

# Attention Alters Feature Space in Motion Processing

Marc Zirnsak<sup>1</sup> and Fred H. Hamker<sup>1,2</sup>

<sup>1</sup>Department of Psychology, Institute II, Westfälische Wilhelms-University, 48149 Münster, Germany, and <sup>2</sup>Department of Computer Science, Artificial Intelligence, Chemnitz University of Technology, 09107 Chemnitz, Germany

Attending to a particular feature modulates the activity of neurons throughout the visual field with the result that relevant features are enhanced while irrelevant features are suppressed. Do these modulatory influences merely lead to a gating of relevant features, or does attention have a direct impact on the representation of feature space, leading to a different percept depending on the content of attention? We observed that direction estimates of the static motion aftereffect drastically change when human observers attend to a stimulus whose motion direction differs from the one of the adaptor. This observation suggests that feature-based attention might operate by local magnifications of feature space between relevant and irrelevant features.

## Introduction

Visual attention refers to our ability to focus on only a small part of the vast amount of information obtained from our visual environment. The neural basis of this ability relies on gain control, that is, mechanisms that alter the input–output ratio of individual neurons in visual brain areas. Attending to a particular location in visual space enhances neural responses to stimuli presented at this location, known as spatial attention (for review, see Reynolds and Chelazzi, 2004). Similar effects have been reported if attention is directed toward a specific, nonspatial attribute (Chelazzi et al., 1993, 1998; Motter, 1994a,b; Treue and Martínez Trujillo, 1999; Martínez-Trujillo and Treue, 2004; Bichot et al., 2005). This feature-based attention is proposed to operate globally throughout the whole visual field, that is, the activity of neurons is enhanced if their tuning characteristics match with the attended feature even if their receptive field location is distinct from the present focus of spatial attention. Moreover, these modulations are found whether the attended feature is currently present or not. For instance, Chelazzi et al. (1993, 1998) reported that the activity of inferior temporal cortex neurons was higher if the monkey had to remember a preferred stimulus than if it had to remember a less preferred one. Similarly, Treue and Martínez Trujillo (1999) demonstrated that the response of direction-selective neurons in the middle temporal area (area MT) to their preferred direction was enhanced, relative to a neutral condition, if the monkey attended to the same direction of motion that was presented outside the receptive field of the recorded neuron. Furthermore, the response to the preferred direction was suppressed if the monkey attended to the anti-preferred direction of the recorded neuron.

Related effects have also been observed in humans. In Saenz et al. (2002), the general design used by Treue and Martínez Trujillo (1999) was adapted and used in a functional magnetic resonance imaging study, replicating their findings in terms of the strength of the hemodynamic response. Furthermore, Boynton et al. (2006) demonstrated that the motion aftereffect measured at a stimulus location different from the attended stimulus was stronger if both stimuli shared the same than if they had the opposite direction of motion. Finally, in Saenz et al. (2003), subjects had to detect a change in speed of two spatially separated motion stimuli. Detection performance was significantly better when both stimuli moved in the same rather than in opposite directions.

While the above-described observations are in accordance with a global upregulation and downregulation of the neural representation of features according to the feature similarity gain model of attention (Treue and Martínez Trujillo, 1999), it remains unknown whether feature-based attention has a direct impact on feature space. That is, does feature-based attention merely scale the population response, resulting in a more active representation of relevant features without altering their identity, or does it in addition alter the encoded feature at the level of neural populations?

## Materials and Methods

If feature-based attention indeed alters feature space, the static motion aftereffect (SMAE) caused by an unattended motion stimulus will change when attention is directed to a different direction of motion (Fig. 1A). The SMAE is the consequence of viewing a moving pattern for a prolonged period and designates the experience of motion while observing a stationary pattern. Thereby, the perceived direction of the stationary pattern points in the opposite direction of the moving pattern, due to an unbalanced population response to the stationary pattern presumably caused by neural adaptation of motion-sensitive cells preferring the previously viewed direction of motion (for review, see Mather et al., 2008). We will use this relationship to infer the direction of motion that was encoded by the neural population. We hypothesize that attending to a certain direction of motion will alter the encoded motion direction of an additional dissimilar stimulus, resulting either in an attraction toward or a repulsion away from the attended direction. However, if both stimuli share the same direction, no change in the encoded direction of the

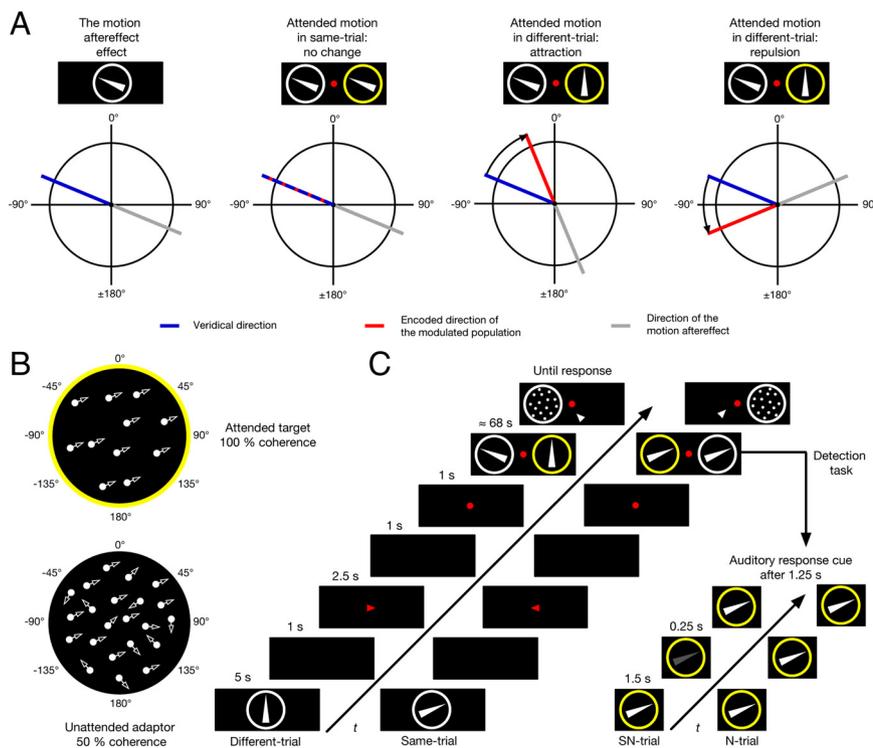
Received July 22, 2009; revised March 18, 2010; accepted March 29, 2010.

F.H.H. was supported by the German Science Foundation (DFG HA 2630/6-1), the European Commission (FP7-ICT: Eyeshots), and the Federal Ministry of Education and Research (BMBF 01GW0653). We are grateful to Tirin Moore, Rufin VanRullen, and Nick Steinmetz for valuable comments on a previous version of the manuscript.

Correspondence should be addressed to Marc Zirnsak at his present address: Department of Neurobiology, Stanford University School of Medicine, Stanford, CA 94305. E-mail: mzirnsak@stanford.edu.

DOI:10.1523/JNEUROSCI.3543-09.2010

Copyright © 2010 the authors 0270-6474/10/306882-09\$15.00/0



**Figure 1.** Hypotheses, stimuli, and experimental protocol. **A**, Illustration of the motion aftereffect and hypotheses. From left to right: Viewing a given direction for a prolonged time leads to the subsequent perception of motion in the opposite direction. If two motion patterns with the same direction are within the visual field, attending to one of them should not systematically alter the encoded direction of the other. However, if the motion direction in both patterns differ, directing attention to one of them should alter the encoded direction of the other and thus cause a change of the motion aftereffect, which was used in this study to infer the encoded direction. **B**, Illustration of RDKs used in the experiment. The yellow circle indicates the attended target RDK. The circle was not present in the experiments. **C**, In different-trials, the target was always directed upward ( $0^\circ$ ), while the dots in the adaptor were moving in a different direction. In same-trials, both target and adaptor were moving in the same direction. Yellow circles indicate the attended site. The initial cue provided the subjects information about the target motion direction. During the adaptation period, the subjects' task was the detection of a recurring luminance change as illustrated in the lower right. During the so-called signal-plus-noise trials (SN-trials), the luminance of the target RDK changed. In so-called noise trials (N-trials) the luminance of the target remained constant. After multiple detection trials in the prolonged adaptation period, the static test pattern was presented.

unattended stimulus should be observed as indicated by the measured SMAE, given that both are encoded veridically or at least equally.

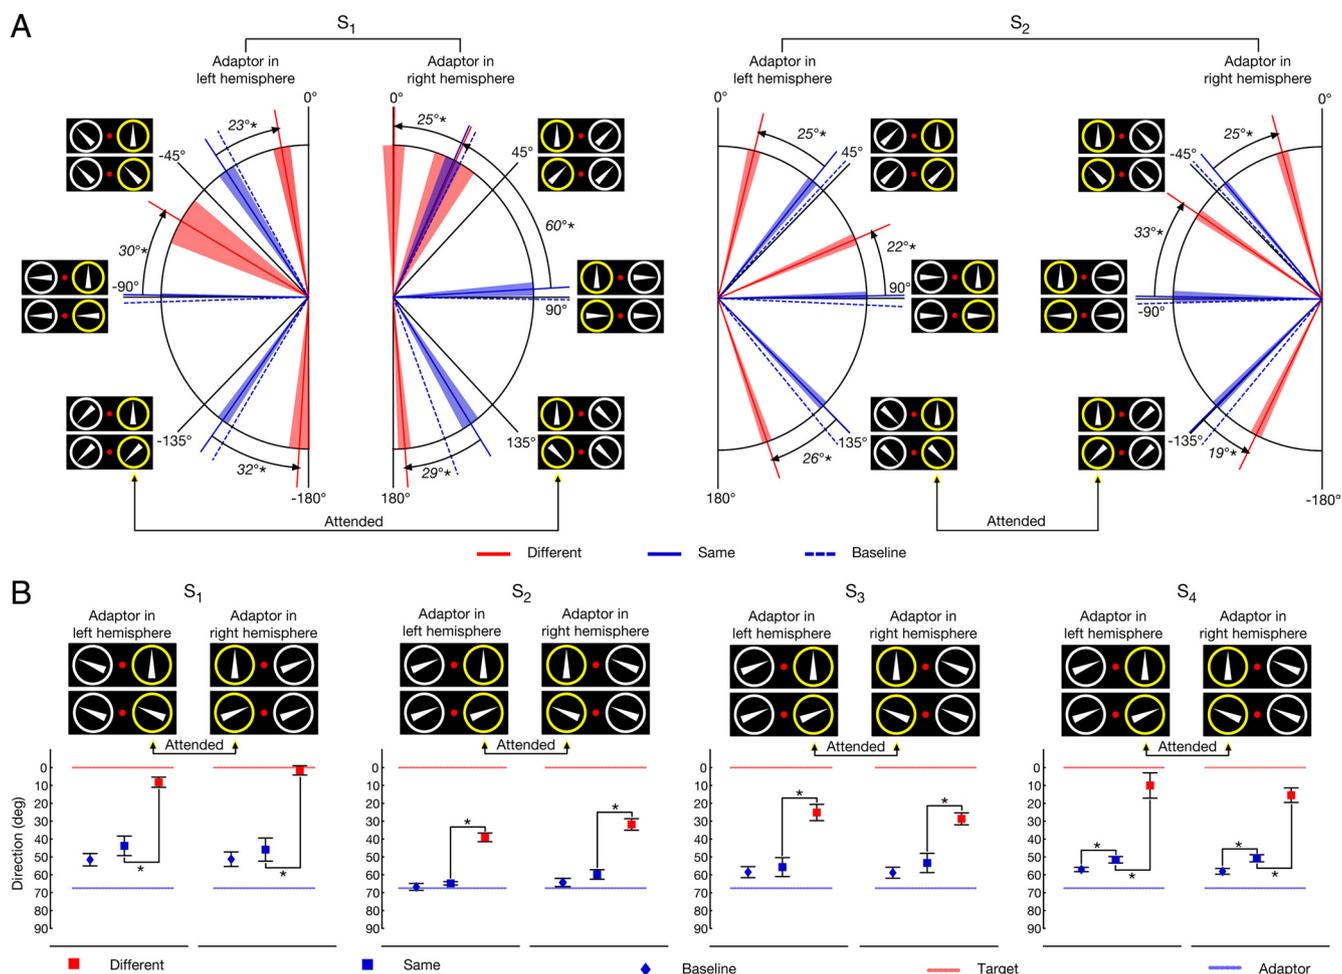
**Main experiment.** While maintaining central fixation subjects attended to a peripheral random-dot kinematogram (RDK), which we will refer to as the “target.” Simultaneously, a second RDK was presented in the opposite hemifield, which we will call the “adaptor” (Fig. 1B). While ignoring the adaptor, the subjects' task was the detection of a small recurring luminance change of the target. After an adaptation period, the target disappeared and the adaptor stopped moving, providing the stationary test pattern, that is, a stationary random-dot cloud with identical characteristics to the RDK except for the motion, to determine the subjects' perceived aftereffect. Two types of main trials were conducted. In so-called “same-trials,” target and adaptor always moved in the same direction. In “different-trials,” the target always moved upward while the adaptor moved in a different direction (Fig. 1C).

Four subjects ( $S_1, \dots, S_4$ ), two female, with normal or corrected-to-normal vision volunteered in the experiments. With the exception of  $S_2$ , who was one of the authors, all subjects were naive with respect to the purpose of the study. Subjects were instructed to attend to the target's motion direction. This task was supported by letting the subjects detect a luminance change of the target. In fact, we regard the instruction to focus attention on the motion direction as the crucial manipulation with respect to the goal of the study, while the luminance detection was primarily used as an indicator to infer the direction of subjects' space-based attention. Before the experiment, subjects were made familiar with the SMAE in both central and peripheral viewing con-

ditions using the same stimuli as described below with 100% coherence.

In the main experiment, the target RDK consisted of 100 and the adaptor RDK consisted of 200 white dots with 100% and 50% coherence, respectively. Note that the coherence of the adaptor was reduced to decrease the population response (Britten and Newsome, 1998) since, for reasons described in supplemental material, section 1.1 (available at [www.jneurosci.org](http://www.jneurosci.org)), we hypothesize that the strength of the feature-based attention would be inversely dependent on the strength of the bottom-up driven activity. Dots of both RDKs were  $0.28^\circ$  in diameter, had a luminance of  $85.6 \text{ cd/m}^2$  (MINOLTA LS-110), and moved in a circular aperture of  $8^\circ$  in diameter with a velocity of  $7.6^\circ/\text{s}$ . The random part of the adaptor changed direction every two frames. All dots had unlimited lifetime as long as they stayed inside the aperture. When a dot left the aperture, it was randomly placed at a new position at the border, and it was ensured that the dot would move inside the aperture again. The apertures were centered  $9^\circ$  left and right to a small ( $0.37^\circ$  in diameter) red fixation point with a luminance of  $20.2 \text{ cd/m}^2$  placed in the center of the screen on a black background with a luminance of  $2.2 \text{ cd/m}^2$ . Stimuli were presented on a cathode ray tube monitor (EIZO FlexScan F930) of 40 cm in width and 30 cm in height with a resolution of  $1856 \times 1392$  pixels and 85 Hz driven by an NVIDIA GeForce 7300 GT graphic card placed in a Mac Pro 1.1. Stimuli were generated with customized software programmed in C using OpenGL. Subjects were placed 53 cm in front of the monitor. A chin rest was used to stabilize the head. The room was dimly illuminated with 5 lx.

A single trial was structured as follows. To cue the target direction, the RDK to be attended was first presented centrally for 5 s. After the presentation of a blank (1 s), a small red arrow cued the attended site for 2.5 s. After another blank (1 s), the fixation point appeared, and 1 s later both target and adaptor were shown. To support the verbal instruction of directing attention to the target motion across multiple trials, subjects were also required to detect a luminance change. After 1.75 s, the luminance of the target was decreased in half of the detection trials for 0.25 s. A short auditory response cue was provided 1.25 s later. After the response of the subject, by pressing one of the two mouse buttons, feedback was given by the same auditory cue for a correct answer and the next detection trial was started. After 20 of these detection trials (end of adaptation period with an average duration of 68 s), the target disappeared and the adaptor stopped moving, thereby providing the stationary test pattern. If the subjects perceived an aftereffect, they indicated its direction by a mouse click relative to the fixation point. Note that this rather long adaptation period was chosen to obtain a high number of trials with perceived aftereffects. After a blank of 2.5 s, the next trial began. During three consecutive trials, the direction of the target and the adaptor never changed. The same is true for the attended site. Note that if detection performance of the luminance change as determined across these three adaptation trials fell outside a predefined range of  $1.35 < d' < 2.56$ , which corresponds to unbiased percentage correct range of  $75\% < p(c)_{\text{max}} < 90\%$ , all three responses were discarded. The initial value of the luminance change was obtained with the method of constant stimuli for each subject, where the value corresponding to  $d' = 1.69$  was chosen. If performance fell outside the predefined range, these values were adjusted in small steps. Information regarding the size of the luminance changes is shown in supplement



**Figure 2.** Dependency of the SMAE on the similarity of the attended stimulus and the adaptor. **A**, Perceived direction of SMAE for two subjects for six adaptor directions ( $-135^\circ$ ,  $-90^\circ$ ,  $-45^\circ$ ,  $45^\circ$ ,  $90^\circ$ ,  $135^\circ$ ). For  $S_1$ , the adaptor was always directed outward with respect to the fovea, and for  $S_2$ , it was always directed inward. The panel at each of the six adaptor directions indicates the presented motion directions in the same- and different-conditions. Mean baselines for each direction are indicated by the dashed blue lines. Mean SMAEs from same-trials are indicated by the solid blue lines. Mean SMAEs of different-trials are indicated by the red lines. For better readability, the indicated SMAEs are transformed to the direction of motion that would have caused them, assuming a simple reversal (Fig. 1A). The colored region around the respective means represents 95% confidence intervals. Arrows denote the differences between the same-condition and the different-condition. Asterisks indicate statistical significance. Yellow circles denote the attended target RDK. **B**, Results for all subjects ( $-67.5^\circ$ ,  $67.5^\circ$ ). For  $S_1$ , the adaptor was always directed outward. For all other subjects, it was directed inward. Mean SMAEs of different-conditions are indicated by red squares. Mean SMAEs of same-conditions are indicated by blue squares, and mean baselines are indicated by blue diamonds. SMAEs are shown as the unsigned direction that would have caused them. Dashed blue lines indicate the unsigned adaptor direction, and dashed red lines indicate the attended direction. Error bars denote 95% confidence intervals.

tal Table S1 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). After a 3 min break, the attended site was changed and the direction of the target and the adaptor was also changed. Note that this rather long period was chosen to avoid transfer effects of adaptation across conditions and to give subjects time to recover, since the task was experienced by all subjects as quite exhaustive. An experimental session consisted of 18 trials lasting  $\sim 50$  min.

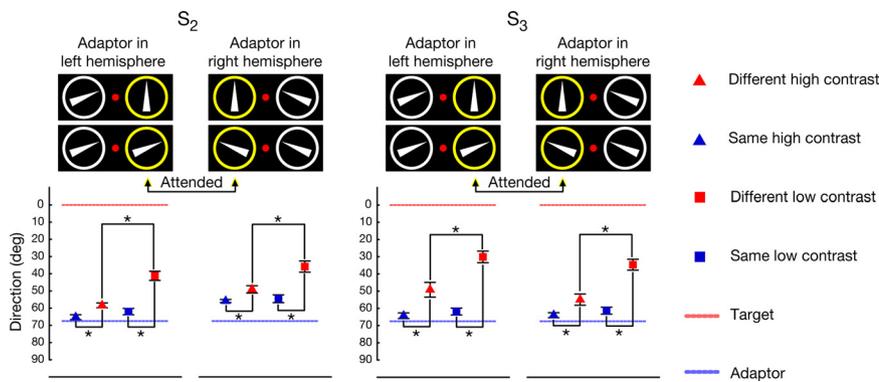
Same-trials and different-trials were randomized within and across sessions. In same-trials, both target and adaptor had the same direction. In different-trials, target direction was always upward ( $0^\circ$ ), and as in same-trials, the direction of the adaptor was randomly selected out of 10 possible directions ( $-135^\circ$ ,  $-112.5^\circ$ ,  $-90^\circ$ ,  $-67.5^\circ$ ,  $-45^\circ$ ,  $45^\circ$ ,  $67.5^\circ$ ,  $90^\circ$ ,  $112.5^\circ$ ,  $135^\circ$ ). For  $S_1$ , the adaptor was always directed outward with respect to the fovea, and for all other subjects, it was directed inward. Note that only for  $S_1$  and  $S_2$  were all directions used.  $S_3$  and  $S_4$  performed a subset consisting of  $-67.5^\circ$  and  $67.5^\circ$ .

To test whether the encoded direction does not change for the same-condition, as hypothesized, we conducted an additional “baseline condition.” This was done because the observed SMAEs of the same-condition often deviate from a veridical representation of the adaptor RDK’s direction. The baseline was measured in separate sessions. Instead of detecting

a luminance change, subjects had to simply indicate the direction of the target RDK. The general procedure was the same as described above with the exception that the central presentation of the target direction and the 3 min break were omitted. Furthermore, instead of lasting for  $\sim 1$  min as in both the same- and different-conditions, a trial was finished as soon as the subjects responded, usually after a few seconds.

In total, the whole experiment lasted over several weeks until for each subject 15 measures were available for each direction in all types of trials. Altogether,  $>50$  h of measurement were required. Approximately 22% of all the trials in the same- and different-conditions were excluded because the detection performance of the subjects fell outside the predefined range. The performance of the subjects is shown in supplemental Table S2 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). From the remaining trials, a further 10% had to be excluded because subjects indicated that they did not perceive an aftereffect. Finally,  $<0.01\%$  of all these trials were identified and excluded as outliers, that is, the indicated direction fell outside the region of  $\pm 2.5$  SDs around the respective mean. The results are summarized in Figure 2 and supplemental Table S3 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material).

**Control experiment 1: contrast dependency.** In this control experiment, we tested the influence of contrast on the observed effect using a high-



**Figure 3.** Contrast dependency of the SMAE. For both subjects the adaptor ( $-67.5^\circ, 67.5^\circ$ ) was always directed inward with respect to the fovea. For the high-contrast condition, mean SMAEs of different-conditions are indicated by red triangles and mean SMAEs of same-conditions are indicated by blue triangles. For the low-contrast condition, mean SMAEs of different-conditions are indicated by red squares and mean SMAEs of same-conditions are indicated by blue squares. Dashed blue lines indicate the unsigned adaptor direction, and dashed red lines indicate the attended direction. Error bars denote 95% confidence intervals. For both subjects, differences between the same-condition and the different-condition are statistically significant in the high- and low-contrast conditions as indicated by asterisks. The low-contrast condition results in a significantly stronger attraction than the high-contrast condition.

contrast and a low-contrast condition. That is, we wondered whether the strength of the observed effect depends on stimulus contrast, as is suggested by electrophysiological experiments (Reynolds et al., 2000; Martínez-Trujillo and Treue, 2002). In these studies it was reported that the influence of spatial attention on the gain of single cells diminishes for high-contrast stimuli. If a related mechanism is working in the feature-based domain as well, we predict less pronounced distortions of the population response for high-contrast than for low-contrast stimuli and thus a reduced difference of the perceived aftereffect between the same- and different-conditions.

Subjects  $S_2$  and  $S_3$  participated in this control experiment. The general procedure and the task were identical to the main experiment, but here, both stimuli had a coherence of 100%. This was done to exclusively test the influence of stimulus contrast on the observed effect. Given equal coherence, a low-contrast stimulus should evoke weaker responses than a high-contrast stimulus. In the high-contrast condition, the dots of the target and the adaptor were white with a luminance of  $85.6 \text{ cd/m}^2$  on a black background ( $2.2 \text{ cd/m}^2$ ), resulting in a Weber contrast of 37.91 for a single dot and for the whole RDK in a root mean square (RMS) contrast (Moulden et al., 1990; Seitz et al., 2006; Martínez-Trujillo and Treue, 2002) of 35.87. In the low-contrast condition, the dots of the target were also white with the same luminance as stated above, but the dots of the adaptor were gray with a luminance of  $5.2 \text{ cd/m}^2$ , resulting in a Weber contrast of 1.36 and a RMS contrast of 1.29. Two directions ( $-67.5^\circ, 67.5^\circ$ ) were used. Information about the size of the luminance change in the detection task is given in supplemental Table S4 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). Approximately 20% of all the trials were excluded because the luminance detection performance of the subjects fell outside the predefined range (see above, Main experiment). The performance of the subjects is shown in supplemental Table S5 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). Less than 0.02% of the remaining trials were identified and excluded as outliers, that is, the indicated direction fell outside the region of  $\pm 2.5$  SDs around the respective mean. The results are summarized in Figure 3 and supplemental Table S6 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material).

**Control experiment 2: unattended motion.** In this control experiment, we tested whether there is still a difference between the same-condition and the different-condition when attention is withdrawn from the target. This was done to exclude an interpretation of our results that is based on a simple weighting of both stimuli as would be predicted by the biased competition framework (Desimone and Duncan, 1995) if both stimuli were processed by the same population of neurons.

Subjects  $S_2$ ,  $S_3$ , and  $S_4$  participated in this control. Both the general procedure and the stimuli were identical to the main experiment. The

initial presentation of the to be attended target direction was replaced by a blank of the same duration (5 s). Instead of detecting a luminance change of the target, subjects looked for a luminance change of the central fixation point while ignoring both RDKs. The fixation point was white with a luminance of  $85.6 \text{ cd/m}^2$ . The magnitude of the luminance change required central fixation, that is, it was assured that detection performance was at chance level if the subject fixated one of the two RDKs. The luminance change was  $8.4 \text{ cd/m}^2$  for  $S_2$  and  $S_3$ , and  $12.7 \text{ cd/m}^2$  for  $S_4$ . The performance of the subjects is shown in supplemental Table S7 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). Approximately 5% of all the trials were excluded because subjects indicated that they did not perceive an aftereffect. No outliers were detected. The results are summarized in Figure 4 and supplemental Table S8 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material).

**Statistical analysis.** Since the variability of the data is reasonably low ( $SD = 5.57$ ,  $SD_{SD} = 3.19$ ,  $min = 1.27$ ,  $max = 18.48$ ), ordinary statistics were used. It is assumed that all measurements are statistically independent. For all comparisons, Welch's generalization of the independent two-sample  $t$  test for unequal variances was used. All tests were conducted two sided. The test-wise  $\alpha'$  level was adjusted due to Bonferroni correction so that  $\alpha' = \alpha/m$ , where  $m$  is the number of comparisons and  $\alpha = 0.05$ . The global null hypothesis is rejected if there exists at least one out of  $m$  multiple comparisons that is significant at the adjusted  $\alpha'$  level.

In the main experiment, we tested separately for all subjects whether there was a statistically significant difference in the indicated SMAE directions between the same-condition and the different-condition and whether there was a difference between the same-condition and the baseline. This was done for all directions, consisting in each case of 10 comparisons for  $S_1$  and  $S_2$  ( $\alpha' = 0.005$ ) and 2 comparisons for  $S_3$  and  $S_4$  ( $\alpha' = 0.025$ ). Statistics are summarized in supplemental Table S3 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material).

In control experiment 1, tests were conducted to determine whether there was a difference between the same-condition and the different-condition for the high- and low-contrast conditions or between the different-conditions of the high- and the low-contrast conditions. In each case, two comparisons were made per subject. Statistics are summarized in supplemental Table S6 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material).

In control experiment 2, we tested whether there was a difference between the indicated SMAE directions of the same-condition and the different-condition. For each subject, two comparisons were made. Statistics are summarized in supplemental Table S8 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material).

**Model.** To explain the observed results, a "simple model" was set up in which the impact of attention on the population response was simulated. The model can be considered as a static, simplified version of earlier formalizations (Hamker, 2005, 2007). Thereby, the influences of attention were simulated in explicit terms. In general, feature-based attention is assumed to modulate the gain of individual neurons via feature specific feedback. This feedback is likely a combination of several signals originating at several stages in the cortical hierarchy ranging from high-level prefrontal to more low-level areas (Hamker, 2004).

Let  $F = [0, 2\pi)$  be the feature space. Tuning curves are described by Gaussian functions located in  $F$ . For a given stimulus, the unmodulated or input activity  $r_i^{\text{in}}$  of the  $i$ th cell is as follows:

$$r_i^{\text{in}} = b_0 + b_1 \exp\left(-\frac{\phi_i^2}{2(\sigma_{\text{TC}})^2}\right)$$

$$\phi_i = ((p_s - c_i^{\text{TC}} + \pi) \bmod 2\pi) - \pi,$$

with  $\phi \in [-\pi, \pi)$ . The preferred direction, that is, the center of the tuning curve in feature space, is denoted by  $c_i^{\text{TC}} \in F$ . The direction of the stimulus is denoted by  $p_s \in F$ . The width of the tuning curve is determined by  $\sigma_{\text{TC}} = 0.52$ ,  $b_0 = 1$  is the baseline activity, and  $b_1 = 10$  determines the height of the Gaussian. Note that cells have a bandwidth of  $68^\circ$ , which is the average bandwidth of motion selective cells as measured in V1 (Albright, 1984), as determined by the half-maximum width after subtraction of the baseline activity  $b_0$ .

The influence of attention on the population response is formalized as a difference of Gaussians, as follows:

$$a_i = \exp\left(-\frac{\phi_i^2}{2(\sigma_A)^2}\right) - c \exp\left(-\frac{\phi_i^2}{2(3\sigma_A)^2}\right)$$

$$\phi_i = ((p_i - c_A + \pi) \bmod 2\pi) - \pi,$$

where  $c_A \in F$  denotes the attended direction and  $p_i$  is the position of the  $i$ th cell in  $F$ . The extend of the central excitatory and the surrounding inhibitory region is determined by the constant  $c$  and by  $\sigma_A$ , which was set to  $\sigma_A = 0.52$  for  $S_1$  and  $\sigma_A = 0.48$  for  $S_2$ . Note that  $c = 0.9$  for subject  $S_1$  and  $c = 0.8$  for subject  $S_2$ . Further note that a simple Gaussian profile with central excitation and surrounding inhibition is inconsistent with the observed results, since it can account for attraction but not for the repulsion effect (see supplemental Fig. S2, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material).

Finally, the modulated response  $r_i^{\text{gain}}$  of a given cell  $i$  is as follows:

$$r_i^{\text{gain}} = r_i^{\text{in}} (1 + w a_i),$$

where  $w$  is a weight that was determined by the data ( $w = 4$  for subject  $S_1$  and  $w = 2.5$  for subject  $S_2$ ). Note that if  $r_i^{\text{gain}} < 0$ , then it was set to  $r_i^{\text{gain}} = 0$ . For the simulations, the whole population consisted of 360 cells that were equally spaced in  $F$ .

With the above-described model, we are now able to simulate a population response for a given direction of motion for both the modulated and unmodulated conditions. To decode the direction represented by the population response, we computed the population vector of the neural activity in  $F$  as follows:

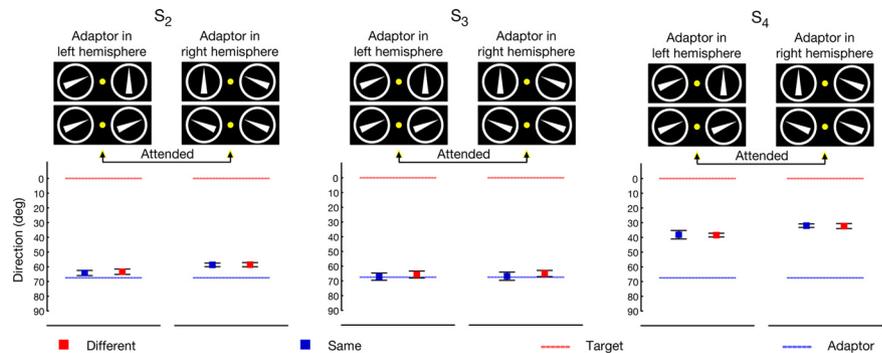
$$x = \sum_i r_i^{\text{gain}} \cos(c_i^{\text{TC}})$$

$$y = \sum_i r_i^{\text{gain}} \sin(c_i^{\text{TC}}).$$

The decoded direction is then given in Cartesian coordinates by the vector  $\mathbf{v} = (x, y)$ . For an overview of decoding methods, see, for example, Seung and Sompolinsky (1993) and Salinas and Abbott (1994).

We obtained the parameters by fitting the model to the data of subjects  $S_1$  and  $S_2$  to provide a qualitative explanation of the observed effect. Therefore, fitting was done initially by hand, followed by a more sophisticated exploration of a restricted part of the parameter space. Furthermore, to keep the description as simple as possible, we did not account for the apparent asymmetries of the data in the two hemispheres (see Fig. 2A).

For the illustration of the tuning curve shifts, a second population of cells was simulated. These cells pool the activity of the first population in feature space, that is, they are driven by a simple weighted sum of afferent responses (Rust et al., 2006). The bandwidth of the tuning curves of the second population is  $83^\circ$ , which is the average bandwidth of motion selective cells as measured in MT (Albright, 1984). The modulated ( $w > 0$ ) and unmodulated ( $w = 0$ ) tuning curves were normalized after the subtraction of the baseline, which was in this case defined as the minimal activity of a given tuning curve. Processing capacity is defined as the number of tuning curves  $n_{\text{TC}}$  for which the response to a given



**Figure 4.** Results of the unattended motion condition. Subjects attended to the fixation point as indicated by the yellow color. For all subjects, the adaptor ( $-67.5^\circ, 67.5^\circ$ ) was always directed inward with respect to the fovea. Mean SMAEs of different-conditions are indicated by red squares. Mean SMAEs of same-conditions are indicated by blue squares. Dashed blue lines indicate the unattended adaptor direction and dashed red lines indicate the attended direction. Error bars denote 95% confidence intervals. For all subjects, none of the differences between the same-condition and the different-condition are statistically significant.

direction exceeds half of the maximum activity of the respective cell [please refer also to Hamker et al. (2008) for the concept of processing capacity] and  $\Delta n_{\text{TC}}$  is the change in capacity in the modulated case with respect to the unmodulated one. For the reported simulations (Fig. 5F), the parameters of  $S_2$  were used. Note that the same is true for the reported distortions of the population response (Fig. 5C,D).

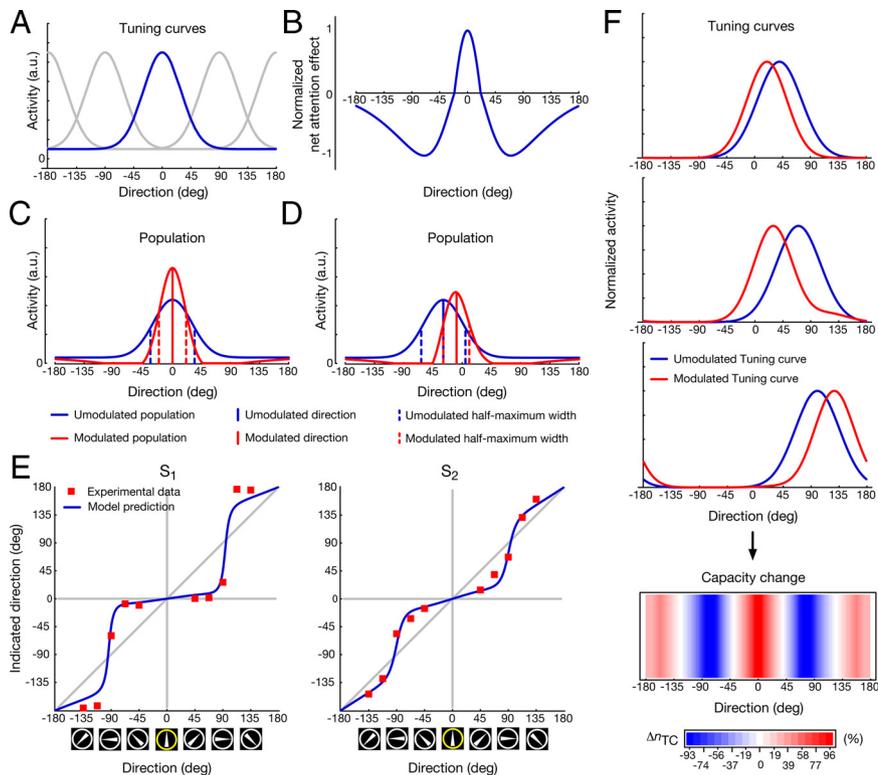
Finally, as stated in the Introduction, feature-based gain modulations have been also reported during remembering previously attended features (Chelazzi et al., 1993, 1998). Thus, the model was extended to simulate additional gain modulations that might occur during the perception of the SMAE after the adaptation and thus would potentially be able to influence the perceived SMAE direction (see supplemental material, section 1.2, available at [www.jneurosci.org](http://www.jneurosci.org)). While this does not change the general assumption that feature-based attention is able to alter the represented direction of a motion stimulus already during the encoding, it has implications for the assumed gain profile. That is, simulating an additional gain modulation allows the “extended model” to be fitted to the data using a simple Gaussian modulatory signal consisting of central excitation and lateral inhibition as reported by Martinez-Trujillo and Treue (2004).

## Results

We observe that attending to a peripherally presented direction of motion dynamically alters the direction of the SMAE, indicative of a change in the representation of the unattended stimulus in feature space. For a given adaptor, the perceived direction of the SMAE can deviate  $>30^\circ$  from the expected direction. Figure 2 summarizes our main results. Note that the indicated SMAEs are transformed to the direction of motion that would have caused them, that is, the stimulus representation in the adaptation phase, assuming a simple reversal (Fig. 1A).

As is apparent in Figure 2, SMAEs of the same-condition often deviate from veridical but are often close to the baseline, that is, the perceived direction of the target when no additional task was required. When the motion direction of the adaptor differs from the one of the target (different-condition), however, attention to the target changes the encoded direction for the adaptor. Take subject  $S_1$  for example (Fig. 2A): when the adaptor motion is to the left ( $-90^\circ$ ), attending a target with leftward motion indicates that the encoded adaptor motion is also leftward. The encoded motion of the same physical stimulus, however, drastically changes to approximately  $-60^\circ$  when attention is directed to a target with upward motion.

Two patterns are observed, attraction and repulsion. For adaptor directions differing up to  $90^\circ$  from the target direction, the resulting SMAEs seem to arise from adaptor directions being



**Figure 5.** Simple model. **A**, Tuning curves of sensory neurons located on a circle in arbitrary units (a.u.). The bandwidth of the tuning curves is equal to the average bandwidth of motion selective cells measured in V1 (Albright, 1984). **B**, The net effect of attention. Directions close to the attended direction ( $0^\circ$ ) are enhanced, while directions farther away are suppressed. **C**, Consistent with experimental data (Martínez-Trujillo and Treue, 2004), the population response is sharpened as indicated by the half-maximum response when the stimulus direction is equal to the attended direction. **D**, For a stimulus with a direction of  $-30^\circ$ , the population response is distorted, that is, the population vector of the neural activity is attracted toward the attended direction. **E**, Indicated SMAEs of subject  $S_1$  and  $S_2$  for 10 directions ( $-135^\circ$ ,  $-112.5^\circ$ ,  $-90^\circ$ ,  $-67.5^\circ$ ,  $-45^\circ$ ,  $45^\circ$ ,  $67.5^\circ$ ,  $90^\circ$ ,  $112.5^\circ$ ,  $135^\circ$ ) denoted by the red squares shown together with the model fits (blue line) as a function of the veridical direction. Note again that SMAEs are shown as the direction that would have caused them, assuming a direct reversal. While directions near the attended direction ( $0^\circ$ ) are attracted, directions farther away are repelled. **F**, Exemplary tuning curve shifts of three model neurons. The center of the unmodulated tuning curve is  $40^\circ$ ,  $70^\circ$ , and  $100^\circ$  from top to bottom. While the center of the tuning curve shifts toward the attended feature for tuning curves with unmodulated centers close to the attended direction, it is shifted away for tuning curves with unmodulated centers farther away. This results in an increase of tuning curves that are processing the attended and the opposite direction. However, since the tuning curves with preferred direction farther away from the attended direction are suppressed, an effective increase in processing occurs at the attended direction, that is, tuning curves of neurons shift their preferred direction such that the whole population preferentially processes the attended properties.

closer to the target than they actually are, that is, they are attracted. This can be seen for the  $-90^\circ$ ,  $-45^\circ$ ,  $45^\circ$ , and  $90^\circ$  adaptor directions in Figure 2A and for both adaptor directions ( $-67.5^\circ$  and  $67.5^\circ$ ) in Figure 2B. For adaptor directions farther away from the target, we find the opposite, that is, the adaptor directions are repelled with respect to the target, as can be seen for the  $-135^\circ$  and the  $135^\circ$  adaptor direction in Figure 2A. Thus, the distance in feature space between attracted and repelled directions is increased.

Although there seems to be a trend of the measured baselines to lie below the SMAEs of the same-condition (Fig. 2B), they do not reach statistical significance (see Material and Methods for statistical details), with the exception of subject  $S_4$ . The significant deviations for  $S_4$  suggest that feature-based attention leads to a change of the encoded direction dependent on the actual neural representation rather than the direction that is inferred by the physical properties of the stimuli. Given the physical properties, one would not expect a difference between the baseline and the SMAEs of the same-condition. However, for  $S_4$  the encoded direction does largely deviate from veridical if both motion stimuli

are unattended (Fig. 4). In the attended cases, they seem attracted toward the perceived direction of the attended target and are thus represented closer to veridical. Regardless, for all subjects all SMAEs of the different-condition do significantly deviate from their respective SMAEs of the same-condition.

It seems unlikely that a response bias can account for the observed pattern of attraction and repulsion. If there was a simple tendency of subjects to indicate the direction of the SMAEs as being closer to the attended target, we should not have found repulsion. Moreover, the magnitude of the effect inversely scales with the contrast of the adaptor as revealed by a control experiment in which we used high-contrast and low-contrast adaptors with 100% coherence (see Materials and Methods). For both the high- and the low-contrast conditions, we find significant deviations, but the magnitude in the high-contrast condition is significantly smaller than the one in the low-contrast condition (see Fig. 3 and supplemental Table S6, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). Although not yet tested in the feature-based domain, we interpret this contrast dependency as support for an attentional explanation since the magnitude of attentional effects has been previously reported to rely on stimulus contrast in the space-based domain (Reynolds et al., 2000; Martínez-Trujillo and Treue, 2002). Thus, if attention in both domains is characterized by related mechanisms at the single-cell level, we predict that gain modulations induced by feature-based attention will diminish for high-contrast stimuli.

Another potential caveat for an explanation based on feature-based attention would arise if both target and adaptor were simultaneously processed by a single population of cells with large receptive fields covering both stimuli. If this were true, attention would simply increase the contribution of the attended stimulus to the net SMAE (Alais and Blake, 1999) as known from the biased competition framework (Desimone and Duncan, 1995). While this again cannot explain the repulsion effect, it could be a possible explanation for the observed attraction. Thus, a second control experiment was conducted in which subjects had to detect a recurring luminance change of the fixation point and to ignore both target and adaptor (see Materials and Methods). Given the assumption that target and distractor are both processed simultaneously by the same cells, the biased competition framework predicts that SMAEs in the different-condition lie somewhere in between the SMAEs of the target and the adaptor. However, this is not what we observe. Consistent with a feature-based attention explanation, the data show that when subjects ignore the target, the effect disappears and no significant differences between the same- and the different-condition are observed (see Fig. 4 and supplemental Table S8, available at [www.jneurosci.org](http://www.jneurosci.org) as supple-

mental material). Thus, we conclude that both stimuli are not simultaneously processed by a single neural population with large receptive fields.

Finally, our results are also unlikely to result from eye movements as revealed by the analysis of the detection performance. The most reasonable case of breaking fixation for subjects would be to fixate the target to maximize the detection of the luminance change. However, the subjects performed much better in the experiments in which they were allowed to fixate the target than in the experiments in which they were required to keep central fixation (see supplemental material, section 1.3, available at [www.jneurosci.org](http://www.jneurosci.org)). Consistent with this, a subsequently conducted control in which eye movements were recorded revealed an average accuracy of  $0.53^\circ$  relative to the fixation point (see supplemental material, section 1.3, available at [www.jneurosci.org](http://www.jneurosci.org)).

Instead, our results can be explained by a class of models in which properties of the attended stimulus are fed back to early sensory areas and induce modulations of the population response. Initially we will focus exclusively on gain modulations during the adaptation phase (Fig. 5), which will be referred to as the “simple model.” While such attentional gain modulations relate to earlier models of feature-based attention (Hamker, 2006; Ardid et al., 2007; Reynolds and Heeger, 2009), feature-based distortions of the population responses have not been explicitly predicted, except for a study simulating a masking phenomenon known as feature inheritance (Hamker, 2007). The general idea of this model is that the gain resulting from feature-based attention is not uniform across the whole population, but rather peaks around the attended feature, which in turn distorts the population response such that it is effectively shifted toward the attended feature (Hamker, 2007). Our present model evaluation revealed that the net gain modulation ought to have a center-surround profile that follows a difference of Gaussians, that is, neurons that prefer directions close to the attended target are enhanced while neurons that prefer directions farther away are suppressed. Thus the strongest suppression is found for neurons preferring intermediate directions with respect to the attended one (Fig. 5B). Consistent with electrophysiological recordings (Martinez-Trujillo and Treue, 2004), this results in a sharpening of the population response if target and adaptor have the same direction (Fig. 5C) as determined by the half-maximum width of the response profile (see Material and Methods). For adaptors close to the target, however, the gain modulation leads to a distortion of the population in such a way that the population vector is shifted toward the target direction (Fig. 5D). Similarly, the vector of the population response to adaptors that lie in the suppressive surround is shifted away from the target direction. Altogether the distortions lead to local magnifications in feature space consistent with the response of the subjects (Fig. 5E). While these effects already occur due to simple gain changes of single neurons without shifts in the tuning curve, neurons that compute a weighted sum of the distorted population response ought to also alter their tuning properties, as predicted by the model (Fig. 5F). Tuning curves that are located in feature space close to the center of attention are attracted, whereas others located farther away are repelled. However, due to the suppressive effects surrounding the center of attention an effective increase in processing occurs at the attended direction. Thus, the whole population preferably processes the attended feature. Furthermore, such an increase in the number of neurons that are preferentially processing a common part of feature space has been theoretically shown to enhance coding accuracy (Abbott and Dayan, 1999; Shamir and Sompolinsky, 2006).

So far we have only considered the possibility of gain modulations during the adaptation phase, that is, the period during which both the adaptor and target RDK were presented. However, additional gain modulations could take place during the period during which the static RDK is shown and thus have a direct influence on the SMAE. Such modulations might be caused by the imagined target (Schlack and Albright, 2007) or in more general terms by a persistent representation of an already vanished stimulus (Chelazzi et al., 1993, 1998). Thus, we extended the simple model to simulate this possibility, that is, two subsequent gain modulations, one from feature-based attention during the adaptation phase and the other in the SMAE phase caused by the remembered target. Note that this “extended model” required us to explicitly simulate the static motion aftereffect (see supplemental material, section 1.2, available at [www.jneurosci.org](http://www.jneurosci.org)). As can be seen in supplemental Figure S4A (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material), this model is consistent with the experimental observations as well. Furthermore, in line with Martinez-Trujillo and Treue (2004), the fit was achieved by assuming a simple Gaussian profile for both signals (supplemental Fig. S4E, G, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material) instead of a difference of Gaussians as required for the simple model. However, the resulting gain profile is substantially smaller than expected from the data of Martinez-Trujillo and Treue (2004) (see supplemental material, section 1.4, available at [www.jneurosci.org](http://www.jneurosci.org)), which might be due to task differences as described in the Discussion. Furthermore, this model predicts a strong collapse of feature space during the encoding phase (supplemental Fig. S4B, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). While this prediction might be due to the specific assumptions made, for instance, the way we simulated the SMAE in the light of insufficient electrophysiological constraints, the collapse of feature space can be avoided with the extended model by assuming a difference-of-Gaussian profile for both modulatory signals (supplemental Fig. S4D, F, H, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). Thus, given the present uncertainties in the model assumptions, the extended explanation of our observations can be considered as an alternative interpretation.

## Discussion

Attention is well known to integrate features, allowing for a coherent percept (Lu and Sperling, 1995; Treisman, 2006). Our study further reveals that in such an integration process, represented features are not fixed entities of the visual input, but rather dependent on the attentional state. Feature-based attention alters the neural representation of feature space as indicated by a change in direction of the SMAE dependent on the attended feature in the adaptation phase. These observations conceptually extend the feature-similarity gain model of attention (Treue and Martinez Trujillo, 1999). At the level of single cells, attention does still result in a simple scaling of the neural activity, which according to the feature-similarity model depends on a gain factor as a function of the similarity between the attended feature and the preferred feature of the particular neuron. However, since these changes in gain are not uniform, the modulations lead to distortions at the level of the population response, which changes the encoded feature. Thus, while previous experiments have only reported an upregulation or downregulation of the neural population response, which simply alters the strength of the encoded entity (Treue and Martínez Trujillo, 1999; Saenz et al., 2002; Liu et al., 2007), our observations suggest modulations of the population response governed by the feature-similarity principle, that

in addition enable the visual system to dynamically separate relevant from irrelevant features, that is, to increase the distance of the respective features in feature space. Furthermore, distorted population responses lead to tuning curve changes in neurons that are driven by these responses, for example, neurons that compute a weighted sum of the population response (Rust et al., 2006). This could explain the change in the spectral tuning of V4 neurons during a naturalistic visual search task, where cells became more sensitive to features of the attended search target (David et al., 2008). Thus, our model links physiological observations with perceptual experience.

As far as the origin of the distorted percept is concerned, it seems that a widely distributed network is involved in generating motion aftereffects (Taylor et al., 2000). However, as for the SMAE, both psychophysical (Moulden, 1980) and electrophysiological (Kohn and Movshon, 2003) evidence exist that it is already generated at the level of V1 (for review, see Mather et al., 2008); thus, distortions of the population response during the direction of feature-based attention should also be observable at this early processing stage. As a consequence, this early impact of feature-based attention will influence the representation of stimuli already before the final decision level has been reached. Although the distinction between encoding and decoding might be smoother than traditionally thought, earlier studies that focused on the discrimination of features attributed the misperception of stimuli to the level of decoding (Jazayeri and Movshon, 2007), for example, by optimally chosen top-down signals (Navalpakkam and Itti, 2007).

How does our work relate to the findings of Martinez-Trujillo and Treue (2004)? As already mentioned, their observed net effect can be described by a central excitation and lateral inhibition as well. However, their data suggest that lateral inhibition is strongest opposite to the attended direction, while according to our simple model, the peak inhibition is located much nearer to the attended direction in feature space and then gradually decreases moving farther away. While the extended model allowed us to explain our findings using a simple Gaussian signal, the estimated profile as constrained by our psychophysical data is substantially smaller than the estimated profile constrained by the electrophysiological measurements [see supplemental material, section 1.4 (available at [www.jneurosci.org](http://www.jneurosci.org)), for a fit of our model to the data of Martinez-Trujillo and Treue (2004)]. As outlined above, one explanation regarding these differences might be that our experiments presumably probe feature-based effects in V1, whereas Martinez-Trujillo and Treue (2004) recorded in area MT. Furthermore, the shape of the net attention effect does not necessarily have to be static but might change depending on the task at hand (Navalpakkam and Itti, 2007; Scolari and Serences, 2009). While subjects in our study were instructed to attend to the direction of the target RDK as accurate as possible, monkeys in the study of Martinez-Trujillo and Treue (2004) had to detect either a velocity change or a direction change of the attended stimulus. In the latter case, it could well be possible that the attentional system boosted neurons that were tuned away from the attended feature to facilitate the detection of another direction. Thus, the area of excitation in feature space should be rather broad as reported by Martinez-Trujillo and Treue (2004).

To summarize, according to the proposed mechanisms, the simple and the extended models are qualitatively equivalent in explaining our psychophysical findings of an altered feature space. While both models assume a major impact of feature-based attention in the adaptation phase, where the stimulus is

encoded, the extended model also predicts an influence of the remembered target in the SMAE phase, that is, a direct impact on our conscious percept. Thus, although it seems evident that top-down feature-based signals underlie our findings, the exact stage at which these signals act is a question for future research.

How do our results relate to findings in the spatial attention domain? Hamker et al. (2008) have demonstrated, using a computational model, that spatial attention directed to the saccade target explains the perisaccadic misperception of briefly flashed stimuli, known as compression of visual space (Morrone et al., 1997; Ross et al., 1997). In that model, misperception arises because spatial attention distorts the population response in the spatial domain, which in turn dynamically alters the receptive field profile. As a result, more cells are effectively processing the attended location. Indeed, a number of studies have indicated attentional changes of the receptive field structure in covert (Connor et al., 1996, 1997; Womelsdorf et al., 2006) and overt (Tolias et al., 2001) shifts of attention. Furthermore, although far less in magnitude, some perceptual effects have also been observed during covert attention shifts, such as repulsion (Suzuki and Cavanagh, 1997; Pratt and Turk-Browne, 2003) and attraction (Yamada et al., 2008). Similar to the effects reported in the present study, these different observations might be accounted for by a center-surround profile of the gain modulation (Hopf et al., 2006) as well. Thus, it seems that related mechanisms are involved in the feature and spatial domain of attention.

To conclude, attention has typically been characterized as a mere control of information flow, whereas the observations of perceptual distortions point toward an active recruitment of processing resources. As a result, a stimulus matching the content of attention is encoded more similarly to the attended one, whereas other stimuli can be repelled in feature space depending on the type of net modulatory profile as determined by the interactions at the systems level.

## References

- Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11: 91–101.
- Alais D, Blake R (1999) Neural strength of visual attention gauged by motion adaptation. *Nat Neurosci* 2:1015–1018.
- Albright TD (1984) Direction and orientation selectivity of neurons in visual area MT of the macaque. *J Neurophysiol* 52:1106–1130.
- Ardid S, Wang XJ, Compte A (2007) An integrated microcircuit model of attentional processing in the neocortex. *J Neurosci* 27:8486–8495.
- Bichot NP, Rossi AF, Desimone R (2005) Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* 308:529–534.
- Boynton GM, Ciaramitaro VM, Arman AC (2006) Effects of feature-based attention on the motion aftereffect at remote locations. *Vision Res* 46:2968–2976.
- Britten KH, Newsome WT (1998) Tuning bandwidths for near-threshold stimuli in area MT. *J Neurophysiol* 80:762–770.
- Chelazzi L, Miller EK, Duncan J, Desimone R (1993) A neural basis for visual search in inferior temporal cortex. *Nature* 363:345–347.
- Chelazzi L, Duncan J, Miller EK, Desimone R (1998) Responses of neurons in inferior temporal cortex during memory-guided visual search. *J Neurophysiol* 80:2918–2940.
- Connor CE, Gallant JL, Preddie DC, Van Essen DC (1996) Responses in area V4 depend on the spatial relationship between stimulus and attention. *J Neurophysiol* 75:1306–1308.
- Connor CE, Preddie DC, Gallant JL, Van Essen DC (1997) Spatial attention effects in macaque area V4. *J Neurosci* 17:3201–3214.
- David SV, Hayden BY, Mazer JA, Gallant JL (2008) Attention to stimulus features shifts spectral tuning of V4 neurons during natural vision. *Neuron* 59:509–521.
- Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18:193–222.

- Hamker FH (2004) A dynamic model of how feature cues guide spatial attention. *Vis Res* 44:501–521.
- Hamker FH (2005) The reentry hypothesis: the putative interaction of the frontal eye field, ventrolateral prefrontal cortex, and areas V4, IT for attention and eye movement. *Cereb Cortex* 15:431–447.
- Hamker FH (2006) Modeling feature-based attention as an active top-down inference process. *BioSystems* 86:91–99.
- Hamker FH (2007) The mechanisms of feature inheritance as predicted by a systems-level model of visual attention and decision making. *Adv Cogn Psychol* 3:111–123.
- Hamker FH, Zirnsak M, Calow D, Lappe M (2008) The peri-saccadic perception of objects and space. *PLoS Comput Biol* 4:e31.
- Hopf JM, Boehler CN, Luck SJ, Tsotsos JK, Heinze HJ, Schoenfeld MA (2006) Direct neurophysiological evidence for spatial suppression surrounding the focus of attention in vision. *Proc Natl Acad Sci U S A* 103:1053–1058.
- Jazayeri M, Movshon JA (2007) A new perceptual illusion reveals mechanisms of sensory decoding. *Nature* 446:912–915.
- Kohn A, Movshon JA (2003) Neuronal adaptation to visual motion in area MT of the macaque. *Neuron* 39:681–691.
- Liu T, Larsson J, Carrasco M (2007) Feature-based attention modulates orientation-selective responses in human visual cortex. *Neuron* 55:313–323.
- Lu Z-L, Sperling G (1995) Attention-generated apparent motion. *Nature* 377:237–239.
- Martinez-Trujillo JC, Treue S (2002) Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron* 35:365–370.
- Martinez-Trujillo JC, Treue S (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr Biol* 14:744–751.
- Mather G, Pavan A, Campana G, Casco C (2008) The motion aftereffect reloaded. *Trends Cogn Sci* 12:481–487.
- Morrone MC, Ross J, Burr DC (1997) Apparent position of visual targets during real and simulated saccadic eye movements. *J Neurosci* 17:7941–7953.
- Motter BC (1994a) Neural correlates of attentive selection for color or luminance in extrastriate area V4. *J Neurosci* 14:2178–2189.
- Motter BC (1994b) Neural correlates of feature selective memory and pop-out in extrastriate area V4. *J Neurosci* 14:2190–2199.
- Moulden B (1980) After-effects and the integration of patterns of neural activity within a channel. *Philos Trans R Soc Lond B Biol Sci* 290:39–55.
- Moulden B, Kingdom F, Gatley LF (1990) The standard deviation of luminance as a metric for contrast in random-dot images. *Perception* 19:79–101.
- Navalpakkam V, Itti L (2007) Search goal tunes features optimally. *Neuron* 53:605–617.
- Pratt J, Turk-Browne NB (2003) The attentional repulsion effect in perception and action. *Exp Brain Res* 152:376–382.
- Reynolds JH, Chelazzi L (2004) Attentional modulation of visual processing. *Annu Rev Neurosci* 27:611–647.
- Reynolds JH, Heeger DJ (2009) The normalization model of attention. *Neuron* 61:168–185.
- Reynolds JH, Pasternak T, Desimone R (2000) Attention increases sensitivity of V4 neurons. *Neuron* 26:703–714.
- Ross J, Morrone MC, Burr DC (1997) Compression of visual space before saccades. *Nature* 386:598–601.
- Rust NC, Mante V, Simoncelli EP, Movshon JA (2006) How MT cells analyze the motion of visual patterns. *Nat Neurosci* 9:1421–1431.
- Saenz M, Buracas GT, Boynton GM (2002) Global effects of feature-based attention in human visual cortex. *Nat Neurosci* 5:631–632.
- Saenz M, Buračas GT, Boynton GM (2003) Global feature-based attention for motion and color. *Vision Res* 43:629–637.
- Salinas E, Abbott LF (1994) Vector reconstructing from firing rates. *J Comput Neurosci* 1:89–107.
- Schlack A, Albright TD (2007) Remembering visual motion: neural correlates of associative plasticity and motion recall in cortical area MT. *Neuron* 53:881–890.
- Scolari M, Serences JT (2009) Adaptive allocation of attentional gain. *J Neurosci* 29:11933–11942.
- Seitz AR, Nanez JE Sr, Holloway SR, Watanabe T (2006) Perceptual learning of motion leads to faster flicker perception. *PLoS One* 1:e28.
- Seung HS, Sompolinsky H (1993) Simple models for reading neuronal population codes. *Proc Natl Acad Sci U S A* 90:10749–10753.
- Shamir M, Sompolinsky H (2006) Implications of neuronal diversity on population coding. *Neural Comput* 18: 1951–1986.
- Suzuki S, Cavanagh P (1997) Focussed attention distorts visual space: an attentional repulsion effect. *J Exp Psychol Hum Percept Perform* 23:443–463.
- Taylor JG, Schmitz N, Ziemons K, Grosse-Ruyken ML, Gruber O, Mueller-Gaertner HW, Shah NJ (2000) The network of brain areas involved in the motion aftereffect. *Neuroimage* 11:257–270.
- Tolias AS, Moore T, Smirnakis SM, Tehovnik EJ, Siapas AG, Schiller PH (2001) Eye movements modulate visual receptive fields of V4 neurons. *Neuron* 29:757–767.
- Treisman A (2006) How the deployment of attention determines what we see. *Vis Cogn* 14:411–443.
- Treue S, Martínez Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399:575–579.
- 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.
- Yamada Y, Kawabe T, Miura K (2008) Mislocalization of a target toward subjective contours: attentional modulation of location signals. *Psychol Res* 72:273–280.