UvA-DARE (Digital Academic Repository)

A dissociation of attention and awareness in phase-sensitive but not phase-insensitive visual channels

Brascamp, J.W.; van Boxtel, J.J.A.; Knapen, T.H.J.; Blake, R.

DOI

[10.1162/jocn.2009.21397](https://doi.org/10.1162/jocn.2009.21397)

Publication date

2010

Published in

Journal of Cognitive Neuroscience

[Link to publication](#)

Citation for published version (APA):

Brascamp, J. W., van Boxtel, J. J. A., Knapen, T. H. J., & Blake, R. (2010). A dissociation of attention and awareness in phase-sensitive but not phase-insensitive visual channels. *Journal of Cognitive Neuroscience*, 22(10), 2326-2344. <https://doi.org/10.1162/jocn.2009.21397>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

A Dissociation of Attention and Awareness in Phase-sensitive but Not Phase-insensitive Visual Channels

Jan W. Brascamp¹, Jeroen J. A. van Boxtel²,
Tomas H. J. Knapen³, and Randolph Blake^{1,4}

Abstract

■ The elements most vivid in our conscious awareness are the ones to which we direct our attention. Scientific study confirms the impression of a close bond between selective attention and visual awareness, yet the nature of this association remains elusive. Using visual afterimages as an index, we investigate neural processing of stimuli as they enter awareness and as they become the object of attention. We find evidence of response enhancement accompanying both attention and awareness, both in the

phase-sensitive neural channels characteristic of early processing stages and in the phase-insensitive channels typical of higher cortical areas. The effects of attention and awareness on phase-insensitive responses are positively correlated, but in the same experiments, we observe no correlation between the effects on phase-sensitive responses. This indicates independent signatures of attention and awareness in early visual areas yet a convergence of their effects at more advanced processing stages. ■

INTRODUCTION

Selective attention and visual awareness seem to go hand in hand: the contents of our visual awareness are strongly dominated by the objects to which we attend. By the same token, salient visual stimuli may fail to reach awareness when attention is impaired, either experimentally (Simons & Chabris, 1999; Mack & Rock, 1998) or pathologically (Driver & Vuilleumier, 2001). On a neural level, cortical responses to a visual stimulus are often enhanced when that stimulus is the object of attention (e.g., Treue & Maunsell, 1996; Motter, 1993; Moran & Desimone, 1985) and similarly when that stimulus reaches conscious awareness (e.g., Grunewald, Bradley, & Andersen, 2002; Leopold & Logothetis, 1996; Logothetis & Schall, 1989), further reinforcing the parallel between the two phenomena.

This partnership between attention and awareness, however, is not inviolable (Koch & Tsuchiya, 2007; Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006; Lamme, 2003; Tong, 2003; Crick & Koch, 1998). When experimental manipulations, such as masking, erase a visual stimulus from an observer's awareness, the instruction to pay attention to the location or features of the unperceived stimulus can still produce a measurable boost in its processing, indicating effective attention without awareness (Shin, Stolte, & Chong, 2009; Bahrami, Lavie, & Rees, 2007; Kanai, Tsuchiya, & Verstraten, 2006; Montaser-Kouhsari & Rajimehr, 2005). Conversely, some stimulus aspects, such as color, can be perceived while at-

tention is engaged elsewhere (Braun & Julesz, 1998; Rock, Linnett, Grant, & Mack, 1992). In terms of the firing behavior of visual neurons, the overall similarity between attention and awareness is less compelling within the primary visual cortex (V1) and the LGN. Within these early stages of visual processing, attention robustly modulates neural firing (Mcalonan, Cavanaugh, & Wurtz, 2008; McAdams & Reid, 2005; Ito & Gilbert, 1999; Roelfsema, Lamme, & Spekreijse, 1998; Motter, 1993), whereas awareness-related response modulation is weak (V1; Maier et al., 2008; Wilke, Logothetis, & Leopold, 2006; Leopold & Logothetis, 1996) or absent (LGN; Lehky & Maunsell, 1996; see also Sengpiel, Blakemore, & Harrad, 1995). The signatures of attention and awareness in the early visual brain thus appear distinct in terms of neural firing activity. It deserves mention that fMRI studies do show hemodynamic changes accompanying awareness manipulations in these early areas (Haynes, Deichmann, & Rees, 2005; Wunderlich, Schneider, & Kastner, 2005; Tong & Engel, 2001; Polonsky, Blake, Braun, & Heeger, 2000), a distinction between fMRI and neurophysiology that we will return to in the Discussion section. Perhaps the most direct demonstration of a separation between attention and awareness is their opposite effects on the formation of negative afterimages (Koch & Tsuchiya, 2007). Being aware of a visual stimulus causes that stimulus to leave a stronger afterimage (Gilroy & Blake, 2005; Tsuchiya & Koch, 2005), whereas attending to a stimulus results in its producing a weaker afterimage (Wede & Francis, 2007; Suzuki & Grabowecky, 2003).

The view of attention and awareness as closely associated yet separable phenomena invites the question of

¹Vanderbilt University, Nashville, TN, ²California Institute of Technology, ³Universite Paris Descartes, ⁴Seoul National University

how the two are related. In particular, we wondered how their separate natures are expressed in empirical observations made while manipulating either attention or awareness experimentally. Does the similarity between the neural concomitants of impairing either attention or awareness observed in many higher visual brain areas reflect their fundamental association? If so, what does it mean when this similarity in neural concomitants is weakened in early visual areas? And finally, how can the neural signatures of attention and awareness become so disparate as to yield opposite effects on a perceptual phenomenon so basic as the formation of visual afterimages?

To address these related questions, we developed a new psychophysical technique. Although staying close to known behavioral paradigms used in research on attention and awareness, this technique is exceptional in that it allows perceptual effects of attention and awareness to be traced back to the activity of specific classes of visual neurons: those sensitive to spatial phase and those insensitive to spatial phase. This specificity may prove revealing because phase sensitivity is characteristic of early visual areas such as LGN and V1 (e.g., Levitt, Schumer, Sherman, Spear, & Movshon, 2001; Hubel & Wiesel, 1968), where the neural signatures of selective attention and visual awareness differ most. Moreover, as detailed below, the striking effects of attention and awareness on the formation of afterimages may hinge on distinct contributions of phase-sensitive and phase-insensitive neurons, which can now be addressed. Our approach is furthermore characterized by maximally similar stimuli and methods used during both attention and awareness manipulations. This allows us to go beyond a qualitative comparison between the effects of attention and awareness and look for closer associations.

METHODS

Apparatus and Observers

Seven observers, four naive, participated in our experiments. Observers were seated in a darkened room and viewed visual stimuli on a gamma-linearized 1280 × 960 pixels CRT screen with a refresh rate of 120 Hz, through a mirror stereoscope at a visual distance of 81 cm. Background luminance and mean stimulus luminance were 43.5 cd/m². Observers were instructed to always keep their eyes directed at a white (87.0 cd/m²) fixation dot (radius = 0.04°) positioned in the center of a finely drawn circular edge (87.0 cd/m²; radius = 6.0°) to aid fusion.

Main Experiment

A commonly applied classification of visual neurons, dating back to the seminal work by Hubel and Wiesel (1962) and,

subsequently, Enroth-Cugell and Robson (1966), is that of phase-sensitive and phase-insensitive cells. Phase-sensitive neurons are characterized by their selectivity for the exact positioning of a stimulus within their receptive field. If the stimulus is a conventional sine-wave grating, this results in sensitivity to the spatial phase of the grating. A phase-sensitive neuron will respond to a given stimulus if it is presented at a particular, optimal, location within the neuron's receptive field but will decrease its response when the positioning within the receptive field changes. Neurons in the second category, phase-insensitive neurons, respond to an appropriate stimulus irrespective of its exact positioning within the receptive field. Although the degree of phase sensitivity may vary continuously within a given neural population (e.g., Levitt et al., 2001), it is useful to think of phase-sensitive and phase-insensitive cells as distinct categories when conceptualizing neural events that underlie vision (e.g., Hubel & Wiesel, 1962). Similarly, and although we acknowledge the nuances that apply, we will for clarity refer simply to "phase-sensitive cells" and "phase-insensitive cells" within the context of the present work.

Phase sensitivity is a hallmark of early stages of visual processing, being strong in primates subcortically (White, Sun, Swanson, & Lee, 2002; Levitt et al., 2001; Xu et al., 2001) and in the category of simple cells in V1 (Hubel & Wiesel, 1962). Phase insensitivity, conversely, is pronounced in extrastriate visual cortex (Levitt, Kiper, & Movshon, 1994; Maunsell & Van Essen, 1983) and in V1 complex cells (Hubel & Wiesel, 1962), but it is weak subcortically.

Our method was designed to independently evaluate perceptual manifestations of phase-sensitive cell activity and of phase-insensitive cell activity, without these separate effects getting entangled. Psychophysical methods do not ordinarily provide separate windows on the activity of both kinds of neural populations, as perception relies on both classes of neurons. Here we achieve a separation, however, by combining two classic psychophysical phenomena in a new manner. The two phenomena are threshold elevation and negative afterimages. Both reflect adaptation that accompanies neural responses to visual stimulation. They therefore provide a window on neural activity associated with vision. Neither by itself, however, allows this activity to be pinpointed to a particular neural class. Threshold elevation refers to impaired visual detection of a pattern following prolonged inspection of that or a similar pattern (e.g., Blakemore & Campbell, 1969). It is arguably associated with adaptation of various neural populations, both phase sensitive and phase insensitive (e.g., Burbeck, 1986; Smith, 1977). Afterimages are the illusory "photo negatives" one may perceive after prolonged exposure to unchanging visual input, for instance a dark spot that is perceived for some time after staring at a light bulb. Classic work has provided evidence that part of the adaptation underlying afterimages occurs in the retinal photoreceptors (Sakitt, 1976; Brindley, 1962; Craik, 1940),

prompting the textbook view of afterimages as an exclusively retinal phenomenon. More recent work, however, has identified contributions of adaptation at more central locations as well, including the visual cortex itself (Gilroy & Blake, 2005; Tsuchiya & Koch, 2005; Suzuki & Grabowecky, 2003; Shimojo, Kamitani, & Nishida, 2001; Anstis, Rogers, & Henry, 1978; Virsu & Laurinen, 1977).

Afterimages and threshold elevation are each other's complement, in the sense that afterimages constitute perception of something that is not there, whereas threshold elevation causes a failure to perceive something that is there. This complementary relation plays a central role in the method we devised. As elaborated in the next paragraph, afterimages likely arise at processing stages that are sensitive to spatial phase, after which their neural signal may pass through phase-insensitive stages (Wede & Francis, 2007; Suzuki & Grabowecky, 2003; see also Alonso & Martinez, 1998). The latter, although not the basis of afterimages, do modulate afterimage visibility. Following this reasoning, the "true" strength of an afterimage at its neural source forms a probe into adaptation at phase-sensitive processing stages. Phase-insensitive adaptation, on the other hand, is manifested in the component of threshold elevation that impairs the visibility of an afterimage once it has formed. The obvious difficulty is to distinguish these two contributions on the basis of the phenomenal experience of afterimages. This is what our method was designed to do.

From the outset we acknowledge that perceptual experiments such as ours do not measure responses of neurons directly. Still, it is possible to infer response properties of neurons that underlie the perceptual effects we measure, and this is a principle aim of our study. Accordingly, in the rest of this article, we will for sake of clarity refer to "phase-sensitive channels" and "phase-insensitive channels," thus drawing a parallel with this widely used classification of visual neurons in neurophysiology while reminding the reader of the distinction between the neural activity and the perceptual outcome we measure.

An origin of afterimages in phase-sensitive channels can be inferred from the nature of afterimages as "photo negatives" of their physical inducers, having the same spatial location and opposite spatial phase. These characteristics require an origin in channels whose response in fact varies with spatial phase. That phase-insensitive channels can and do modulate the visibility of an afterimage follows from the finding that static inducers simultaneously cause afterimages as well as threshold elevation that impairs the visibility of those same afterimages (Wede & Francis, 2007; Suzuki & Grabowecky, 2003; Georgeson & Turner, 1985). This component of threshold elevation necessarily arises in phase-insensitive channels because no phase-sensitive channel that adapts to the inducing image can also be involved in detecting the subsequent afterimage, for the inducer and the afterimage are of opposite phase.

We developed a method built on signal detection theory to simultaneously gauge both the "true strength" of an afterimage before it passes through phase-insensitive stages

and the strength of the component of threshold elevation that impairs afterimage visibility to an observer. We use these two measures to gauge the adaptation state of phase-sensitive and phase-insensitive channels, respectively.

Figure 1A, left, shows a typical trial sequence. Observers maintained strict fixation while a stationary "inducing" grating (0.67 Michelson contrast), windowed by a Gaussian aperture ($\sigma = 0.68^\circ$), was presented for 4 sec; a period of stimulation that is sufficient to induce both an afterimage and an elevation in threshold. At the end of this 4-sec induction period, the fixation dot turned from white to red and the inducing grating was ramped off (over 250 msec) into a "nulling" grating, designed to cancel the afterimage as it arose. This nulling image was a negative of the afterimage or, in other words, a low-contrast version of the inducing grating itself. A mask pattern subsequently replaced the nulling grating 1 sec after its appearance. This moment was marked by a tone. In cases where a nulling grating like ours is approximately comparable in strength to the afterimage it is designed to cancel, neither the afterimage nor the nulling grating is seen when the two are superimposed, the perceptual result being nothing but a blank screen (Kelly & Martinez-Uriegas, 1993; Georgeson & Turner, 1985; Leguire & Blake, 1982). We quantified this by asking observers whether a grating was visible the moment before the mask appeared or, in other words, whether the mask replaced a grating or a blank screen in their perception. The clear temporal separation between the color change of the fixation dot at the end of the inducer period and the auditory signal that marked the perceptual task 1 sec later minimized the possibility of mistaking the inducer grating itself for an imperfectly canceled afterimage during this task. Observers were instructed to ignore perceptual changes, if any, that might occur during the second leading up to the task.

Over repeated trials, we varied the contrast of the nulling grating to generate the sort of response pattern shown in Figure 1A, right. Observers reported perceiving a grating both when the nulling image was very weak and when it was relatively strong. In the former case, it was the afterimage that was seen whereas in the latter case it was the nulling image itself that was seen, as symbolized by the gradient displayed above the plot. Between these extremes lies a range of nuller contrasts where observers frequently reported stimulus absence, leading to a U-shaped curve. The low central region of this curve corresponds to nuller strengths that are similar to the afterimage strength, resulting in near-perfect cancellation. (As is true for all detection tasks, the underlying psychometric function does not bend sharply but, instead, gradually.)

The location of the center of this U-shape is our first variable of interest. It provides a measure of the "true" afterimage strength that is not influenced by threshold elevation, that is, before passage through phase-insensitive stages (Figure 1B, left). Certainly, whether a particular combination of nuller and afterimage is detected depends on threshold elevation. However, this is true both when the nuller is weaker than the afterimage, determining

Figure 1. Experimental method. (A) Basic trial sequence (left). Observers viewed a sine wave grating for 4 sec while keeping steady fixation. Then, as the fixation dot changed color, this inducing grating was replaced with a “nulling” grating designed to cancel the inducer’s afterimage. One second later this nuller was replaced with a mask at the sound of a tone. The nuller’s contrast was systematically varied between trials. Observers reported on every trial whether they saw a grating at the location of the nulling grating, the instant before the mask appeared. The plot on the right shows a typical response pattern. Observers perceive a grating (the afterimage) when the nuller is weak and they also perceive a grating (the nuller itself) when the nuller is relatively strong. At intermediate nuller contrasts, however, the nuller and the afterimage cancel out and observers often do not perceive any grating. (B) Interpretation of the data (see text for details). The location of the U-shaped detection curve indicates the “true contrast” of the afterimage, which is associated with adaptation of phase-sensitive channels (left). The width of the U-shaped detection curve provides a measure of contrast detection threshold and is interpreted as a measure of adaptation in phase-insensitive channels (middle). The slope of the curve’s legs is associated with the level of internal noise (right). (C) Manipulation of attention and awareness. To vary the degree of attention allocation to the inducing grating observers were instructed to attend either the grating or a fixation dot located elsewhere (left). To modulate awareness of the inducing grating, it either was or was not rendered invisible using interocular suppression (right).

the left leg of the curve, and when the nuller is stronger than the afterimage, determining the right leg. We performed a control experiment to confirm such a symmetrical influence of threshold elevation on both sides of the curve and thus its irrelevance for the location of its center (see Appendix). Threshold elevation, conversely, is reflected in the distance between the left and right leg of the U-shape; our second variable of interest. In the presence of threshold elevation, a larger deviation between the strength of the afterimage and the nuller is required for their combination to be detected, leading to wider U-curves (Figure 1B, middle). Figure 1B, right, illustrates the influence of internal noise, the final parameter shaping the detection curve but otherwise not relevant for our analyses. We quantified both the location of the center of the U-shape and the width of the U-shape by fitting the function $P(C_N, C_D) = 1 - \Phi_{|C_N - C_A|, \sigma}(C_D)$ to detection data from each condition and observer separately. Here $P(C_N, C_D)$ is the probability of detection given nuller contrast C_N and contrast detection threshold C_D , Φ is a cumulative Gaussian func-

tion, C_A is true afterimage contrast, and σ is the standard deviation of the Gaussian (see Appendix for derivation). In this manner, we thus obtained estimates of both the true afterimage contrast at its neural origin and the contrast detection threshold.

Observers performed the nulling task for four different, randomly interleaved conditions designed to manipulate attention and awareness (Figure 1C), indicated by a letter at the start of each trial. In the two “attention” conditions, observers directed their attention (but not their fixation) either toward the inducing grating, which was presented outside fixation, or toward the fixation point while ignoring the inducing grating. In the two “awareness” conditions, the inducing grating was presented monocularly and either with or without an incompatible pattern presented at the same location in the other eye. For this purpose, we used a checkerboard pattern with black (~ 0 cd/m²) and white (87.0 cd/m²) square checks of side 0.15°, windowed using a Gaussian aperture ($\sigma = 0.75^\circ$). The pattern reversed contrast polarity every 100 msec and

revolved 1° around its center every 10 msec, in a direction that was chosen randomly every trial. The high contrast and dynamic nature of this pattern ensured that the inducing grating was rendered invisible to the observer, who perceived only the dynamic suppressor (Tsuchiya & Koch, 2005). The suppressor, like the inducer, was ramped off over 250 msec at the end of the adaptation period.

To facilitate comparison between manipulations of attention and awareness, the stimulus layout was identical for all four conditions; that is, on every trial, the inducing grating was presented to only one randomly chosen eye and always at the same eccentricity (2.8°). Moreover, by confining the perceptual judgment of grating presence or absence to the very last moment before mask onset in the manner described above, we ensured ample time for observers to direct their attention to the location of the nuller (a full second between the moment the fixation dot changed color and the moment the tone sounded), regardless of the attention instructions that applied while the inducer was present. To further ensure maximal similarity between the four conditions and also to keep task demand at a moderate level, we chose not to include additional tasks for controlling observers' attention during adaptation. Previous work has shown mere attention instructions to be adequate in similar paradigms (Suzuki & Grabowecy, 2003), as confirmed by the fact that our attention manipulations yielded significant effects. For the same reasons, we did not ask observers to monitor the effectiveness of perceptual suppression on-line. Instead, we performed a separate experiment beforehand using the same stimuli, establishing that for these observers perceptual suppression was broken less than 0.5% of the time.

The angular position of the inducer shifted clockwise by 45° on every consecutive trial to avoid between-trial buildup of adaptation. The nulling grating was presented to the same eye as the inducer. Following indications that afterimage formation may depend on the spatial frequency of the inducing grating (e.g., Georgeson & Turner, 1985; Leguire & Blake, 1982), we performed our experiment at three different spatial frequencies (0.66, 2.0, or 3.3 cycles/degree).

Fits were obtained by minimizing the sum of squared errors using a Levenberg–Marquardt algorithm. Occasionally, detection data would not display a clear dip in the center of the U-shape, presumably because eye movements had disrupted the nulling method or because the dip was narrow and fell between the nuller contrasts we sampled. The 8 of 84 cases in which this happened were characterized by fits that converged on a negative value for the detection threshold. These were not included in our analyses.

Adaptation to the Suppressor

We performed a control experiment to investigate the effect of adaptation to the suppressor on threshold elevation in the absence of an inducer. Here observers were exposed to a suppressor pattern in isolation and subse-

quently reported detection of a physical grating. This condition was paired with a baseline condition in which the suppressor pattern was replaced by nothing but a blank screen, followed by the same detection task. The stimulus dimensions and trial structure of this experiment were identical to those of the main experiment, and from an observer's perspective, the condition with a suppressor in isolation was identical to the main experiment condition including a suppressor and an inducer. The only addition was a white dial (87.0 cd/m²; length = 0.13°) extending from fixation during adaptation to indicate target location, which was not necessarily obvious in trials with neither an inducer nor a suppressor. Detection targets were presented monocularly to the eye contralateral to the suppressor, if present.

Afterimage Duration

Finally, we performed an experiment measuring the perceived duration of afterimages following exposure to an inducer. Stimuli and trial structure were again identical to those of our main experiment, except that no nuller was presented, and the nulling task was replaced by a judgment of afterimage duration. In addition, we added an inducer spatial frequency of 4.6 cycles/degree, resulting in four different frequencies in total. Inducers of such a high spatial frequency do not lend themselves well to nulling tasks, which require an exact overlap between the nuller and the afterimage, but they are suitable for duration measures.

RESULTS

Effects on Phase-sensitive Channels

Figure 2 shows the results of the attention conditions (left column) and the awareness conditions (right column). By means of illustration, Figure 2A displays the raw outcome (dots) for one combination of observer and spatial frequency as well as the fitted curves. It can be seen that the width and the position of the curve are influenced, both by attending to the inducer (left panel; light curve vs. dark curve) and by being aware of the inducer (right panel; light curve vs. dark curve). To highlight how this affects the variables of interest, the bars drawn on the *x*-axis indicate in corresponding colors the locations of the centers of the U-shaped curves as well as intervals that span one detection threshold to both sides of each center. We note that, in these examples, the effects of attention and awareness on the left leg of the U-shape seem small. This is because the symmetrical increase in curve width, caused by threshold elevation (Figure 1B, middle panel), is counteracted in this region by the overall rightward shift that occurs simultaneously (Figure 1B, left panel). We performed a control experiment (see Appendix) that rules out an alternative explanation, that threshold elevation would increase curve width asymmetrically on only the right side.

Figure 2. Result of the main experiment. (A) Effects of manipulating attention (left) and awareness (right) on the outcome of our nulling task for one observer and one inducer spatial frequency. Dots are empirical data points; curves show the fits on which we base our further analyses. Both attention and awareness influence the position and width of the detection curve, as summarized by the position and width of the bars drawn on the *x*-axis. (B) The true contrast of the afterimage, measured by the position of the U-shaped detection curve, averaged over all seven observers. The true contrast increases when attention is allocated to the inducing grating (left; light vs. dark curve) and also when the inducing grating reaches conscious awareness (right; light vs. dark curve). This is true for all inducer spatial frequencies tested (*x*-axis). (C) Observer-averaged contrast detection thresholds, measured by the width of the U-shaped detection curve, are more elevated following attention to the inducing grating (left; light curve) than following attention to the fixation point (left; dark curve). Perceptually suppressing the inducing grating, on the other hand, has a weaker effect that varies with inducer spatial frequency (right).

Figure 2B summarizes effects on the center position, averaged over all seven observers. Because this position denotes the contrast of the afterimage before any effects of threshold elevation (Figure 1B, left panel), we will refer to it as “true afterimage contrast.” A rightward shift of the center position corresponds to an increase in this true contrast, whereas a leftward shift denotes a decrease. The left panel of Figure 2B shows that directing attention to the inducing grating tends to increase the true contrast of its afterimage (*y*-axis). This trend is consistent across three different inducer spatial frequencies (*x*-axis) and significant when combining across spatial frequencies, ANOVA with factors spatial frequency ($n = 3$) and attention condition ($n = 2$), $F = 6.2$ and $p < .01$ for the latter factor. Figure 2B, right panel, shows a similar

result for the awareness manipulation, with a significantly higher true afterimage contrast following a visible inducing grating than following a perceptually suppressed grating, ANOVA with factors spatial frequency ($n = 3$) and awareness condition ($n = 2$), $F = 83.0$ and $p \ll .01$ for the latter factor. As afterimage formation can be attributed to adaptation of phase-sensitive channels, these results indicate that the inducing gratings cause more adaptation in phase-sensitive channels, both when the gratings are attended and when they are consciously perceived.

Effects on Phase-insensitive Channels

Figure 2C shows the accompanying effects observed in these same experiments on the width of the U-shaped

curve (Figure 1B, middle panel). This variable reflects the contrast difference between afterimage and nullo image that is required for an observer to detect their combined sum (Figure 1A, right panel), thus providing a measure of contrast detection threshold. The left panel shows that contrast detection thresholds (y -axis) are significantly higher following adaptation to a grating that is attended than to a grating that is not attended, ANOVA with factors spatial frequency ($n = 3$) and attention condition ($n = 2$), $F = 8.5$ and $p < .01$ for the latter factor.

The right panel of Figure 2C shows how perceptual suppression of our inducing grating affected subsequent contrast detection thresholds. The pattern of results is quite different from that observed for the attention manipulation (left panel). Here, the effects on contrast detection threshold are smaller, and they vary across spatial frequencies of the inducing grating. At a low spatial frequency perceptual suppression tends to increase contrast detection thresholds (leftmost pair of points), whereas at higher spatial frequencies detection thresholds tend to be reduced following perceptual suppression (middle and right pairs of points). The latter pattern, but not the former, would be similar to the attention result (left panel). These effects of perceptual suppression thus vary with spatial frequency and are, overall, not statistically significant.

Possible Influence of the Suppressor

We were initially surprised by the indication that perceptual suppression influences threshold elevation in a manner that depends on inducer spatial frequency (Figure 2C, right panel). In particular, we had not expected any trend toward strengthened threshold elevation following perceptual suppression, as is apparent at the lowest inducer spatial frequency (Figure 2C, right panel, leftmost pair of data points). If we interpret this result in terms of adaptation to the inducing grating, this result would suggest that the inducing grating elicits stronger threshold elevation when it is not perceived, which disagrees with previous results (e.g., Blake, Tadin, Sobel, Raissian, & Chong, 2006). This motivated us to explore an alternative account. We considered the possibility that detection thresholds may be influenced, not only by adaptation to the inducing grating but also by adaptation to the dynamic pattern that is presented to the other eye to elicit perceptual suppression. Such an effect would contribute to threshold elevation in the perceptual suppression condition, potentially influencing our results.

To test this alternative account, the same observers who participated in our main experiment performed a control experiment to estimate the amount of threshold elevation that a suppressor might cause on its own. Observers adapted to the dynamic suppressor pattern, but this time without any inducing grating in the other eye (Figure 3A). As the dynamic pattern was removed, it was immediately replaced by a physical grating presented to the other eye,

Figure 3. Control experiment to investigate adaptation to the dynamic suppressor. (A) Trial structure. Observers adapted to a suppressor in isolation and then detected a physical grating. This condition was compared with one in which nothing was shown during the adaptation period. (B) Detection was impaired following adaptation to the suppressor (black solid curve) relative to the condition without an adapting stimulus (gray solid curve). The dotted heavy lines replot the results from our main experiment (Figure 2C, right), showing detection thresholds following adaptation to an inducing grating (light dotted curve) and to a perceptually suppressed inducing grating (dark dotted curve). (C) Approximation of the effect of perceptual suppression per se on threshold elevation while controlling for adaptation to the suppressor pattern. Threshold elevation in the perceptual suppression condition is expressed relative to the control condition involving adaptation to a suppressor pattern in isolation (dark bars). Threshold elevation in the condition without perceptual suppression is expressed relative to the control condition with no adapting pattern (light bars). This analysis points to consistently more threshold elevation by a nonsuppressed inducer than by a perceptually suppressed inducer.

and observers were asked to report detection of this grating. From an observer's perspective, this condition was indistinguishable from the perceptual suppression condition of our main experiment, although no afterimages were involved. The control experiment also included a second, baseline condition in which observers viewed just a blank screen during the adaptation period, and then similarly reported detection of a physical grating. By systematically varying the contrast and spatial frequency of the detection target, we obtained measures of detection performance at various spatial frequencies, with or without prior adaption to the dynamic pattern.

Figure 3B shows results from the two conditions in this control experiment. Following adaptation to just the dynamic suppressor (black curve), detection thresholds are increased relative to the baseline condition without an adapting pattern (gray curve), ANOVA with factors spatial frequency ($n = 3$) and condition ($n = 2$), $F = 21.9$ and $p \ll .01$ for the latter factor. This difference is largest for the lowest spatial frequency of the detection target (two-tailed paired t test; $df = 6$, $t = 6.7$, $p \ll .01$) and is much reduced at the two higher spatial frequencies ($df = 6$, $t = 3.6$, $p < .05$ and $df = 6$, $t = 2.5$, $p < .05$, respectively). This confirms that the dynamic pattern by itself has the potential to raise detection thresholds.

The result of Figure 3B indicates that, in our main experiment, detection thresholds following perceptual suppression likely reflect adaptation to the inducing grating as well as adaptation to the suppressor itself. Because we are interested primarily in adaptation to the inducing grating and how this is affected when the grating's visibility is impaired, we reanalyzed the data from our main experiment to isolate this component. As a first step, the heavy dotted lines in Figure 3B replot the results from the conditions in our main experiment that addressed the effects of perceptual suppression, previously shown Figure 2C, right panel. The comparison of those conditions with our control conditions is revealing. The light dotted curve, associated with the main experiment condition without perceptual suppression, runs consistently above the gray curve, associated with the control condition without any adapting stimulus. This indicates that the inducing grating causes threshold elevation when it is consciously perceived. The dark dotted curve, however, which is associated with the main experiment condition using a perceptually suppressed grating, runs remarkably close to the black curve that denotes the control condition with a suppressor in isolation. This indicates that a perceptually suppressed grating causes only little threshold elevation on top of that caused by the suppressor itself.

The comparison of Figure 3B strongly suggests that suppressing an image from awareness reduces the amount of threshold elevation it causes, although this may not always be obvious when using a suppressor that, itself, causes threshold elevation. Figure 3C shows an approximation of the true effect of perceptual suppression on threshold

elevation in our main experiment while attempting to compensate for adaptation to the suppressor. We recalculated the threshold elevation data displayed in Figure 2C, right panel, as the logarithm of the amount of threshold elevation relative to a baseline ($\log(\text{threshold}) - \log(\text{baseline})$). This is a standard measure for quantifying threshold elevation (e.g., Snowden & Hammett, 1996; Bjørklund & Magnussen, 1981), and it is useful here because it provides the liberty of choosing different baselines for the conditions with and without perceptual suppression. To control for the adapting effects of the suppressing pattern itself, we calculated threshold elevation in the perceptual suppression condition relative to the control condition with just a suppressor, whereas threshold elevation in the no-suppression condition was calculated relative to the control condition without any adapter. When isolating the effect of the inducing grating in this manner threshold elevation is consistently more pronounced without perceptual suppression (light bars) than with perceptual suppression (dark bars), ANOVA with factors spatial frequency ($n = 3$) and suppression condition ($n = 2$), $F = 11.2$ and $p < .01$ for the latter factor. This trend is present for all three spatial frequencies and it is significant for each of the highest two spatial frequencies individually (two-tailed paired t test, $df = 6$, $t = 2.7$, $p < .05$ for 2.0 cycles/degree and $df = 4$, $t = 2.9$, $p < .05$ for 3.3 cycles/degree).

The values in Figure 3C should not be interpreted quantitatively because it is unsure whether and how adaptation to the suppressor pattern and to the inducing grating interact in the main experiment. Nevertheless, the results of Figure 3B and C combined corroborate the conclusion that an inducing grating causes less threshold elevation under conditions of perceptual suppression. This indicates that adaptation to the inducer at phase-insensitive stages is reduced when the inducer is suppressed from awareness.

In sum, the left panels of Figure 2 indicate that attending to a stimulus increases the amount of adaptation caused in phase-sensitive channels (panel B) as well as in phase-insensitive channels (panel C). The right panels of Figure 2, combined with Figure 3, show the same for becoming aware of the stimulus, which also enhances adaptation in both types of neural channels.

Predicting Afterimage Durations: A Paradoxical Effect of Awareness

The subjective strength and duration of an afterimage, two commonly used psychophysical measures (e.g., Gilroy & Blake, 2005; Tsuchiya & Koch, 2005; Suzuki & Grabowecky, 2003; Shimojo et al., 2001; Virsu & Laurinen, 1977), depend on the true contrast of that afterimage as well as on the observer's contrast sensitivity (Leguire & Blake, 1982). The duration for which an afterimage remains visible, for instance, is equal to the time it takes for the true afterimage contrast to fall below an observer's contrast detection

threshold. If true afterimage contrast and contrast detection threshold are indeed the main factors underlying perceived afterimage durations, we should be able to predict these durations from the values measured in our main experiment in the following way.

After an inducing image is removed from the screen, the true contrast of its afterimage gradually decays. To a first approximation, the time it takes for this decaying contrast to fall below an observer's contrast detection threshold is proportional to $\log(\frac{C_A}{C_D})$: the logarithm of the ratio between the initial afterimage contrast C_A and the contrast detection threshold C_D . This calculation follows directly from a simplified scenario in which true afterimage contrast starts dropping exponentially upon inducer offset, causing the afterimage to disappear as soon as its true

contrast drops below the detection threshold. (A more complete scenario, where both the afterimage and the detection threshold start decaying upon inducer offset, has too many free parameters to be constrained by our data.) This calculation is not intended to provide quantitative duration estimates. It is very useful, however, for making a comparison between conditions. When we enter the true afterimage contrasts and contrast detection thresholds of Figure 2 into this proportionality, we obtain qualitative predictions of the perceived afterimage durations associated with our manipulation of attention or awareness. These are shown in Figure 4A.

The bottom panels of Figure 4A facilitate comparison by depicting differences between the durations predicted in various conditions. Comparing the condi-

Figure 4. Perceived afterimage duration. (A) The top panels show predictions of perceived afterimage duration, based on measured values of true afterimage contrast and contrast detection threshold (Figure 2), and normalized relative to the “no suppression” condition. The bottom panels show the differences between these predicted values. These predictions indicate that attention to the inducer shortens the resulting afterimage (left panel), whereas awareness of the inducer shortens the afterimage at high inducer spatial frequencies, but lengthens them at low inducer spatial frequencies (right panel). (B) These predictions were confirmed in an experiment where observers reported the perceived duration of afterimages following adaptation periods identical to those in our main experiment.

tion where the inducer is attended with the condition where it is not attended, our predictions show that attention to the inducer should reduce the duration of the subsequent afterimage (Figure 4A, left column). This predicted decrease in duration is counterintuitive, in the sense that it accompanies an increase in adaptation (Figure 2). It is, however, in agreement with experimental findings (Wede & Francis, 2007; Suzuki & Grabowecky, 2003).

For inducers of low spatial frequency, our awareness prediction (right column) is also consistent with experimental findings (Gilroy & Blake, 2005; Tsuchiya & Koch, 2005). At those spatial frequencies, awareness of the inducer is predicted to lengthen the resulting afterimage. The most interesting prediction, however, occurs for the awareness manipulation at high spatial frequencies. There, our calculations indicate that consciously perceived inducers should leave briefer afterimages than perceptually suppressed inducers. This is a paradoxical outcome that, moreover, is at odds with experimental findings so far (Gilroy & Blake, 2005; Tsuchiya & Koch, 2005).

To verify our theoretical predictions, we performed an additional experiment where observers reported the perceived duration of afterimages following exposure to an inducing grating. Experimental conditions were the same as in our main experiment, the only difference being that the nulling task was replaced by a duration judgment.

Figure 4B depicts measured afterimage durations, for the same seven observers that also participated in our main experiment plus one additional naive observer. These measurements confirm the predictions (Figure 4A) made based on their detection thresholds and true afterimage contrasts. Attention to the inducer is shown to reduce the duration of the resulting afterimage (left panel), ANOVA with factors spatial frequency ($n = 4$) and attention condition ($n = 2$), $F = 4.9$ and $p < .05$ for the latter factor. Awareness of the inducer, on the other hand, is shown to have different effects at different spatial frequencies (right panel), ANOVA with factors spatial frequency ($n = 4$) and awareness condition ($n = 2$), $F = 10.1$ and $p \ll .01$ for the interaction between the two factors. Indeed, awareness of the inducer increases afterimage durations at low spatial frequencies (right panel; two-tailed paired t test on the 0.66 cycle/degree data, $df = 7$, $t = 7.8$, $p \ll .1$) but reduces afterimage durations at higher spatial frequencies (two-tailed paired t test on the 3.3 cycle/degree data, $df = 7$, $t = 2.7$, $p < .05$; the 2.0 cycle/degree data not significant). Results at an even higher spatial frequency (not shown; 4.6 cycles/degree) closely matched those at 3.3 cycles/degree, with significant reductions of afterimage duration following both attention and awareness (two-sided paired t test, $df = 7$, $t = 2.5$, $p < .05$ for attention and $df = 7$, $t = 3.0$, $p < .05$ for awareness).

The reversal with spatial frequency of the effect of awareness, observed in the right panels of Figure 4, was further replicated in an experiment that used a matching

procedure to determine subjective afterimage strength instead of afterimage duration (see Appendix), thus confirming its robustness.

The agreement between predictions and observations in Figure 4 supports the idea that the driving factors underlying subjective afterimage measures are true afterimage contrast and contrast detection threshold, and it confirms that the nulling method of our main experiment successfully isolated those factors. In addition, it aids the interpretation of existing work centered on subjective afterimage measures (see Discussion).

A Dissociation of the Effects of Attention and Awareness

Our results indicate that attention and awareness have similar effects on the neural processing of a stimulus. They both increase the true contrast of its afterimage (Figure 2), and they both increase the amount of contrast threshold elevation the stimulus causes (Figures 2 and 3). This indicates that both phase-sensitive channels and phase-insensitive channels adapt more strongly when their input is attended as well as when it is consciously perceived.

To investigate the extent to which these qualitative similarities reflect a more fundamental association, we analyzed individual differences between our observers in the measured effects. Figure 5A displays the effects of awareness (x -axis) and attention (y -axis) on true afterimage contrast, as measured in our main experiment (Figure 2, left column). The effect of being aware of the inducing grating was calculated for each observer simply by subtracting the true afterimage contrast observed in the perceptual suppression condition from the true afterimage contrast observed in the condition without perceptual suppression, averaging across spatial frequencies. Attention effects were calculated in an analogous manner by subtracting the true contrasts with and without attention.

Figure 5A displays no correlation between the extent to which awareness affects true afterimage contrast and the extent to which attention does, thus providing no evidence that the two phenomena share a neural basis at phase-sensitive stages (linear regression; $p = .19$ and adjusted $r^2 = 0.18$).

Figure 5B shows the outcome of the same analysis, performed on contrast detection data. The axes depict differences between the values of log threshold elevation observed in various conditions, first displayed in Figure 3C. Contrary to true afterimage contrast, this measure does reveal a correlation between the effects of attention and awareness (linear regression; $p = .01$ and adjusted $r^2 = 0.68$). This correlation suggests a direct association between the effects of attention and awareness in phase-insensitive channels in our experiment.

The correlation of Figure 5B does not depend on our use of log threshold elevation as a metric of contrast sensitivity. We observe a similar correlation when using the

Figure 5. Correlation between the effects of attention and awareness. (A) The extent to which awareness of the inducer influences true afterimage contrast (x-axis) shows no correlation with the extent to which attention to the inducer does (y-axis), when comparing different observers (signified by individual dots). (B) In a similar analysis, the effects of attention and awareness on threshold elevation do show a correlation. This plot displays differences in log (threshold elevation), and similar correlations are observed when using alternative contrast sensitivity measures.

raw contrast detection thresholds of Figure 2, right column, instead (linear regression; $p = .04$ and corrected $r^2 = 0.51$), or when using threshold elevation without taking the logarithm (linear regression; $p = .03$ and adjusted $r^2 = 0.57$).

Before ascribing the lack of correlation in Figure 5A to a true difference between the effects of attention and awareness at phase-sensitive stages, one needs to consider whether the null result could, instead, arise from a lack of statistical power that causes us to overlook a correlation that is in fact there. As a first step to investigate this issue, we compared the slopes of the regression lines in Figure 5A and B and found that they are significantly different (t test where $t = \frac{\text{slope}_1 - \text{slope}_2}{\text{SE}_{\text{slope}_1 - \text{slope}_2}}$ and $df = (n_1 - 2) + (n_2 - 2)$, e.g., Glantz, 2005; $df = 10$, $t = 2.54$, $p < .05$). This result demonstrates that, in any case, our data do provide sufficient statistical power to distinguish the two slopes. As a second step, we performed an error analysis to determine the amount of uncertainty on our measurements of true afterimage contrast and contrast threshold elevation, respectively. As the two regression lines of Figure 5 are each based on the same amount of data, a lack of statistical power would be a more likely explanation if the data underlying Figure 5A, where no correlation is evident, were noisier than the data underlying Figure 5B, which shows a significant correlation. We therefore performed a bootstrap analysis in which we randomly drew observations with replacement from each subject-condition of our original data set and then fitted detection curves to these simulated data sets. This analysis revealed an average standard error on our measure of true afterimage contrast that was in fact smaller (0.022 contrast units) than the average standard error on our measure of contrast detection threshold (0.032 contrast units). In other words, the uncertainty on

the measure that did yield a positive correlation (detection threshold) was larger than the uncertainty on the measure that showed no correlation (true afterimage contrast). This renders it unlikely that the lack of correlation in the latter measure is due to a lack of statistical power. As a final test, we analyzed the data of a control experiment in search of a correlation between the effects of attention and awareness on true afterimage contrast. In this experiment (originally performed for a different reason; see Figure A2), we obtained estimates of true afterimage contrast (but not detection threshold) from six observers using a different nulling method than the one employed in our main experiment. The outcome of this experiment was similar in many respects to that of our main experiment, the relevant point here being that the effects of attention and awareness on true afterimage contrast again did not show any sign of correlation (linear regression; $p = .50$ and adjusted $r^2 = -0.10$). In sum, the above analyses combined indicate that our data do provide sufficient statistical power to identify a correlation if it were there, and they support the interpretation that the lack of correlation in Figure 5A reflects a true difference in the effects of attention and awareness in phase-sensitive channels.

DISCUSSION

Exposure to a static image causes an afterimage as well as a reduction in contrast sensitivity that affects the visibility of that same afterimage. We devised a method to simultaneously determine both the true contrast of an afterimage as it exists at its neural origin and the magnitude of the sensitivity reduction that impairs afterimage visibility. One characteristic of afterimages is that their exact spatial layout is dictated by the spatial layout of their inducers. On

the basis of this spatial specificity, we argue that afterimages arise in phase-sensitive channels: the only channels whose responses display sufficient spatial specificity. Consequently, we treat true afterimage contrast as a gauge on adaptation in phase-sensitive channels. Conversely, we argue that the visibility of afterimages (but not their true contrast) is modulated by adaptation in phase-insensitive channels that are involved in processing the afterimage after it has formed. Indeed, the fact that the sensitivity reduction that an inducer causes can affect the visibility of its own afterimage implies the involvement of phase-insensitive channels: the only channels responsive to both the inducer and the afterimage, which are of opposite phase. We treat this sensitivity reduction, therefore, as a probe into adaptation of phase-insensitive channels.

Using our method, we provide evidence that suppressing a visual stimulus from awareness lessens the buildup of adaptation in response to the stimulus, both in phase-sensitive channels and in phase-insensitive channels. In addition, our results indicate that diverting the observer's attention away from the stimulus similarly lessens adaptation accumulation in both types of channels. When comparing between observers, we find a significant, positive correlation between the effects of attention and awareness on threshold elevation but no correlation between their effects on true afterimage contrast. This suggests that the signatures of attention and awareness are unrelated at the level of phase-sensitive channels but not at the level of phase-insensitive channels.

A Dissociation and a Correlation between the Effects of Attention and Awareness

The indication that the effects of attention and awareness are uncorrelated at phase-sensitive stages may reflect the manner in which awareness is modulated in paradigms involving interocular suppression. These rely on a bottom-up conflict between stimuli presented to the two eyes, which may well engage neural interactions at early processing stages where monocular signals meet (e.g., Lee, Blake, & Heeger, 2004; Polonsky et al., 2000). Similar interactions are not necessarily engaged by our attention instruction, which arguably draws on top-down modulation of sensory processing by signals originating from parietal and frontal areas (e.g., Kanwisher & Wojciulik, 2000). As phase-sensitive neurons are primarily found at the same early processing stages where monocular signals interact (e.g., Levitt et al., 2001; Hubel & Wiesel, 1962), this may explain the lack of correlation between the effects of attention and awareness on phase-sensitive channels suggested by our experiment.

In contrast, our results do indicate a correlation between the effects of attention and awareness on adaptation in phase-insensitive channels. This suggests that our awareness manipulation results in changes in the activity of phase-insensitive channels that are akin to the effects

of withdrawing attention. One scenario consistent with this finding is that the suppressor renders the inducer invisible by ways unrelated to attention, after which this invisibility causes observers to stop directing attention to the inducer. An alternative possibility is that the mechanism by which interocular suppression renders targets invisible inherently involves interference with attention allocation to the target. Perhaps, in other words, we observe a correlation because invisibility during interocular suppression is in fact a type of inattentional blindness. This latter scenario is consistent with evidence that attentional selection is one of the main determinants of the contents of conscious awareness (Koch & Tsuchiya, 2007; Driver & Vuilleumier, 2001; Rees & Lavie, 2001).

Significance for the Interpretation of Single Cell and fMRI Results

The robust effect of perceptual suppression of an inducing grating on both true afterimage contrast and contrast detection threshold is indicative of an effect at early stages of visual processing. Contrast thresholds are modulated by adaptation of neurons in V1 (Sclar, Lennie, & DePriest, 1989; Movshon & Lennie, 1979). True afterimage contrast, in turn, is arguably an indicator of neural adaptation in V1 and in subcortical structures such as the LGN, where neurons exhibit much stronger phase sensitivity (White et al., 2002; Levitt et al., 2001; Xu et al., 2001; Hubel & Wiesel, 1962) than at more advanced processing stages (Levitt et al., 1994; Maunsell & Van Essen, 1983). This interpretation fits well with other recent psychophysical studies involving visual aftereffects (e.g., Van Boxtel & Koch, 2009; Blake et al., 2006; Kanai et al., 2006; Gilroy & Blake, 2005; Tsuchiya & Koch, 2005) and also with fMRI results that indicate robust response reduction in early visual areas during perceptual suppression (Haynes et al., 2005; Wunderlich et al., 2005; Tong & Engel, 2001; Polonsky et al., 2000). This view, however, does not dovetail very well with single-unit recording results from alert, behaving monkeys, which generally indicate little response modulation in lower visual areas accompanying perceptual suppression (Maier et al., 2008; Wilke et al., 2006; Lehky & Maunsell, 1996; Leopold & Logothetis, 1996).

Various reasons have been offered to account for the differences between fMRI and physiological measurements during perceptual suppression (Maier et al., 2008), one being that perceptual suppression may provide an instance where modulatory signals into a brain area become decoupled from spiking activity (see also Wilke et al., 2006; Fries, Schröder, Roelfsema, Singer, & Engel, 2002). As fMRI BOLD responses have been argued to register synaptic activity more than spiking activity (Viswanathan & Freeman, 2007; Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001), this might cause modulated BOLD activation without proportionate changes in spiking activity. Psychophysical results now add an interesting piece to this puzzle:

Whatever aspect of neural responses is being modulated in early visual areas during perceptual suppression, that aspect accounts for a sizable fraction of visual adaptation. Given that adaptation is a key element of neural information processing (Kohn, 2007; Schwartz, Hsu, & Dayan, 2007), psychophysical observations thus underscore that awareness-related modulation in those early areas is functionally significant and alters the information represented in those areas. It is conceivable that altered adaptation during perceptual suppression involves a dissociation between neural adaptation and neural spiking, similar to that observed in other situations (e.g., Crowder et al., 2006).

Opposite Effects of Attention and Awareness?

Previous findings have shown opposite effects of attention and awareness on the perception of subsequent afterimages (Wede & Francis, 2007; Gilroy & Blake, 2005; Tsuchiya & Koch, 2005; Suzuki & Grabowecky, 2003; see also Lou, 2001), suggestive of a profound distinction between attention and awareness (Koch & Tsuchiya, 2007). Although we successfully replicated the findings that underlie this suggestion, our results indicate that attention and awareness in fact affect afterimage formation in a very similar fashion.

The effects of attention and awareness are similar in that they both raise the true contrast of the resulting afterimage (Figure 2) as well as raising the amount of contrast threshold elevation the inducer causes (Figures 2 and 3). Opposite effects of attention and awareness are only observed when using measures that reflect the joint outcome of both these factors, such as perceived afterimage duration or subjective afterimage intensity. These measures increase when the true afterimage contrast is raised, but they decrease when the contrast detection threshold is raised. Because, in reality, both factors are raised, the relative strength of the two factors becomes important (Figures 2 and 4). In other words, opposite effects of attention and awareness on subjective afterimage intensity or duration are not indicative of opposite effects on a neural level, but of an altered balance between counteracting effects.

In addition, our results indicate that the documented effect of awareness, namely, an increase in the perceived duration and subjective intensity of the afterimage, is limited to the low inducer spatial frequencies used in the original studies (0.5–0.6 cycles/degree in Tsuchiya & Koch, 2005; 1.0 cycle/degree in Gilroy & Blake, 2005). At higher spatial frequencies, the effect of awareness reverses, causing a decrease in subjective afterimage intensity and duration. At these higher spatial frequencies, therefore, the effects of attention and awareness are in qualitative agreement even at the level of subjective afterimage intensity and duration (Figure 4 and Appendix). In our experiments, the effect of awareness manipulations is influenced by adaptation to the suppressor, particularly when the inducer has a low spatial frequency (Figures 2–4). The importance of this component may depend on factors such as the spatial

frequency content of the suppressor (e.g., Blakemore & Campbell, 1969), and a control experiment by Tsuchiya and Koch (2005) provides evidence against a key role of adaptation to the suppressor at their stimulus settings (their Figure 3). Nevertheless, our results demonstrate that multiple components contribute to the effects of perceptual suppression on subsequent afterimage perception and therefore call for caution when interpreting these effects.

Previous studies have already inferred that augmented contrast threshold elevation is responsible for the reduction in subjective afterimage duration and intensity following attention to the inducer (Wede & Francis, 2007; Suzuki & Grabowecky, 2003). Like ours, those studies have argued for a distinction between what we are calling phase-sensitive channels and phase-insensitive channels in interpreting such effects. Our work builds on this notion and provides an empirical confirmation that altered threshold elevation is responsible for subjective afterimage weakening following attention. In addition, we provide direct measurements of changes in threshold elevation associated with manipulations of attention and awareness as well as of accompanying changes in true afterimage contrast that in previous work remained obscured by changes in detection threshold.

Conclusion

Our results indicate that manipulations of attention and awareness have similar effects on adaptation in visual neurons while explaining why previous findings suggest that they do not. Our results also show that the impression from physiological work, that awareness-related response modulation is minimal in lower visual areas, is not borne out in psychophysical adaptation measures, which indicate that functionally significant response modulation occurs in those areas. In spite of an overall similarity in the effects of attention and awareness, we observe a direct correlation only in the measure that reflects adaptation in phase-insensitive channels. In contrast, our data indicate that adaptation of phase-sensitive channels is differentially modulated by attention and by awareness. These results are consistent with a dissociation of attention and awareness in early, but not later, visual areas.

APPENDIX

Derivation of an Equation for the Detection Curve

Our approach assumes that the nuller image and the neural source of the afterimage are combined at some phase-sensitive stage, after which their combined signal passes through a phase-insensitive stage that may modulate visibility, a scenario schematically depicted in Figure A1A. This choice for a serial organization was motivated by the fact that predominantly phase-sensitive stages precede predominantly phase-insensitive stages in the visual hierarchy (e.g., White et al., 2002; Levitt et al., 2001, 1994; Maunsell

& Van Essen, 1983). Even within V1, where phase-sensitive simple cells and phase-insensitive complex cells are both common, the response profiles of complex cells arise from phase-sensitive inputs from LGN and from V1 simple cells, thus roughly maintaining the same hierarchical sequence (Rust, Schwartz, Movshon, & Simoncelli, 2005; Alonso & Martinez, 1998; Ferster & Lindström, 1983; Movshon, Thompson, & Tolhurst, 1978).

In standard signal detection theory, a sensory signal is represented as a distribution on some intensity axis, which describes the probabilities associated with all possible intensities the signal can take on (e.g., Green & Swets, 1988). Under the simplest assumption that the signal from the nulloer image and the signal at the afterimage source are combined linearly, their combined signal can be represented as a Gaussian distribution whose mean equals the absolute difference between the expected values of the two constituent signals (Figure A1B, left). We take the difference here because the nulloer image and the afterimage have opposite polarity.

$$p(I) = \varphi_{|\mu_N - \mu_A|, \sigma}(I). \quad (1)$$

Here $p(I)$ denotes the probability density at intensity I , φ is a Gaussian distribution, μ_N and μ_A indicate the expected values of the nulloer intensity and the signal intensity at the afterimage source, respectively, and the standard deviation σ quantifies the amount of uncertainty in the combined signal.

We again follow standard signal detection theory by positing a decision criterion D on this axis that denotes the intensity above which the observer reports seeing a stimulus and below which he or she reports not seeing a stimulus (Figure A1B, left). The probability of detection is then represented by the area under the distribution to the right of this criterion value (shaded area in Figure A1B, left) or 1 minus the area left of the decision criterion:

$$P(\mu_N, D) = 1 - \Phi_{|\mu_N - \mu_A|, \sigma}(D) \quad (2)$$

Here $P(\mu_N, D)$ indicates detection probability for a given nulloer intensity μ_N and criterion intensity D . $\Phi(D)$ denotes a cumulative Gaussian distribution function at intensity D , which, by definition, equals the integral of the accompanying Gaussian function from minus infinity to D .

In our experiment, we vary the contrast of the nulloer grating. To calculate detection probability as a function of the combination of this nulloer contrast and the contrast detection threshold, we replace the arbitrary intensity units in Equation 2 with physical intensity measures of contrast:

$$P(C_N, C_D) = 1 - \Phi_{|C_N - C_A|, \sigma}(C_D) \quad (3)$$

Here C_N and C_A denote the contrast of the nulloer and the true contrast of the afterimage (at its source in phase-sensitive channels), respectively, and C_D is the contrast detection threshold. We obtain our measures of detection threshold and true afterimage contrast by fitting this

Figure A1. Derivation of a detection curve. (A) We assume a serial organization, where the nulloer image and the neural source of the afterimage are first combined at phase-sensitive stages, after which their combined signal passes through phase-insensitive stages. (B, left) The combined signal can be represented as a Gaussian distribution on an intensity axis, centered on the absolute difference between the mean intensities of the nulloer signal and the afterimage signal (dashed vertical line). The probability of detection is then denoted by the area under the distribution (shaded gray) that lies to the right of a decision criterion (vertical line). (B, right) Replacing abstract “intensity” with physical contrast we obtain a curve that denotes this probability as a function of nulloer contrast. We obtain estimates of the true afterimage contrast and the detection threshold by fitting this curve to our data.

Figure A2. Test for asymmetric effects of threshold elevation.

(A) Using inducers and nulling gratings identical to our main experiment, observers judged the spatial phase of the nuller–afterimage combination (left), resulting in a systematic shift in reported phase with increasing nuller contrast (right). (B) If contrast threshold elevation specifically affected visibility of nuller–afterimage combinations that share the spatial phase of the inducer, this would result in an asymmetric response curve (left). No such asymmetry is detectable in our data (right).

equation to our experimental data, with free parameters C_A , C_D , and σ .

Response Bias

If cognitive rather than sensory factors cause differences in decision criterion between conditions, this amounts to a response bias. When signal detection theory is applied to detection of physical stimuli, sensory and cognitive factors are separated by including blank trials in the experiment, during which no stimulus is presented (e.g., Green & Swets, 1988). A systematic approach using blank trials is not possible here because we do not know which nuller contrast will exactly cancel the afterimage on a given trial, as would be required for a signal strength of zero. An alternative kind of blank trial, lacking both the afterimage and the nuller image, is only informative in the condition involving a suppressor, as observers would notice the absence of an inducer in the remaining conditions. Our perceptual suppression condition did in fact involve randomly interspersed blank trials of this second kind, and the fact that our observers erroneously reported detecting a grating in only three of 103 such blank trials instills confidence that they reported perception accurately.

Nevertheless, we were careful to examine the possibility that cognitive rather than sensory factors could explain our results. We found several aspects of our data that are inconsistent with this idea.

Our data show that attention to the inducer elevates subsequent detection thresholds (Figure 2). Furthermore, the prediction that subjective afterimage duration would

be reduced by this threshold elevation was confirmed experimentally (Figure 4). This was true for both our naive observers and our nonnaive observers. An explanation in terms of response bias would involve the assumption that attention to the inducer biases observers to report subsequent afterimages to be briefer. It is known, however, that observers do not expect afterimages to be briefer after attending to the inducer (Suzuki & Grabowecy, 2003). Moreover, the same attention-mediated reduction in subjective afterimage duration has been observed in experiments similar to ours that used measures that do not depend on criterion (Suzuki & Grabowecy, 2003), rendering it unlikely that criterion is the cause in our case.

Our observations regarding the effects of perceptual suppression of the inducer on detection threshold are also at odds with bias-based explanations. Our raw measure of detection threshold (Figure 2) is consistent with our observations on perceived afterimage duration (Figure 4). Both measures show opposite effects in the perceptual suppression condition, depending on the spatial frequency of the inducer. The hypothesis that these measures reflect a response bias thus requires the assumption that perceptual suppression of the inducer causes opposite biases, depending on the spatial frequency of the inducer. We see no justification for this assumption.

No Evidence for an Influence of Phase-sensitive Channels on Threshold Elevation

We performed a control experiment to verify that threshold elevation has an equal influence on the left leg and the

right leg of the U-shaped curve and thus does not change the position of the curve's center. This is related to the notion that threshold elevation in our paradigm reflects adaptation of phase-insensitive channels. A significant contribution of phase-sensitive channels to threshold elevation could result in a larger influence of threshold elevation on the right leg of the U-shaped curve because in that region the phase of the nuller–afterimage combination matches the phase of the inducer to which the channels adapted.

The trial sequence in this experiment was almost identical to that of the main experiment, using the same inducers and nulling stimuli (Figure A2A, left panel). Our manipulations of attention and awareness were also the same, and we used the same three spatial frequencies. Instead of performing a detection task, however, our six observers (five naive; all were participants of the main experiment) were instructed to report the spatial phase of the nuller–afterimage combination. Following Kelly and Martinez-Uriegas (1993), we added a thin red marker line through the center of the nulling grating, thus reducing the task to a report of which side of the line was flanked by the darker grating band (Figure A2A, left panel). We found that we could facilitate this judgment further by orienting the stimulus gratings parallel to the imaginary line connecting the stimulus center to fixation. Otherwise the stimuli in this experiment were identical to those of the main experiment.

The right panel of Figure A2A displays a typical response pattern obtained by varying the contrast of the nulling grating. The location along the curve where the observer is

equally likely to report either side of the line to be flanked by the darker band marks the point where nuller contrast and true afterimage contrast are equal ($C_N = C_A$).

If threshold elevation had a stronger impact on the visibility of nuller–afterimage combinations that have the same spatial phase as the inducer, this would specifically affect the right side of the response curve, both in our main experiment and in this control experiment. In the control experiment, this would compel observers to guess their answer more often for those nuller–afterimage combinations specifically, thus resulting in a shallower slope of the rightward section of the curve. Figure A2B, left panel, illustrates this scenario.

To test for this kind of asymmetry in the response curve, we fitted our data with a combination of two half cumulative Gaussian functions, joined together at their means (Figure A2B, left panel). Free parameters were the standard deviation of the left half section (σ_L), the standard deviation of the right half section (σ_R), and the location of the mean. Although the idea of a sharp kink separating the left and right section of the response curve is probably a simplification, any asymmetry should show up as a difference between σ_R and σ_L . Specifically, an influence of phase-sensitive channels on detection threshold in this experiment should show up as $\sigma_R > \sigma_L$.

The right panel of Figure A2B displays the distribution of all differences between σ_R and σ_L for our subject conditions. There is no sign that the distribution would be skewed toward values greater than zero, as would be

Figure A3. The effect of perceptually suppressing the inducer on the subjective strength of the afterimage. (A) Trial structure. Two inducers were presented on opposite sides of fixation, one perceptually suppressed, and the other visible. Observers judged which of the resulting afterimages was stronger. Stimuli were sine waves windowed with a Gaussian aperture ($\sigma = 0.68^\circ$), presented at 2.5° eccentricity. (B) Typical response pattern as a function of the contrast of the visible inducer (left panel). When the visible inducer and the invisible inducer are of equal contrast, the afterimage of the visible inducer is usually judged to be stronger for low inducer spatial frequencies, but the converse is true for high inducer spatial frequencies (right panel). This reversal of the effect of perceptual suppression on subjective strength is analogous to the reversal found with perceived duration (Figure 4).

expected if $\sigma_R > \sigma_L$, and the distribution mean is not significantly different from zero (two-sided t test; $df = 61$ after removing 10 nonconverging fits, $t = 1.4$, $p = .18$). This indicates that threshold elevation does not differentially affect either side of the response curve in our control experiment renders it unlikely that it does in our main experiment, which used the exact same stimuli. It thus supports the idea that phase-sensitive channels do not significantly influence threshold elevation in our tasks.

Subjective Afterimage Strength

Like perceived afterimage duration, the subjective intensity of an afterimage depends on a combination of true afterimage contrast and contrast sensitivity. Afterimage weakening following perceptual suppression of the inducer has previously been reported using a subjective duration measure (Gilroy & Blake, 2005) as well as using a subjective intensity measure (Tsuchiya & Koch, 2005). We found that the effect on perceived duration reverses with spatial frequency (Figure 4) and wished to test whether the same applies to subjective strength.

Our trial sequence followed that of Tsuchiya and Koch (2005). It is illustrated in Figure A3A. Two inducing gratings were presented on opposite sides of fixation, one being perceptually suppressed whereas the other one remained visible. Following a 4-sec adaptation period, observers reported which afterimage appeared stronger: the one left by the visible inducer or the one left by the invisible inducer. As in our main experiment, we used inducing gratings with various spatial frequencies.

Figure A3B, left panel, shows a typical response pattern, obtained by systematically varying the contrast of the visible inducer (the contrast of the perceptually suppressed inducer remained fixed at 0.52 Michelson). The proportion of trials that the afterimage of the visible inducer was judged as stronger increased with the contrast of the visible inducer. Our variable of interest is this proportion for the situation where the visible inducer and the perceptually suppressed inducer have physically equal contrasts.

This proportion for physically equal inducers is shown in the right panel of Figure A3B, averaged across seven observers (four naive; three also participated in our main experiment). For inducers of low spatial frequency, the fraction lies above 0.5 (two-sided t test, $df = 6$, $t = 17.5$, $p \ll 0.01$ for the 0.2 cycle/degree condition), indicating that perceptual suppression of the inducer subjectively weakens the resulting afterimage. At higher spatial frequencies, however, this effect reverses (two-sided t test, $df = 6$, $t = 2.8$, $p < .05$ for the 3.0 cycle/degree condition; intermediate condition not significant), showing that afterimages of perceptually suppressed inducers are subjectively stronger. The same reversal we observed for perceived duration (Figure 4B) thus applies to subjective intensity as well.

Given that subjective afterimage intensity depends on suprathreshold contrast sensitivity rather than on the level of the contrast detection threshold itself, we cannot provide a direct prediction of subjective afterimage intensity by combining measurements of true afterimage contrast and contrast detection threshold, as we could in the case of perceived afterimage duration (Figure 4). Nevertheless, subjective afterimage intensity is comparable to perceived afterimage duration in the respect that it depends on a combination of true afterimage contrast and contrast sensitivity. One implication of this is that adaptation to the suppressor pattern, an important factor explaining why perceptual suppression reduces perceived afterimage duration at low spatial frequencies (Figure 3), is also a likely contributor to the reduction of subjective afterimage intensity that we observe at low spatial frequencies in the present experiment (Figure A3).

Acknowledgments

This work was supported by Rubicon grants from the Netherlands Organisation for Scientific Research (J. B. and J. v. B.) and NIH EY13358 (R. B.).

Reprint requests should be sent to Jan W. Brascamp, Vanderbilt University, Department of Psychology, PMB 407817, 2301 Vanderbilt Place, Nashville, TN 37240, or via e-mail: j.w.brascamp@gmail.com.

REFERENCES

- Alonso, J., & Martinez, L. (1998). Functional connectivity between simple cells and complex cells in cat striate cortex. *Nature Neuroscience*, 1, 395–403.
- Anstis, S., Rogers, B., & Henry, J. (1978). Interactions between simultaneous contrast and coloured afterimages. *Vision Research*, 18, 899–911.
- Bahrami, B., Lavie, N., & Rees, G. (2007). Attentional load modulates responses of human primary visual cortex to invisible stimuli. *Current Biology*, 17, 509.
- Björklund, R., & Magnussen, S. (1981). A study of interocular transfer of spatial adaptation. *Perception*, 10, 511–518.
- Blake, R., Tadin, D., Sobel, K. V., Raissian, T. A., & Chong, S. C. (2006). Strength of early visual adaptation depends on visual awareness. *Proceedings of the National Academy of Sciences, U.S.A.*, 103, 4783–4788.
- Blakemore, C., & Campbell, F. (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology*, 203, 237–260.
- Braun, J., & Julesz, B. (1998). Withdrawing attention at little or no cost: Detection and discrimination tasks. *Perception & Psychophysics*, 60, 1–23.
- Brindley, G. (1962). Two new properties of foveal after-images and a photochemical hypothesis to explain them. *Journal of Physiology*, 164, 168–179.
- Burbeck, C. (1986). Negative afterimages and photopic luminance adaptation in human vision. *Journal of the Optical Society of America*, 3, 1159–1165.
- Craik, K. (1940). Origin of visual after-images. *Nature*, 145, 512.
- Crick, F., & Koch, C. (1998). Consciousness and neuroscience. *Cerebral Cortex*, 8, 97.
- Crowder, N., Price, N., Hietanen, M., Dreher, B., Clifford, C., & Ibbotson, M. (2006). Relationship between contrast

- adaptation and orientation tuning in V1 and V2 of cat visual cortex. *Journal of Neurophysiology*, 95, 271.
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences*, 10, 204–211.
- Driver, J., & Vuilleumier, P. (2001). Perceptual awareness and its loss in unilateral neglect and extinction. *Cognition*, 79, 39–88.
- Enroth-Cugell, C., & Robson, J. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *Journal of Physiology*, 187, 517–552.
- Ferster, D., & Lindström, S. (1983). An intracellular analysis of geniculocortical connectivity in area 17 of the cat. *Journal of Physiology*, 342, 181–215.
- Fries, P., Schröder, J.-H., Roelfsema, P., Singer, W., & Engel, A. (2002). Oscillatory neuronal synchronization in primary visual cortex as a correlate of stimulus selection. *Journal of Neuroscience*, 22, 3739–3754.
- Georgeson, M., & Turner, R. (1985). Afterimages of sinusoidal, square-wave and compound gratings. *Vision Research*, 25, 1709–1720.
- Gilroy, L., & Blake, R. (2005). The interaction between binocular rivalry and negative afterimages. *Current Biology*, 15, 1740–1744.
- Glantz, S. A. (2005). *Primer of biostatistics* (6th ed.). New York: McGraw-Hill.
- Green, D. M., & Swets, J. A. (1988). *Signal detection theory and psychophysics*. New York: John Wiley & Sons Ltd.
- Grunewald, A., Bradley, D., & Andersen, R. (2002). Neural correlates of structure-from-motion perception in macaque V1 and MT. *Journal of Neuroscience*, 22, 6195–6207.
- Haynes, J.-D., Deichmann, R., & Rees, G. (2005). Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus. *Nature*, 438, 496–499.
- Hubel, D., & Wiesel, T. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, 160, 106–154.
- Hubel, D., & Wiesel, T. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 195, 215–243.
- Ito, M., & Gilbert, C. (1999). Attention modulates contextual influences in the primary visual cortex of alert monkeys. *Neuron*, 22, 593–604.
- Kanai, R., Tsuchiya, N., & Verstraten, F. (2006). The scope and limits of top-down attention in unconscious visual processing. *Current Biology*, 16, 2332–2336.
- Kanwisher, N., & Wojciulik, E. (2000). Visual attention: Insights from brain imaging. *Nature Reviews Neuroscience*, 1, 91–100.
- Kelly, D., & Martinez-Uriegas, E. (1993). Measurements of chromatic and achromatic afterimages. *Journal of the Optical Society of America. A*, 10, 29–37.
- Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: Two distinct brain processes. *Trends in Cognitive Sciences*, 11, 16–22.
- Kohn, A. (2007). Visual adaptation: Physiology, mechanisms, and functional benefits. *Journal of Neurophysiology*, 97, 3155–3164.
- Lamme, V. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, 7, 12–18.
- Lee, S.-H., Blake, R., & Heeger, D. (2004). Traveling waves of activity in primary visual cortex during binocular rivalry. *Nature Neuroscience*, 8, 22–23.
- Leguire, L., & Blake, R. (1982). Role of threshold in afterimage visibility. *Journal of the Optical Society of America*, 72, 1232–1237.
- Lehky, S., & Maunsell, J. (1996). No binocular rivalry in the LGN of alert macaque monkeys. *Vision Research*, 36, 1225–1234.
- Leopold, D., & Logothetis, N. (1996). Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature*, 379, 549–553.
- Levitt, J., Kiper, D., & Movshon, J. (1994). Receptive fields and functional architecture of macaque V2. *Journal of Neurophysiology*, 71, 2517–2542.
- Levitt, J., Schumner, R., Sherman, S., Spear, P., & Movshon, J. (2001). Visual response properties of neurons in the LGN of normally reared and visually deprived macaque monkeys. *Journal of Neurophysiology*, 85, 2111.
- Logothetis, N., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412, 150–157.
- Logothetis, N., & Schall, J. (1989). Neuronal correlates of subjective visual perception. *Science*, 245, 761–763.
- Lou, L. (2001). Effects of voluntary attention on structured afterimages. *Perception*, 30, 1439–1448.
- Mack, A., & Rock, I. (1998). Inattention blindness: Perception without attention. In R. D. Wright (Ed.), *Visual attention* (pp. 55–76). Oxford: Oxford University Press.
- Maier, A., Wilke, M., Aura, C., Zhu, C., Ye, F., & Leopold, D. (2008). Divergence of fMRI and neural signals in V1 during perceptual suppression in the awake monkey. *Nature Neuroscience*, 11, 1193–1200.
- Maunsell, J., & Van Essen, D. (1983). Functional properties of neurons in middle temporal visual area of the macaque monkey: II. Binocular interactions and sensitivity to binocular disparity. *Journal of Neurophysiology*, 49, 1148–1167.
- McAdams, C., & Reid, R. (2005). Attention modulates the responses of simple cells in monkey primary visual cortex. *Journal of Neuroscience*, 25, 11023–11033.
- Mcalonan, K., Cavanaugh, J., & Wurtz, R. (2008). Guarding the gateway to cortex with attention in visual thalamus. *Nature*, 456, 391.
- Montaser-Kouhsari, L., & Rajimehr, R. (2005). Subliminal attentional modulation in crowding condition. *Vision Research*, 45, 839–844.
- Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, 229, 782–784.
- Motter, B. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *Journal of Neurophysiology*, 70, 909–919.
- Movshon, J., & Lennie, P. (1979). Pattern-selective adaptation in visual cortical neurones. *Nature*, 278, 850–852.
- Movshon, J., Thompson, I., & Tolhurst, D. (1978). Receptive field organization of complex cells in the cat's striate cortex. *Journal of Physiology*, 283, 79.
- Polonsky, A., Blake, R., Braun, J., & Heeger, D. (2000). Neuronal activity in human primary visual cortex correlates with perception during binocular rivalry. *Nature Neuroscience*, 3, 1153–1159.
- Rees, G., & Lavie, N. (2001). What can functional imaging reveal about the role of attention in visual awareness? *Neuropsychologia*, 39, 1343–1353.
- Rock, I., Linnett, C., Grant, P., & Mack, A. (1992). Perception without attention: Results of a new method. *Cognitive Psychology*, 24, 502–534.
- Roelfsema, P., Lamme, V., & Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. *Nature*, 395, 376–381.
- Rust, N., Schwartz, O., Movshon, J., & Simoncelli, E. (2005). Spatiotemporal elements of macaque v1 receptive fields. *Neuron*, 46, 945–956.

- Sakitt, B. (1976). Psychophysical correlates of photoreceptor activity. *Vision Research*, 16, 129–140.
- Schwartz, O., Hsu, & Dayan, P. (2007). Space and time in visual context. *Nature Reviews Neuroscience*, 8, 522–535.
- Scialo, G., Lennie, P., & DePriest, D. (1989). Contrast adaptation in striate cortex of macaque. *Vision Research*, 29, 747–755.
- Sengpiel, F., Blakemore, C., & Harrad, R. (1995). Interocular suppression in the primary visual cortex: A possible neural basis of binocular rivalry. *Vision Research*, 35, 179–195.
- Shimojo, S., Kamitani, Y., & Nishida, S. (2001). Afterimage of perceptually filled-in surface. *Science*, 293, 1677–1680.
- Shin, K., Stolte, M., & Chong, S. (2009). The effect of spatial attention on invisible stimuli. *Attention, Perception and Psychophysics*, 71, 1507–1513.
- Simons, D., & Chabris, C. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28, 1059–1074.
- Snowden, R., & Hammett, S. (1996). Spatial frequency adaptation: Threshold elevation and perceived contrast. *Vision Research*, 36, 1797–1809.
- Suzuki, S., & Grabowecky, M. (2003). Attention during adaptation weakens negative afterimages. *Journal of Experimental Psychology: Human Perception and Performance*, 29, 793–807.
- Tong, F. (2003). Primary visual cortex and visual awareness. *Nature Reviews Neuroscience*, 4, 219–229.
- Tong, F., & Engel, S. (2001). Interocular rivalry revealed in the human cortical blind-spot representation. *Nature*, 411, 195–199.
- Treue, S., & Maunsell, J. (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature*, 382, 539–541.
- Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience*, 8, 1096–1101.
- Van Boxtel, J. J., & Koch, C. (2009). Afterimage duration and its modulation by attention and consciousness. *Journal of Vision*, 9, 121.
- Virsu, V., & Laurinen, P. (1977). Long-lasting afterimages caused by neural adaptation. *Vision Research*, 17, 853–860.
- Viswanathan, A., & Freeman, R. (2007). Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. *Nature Neuroscience*, 10, 1308–1312.
- Wede, J., & Francis, G. (2007). Attentional effects on afterimages: Theory and data. *Vision Research*, 47, 2249–2258.
- White, A., Sun, H., Swanson, W., & Lee, B. (2002). An examination of physiological mechanisms underlying the frequency-doubling illusion. *Investigative Ophthalmology & Visual Science*, 43, 3590–3599.
- Wilke, M., Logothetis, N., & Leopold, D. (2006). Local field potential reflects perceptual suppression in monkey visual cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, 103, 17507–17512.
- Wunderlich, K., Schneider, K., & Kastner, S. (2005). Neural correlates of binocular rivalry in the human lateral geniculate nucleus. *Nature Neuroscience*, 8, 1595–1602.
- Xu, X., Ichida, J., Allison, J., Boyd, J., Bonds, A., & Casagrande, V. (2001). A comparison of koniocellular, magnocellular and parvocellular receptive field properties in the lateral geniculate nucleus of the owl monkey (*Aotus trivirgatus*). *Journal of Physiology*, 531, 203–218.