

Effects of feature-based attention on the motion aftereffect at remote locations

A. Cyrus Arman^{a,c}, Vivian M. Ciaramitaro^{a,b}, Geoffrey M. Boynton^{a,*}

^a *The Salk Institute, 10010 N. Torrey Pines Road, La Jolla, CA 92037-1099, USA*

^b *Department of Psychology, University of California, San Diego, USA*

^c *University of Southern California, Neuroscience Graduate Program, USA*

Received 16 October 2005; received in revised form 28 February 2006

Abstract

Previous studies have shown that attention to a particular stimulus feature, such as direction of motion or color, enhances neuronal responses to unattended stimuli sharing that feature. We studied this effect psychophysically by measuring the strength of the motion aftereffect (MAE) induced by an unattended stimulus when attention was directed to one of two overlapping fields of moving dots in a different spatial location. When attention was directed to the same direction of motion as the unattended stimulus, the unattended stimulus induced a stronger MAE than when attention was directed to the opposite direction. Also, when the unattended location contained either uncorrelated motion or had no stimulus at all an MAE was induced in the opposite direction to the attended direction of motion. The strength of the MAE was similar regardless of whether subjects attended to the speed or luminance of the attended dots. These results provide further support for a global feature-based mechanism of attention, and show that the effect spreads across all features of an attended object, and to all locations of visual space.

© 2006 Elsevier Ltd. All rights reserved.

Keywords: Attention; Motion; Color; Feature; Motion aftereffect

1. Introduction

‘Attention’ is typically defined as the process of intentionally selecting a subset of all possible sensory information for further processing. Electrophysiological studies in the macaque monkey show that attention to a spatial location (spatial attention) generally enhances the responses of neurons with receptive fields at the attended location (e.g., Connor, Preddie, Gallant, & Van Essen, 1997).

More recently, however, electrophysiological (Bisley, Zaksas, Droll, & Pasternak, 2004; Martinez-Trujillo & Treue, 2004; Treue & Martinez Trujillo, 1999), neuroimaging (Saenz, Buracas, & Boynton, 2002; Sohn, Chong, Pappathomas, & Vidnyanszky, 2005), and behavioral studies (Melcher, Pappathomas, & Vidnyanszky, 2005; Saenz, Bura-

cas, & Boynton, 2003) show that attention to a feature, such as the direction of motion, at one location can influence the responses of neurons with receptive fields well outside the focus of attention. This can be called “global feature-based attention” because it influences the representation of stimuli globally, throughout the visual field.

Global feature-based attention might support processes such as perceptual grouping and visual search. For example, the spreading of feature-based attention can explain the gestalt principle of common fate, whereby objects that move together tend to be perceived as part of a common object. The spread of feature-based attention could aid the process of searching for a red book on a bookshelf (for example) by enhancing the representation of red objects throughout the visual field, and thereby increasing their salience as potential targets for saccades.

Global feature-based attention is inherently difficult to measure behaviorally because it is not possible to ask subjects to make direct perceptual judgments about

* Corresponding author. Fax: +1 858 455 7938.

E-mail address: boynton@salk.edu (G.M. Boynton).

URL: www.sn1-b.salk.edu (G.M. Boynton).

unattended stimuli. Behavioral measurements of the representation of unattended stimuli therefore require indirect methods such as priming (Melcher et al., 2005) or adaptation (Rezec, Krekelberg, & Dobkins, 2004). In the present study, as in a recently published study (Sohn et al., 2005), we used the motion aftereffect (MAE) as an indirect way to measure the effects of global feature-based attention on the response to unattended stimuli.

The MAE is an illusion in which viewing motion in one direction for several seconds subsequently induces the illusory percept of motion in the opposite direction. The logic of our approach is that the strength of the MAE should reflect the strength of the neuronal response to the unattended stimulus that induced it (Rezec et al., 2004).

A number of studies have shown that the strength of the MAE can be modulated by attention. Studies have quantified such effects by measuring the duration of the MAE (Chaudhuri, 1990; Rezec et al., 2004), the strength of motion in the opposite direction required to null the MAE (Lankheet & Verstraten, 1995), or the bias induced in the perceived direction of the MAE as a function of stimulus coherence (Alais & Blake, 1999). Chaudhuri showed that the duration of the MAE induced by a moving textured background was reduced when subjects performed an alphanumeric discrimination task at fixation compared to fixating passively. Rezec et al. (2004) also found a reduction in the MAE duration when subjects performed a vowel detection task at fixation compared to performing a speed judgment on the adapting dot stimulus. Furthermore, if two stimuli share the same spatial location, attending to one of two superimposed moving dot patterns can induce a MAE opposite to the direction of the attended field, whereas no MAE is observed under conditions of passive viewing (Lankheet & Verstraten, 1995). Thus, distinct motion mechanisms can be selectively adapted based on the stimulus attended, even for stimuli sharing the same spatial location. Our experiments expand upon these basic findings by examining how the features shared between attended and unattended stimuli can alter the strength of the MAE induced at a distant, unattended spatial location.

We show here that the MAE induced by an unattended stimulus is enhanced when attention is directed to a spatially distant stimulus that has the same direction of motion as the unattended stimulus. This provides further evidence that attention to a particular direction of motion at one location in the visual field enhances the response of neurons selective to that direction of motion throughout the visual field. Furthermore, our results show that this spread of attention does not require subjects to be performing a motion-related task, and that these effects spread to all areas of the visual field, regardless of the stimulus placed at the unattended location.

2. Methods

2.1. Subjects

All subjects had normal or corrected-to normal vision and ranged in age between 22 and 35 years. Five subjects participated in each of the three experiments. Two of the participants from Experiments 3 also participated in Experiments 1 and 2. Informed consent was obtained from all participants after they received a written explanation of the nature of the study. The experiment followed the tenets of the Declaration of Helsinki, and conformed to a protocol approved by the Salk Institute's Institutional Research Board (IRB).

2.2. Stimuli

Stimuli were presented using a gamma corrected NEC (LT157) LCD projector (1024 × 768, 60 Hz), that projected the image onto a back-projection screen placed 65 cm from the viewer's chin-rest. All stimuli were generated on a Macintosh G4 Powerbook using MATLAB (version 5) and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Stimuli consisted of moving dot patterns placed inside 4 deg radius circular apertures centered 9 deg to the left and right of fixation. All fields consisted of 50 dots, with each dot subtending 17.92 min of visual angle. Dots had a luminance of 2925 cd/m² and were presented on a black background. All dots moved at a baseline speed of 2 deg/s plus a random increment of up to 0.5 deg/s. Responses were collected on a standard numeric keypad.

2.3. Procedure

Each block consisted of 54 or 70 trials. As shown in Fig. 1, each trial consisted of an adapting phase and a test phase. Stimuli and tasks for the adapting phase were nearly identical to those used in a previous fMRI study (Saenz et al., 2002). During the adapting phase, subjects were cued to

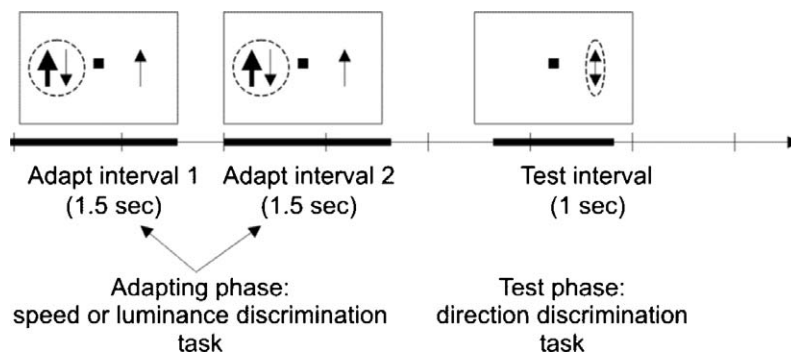


Fig. 1. Basic trial structure. Trials began with an adaptation phase in which subjects perform a 2IFC speed or luminance discrimination task on one field of overlapping dots on one side of fixation (upward moving dots on the left in this example). An unattended stimulus is presented in the contralateral hemifield (upward moving dots in this example). The MAE induced by this unattended stimulus is assessed during the test phase in which subjects indicate whether a slowly moving field of dots is moving upward or downward.

attend to one of two overlapping fields of coherently moving dots on one side of fixation (the left side in the example in Fig. 1, indicated by the arrows). Stimuli were presented in two 1500 ms intervals, while subjects performed either a two-interval forced-choice (2AltFC) speed discrimination task (Experiment 1) or a 2AltFC luminance discrimination task (Experiment 2) on one of the two overlapping fields of dots. The speed or luminance increment was varied on a trial-by-trial basis using a three-down one-up staircase procedure which leads to asymptotic performance at around 80% correct. Each staircase generated a psychometric function that was fit with a Weibull function using a maximum likelihood procedure to determine the speed or luminance increment corresponding to 80% correct performance.

While performing the task on the attended field of dots, subjects ignored the field of dots on the opposite side of fixation. The baseline speed and luminance of stimuli presented at the attended and ignored locations varied independently from each other so that subjects could not compare between stimuli to perform the task. For a given block, during the adapting phase, the stimulus on the unattended side contained one of the following: (1) motion in the *same* direction as the attended field of dots, (2) motion in the *opposite* direction as the attended field of dots, (3) *uncorrelated* motion, or (4) *no stimulus* at all (Experiment 3 only).

Immediately after the adapting phase of each trial, a motion nulling procedure was used during a test phase that measured the MAE induced at the unattended location. A slowly moving upward, downward or stationary field of dots was presented on the previously unattended side for one second. Subjects then made a two-alternative forced-choice decision about whether the dots appeared to move upward or downward. The true motion of the dots was chosen from a set of seven speeds centered on zero (−0.6, −0.4, −0.2, 0, 0.2, 0.4, and 0.6 deg/s) for Experiments 1 and 2, and (−.3, −.2, −.1, 0, .1, .2, and .3) for Experiment 3, which were presented in a pseudorandom order using a method of constant stimuli (10 presentations of each of the seven speeds over the 70 trials of the block). The responses of the subjects to these test stimuli were fit using a maximum likelihood procedure with a cumulative normal function. We estimated the speed of motion that nulls the MAE induced at the unattended location during the adapting phase by interpolating to find the speed for which subjects were equally likely to report that the dots were moving upwards and downwards.

If subjects did not respond within 200 ms after the offset of the stimuli in either the adapting phase or the test phase, the next phase or the next trial began regardless in order to maintain a constant 5 s inter-trial interval. This ensured that the timing of the trials and the state of adaptation was not dependent upon the timing of the subjects' responses.

In Experiment 1, subjects performed a speed discrimination task during the adapting phase, while either an upward, downward or uncorrelated stimulus was presented in the unattended location. Experiment 2 was the same as Experiment 1 except that subjects performed a luminance discrimination task during the adapting phase. In Experiment 3, subjects performed a speed discrimination task during the adapting phase, while no unattended (adapting) stimulus was presented in the opposite hemifield.

In Experiments 1 and 2 each condition was repeated 20 times across 5 sessions for each subject. The attended side, the attended direction of motion, and the adaptation condition were counterbalanced across experimental blocks. In Experiment 3 blocks using the single 'no-adapting-stimulus' condition were conducted 60 times across 5 sessions for each subject; the attended side and the attended direction of motion were counterbalanced across runs.

Eye movements can play a role in inducing MAEs for moving stimuli (Anstis & Gregory, 1965; Chaudhuri, 1991; Mack et al., 1987). When a stimulus moves across the eye, the resulting retinal motion can increase the likelihood of a nystagmus-like or pursuit-like eye movement being generated to stabilize the image on the retina. Such eye movements not only affect the retinal stimulus being adapted but can also induce non-visual signals that can create their own MAE. Thus, an "extra retinal" MAE can be induced even when subjects are allowed to track the adapting stimulus, despite retinal slip being minimal and even in the wrong direction to account for the direction of the induced MAE (Morgan, Ward, & Brussell, 1976).

Although we did not track eye movements in our experiments we did use peripheral adapting stimuli and a fixation target which can effectively control eye movements, eliminating an extra-retinal MAE, and if anything reversing it (Freeman, Sumnall, & Snowden, 2003). Furthermore, if the motion of the attended stimulus elicited eye movements, the resulting extra-retinal MAE should have been the same irrespective of the test stimulus at the unattended location and therefore no global feature-based attention effects would have been observed.

3. Results

3.1. Experiment 1: Speed discrimination task

Experiment 1 measured the strength of the MAE at an unattended location using stimuli similar to those in a previously published fMRI study as adapting conditions (Saenz et al., 2002). Subjects attended to the hemifield containing overlapping fields of upward and downward moving dots and performed a speed discrimination task on the field moving in the cued direction. Dots in the unattended field could have motion that was the *same* as the attended field, *opposite*, or *uncorrelated*.

Average results across the five subjects are shown in Fig. 2A, and individual subject results are shown in Figs. 2B–F. Error bars represent the standard error of the mean across subjects (Fig. 2A) or across repeated measurements within subjects (Figs. 2B–F).

For the *same* and *opposite* conditions, positive values are defined as nulling speeds in the opposite direction as the adapting stimulus. That is, the height of the bar graphs for the *same* and *opposite* conditions indicates the strength of the MAE induced by the unattended stimulus. The strength of the MAE is significantly stronger for the *same* condition than for the *opposite* condition (two-way ANOVA $p < .001$). This implies that the neuronal response to the unattended coherent field of dots is stronger when its direction of motion matches the direction of the attended motion, and weaker when its direction of motion does not match. Since the attended stimulus contained overlapping fields moving in both directions, this result must be due to the effects of feature-based attention, and not to bottom-up stimulus driven effects.

In the *uncorrelated* condition, a positive value indicates a MAE in the opposite direction to the attended direction. The unattended uncorrelated stimulus induced a significant MAE in the opposite direction to the attended direction (t test, $p < .05$). Since the adapting stimulus contained no net motion, this result again must be due to the attentional state of the observer, and not to stimulus-driven effects.

One possible explanation for these results is that spatial attention could have been inadvertently drawn to the adapting stimulus, with a greater spread to the adapting stimulus during the *same* condition than the *opposite* condition. If this were the case, and spatial attention is a limited resource, then we might expect poorer performance on the speed discrimination task (higher speed discrimination thresholds) on the attended side during the *same* than the *opposite* condition.

Fig. 3 shows speed discrimination thresholds during the adapting phase of the trials across the three conditions.

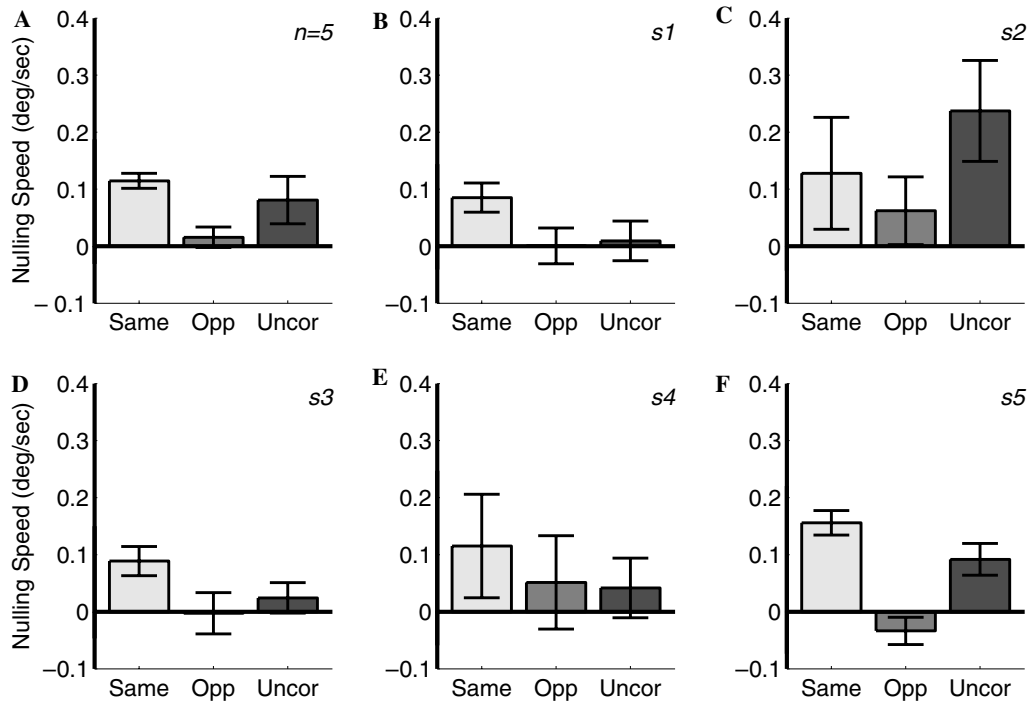


Fig. 2. MAE measurements from Experiment 1 (speed discrimination task during the adapting phase). Light, medium and dark shaded bars indicate the MAE nulling speed and direction induced by the unattended stimulus when moved in the *same* direction as the attended stimulus, the opposite direction, or contained uncorrelated motion, respectively. Positive values in the *same* and *opposite* conditions indicate nulling speeds in the opposite direction as the adapting stimulus. Positive values in the *uncorrelated* condition indicate nulling speeds in the opposite direction as the attended direction of motion. (A) Average across all five subjects. (B–E) Individual subject results. Error bars indicate standard error of the mean either across subjects (A) or within subjects (B–E).

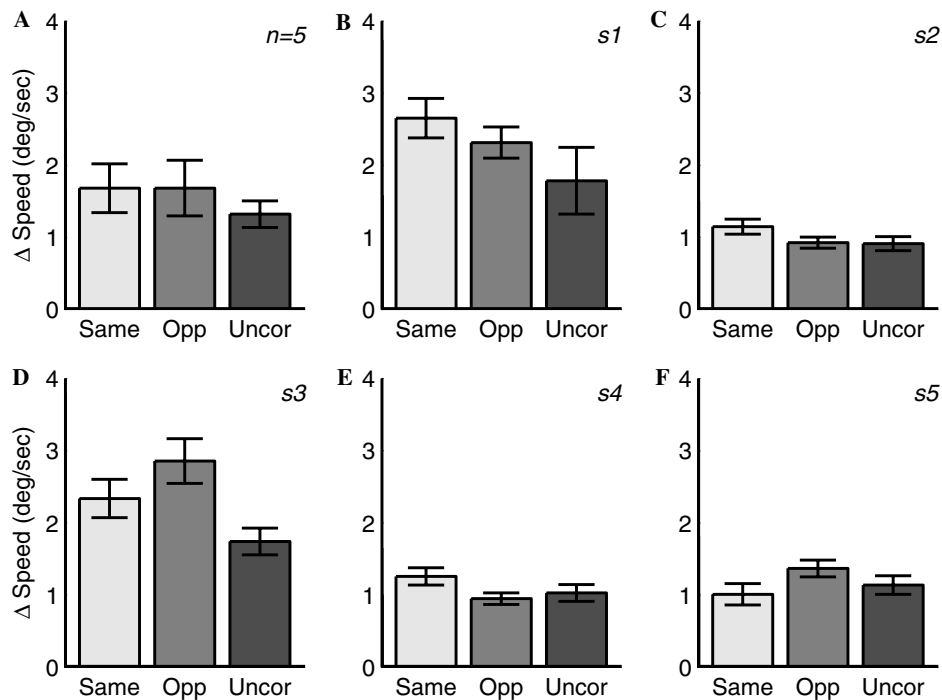


Fig. 3. Speed discrimination thresholds from the attended stimulus during the adapting phase from Experiment 1. (A) Average across all subjects. (B–E) Individual subject results. Error bars indicate standard error of the mean either across subjects (A) or within subjects (B–E).

Fig. 3A shows the means across the five subjects, and Figs. 3B–E show results from individual subjects. An analysis of variance shows no significant difference between the *same*

and *opposite* conditions, indicating that the unattended stimulus had no influence on the task at the attended location ($p = .702$). This suggests that it is unlikely that the

adaptation effects of Experiment 1 are due to the inadvertent allocation of spatial attention.

3.2. Experiment 2: Luminance discrimination task

In Experiment 2, the stimuli and two of the subjects were the same as for Experiment 1, but this time subjects performed a luminance discrimination task on one of the two overlapping fields of moving dots. The strength of the MAE induced by the unattended stimuli is shown in Fig. 4. Results are nearly identical to those measured using the speed discrimination task (Fig. 2). Again, an analysis of variance shows that the MAE induced by the *same* condition is greater than that induced by the *opposite* condition ($p < .001$). Therefore, performing a task unrelated to motion induces the same global feature-based effect for motion at the unattended location.

Fig. 5 shows luminance discrimination thresholds measured during the adapting phase, as shown in Fig. 3. An analysis of variance shows no significant difference between the luminance discrimination thresholds in the *same* and *opposite* conditions. This means that the unattended stimulus had no influence on performance of the task at the attended location (ANOVA, $p = .702$), and indicates that the unattended stimulus did not systematically attract spatial attention away from the task.

3.3. Experiment 3: No adapting stimulus

In Experiment 3, subjects performed a speed discrimination task during the adapting phase (as in Experiment 1),

but no stimulus was presented in the contralateral hemifield. Fig. 6 shows the strength of the MAE. As for the *uncorrelated* condition, a positive value indicates an MAE in the opposite direction to the attended direction of motion. Fig. 6A shows the mean across five subjects, and Figs. 6B–E show the results from the five individual subjects. Surprisingly, attention to one direction of motion in one hemifield induced a significant MAE in the opposite hemifield even without the presence of an inducing stimulus ($p < .05$).

4. Discussion

4.1. Neural mechanisms underlying the MAE

Our interpretation of our results relies on the assumption that the strength of the MAE reflects the strength of the response to the unattended adapting stimulus. This follows from most explanations of the MAE (and other after-effects), which propose that a fatigue in the population of neurons selective to the adapting stimulus leads to an imbalance in oppositely tuned mechanisms. This leads to a subsequent percept that favors the non-adapted subpopulation of neurons (Wohlgemuth, 1911). One model for this posits mutually suppressive neurons tuned to different directions of motion, in which adaptation to one direction of motion reduces the suppression of neurons tuned to other directions (Grunewald & Lankheet, 1996). Thus, adapting to one direction of motion can lead to a subsequent increase in the baseline responses to previously unstimulated neurons, with the greatest disinhibition for

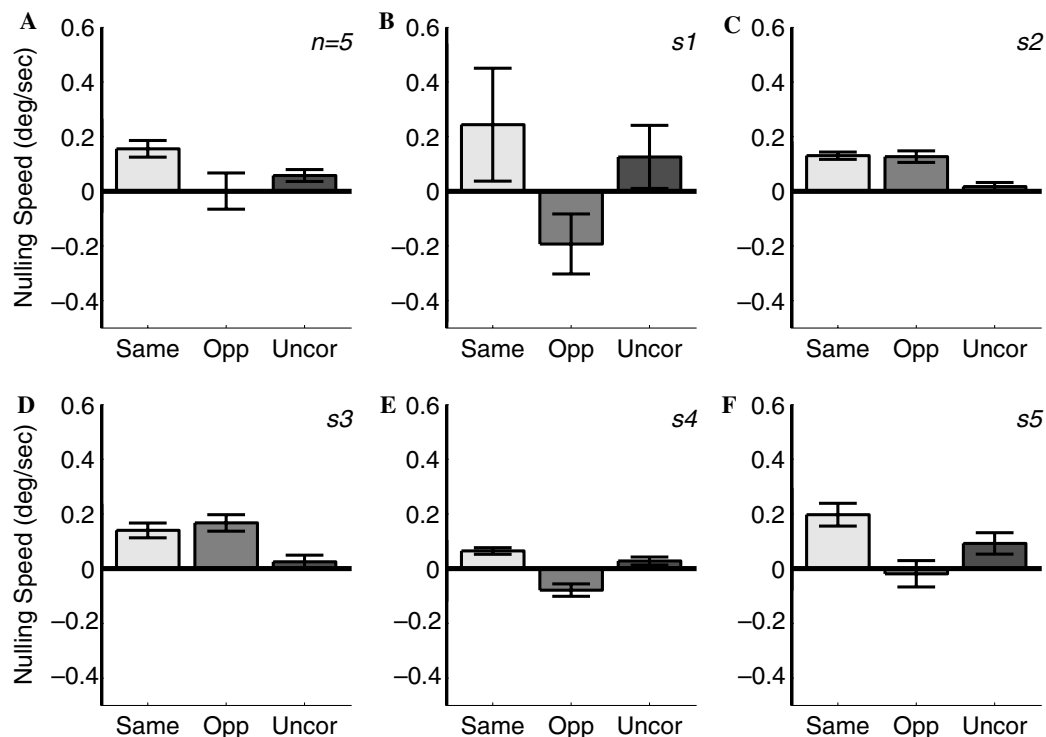


Fig. 4. MAE measurements from Experiment 2 (luminance discrimination task during the adapting phase). Conventions are the same as Fig. 2.

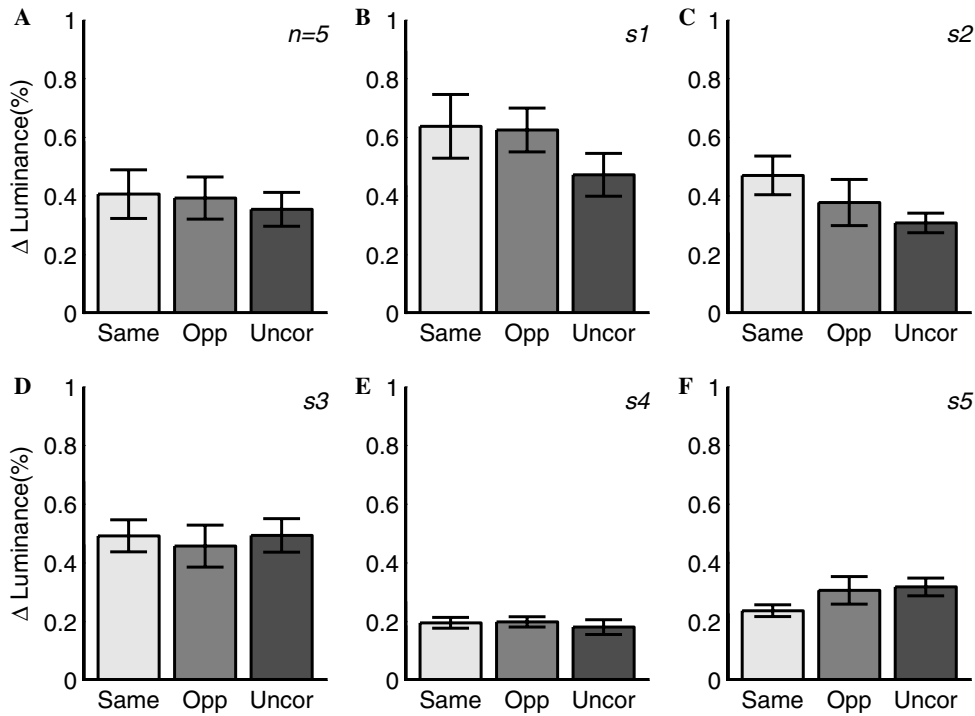


Fig. 5. Luminance discrimination thresholds from the attended stimulus during the adapting phase from Experiment 2. Figure conventions are the same as Fig. 3.

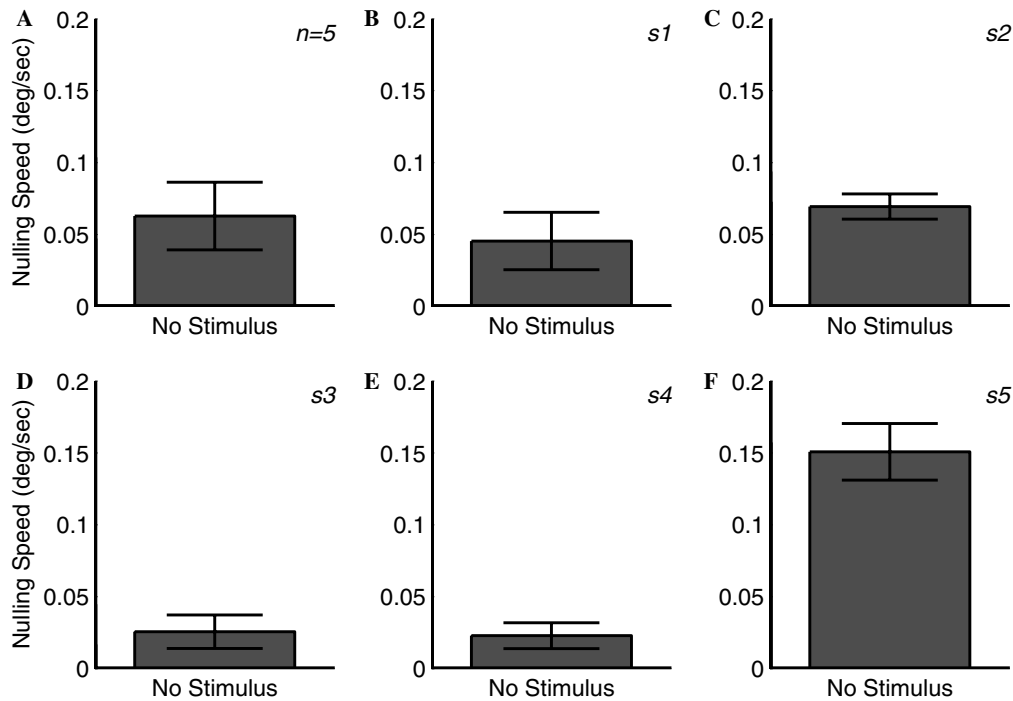


Fig. 6. MAE measurements from Experiment 3 (speed discrimination task with no adapting stimulus in the unattended location). Positive values indicate the nulling speed in the opposite direction of the attended direction of motion.

neurons tuned to the direction opposite the adapting stimulus.

Empirical results support the idea that the strength of the MAE grows with the strength of the response to the adapting stimulus. For example, the duration of the MAE

increases with stimulus contrast and with spatial attention (Rezec et al., 2004; see Pantle, 1998 for Section 3). It is, of course, unlikely that there is a truly linear relationship between the neuronal firing rate to the adapting stimulus and our particular measure of MAE strength (nulling

speed). In fact, it has been shown that the relationship between different measures of the MAE, such as duration or nulling speed, is not linear (Hershenson, 1989). Fortunately, our interpretation only requires that the strength of the MAE be monotonic with the strength of the response to the adapting stimulus.

4.2. The feature-similarity gain model of attention

An earlier electrophysiological study in macaque area MT showed that the firing rate of neurons with receptive fields outside the focus of attention was greater when attention was directed elsewhere to a stimulus that moved in the preferred direction of the neuron (Treue & Martinez Trujillo, 1999). A subsequent study showed that varying the attended direction of motion acts as a multiplicative gain factor on the neuronal responses to the unattended stimulus (Martinez-Trujillo & Treue, 2004). This gain factor can be excitatory (greater than 1) when attention is directed to the preferred direction of the neuron, or inhibitory (less than 1) when attention is directed to the anti-preferred direction. The authors describe their results in terms of a ‘feature similarity gain model’ of attention in which the strength of feature-based attentional modulation depends on the similarity between the attended direction of motion and the cell’s preferred direction (see Boynton, 2005, & Treue, 2001, for reviews).

The feature similarity gain model considered in the context of a population response can predict our behavioral results. Experiment 1 shows that stronger MAEs occur when the attended direction was the same, rather than the opposite direction of the adapting stimulus. In the *same* condition of our experiment, attention to the same direction of motion as the unattended stimulus should cause an increase in response in the subpopulation of neurons tuned to that direction of motion. In the *opposite* condition the response of the neurons tuned to the unattended stimulus should be suppressed. If we assume that the strength of the MAE reflects the strength of the neuronal response to the unattended stimulus (see Section 3 below), then it follows that the MAE induced by the unattended stimulus in the *same* condition should be larger than in the *opposite* condition.

A recent neuroimaging experiment using stimulus conditions similar to those used here also supports this feature similarity gain model by showing that the fMRI response in early retinotopic visual areas to an unattended stimulus was greater when that stimulus shared the same direction of motion or color as an attended stimulus in the contralateral visual field (Saenz et al., 2002).

The *uncorrelated* condition should induce some response in all direction-selective neurons. However, according to the feature similarity gain model, attention should selectively enhance the response of the subset of neurons tuned to the attended direction, and suppress those tuned to the opposite direction. This explains our observation of an MAE in the opposite direction to the attended motion in the uncorrelated condition.

A related study showed such an effect of attention on the MAE at the attended location. Searching for a brief pulse of weakly correlated motion in a particular direction can bias the direction of the MAE induced by an overlapping, unattended correlated motion stimulus. Even though the attended motion had low signal strength, attention was still able to bias the direction of the MAE, presumably by selectively enhancing a subset of neurons tuned to the attended direction (Alais & Blake, 1999).

4.3. The integrated competition model of attention

Experiment 2 shows that the MAE results measured in Experiment 1 do not depend on the subject performing a motion task. While the luminance discrimination task in Experiment 2 required segregating the upward and downward moving dots, observers did not have to make a decision about the motion of the attended stimulus. Nonetheless, we observed motion aftereffects that were as large as those found for the speed discrimination task. It appears that attention to any feature of a field of moving dots induces a spread of direction-selective attention throughout the visual field.

This is consistent with Duncan’s (1996) ‘integrated-competition’ model that states that when attention is directed to one feature of an object, all of its features will tend to become dominant in their respective cortical modules. Recent brain imaging and behavioral results support the integrated competition model. For example, using priming as a way to measure the representation of unattended stimuli, a recent behavioral study showed that attention to one feature of a stimulus automatically leads to a spread of attention to other features of the attended stimulus (Melcher et al., 2005). In a study closely related to ours, attention directed to the color of a stimulus in one location increased the duration of the MAE induced by an unattended moving stimulus when it shared the attended color (Sohn et al., 2005). The same study showed that fMRI responses in area MT+ to an unattended stimulus were enhanced when it had the same color as the non-moving attended stimulus.

4.4. The ‘remote’ or ‘phantom’ MAE

Experiment 3 showed that attention to a particular direction of motion induced a MAE in a region in the opposite hemifield even without an adapting stimulus. This is not the first example of the MAE spreading beyond an inducing stimulus. A relatively strong ‘phantom’ or ‘remote’ MAE was produced with adapting stimuli that induced an illusory moving stimulus in the unstimulated region (von Grunau, 1986; Weisstein, Maguire, & Berbaum, 1977; Zaidi & Sachtler, 1991). Other studies showed that a phantom MAE could be induced even without perceived motion in the unstimulated region. Unlike our study, these studies induced a global pattern of motion in which stimulated and unstim-

ulated regions were either overlapping (Bonnet & Pouthas, 1972) or were close to overlapping (Bex, Metha, & Makous, 1999; Price, Greenwood, & Ibbotson, 2004; Snowden & Milne, 1997). Because the stimulated regions were adjacent to where the phantom MAE was measured, the phantom MAE effect has often been attributed to direction-selective neurons with large receptive fields that covered both stimulated and unstimulated regions. However, our stimuli were distantly located in opposite peripheral hemifields, making this explanation of our results less plausible. Most importantly, while it has been shown that attention can affect the MAE at the adapting location (Chaudhuri, 1990; Lankheet & Verstraten, 1995), our study shows that attention alone can influence the strength and direction of the remote or phantom MAE.

The remote or phantom adaptation effects measured in Experiment 3 are evidence of direction-selective adaptation of neurons with receptive fields that do not overlap with the adapting stimulus. One possible explanation for this is that attention to a particular direction of motion in one location may enhance baseline-firing rates within all neurons selective to that direction of motion, regardless of their receptive field location and the presence or absence of a stimulus in their receptive field. This change in baseline firing rate could then induce adaptation in this sub-population, resulting in a MAE in the opposite direction to the attended motion. Thus, the same feature-similarity gain model (Treue, 2001) could explain both our results when a stimulus is present well as when a stimulus is absent.

There is electrophysiological and neuroimaging evidence to support this possibility. In a delayed match-to-sample task, the firing rates of macaque MT neurons were found to be mediated by whether their tuning matched the direction of the stimulus, during the early part of a 1500 ms delay period in which no stimulus was present (Bisley et al., 2004; Pasternak & Greenlee, 2005). Moreover, spatial attention produces robust fMRI responses in early retinotopic visual areas even in the absence of a stimulus (Kastner, Pinsk, De Weerd, Desimone, & Ungerleider, 1999; Ress, Backus, & Heeger, 2000), and this change in baseline response is affected by the feature that the subject is preparing to detect (Chawla, Rees, & Friston, 1999).

It is possible that a release from adaptation at remote unattended locations could be explained by the same mutually suppressive mechanism posited earlier to account for adaptation effects at the attended location. Interestingly, if these suppressive mechanisms are broadly tuned and operate widely across spatial locations, this could explain why remote MAEs have been found to be more broadly tuned than stimulus-driven MAEs (Price et al., 2004).

Acknowledgments

The authors would like to thank Ione Fine, John Serences, Minna Ng, Nils Pfeiffer, and Melissa Saenz for helpful suggestions on this manuscript. Supported by National Institutes of Health Grant EY-012925.

References

- Alais, D., & Blake, R. (1999). Neural strength of visual attention gauged by motion adaptation. *Nature Neuroscience*, 2(11), 1015–1018.
- Anstis, S. M., & Gregory, R. L. (1965). The aftereffect of seen motion: The role of retinal stimulation and of eye movement. *Quarterly Journal of Experimental Psychology*, 17, 173–174.
- Bex, P. J., Metha, A. B., & Makous, W. (1999). Enhanced motion aftereffect for complex motions. *Vision Research*, 39(13), 2229–2238.
- Bisley, J. W., Zaksas, D., Droll, J. A., & Pasternak, T. (2004). Activity of neurons in cortical area MT during a memory for motion task. *Journal of Neurophysiology*, 91(1), 286–300.
- Bonnet, C., & Pouthas, V. (1972). Apparent size and duration of a movement after-effect. *Quarterly Journal of Experimental Psychology*, 24(3), 275–281.
- Boynton, G. M. (2005). Attention and visual perception. *Current Opinion in Neurobiology*, 15(4), 465–469.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436.
- Chaudhuri, A. (1990). Modulation of the motion aftereffect by selective attention. *Nature*, 344(6261), 60–62.
- Chaudhuri, A. (1991). Eye movements and the motion aftereffect: Alternatives to the induced motion hypothesis. *Vision Research*, 31(9), 1639–1645.
- Chawla, D., Rees, G., & Friston, K. J. (1999). The physiological basis of attentional modulation in extrastriate visual areas. *Nature Neuroscience*, 2(7), 671–676.
- Connor, C. E., Preddie, D. C., Gallant, J. L., & Van Essen, D. C. (1997). Spatial attention effects in macaque area V4. *Journal of Neuroscience*, 17(9), 3201–3214.
- Duncan, J. (1996). In J. L. McClelland (Ed.), *Attention and performance*, XVI (pp. 549–578). Cambridge, MA: MIT Press.
- Freeman, T. C., Sumnall, J. H., & Snowden, R. J. (2003). The extra-retinal motion aftereffect. *Journal of Vision*, 3(11), 771–779.
- Grunewald, A., & Lankheet, M. J. (1996). Orthogonal motion after-effect illusion predicted by a model of cortical motion processing. *Nature*, 384(6607), 358–360.
- Hershenson, M. (1989). Duration, time constant, and decay of the linear motion aftereffect as a function of inspection duration. *Perception & Psychophysics*, 45(3), 251–257.
- Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1999). Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron*, 22(4), 751–761.
- Lankheet, M. J., & Verstraten, F. A. (1995). Attentional modulation of adaptation to two-component transparent motion. *Vision Research*, 35(10), 1401–1412.
- Mack, A., Goodwin, J., Thordarsen, H., Benjamin, D., Palumbo, D., & Hill, J. (1987). Motion aftereffects associated with pursuit eye movements. *Vision Research*, 27(4), 529–536.
- Martinez-Trujillo, J. C., & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology*, 14(9), 744–751.
- Melcher, D., Papathomas, T. V., & Vidnyanszky, Z. (2005). Implicit attentional selection of bound visual features. *Neuron*, 46(5), 723–729.
- Morgan, M. J., Ward, R. M., & Brussell, E. M. (1976). The aftereffect of tracking eye movements. *Perception*, 5(3), 309–317.
- Pantle, A. (1998). How do measures of the motion after effect measure up. In G. Mather, F. Verstraten, & S. Anstis (Eds.), *The motion aftereffect: A modern perspective* (pp. 23–39). Cambridge, MA: MIT Press.
- Pasternak, T., & Greenlee, M. W. (2005). Working memory in primate sensory systems. *Nature Review Neuroscience*, 6(2), 97–107.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Price, N. S., Greenwood, J. A., & Ibbotson, M. R. (2004). Tuning properties of radial phantom motion aftereffects. *Vision Research*, 44(17), 1971–1979.
- Ress, D., Backus, B. T., & Heeger, D. J. (2000). Activity in primary visual cortex predicts performance in a visual detection task. *Nature Neuroscience*, 3(9), 940–945.

- Rezec, A., Krekelberg, B., & Dobkins, K. R. (2004). Attention enhances adaptability: Evidence from motion adaptation experiments. *Vision Research*, 44(26), 3035–3044.
- Saenz, M., Buracas, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience*, 5(7), 631–632.
- Saenz, M., Buracas, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision Research*, 43(6), 629–637.
- Snowden, R. J., & Milne, A. B. (1997). Phantom motion after effects—evidence of detectors for the analysis of optic flow. *Current Biology*, 7(10), 717–722.
- Sohn, W., Chong, S. C., Papathomas, T. V., & Vidnyanszky, Z. (2005). Cross-feature spread of global attentional modulation in human area MT+. *Neuroreport*, 16(12), 1389–1393.
- Treue, S. (2001). Neural correlates of attention in primate visual cortex. *Trends in Neurosciences*, 24(5), 295–300.
- Treue, S., & Martinez Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, 399(6736), 575–579.
- von Grunau, M. W. (1986). A motion aftereffect for long-range stroboscopic apparent motion. *Perception & Psychophysics*, 40(1), 31–38.
- Weisstein, N., Maguire, W., & Berbaum, K. (1977). A phantom-motion aftereffect. *Science*, 198(4320), 955–958.
- Wohlgemuth, A. (1911). On the after-effect of seen movement. *British Journal of Psychology, Monograph Supplement*, 1, 1–117.
- Zaidi, Q., & Sachtler, W. L. (1991). Motion adaptation from surrounding stimuli. *Perception*, 20(6), 703–714.