of V1. The arrows indicate the anterior - posterior (A) and superior - inferior (S) axes. See Figure S1 for all subjects. **B** The anatomical image in A, but now overlaid with normalized cortical depth estimates. **C** The anatomical image in A, but now overlaid with preferred eccentric position estimates from the attend ipsilateral (left panel) and attend contralateral (right panel) conditions. **D** Averaged fMRI response amplitude across all subjects as a function of cortical depth for the hemispheres contralateral (solid line) and ipsilateral (dashed line) to the attended location **E** Preferred eccentric position changes as a function of cortical depth averaged across all subjects (solid black line), accounted for global difference in mean preferred eccentric position change. Thin gray lines represent the data from individual subjects, corrected for global difference in mean preferred eccentric position change. Error bars in D and E represent the standard error of the weighted mean across subjects per bin, determined by bootstrapping (1000 iterations). We find a significant negative slope across cortical depth, indicating larger preferred eccentric position changes in the deep cortical portion, near the gray matter/white matter (GM/WM) boundary.

Eye movements do not produce variation in pRF attraction across cortical depth

One factor that potentially could confound the results is that subjects made involuntary eye movements towards the attended target during task performance (Figure 3.5A and Supplemental Figure 3.2). These eye movements towards the attended targets could potentially explain the preferred eccentric position changes.

First, we measured subjects' eye-movements in an identical setting outside the scanner. These measured eye movements recorded a bias in horizontal gaze

position towards the attended target of 0.046 degrees visual angle per condition, yielding a total bias of 0.092 degrees between the two conditions. This gaze position bias would produce preferred eccentric position changes in the same direction as the attentional pRF attraction. However, the size of the attentional preferred eccentric position changes (0.29 degrees) are much larger. Furthermore, unlike attentional pRF attraction, the effects of eye-movements are similar across the visual hierarchy (Klein et al., 2014) and theoretically also across cortical depth.

To evaluate whether eye-movements can explain the preferred eccentricity change, we correlated subjects' average horizontal gaze position bias with the preferred eccentricity change in V1. This correlation is significant ($r = 0.66$, $p = 0.039$) but the effect is driven by one outlier (S10). S10's gaze position bias was almost twice the size of all other subjects. Removal of S10, removes the correlations between average gaze position bias and preferred eccentricity change ($r = 0.34$, $p = 0.37$). More importantly, the correlation of the average gaze position with the change of preferred eccentricity change across cortical depth is not significant with or without the outlier. Furthermore, removal of the outlier still maintains our main effect of significant variation of preferred eccentric position changes across cortical depth (slope coefficient $\sim$ 0.06 degrees visual angle, $p = < 0.001$). Therefore, eye-movements can contribute to preferred eccentric position changes, but they do not produce the variation in preferred eccentric position changes across cortical depth in V1.

Furthermore, we conducted a control experiment as well as a simulation to evaluate whether eye-movements can explain the variation in preferred eccentric position across cortical depth. In the control experiment five subjects from the main experiment changed their gaze position between experimental runs rather than changing the location they attended. The total change in gaze position between conditions was 0.2 degrees, which is about twice the size of the gaze position bias between conditions measured prior to the scanning sessions. In the simulation, we introduced a change in gaze position to match the effect size of the attentional modulation (see Materials and Methods).Both measured and simulated eye movement-related preferred eccentric position changes did not produce a significant slope across cortical depth (weighted linear regression, slope coefficients: 0.0086 ($p = 0.469$), Figure 3.5B solid gray line and 0.04283 ($p = 0.369$), Figure 3.5B dashed gray line, for measured and simulated data respectively). Thus, the main experiment but not the control experiments show a significant variation across cortical depth. The latter is the case even if we restrict the main experiment to the subjects that participated in the control experiment.

Finally, the profile of fMRI response amplitude measured in the eye movement control experiment is very similar to the profile from the main attention

experiment for the same subjects (Figure 3.5C, dashed black line and solid black line respectively). This demonstrates that the different profiles of preferred eccentric position change are not likely to be due to differences in fMRI response amplitude profiles between the two experiments.

In sum the correlations with eye-position, control experiments and simulations show that eye movements did not produce the cortical depth dependent effect on preferred eccentric position change we measured in the attention conditions.

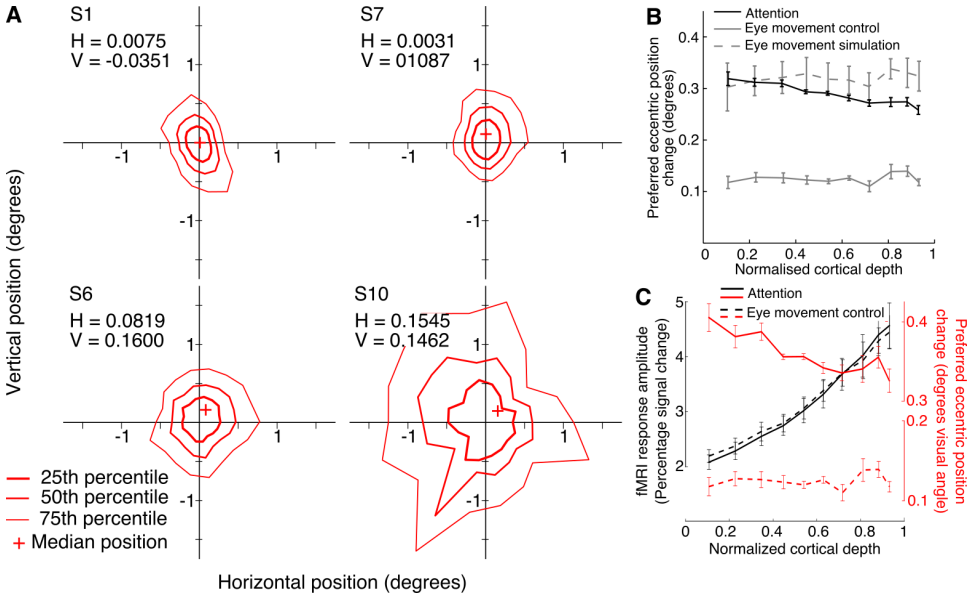

Figure 3.5 Subjects eye movements and their effect on the profile of preferred eccentric position change across cortical depth. **A** Distribution of eye positions relative to the fixation point during task performance for the two subjects with the smallest (S1 and S7) and two subjects with the largest (S6 and S10) gaze position bias. All gaze positions are arranged such that the attended location is always right of the center of the graph, at 6.3° visual angle. Red lines mark the 25th, 50th and 75th percentile of the gaze positions, the plus sign marks the median off set of the distribution. See Supplemental Figure 2 for all subjects. **B** Preferred eccentric position changes produced by attention (solid black line, same data as in Figure 3.4E), compared to preferred eccentric position changes produced by eye movements, both measured (solid gray line) and simulated (dashed gray line), as a function of cortical depth. Whereas attention produced a negative slope, eye movements did not. **C** fMRI response amplitude (black lines) and preferred eccentric position changes (red lines), from the main attention experiment (solid lines) and eye movement control experiment (dashed lines) as a function of normalized cortical depth. The data from the main experiment is from the subjects that were also included in the eye movement control experiment. Despite the similar profiles of fMRI response amplitude for both experiments, the profiles of preferred eccentric position change are very different. All error bars represent the standard error of the weighted mean per bin, across subjects, determined by bootstrapping (1000 iterations).

pRF attraction is independent from fMRI response amplitude

Here we investigate whether changes in response amplitude are responsible for changes in pRF attraction. This is particularly relevant for cortical depth measurements as response amplitude varies with cortical depth (Duvernoy, Delon, & Vannson, 1981; De Martino et al., 2013).

FMRI response amplitude increased with cortical depth (Figure 3.4D) and this increase differed between the contralateral and ipsilateral hemispheres in the attention experiment (Figure 3.6A solid black line; weighted linear regression, slope coefficient: 0.26, $p < 0.001$). Thus, fMRI response amplitude changed in two important ways: 1. it increased from deep to superficial cortical portions, as expected from our GE sequence, and 2. this increase differed between contralateral and ipsilateral hemispheres. If the preferred eccentric position changes are related to fMRI response amplitude, these changes in response amplitude can be a potential confound.

In order to determine whether response amplitude is a potential confound, we assessed the relationship between preferred eccentric position changes and fMRI

response amplitude. Preferred eccentric position changes across cortical depth are negatively correlated to fMRI response amplitude in the attention experiment, with averaged fMRI response amplitude increasing and preferred eccentric position change decreasing towards superficial portions (Figure 3.5C; solid lines). Thus, increased fMRI signals do not yield increased preferred position changes.

Alternatively, if low signal amplitudes would produce larger preferred position changes, we should measure a correlation between these two quantities in the eye movement control experiment as well. However, this is not the case. We do not find such a correlation ($p = 0.47$), despite the similarity in fMRI response profiles between the attention and control experiment (Figure 3.5C, black lines).

In addition, if preferred eccentric position changes are dependent on fMRI response amplitude, variation in response amplitude would bring about variation in preferred eccentric position changes as well. In this case, the difference in fMRI response increases between conditions in the attention experiment (Figure 3.6A) could underlie the profile of preferred eccentric position change (Figure 3.4E). If this were the case, the same would be true for the eye movement control experiment. However, we measured a similar difference in the increase in fMRI response amplitude across cortical depth in this experiment as we did in the attention experiment (Figure 3.6A; weighted linear regression, slope coefficient: 0.094, $p < 0.001$). Moreover, the control experiment revealed no variation in preferred eccentric position change (Figure 3.5B, gray lines).

Finally, we also simulated the effect of fMRI response amplitude on preferred eccentric position changes. We created two fMRI data sets and introduced a preferred eccentric position change between these two sets. Next, we computed the preferred eccentric position change between the sets for a range of response amplitudes (1 - 10 %-percent signal change) and added normal distributed, random noise to the data. We then bootstrapped (1000 iterations) the average preferred eccentric position change and the 95% confidence interval as a function of response amplitude. This simulation did not reveal a systematic bias of preferred eccentric position change as a function of fMRI response amplitude (Figure 3.6B).

In sum, fMRI response amplitude changed across cortical depth as expected, but also differed between conditions. However, this difference was not specific in the attention experiment, but also present in the eye movement control experiment. Importantly, neither our data nor our simulation support the possibility that these changes in fMRI response amplitude would produce the profile of preferred eccentric position changes as measured in the attention experiment.

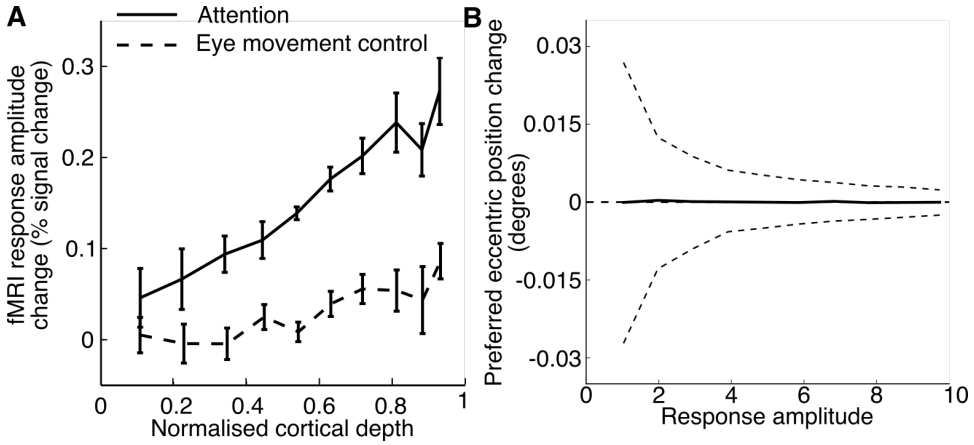

Figure 3.6 fMRI response amplitude changes. **A** fMRI response amplitude change between contralateral and ipsilateral hemispheres for the attention experiment (solid line) and eye movement control experiment (dashed line) as function of cortical depth. For both experiments, we find that response amplitude changes increase with cortical depth. **B** Simulated preferred eccentric position change as a function of response amplitude. See text for details. We bootstrapped the average preferred eccentric position change (solid line) and the 95% confidence interval (dashed lines) as a function of response amplitude. The preferred eccentric position changes are relative to the simulated change, 0.29 degrees, which is the same as the averaged preferred eccentric position change measured in the attention experiment (Figure 3.3D). Error bars in A represent the standard error of the weighted mean per bin, determined by bootstrapping (1000 iterations). This simulation reveals no systematic bias of preferred eccentric position change as a function of fMRI response amplitude.

Discussion

We used sub-millimeter, ultra high-field fMRI to assess attentional pRF attraction across cortical depth in human V1. We measured pRF attraction through changes in preferred eccentric position between two attention conditions that differed in the location at which voluntary spatial attention was directed. We extracted the profile of pRF attraction across cortical depth and found pRF attraction at every cortical portion, although it was largest in the deep cortical portion, near the white/gray matter boundary and decreased towards superficial portions. The results of a control experiment demonstrates that eye movements towards the attended targets yielded preferred eccentric position changes that do not vary across cortical depth

and thus cannot account for the profile we measure. Moreover, our approach focused on attentional modulations of preferred eccentric position, rather than fMRI response amplitude. As such, our results are not confounded by the response amplitude variation across cortical depth resulting from the vascular properties of the cortex (Duvernoy, Delon, & Vannson, 1981; De Martino et al., 2013).

By applying computational models of attention to the known neuroanatomical organization of V1, we hypothesize that a feed forward mechanism would yield pRF attraction that does not vary across cortical depth (Figure 3.2D) and a feedback mechanism would yield pRF attraction limited to either deep or superficial cortical portions, or both (Figure 3.2E). Therefore, we interpret our results as providing evidence that a combination of a feed forward and a feedback mechanisms underying pRF attraction in V1. We propose that the feedback component specifically targets deep cortical portions.

We speculate that response modulations at the level of LGN produce pRF attraction in V1 central cortical portions (Felleman and Van Essen, 1991; McAdams and Maunsell, 1999; O'Connor et al., 2002; McAlonan et al., 2008). Following the flow of feed forward information across cortical depth, this pRF attraction will be inherited by deep and superficial cortical portions (Fitzpatrick, Lund, & Blasdel, 1985; Maunsell & Gibson, 1992; Yoshioka, Levitt, & Lund, 1994; Usrey & Fitzpatrick, 1996; Callaway, 1998; Briggs & Callaway, 2001; Self, Kerkoerle, Supèr, & Roelfsema, 2013; Fracasso et al., 2016). The stronger pRF attraction in deep cortical portions cannot be explained by this feed forward mechanism and we suggest it to be the result of feedback processing. Likely sources of this feedback component are higher visual areas (Rockland and Pandya, 1979; Felleman and Van Essen, 1991).

We find a stronger pRF attraction deeper cortical portions but not in superficial portions. This is an apparent contradiction with the presence of feedback afferents in superficial layers. There are several possible explanations. First, pRF attraction may reflect a specific type of feedback in which deep and superficial afferents may have different functional specializations. In recent years, a variety of cortical depth dependent effects on responses by feedback processing in general and endogenous attention specifically have been reported (Self et al., 2013; Muckli et al., 2015; Kok et al., 2016; Hembrook-Short et al., 2017; Kerkoerle et al., 2017; Nandy et al., 2017; Self and Roelfsema, 2017). The specific targeting of deep cortical portions in pRF attraction is consistent with the overall picture that attentional modulation is selective and differs between cortical layers and cell types (Hembrook-Short et al., 2017; Nandy et al., 2017). Alternatively, attentional modulation across cortical depth may depend on the match between task demands and neural tuning properties. In this reasoning, pRFs that were better suited to perform our contrast discrimination task were attracted more. Possibly, these pRFs may be more dominant in deeper

cortical portions, explaining why we measured stronger pRF attraction here (Hembrook-Short et al., 2017).

One surprising aspect of our results is that pRF attraction is not integrated across cortical depth to yield the same amount of attraction at every depth. Apparently, pRFs are attracted to varying degrees across cortical depth. As a result, the spatial location that produces the strongest response changes from one cortical portion to the other. From the perspective of the computational aims in V1, this may seem counterproductive. Although we do not know what the computational consequences of this result are, we have reported a similar effect across the visual hierarchy previously (Klein et al., 2014). Here, pRF attraction varied between different visual field maps, apparently misaligning pRFs between different stages of the hierarchy.

The attention field model predicts that pRF attraction is a function of pRF size and attention field size (Equation 3.4). Specifically, larger pRFs will produce a stronger attraction. PRF size will vary with eccentricity (Hubel and Wiesel, 1962; Dumoulin and Wandell, 2008) and cortical depth (Fracasso et al., 2016). However, we showed that pRF attraction only varies with pRF size if the attention field directly interacts with the pRFs, pRF attraction does not vary with pRF size if the attraction is inherited. Nevertheless, one could ask whether we can measure pRF attraction as a function of pRF size. Unfortunately, we cannot. First, we focused on measuring pRF position by using an expanding ring stimulus, which is not suitable to reliably measure pRF size (Dumoulin and Wandell, 2008). Furthermore, due to the expanding ring stimulus we cannot measure pRF positions outside the stimulus range. pRFs that are centered beyond our stimulus range, but still overlap with some of the stimulus' positions, will appear as if they lie at the edge of our stimulus. This property limits the preferred eccentric position changes we can measure for pRFs near the stimulus edge and complicates interpreting the profile of pRF attraction across eccentricity. Note that this stimulus artifact does not limit the overall pRF attraction that can be measured, which still allows us to draw conclusions about the cortical depth dependency of pRF attraction.

We find a large inter-subject variability of preferred eccentric position change across V1 (Figure 3.3D). This variability may have several origins. First, we know that pRF size typically varies between subjects by a factor of at least 2 (Harvey and Dumoulin, 2011). Therefore, variation in pRF size between subjects is likely to contribute to the variation in pRF attraction between subjects (Figure 3.2D). Second, variation in attention field size between subjects can produce variation in pRF attraction. Although we tailored task difficulty to yield similar performance across subjects, subjects may still display different task performance and effort. Finally,

confounding factors, such as the variability in fixation bias can also contribute to the variability in measured pRF attraction - but not as a function of cortical depth.

We observed a decrease in pRF attraction across cortical depth, which we assessed assuming a linear relation between cortical depth and pRF attraction (Figure 3.4D). However, we hypothesized that a feedback contribution to pRF attraction specifically in deep cortical portions would manifest as a stronger pRF attraction in this portion followed by a reduced, constant attraction across central and superficial cortical portions (Figure 3.2E). We emphasize, however, that the aim of our hypothesized profiles was give a qualitative overview of the expected results, not to predict the exact shape of pRF attraction across cortical depth. Methodological issues related to fMRI, such as partial volume effects and the BOLD spread function, will smooth the profile of pRF attraction and obscure its exact shape across cortical depth.

We also found that fMRI response amplitude changed in the attention experiment in two main ways: 1. it increased from deep to superficial cortical portions and 2. this increase differed between contralateral and ipsilateral hemispheres. The difference in increase between the contralateral and ipsilateral hemispheres seems to suggest that spatial attention increases fMRI responses near the attended location in a cortical depth dependent manner. However, as we found a similar profile for the eye movement control experiment, we cannot attribute this effect to spatial attention. Importantly, data from the eye movement control experiment and an additional simulation demonstrates that preferred eccentric position changes are independent from fMRI response amplitude. As such, changes in response amplitude do not underlie the profile of preferred eccentric position change in the attention experiment.

Finally, we have several reasons to exclude methodological issues concerning sub-millimeter fMRI, such as head motion and misalignment, as a possible explanation for our results. First, we collected all the experimental data for each subject in a single scanning session, with the left and right conditions alternating between scans. As such, the data from both attention conditions are affected similarly by head motion and distortions of the functional volumes. Second, we used the same alignment between the functional and anatomical images for both left and right experimental conditions. Although we took great care to coregister the anatomical and functional volumes as accurate as possible (Figure 3.4 and Supplemental Figure 3.1), some coregistration inaccuracies may still be present. In that case, these inaccuracies would affect the data from both conditions equally. We point out that the fMRI response profile measured for the attention experiment is very similar to that of the control experiment (Figure 3.5C, black lines). This demonstrates that our approach is accurate enough to yield highly reproducible

outcomes. Third, the data for the eye movement control experiment was acquired, pre-processed and analyzed in the same way as the data for the main experiment. However, in contrast to the main experiment, the control experiment did not reveal any significant variation of preferred eccentric position change across cortical depth, demonstrating that this variation is specific to the attention conditions in the main experiment.

Conclusions

In conclusion, we examined the influence of voluntary spatial attention on pRF positions across cortical depth in human V1. As we specifically focused on pRF position attraction, we avoided that our results would potentially be confounded by factors such as fMRI response amplitude. We observe pRF attraction in every cortical portion (deep, center and superficial) with the attraction being largest in the deep cortical portion, near the gray/white matter boundary. We speculate that this profile is best explained by a combination of a feed forward and a feedback mechanism underlying pRF attraction, with the feedback component operating stronger in deep cortical portions. As such, our study highlights the utility of high-resolution functional imaging in providing insights in processes underlying attentional modulations of responses in early visual cortex.

Supplementary material

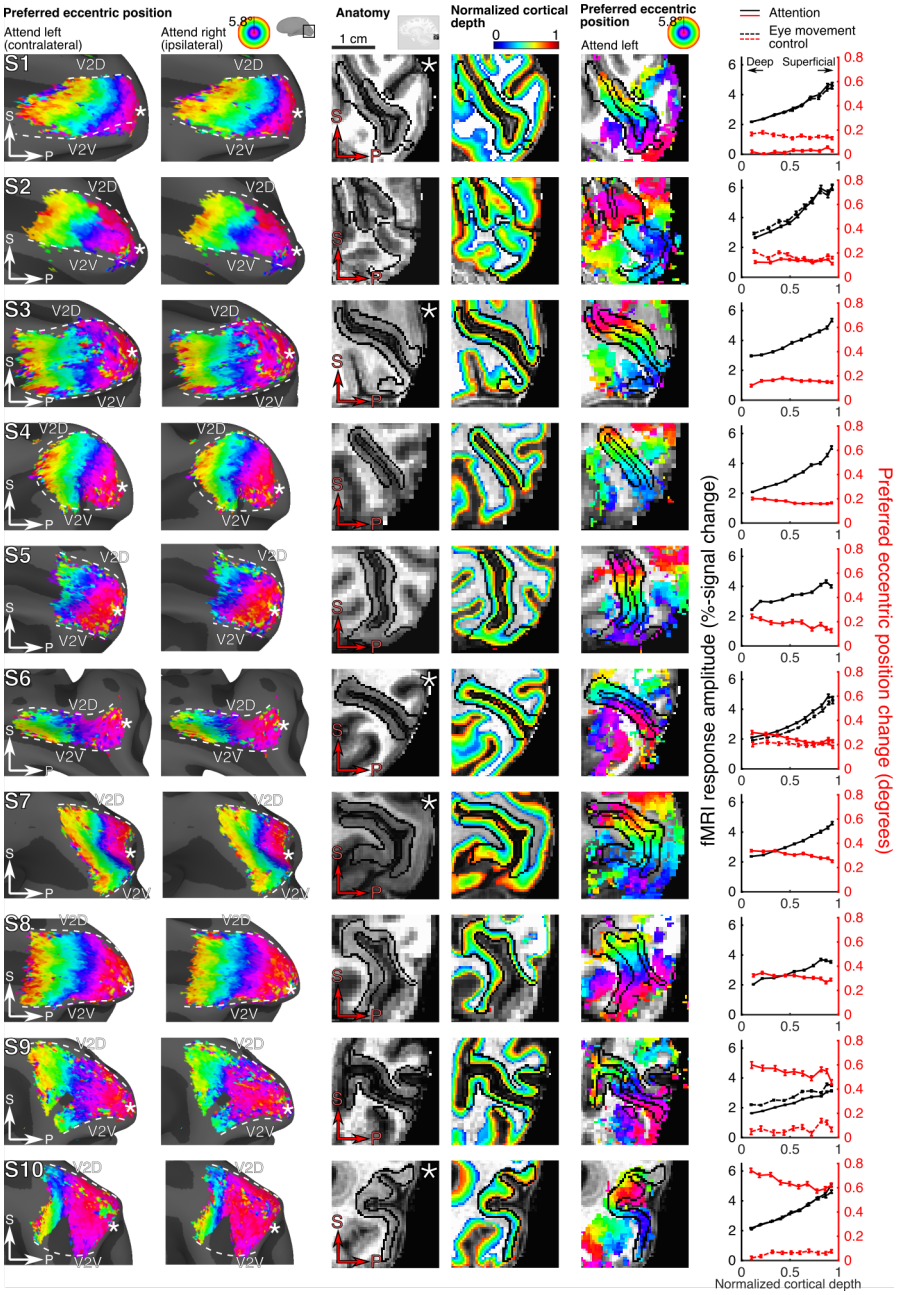

Figure S3.1 Data from all individual subjects. Preferred eccentric positions from V1 overlaid on a reconstructed cortical surface of the right hemisphere (inset), for the attend left (**First column**, contralateral hemisphere) and the attend right (**Second column**, ipsilateral hemisphere) conditions. Preferred eccentric positions gradually increase along the posterior - anterior axis. White dashed lines mark the boundary between V1 and V2 ventral (V2v) and V2 dorsal (V2d). The white asterisk marks the foveal representation. The arrows indicate the posterior - anterior (P) and superior - inferior (S) axis. **Third column** Anatomical images for all of our subjects. The black outline marks the borders of V1. The arrows indicate the anterior - posterior (P) and superior - inferior (S) axes. Scale length indicates 1cm. **Fourth column** The anatomical image in the third column, but now overlaid with normalized cortical depth estimates. Normalized depths range between 0 and 1, where 0 marks the gray/white matter boundary (deep) and 1 the pial surface (superficial). **Fifth column** The anatomical image in the third column, but now overlaid with preferred eccentric position estimates from the attend left condition. **Sixth column** fMRI response amplitude (left axis, black) and preferred eccentric position change (right axis, red) as a function of normalized cortical depth from the main experiment (solid lines) and the eye movement control experiment (dashed lines, S1, S2, S6, S9 and S10). As expected, fMRI response amplitude increase with cortical depth, whereas preferred eccentric position changes tend to decrease with cortical depth (main experiment) or are relatively constant across cortical depth (eye movement control). Error bars represent the standard error of the mean per bin.

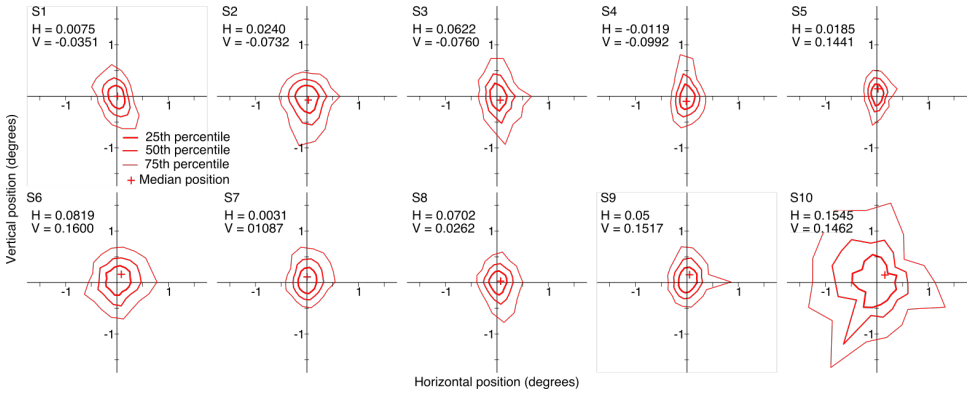

Figure S3.2. Gaze position bias for individual subjects
Distribution of gaze positions relative to the fixation point during task performance for every subject separately. All gaze positions are arranged such that the attended location is always right of the center of the graph, at 6.3° VA. Red lines mark the 25th, 50th and 75th percentile of the gaze positions, the plus sign marks the median off set of the distribution.

Chapter 4

Perceptual consequences of attentional magnification

Published as:
Klein, B. P., Paffen, C. L. E., te Pas, S. F., & Dumoulin, S. O. (2016). Predicting bias in perceived position using attention field models. *Journal of Vision,* 16(7):15, 1–15, doi:10.1167/16.7.15

BPK, CLEP, SFP and SOD designed research. BPK collected the data. BPK analyzed the data. BPK, CLEP, SFP, and SOD wrote the paper.

Abstract

Attention is the mechanism through which we select relevant information from our visual environment. We have recently demonstrated that attention attracts receptive fields across the visual hierarchy (Klein et al., 2014). We captured this receptive field attraction using an attention field model. Here, we apply this model to human perception: We predict that receptive field attraction results in a bias in perceived position, which depends on the size of the underlying receptive fields. We instructed participants to compare the relative position of Gabor stimuli, while we manipulated the focus of attention using exogenous cueing. We varied the eccentric position and spatial frequency of the Gabor stimuli to vary underlying receptive field size. The positional biases as a function of eccentricity matched the predictions by an attention field model, whereas the bias as a function of spatial frequency did not. As spatial frequency and eccentricity are encoded differently across the visual hierarchy, we speculate that they might interact differently with the attention field that is spatially defined.

Introduction

Visual attention is the mechanism through which we select relevant information from our visual environment. Selection of visual information can be based on several of its features, one of which is the location of the information in the visual field, also known as visual spatial attention. Recent models of visual attention describe attention as a Gaussian component that amplifies neural responses (i.e., the attention field). The neural responses in the absence of attention are typically also conceptualized as Gaussians. Consequently, the effect of attention on neural processing is modeled as a Gaussian multiplication representing response amplification (Womelsdorf et al., 2008; Reynolds and Heeger, 2009). This multiplicative amplification may be followed by a divisive process representing response normalization (Reynolds and Heeger, 2009)

Using functional magnetic resonance imaging (fMRI), we have recently shown that attention attracts receptive fields across the visual hierarchy in humans (Klein et al., 2014). We captured this attraction with an attention field model that includes a multiplication between two Gaussian components (i.e., the attention field and the receptive fields) without normalization (Klein et al., 2014). This revealed that the size of the attention field was relatively constant across the visual cortex even though the amount of receptive field attraction increased up the visual hierarchy. As the attention field model predicts a larger attraction for larger receptive fields, we suggested that the increase in the receptive field size up the visual hierarchy was the dominant cause of the increase in receptive field attraction.

On a perceptual level, attention biases the perceived position of stimuli outside the focus of attention away from the attended location (Suzuki and Cavanagh, 1997; Shim and Cavanagh, 2005). Assuming that the perceived position of a stimulus is related to its response distribution across a population of neurons tuned for spatial position, a shift of the distribution away from the attended location may produce a bias in perceived position. Suzuki and Cavanagh (1997) proposed three mechanisms through which attention could produce such a shift: (a) response suppression, (b) receptive field shrinkage, and (c) receptive field attraction. Our attention field model, as discussed above, offers a mathematical formalization of receptive field attraction. Therefore, it provides a quantitative framework to assess the contribution of receptive field attraction to the positional bias. We propose that attention attracts receptive fields, moving them from their original location to a position closer to the focus of attention. Consequently, any stimulus located near the focus of attention is now covered by receptive fields that were originally positioned further away from the attended location. This shifts the stimulus's response distribution away from the attended location, which is accompanied by a shift in the

perceived position of the stimulus away from the attended location. As our attention field model predicts a larger attraction for larger receptive fields, we expect a larger positional bias when the underlying receptive fields are larger. Specifically, as receptive field size increases with eccentricity (Hubel and Wiesel, 1962; Van Essen et al., 1984; Dumoulin and Wandell, 2008; Harvey and Dumoulin, 2011) and tuning to lower spatial frequencies (Jones and Palmer, 1987) we expect that attention induces a larger positional bias for stimuli presented at higher eccentricities as well as stimuli containing lower spatial frequencies.

We examined the magnitude of the positional bias as a function of stimulus eccentricity and spatial frequency in two experiments, using an exogenous cueing paradigm to manipulate the focus of attention (Posner, 1980; Carrasco et al., 2004). We measured positional biases by instructing participants to compare the distance between two pairs of Gabor stimuli, one of which was centered on the previously cued location. Based on our model, we expect that when the cued location falls between the two Gabors making up the pair, both Gabors are perceived to be further away from the cued location. This in turn would lead to an increase in the perceived distance between the Gabors of the cued pair. In the first experiment, we varied the eccentricity at which the cue and Gabor pairs were presented and found that positional biases indeed increased with eccentricity, as expected from our model. In the second experiment, we varied the spatial frequency of the Gabor stimuli but did not find any variation in the positional biases across spatial frequency. In sum, the attention field model can explain the bias in perceived position of stimuli outside the focus of attention as a function of eccentricity, but not spatial frequency. As spatial frequency and eccentricity are encoded differently across the visual hierarchy, we speculate that they might interact differently with the spatially defined attention field.

Methods

Experiment 1: Does the bias in perceived position scale with eccentricity?

The attention field model we used to capture receptive field attraction in human participants (Klein et al., 2014) predicts a larger attraction for larger receptive fields. Because it is known that the average receptive field size increases with eccentricity (Hubel and Wiesel, 1962; Van Essen et al., 1984; Dumoulin and Wandell, 2008; Harvey and Dumoulin, 2011) we expect a larger positional bias for stimuli presented at higher eccentricities.

Participants. Six participants participated in this experiment (all men, age range 26–33). Five participants were naive as to the purpose of the experiment. All participants had normal or corrected-to-normal visual acuity and gave informed consent. Four participants in this experiment participated in at least one other experiment as well.

Experimental setup. Stimuli were generated in Matlab (MathWorks, Natick, MA) using the PsychToolbox (Brainard, 1997; Pelli, 1997) on a PowerMac G5 with an ATI Radeon 9800 XT graphics card. The stimuli were presented on a LaCie 321 monitor (native resolution: 1,600 3 1,200; refresh rate: 60 Hz), with a mean luminance background of ~75 cd/m^2. Participants viewed the stimuli from a distance of 90 cm while their head was supported by a chin rest.

Stimuli and task. Throughout the experiment, a black cross (0.2° x 0.2°) positioned at the center of the screen served as a fixation point. Each trial started with a precue period (lasting 750 ms) during which only the fixation cross was visible. On half of the trials, this period was followed by an exogenous cue (black dot, 0.25° in diameter), which flashed on the screen for 50 ms, at 3°, 6°, or 9° either left or right from the fixation cross (cue trials). On the remaining half of the trials, only the fixation cross was presented (no cue trials). After an interstimulus interval (ISI) of 67 ms, four Gabors (45% Michelson contrast), arranged in two pairs, were presented. The two Gabors within a pair were horizontally aligned, and each pair was centered on the horizontal meridian. During cueing trials, the pairs were presented at the same distance from the fixation cross as the cue. The Gabors consisted of a Gaussian filtered (standard deviation: 0.33°, truncated at 1.8°) and horizontally oriented cosine grating (4.75 cycles per degree [cpd]; Figure 4.4.1A). The vertical center-to-center distance between the Gabors within one of the two pairs was fixed at 3° (standard pair), whereas the vertical distance within the other pair was systemically varied (test pair, see below). After presentation of the Gabor pairs, the fixation cross turned white, prompting the participant to respond, after which the next trial started. Both the standard and test pair had a 50% chance of being presented at the previously cued location (cue standard or cue test, respectively). In total, there were 2 (cue/no cue) x 2 (cue standard/cue test) x 2 (cue left/cue right) x 3 (distance from fixation)
= 24 conditions.

Participants had to report which of the two Gabor pairs they perceived as vertically closer together by pressing either ‘‘F’’ or ‘‘J’’ for the left or right Gabor pair, respectively. Note that we expect the Gabor pair on the cued side to be perceived further apart (see model below): Participants with a response bias for the cued side

would thereby be more likely to show an effect opposite to the expected direction. The vertical distance between the test pair was controlled by a Quest staircase (Watson and Pelli, 1983) for each of the 24 conditions separately. Each staircase terminated after 20 trials and was used to initially estimate the point of subjective equality (PSE): the distance at which the test pair appeared as far apart as the standard pair, for each condition separately. Subsequently, 20 trials were presented for which the vertical distance between the test pairs was randomly jittered around the estimated PSE. The amount of jitter was randomly sampled from a uniform distribution, limited between -0.4° and 0.4°. In sum, 40 trials were presented per condition, resulting in 960 trials per session. Participants completed two sessions, totaling 1,920 trials per subject.

Analysis. We combined the presented vertical distances between the test pair and the participant's corresponding responses from the Quest and jitter trials and collapsed the data from both sessions. We fitted cumulative Gaussians (the distance x at p = 0.5: μ and the slope: σ) to each of the no-cue conditions. The value for μ was taken as the distance at which the test pair appears as far apart as the standard pair (PSE). Using bootstrapping (1,000 iterations), we generated a distribution of both fit parameters (μ and σ) and computed the median PSE and slope using these distributions. The median PSE served as a baseline measurement for each cue condition separately. Next, we corrected for any possible bias in perceived distance between the Gabor stimuli that was not due to cueing by subtracting the baseline PSEs obtained for each no-cue condition from the distances presented in the corresponding cue condition. Subsequently, we collapsed the vertical distances and responses from the cue left and cue right conditions. We then fitted cumulative Gaussians to the corrected distances and responses of each cue condition (cue test and cue standard conditions separately) and bootstrapped (1,000 iterations) the median and 95% confidence interval of both fit parameters. In addition, we collapsed the cue test and cue standard conditions together and fitted cumulative Gaussians to the collapsed data. Collapsing was done in such a way that a positive value indicates an increase in perceived distance between the Gabor stimuli due to cueing. Again, bootstrapping was used to obtain the median and 95% confidence interval of the PSE and slope.

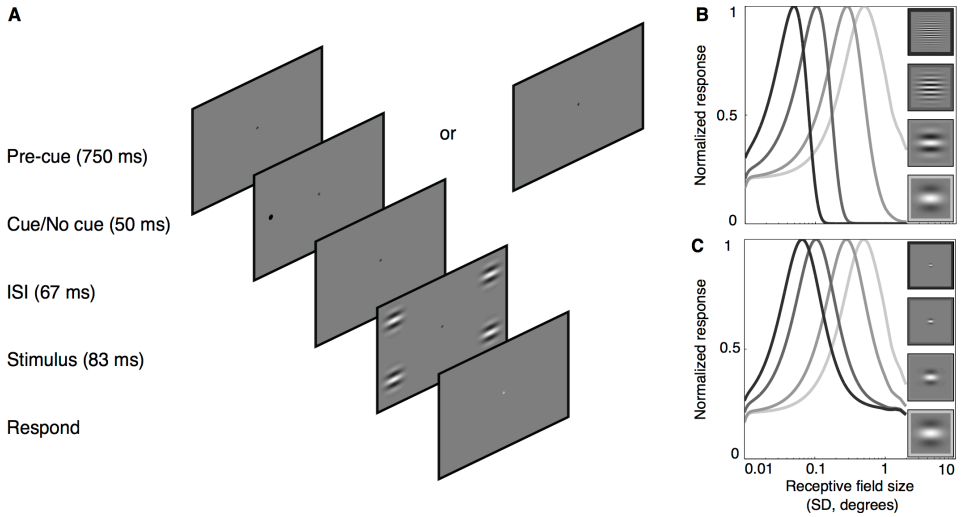

Figure 4.1. Stimulus and experimental procedure
A During the precue period (750 ms), only the fixation cross was presented. On half of the trials, this period was followed by a cue (50 ms), left or right of fixation. On the remaining half of the trials, only the fixation cross was presented during this period. After an interstimulus interval (ISI; 67 ms), two Gabor pairs were flashed on the screen (83 ms), left and right of fixation, centered on the cued location and the location opposite to that. After this period, the fixation cross turned white to prompt the participant to respond. Participants had to report which of the two Gabor pairs they perceived closer together (vertically) by pressing the ''F'' or ''J'' button for left and right, respectively. **B** Estimated responses to the Gabors used in Experiment 2 (insets) across a range of receptive field sizes (standard deviation). Responses are computed by convolving the Gabors with a filter model of spatial frequency–selective receptive fields. **C** Similar to B but now for the Gabors used in one of the control experiments (insets). The ratio between the spatial frequency and Gaussian standard deviation of the Gabors is kept constant, resulting in a constant width of the estimated responses across receptive field size.

Experiment 2: Does the bias in perceived position scale with spatial frequency?

Tuning to lower spatial frequencies is related to larger receptive fields (Jones and Palmer, 1987). As the attention field model predicts a larger attraction for larger receptive fields, we expect that a positional bias resulting from this attraction increases with lower spatial frequencies as well.

Participants. Six participants participated in this experiment (1 woman, age range 25–30). All participants were naive as to the purpose of this experiment. In addition, two participants completed an experiment designed to estimate the contrast differences needed between spatial frequencies to correct for differences in contrast sensitivity (see below). One participant was naive as to the purpose of this experiment. All participants had normal or corrected-to-normal visual acuity and gave informed consent. All participants participated in at least one other experiment.

Stimulus and task. This experiment was similar to the eccentricity experiment, but now the cue and all Gabor pairs were presented at 3° distance, left and right from fixation. In addition, the stimuli were viewed from a 180-cm distance instead of 90 cm. The Gabors varied in spatial frequency between trials: either 1, 1.75, 4.75, or 10 cpd. Both pairs within the same trial had the same spatial frequency. This 10-fold range (~3.3 octaves) in spatial frequencies ensured stimulation of different-sized receptive fields (DeValois et al., 1982). We verified this by convolving the Gabors with a filter model resembling spatial frequency selective receptive fields (Jones and Palmer, 1987, Figure 4.1B). The contrasts of the Gabors were 17%, 21%, 33%, and 60% Michelson for the 1, 1.75, 4.75, and 10 cpd, respectively. These values were chosen to make all Gabors appear similar in contrast and were estimated in a separate experiment (see below). This experiment had 32 conditions and a total of 2,560 trials. Participants completed the experiment in two sessions.

To account for differences in perceived contrast across the range of spatial frequencies used in this experiment, we adjusted the contrast of the 1, 1.75, and 4.75 cpd Gabors to appear similar in contrast to that of the 10 cpd Gabor, which was fixed at 60% Michelson. The contrasts needed to achieve this were estimated in a separate experiment. This experiment used a similar procedure as the main experiments but differed in several ways. First, a cue was never presented prior to the Gabor pairs. Second, the Gabor pairs within one trial differed in spatial frequency and contrast: One was always 10 cpd and set to 60% Michelson contrast (standard pair), and the other pair was one of the remaining spatial frequencies, and its contrast was systematically varied (test pair). Participants were instructed to report which of the two Gabor pairs appeared higher in contrast by either pressing ''F'' or ''J'' for the left or right Gabor pair, respectively. Initially, Quest staircases were used to control the contrast of the test pair and estimate the contrast at the PSE after 25 trials, for each condition separately. Subsequently, 75 trials were presented in which the amount of contrast of the test pair was randomly jittered around the PSE estimated for the corresponding condition. The amount of jitter was randomly sampled from a uniform distribution, limited between -15% and 15% Michelson. This experiment

had 2 (standard left/standard right) x 3 (each of the three remaining spatial frequencies) = 6 conditions. One hundred trials were presented for each condition, totaling 600 trials per experiment. Two participants completed the experiment. Data were analyzed by collapsing the responses and contrasts from the standard left and standard right conditions together, fitting cumulative Gaussians to the collapsed data and bootstrapping the median PSE values.

Control Experiment 1: Does the size of the Gabor influence the results?

In Experiment 2, we increased the spatial frequency of Gabors that had a constant size in the spatial domain. Doing so introduces more cycles of the cosine with higher spatial frequencies, narrowing the Gabors in the frequency domain. In this control experiment, we examine whether this confounding factor may have affected the results in Experiment 2 by fixing the size of the Gabors in the frequency domain and thus varying their size in the spatial domain.

Participants. Six participants participated in this experiment (1 woman, age range 25–31). Five participants were naive as to the purpose of this experiment. In addition, two participants completed an experiment designed to estimate the contrast differences needed between spatial frequencies to correct for differences in contrast sensitivity. One participant was naive as to the purpose of this experiment. All participants had normal or corrected-to-normal visual acuity and gave informed consent. All participants participated in at least one other experiment.

Stimulus and task. This experiment was similar to the spatial frequency experiment, except that the standard deviation of the Gabor's Gaussian filter was a set to one third of the Gabor's wavelength. This reduces the size of the Gabors with increasing spatial frequency and made the 10 cpd Gabors difficult to observe. We replaced this condition with the 7.5 cpd condition, reducing the range of spatial frequencies to ~2.9 octaves. This range is still sufficient to stimulate different-sized receptive fields between conditions (DeValois et al., 1982). We verified the size of the Gabors and the separation of the peaks in the frequency domain by applying the filter model (Jones and Palmer, 1987, Figure 4.1C).The Gabors' standard deviations were 0.33°, 0.19°, 0.07°, and 0.04°, and their contrasts were 23%, 26%, 41%, and 60% Michelson for 1, 1.75, 4.75, and 7.5 cpd, respectively. The contrast values were chosen to make all Gabor stimuli look similar in contrast and were estimated in an experiment identical to the contrast experiment conducted prior to Experiment 2. Participants completed 2,560 trials for this experiment, which were distributed across two sessions.

Control Experiment 2: Was spatial attention prompted by presenting a cue?

Effects of exogenous attention on perception should peak about 80 to 130 ms after the onset of the cue (Nakayama and Mackeben, 1989; Cheal and Lyon, 1991). Our initial ISI (67 ms) ensured a stimulus presentation in this window. If attention does indeed underlie our results, effects of cuing should become weaker on both shorter and longer intervals. Note that we do not expect the effect of cueing to completely disappear on both shorter and longer intervals. The positional bias can still be present even when the cue and stimulus are simultaneously presented or the cue and stimulus are separated by 1500 ms (Suzuki and Cavanagh, 1997). The important thing to notice here is that the increase in perceived distance due to cueing should be reduced for shorter and longer intervals. In this experiment, we verify this prediction.

Participants. Six participants participated in this experiment (1 woman, age range 22–31). Five participants were naive as to the purpose of this experiment. All participants had normal or corrected-to-normal visual acuity and gave informed consent. Four participants participated in at least one other experiment.

Stimulus and task. This experiment was similar to Experiment 1, but now the cue and Gabor stimuli were always presented at 3° left and right from fixation and the interval between cue offset and stimulus onset was either 0 ms, 67 ms, or 600 ms. The contrast was set to 33% Michelson. Participants completed two sessions, totaling 1,920 trials per subject.

Attention field model

The attention field model is summarized in Equation 4.1 below and is identical to the model that captures receptive field attraction in human subjects (Klein et al., 2014) and in macaques (Womelsdorf et al., 2008). The model consists of a multiplication between two Gaussian components but has no normalization process as in the attention field model proposed by Reynolds and Heeger (2009)

In the model, the onset of the cue attracts attention toward its location, which we model as a Gaussian-shaped attention field at the cued location. This attention field attracts receptive fields toward the cued location. We model the receptive field as a Gaussian as well; therefore, the receptive field attraction is given by a Gaussian multiplication:

$$\text{RF attraction} = \mu_{or} - \frac{\mu_{or}\sigma_{af}^2}{\sigma_{af}^2 + \sigma_{rf}^2} \qquad \text{Equation 4.1}$$

where σ_{rf} and σ_{af} represent the size of the receptive field and the size of the attention field, respectively (i.e., the standard deviation of the Gaussian components). Furthermore, μ_{or} represents the receptive field's original location relative to the attention field.

All parameters in Equation 4.1 needed to obtain a prediction for the receptive field attraction, except for the attention field size, can be derived from our stimulus design and data. We model the original position of the receptive field relative to the attention field as the center-to-center distance between the cue and one of the Gabors, plus the receptive field attraction (i.e., half the increase in perceived distance measured for a specific condition). Estimating the receptive field sizes underlying the increase in perceived distance requires us to make assumptions about the relation between receptive field size and eccentricity and the relation between receptive field size and spatial frequency tuning. Receptive field size increases with eccentricity (Hubel and Wiesel, 1962; Van Essen et al., 1984; Dumoulin and Wandell, 2008; Harvey and Dumoulin, 2011) which we model as a linear function:

$$\text{Receptive field width} = a + bx \qquad \text{Equation 4.2}$$

where x stands for the eccentricity. Furthermore, a and b are the intercept and slope coefficients, which we estimate to be 0.58 and 0.15, respectively. This estimate is based on data from V1, obtained in our fMRI study (Klein et al., 2014). Spatial frequency tuning is related to receptive field size: Larger receptive fields are tuned to lower spatial frequencies. Using a biological plausible relation between receptive field size and spatial frequency tuning (Jones and Palmer, 1987) we obtain an estimate of the receptive field size tuned to the spatial frequency in every spatial frequency condition:

$$\text{Receptive field width} = \frac{1}{\omega 3.3} \qquad \text{Equation 4.3}$$

where ω represents the spatial frequency of the Gabors. As we know the receptive field attraction in every condition (half of the increase in perceived distance) and

have an estimate of the receptive field size and original receptive field location in every condition, we can compute the attention field size necessary to produce the increase in positional bias in every condition separately:

$$\sigma_{af} = \sigma_{or}\sqrt{\frac{\mu_{at}}{RF\ attraction}} \qquad \text{Equation 4.4}$$

We assess how model predictions for the increase in perceived distance compare to the positional biases we measure. As we do not have an independent estimate of the attention field size, we choose to make a qualitative rather than quantitative comparison. We compute the attention field size needed to produce the increase in perceived distance measured in every condition of the eccentricity and spatial frequency experiment (Equation 4.4). We use these attention field sizes to compute the receptive field attraction across a range of receptive field sizes (Equation 4.1) and multiply this by 2 to give the predicted increase in perceived distance. Doing so illustrates how well model predictions using the attention field size measured in one condition compare to perceived distance increments in the remaining conditions.

Results

The attention field model predicts an increase in perceived distance between the Gabors on the cued side

Our predictions are based on the following logic: The onset of the cue will attract attention to its location, leading to a receptive field attraction toward the cued location. As a consequence of the attraction, receptive fields whose original location was at the side of the Gabor's center opposite to the cued location (original RF location, Figure 4.2A) will move to a location that is aligned with the Gabor's center (RF location with attention field, Figure 4.2A). This receptive field attraction moves the cortical representation (solid and dashed circles, Figure 4.2B) of both Gabors within the pair away from the cued location (change in cortical location, Figure 4.2B). This change in cortical location increases the distance between the representations of the Gabors within the pair, leading to an increased perceived distance between the Gabors on the cued side. An attention field model predicts that this increase in perceived distance depends on the size of both the attention field and the receptive fields underlying the positional bias (Figure 4.2C). In sum, the

attention field model predicts larger increases in perceived distance with larger receptive fields and smaller attention fields.

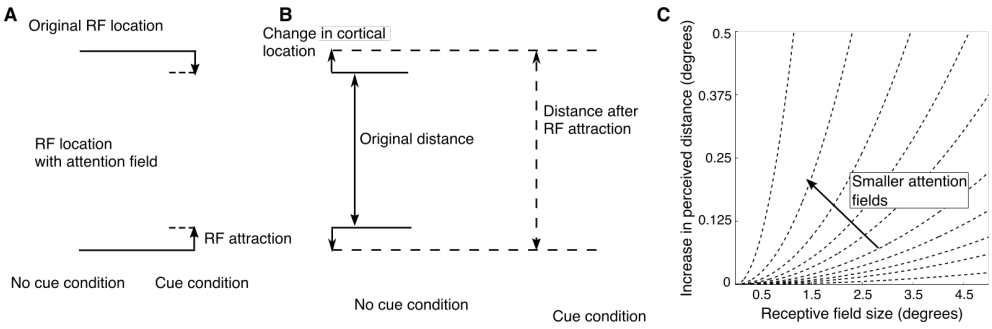

Figure 4.2. Proposed model. **A** When no cue is presented, receptive fields are aligned with their original location (solid circles). The onset of a cue attracts attention toward its location, which we model as a Gaussian-shaped attention field at the cued location (black dot). The interaction between the attention field and receptive field results in a receptive field attraction toward the cued location (RF attraction). **B** Without a cue, the cortical representation of the Gabors (solid circles) corresponds to the location at which the Gabors were presented. Consequently, the distance between the representations of Gabors within a pair corresponds to their physical distance in visual space (original distance). Because of the receptive field attraction toward the attention field (black dot), receptive fields whose original location is further away from the cued location now cover the Gabors. As a consequence, the cortical representations move away from the cued location at the center of the Gabor pairs (change in cortical location), resulting in an increase in the distance between the cortical representations (distance after RF attraction). We propose that this increase in distance between the cortical representations leads to an increase in perceived distance between the Gabors on the cued side. **C** The attention field model predicts an increase in perceived distance between the Gabors on the cued side. The predicted increases are larger for larger receptive fields (x axis) and smaller attention field sizes (separate lines).

Experiment 1: The bias in perceived position scales with eccentricity

To assess the effect of cueing on the perceived distance between the two Gabors within a pair, we fitted cumulative Gaussians to the corrected distances (see the Methods section) and responses from the cue conditions. Figure 4.3A shows the median fit and the 95% confidence intervals obtained by bootstrapping for the cue standard and cue test conditions separately. The data are taken from one subject and one eccentricity condition (6°). We plot the probability that the participant reports seeing the test pair closer together than the standard pair as a function of the difference in distance between the test pair and standard pair. As expected, the curve fitted for the cue test condition (black lines) is shifted to the left, demonstrating that when presented at a previously cued position, the test pair has to be closer together to appear as far apart as the standard pair. An effect in the opposite direction is observed for the cue standard condition (red lines). This demonstrates that the participant perceived the Gabor pair on the cued side farther apart than the pair on the noncued side. Figure 4.3B shows the difference in distance between the test and standard pair at the PSE and the 95% confidence intervals for both the cue test and cue standard conditions (dark and light gray bars, respectively) for all three eccentricity conditions for the same participant. The effects of cueing on the perceived distance between the Gabor pairs are in the expected direction for both the cue test and cue standard conditions. To assess the magnitude of the effect of cueing across eccentricity conditions, we collapsed the cue test and cue standard conditions (Figure 4.3C, D) for this participant. Figure 4.3C shows the probability that the participant reports seeing the noncued pair closer together than the cued pair, as a function of the difference in distance between the cued and noncued pair. The curve is shifted to the right, meaning that the noncued pair had to be further apart than the cued pair, to appear equally far apart. Figure 4.3D shows the difference in distance between the cued and noncued pair at the PSE for the collapsed data for all eccentricity conditions. We interpret this as the magnitude of the effect of cueing on the perceived distance between the Gabor pair across eccentricity. Figure 4.3E shows the magnitude of the effect of cueing on the perceived distance between the Gabor pair across eccentricity, for all participants separately (gray lines) and the average across all participants (solid line), together with the standard error of the mean (SEM, gray area). A one-way repeated-measures analysis of variance (ANOVA) revealed a significant effect of eccentricity on the magnitude of the cueing effect, $F(2, 10) = 7.98$, p , 0.01. A linear regression confirmed a significant increase of the cueing effect with eccentricity (slope coefficient = 0.0244, $p = 0.01$).

To illustrate how model predictions for increases in perceived distance resemble the perceived distance increments measured in this experiment, we compute the attention field size necessary to produce the increase in perceived distance for every eccentricity condition separately (Equation 4.4). The computed attention field sizes are 4.54° for the 3° condition and 6.10° for both the 6° and 9° condition. We used these attention field sizes to compute the predicted increase in perceived distance across a large range of receptive field sizes. Figure 4.3F presents the measured effect of cueing on the perceived distance between the Gabor pairs (solid line, average across subjects, together with the SEM [gray area]) as a function of estimated receptive field size (Equation 4.2). The dashed lines represent model predictions of perceived distance increments, computed for each of the three attention field sizes. The predicted increases in perceived distance follow a similar increase with eccentricity as the measured increase in perceived distance.

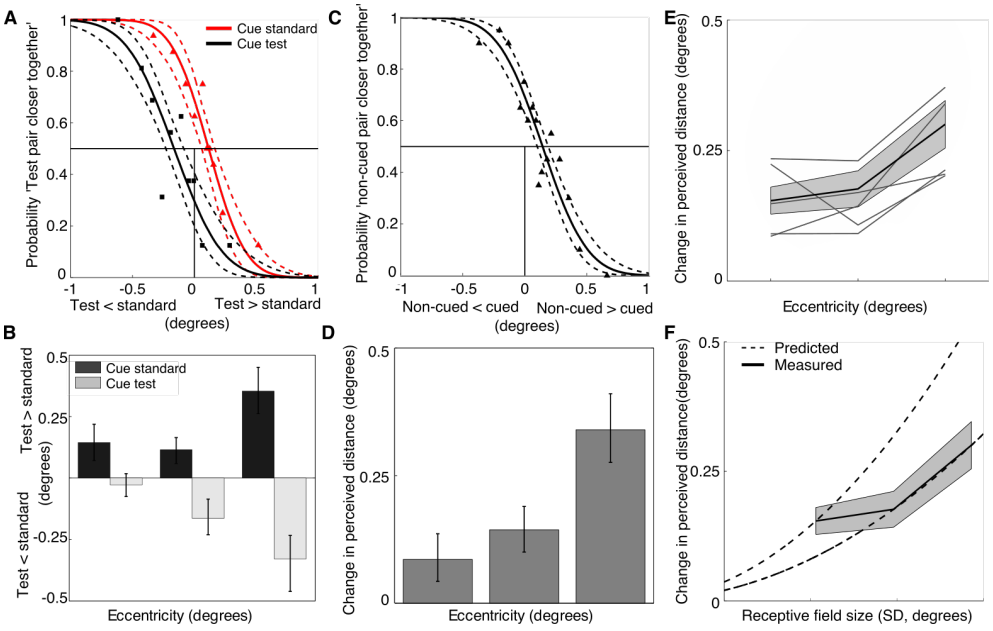

Figure 4.3 Results eccentricity experiment. **A** The probability that a participant reports seeing the test pair closer together than the standard pair, as a function of the difference in distance between the two pairs. Curves are the median cumulative Gaussian fits (solid lines) and the 95% confidence intervals (dashed lines) for the cue standard and cue test conditions separately. The data are from a single participant collected for the 6° condition. When the test pair is cued, it has to be closer together to appear as far apart as the standard pair (black curves). When the standard pair is cued, the test pair has to be further apart to appear as far apart as the standard pair (red curves). This indicates an increase in perceived distance when the Gabor pair appears at a previously cued location. The solid black, thin lines mark the p = 0.50 (PSE) and the point of equal presented distance between the test and standard pair. **B** The difference in distance between the test and standard pair at the PSE of every eccentricity condition separately. Data are from the same subject as in A. **C** Similar to A but now the cue standard and cue test conditions are collapsed. The plot shows the probability that the participant reports seeing the noncued pair closer together than the cued pair, as a function of the difference in distance between the cued and noncued pair. **D** Similar to B but now the cue standard and cue test conditions are collapsed. This shows that the magnitude of the effect of cueing on the perceived distance between the cued pair increases with increasing eccentricity. **E** The magnitude of the effect of cueing on the perceived distance between the cued pair for every participant separately (thin gray lines). The thick solid line marks the average across all participants; the gray area represents the standard error of the mean (SEM). **F** The average effect of cueing on the perceived distance between the Gabor stimuli across participants (solid line), together with the SEM (gray area) as function of the estimated receptive field size underlying the increases in perceived distance. Dashed lines represent increases in perceived distance predicted by the attention field model. All error bars in panels B and D are 95% confidence intervals obtained by bootstrapping. The triangles and squares in panel A and C represent binned averages of the raw data; cumulative Gaussians were fit on the raw, unbinned data. Cueing results in larger perceived distances between the Gabor pair, and this effect increases with eccentricity, as expected based on predictions from our attention field model

Experiment 2: Perceived distance does not scale with spatial frequency

The data from this experiment were analyzed in the same way as those for the eccentricity experiment. Figure 4.4A shows the probability that a participant reports seeing the test pair closer together than the standard pair, as a function of the difference in distance between the test and standard pair. Data are from the 4.75 cpd

condition collected for one participant. Figure 4.4B shows the difference in distance between the test and standard pair at the PSE and the 95% confidence intervals for both the cue test and cue standard conditions (dark and light gray bars, respectively) for all four spatial frequency conditions from the same participant. We collapsed the cue test and cue standard conditions in Figure 4.4C and D for the same participant. This reveals the magnitude of the effect of cueing on the perceived distance between the Gabor pairs. Figure 4.4E shows the magnitude of the effect of cueing on the perceived distance between the Gabor pair across spatial frequency for all participants separately (dashed lines) and the average across participants (solid line), together with the SEM (gray area). A one-way repeated-measures ANOVA revealed a significant effect of spatial frequency on the magnitude of the cueing effect, $F(3, 15) = 5.31$, $p = 0.011$, but post hoc testing revealed no significant differences between the spatial frequency conditions.

Again, we illustrate how model predictions for increases in perceived distance resemble those measured in this experiment. Using an estimate for the receptive field size tuned to the spatial frequency in every condition separately (Equation 4.3), we compute the attention field size for every condition (Equation 4.4). The computed attention field sizes are 1.54°, 0.93°, 0.28°, and 0.14° for the 1, 1.75, 4.75, and 10 cpd conditions, respectively. We used these attention field sizes to compute the predicted increase in perceived distance for every condition, across a large range of receptive field sizes. Figure 4.4F presents the measured effect of cueing on the perceived distance between the Gabor pairs (solid line, average across subjects, together with the SEM [gray area]). The dashed lines represent model predictions of perceived distance increments, computed for each of the three attention field sizes. An attention field model using these attention field sizes clearly predicts an increase of the increase in perceived distance between the Gabor stimuli that is not present in our data.

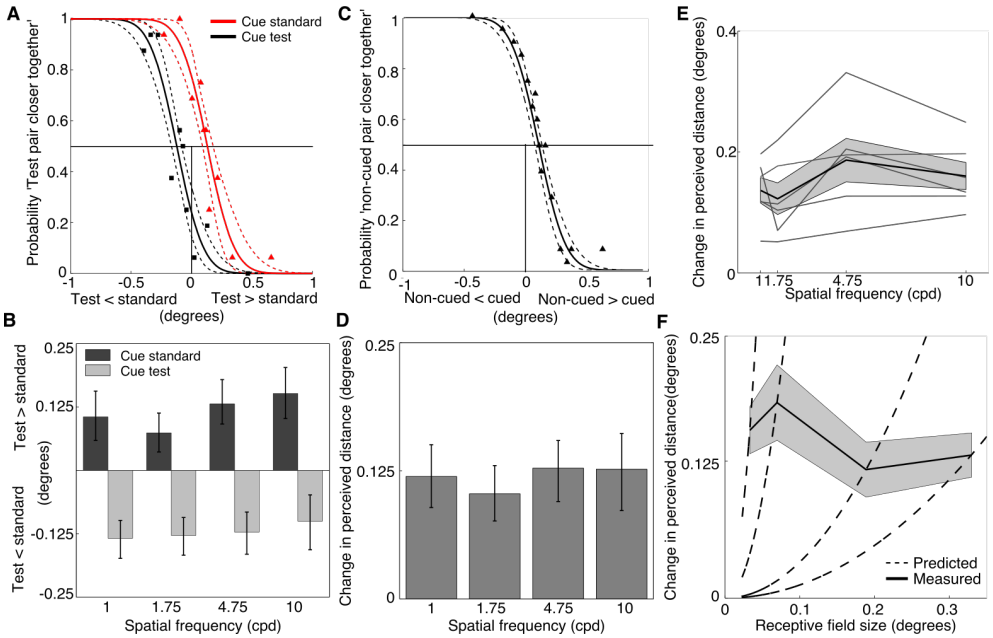
A
Probability 'Test pair closer together'
Cue standard
Cue test
Test < standard
Test > standard
(degrees)
C
Probability 'non-cued pair closer together'
Non-cued < cued
Non-cued > cued
(degrees)
E
Change in perceived distance (degrees)
Spatial frequency (cpd)
B
Test > standard
(degrees)
Test < standard
Cue standard
Cue test
Spatial frequency (cpd)
D
Change in perceived distance (degrees)
Spatial frequency (cpd)
F
Change in perceived distance(degrees)
Predicted
Measured
Receptive field size (degrees)

Figure 4.4 Results spatial frequency experiment. **A** The probability that a participant reports seeing the test pair closer together than the standard pair, as a function of the difference in distance between the two pairs. Curves are median cumulative Gaussian fits (solid lines) and the 95% confidence interval (dashed lines) for the cue standard and cue test conditions separately. The data are from a single participant collected for the 4.75 cpd condition. The solid black, thin lines mark the p = 0.50 (PSE) and the point of equal presented distance between the test and standard Gabor pairs. **B** The difference in distance between the test and standard pair at the PSE of every spatial frequency condition separately. Data are from the same subject as in A. **C** Similar to A but now the cue standard and cue test conditions are collapsed. The plot shows the probability that the participant reports seeing the noncued pair closer together than the cued pair, as a function of the difference in distance between the cued and noncued pair. **D** Similar to B but now the cue standard and cue test conditions are collapsed. This shows that the magnitude of the effect of cueing on the perceived distance between the cued pair is about constant across spatial frequencies. **E** The magnitude of the effect of cueing on the perceived distance between the cued pair for every participant separately (thin gray lines). The thick solid line marks the average across all participants; the gray area represents the standard error of the mean (SEM). **F** The average effect of cueing on the perceived distance between the Gabor stimuli across participants (solid line), together with the SEM (gray area) as function of the estimated receptive field size underlying the increases in perceived distance. Dashed lines represent increases in perceived distance predicted by the attention field model. All error bars in panels B and D are 95% confidence intervals obtained by bootstrapping. The triangles and squares in panel A and C represent binned averages of the raw data; cumulative Gaussians were fit on the raw, unbinned data. Based on our model, we expected that increasing the spatial frequency of the Gabors leads to a smaller increase in perceived distance. However, we find that the positional bias does not vary as a function of Gabor spatial frequency.

Control experiments: Perceived distance increments depend on stimulus timing but are independent from stimulus size

Data analyses for the control experiments were similar to that of the main experiments. Figure 4.5A shows the increase in perceived distance between the cued pair when controlling for the size of the Gabors in the frequency domain. Gray lines mark the results for each participant separately; the solid black line and the gray area mark the average across participants and the SEM, respectively. A one-way

repeated-measures ANOVA reveals no significant differences between the spatial frequency conditions, $F(3, 15) = 1.58$, $p = 0.235$.

Figure 4.5B shows the increase in perceived distance as a function of ISI for every participant separately (gray lines) and the average across all participants (black line), together with the SEM (gray area). A one-way repeated-measures ANOVA shows a significant effect of ISI on the perceived distance between the Gabors, $F(2, 10) = 18.849$, $p = ,10^{-6}$. Post hoc testing reveals a significant reduction in the effect of the cue for the ISI = 600 ms, compared with the ISI = 0 ms and ISI = 67 ms condition ($p = 0.024$ and 0.02, respectively). In addition, a marginally significant reduction of the increase in perceived distance was found for the ISI = 0 ms compared with the ISI = 67 ms condition ($p = 0.054$). All p values are corrected using false discovery rates (Benjamini and Hochberg, 1995). The increases in perceived distance do not depend on the size of the Gabors, but do depend on the interval between the cue and Gabors.

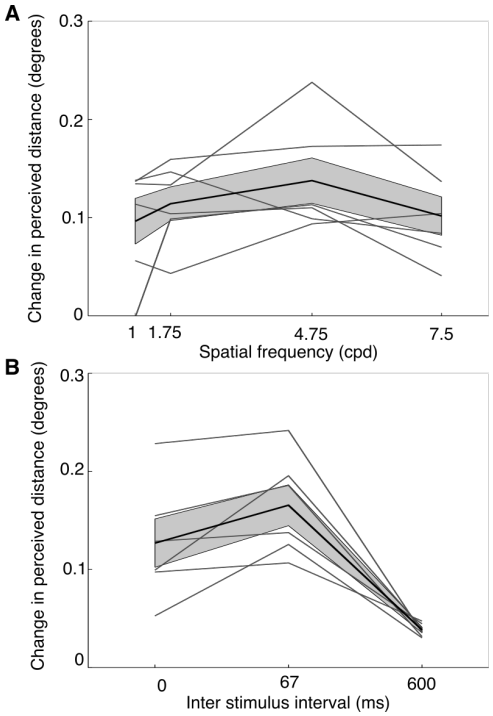

Figure 4.5 Results control experiments. **A** The magnitude of the effect of cueing on the perceived distance between the cued Gabor pair when controlling for the size of the Gabor stimuli in the frequency domain. Gray lines represent data from every participant separately. The black line represents the average across subjects, the gray area the SEM. **B** Similar to A but now when the timing between the cue and Gabors is varied. Data presented in this figure are analyzed in the same way as those presented in Figures 4.3 and 4.4. The increase in perceived distance does not depend on the size of the Gabors but does depend on the timing between the cue and Gabors.

Discussion

We have demonstrated that the positional bias of Gabors outside the focus of attention, measured by an increase in perceived distance between the two Gabors, increases with eccentricity. Perceived distance increments, however, do not vary with the spatial frequency content of the Gabors. Control experiments verify that attention underlies the observed positional bias in this experiment. In addition, the size of the Gabors in the frequency or the spatial domain did not affect the magnitude of the bias.

Several models have been proposed that can account for the positional bias examined here. Suzuki and Cavanagh (1997) proposed three mechanisms through which attention can lead to a positional bias, namely, response suppression, receptive field shrinkage, and receptive field attraction. More recently, Baruch and Yeshurun (2014) proposed the attentional attraction field model. This model explains the positional bias as a consequence of receptive field attraction. Our model as presented here explains the positional bias as a consequence of receptive field attraction as well. However, our model includes receptive field size as an important parameter, which is not the case for the models proposed by Suzuki and Cavanagh (1997) and Baruch and Yeshurun (2014). Here, we attempted to examine the dependence of the positional bias on receptive field size, a key prediction from our model. Our experiments were not designed to distinguish between different models, and as such, our results do not permit differentiation between models.

Within our model, the positional bias is the consequence of an interaction between the attention field, representing the strength of attentional modulation and the receptive field. We find that the positional bias increases with eccentricity, which we interpret as a consequence of larger receptive fields at higher eccentricities. Within this interpretation, larger receptive fields are more susceptible to attentional modulation and are thus attracted more. This assumes that variation in attentional modulation, and thus attention field size, between the conditions in the eccentricity experiment cannot explain the increase in positional bias. We have several reasons to assume that this is indeed the case.

First, the size of the attention field typically varies with experimental parameters, such as spatial uncertainty (Belopolsky et al., 2007; Hernández et al., 2010; Herrmann et al., 2010) and may depend on the specific task and stimulus used in an experiment (Intriligator and Cavanagh, 2001). Each condition in our eccentricity experiment used the same stimulus and participants performed the same task, only the eccentricity at which the stimulus was presented varied. Moreover, stimuli were equally likely to appear at each of the three eccentricities, thus minimizing variation in spatial uncertainty about the stimulus position between

conditions. Therefore, we expect little variation in attention field size between our conditions because of variation in spatial uncertainty and stimulus or task characteristics. Second, in our model, the attention field represents a Gaussian-shaped multiplicative amplification of neural responses (Womelsdorf et al., 2008; Reynolds and Heeger, 2009; Klein et al., 2014). Within this framework, stronger attentional modulation is predicted when response amplification increases and when this amplification is spread over a narrower range of visual space. Thus, if attention elicits stronger response amplification at higher eccentricities, or if this amplification is spread over a smaller extent of visual space, the attentional modulation may be stronger at higher eccentricities. This could then explain the larger positional bias we measure at higher eccentricities. However, fMRI studies suggest that higher eccentricities are related to less response amplification in most, but not all, visual areas (Bressler et al., 2013) and possibly to an increased spread of this amplification (Puckett and DeYoe, 2015). Although far from conclusive, within the framework of our model, this suggests a weaker attentional modulation of receptive fields at higher eccentricities, which would lead to a reduced attraction of receptive fields and thus predict a smaller positional bias. In sum, we expect little variation in attention field size between the conditions of our eccentricity experiment due to task and stimulus characteristics. In addition, higher eccentricities may be related to reduced attentional modulation of receptive field position. Therefore, we assume that variation in attention field size cannot account for the increase in positional bias with eccentricity but that increases in receptive field size underlie this result.

The increase of receptive field size with eccentricity accounts for other eccentricity-dependent effects of attention as well. For example, whether attention improves or impairs performance on a texture segmentation task depends on the match between the spatial scale of the segmentation task and the receptive field size at the eccentricity at which the task is presented (Yeshurun and Carrasco, 1998). We varied the eccentricity of both the cue and target to examine the effects of eccentricity on the magnitude of the positional bias. Other studies manipulated only the eccentricity of the cue to examine how the positional bias varies with cue–target distance. They found that the positional bias peaks for cue–target distances of approximately 2°–4° and decreases for shorter and longer distances (Suzuki and Cavanagh, 1997; Kosovicheva et al., 2010). As presented here, our model predicts an increasing positional bias with increasing distance but not a decrease. To be able to account for a decreasing effect with distance, our model can be extended to include a parameter that attenuates receptive field attraction as a function of distance to the attended location.

Although an attention field model predicts variation of the positional bias with spatial frequency, we did not find such an effect. There can be several explanations for the lack of variation in perceived position across spatial frequency. First, it should be noted that the model predictions for the spatial frequency experiment are derived using another model describing the relation between spatial frequency and receptive field size in V1 (Jones and Palmer, 1987). This assumes that changes in receptive field position at the level of V1 affect the positional bias examined here. However, some psychophysical evidence suggests that the positional bias may be induced at levels prior to binocular integration in V1 (DiGiacomo and Pratt, 2012). Thus, varying the spatial frequency of our stimulus may lead to changes in underlying receptive field size and consequently to variation in receptive field attraction at the level of V1. However, it is possible that these changes at the level of V1 are not translated to changes in the positional bias. In contrast, every visual field map across the visual processing hierarchy displays an eccentricity by receptive field size relationship. Consequently, increasing the eccentricity at which a stimulus is presented would vary the underlying receptive field size at every level of the visual hierarchy and thus induce variation in receptive field attraction throughout the visual hierarchy. This makes it more likely that the level of visual processing at which the positional bias is induced shows variation in receptive field attraction as well. Thus, the different ways by which eccentricity and spatial frequency are encoded across the visual hierarchy may have led to the difference between the results of the eccentricity and spatial frequency experiment.

Second, the attention field model described in this study is based on fMRI results obtained using endogenous attention. In this study, however, we test its perceptual consequences with exogenous attention. Exogenous and endogenous attention might be controlled differently in the primate brain (Corbetta and Shulman, 2002; Buschman and Miller, 2007; Busse et al., 2008) and may have some different effects on perception (Briand and Klein, 1987; Briand, 1998; Lu and Dosher, 2000; Yeshurun et al., 2008). However, we do not believe that differences between exogenous and endogenous attention are responsible for the discrepancy between the predicted and measured positional biases in Experiment 2. Both exogenous and endogenous attention affect visual processing as early as the primary visual cortex (Motter, 1993; Gandhi et al., 1999; Somers et al., 1999; Liu et al., 2005) where spatial frequency–selective cells have been well described (Jones and Palmer, 1987). In addition, an attention field model captures the effects of both exogenous and endogenous attention on contrast sensitivity in human perception (Herrmann et al., 2010) suggesting that they affect visual processing in very similar ways. Moreover, positional biases have been found using both exogenous and endogenous attention (Suzuki and Cavanagh, 1997). Finally, using exogenous attention, we did

find an increase in positional biases with eccentricity in the current study, as predicted by the model.

Third, if different spatial frequencies are affected by a positional bias to differing extents, this may brake phase alignment across different spatial frequencies and disrupt accurate edge detection. Therefore, the visual system could correct for spatial misalignment across spatial frequencies. This correction may be conceptually similar to the correction imposed on retinal motion due to eye movements (Sommer and Wurtz, 2008; Wurtz, 2008). Therefore, local mechanisms may correct for the spatial misalignment of different spatial frequencies.

Fourth, it is possible that the magnitude of the positional bias we measure here is not related to variation in receptive field shifts but to uncertainty about the distance difference between the Gabor stimuli. Within this framework, participants are more susceptible to a positional bias when they are more uncertain about the distance difference between the two Gabor stimuli. This task uncertainty is represented by the slope (sigma) of the cumulative Gaussians we fit to our data. Indeed, we find that the slopes of our functions are significantly related to the measured bias in both the eccentricity and spatial frequency experiment but not in the control experiment, where we varied the spatial extent of the Gabor stimuli. Therefore, position uncertainty can explain some, but not all, of our results.

Finally, it is important to consider that our task was spatial in nature. Our cue was spatially localized, and we examined its effect on the spatial position of the Gabor stimuli. In the spatial frequency experiment, we manipulated the spatial frequency of the Gabor stimuli, a stimulus dimension that is orthogonal to and independent of the stimulus position. This discrepancy between the nature of the task and the nature of the varied stimulus property could be responsible for the lack of an effect of spatial frequency on perceived position. Thus, the attention field we employed may operate only in the spatial domain and not in the spatial frequency domain. Similarly, attention may operate independently from the spatial domain in feature space (David et al., 2008; Hayden and Gallant, 2009; Çukur et al., 2013)

Conclusions

We examined biases in the perceived position of objects surrounding the focus of spatial attention. An attention field model predicts a larger positional bias when underlying receptive fields are larger. We found that positional biases increase with eccentricity but are relatively constant across a range of spatial frequencies. We have discussed several reasons that might explain why positional biases do not vary with spatial frequency. The discussed mechanisms can potentially be incorporated into an attention field model to increase its explanatory power in the perceptual

domain. The attention field model is therefore a useful way to bridge neural mechanisms of attention with perception.

Chapter 5

Numerosity tuning in human visual cortex

Published as:
Harvey, B.M., Klein, B.P., Petridou, N., & Dumoulin, S.O. (2013).
Topographic Representation of Numerosity in the Human Parietal Cortex. *Science*, 341, 1123. doi:10.1126/science.1239052

BMH, BPK, and SOD designed research. BMH, BPK, and NP collected the data. BMH and BPK analyzed the data. BMH and SOD wrote the paper.

Abstract

Numerosity, the set size of a group of items, is processed by the association cortex, but certain aspects mirror the properties of primary senses. Sensory cortices contain topographic maps reflecting the structure of sensory organs. Are the cortical representation and processing of numerosity organized topographically, even though no sensory organ has a numerical structure? Using high-field functional magnetic resonance imaging (at a field strength of 7 teslas), we described neural populations tuned to small numerosities in the human parietal cortex. They are organized topographically, forming a numerosity map that is robust to changes in low-level stimulus features. The cortical surface area devoted to specific numerosities decreases with increasing numerosity, and the tuning width increases with preferred numerosity. These organizational properties extend topographic principles to the representation of higher-order abstract features in the association cortex.

Humans and many other animals use numerosity to guide behavior and decisions (Dehaene, 1997; Nieder and Miller, 2003; Burr and Ross, 2008; Brannon and Terrace, 2016). Numerosity perception becomes less precise as the size of numbers increases (Whalen et al., 1999; Nieder and Miller, 2003, 2004; Piazza et al., 2004; Tokita and Ishiguchi, 2010) and is particularly effective for small numbers (Saltzman and Garner, 1948). Animals, infants, and tribes with no numerical language perceive numerosity (Dehaene et al., 1999, 2008; Cantlon et al., 2009; Brannon and Terrace, 2016) although they cannot count or use symbolic representations of number. Thus, numerosity processing is an evolutionarily preserved cognitive function, distinct from counting and humans' unique symbolic and mathematical abilities (Dehaene et al., 1999). Because aspects of numerosity processing mirror primary sensory perception, it has been referred to as a "number sense" (Dehaene, 1997; Burr and Ross, 2008)

The primary sensory and motor cortices in the brain are organized topographically. Is the neural organization for numerosity similarly topographic? The neural representation of numerosity resides in higher-order association cortices, including the posterior parietal cortex. Human functional magnetic resonance imaging (fMRI) consistently identifies this region as particularly responsive to numerosity manipulations (Dehaene et al., 1999; Pinel et al., 1999; Piazza et al., 2004; Eger et al., 2009) and in similar regions, macaque neurophysiology describes neurons tuned to visual numerosity (Nieder et al., 2002; Nieder and Miller, 2003, 2004; Tokita and Ishiguchi, 2010). Both human fMRI and macaque neurophysiological response properties are closely linked to behavioral numerosity performance (Nieder and Miller, 2003; Piazza et al., 2004)

We elicited responses to visual patterns with varying numerosity in study participants, while acquiring high-field (7 teslas) fMRI data. Changing numerosity in a visual display affects visual features such as luminance, contrast, density, and total edge length. Therefore, establishing numerosity selectivity requires several control conditions (Fig. 5.1A and figs. S5.1 and S5.2) (Nieder et al., 2002). Consequently, we included conditions in which total dot area ("constant area" condition), individual dot size ("constant dot size"), or total dot circumference ("constant circumference") were constant. A further condition contained much higher dot pattern density ("high density"). Finally, to check generalization to other objects, we replaced dots with different shapes ("variable features"). During stimulus presentation, participants reported when dots were shown in white rather than black (10% of presentations). No numerosity judgments were required. Participants performed above 90% correct.

The displayed numerosity varied systematically within an fMRI scan (Fig. 5.1B, top inset). This stimulus elicited remarkably different response profiles at

different recording sites (Fig. 5.1B), despite having similar hemodynamic response functions (fig. S5.6). We summarized these fMRI signals using numerosity-tuned neural models (Fig. 5.1C and fig. S5.4). These describe Gaussian functions in logarithmic numerosity space, following behavioral (Nieder and Miller, 2003), computational (Dehaene and Changeux, 1993), neuroimaging (Piazza et al., 2004), and neurophysiological (Nieder and Miller, 2003, 2004) results (fig. S5.5). The models have two parameters: preferred numerosity and tuning width (the numerosity range to which the population responds). This analysis is analogous to conventional population receptive field analysis in the visual cortex (Dumoulin and Wandell, 2008). These models explain much of the signal variance (R^2), summarizing fMRI responses with two parameters. They capture similar amounts of variance for both example response profiles in Fig. 5.1B, explaining time course differences by different numerosity tunings (Fig. 5.1C).

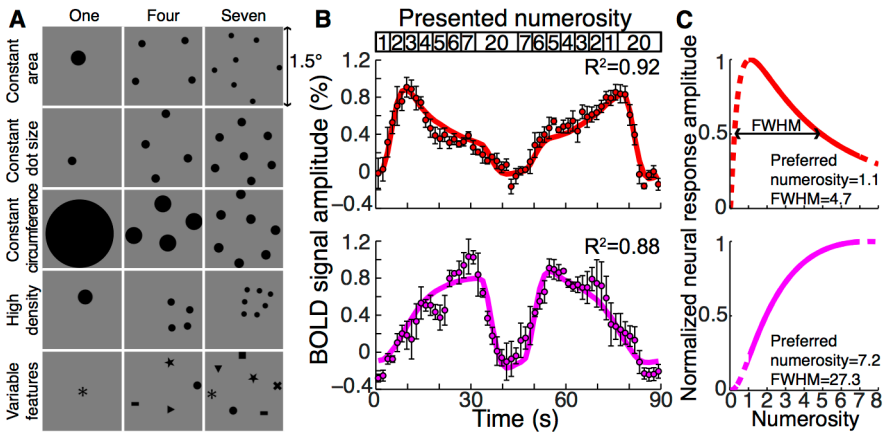

Fig. 5.1 Stimuli, responses, and neural population tuning. **A** Illustration of stimulus conditions, with examples representing different numerosities. **B** Two example fMRI time courses from sites in the posterior parietal cortex, separated by about 2 cm, elicited by the numerosity stimulus sequence (top inset). BOLD, blood oxygen level–dependent. Points represent mean response amplitudes; error bars represent the standard error over repeated runs. In the upper panel, the largest response amplitude occurs after the presentation of low numerosities, whereas in the lower panel the largest response occurs with higher numerosities, considering the hemodynamic response delay. To quantify these differences, we developed a novel data-analysis method that extracts numerosity tuning from the time courses, following methods we developed in the visual cortex (Dumoulin and Wandell, 2008)(Fig. S5.4). The numerosity model captures about 90% of the variance (R^2) in the time courses, as indicated by the colored lines. **C** Representation of the neural model that best fits each time course. The model describes a Gaussian tuning function in logarithmic numerosity space with two parameters: preferred numerosity and tuning width defined by the full width at half maximum (FWHM). Different model parameters explain the differences seen in B, capturing a similar amount of the variance.

A specific region in the posterior parietal cortex was highlighted, where the models captured much response variance in all stimulus conditions (Fig. 5.2A and fig. S5.7). This region's position was consistent between the eight participants, in the posterior superior parietal lobule, centered at mean (SD) Montreal Neurological Institute x,y,z coordinates of 23 (4), –60 (7), 60 (7) (Collins et al., 1994) and closely matches previous reports of a region responding strongly to numerosity manipulations (Dehaene et al., 1999; Pinel et al., 1999; Piazza et al., 2004; Eger et al., 2009)

Projecting each recording site's preferred numerosity onto the unfolded cortical surface revealed an orderly topographic map (Fig. 5.2B). Medial and lateral regions preferred low and high numerosities, respectively. The topographic progression and its direction were consistent between participants and stimulus conditions (fig. S5.8). Numerosity selectivity was also present in the left hemisphere (fig. S5.9) and in neighboring regions of the right hemisphere, but with lower variance explained and less clear, less consistent topographic structure. To quantify the numerosity organization, we sorted recording sites within this region by their distance from the borders representing lowest and highest numerosities (white lines in Fig. 5.2B). We then plotted preferred numerosity against cortical distance for individual stimulus conditions and their average (Fig. 5.2C and fig. S5.10).

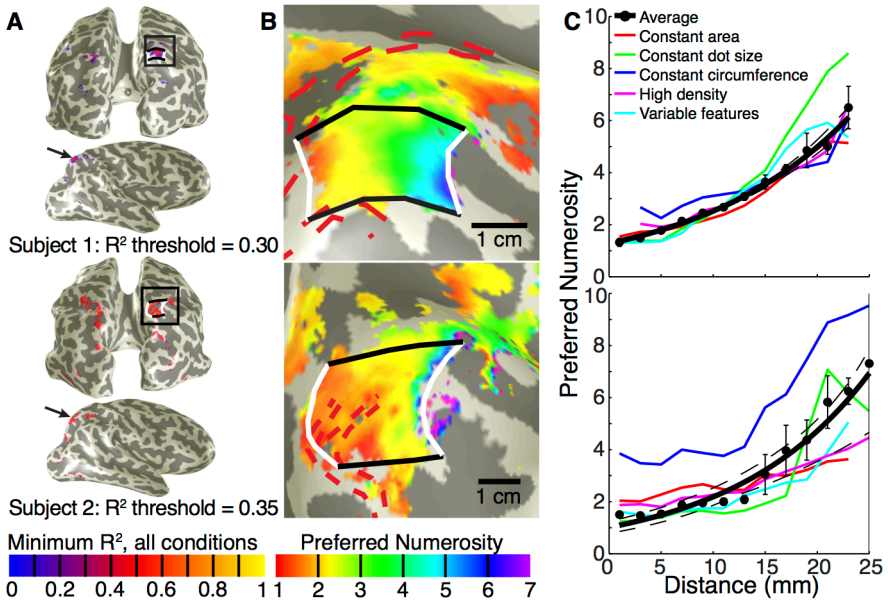

Fig. 5.2 Topographic representation of numerosity **A** The variance explained by the model (R^2) highlighted a region in the right parietal cortex where neural populations demonstrated numerosity tuning in all stimulus conditions (Fig. 5.1A). The black square is enlarged in B. **B** Numerosity preferences for data averaged from all stimulus conditions, showing preferred numerosity increasing from the medial to lateral ends (white lines) of the region of interest (ROI) (black and white lines). Areas of low signal intensity, corresponding to pial surface veins (red dashed lines, fig. S5.3), were excluded from further analysis (Winawer et al., 2010). **C** Numerosity preference progression from medial to lateral along the ROI for all conditions. All recording sites were organized by their distances from the two white lines. Dots represent the mean preferred numerosity in each distance bin, with error bars showing the standard error. We fitted the binned points with a logarithmic function (solid black line), with 95% confidence intervals to the fit (dashed black lines) determined by bootstrapping. More cortical area is devoted to lower number; i.e., cortical magnification decreases at higher numerosity. Different stimulus conditions are represented as colored lines joining the condition-specific bin means.

Numerosity preference was organized topographically in all stimulus conditions, so numerosity preference is significantly correlated between conditions (Fig. 5.3A and fig. S5.11A). However, absolute numerosity preference varied with stimulus condition (Fig. 5.3B and fig. S5.11B), which is consistent with results from single

macaque neurons (Nieder and Miller, 2004). In particular, the constant circumference condition differs from other conditions. It has very different dot sizes from other conditions (Fig. 5.1A and fig. S5.2), and we propose that dot size interacts with numerosity preference, because line length–selective neurons are found with numerosity-selective neurons in the macaque posterior parietal lobe (Tudusciuc and Nieder, 2007).

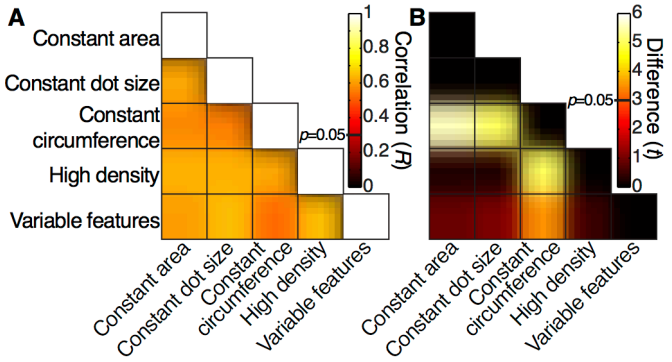

Fig. 5.3 Comparison of numerosity preferences across recording points in different stimulus conditions, averaged across participants. **A** Because numerosity preferences are topographically organized in all stimulus conditions, they are always correlated. **B** However, preferences change with stimulus conditions, with preference increasing particularly in the constant circumference condition. *t*, *t*-statistic.

The rate of numerosity preference change with distance increased with numerosity; i.e., the cortical magnification factor decreases (Fig. 5.2C). Thus, more cortical surface area represents lower than higher numerosities (fig. S5.12). Similarly, more macaque posterior parietal neurons prefer low than high numerosities (Nieder and Miller, 2004). Such overrepresented parts of other topographic maps also show more precise response selectivity than elsewhere in the map; i.e., tuning width is smaller. Tuning widths changed across the topographic map along with preferred numerosity (Fig. 5.4A and fig. S5.13A). Population tuning width increased with preferred numerosity (Fig. 5.4B and fig. S5.13B), in line with single macaque neuron (Nieder and Miller, 2003, 2004), neuroimaging (Piazza et al., 2004) and behavioral results (Whalen et al., 1999; Piazza et al., 2004; Tokita and Ishiguchi, 2010). However, population tuning widths are larger than macaque single-neuron tuning widths (fig. S5.14). This may arise from differences in neural

population size (~400,000 neurons in our recording points), the scatter of response preferences at a single cortical location, hemodynamic properties, interpolation steps in fMRI analyses, task differences, and/or species differences (Dumoulin and Wandell, 2008).

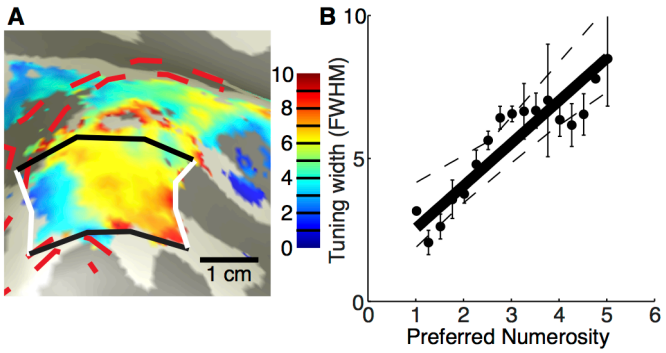

Fig. 5.4 The progression of population tuning width (see Fig. 5.1C) across the cortical surface A and with preferred numerosity B for one representative participant. Dots represent mean tuning widths in each preferred numerosity bin, and error bars represent standard errors. Dashed lines represent 95% confidence intervals of the fit (solid line) to the bin means. Tuning width increases with preferred numerosity for all participants (fig. S5.13B).

Neuroimaging studies consistently show that this part of the parietal cortex responds to numerosity manipulations (Dehaene et al., 1999; Pinel et al., 1999; Piazza et al., 2004; Eger et al., 2009), and parietal lesions can cause number-processing deficits (Dehaene and Cohen, 1991). Macaque neurophysiology demonstrates numerosity tuning in single neurons in a similar parietal region (Nieder and Miller, 2004), and human neural adaptation properties suggest that numerosity-tuned populations exist in the same area, with tuning widths increasing with preferred numerosity (Piazza et al., 2004). We extended these observations by directly measuring numerosity tuning in the human cortex and describing a topographic organization of numerosity, a numerosity cortical magnification factor, and a relationship between preferred numerosity and numerosity tuning width. Based on similar behavioral performance and cortical location of numerosity-selective populations in humans and macaques, we expect similar topographic organization in macaques. The spatial scale of the topographic organization is several centimeters. Consequently, methodological limitations of single-neuron recordings may have prevented its identification; i.e., at single-neuron resolution,

topography may be obscured by the scatter of response properties, broad single-neuron tuning, neurons with other response properties, and an unknown direction of topography change. However, both methodologies are complementary, and our measurements are consistent with neurophysiology. Both support numerosity tuning, albeit at different scales, in similar parts of the brain, with more neurons tuned to smaller numerosities and increases in tuning width with preferred numerosity. These properties are analogous to organization properties of the sensory and motor cortices and may underlie the decreased precision at higher numerosities that is commonly seen in human and animal behavior (Dehaene et al., 1999; Pinel et al., 1999; Whalen et al., 1999; Nieder and Miller, 2003; Piazza et al., 2004; Eger et al., 2009)

Our numerosity-selective responses cannot be explained by other visual attributes of the stimulus. First, tuning and topographic structure were found using stimuli controlled for low-level features. Second, responses in visual field maps such as V1 cannot be captured by the numerosity model but follow stimulus contrast energy (fig. S5.15). Third, parietal visual field map borders (Swisher et al., 2007) did not correspond to numerosity map borders and their relative positions varied considerably between participants (fig. S5.16). In macaques, over 80% of single neurons here show no numerosity selectivity (Nieder and Miller, 2004; Tudusciuc and Nieder, 2007) so independent representations of numerosity and visual space may exist in one cortical region, represented by different neurons. Alternatively, these populations have large visual receptive fields and may be tuned to numerosity presented anywhere within this receptive field. Interactions between overlapping numerosity and visuospatial representations may underlie the cognitive spatial "number line" (Walsh, 2003; Dehaene et al., 2008). However, we find no consistent relationship between numerosity and visuospatial responses.

What is the nature of the numerosity representation? We found no number-tuned responses for Arabic numerals (fig. S5.17), suggesting that neurons here do not respond to symbolic number representations. We propose that current biologically plausible computational models of numerosity processing, driven by visual features, can produce the numerosity selectivity we see (Dehaene and Changeux, 1993; Dakin et al., 2011). Some models suggest that (as we find) numerosity selectivity depends on stimulus features, such as dot size (Dakin et al., 2011). Computational models of numerosity extraction may thus explain these differences in numerosity tuning, consistent with behavioral results (Tokita and Ishiguchi, 2010; Dakin et al., 2011)

Numerosity processing and its cortical organization may be fundamental to human abilities in mathematics and economics. Although numerosity judgments and complex mathematical abilities rely on different processes, individual differences in

these abilities are correlated (Halberda et al., 2008). Macaques and young children can perform simple, approximate addition and subtraction (Starkey, 1992; Cantlon and Brannon, 2007). In macaques, the parietal and prefrontal cortices contain neurons responding specifically during simple mathematical tasks, together with numerosity-selective neurons (Bongard and Nieder, 2010). Associations between visual numerosity and symbolic number representations develop early in life (Cantlon et al., 2009). Numerosity, number, and size are fundamental to our understanding of magnitude and quantity and underlie higher-level concepts of value (Walsh, 2003).

Our results demonstrate that topographic representations, common in the sensory and motor cortices, can emerge within the brain to represent abstract features such as numerosity. Similarities in cortical organization suggest that the computational benefits of topographic representations, for example efficiency in wiring (Ramon y Cajal et al., 1999; Chen et al., 2006) apply to higher-order cognitive functions and sensory-motor functions alike. As such, topographic organization may be common in higher cognitive functions. On the other hand, topographic organization supports the view that numerosity perception resembles a primary sense (Dehaene, 1997; Burr and Ross, 2008). These views are not mutually exclusive, but both challenge the established distinction between primary topographic representations and abstracted representations of higher cognitive functions.

Supplementary material

Subjects

We present data from eight subjects (two female, age range 19-38 years). Two were left-handed. All were well educated, with good mathematical abilities. All had normal or corrected to normal visual acuity. All were trained with tasks requiring numerosity judgments before scanning. Three subjects were naïve to the aims of the experiment. All experimental procedures were cleared by the ethics committee of University Medical Center Utrecht. A total of thirteen subjects were recruited. However, five subjects were excluded from analysis and no complete data-set was recorded because either (a) after the first scanning session it was clear that the position of the superior sagittal sinus prevented imaging of the numerosity map (3 subjects), or due to excessive movement in early recordings (1 subject), or the subject chose not to complete all scanning sessions (1 subject).

Numerosity stimuli

Visual stimuli were presented by back-projection onto a 15.0x7.9 cm screen inside the MRI bore. The subject viewed the display through prisms and mirrors, and the total distance from the subject's eyes (in the scanner) to the display screen was 41 cm. Visible display resolution was 1024x538 pixels.

The stimuli were generated in Matlab using the PsychToolbox (Brainard, 1997; Pelli, 1997). A large diagonal cross, composed of thin red lines, crossed the entire display, a design that allows very accurate fixation (Schira et al., 2009). Subjects were asked to fixate the intersection of the cross. Stimuli consisted of groups of dots presented in the central 0.75° (radius) of the visual field (Fig S5.1). This small stimulus area was chosen to decrease the need to make eye movements towards the individual stimulus items and to minimize the cortical surface extent of the visually-responsive part of the brain activated by presentation of the stimulus, avoiding confusion between spatially-specific responses and numerosity-specific responses.

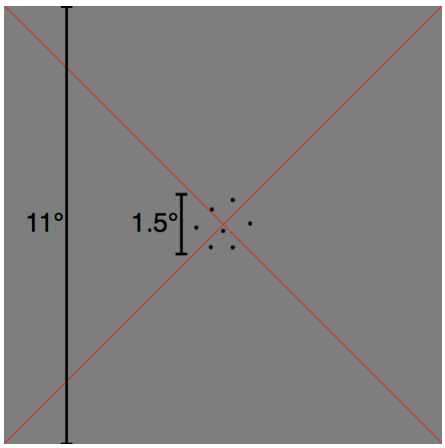

Fig S5.1. A full example stimulus as seen by the subject in the scanner. The dot pattern covered the central 1.5° (visual angle) diameter within an 11° diameter mean-luminance (gray) field. A large, thin, red fixation cross passes diagonally through the center of the display, and through the center of the dot pattern. Subjects fixated the intersection of the cross.

Dots were randomly positioned at each presentation so that each dot fell entirely within the central 0.75°. As such, for all numerosities contrast energy was distributed equally across the stimulus area, avoiding any links between numerosity and the visual field position of the contrast energy. Each of the many numerosity presentations whose responses were averaged together (192 presentations for each numerosity in each condition) contained dots placed in a new, random position, so no specific visual position was associated with any numerosity. To prevent perceptual grouping, individual items were distributed roughly homogenously

across the stimulus area (except for the high density condition described below) (see Fig S5.2). We used various stimulus configurations (Fig 5.1A, S5.2) (Nieder et al., 2002) to ensure that responses to low-level visual features of the stimulus did not follow the same time course in different conditions. The first stimulus configuration ('constant area' condition) kept the total surface area of all of the dots combined constant across numerosities, ensuring equal luminance across conditions. The second ('constant dot size') kept individual dot size constant. The third ('constant circumference') kept the total circumference constant, ensuring equal edge density. The fourth condition ('high density') contained the same dots as the first, but at higher density, with all dots falling within a 0.375° radius circle that was randomly placed inside the stimulus area. The final condition ('variable features') used different shapes as stimulus elements, ensuring that results generalize to objects other than dots.

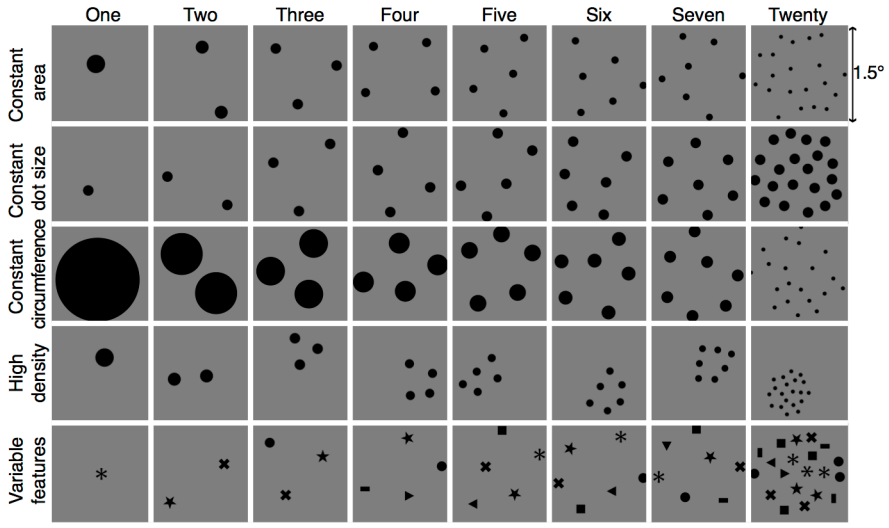

Fig S5.2. Stimulus conditions used, with examples of stimuli representing every numerosity used (comparable to Fig 5.1A). Item position was randomly assigned every stimulus presentation.

All patterns were presented as black dots (or shapes) on a gray background. Dot patterns were presented briefly (300 ms) to ensure subjects did not have time to count (Saltzman and Garner, 1948). This was repeated every 750 ms, each time with a new random pattern presented, with 450 ms presentation of a uniform gray background between dot pattern presentations. This was repeated six times, over

4500 ms, corresponding to three fMRI volume acquisitions (TRs), before the numerosity changed. On 10% of pattern presentations, the dots were shown in white instead of black. Subjects were instructed to press a button when this happened to ensure they were paying attention to the patterns during fMRI acquisition, and responded on 90-100% of white dot (or shape) presentations within each scanning run. No numerosity judgments were required. A recent macaque result also demonstrates numerosity-selective responses in untrained macaques that were not performing a task (Viswanathan and Nieder, 2013), suggesting that these responses do not require attention to numerosity.

The numerosities one through seven were shown as the main stimulus (Fig 5.1B, top inset), first presented in ascending order, followed by a longer period (27 seconds) where the stimulus contained twenty dots, followed by the numerosities in descending order, followed by another long period of twenty dots. This sequence was repeated four times in each scanning run. The long period of twenty dots had a similar function to the blank periods used in visual field mapping stimuli in population receptive field experiments, allowing us to distinguish between very small and very large tuning widths (Dumoulin and Wandell, 2008), i.e. between populations which responded at all times equally to numerosities 1-7 and those that never responded. During this period, little neural response was expected from numerosity-selective neurons with small preferred numerosities, as such a large numerosity should be well outside of the numerosity range that elicits strong responses. This allows hemodynamic responses to return back to baseline between blocks of changing numerosity. However, using twenty dots (rather than zeros dots) provides a stronger visual stimulus than the other numbers. As such, neural populations responding to the contrast energy of the stimulus should respond most strongly during the twenty dot presentation, avoiding confusion with populations preferring a specific large numerosity.

Numerosity stimuli were presented many times in quick succession between each numerosity change, which is required to ensure strong fMRI responses and accurate measurements of response properties. Unfortunately, as in many fMRI experiments, this stimulus sequence is likely to cause some adaptation to the presented numerosity (Piazza et al., 2004; Burr and Ross, 2008). By using stimuli with ascending and descending numerosity changes mixed together, adaptation effects on preferred numerosity estimates are counterbalanced by previously presenting numerosities that give higher and lower responses than the currently presented numerosity. As the model must fit both of these response sequences with one set of tuning parameters, the resulting tuning estimates reflect responses preceded by both high and low responses, as such the effect of adaptation on tuning estimates is likely to be small.

Eye tracking

Differences in the extent of eye movements during presentation of different numerosities could confuse numerosity-selective responses with eye movement selective responses. This is a particular concern as the posterior parietal cortex is involved in generating eye movements (Merriam and Colby, 2005; Williams and Smith, 2010; Harvey and Dumoulin, 2011; Wardak et al., 2011). Our subjects were experienced in visual fMRI experiments requiring accurate fixation (Harvey and Dumoulin, 2011). Subjects were also instructed on the importance of maintaining fixation. Finally, we presented a large fixation cross on the stimulus (Fig S5.1) to minimize eye movements (Schira et al., 2009).

Nevertheless eye movements for all subjects were measured outside of the scanner using a highly accurate Eyelink II system. This allowed the distribution of fixation positions to be measured with the same stimulus and task as was shown in the scanner. The distribution of eye positions did not vary with numerosity, and was in all cases very small (less than 0.27° standard deviation in the distribution of eye positions). This result is similar to eye movements recorded in our previous studies (Harvey and Dumoulin, 2011).

MRI acquisition

Anatomical MRI data for Subjects 1, 2, 6 and 7 were acquired on a 3T scanner with a Quasar Dual gradient set. T1-weighted anatomical MRI data were acquired at a resolution of 0.75x0.75x0.8 mm. Repetition time (TR) was 10.029 ms, echo time (TE) was 4.6 ms, and flip angle was 8 degrees.

Anatomical MRI data for Subjects 3, 4, 5 and 8 were acquired on a 7T scanner using a 32 channel head coil. T1-weighted anatomical MRI data were acquired at a resolution of 0.5x0.5x0.8 mm. Repetition time (TR) was 7 ms, echo time (TE) was 2.84 ms, and flip angle was 8 degrees.

Functional T2*-weighted 2D echo planar images were acquired on a 7T scanner using a 32 channel head coil at a resolution of 1.98x1.98x2.00 mm, with a field of view of 190x190x50 mm. TR was 1500 ms, TE was 25 ms, and flip angle was 80 degrees. The acquired volume covered the entire occipital lobe, most of the parietal lobe, and the posterior temporal lobe. Functional runs were each 248 time frames (372 seconds) in duration, of which the first eight time frames (12 seconds) were discarded to ensure the signal was at steady state. Four repeated runs were acquired within the same session for each stimulus condition. Within each session, data for two conditions were acquired, in interleaved order. Other stimulus conditions were acquired on different days. Visual field mapping used the same

scanning protocol. Data were acquired in a separate session, with eight repeated runs.

Preprocessing of anatomical and functional images

Functional MRI analysis was performed in the mrVista software package, which is freely available at (http://white.stanford.edu/software/). T1-weighted anatomical scans were automatically segmented using Freesurfer (Dale et al., 1999) and then hand-edited to minimize segmentation errors (Teo et al., 1997) using ITK-SNAP (Yushkevich et al., 2006). This provided a highly accurate description of the cortical surface, an anatomical segmentation space used for analysis of cortical organization. The cortical surface was reconstructed at the gray-white matter border and rendered as a smoothed 3D surface (Wandell et al., 2000). Head movement and motion artifacts between and within functional scans were measured and corrected for (Nestares and Heeger, 2000). Functional data were aligned to the anatomical scans (Nestares and Heeger, 2000) and interpolated to the anatomical segmentation space. Data from several sessions, resulting from all stimulus conditions, was imported into the same anatomical segmentation space. FMRI time series data from all stimulus conditions were averaged together to produce a data set with a very strong signal. Across cortical thickness, data from all recording points (voxels) were collapsed and averaged onto the nearest point on the cortical surface. This increased signal strength and formed a (folded) two-dimensional representation of the gray matter. All results were also analyzed using the same methods without collapsing onto a two-dimensional surface, which produced very similar results. Note that the acquired data were averaged, rather than the model parameters. Data from each individual condition was analyzed separately.

Exclusion of vein artifacts

A large draining venous system, consisting of the superior sagittal sinus and its branches, runs over the medial posterior parietal lobe. Draining veins spatially and temporally distort the fMRI signal around this vein (Winawer et al., 2010; Olman et al., 2012). Theoretically, measurements from high field MRI scanners (we used 7T here) are less susceptible to signals originating in large veins (Yacoub and Hu, 2001). Nevertheless, we identified these locations by the mean signal intensity of the BOLD signal (Fig S5.3) (Yacoub and Hu, 2001; Harvey and Dumoulin, 2011) and excluded these from further analysis. Similar results are found if these vein regions are included in the analysis. However, this procedure excludes regions where the

topographic representation and tuning width measurements are distorted (Fig 5.2B, 5.4A, S5.8, S5.9 and S5.13A).

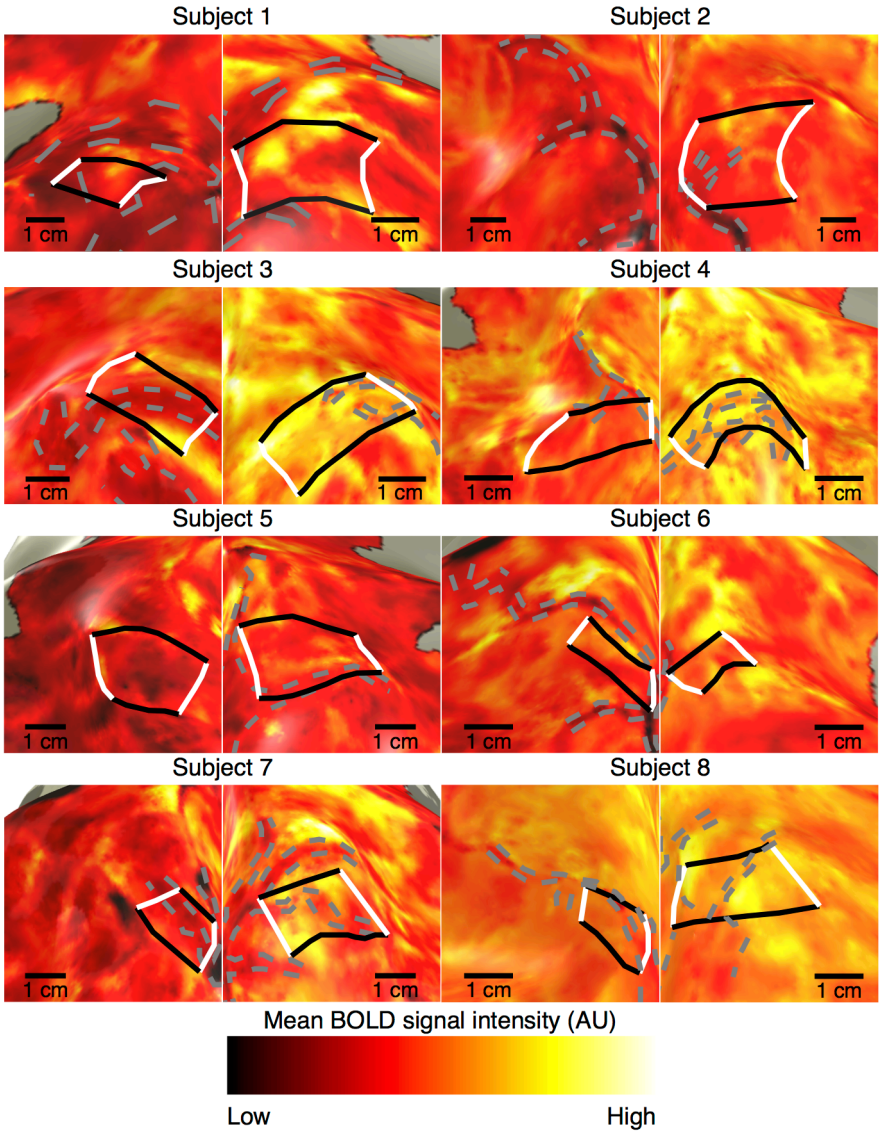

Fig S5.3. BOLD signal strength at each recording point in the average data, rendered on inflated cortical surfaces showing the same views as Fig S5.8 (right) and Fig S5.9 (left). Large draining veins on the pial surface, the superior sagittal sinus and its branches, can be seen as areas of low signal strength. These are outlined with gray dashed lines, which correspond to the red dashed lines seen in other renderings of the cortical surface. Data from these areas are distorted and the blood flow and oxygenation here results from neural activity elsewhere. Based on a subject-specific threshold of minimum signal strength in the average data, such recording points are excluded from analysis of preferred number and tuning width. The lines highlighting these areas are for illustration only.

fMRI data-analysis

Population numerosity tuning was estimated from the fMRI data and the time course of presented numerosity in a similar manner to methods we developed to estimate population receptive field properties in human visual cortex (Dumoulin and Wandell, 2008). Population receptive field models describe the aggregate visual receptive field of the neural population within each fMRI recording site. Here we measured population tuning to a one-dimensional stimulus property (numerosity) rather than the two-dimensional stimulus property of the population receptive field approach (visual field position). However, the principle of the approach is the same as in population receptive field modeling (summarized in Fig S5.4). A forward model predicts neuronal responses at each stimulus time point depending on which numerosity was shown. The model describes tuning in logarithmic numerosity space using a Gaussian characterized by a preferred numerosity (mean of the Gaussian distribution) and tuning width (standard deviation of the Gaussian). By examining the overlap of the stimulus at each time point with this tuning model, a prediction of the neuronal response time course is generated. By convolving this with a hemodynamic response function (HRF), a predicted fMRI time course is generated. The predicted fMRI time courses were generated for all combinations of a large range of preferred numerosity and tuning width parameters. For each recording point, the parameters were chosen from the prediction that fit the data most closely by minimizing the sum of squared errors (R^2) between the predicted and observed fMRI time series.

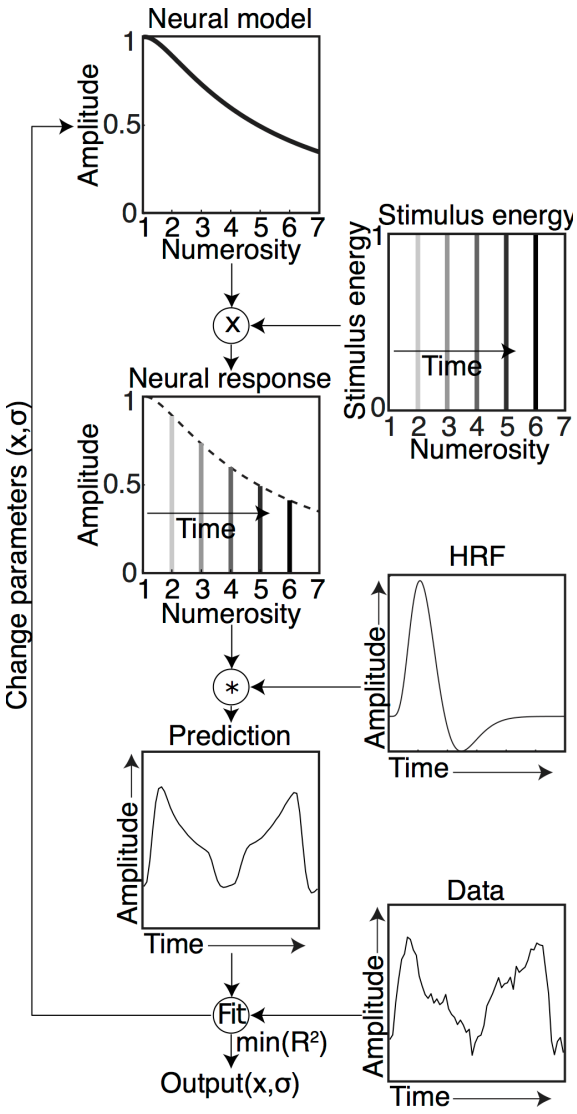

Fig S5.4. A flow chart describing the model fitting principle for a single recording point (fMRI voxel). A neural model describing one possible neural response configuration, with a particular set of parameters describing the preferred numerosity and tuning width of the neural population within a voxel. The predicted neural response over time is calculated by multiplying the presented numerosity at each time point by the normalized response at that numerosity in the neural model (X). This predicted neural response time course is convolved with a hemodynamic response function (∗) to produce a predicted fMRI time course, given this set of neural model parameters. Best fitting model parameters are found by minimizing the difference between the predicted and the recorded data.

The numerosity tuning model parameters extended beyond the stimulus range, allowing model fit parameters beyond this range. This allows us to be confident that returned parameters within the stimulus range are reported accurately, rather than the best fit of a limited set. However, recording points with preferred numerosities modeled outside the stimulus range must be treated with caution. Such recording points prefer a numerosity that we did not display, so responses monotonically increase or decrease across the stimulus range. As such we have little confidence

that the preferred numerosity estimate is correct. Recording points whose preferred numerosity was outside the stimulus range were not labeled on cortical surface renderings and were excluded from further analysis.

The Gaussian tuning function in logarithmic numerosity space was motivated by behavioral (Nieder and Miller, 2003), computational (Dehaene and Changeux, 1993), neuroimaging (Piazza et al., 2004) and neurophysiological (Nieder and Miller, 2003, 2004) results. Fitting a tuning function using a Gaussian in linear numerosity space gave similar estimates of preferred numerosity but with decreased variance explained (Fig S5.5). A Gaussian in logarithmic space, when converted into linear space, is asymmetrical, skewed towards higher numerosities (see Fig 5.1C). Consequently, a Gaussian fit in linear space tends to bias preferred numerosity estimates slightly towards higher numbers. Despite this, topographic representations are well maintained. However, the tuning width estimate is strongly and systematically affected by this asymmetry. For recording points preferring numerosities low in the stimulus range, the steeper low end of the response is truncated, so that the fit relies largely in the shallower side above the preferred numerosity, overestimating the tuning width. For recording points preferring numerosities high in the stimulus range, the fit is determined from the steeper change seen on the low side, underestimating tuning width.

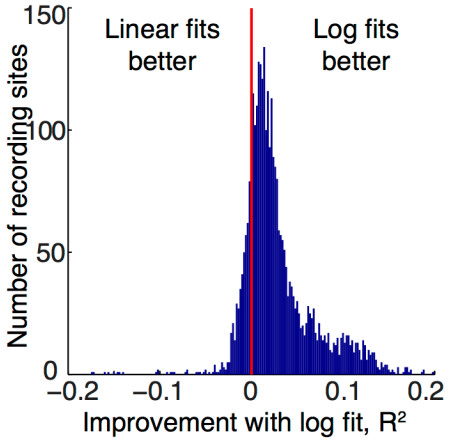

Fig S5.5. Neural models describing a Gaussian in logarithmic space fit the data better than those describing a Gaussian in linear space in 83% of voxels from the grouped data from all subjects (paired *t*-test, $p<0.00001$). These comparisons are made for models using the average data over all conditions, for all recording points within the ROIs shown in Fig S5.8.

We estimate the HRF parameters across the whole acquired fMRI volume from the data, using a near-identical procedure we employ in visual cortex (Harvey and Dumoulin, 2011; Hoffmann et al., 2012; Takemura et al., 2012). Briefly, by having the stimulus pass through the numerosity range in both ascending and descending directions, we can derive the HRF properties. We estimated the HRF parameters by

comparing predicted time-series and actual time-series and chose the HRF parameters that minimized the difference between predicted and actual time series over the entire volume for each condition. Next, we averaged the HRF parameters determined from each condition's data and used those HRF parameters to re-estimate the numerosity tuning models. This procedure improved the goodness of fit and ensured the same HRF is used in modeling responses to all conditions and at every point in the brain. Very similar results, though with less good model fits, were obtained by fitting the data using a canonical HRF (Engel et al., 1997; Friston et al., 1998; Worsley et al., 2002) again analogous to results we have obtained in visual cortex (Harvey and Dumoulin, 2011). In order to exclude the possibility that the differences seen in fMRI time series (Fig 5.1B) result from differences in hemodynamics, we measured hemodynamic response functions from the same example sites (Fig S5.6). To elicit these responses, we showed patterns with three dots, to which both sites respond strongly (about 70% of their maximum neural response), for 4.5 seconds with the stimulus properties as used in the main experiment. We then showed patterns with twenty dots, to which neither site responds strongly, for 25.5 seconds, allowing the response to return to the baseline that we used in the main experiment. Both sites yielded very similar hemodynamic response functions.

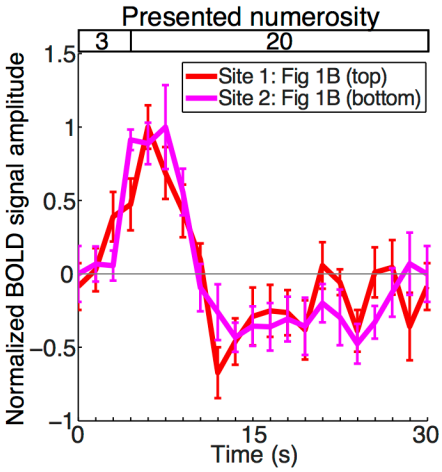

Fig S5.6: Hemodynamic response functions from the two sites shown in Fig 5.1B. These sites respond very differently to the stimulus used in the main experiment. Yet both sites have very similar hemodynamic response functions, demonstrating that hemodynamic differences cannot explain the different responses to our numerosity stimulus.

ROI definitions

We determined which parts of the cortex responded in a numerosity-selective way in all conditions using a conjunction analysis by taking the minimum goodness of fit (R^2) of the numerosity tuning model in any condition. The resulting minimum R^2

values were rendered onto an inflated cortical surface (Fig 5.2A, S5.7). This analysis highlighted a consistent region in the posterior parietal lobe in all subjects. This region formed the basis our region of interest (ROI). We then rendered the preferred numerosities of each recording site on the cortical surface (Fig S5.8, 5.2B). In every subject and stimulus condition, we found a topographic representation of preferred numerosity in the same medial-to-lateral direction within the region of consistently good model fits seen in Fig S5.7. The model from the average data consistently gave the clearest and largest region of topographic representation, covering the topographically organized regions in all other conditions, so we used this data to define our ROI. Medial and lateral borders of the ROI each followed lines of equal preferred numerosity at the low and high ends of the preferred numerosity range seen in each subject. Anterior and posterior borders describe the edges of the topographic organization, which typically also coincided with decreases in the goodness of model fits.

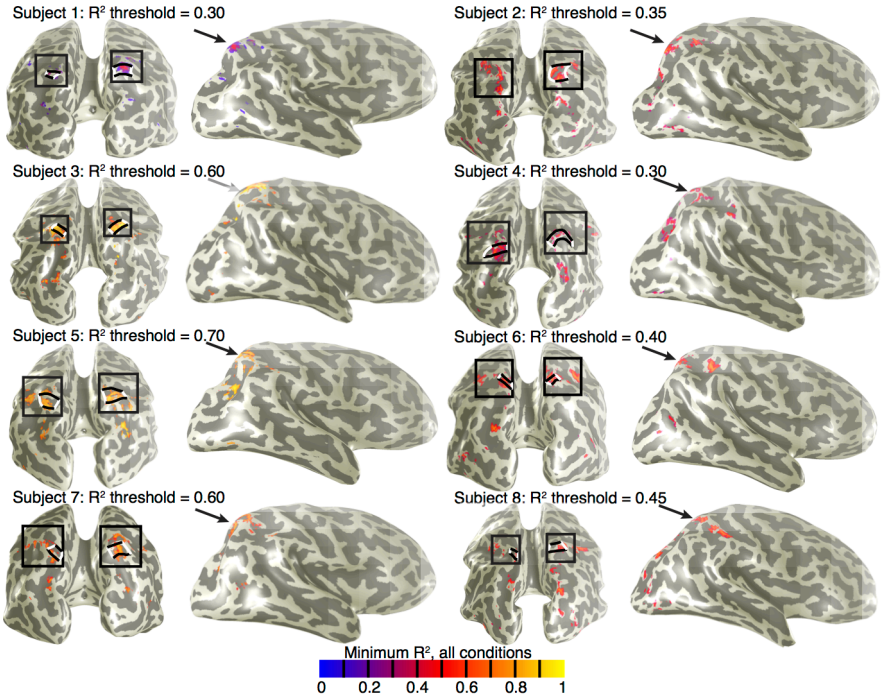

Fig S5.7. Conjunction analysis showing regions where models fit the data well in all stimulus conditions: regions with clear numerosity tuning despite changes in low-level features of the stimulus between conditions. Colors indicate the minimum goodness of fit (R^2) in any stimulus condition, rendered on an inflated cortical surface viewed from the back (left) and right side (right). A clear region of good fits is visible in the posterior parietal lobe in all subjects, and its position relative to the surrounding sulci is consistent between subjects. The areas shown in Fig S5.8 and S5.9 are shown for each subject as black boxes on the right and left hemispheres respectively. In each of these, the regions of interest that have clear topographic organization (Fig S5.8 and S5.9) are shown as black and white outlines, which correspond well to the areas of good model fits shown here. Note that data whose R^2 is below a certain threshold are not shown. This threshold varies between subjects due to variations in data quality. Arrows represent the point of view from which the left panel (back view) is taken.

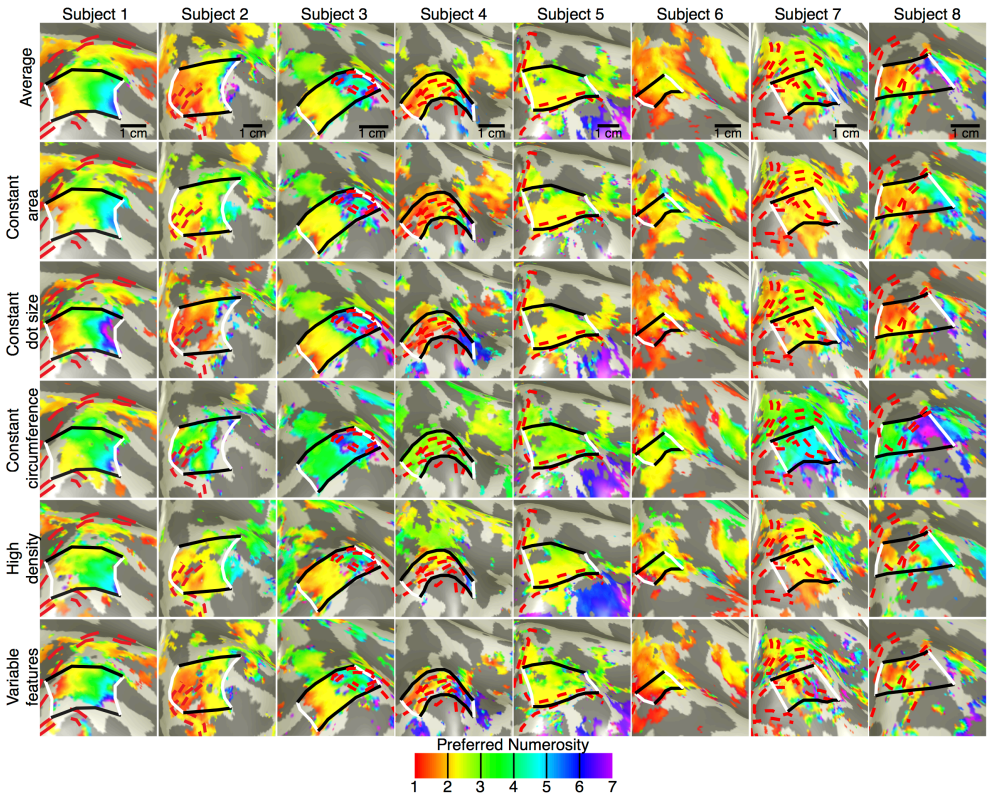

Fig S8. Preferred numerosity varies across the cortical surface of the right hemisphere in all stimulus conditions. Details of the area with the best minimum R2 values across all stimulus conditions, shown in Fig S7. Colors represent different preferred numerosities rendered on an inflated back view of the cortical surface. An area of clear topographic representation in all stimulus conditions is defined as an ROI in black and white. The borders of this area representing low and high preferred numerosities are shown as white lines at the medial and lateral ends of the ROI, on the left and right respectively. The posterior and anterior borders of this topographic representation are shown as black lines. Data is thresholded based on goodness of fit: In the average data, recording points where R2 is below 0.5 are not shown; for all individual conditions the R2 threshold is 0.3. Recording points where model fits a preferred numerosity outside of the stimulus range are not shown. Dashed red lines outline distortions in the data caused by the presence of large veins on the pial surface (Fig S3).

Beyond this ROI, there is also evidence of numerosity-selective responses anterior and lateral to the ROI, i.e. in the anterior intra-parietal sulcus (IPS). Organization here was less consistent than within the ROI (Fig S5.8) and model fits were less consistently good between stimulus conditions (Fig S5.7). The extent of our recording volume did not allow a complete view of organization in the anterior IPS, though many studies find that this area also responds strongly to numerosity manipulations (Dehaene et al., 1999; Pinel et al., 1999; Piazza et al., 2004; Eger et al., 2009). We also rendered preferred numerosity at each cortical surface location in the left posterior parietal lobe on the cortical surface (Fig S5.9) and, where a topographic progression was visible, examined how preferred numerosity changed with distance across this region. Model fits were less good here (Fig S5.7), but draining vein artifacts obscure much of the cortical organization. One subject showed no evidence of topographic structure in the left parietal lobe at all (Subject 2, Fig S5.9). Organization here was also less consistent between conditions, and confidence intervals on the fit slope parameter included negative slopes for at least one condition in every subject. While there is at least some evidence of topographic organization here, it is possible that this results from low-level stimulus features that co-varies with numerosity in some conditions.

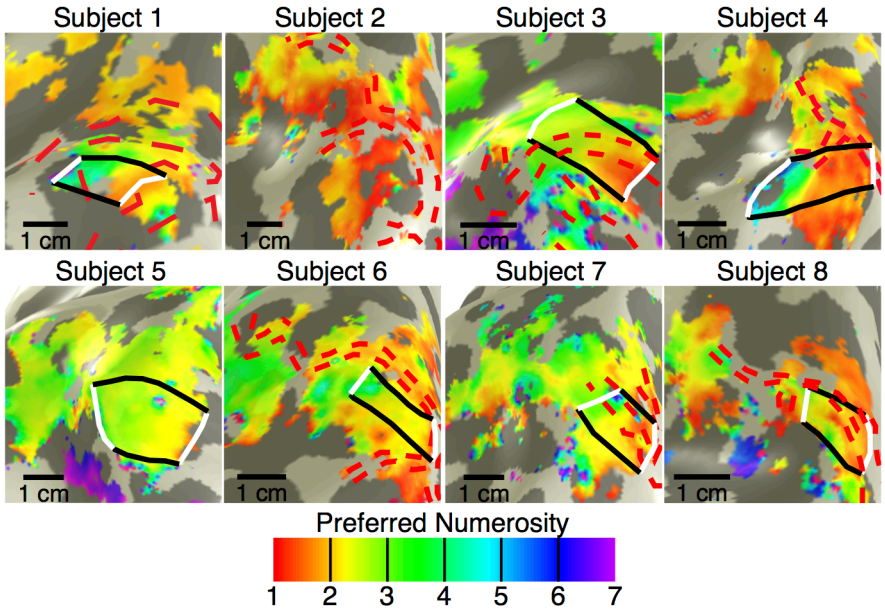

Fig S9. Preferred numerosity determined from averaged data varies across the cortical surface of the left hemisphere, as described for Fig S8. The organization is less clear than seen in the right hemisphere, and areas of topographic organization are less consistent between subjects. In Subject 2's left hemisphere, no area of topographic organization is seen, so no ROI is defined.

Analysis of changes across the ROI

We defined the medial and lateral edges of each numerosity map ROI as lines and calculated the distance along the cortical surface from each point in each ROI to the nearest point on each of these lines (white lines in Fig 5.2B, S5.8). The ratio between the distances to each end gives a normalized distance along the ROI in the primary direction of numerosity change. We multiplied this normalized distance by the mean length of the ROI in this direction. The result gives a measure of the distance along the ROI for each recording point. This distance measure will deviate slightly from the actual distance along the ROI, particularly if the ROI is not rectangular. However, it has several advantages: it includes the relative distance to both edges (rather than the absolute distance to one end only); it forces the distance between the edges to be constant for all edge positions (as the numerosity difference is constant); it increases statistical power by including measures from every point in the ROI (unlike

measuring the progression along a single line ROI); and it provides a distance measure in millimeter units (rather than a normalized measure only).

For every 2 mm increase in distance along the ROI, we formed a bin of recording points, then calculated the mean and standard error of the preferred numerosity of the points within the bin, excluding points where model fits had R^2 values below 0.3. We then fit lines to bootstrapped samples of the binned points that were the best fitting straight lines in logarithmic space. The fit variables were a slope and an intercept term. Logarithmic fits gave significantly higher R^2 values than straight line fits for all subjects, determined by bootstrapping, and visual inspection of the fits here demonstrates that these fits describe the binned data well (Fig 5.2C, S5.10). We weighted the residual of each bin by the inverse standard error of the mean when minimizing the R2 of each fit. As such, bins with a high standard error of the mean preferred numerosity influenced fits less than bins with a small standard error. From this set of fits, we took the median of each fit parameter as the best fit. We determined 95% confidence intervals by plotting all lines generated during bootstrapping iterations and finding the 2.5% and 97.5% percentiles on values for these fits.

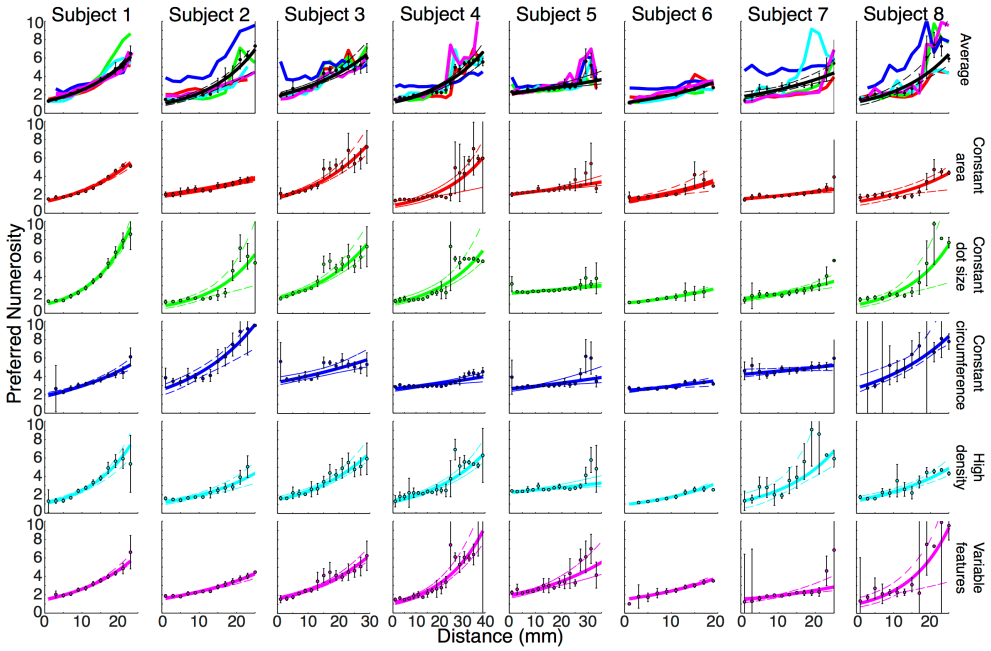

Fig S5.10. Progression of preferred numerosity with distance along the right hemisphere ROI in all stimulus conditions and the average data for each subject. Points represent the mean preferred numerosity in each distance bin, with error bars representing the standard error. Solid lines are the fit to the bin means, weighted by the inverse of the standard error of each bin. These fit lines are straight in logarithmic space. Dashed lines represent 95% confidence intervals determined by bootstrapping fits to the binned points.

The 95% confidence intervals for the slope of these fits (Fig S5.10, 5.2C) were above zero for all subjects and conditions, indicating a significant progression of preferred numerosity across each ROI. As numerosity preference changes across the ROI in the same direction in all stimulus conditions, any pair of stimulus conditions should have correlated changes within the ROI. Indeed, we find that all pairs of conditions are well correlated, with very few exceptions that were not consistent between subjects (Fig S5.11A, 5.3A). However, numerosity preference is not invariant to stimulus condition (Fig S5.11B, 5.3A), as shown by a paired *t*-test of numerosity preferences at all recording points within the ROI. This is consistent with results from single macaque neurons (Nieder and Miller, 2004). The constant circumference condition in particular differs from other conditions, often giving higher numerosity preferences. To average correlation strength across subjects (Fig 5.3A), we took the average regression coefficient (R). To average the t statistic (Fig 5.3B), we calculated t using the mean difference between conditions across subjects, the Pythagorean mean of the standard error, and the mean number of measurements (which differed between subjects as area size differed). To determine the R and t values corresponding to significant correlations and differences, we used Bonferroni correction for ten comparisons (i.e. the number of comparisons made between different conditions).

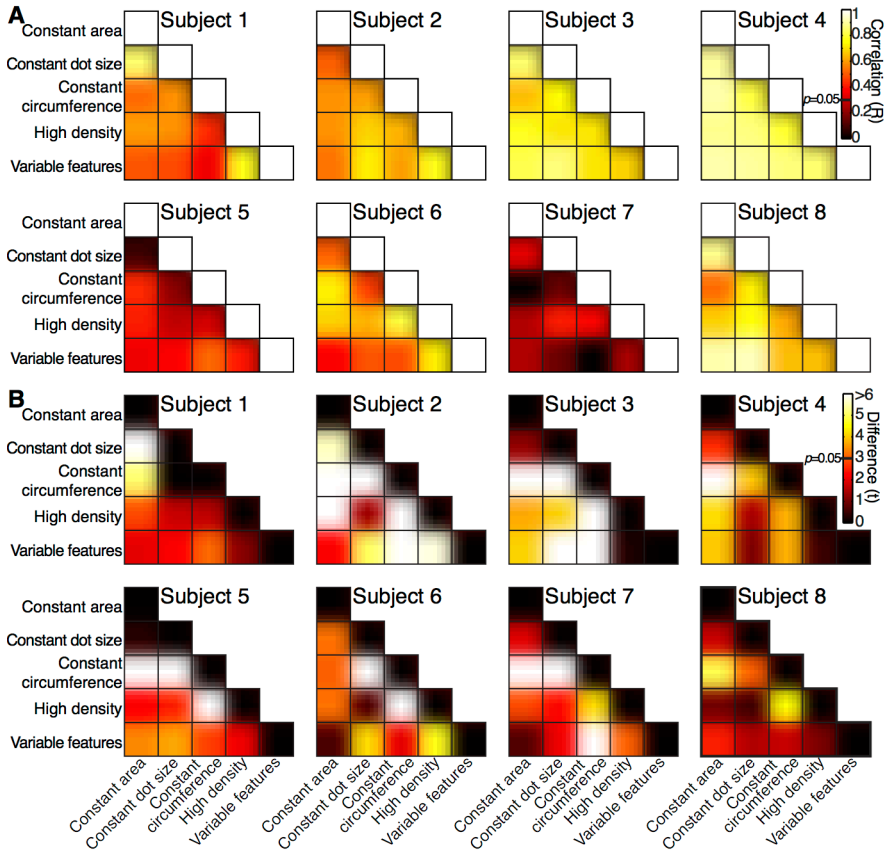

Fig. S5.11. Comparison of numerosity preferences across recording points in different stimulus conditions for every subject. **A** As numerosity preferences are topographically organized in all stimulus conditions, they are correlated between all stimulus conditions. **B** However, preferences change with stimulus conditions, increasing particularly in the constant circumference condition. Significance levels are corrected for multiple comparisons.

We then determined the rate of change of numerosity with distance across the ROI for each bootstrap fitting iteration used in Fig S5.10 on the averaged data. This gave a cortical magnification factor for numerosity in the primary direction of numerosity preference change across the ROI, along with confidence intervals from the group of bootstrap iterations (Fig S5.12). Cortical magnification factor is commonly determined in visual cortex by similar methods (Engel et al., 1994; Sereno et al.,

1995; Smith et al., 2001; Duncan and Boynton, 2003; Larsson and Heeger, 2006; Harvey and Dumoulin, 2011)

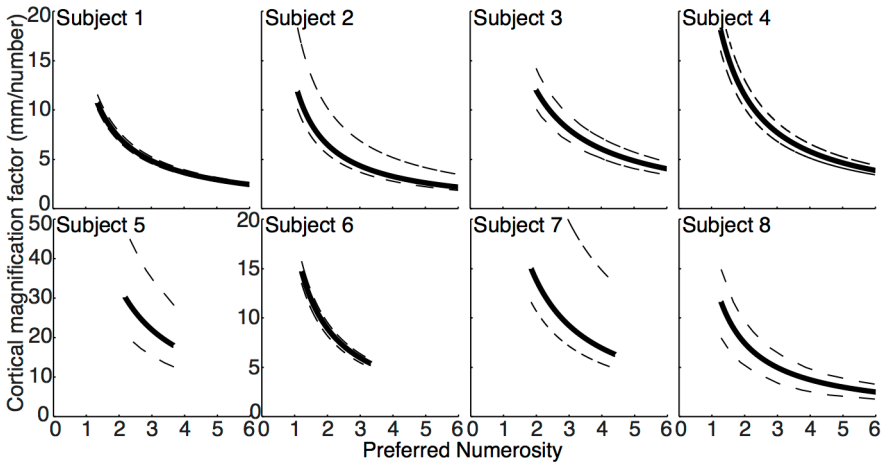

Fig S5.12. Cortical magnification factor of each numerosity in the right hemisphere ROI, determined from the averaged data. As numerosity increases, the distance along the ROI corresponding to a single number increase in preferred numerosity decreases. Fits are derived from the fits shown in the top panels of Fig S5.9. Dashed lines represent 95% confidence intervals determined by bootstrapping and refitting the bins in Fig S5.9.

Tuning width also changed across the cortical surface (Fig 5.4A, S5.13A), increasing with increased preferred numerosity. Data here are again shown for the average of all stimulus conditions. To determine how tuning width changed with preferred numerosity, we binned data based on preferred numerosities between one and five, with numerosity increments of 0.25 between bins. Few recording sites with preferred numerosities above five were seen, and near the edge of the stimulus range unstable fits are common as there is little information in the data set to distinguish between different tuning widths (Dumoulin and Wandell, 2008; Harvey and Dumoulin, 2011). Where the ROI did not contain preferred numerosities over this whole range, fits were limited to the range of preferred numerosities seen. Lines and confidence intervals were fit as described above, though here using a linear fit (Fig 5.4B, S5.13B).

As with changes in preferred numerosity with distance across the ROI, for all subjects the confidence intervals of tuning width at the high end of the preferred

numerosity range are above those at the low end, demonstrating a significant increase in tuning width with preferred numerosity.

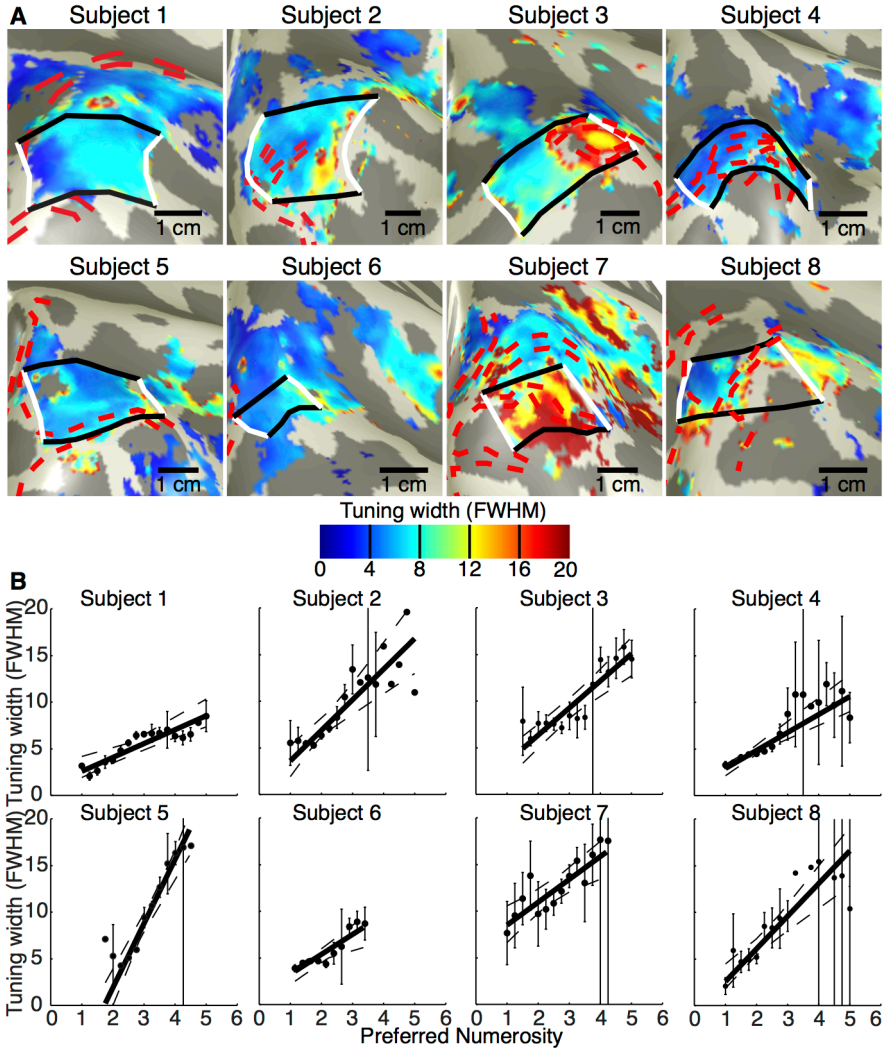

Fig S5.13. Tuning width changes across the cortical surface with preferred numerosity. **A** Change in tuning width across the cortical surface for averaged data, showing the same views seen in Fig S5.8. Tuning width increases from the medial to the lateral ends of the ROI. Data is thresholded based on goodness of fit: recording points where R^2 is below 0.5 are not shown. **B** Tuning width increases shown as a function of preferred numerosity. Recording points are binned based on preferred numerosity. Points represent the mean tuning width in each bin, error bars represent the standard error. Solid lines are the linear fit to the bins, weighted by the inverse of the standard error of each bin. Dashed lines represent 95% confidence intervals determined by bootstrapping fits to the binned points. Note that the scales used in both panels differ from those used in Fig 5.4.

Comparisons between population tuning widths and macaque single neuron tuning

Population tuning width is larger than the tuning width of a single neuron (Fig S5.14). As such, the tuning widths we report should not be taken as measures of single neuron tuning. Many factors contribute to the difference between population and single neuron measures (Smith et al., 2001; Dumoulin and Wandell, 2008). Yet population tuning width measures may be proportional to single neuron tuning widths as most of the factors contributing to population tuning width are proportional to single neuron tuning width (Harvey and Dumoulin, 2011). First, the neural population in an fMRI recording point is far larger than a single neuron. At the acquired functional resolution, one recording point contains a population of about 400,000 neurons, which increases the aggregate tuning width of the population (Rockel et al., 1980; Leuba and Garey, 1989). Second, at any single cortical location, a range of response preferences (scatter) is found, again increasing the aggregate tuning width. Third, the hemodynamic response also has a spatial component that reflects activity outside the recording site because part of the blood flow at any site comes from outside that site. Fourth, subject motion and various spatial interpolation steps will also increase the recorded population. Fifth, there are differences in the task, specifically our subjects are not making any numerosity judgments. Task and related attentional demands have known effects on tuning width, at least in visual cortex. Last, there may be differences between humans and macaques.

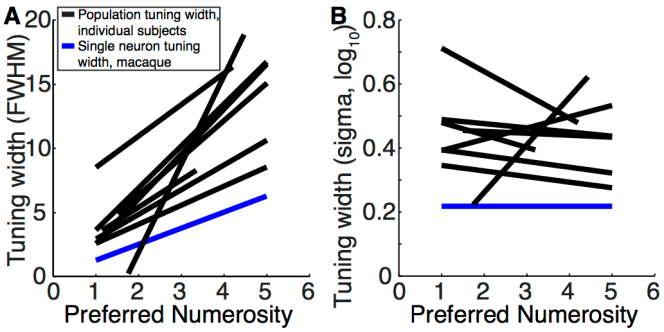

Fig S5.14. Comparison of population tuning width from our subjects (black lines) to reports from macaque single neurons (blue lines, (Nieder and Miller, 2003)), presented as FWHM in linear space, the tuning width measure we use **A** and sigma in log space, the tuning width measure given for the macaque data **B**.

Individual differences

Individual differences in the clarity of map organization are apparent in the data, but we could not see any differences in clarity that could not be explained by differences in the relative locations of the numerosity map and the superior sagittal sinus.

Differences between subjects are evident in map size (Figs S8, S5.10), range of numerosity preferences in the map (Figs S8, S5.10), tuning width (Figs S5.13, S5.14), consistency between stimulus conditions (Fig S5.11), and topographic organization in the left hemisphere (Fig S5.9) and in the right posterior parietal lobe outside our ROI (Fig S5.8). Similar individual differences are present in the size of the visual cortex and its population tuning widths (Dougherty et al., 2003; Harvey and Dumoulin, 2011), and these have been linked to differences in performance in visual tasks (Duncan and Boynton, 2003; Schwarzkopf et al., 2011). It is tempting to speculate that individual differences in numerosity map structure may have perceptual consequences, such as differences in subitizing range. However, far more subjects would be needed to examine this question. It also seems possible that these differences may be related to inter-subject differences in experience with numerosity tasks. All of our subjects were experienced with numerosity tasks, and the extent of experience does not have any obvious relationship to individual differences in map structure.

Early visual cortex responses

Changing numerosity in a visual display affects visual features such as luminance, contrast, and spatial frequency, so establishing numerosity selectivity requires several control conditions in which these low-level features do not change with numerosity. On the other hand, early visual cortex is sensitive to these low-level features. Therefore, examining early visual cortex is useful to compare low-level visual responses to numerosity responses. These should differ if numerosity responses are driven by numerosity rather than low-level visual features. Recordings of responses to our numerosity stimulus from the central visual field representation in V1 show very different time courses to those in numerosity-selective regions (Fig S5.15). Unlike numerosity-selective responses (Fig 5.1B), responses amplitude is high when twenty dots are shown, but low when small numerosities are presented. These V1 response profiles cannot be captured by a neural model of numerosity selectivity. Instead, responses agree with conventional descriptions of primary visual cortex response properties, following the contrast energy in the visual stimulus. Thus, early visual cortex and numerosity selective cortex response profiles differ. This provides further evidence that low-level visual features do not drive the response profile in numerosity selective cortex.

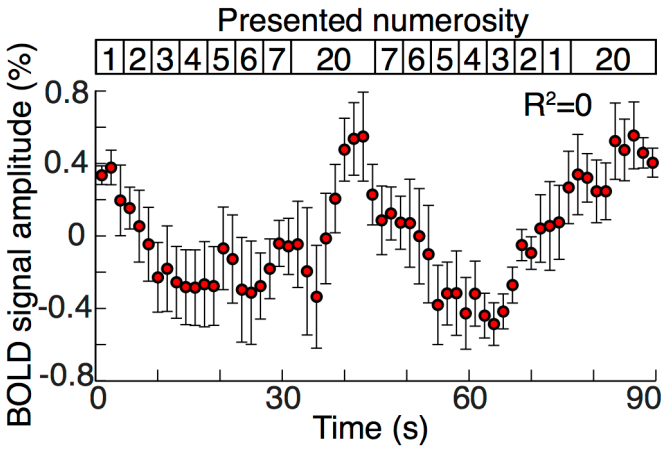

Fig S5.15. Response time course of a recording point in primary visual cortex to the numerosity stimulus, averaged over all conditions (top). The response is modulated by the stimulus sequence, but the response is highest when twenty dots are shown, rather than during presentation of small numerosities (as seen in the posterior parietal cortex, Fig 5.1B). These responses fit well with the contrast content of the stimulus, which is greatest when twenty dots are shown, in agreement with conventional descriptions of primary visual cortex response properties. The response is not strongly modulated by presentation of different small numerosities.

Visual field mapping stimuli

As the posterior parietal cortex contains several visual field maps (Schluppeck et al., 2005; Silver et al., 2005; Swisher et al., 2007; Konen and Kastner, 2008) we also acquired visual field mapping responses to ensure we did not confuse visual field position specific responses with numerosity-specific responses and to examine the relationship between these maps. The mapping paradigm was almost identical to that described in previous studies (Dumoulin and Wandell, 2008; Amano et al., 2009; Levin et al., 2010; Winawer et al., 2010; Harvey and Dumoulin, 2011; Hoffmann et al., 2012; Zuiderbaan et al., 2012). The stimulus consisted of drifting bar apertures at various orientations, which exposed a 1/f filtered noise pattern, with a contrast energy distribution that is inversely related to spatial frequency, as seen in natural images. This pattern was presented twice in each volume acquisition while the subject performed an attentionally demanding contrast discrimination task on a similar pattern presented at fixation. This design allowed us to strongly focus subjects' attention on the fixation point during all visual stimulation, which is a particular concern for the visual field maps in the parietal lobe that are strongly affected by attention (Schluppeck et al., 2005; Silver et al., 2005; Swisher et al., 2007; Konen and Kastner, 2008). The stimulus had a radius of 5.5°, larger than the numerosity stimulus (0.75° radius). Two diagonal red lines, which intersected at the center of the display, were also presented throughout the entire scanning run to provide a target for accurate fixation (Schira et al., 2009) even when the subject was not performing the task.

fMRI analysis for visual field mapping

Visual field mapping data were analyzed following a standard population receptive field analysis, as described elsewhere (Dumoulin and Wandell, 2008; Harvey and Dumoulin, 2011). We identified visual field map borders based on reversals in polar angle of visual field position preference (Sereno et al., 1995) and identified

particular visual field maps in the parietal lobe with reference to previous studies (Schluppeck et al., 2005; Swisher et al., 2007; Konen and Kastner, 2008; Bressler et al., 2013) (Fig S5.16). While visual field maps overlap with the numerosity map, there is no clear relationship between them.

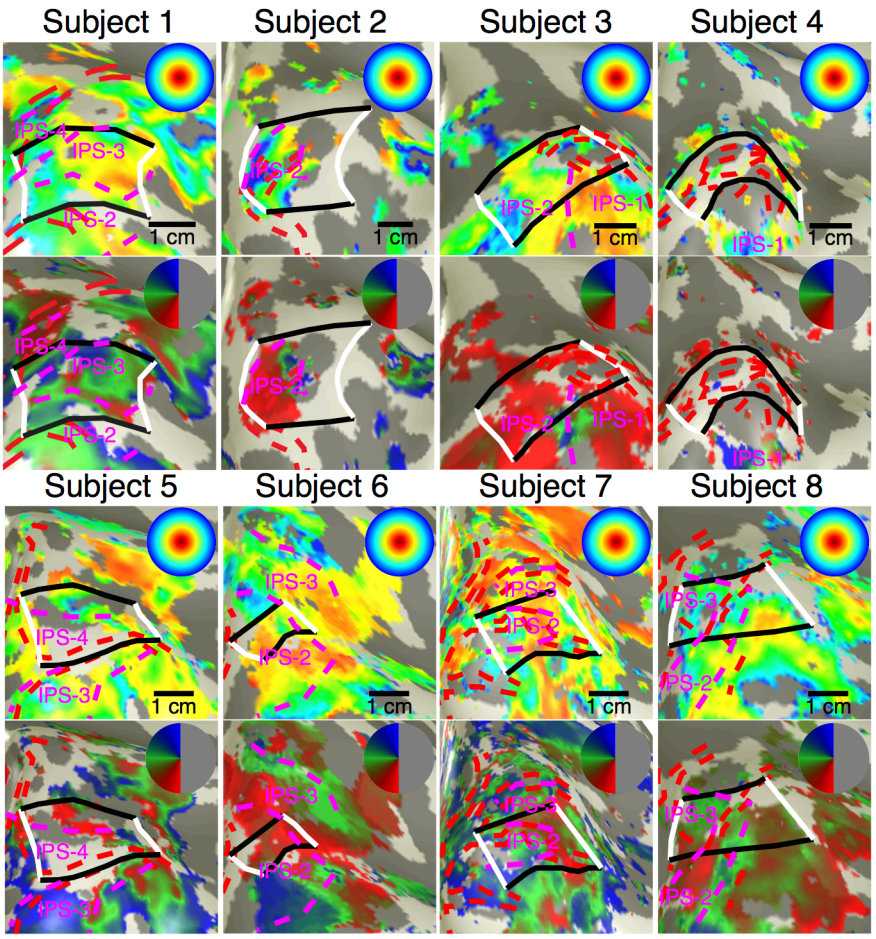

Fig S5.16. Visual field map representations around the numerosity map, showing the same views seen in Fig S5.8 and S5.13A. Borders between visual field maps are marked by dashed purple lines. Dashed red lines show locations of veins and black and white lines show the borders of numerosity map ROIs, as in previous figures. While visual field maps overlap with the numerosity map, there is no clear relationship between them. Insets demonstrate relationship between displayed color and preferred visual field position in visual field eccentricity and polar angle.

Response to Arabic numerals

We recorded responses to Arabic numerals presented and analyzed in the same paradigm as the numerosity stimulus. We find no evidence of numeral-tuned responses in the numerosity map. A representative example time series is shown in Fig S5.17. Numeral tuned models do not explain a significant amount of variance in this response (R^2=0.03), while the same site shows clear numerosity-tuned responses (R^2=0.83). The mean(sd) R^2 value within the ROI was 0.08(0.05) for numerals and 0.67(0.19) for numerosity dot patterns.

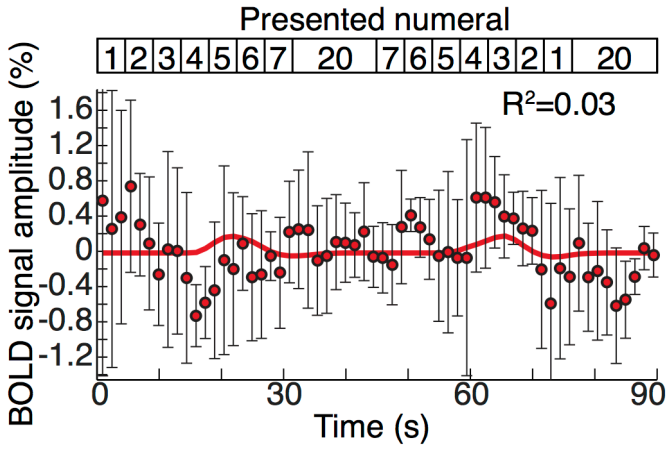

Fig S5.17. Responses to Arabic numerals presented and analyzed in the same paradigm as the numerosity stimulus. Models tuned to numeral preferences do not explain a significant amount of the signal variance (fits shown as red line), while this site shows clear numerosity tuning. Error bars show the standard error of the mean response to repeated runs (points).

Chapter 6

General discussion

In the introduction of this thesis, we argued that attention is a flexible mechanism that allocates the limited neural resources for processing visual information to information that is relevant for the observer. As such, attention has a strong influence on neural responses, affecting a variety of aspects of visual processing and perception. Conceptually, the effect of attention on visual processing can be described as a magnifying glass. Shortly, attention enlarges the representation of the attended information, at the expense of the representation of the information that is not attended. However, the empirical support for this concept is currently limited. In this thesis, we assessed if attention indeed acts as a magnifying glass on visual information in human visual cortex.

According the attention field model (Womelsdorf et al., 2008; Reynolds and Heeger, 2009) attention acts as a magnifying glass by modulating the selectivity of neurons towards the attended information. Importantly, in order to achieve this result, attention has to modulate the selectivity of all neurons that are selective for the attended information. However, current support for attentional modulation of neural selectivity comes from neurophysiological studies in non-human primates that can only measure from a limited number of neurons. We utilized several methodological advances in neuroimaging that allowed us to examine the attentional modulation of neural selectivity across large parts of visual cortex in humans and thus assess if attention indeed acts like a magnifying glass on visual information.

In Chapter 2 we demonstrated that spatial attention indeed acts like a magnifying glass on visual information. Furthermore, as will argue below, the results from this chapter also raise the possibility that spatial attention acts like a constant influence on visual processing across the entire visual cortex. In Chapter 3 we used high-resolution fMRI to show that a mixture of feed forward and feedback mechanisms underlie this attentional magnification of visual processing in human visual cortex. In Chapter 4, we related attentional magnification on visual processing to changes in human perception. We found that spatial attention produces perceptual magnification. This perceptual effect was captured by the attention field model that also captures attentional magnification of spatial position. Finally, in Chapter 5 we extended the biological models used in Chapter 2 to a novel visual feature: numerosity. This chapter demonstrates that the organization of numerosity processing is very similar that of spatial processing. This suggest that numerosity processing might be affected by attention in very similar was and that the concept of a magnifying glass is also present in stimulus dimensions that are orthogonal to spatial position. We will discuss these results in more detail below.

Spatial attention acts as a magnifying glass

In Chapter 2 we assessed whether spatial attention acts like a magnifying glass on visual processing. Specifically, we assessed how spatial attention modulates neural selectivity for spatial position across human visual cortex. To this end, we combined the attention field model with population receptive field (pRF) modeling (Dumoulin and Wandell, 2008). In this case, the attention field model predicts that spatial attention attracts pRFs towards its location. With pRF attraction, spatial attention modulates the spatial selectivity of neural populations towards the attended information and thus produces attentional magnification. If spatial attention acts like a magnifying glass, pRF attraction has to be present across the entire visual field and visual cortex. In line with these predictions by the attention field model, we measured pRF attraction across the visual field and visual cortex. These results provide support for the attentional magnifying glass.

Although we did not measure neural receptive fields directly, we believe that our conclusions regarding pRF attraction from Chapter 2 also apply to neural receptive fields. The pRF is the aggregate of the receptive fields of all neurons within a fMRI voxel. As such, both the receptive field and the pRF are a measure of the spatial selectivity of a neural population, but only at different spatial scales; the pRF position that we measure for every voxel will reflect the averaged receptive field position of all neurons within a voxel. Therefore, the pRF attraction that we measure here will also apply to individual neurons in the visual system.

Finally, note that we did not find any reliable changes in response amplitude related to pRF attraction. However, the attention field model predicts that spatial attention both attracts pRFs: the magnifying glass, and increases neural responses: the spotlight. This discrepancy can have many sources, like the visual stimulus we used or the specific analysis that was adopted. It is up to future studies to examine the exact relation between changes in response amplitude and pRF attraction.

Spatial attention acts as a single constant influence on visual processing

Another prediction of the attention field model is that attentional magnification is related to the size of the attracted pRFs. More specifically, larger pRFs should be attracted more, which in turn should lead to a larger magnification of the attended information. In Chapter 2, we measured increasingly stronger attentional magnification up the visual hierarchy. As pRF sizes increase up the visual hierarchy as well, the stronger magnification in later visual field maps was expected based on the attention field model. However, this does not reveal the extent in which pRF size affects the amount of magnification.

Within the experimental design we used in Chapter 2, the attention field model predicts that attentional magnification is a function of both the size of the pRF and the strength of the attentional influence, measured with the size of the attention field. Therefore, the increase in attentional magnification up the visual hierarchy could also be the result of a stronger influence of attention at later stages of visual processing.

We designed our experiment in Chapter 2 to be able to isolate the strength of the attentional influence on pRF position at every stage of visual processing. Doing this revealed that variation in pRF size is indeed the main factor underlying variation in attentional magnification. The strength of the attentional influence did not change much between different visual field maps. Indeed, our results were equally well explained by assuming a constant attentional influence across the entire visual hierarchy. This suggests that although spatial attention magnifies the attended information increasingly stronger up the visual hierarchy, it does this because of the larger pRFs at later stages of visual processing. The attentional influence itself appears to be a constant influence across the entire visual cortex.

Both feed forward and feedback mechanisms underlie attentional magnification

In Chapter 3 we assessed if the attentional magnification measured in Chapter 2 is implemented by either a feed forward or feedback mechanism in human primary visual cortex (V1). Mechanistic models broadly suggest two types of mechanisms that can underlie this magnification. Some models suggest that attentional magnification is implemented via feedback either progressing backwards across the visual hierarchy or originating from a single part of the visual system, affecting al stages simultaneously (Compte and Wang, 2006; Buffalo et al., 2010; Saalmann and Kastner, 2011; Bobier et al., 2014; Miconi and VanRullen, 2016). Alternatively, feed forward mechanisms have been proposed to underlie attention magnification as well. Here, magnification at one stage of the visual hierarchy results from attentional changes at an earlier stage of visual processing (Olshausen et al., 1993; McAdams and Maunsell, 1999; O'Connor et al., 2002; McAlonan et al., 2008; Briggs et al., 2013).

In Chapter 3, we contrasted these two mechanisms. We applied the attention field model to the known flow of feed forward and feedback information across cortical depth in primary visual cortex (V1). This yielded
distinctive predicted profiles of attentional magnification across cortical depth for a feed forward and a feedback mechanism. In addition, we derived two important points concerning attentional magnification across cortical depth: 1. A feed forward mechanism can not produce variation in attentional magnification across cortical depth. 2. The contribution of a feedback mechanism would manifest itself as

stronger magnification in certain portions of cortical depth. We measured the magnification across cortical depth in V1, using sub-millimeter (0.7mm isotropic) fMRI. This revealed that magnification occurs across cortical depth, but that it is stronger in cortical portions near the gray/white matter boundary. This specific pattern was best explained by a combination of feed forward and feedback mechanism contributing to attentional magnification.

It is important to note that the point above stating that a feed forward mechanism can not produce variation in attentional magnification not only applies to magnification across cortical depth, but to magnification across the visual hierarchy as well. So, if spatial attention would implement magnification only in one visual field map, the feed forward flow of information up the visual hierarchy would not produce a larger magnification in later visual field maps. In other words, if attention is implemented at only one stage of visual processing, we would not expect to measure an increase in magnification up the visual hierarchy from a feed forward mechanism only. However, we demonstrated in Chapter 2 that attentional magnification *does* increase up the visual hierarchy. We speculate that this means that spatial attention is implemented at every stage of visual processing separately. Consequently, the combination of feed forward and feedback mechanisms underlying attentional magnification in V1 may be applied at every stage of visual processing. This speculation is consistent with earlier ideas by Buschman and Kastner (2015).

Attentional magnification affects perception near the attended location

In Chapter 4 we linked the attentional magnification we measured in Chapter 2 to perception. Given the properties of attentional magnification, we expected two main effects of spatial attention on human perception. First, objects surrounding the location of spatial attention should be perceived to be further apart than they really are. Second, this bias should depend on the receptive field size of the neurons processing the objects. Specifically, larger receptive fields should lead to a larger bias. It is interesting to realize that these predictions concern a distortion of a perceptual property, i.e. spatial position, rather than an improvement, as is typically observed with attention.

To examine this, we measured the perceived distance between two objects located at opposite sides of the location of attention (Figure 4.1A). We compared the perceived distance between the objects surrounding the location of attention to the perceived distance between two identical objects that did not surround the attended location. We manipulated the receptive field sizes that processed the objects either

by varying the eccentric position at which they are presented, or by changing their spatial frequency content.

The results show that, consistent with our expectations, the perceived distance between the objects increased when the objects surrounded the focus of attention. Importantly, we also measured an increase of this effect when the objects were presented further away from the center of the visual field. However, the increase in perceived distance did not change with the spatial frequency of the objects. How can we reconcile these different results?

First, it should be noted that the relation between spatial frequency selectivity and receptive field size is only well described in one visual field map; V1 (Jones and Palmer, 1987), whereas the relation between receptive field size and eccentricity is present in every visual field map. Consequently, if the bias in perceived position is produced by changes in neural processing at stages other than V1 (DiGiacomo and Pratt, 2012), changing the spatial frequency of the object may not lead to a change in the receptive field sizes underlying the bias in perceived position of the objects. In contrast, increasing the eccentricity at which a stimulus is presented would vary the underlying receptive field size at every level of the visual hierarchy and thus induce variation in receptive field attraction throughout the visual hierarchy. This makes it more likely that the level of visual processing at which the positional bias is induced shows variation in receptive field attraction as well. Thus, the different ways by which eccentricity and spatial frequency are encoded across the visual hierarchy may have led to the difference between the results of the eccentricity and spatial frequency experiment.

A second point to note is that the attention field model we used in Chapter 4 is based on fMRI results obtained using endogenous attention. In the study assessing perception, we tested its perceptual consequences with exogenous attention. However, we do not believe that using exogenous attention instead of endogenous attention is responsible for the discrepancy between the predicted and measured positional biases in the spatial frequency experiment. Both endogenous and exogenous attention affect visual processing as early as the primary visual cortex (Motter, 1993; Gandhi et al., 1999; Somers et al., 1999; Liu et al., 2005) where spatial frequency–selective cells have been well described (Jones and Palmer, 1987). In addition, an attention field model captures the effects of both exogenous and endogenous attention on contrast sensitivity in human perception (Herrmann et al., 2010) suggesting that they affect visual processing in very similar ways. Moreover, positional biases have been found using both exogenous and endogenous attention (Suzuki and Cavanagh, 1997). Finally, using exogenous attention, we *did* find an increase in positional biases with eccentricity in the current study, as predicted by the model.

Finally, it is important to consider that our task was spatial in nature. Our cue was spatially localized, and we examined its effect on the spatial position of the Gabor stimuli. In the spatial frequency experiment, we manipulated the spatial frequency of the Gabor stimuli, a stimulus dimension that is independent from the stimulus' position. This discrepancy between the nature of the task and the nature of the varied stimulus characteristic could be responsible for the lack of an effect of spatial frequency on perceived position. It may be possible that the magnification in the spatial domain is not affected by modulations in the spatial frequency domain. Consistent with this possibility, attentional effects on the processing of spatial information are independent from its effects on the processing of other visual characteristics (David et al., 2008; Hayden and Gallant, 2009; Çukur et al., 2013). This suggests that the attentional magnification we measured in the spatial domain may be independent from attentional magnification in other visual characteristics. However, to examine if this is indeed the case using our approach, we need to develop biological models that capture the selectivity for these visual characteristics first.

Numerosity selective neural populations in human visual cortex

In Chapter 5, we modified the biological models used in Chapter 2 to describe numerosity selectivity. This approach revealed several highly interesting results: 1 Neural populations in the human brain are selective for numerosity. This is very much like the population receptive fields described in the introduction, but now populations have a specific preferred numerosity rather than a preferred spatial position. 2. These neural populations are organized in a topographic map. This is analogous to the retinotopic organization of visual field maps. However, instead of preferring neighboring visual field positions, neighboring populations in this case prefer similar numerosities. 3. Neural populations preferring higher numerosities respond to a wider range of numerosities. This is similar to the increase in receptive field size with increasing distance from the center of the visual field (eccentricity). 4. More neural populations are devoted to processing lower quantities than to processing higher quantities. This is similar to the cortical magnification factor, a fundamental property of visual field maps.

The similarities between visual spatial selectivity and numerosity selectivity are striking and allow for a straightforward application of the attention field model to numerosity processing. When applied to numerosity processing, the attention field model predicts that attention should act as a magnifying glass on numerosity as well. More specifically, the attention field model predicts that attention should modulate the sensitivity of numerosity selective neural populations towards the

attended quantity and thus produces an enlarged representation of the attended numerosity. Furthermore, the strength of this magnification should depend on the range of numerosities (i.e. the tuning width) for which a neural population is selective. Finally, this magnification in numerosity processing can affect the perceived numerosity as well: perceived numerosities would be repelled from the numerosity that is attended. Again, the magnitude of this effect would depend on the width of numerosity selectivity of a neural population underlying this effect.

As we now have the tools to describe attentional magnification for several visual characteristics reliably, we can gain insights in the extent in which the magnification of one characteristic is independent from the magnification of another. For example, if we attend a numerosity presented at a certain location, we can assess if varying attention for the location would change the magnification of numerosity. As numerosity likely shares a common neural mechanism with object size and even duration, this idea can possibly be extended to these qualities as well (Walsh, 2003; Harvey et al., 2015).

Finally, numerosity can be presented reliably via multiple modalities: visually through the amount of dots, auditory by a number of beeps, somatosensory by air puffs on the skin. This is in contrast to spatial position, which is best and most reliably presented through the visual sense. If the processing of numerosity we described in Chapter 5 proves to be independent of the modality used to present the numerosity, we could move even beyond the visual sense and investigate how attentional magnification operates across different modalities. For example, we could examine how attention manipulated via one modality would affect the representation of a numerosity presented via another modality.

Conclusions

In conclusion, we examined in this thesis if attention acts like a magnifying glass on visual processing and draws neural resources towards the attended information across visual cortex. We examined this specifically in the spatial domain. We used a unique approach that relied on novel fMRI methods and combined biological models of visual processing with computational models of attention. Doing so, we found that spatial attention indeed acts like a magnifying glass. By acting as a unified, global influence on visual processing, spatial attention enlarges the representation of attended information at the expense of information that is not attended. In order to achieve this, spatial attention relies on both feed forward and feedback mechanisms to modulate visual processing across the entire visual cortex. We demonstrated that spatial attention not only acts as a magnifying glass in visual processing, but also in human perception. The attended location appears larger than it actually is and

objects surrounding the attended location are perceived to be further apart than they really are. Finally, we paved the way for future studies to extend our approach to numerosity processing by developing biological models that capture selectivity for numerosity. This demonstrated that the processing of numerosity is very much alike the processing of spatial position and suggests that they may be affected by attention in very similar ways as well. Future studies are needed to examine how attention acts in other visual domains like numerosity, how this affects perception and how the effects of attention interact across different stimulus characteristics.

Appendix

References

Amano, K., Wandell, B.A., 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. doi:10.1152/jn.00102.2009

Anton-Erxleben, K., 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., Kastner, S., 2011. Visuotopic Organization of Macaque Posterior Parietal Cortex: A Functional Magnetic Resonance Imaging Study. J. Neurosci. 31:2064–2078. doi:10.1523/JNEUROSCI.3334-10.2011

Baruch, O., Yeshurun, Y., 2014. Attentional attraction of receptive fields can explain spatial and temporal effects of attention. Vis. cogn. 22:704–736. doi:10.1080/13506285.2014.911235

Bazin, P.L., Pham, D.L., 2007. Topology-preserving tissue classification of magnetic resonance brain images. IEEE Trans. Med. Imaging 26:487–496. doi:10.1109/TMI.2007.893283

Belopolsky, A. V, Zwaan, L., Theeuwes, J., Kramer, A.F., 2007. The size of an attentional window modulates attentional capture by color singletons. Psychon. Bull. Rev. 14:934–938. doi:10.3758/BF03194124

Benevento, L.A., Rezak, M., 1976. The cortical projections of the inferior pulvinar and adjacent lateral pulvinar in the rhesus monkey (macaca mulatta): An autoradiographic study. Brain Res. 108:1–24. doi:10.1016/0006-8993(76)90160-8

Benjamini, Y., Hochberg, Y., 1995. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. J. R. Stat. Soc. 57:289–300.

Blasdel, G.G., Lund, J.S., 1983. Termination of afferent axons in macaque striate cortex. J. Neurosci. 3:1389–1413.

Bobier, B., Stewart, T.C., Eliasmith, C., 2014. A Unifying Mechanistic Model of Selective Attention in Spiking Neurons. PLoS Comput. Biol. 10. doi:10.1371/journal.pcbi.1003577

Bongard, S., Nieder, A., 2010. Basic mathematical rules are encoded by primate prefrontal cortex neurons. Proc. Natl. Acad. Sci. 107:2277–2282. doi:10.1073/pnas.0909180107

Brainard, D.H., 1997. The Psychophysics Toolbox. Spat. Vis. 10:433–436.

Brannon, E.M., Terrace, H.S., 2016. Ordering of the Numerosities 1 to 9 by Monkeys Published by : American Association for the Advancement of Science Stable URL : http://www.jstor.org/stable/2899267

Brefczynski, J.A., DeYoe, E.A., 1999. A physiological correlate of the “spotlight” of visual attention. Nature 2:370–374.

Bressler, D.W., Fortenbaugh, F.C., Robertson, L.C., Silver, M.A., 2013. Visual spatial attention enhances the amplitude of positive and negative fMRI responses to visual stimulation in an eccentricity-dependent manner. Vision Res. 85:104–112. doi:10.1016/j.visres.2013.03.009

Briand, K.A., 1998. Feature integration and spatial attention: more evidence of a dissociation between endogenous and exogenous orienting. J. Exp. Psychol. Hum. Percept. Perform. 24:1243–1256. doi:10.1037//0096-1523.24.4.1243

Briand, K.A., Klein, R.M., 1987. Is Posner's "beam" the same as Treisman's "glue"?: On the relation between visual orienting and feature integration theory. J. Exp. Psychol. Hum. Percept. Perform. 13:228–241. doi:10.1037/0096-1523.13.2.228

Briggs, F., Callaway, E.M., 2001. Layer-specific input to distinct cell types in layer 6 of monkey primary visual cortex. J. Neurosci. 21:3600–3608. doi:21/10/3600 [pii]

Briggs, F., Mangun, G.R., Usrey, W.M., 2013. Attention enhances synaptic efficacy and the signal-to-noise ratio in neural circuits. Nature 499:476–80. doi:10.1038/nature12276

Britten, K.H., 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., Desimone, R., 2010. A backward progression of attentional effects in the ventral stream. Proc. Natl. Acad. Sci. U. S. A. 107:361–365. doi:10.1073/pnas.0907658106

Burr, D., Ross, J., 2008. A Visual Sense of Number. Curr. Biol. 18:425–428. doi:10.1016/j.cub.2008.02.052

Buschman, T.J., Kastner, S., 2015. From Behavior to Neural Dynamics: An Integrated Theory of Attention. Neuron 88:127–144. doi:10.1016/j.neuron.2015.09.017

Buschman, T.J., Miller, E.K., 2007. Top-Down Versus Bottom-Up Control of Attention in the Prefrontal and Posterior Parietal Cortices. Science (80-.). 315:1860–1862. doi:10.1126/science.1138071

Busse, L., Katzner, S., Treue, S., 2008. Temporal dynamics of neuronal modulation during exogenous and endogenous shifts of visual attention in macaque area MT. Proc. Natl. Acad. Sci. 105:16380–16385. doi:10.1073/pnas.0707369105

Callaway, E.M., 1998. Local Circuits in Primary Visual Cortex of the Macaque Monkey. Annu. Rev. Neurosci. 21:47–74. doi:10.1146/annurev.neuro.21.1.47

Cantlon, J.F., Brannon, E.M., 2007. Basic math in monkeys and college students. PLoS Biol. 5:e308. doi:10.1371/journal.pbio.0050328

Cantlon, J.F., Libertus, M.E., Pinel, P., Dehaene, S., Brannon, E.M., Pelphrey, K.A., 2009. The neural development of an abstract concept of number. J. Cogn. Neurosci. 21:2217–2229. doi:10.1162/jocn.2008.21159.

Carrasco, M., 2011. Visual attention: The past 25 years. Vision Res. 51:1484–1525.
Carrasco, M., 2006. Covert attention increases contrast sensitivity: psychophysical, neurophysiological and neuroimaging studies. Prog. Brain Res. 154:33–70.
Carrasco, M., Ling, S., Read, S., 2004. Attention alters appearance. Nat. Neurosci. 7:308–313. doi:10.1038/nn1194
Cheal, M.L., Lyon, D.R., 1991. Importance of precue location in directing attention. Acta Psychol. (Amst). 76:201–211.
Chen, B.L., Hall, D.H., Chklovskii, D.B., 2006. Wiring optimization can relate neuronal structure and function. Proc. Natl. Acad. Sci. 103:4723–4728. doi:10.1073/pnas.0506806103
Collins, D.L., Neelin, P., Peters, T., Evans, A.C., 1994. Automatic 3D Intersubject Registration of MR Volumetric Data in Standardized Talairach Space. J. Comput. Assist. Tomogr. 18:192–205.
Compte, A., Wang, X.J., 2006. Tuning curve shift by attention modulation in cortical neurons: A computational study of its mechanisms. Cereb. Cortex 16:761–778. doi:10.1093/cercor/bhj021
Connor, C.E., Gallant, J.L., Preddie, D.C., 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., Van Essen, D.C., 1997. Spatial Attention Effects in Macaque Area V4. J. Neurosci. 17:3201–3214.
Cook, E.P., 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. doi:22/5/1994 [pii]
Corbetta, M., Miezin, F.M., Dobmeyer, S., Shulman, G.L., Petersen, S.E., 1990. Attentional Modulation of Neural Processing of Shape, Color, and Velocity in Humans. Science (80-.). 248:1556–1559.
Corbetta, M., Shulman, G.L., 2002. Control of Goal-Directed and Stimulus-Driven Attention in the Brain. Nat. Rev. Neurosci. 3:201–215. doi:10.1038/nrn755
Cox, R.W., 1996. AFNI : Software for Analysis and Visualization of Functional Magnetic Resonance Neuroimages 162–173.
Çukur, T., Nishimoto, S., Huth, A.G., Gallant, J.L., 2013. Attention during natural vision warps semantic representation across the human brain. Nat. Neurosci. 16:763–770. doi:10.1038/nn.3381
Dakin, S.C., Tibber, M.S., Greenwood, J.A., Kingdom, F.A.A., Morgan, M.J., 2011. A common visual metric for approximate number and density. Proc. Natl. Acad. Sci. 108:19552–19557. doi:10.1073/pnas.1113195108
Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical Surface-Based Analysis. Neuroimage 9:179–194. doi:10.1006/nimg.1998.0395

David, S. V., Hayden, B.Y., Mazer, J. a., Gallant, J.L., 2008. Attention to Stimulus Features Shifts Spectral Tuning of V4 Neurons during Natural Vision. Neuron 59:509–521. doi:10.1016/j.neuron.2008.07.001

De Martino, F., Zimmermann, J., Muckli, L., Ugurbil, K., Yacoub, E., Goebel, R., 2013. Cortical Depth Dependent Functional Responses in Humans at 7T: Improved Specificity with 3D GRASE. PLoS One 8:e60514. doi:10.1371/journal.pone.0060514

Dehaene, S., 1997. The Number Sense: How the Mind Creates Mathematics. Oxford Univ. Press, New York.

Dehaene, S., Changeux, J.P., 1993. Development of Elementary Numerical Abilities: A Neuronal Model. J. Cogn. Neurosci. 5:390–407.

Dehaene, S., Cohen, L., 1991. Two mental calculation systems: a case study of severe acalculia with preseved approximation. Neuropsychologia 29:1045–1074.

Dehaene, S., Izard, V., Spelke, E., Pica, P., 2008. Log or Linear ? Distinct Intuitions of the number scale in Western and Amazonian indigene cultures. Science. 320:1217–1220. doi:10.1126/science.1156540

Dehaene, S., Spelke, E., Pinel, P., Stanescu, R., Tsivkin, S., 1999. Sources of Mathematical Thinking: Behavioral and Brain-Imaging Evidence. Science. 284:970–974. doi:10.1126/science.284.5416.970

DeValois, R.L., Albrecht, D.G., Thorell, L.G., 1982. Spatial frequency selectivity of cells in macaque visual cortex. Vision Res. 22:545–559. doi:10.1016/0042-6989(82)90113-4

DeYoe, E.A., Carman, G.J., Bandettini, P., Glickman, S., Wieser, J., Cox, R., Miller, D., Neitz, J., 1996. Mapping striate and extrastriate visual areas in human cerebral cortex. Proc. Natl. Acad. Sci. U. S. A. 93:2382–6. doi:10.1073/pnas.93.6.2382

DiGiacomo, A., Pratt, J., 2012. Misperceiving space following shifts of attention: Determining the locus of the attentional repulsion effect. Vision Res. 64:35–41. doi:10.1016/j.visres.2012.05.009

Dougherty, R.F., Koch, V.M., Brewer, A.A., Fischer, B., Modersitzki, J., Wandell, B.A., 2003. Visual field representations and locations of visual areas V1/2/3 in human visual cortex. J. Vis. 3:586–598. doi:10.1167/3.10.1

Dumoulin, S.O., Fracasso, A., van der Zwaag, W., Siero, J.C.W., Petridou, N., 2017. Ultra-high field MRI: Advancing systems neuroscience towards mesoscopic human brain function. Neuroimage. doi:10.1016/j.neuroimage.2017.01.028

Dumoulin, S.O., Hess, R.F., May, K.A., Harvey, B.M., Rokers, B., Barendregt, M., 2014. Contour extracting networks in early extrastriate cortex. J. Vis. 14:1–14.

Dumoulin, S.O., Wandell, B.A., 2008. Population receptive field estimates in human visual cortex. Neuroimage 39:647–660. doi:10.1016/j.neuroimage.2007.09.034

Duncan, R.O., Boynton, G.M., 2003. Cortical magnification within human primary visual cortex correlates with acuity thresholds. Neuron 38:659–671. doi:10.1016/S0896-6273(03)00265-4

Duvernoy, H.M., Delon, S., Vannson, J.L., 1981. Cortical blood vessels of the human brain. Brain Res. Bull. 7:519–579. doi:10.1016/0361-9230(81)90007-1

Eger, E., Michel, V., Thirion, B., Amadon, A., Dehaene, S., Kleinschmidt, A., 2009. Deciphering Cortical Number Coding from Human Brain Activity Patterns. Curr. Biol. 19:1608–1615. doi:10.1016/j.cub.2009.08.047

Engel, S.A., Glover, G.H., Wandell, B.A., 1997. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. Cereb. Cortex 7:181–192. doi:10.1093/cercor/7.2.181

Engel, S.A., Rumelhart, B.A., Wandell, B.A., Lee, A.T., Glover, G.H., Chichilnisky, E.J., Shadlen, M.N., 1994. fMRI of human visual cortex. Nature 369:525.

Eriksen, C.W., Hoffman, J.E., 1972. Some characteristics of selective attention in visual perception determined by vocal reaction time. Percept. Psychophys. 11:169–171. doi:10.3758/BF03210367

Eriksen, C.W., St. James, J.D., 1986. Visual attention within and around the field of focal attention: A zoom lens model. Percept. Psychophys. 40:225–240. doi:10.3758/BF03211502

Eriksen, C.W., Yeh, Y.-Y., 1985. Allocation of attention in the visual field. J. Exp. Psychol. Hum. Percept. Perform. doi:10.1037/0096-1523.11.5.583

Fedorov, A., Beichel, R., Kalpathy-Cramer, J., Finet, J., Fillion-Robin, J.C., Pujol, S., Bauer, C., Jennings, D., Fennessy, F., Sonka, M., Buatti, J., Aylward, S., Miller, J. V., Pieper, S., Kikinis, R., 2012. 3D Slicer as an image computing platform for the Quantitative Imaging Network. Magn. Reson. Imaging 30:1323–1341. doi:10.1016/j.mri.2012.05.001

Felleman, D.J., Van Essen, D.C., 1991. Distributed hierachical processing in the primate cerebral cortex. Cereb. Cortex 1:1–47.

Fitzpatrick, D., Lund, J.S., Blasdel, G.G., 1985. Intrinsic connections of macaque striate cortex: afferent and efferent connections of lamina 4C. J. Neurosci. 5:3329–3349.

Fracasso, A., Petridou, N., Dumoulin, S.O., 2016. Systematic variation of population receptive field properties across cortical depth in human visual cortex. Neuroimage 139:427–438. doi:10.1016/j.neuroimage.2016.06.048

Friston, K.J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M.D., Turner, R., 1998. Event-Related fMRI: Characterizing Differential Responses. Neuroimage 7:30–40. doi:10.1006/nimg.1997.0306

Gandhi, S.P., Heeger, D.J., Boynton, G.M., 1999. Spatial attention affects brain activity in human primary visual cortex. Proc. Natl. Acad. Sci. 96:3314–3319.

Glover, G.H., 1999. Deconvolution of Impulse Response in Event-Related BOLD fMRI. Neuroimage 9:416–429. doi:10.1006/nimg.1998.0419

Haacke, E. M., Brown, R. W., Thompson, M. R., & Venkatesan, R., 1999. *Magnetic resonance imaging: physical principles and sequence design* (Vol. 82). New York:: Wiley-Liss.

Haak, K. van, Winawer, J., Harvey, B.M., Renken, R., Dumoulin, S.O., Wandell, B.A., Cornelissen, F.W., 2013. Connective field modeling. Neuroimage 66:376–384.

Haenny, P.E., 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. doi:10.1007/BF00247569

Halberda, J., Mazzocco, M.M.M., Feigenson, L., 2008. Individual differences in non-verbal number acuity correlate with maths achievement. Nature 455:665–668. doi:10.1038/nature07246

Han, X., Pham, D.L., Tosun, D., Rettmann, M.E., Xu, C., Prince, J.L., 2004. CRUISE: Cortical reconstruction using implicit surface evolution. Neuroimage 23:997–1012. doi:10.1016/j.neuroimage.2004.06.043

Harvey, B.M., 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. doi:10.1523/JNEUROSCI.2572-11.2011

Harvey, B.M., Fracasso, A., Petridou, N., Dumoulin, S.O., 2015. Topographic representations of object size and relationships with numerosity reveal generalized quantity processing in human parietal cortex 112:13525–13530. doi:10.1073/pnas.1515414112

Hayden, B.Y., Gallant, J.L., 2009. Combined effects of spatial and feature-based attention on responses of V4 neurons. Vision Res. 49:1182–1187. doi:10.1016/j.visres.2008.06.011

Hembrook-Short, J.R., Mock, V.L., Briggs, F., 2017. Attentional Modulation of Neuronal Activity Depends on Neuronal Feature Selectivity. Curr. Biol. 27:1–10. doi:10.1016/j.cub.2017.05.080

Hernández, M., Costa, A., Humphreys, G.W., 2010. The size of an attentional window affects working memory guidance. Attention, Perception, Psychophys. 72:963–972. doi:10.3758/APP.72.4.963

Herrmann, K., Montaser-Kouhsari, L., Carrasco, M., Heeger, D.J., 2010. When size matters: attention affects performance by contrast or response gain. Nat. Neurosci. 13:1554–1559. doi:10.1038/nn.2669

Heuer, H.W., Britten, K.H., 2002. Contrast dependence of response normalization in area MT of the rhesus macaque. J. Neurophysiol. 88:3398–3408. doi:10.1152/jn.00255.2002

Hoffmann, M.B., Kaule, F.R., Levin, N., Masuda, Y., Kumar, A., Gottlob, I., Horiguchi, H., Dougherty, R.F., Stadler, J., Wolynski, B., Speck, O., Kanowski, M., Liao, Y.J., Wandell, B.A., Dumoulin, S.O., 2012. Plasticity and Stability of the Visual System in Human Achiasma. Neuron 75:393–401. doi:10.1016/j.neuron.2012.05.026

Hopf, J.-M., Boehler, C.N., Luck, S.J., Tsotsos, J.K., Heinze, H.-J., Schoenfeld, M.A., 2006. Direct neurophysiological evidence for spatial suppression surrounding the focus of attention in vision. Proc. Natl. Acad. Sci. 103:1053–1058. doi:10.1073/pnas.0507746103

Hubel, D.H., Wiesel, T.N., 1974. Uniformity of monkey striate cortex: A parallel relationship between field size, scatter, and magnification factor. J. Comp. Neurol. 158:295–305. doi:10.1002/cne.901580305

Hubel, D.H., Wiesel, T.N., 1972. Laminar and columnar distribution of geniculo-cortical fibers in the macaque monkey. J. Comp. Neurol. 146:421–50. doi:10.1002/cne.901460402

Hubel, D.H., Wiesel, T.N., 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J. Physiol. 160:106–154.

Intriligator, J., Cavanagh, P., 2001. The Spatial Resolution of Visual Attention. Cogn. Psychol. 43:171–216. doi:10.1006/cogp.2001.0755

Johnston, W.A., Heinz, S.P., 1978. Flexibility and Capacity Demands of Attention. J. Exp. Psychol. Gen. 107:420–435. doi:10.1037/0096-3445.107.4.420

Johnston, W., Dark, V., 1986. Selective attention. Annu. Rev. Psychol. 37:43–75. doi:10.1146/annurev.ps.37.020186.000355

Jones, J.P., Palmer, L.A., 1987. An Evaluation of the Two-Dimensional Gabor Filter Model of Simple Receptive Fields in Cat Striate Cortex. J. Neurophysiol. 58:1233–1258.

Kastner, S., De Weerd, P., Desimone, R., Ungerleider, L.G., 1998. Mechanisms of Directed Attention in the Human Extrastriate Cortex as Revealed by Functional MRI. Science (80-.). 282:108–111.

Kay, K.N., Winawer, J., Mezer, A., Wandell, B.A., 2013. Compressive spatial summation in human visual cortex. J. Neurophysiol. 110:481–494. doi:10.1152/jn.00105.2013

Kerkoerle, T. van., Self, M.W., Roelfsema, P.R., 2017. Effects of attention and working memory in the different layers of monkey primary visual cortex. Nat. Commun. 8:13804. doi:10.1038/ncomms13804

Klein, B.P., Harvey, B.M., Dumoulin, S.O., 2014. Attraction of Position Preference by Spatial Attention throughout Human Visual Cortex. Neuron 84:227–237. doi:10.1016/j.neuron.2014.08.047

Klein, B.P., Paffen, C.L.E., te Pas, S.F., Dumoulin, S.O., 2016. Predicting bias in perceived position using attention field models. J. Vis. 16:1–15. doi:10.1167/16.7.15

Kok, P., Bains, L.J., Van Mourik, T., Norris, D.G., De Lange, F.P., 2016. Selective activation of the deep layers of the human primary visual cortex by top-down feedback. Curr. Biol. 26:371–376. doi:10.1016/j.cub.2015.12.038

Konen, C.S., Kastner, S., 2008. Representation of Eye Movements and Stimulus Motion in Topographically Organized Areas of Human Posterior Parietal Cortex. J. Neurosci. 28:8361–8375. doi:10.1523/JNEUROSCI.1930-08.2008

Kosovicheva, A.A., Fortenbaugh, F.C., Robertson, L.C., 2010. Where does attention go when it moves?: Spatial properties and locus of the attentional repulsion effect. J. Vis. 10:33–33. doi:10.1167/10.12.33

Kumano, H., Uka, T., 2010. The Spatial Profile of Macaque MT Neurons Is Consistent With Gaussian Sampling of Logarithmically Coordinated Visual Representation. J. Neurophysiol. 104:61–75.

LaBerge, D., 1983. Spatial extent of attention to letters and words. J. Exp. Psychol. Hum. Percept. Perform. 9:371–379. doi:10.1037/0096-1523.9.3.371

Larsson, J., Heeger, D.J., 2006. Two Retinotopic Visual Areas in Human Lateral Occipital Cortex. J. Neurosci. 26:13128–13142. doi:10.1523/JNEUROSCI.1657-06.2006

Lawrence, S.J.D., Formisano, E., Muckli, L., de Lange, F.P., 2017. Laminar fMRI: Applications for cognitive neuroscience. Neuroimage 1–7. doi:10.1016/j.neuroimage.2017.07.004

Lennie, P., 2003. The Cost of Cortical Computation. Curr. Biol. 13:493–497. doi:10.1016/S0960-9822(03)00135-0

Leuba, G., Garey, L.J., 1989. Comparison of neuronal and glial numerical density in primary and secondary visual cortex of man. Exp. Brain Res. 77:31–38. doi:10.1007/BF00250564

Levin, N., Dumoulin, S.O., Winawer, J., Dougherty, R.F., Wandell, B.A., 2010. Cortical Maps and White Matter Tracts following Long Period of Visual Deprivation and Retinal Image Restoration. Neuron 65:21–31. doi:10.1016/j.neuron.2009.12.006

Liu, T., Pestilli, F., Carrasco, M., 2005. Transient Attention Enhances Perceptual Performance and fMRI Response in Human Visual Cortex. Neuron 45:469–477. doi:10.1016/j.neuron.2004.12.039

Lu, Z.-L., Dosher, B.A., 2000. Spatial attention: different mechanisms for central and peripheral temporal precues? J. Exp. Psychol. Hum. Percept. Perform. 26:1534–1548. doi:10.1037//0096-1S23.26J.1534

Lund, J.S., Lund, R.D., Hendrickson, A.E., Bunt, A.H., Fuchs, A.F., 1975. The origin of efferent pathways from the primary visual cortex, area 17, of the macaque monkey as shown by retrograde transport of horseradish peroxidase. J. Comp. Neurol. 164:287–303. doi:10.1002/cne.901640303

Marques, J.P., Kober, T., Krueger, G., van der Zwaag, W., Van de Moortele, P.F., Gruetter, R., 2010. MP2RAGE, a self bias-field corrected sequence for improved segmentation and T1-mapping at high field. Neuroimage 49:1271–1281. doi:10.1016/j.neuroimage.2009.10.002

Martinez-Trujillo, J.C., Treue, S., 2002. Attentional Modulation Strength in Cortical Area MT. Neuron 35:365–370. doi:10.1016/S0896-6273(02)00778-X

Martínez-Trujillo, J.C., Treue, S., 2004. Feature-Based Attention Increases the Selectivity of Population Responses in Primate Visual Cortex. Curr. Biol. 14:744–751.

Maunsell, J.H., Gibson, J.R., 1992. Visual response latencies in striate cortex of the macaque monkey. J. Neurophysiol. 68:1332–1344. doi:10.1113/jphysiol.1962.sp006837

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

McAdams, C.J., Maunsell, J.H.R., 1999. Effects of Attention on Orientation-Tunnig Functions of Single Neurons in Macaque Cortical Area V4. J. Neurosci. 19:431–441.

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

McMains, S.A., Somers, D.C., 2004. Multiple spotlights of attentional selection in human visual cortex. Neuron 42:677–686. doi:10.1016/S0896-6273(04)00263-6

Merriam, E.P., Colby, C.L., 2005. Active Vision in Parietal and Extrastriate Cortex. Neurosci. 11:484–493. doi:10.1177/1073858405276871

Miconi, T., VanRullen, R., 2016. A Feedback Model of Attention Explains the Diverse Effects of Attention on Neural Firing Rates and Receptive Field Structure. PLoS Comput. Biol. 12:1–18. doi:10.1371/journal.pcbi.1004770

Montijn, J.S., Klink, P.C., van Wezel, R.J.A., 2012. Divisive Normalization and Neuronal Oscillations in a Single Hierarchical Framework of Selective Visual Attention. Front. Neural Circuits 6:1–17. doi:10.3389/fncir.2012.00022

Moran, J., Desimone, R., 1985. Selective Attention Gates Visual Processing in the Extrastriate Cortex. Science (80-.). 229:782–784.

Motter, B.C., 2009. Central V4 Receptive Fields Are Scaled by the V1 Cortical Magnification and Correspond to a Constant-Sized Sampling of the V1 Surface. J. Neurosci. 29:5749–5757.

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.

Muckli, L., De Martino, F., Vizioli, L., Petro, L.S., Smith, F.W., Ugurbil, K., Goebel, R., Yacoub, E., 2015. Contextual Feedback to Superficial Layers of V1. Curr. Biol. 25:2690–2695. doi:10.1016/j.cub.2015.08.057

Nakayama, K., Mackeben, M., 1989. Sustained and transient components of focal visual attention. Vision Res. 29:1631–1647.

Nandy, A.S., Nassi, J.J., Reynolds, J.H., 2017. Laminar Organization of Attentional Modulation in Macaque Visual Area V4. Neuron 93:1–12. doi:10.1016/j.neuron.2016.11.029

Nestares, O., Heeger, D.J., 2000. Robust multiresolution alignment of MRI brain volumes. Magn. Reson. Med. 43:705–715. doi:10.1002/(SICI)1522-2594(200005)43:5<705::AID-MRM13>3.0.CO;2-R

Niebergall, R., Khayat, P.S., Treue, S., Martinez-Trujillo, J.C., 2011a. Expansion of MT Neurons Excitatory Receptive Fields during Covert Attentive Tracking. J. Neurosci. 31:15499–15510. doi:10.1523/JNEUROSCI.2822-11.2011

Niebergall, R., Khayat, P.S., Treue, S., Martinez-Trujillo, J.C., 2011b. Multifocal attention filters targets from distracters within and beyond primate mt neurons' receptive field boundaries. Neuron 72:1067–1079. doi:10.1016/j.neuron.2011.10.013

Nieder, A., Freedman, D.J., Miller, E.K., 2002. Representation of the Quantity of Visual Items in the Primate Prefrontal Cortex. Science (80-.). 297:1708–1711. doi:10.1126/science.1072493

Nieder, A., Miller, E.K., 2004. A parieto-frontal network for visual numerical information in the monkey. Proc. Natl. Acad. Sci. U. S. A. 101:7457–7462. doi:10.1073/pnas.0402239101

Nieder, A., Miller, E.K., 2003. Coding of cognitive magnitude: Compressed scaling of numerical information in the primate prefrontal cortex. Neuron 37:149–157. doi:10.1016/S0896-6273(02)01144-3

O'Connor, D.H., Fukui, M.M., Pinsk, M.A., 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., Yacoub, E., 2012. Layer-specific fmri reflects different neuronal computations at different depths in human V1. PLoS One 7. doi:10.1371/journal.pone.0032536

Olshausen, B.A., Anderson, C.H., Van Essen, D.C., 1993. A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. J. Neurosci. 13:4700–4719. doi:10.1.1.66.2555

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

Petridou, N., Italiaander, M., van de Bank, B.L., Siero, J.C.W., Luijten, P.R., Klomp, D.W.J., 2013. Pushing the limits of high-resolution functional MRI using a simple high-density multi-element coil design. NMR Biomed. 26:65–73. doi:10.1002/nbm.2820

Piazza, M., Izard, V., Pinel, P., Le Bihan, D., Dehaene, S., 2004. Tuning curves for approximate numerosity in the human intraparietal sulcus. Neuron 44:547–555. doi:10.1016/j.neuron.2004.10.014

Pinel, P., Le Clec'H, G., Van de Moortele, P.F., Naccache, L., Le Bihan, D., Dehaene, S., 1999. Event-related fMRI analysis of the cerebral circuit for number comparison. Neuroreport 10:1473–1479.

Posner, M.I., 1980. Orienting of attention. Q. J. Exp. Psychol. 32:3–25. doi:10.1080/00335558008248231

Posner, M.I., Gilbert, C.D., 1999. Attention and primary visual cortex. Proc. Natl. Acad. Sci. 96:2585–2587. doi:10.1073/pnas.96.6.2585

Puckett, A.M., DeYoe, E.A., 2015. The Attentional Field Revealed by Single-Voxel Modeling of fMRI Time Courses. J. Neurosci. 35:5030–5042. doi:10.1523/JNEUROSCI.3754-14.2015

Ramon y Cajal, S., Pasik, P., Pasik, T., 1999. Texture of the Nervous System of Man and the Vertebrates. Springer, New York.

Reynolds, J.H., Heeger, D.J., 2009. The Normalization Model of Attention. Neuron 61:168–185. doi:10.1016/j.neuron.2009.01.002

Rockel, A.J., Hiorns, R.W., Powell, T.P.S., 1980. The basic uniformity in structure of the neocortex. Brain 103:221–244.

Rockland, K.S., Pandya, D.N., 1979. Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. Brain Res. 179:3–20. doi:10.1016/0006-8993(79)90485-2

Saad, Z.S., Glen, D.R., Chen, G., Beauchamp, M.S., Desai, R., Cox, R.W., 2009. Alignment using Local Pearson Correlation. Neuroimage 44:839–848. doi:10.1016/j.neuroimage.2008.09.037.A

Saalmann, Y.B., Kastner, S., 2011. Cognitive and Perceptual Functions of the Visual Thalamus. Neuron 71:209–223. doi:10.1016/j.neuron.2011.06.027

Saltzman, I.J., Garner, W.R., 1948. Reaction time as a measure of span of attention. J. Psychol. 25:227–241.

Schira, M.M., Tyler, C.W., Breakspear, M., Spehar, B., 2009. The Foveal Confluence in Human Visual Cortex. J. Neurosci. 29:9050–9058. doi:10.1523/JNEUROSCI.1760-09.2009

Schluppeck, D., Climcher, P., Heeger, D.J., 2005. Topographic organization for delayed saccades in human posterior parietal cortex. J. Neurophysiol. 94:1372–1384. doi:10.1152/jn.01290.2004.

Schwarzkopf, D.S., Song, C., Rees, G., 2011. The surface area of human V1 predicts the subjective experience of object size. Nat. Neurosci. 14:28–30. doi:10.1038/nn.2706

Self, M.W., Kerkoerle, T. van, Supèr, H., Roelfsema, P.R., 2013. Distinct Roles of the Cortical Layers of Area V1 in Figure-Ground Segregation. Curr. Biol. 23:2121–2129. doi:10.1016/j.cub.2013.09.013

Self, M.W., Roelfsema, P.R., 2017. Paying Attention to the Cortical Layers. Neuron 93:9–11. doi:10.1016/j.neuron.2016.12.032

Self, M.W., van Kerkoerle, T., Goebel, R., Roelfsema, P.R., 2017. Benchmarking laminar fMRI: Neuronal spiking and synaptic activity during top-down and bottom-up processing in the different layers of cortex. Neuroimage 1–12. doi:10.1016/j.neuroimage.2017.06.045

Sereno, M.I., Dale, a M., Reppas, J.B., Kwong, K.K., Belliveau, J.W., Brady, T.J., Rosen, B.R., Tootell, R.B.H., 1995. Borders of Multiple Visual Areas in Humans Revealed by Functional Magnetic Resonance Imaging Borders of Multiple Visual Areas in Humans Revealed by Functional Magnetic Resonance Imaging. Science (80-.). 268:889–893.

Shim, W.M., Cavanagh, P., 2005. Attentive tracking shifts the perceived location of a nearby flash. Vision Res. 45:3253–3261. doi:10.1016/j.visres.2005.05.029

Shipp, S., 2003. The functional logic of cortico-pulvinar connections. Philos. Trans. R. Soc. Lond. B. Biol. Sci. 358:1605–1624. doi:10.1007/s00429-016-1250-9

Silver, M. a, Ress, D., Heeger, D.J., 2005. Topographic Maps of Visual Spatial Attention in Human Parietal Cortex. J. Neurophysiol. 94:1358–1371. doi:10.1152/jn.01316.2004.

Smith, A.T., Singh, K.D., Williams, A.L., Greenlee, M.W., 2001. Estimating Receptive Field Size from fMRI Data in Human Striate and Extrastriate Visual Cortex. Cereb. Cortex 11:1182–1190. doi:10.1093/cercor/11.12.1182

Somers, D.C., Dale, A.M., Seiffert, A.E., Tootell, R.B.H., 1999. Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. PNAS 96:1663–1668.

Sommer, M.A., Wurtz, R.H., 2008. Brain Circuits for the Internal Monitoring of Movements. Annu. Rev. Neurosci. 31:317–338. doi:10.1146/annurev.neuro.31.060407.125627

Spitzer, H., Desimone, R., Moran, J., 1988. Increased Attention Enhances Both Behavioral and Neuronal Performance. Science (80-.). 240:240–338.

Starkey, P., 1992. The early development of numerical reasoning. Cognition 43:93–126. doi:10.1016/0010-0277(92)90034-F

Suzuki, S., Cavanagh, P., 1997. Focused Attention Distorts Visual Space: An Attentional Repulsion Effect. J. Exp. Psychol. Hum. Percept. Perform. 23:443–463.

Swisher, J.D., Halko, M.A., Merabet, L.B., McMains, S.A., Somers, D.C., 2007. Visual Topography of Human Intraparietal Sulcus. J. Neurosci. 27:5326–5337. doi:10.1523/JNEUROSCI.0991-07.2007

Takemura, H., Ashida, H., Amano, K., Kitaoka, A., Murakami, I., 2012. Neural Correlates of Induced Motion Perception in the Human Brain. J. Neurosci. 32:14344–14354. doi:10.1523/JNEUROSCI.0570-12.2012

Teo, P.C., Sapiro, G., Wandell, B. a, 1997. Creating connected representations of cortical gray matter for functional MRI visualization. IEEE Trans. Med. Imaging 16:852–863. doi:10.1109/42.650881

Tokita, M., Ishiguchi, A., 2010. How might the discrepancy in the effects of perceptual variables on numerosity judgment be reconciled? Atten. Percept. Psychophys. 72:1839–1853. doi:10.3758/APP

Treisman, A.M., 1960. Contextual cues in selective listening. Q. J. Exp. Psychol. 12:242–248. doi:10.1080/17470216008416732

Treue, S., 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.

Truong, T., Chen, B., Song, A.W., 2008. Integrated SENSE DTI with correction of susceptibility- and eddy current-induced geometric distortions 40:53–58. doi:10.1016/j.neuroimage.2007.12.001

Tudusciuc, O., Nieder, A., 2007. Neuronal population coding of continuous and discrete quantity in the primate posterior parietal cortex. Proc. Natl. Acad. Sci. U. S. A. 104:14513–8. doi:10.1073/pnas.0705495104

Usrey, W.M., Fitzpatrick, D., 1996. Specificity in the axonal connections of layer VI neurons in tree shrew striate cortex: evidence for distinct granular and supragranular systems. J. Neurosci. 16:1203–1218.

Van Essen, D., 2003. DEL_ME, in: Chalupa, L.M., Werner, J.S. (Eds.), The Visual Neurosciences. MIT Press, Cambridge, MA, pp. 507–521.

Van Essen, D.C., Newsome, W.T., Maunsell, J.H.R., 1984. The visual field representation in striate cortex of the macaque monkey: asymmetries, anisotropies, and individual variability. Vision Res. 24:429–448.

Viswanathan, P., Nieder, A., 2013. Neuronal correlates of a visual "sense of number" in primate parietal and prefrontal cortices. Proc. Natl. Acad. Sci. 110:11187–11192. doi:10.1073/pnas.1308141110

Waehnert, M.D., Dinse, J., Weiss, M., Streicher, M.N., Waehnert, P., Geyer, S., Turner, R., Bazin, P.L., 2014. Anatomically motivated modeling of cortical laminae. Neuroimage 93:210–220. doi:10.1016/j.neuroimage.2013.03.078

Walsh, V., 2003. A theory of magnitude: Common cortical metrics of time, space and quantity. Trends Cogn. Sci. 7:483–488. doi:10.1016/j.tics.2003.09.002

Wandell, B.A., Chial, S., Backus, B.T., 2000. Visualization and measurement of the cortical surface. J. Cogn. Neurosci. 12:739–752. doi:10.1162/089892900562561

Wandell, B.A., Dumoulin, S.O., Brewer, A.A., 2007. Visual Field Maps in Human Cortex. Neuron 56:366–383. doi:10.1016/j.neuron.2007.10.012

Wandell, B.A., Winawer, J., Kay, K.N., 2015. Computational modeling of topographic arrangements in human visual cortex, in: Toga, A.W. (Ed.), Brain Mapping. Academic Press, pp. 651–659. doi:10.1227/01.neu.0000279839.00334.f7

Wardak, C., Olivier, E., Duhamel, J.R., 2011. The relationship between spatial attention and saccades in the frontoparietal network of the monkey. Eur. J. Neurosci. 33:1973–1981. doi:10.1111/j.1460-9568.2011.07710.x

Watson, A.B., Pelli, D.G., 1983. QUEST: A Bayesian adaptive psychometric method. Percept. Psychophys. 33:113–120.

Whalen, J., Gallistel, C.R., Gelman, R., 1999. Nonverbal Counting in Humans: The Psychophysics of Number Representation. Psychol. Sci. 10:130–137. doi:10.1111/1467-9280.00120

Williams, A.L., Smith, A.T., 2010. Representation of Eye Position in the Human Parietal Cortex 104:2169–2177. doi:10.1152/jn.00713.2009.

Winawer, J., Horiguchi, H., Sayres, R.A., Amano, K., Wandell, B.A., 2010. Mapping hV4 and ventral occipital cortex: The venous eclipse. J. Vis. 10:1–1. doi:10.1167/10.5.1

Womelsdorf, T., Anton-Erxleben, K., Pieper, F., Treue, S., 2006. Dynamic shifts of visual receptive fields in cortical area MT by spatial attention. Nat. Neurosci. 9:1156–1160. doi:10.1038/nn1748

Womelsdorf, T., Anton-Erxleben, K., Treue, S., 2008. Receptive Field Shift and Shrinkage in Macaque Middle Temporal Area through Attentional Gain Modulation. J. Neurosci. 28:8934–8944.

Worsley, K.J., Liao, C.H., Aston, J., Petre, V., Duncan, G.H., Morales, F., Evans, A.C., 2002. A General Statistical Analysis for fMRI Data. Neuroimage 15:1–15. doi:10.1006/nimg.2001.0933

Wurtz, R.H., 2008. Neuronal mechanisms of visual stability. Vision Res. 48:2070–2089. doi:10.1016/j.visres.2008.03.021

Yacoub, E., Hu, X., 2001. Detection of the Early Decrease in fMRI in the Motor Area. Magn. Reson. Med. 45:184–190.

Yeshurun, Y., Carrasco, M., 1998. Attention improves or impairs visual performance by enhancing spatial resolution. Nature 396:72–75.

Yeshurun, Y., Montagna, B., Carrasco, M., 2008. On the flexibility of sustained attention and its effects on a texture segmentation task. Vision Res. 48:80–95. doi:10.1016/j.visres.2007.10.015

Yoshioka, T., Levitt, J.B., Lund, J.S., 1994. Independence and merger of thalamocortical channels within macaque monkey primary visual cortex: Anatomy of interlaminar projections. Vis. Neurosci. 11:467–489. doi:10.1017/S0952523800002406

Yushkevich, P.A., Piven, J., Hazlett, H.C., Smith, R.G., Ho, S., Gee, J.C., Gerig, G., 2006. User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. Neuroimage 31:1116–1128. doi:10.1016/j.neuroimage.2006.01.015

Zuiderbaan, W., Harvey, B.M., Dumoulin, S.O., 2012. Modeling center – surround configurations in population receptive fields using fMRI. J. Vis. 12:1–15. doi:10.1167/12.3.10.Introductio

Appendix

Nederlandse samenvatting

Het zicht is het meest belangrijke zintuig in mensen. In ons brein is ongeveer 25% van de cerebrale cortex betrokken bij het verwerken van visuele informatie. Niettemin is de informatieverwerkingscapaciteit van ons brein beperkt en kan niet alle visuele informatie verwerkt worden. Hierdoor is het noodzakelijk om de informatieverwerkingscapaciteit dynamisch te kunnen richten op de verwerking van informatie die op dat moment het meest relevant is voor ons. Visuele aandacht is het mechanisme waardoor wij informatie kunnen selecteren voor verdere verwerking.

Aandachtsselectie kan via twee manieren plaats vinden: onvrijwillig (exogeen) en vrijwillig (endogeen). Bij exogene aandacht trekt de informatie aandacht naar zichzelf toe, bijvoorbeeld wanneer we iets onverwachts zien. In het geval van endogene aandacht richten we onze aandacht vrijwillig op visuele informatie. Daarnaast kunnen we de aandachtsselectie baseren op verschillende aspecten van visuele informatie. Bijvoorbeeld, we kunnen onze aandacht richten op een bepaalde locatie in het gezichtsveld. We spreken in dit geval van spatiële aandacht. We kunnen onze aandacht ook richten op een specifieke kleur, lijnoriëntatie of *numerosity* (de hoeveelheid items in een set). In dit geval spreken we van *feature* aandacht.

Conceptueel gezien kan het proces van aandachtsselectie vergeleken worden met een *spotlight* die visuele informatie *highlight* of een vergrootglas die de geselecteerde informatie uitvergroot. De *spotlight* en vergrootglas concepten van aandacht kunnen vertaald worden naar het effect van aandacht op neurale responsen in ons visuele systeem. In dit geval is de *spotlight* vergelijkbaar met sterkere neurale responsen op de geselecteerde informatie en het vergrootglas met een uitvergroting van de neurale representatie van de geselecteerde informatie.

Opmerkelijk is dat recente computationele modellen van aandacht, de zogenaamde aandachtsveld modellen, voorspellen dat aandacht zowel de neurale responsen versterkt en de representatie van de geselecteerde informatie vergroot. Met andere woorden, de *spotlight* en vergrootglas van aandacht zijn volgens het aandachtsveld model twee zijden van dezelfde medaille. Ondanks dat er op dit moment voldoende empirische ondersteuning is voor aandacht als een *spotlight* is de huidige ondersteuning voor aandacht als een vergrootglas onvoldoende. In dit proefschrift hebben we onderzocht of aandacht inderdaad als een vergrootglas van visuele informatie kan functioneren.

Om dit te onderzoeken bouwen we verder op het aandachtsveld model. Volgens dit model kan aandacht functioneren als een vergrootglas door te interacteren met de selectiviteit van neuronen. Visuele neuronen zijn selectief voor een beperkt bereik van visuele informatie, bijvoorbeeld voor een bepaalde locatie of een bepaalde

numerosity. Als aandacht inderdaad functioneert als een vergrootglas is het cruciaal dat het interacteert met de selectiviteit voor de geselecteerde informatie van alle visuele neuronen. Dus als je je aandacht richt op een bepaalde locatie zou de spatiële selectiviteit van alle neuronen moeten worden gemoduleerd.

In hoofdstuk 2 gaan we na of dit daadwerkelijk het geval is in mensen. We maken gebruik van recent ontwikkelde data-analyses die ons in staat stellen om de spatiële selectiviteit in een groot gedeelte van het visuele brein in mensen te meten. We vinden dat dit inderdaad het geval is: spatiële aandacht moduleert de spatiële selectiviteit in grote delen van het visuele systeem en opereert in dat opzicht als een vergrootglas. Een belangrijke andere bevinding beschreven in dit hoofdstuk is dat de invloed van aandacht even sterk is in ieder deel van het visuele systeem. Dit suggereert dat spatiële aandacht een gecoördineerde, uniforme invloed op spatiële selectiviteit uitoefent.

In hoofdstuk 3 gaan we na hoe deze vergroting van de geselecteerde informatie is geïmplementeerd in het visuele systeem. Eerdere literatuur noemt twee mogelijke mechanismen. De eerste mogelijkheid is een *feed-forward* mechanisme. Binnen dit mechanisme leiden aandachtseffecten in delen van het visuele systeem die informatie als eerste verwerken leiden tot vergroting in delen die de informatie later verwerken. De tweede mogelijkheid is een feedback mechanisme waarbij latere delen van het visuele systeem een vergroting implementeren in vroegere delen. Deze twee mechanismen genereren onderscheidbare voorspellingen met betrekking tot het profiel van de aandachtsvergroting op microniveau in het visuele systeem. In dit hoofdstuk gebruiken we geavanceerde technieken om de aandachtsvergroting op microniveau te kunnen meten in het menselijke visuele systeem en dit vervolgens te vergelijken met de voorspelde profielen. We vinden dat de aandachtsvergroting is geïmplementeerd door een combinatie van een *feed forward* en feedback mechanisme.

Het aandachtsveld model voorspelt niet alleen een vergroting op neuraal niveau, maar ook op perceptueel niveau. In hoofdstuk 4 gaan we na of dit ook daadwerkelijk het geval is. We vinden dat spatiële aandacht inderdaad leidt tot een vergroting in de waarneming rondom de aandachtslocatie. Verder voorspelt het aandachtsveld model dat de mate van deze vergroting door aandacht afhangt van de locatie van aandacht ten opzichte van het centrum van het gezichtsveld en de spatiële frequentie van de stimulus die gebruikt wordt om de vergroting te meten. Dit lijkt inderdaad zo te zijn in het geval van de locatie van aandacht, maar niet voor de spatiële frequentie van de gebruikte stimulus. We speculeren dat deze discrepantie

het gevolg is van de verschillende wijzen waarop het visuele systeem spatiële positie en spatiële frequentie verwerkt.

In hoofdstukken 2, 3 en 4 hebben we spatiële aandacht en vergroting van spatiële posities onderzocht. Echter bevat visuele informatie veel meer aspecten dan spatiële positie alleen en kunnen we onze aandacht ook richten op meer dan alleen spatiële locaties. We kunnen onze aandacht bijvoorbeeld ook richten op een specifieke *numerosity*.

Numerosity refereert aan het aantal items in een set. De accuraatheid waarmee we *numerosity* kunnen waarnemen neemt af wanneer het aantal items in een set toeneemt. Het waarnemen van *numerosity* is een eigenschap die onafhankelijk is van andere numerieke vaardigheden zoals tellen en rekenen. Verder lijkt het verwerken van *numerosity* in sommige opzichten sterk op die van primaire zintuiglijke waarneming, zoals het zicht. Vanwege deze eigenschappen verwachten wij dat de mechanismen betrokken bij de verwerking van *numerosity* sterk lijken op de mechanismen betrokken bij het verwerken van visuospatiële informatie.

In hoofdstuk 5 ontwikkelen we data-analyses waarmee we *numerosity* selectiviteit kunnen meten in het menselijke brein. We gebruiken deze techniek om de verwerking van *numerosity* in mensen in kaart te brengen. Dit laat zien dat *numerosity* verwerking in veel aspecten sterk lijkt op de verwerking van visuospatiële informatie. Door deze overeenstemmingen kunnen de methoden gebruikt in dit proefschrift eenvoudig worden vertaald naar de verwerking van *numerosity*. Op deze wijze baant hoofdstuk 5 de weg voor het toepassen van het aandachtsveld model op meerdere stimulus dimensies en te onderzoeken of aandacht ook in deze dimensies functioneert als een vergrootglas.

Appendix

Dankwoord

Serge, zonder jou was ik nooit zo ver gekomen. Ik vond het erg prettig om met je samen te werken de afgelopen jaren. Bedankt voor al je geduld, wijsheid en geloof in mij.

Susan en Chris, bedankt voor alle keren dat jullie drafts van mijn papers door hebben gelezen en al het meedenken tijdens onze meetings.

Joke, het begon allemaal bij jouw suggestie om even bij Serge Dumoulin langs te gaan.

Ben, thanks for being a very motivated and intelligent colleague. I enjoyed collaborating with you and I was always happy to get your input on some new data or ideas.

Alessio, thanks for your pleasant collaboration. I enjoyed working with you and was always happy to have a coffee and a little chat with you.

Wietske, het was altijd prettig om even met jou te kunnen babbelen, ook over andere dingen dan werk.

Jelle, Akhil, Ada, Jutta, Yuxuan and Carlien, thanks for appreciating my sometimes not so very appropriate humor. And, once again, sorry for all the beeps.

Pim en Marianne, bedankt voor jullie vriendschap en altijd de mogelijkheid om elkaar te kunnen spreken.

Hester en Marcel, bedankt voor al jullie gastvrijheid en altijd even langs te kunnen komen.

Vera en Petrick, bedankt voor de regelmatige 'slechte horror en doorzak avonden/nachten'. Ze zijn erg welkom!

Bernie en Froukje, bedankt voor jullie steun en altijd klaar te staan om op Juliëtte te passen wanneer we er echt even tussen uit moesten

Daniëlle en Michael, bedankt voor alle keren dat jullie voor ons hebben gekookt en we ontspannen met elkaar konden eten. Bedankt voor jullie steun tijdens de moeilijke perioden afgelopen jaren.

Mark en Erna, bedankt voor alle ontspannen vakanties en voor jullie steun in de zware tijden afgelopen jaren.

Wil en Minke, bedankt voor alle steun, gastvrijheid en alle keren dat jullie op Juliëtte wilden passen.

Ma, bedankt voor al uw steun de afgelopen jaren. Bedankt dat u altijd klaarstond om ons te ontvangen of op Juliëtte te passen.

Pa, het spijt mij dat u er niet bij kan zijn. Bedankt voor al uw steun en geloof in mij. Ik denk dat u trots op mij zou zijn geweest.

Neelie, bedankt voor al de ruimte die je me de hebt gegeven en dat ik bij jou altijd mezelf kan zijn. Bedankt dat je er altijd bent om me weer in de goede richting te duwen. Ik hou van je!

Juliëtte, je beseft waarschijnlijk nog niet hoeveel je voor me betekent. Zonder jou was dit werk niet tot stand gekomen. Bedankt voor je lieve onbevangenheid. Ik hou van je!

Appendix

List of publications

Klein BP, Fracasso A, Paffen CLE, te Pas, SF, Dumoulin SO. (under review). Cortical depth dependent population receptive field attraction by spatial attention in human V1.

Klein BP, Paffen CLE, Pas SF, Dumoulin SO. (2016). Predicting bias in perceived position using attention field models. *Journal of Vision* 16(7):15, 1-15

Klein BP, Harvey BM, Dumoulin SO. (2014). Attraction of position preference by spatial attention throughout human visual cortex. *Neuron* 84: 227-237.

Harvey BM, **Klein BP**, Petridou N, Dumoulin SO (2013). Topographic representation of numerosity in the human parietal cortex.. *Science* 341: 1123-1126.

Van Der Stigchel S, Bethlehem RAI, **Klein BP**, Berendschot TTJM, Nijboer T, Dumoulin SO. (2013). Macular degeneration affects eye movement behaviour during visual search. *Frontiers in Psychology* 4: 579.

Appendix

Curriculum Vitae

Barrie P. Klein was born on the 4th of September, 1985 in Goudswaard (Korendijk), the Netherlands. In 2001 he graduated from his secondary education at the Willem van Oranje in Oud-Beijerland. After this period he worked as a butcher for several years. In 2006 he started his study Psychodiagnostisch Werk at Saxion University of Applied Sciences in Deventer. In 2008, he moved on to study Psychology at Utrecht University which he completed in 2011 with highest distinction (*Cum Laude*). He continued his academic education in Utrecht and completed the research master programme Neuroscience & Cognition, again with highest distinction (*Cum Laude*). In 2013 he started as a PhD student at the Department of Experimental Psychology of Utrecht University, under supervision of Prof. dr. S.O. Dumoulin. He will now start as a software engineer at Ordina in Nieuwegein.