Behavioral/Cognitive

Different Signal Enhancement Pathways of Attention and Consciousness Underlie Perception in Humans

Jeroen J.A. van Boxtel

School of Psychological Sciences and Monash Institute of Cognitive and Clinical Neurosciences, Monash University, Clayton 3800 Victoria, Australia, and Division of Biology, California Institute of Technology, Pasadena, California 91125

It is not yet known whether attention and consciousness operate through similar or largely different mechanisms. Visual processing mechanisms are routinely characterized by measuring contrast response functions (CRFs). In this report, behavioral CRFs were obtained in humans (both males and females) by measuring afterimage durations over the entire range of inducer stimulus contrasts to reveal visual mechanisms behind attention and consciousness. Deviations relative to the standard CRF, i.e., gain functions, describe the strength of signal enhancement, which were assessed for both changes due to attentional task and conscious perception. It was found that attention displayed a response-gain function, whereas consciousness displayed a contrast-gain function. Through model comparisons, which only included contrast-gain modulations, both contrast-gain and response-gain effects can be explained with a two-level normalization model, in which consciousness affects only the first level and attention affects only the second level. These results demonstrate that attention and consciousness can effectively show different gain functions because they operate through different signal enhancement mechanisms.

Key words: afterimages; attention; consciousness; contrast response function

---

Significance Statement

The relationship between attention and consciousness is still debated. Mapping contrast response functions (CRFs) has allowed (neuro)scientists to gain important insights into the mechanistic underpinnings of visual processing. Here, the influence of both attention and consciousness on these functions were measured and they displayed a strong dissociation. First, attention lowered CRFs, whereas consciousness raised them. Second, attention manifests itself as a response-gain function, whereas consciousness manifests itself as a contrast-gain function. Extensive model comparisons show that these results are best explained in a two-level normalization model in which consciousness affects only the first level, whereas attention affects only the second level. These findings show dissociations between both the computational mechanisms behind attention and consciousness and the perceptual consequences that they induce.

Introduction

Attention and consciousness are both major cognitive functions that determine visual processing. The relationship between attention and consciousness and their effects on our visual perception have been strongly debated (Wundt, 1874; Iwasaki, 1993; Posner, 1994; Chun and Wolfe, 2000; Dehaene et al., 2006; Koch and Tsuchiya, 2007; Mole, 2008; De Brigard and Prinz, 2010; van

Boxtel et al., 2010a; Cohen et al., 2012; Koch and Tsuchiya, 2012; Prinz, 2012). A commonly held belief is that a tight link exists between attention and consciousness (Dehaene et al., 2006; De Brigard and Prinz, 2010; Cohen et al., 2012; Prinz, 2012). However, other researchers believe that attention and consciousness can be separated experimentally (Kentridge et al., 2004; Wyart and Tallon-Baudry, 2008; Brascamp et al., 2010; van Boxtel et al., 2010b; Watanabe et al., 2011). It should be noted that "attention" here refers to endogenous directed focal attention, not, for example, exogenous attention, and "consciousness" refers to perception as gauged by reported visibility and manipulated through masking.

Evidence that attention and consciousness are separable would find the strongest support from a double dissociation between these two functions (Koch and Tsuchiya, 2007; van Boxtel et al., 2010a; Cohen et al., 2012); that is, when the effects of attention and consciousness go in opposite directions. A rare behavioral example of such a double dissociation comes from a psychophysical paradigm that investigated how afterimage durations are affected

by attention and consciousness (van Boxtel et al., 2010b). It was shown that consciousness increased afterimage durations, whereas attention decreased afterimage durations.

Although this previous research shows strong support for a dissociation between attention and consciousness, it does not provide insight into how this dissociation comes about. To gain such mechanistic understanding, researchers often map out contrast response functions (CRFs) (Albrecht and Hamilton, 1982; Reynolds and Heeger, 2009; Carrasco, 2011). Therefore, to determine whether the effects of attention and consciousness are subserved by similar or different mechanisms, afterimage durations were measured over a range of contrasts to determine behavioral CRFs and investigate the modulations of this curve caused by attention and consciousness.

CRFs of neuronal responses (Albrecht and Hamilton, 1982), as well as behavioral performance (Carrasco, 2011), generally increase in a sigmoidal fashion with increasing stimulus contrast (see Fig. 1A). Attention generally enhances contrast responses, but it can do so in different ways (Reynolds and Heeger, 2009; Carrasco, 2011). For example, attention can shift the contrast response curve to the left, thereby effectively boosting stimulus contrast (Reynolds and Heeger, 2009; Carrasco, 2011). This response modulation is called a contrast gain function and shows the largest effects at intermediate contrasts (Fig. 1A,B). Attention has also been reported to modulate responses multiplicatively. This type of function is called a response gain function, which shows the largest effects at high contrasts (Fig. 1A,B). The influence of consciousness on CRFs is not yet known, although one of the techniques to modulate conscious perception, interocular suppression, often, but not always, produces contrast gain functions (Sengpiel et al., 1998; Watanabe et al., 2004; Li et al., 2005; Bahrami et al., 2008; Yuval-Greenberg and Heeger, 2013).

Importantly, no study has measured and controlled for the effects of both attention and consciousness on CRFs. It is therefore possible that the previously reported effects in CRFs are a mixture of both attention and consciousness. Here, a technique is used that strictly controlled both attention and consciousness modulations, allowing the study of their separate effects on CRFs. Gain functions for both attention and consciousness were thus independently determined to study whether they operate through the same or distinct signal enhancement mechanisms.

Materials and Methods

Experimental design and statistical analysis
The same stimuli and paradigm were used as described previously (van Boxtel et al., 2010b).

Participants
Low spatial frequency (SF). Data were obtained from 17 individuals (nine male/eight female college students). Based on trial inclusion criteria (see below), two participants were excluded because they had missing data in more than two parameter combinations. Therefore, the data from 15 individuals (eight males) were analyzed further.

High SF. Nineteen participants participated. Three participants were excluded because of missing data. Therefore, data from 16 individuals (6 male, 10 female, mean age 20 years/1 month, SD = 1 year/4 months) were analyzed further.

Stimuli
The afterimage inducer was a Gabor patch with a contrast that was drawn from 0.03, 0.06, 0.125, 0.25, 0.50, and 1. The patch was Gaussian windowed ($\sigma = 1.43°$) and had an SF of 0.23 cycles/° or 3 cycles/°, a random orientation, and was presented at 4.9° eccentricity. The mask, which was shown on half of the trials, was of 100% contrast rotated at 120°/s and consisted of a Gaussian windowed ($\sigma = 1.43°$) checkerboard (0.78 cycles/°).

It reversed contrast every 67 ms. Presentation location of the inducer and the mask was shifted by 45° counterclockwise between trials. Background luminance was 49 cd/m². All experiments were performed on a gamma-corrected monitor.

In the low-attention conditions, the attentional task was a rapid serial visual presentation (RSVP) of red letters (font Helvetica, 12 point). These letters were shown for 133 ms, after which they were immediately replaced by the next letter.

Procedure
Each trial had three phases. The first phase was the adaptation phase during which the afterimage inducer was shown in one eye (left and right eye presentation was counterbalanced over trials). To the other eye, the mask was shown in invisible trials, whereas no mask was shown in visible trials. The adaptation phase lasted 4 s. To create conditions in which a low amount of attention was paid to the inducer, subjects were distracted by a RSVP in which they counted the number of X's ($n = 2$–5; participants were told there could be 1–5 in any trial) that appeared in an RSVP stream of nontarget letters (randomly chosen from: M, S, T, A, B, C, D, O, K, P, Y). Subjects did not report visibility in the low-attention trials, avoiding the need to deploy attention to the inducer. In the high-attention trials, the RSVP task was not performed but the letters were shown. Instead, subjects tracked the subjective visibility of the inducer by pressing and releasing a keyboard button.

The second phase was the afterimage phase. Participants pressed a button as soon as they perceived an afterimage and released the button when the afterimage disappeared. Because afterimages in this experiment were perceived instantaneously, afterimage duration was recorded from the start of the afterimage phase until the button was released. Participants pressed the space bar if no afterimage was perceived (this was recorded as a 0 s afterimage duration).

In the third phase, the observer was asked to indicate the number of X's that were counted. The number was indicated with the keypad. In the high-attention condition, this question was skipped by pressing the space bar. All trials were presented in a pseudorandom order and divided over eight blocks separated by brief rest periods.

Data analysis
For the data analysis, trials were excluded when the reported number of X's was off by >1 (except where noted differently). For the trials with the mask, which should lead to invisibility of the inducer stimulus, trials that were nonetheless reported as visible for any length of time (i.e., they broke through suppression) were excluded. These selection criteria led to the inclusion of the following number of trials (mean ± SEM over participants): in the low SF condition: 47.53 ± 0.47 (Visible–High Attention), 35.53 ± 1.40 (Visible–Low Attention), 41.93 ± 2.63 (Invisible–High Attention), 36.87 ± 0.97 (Invisible–Low Attention); and in the high SF condition: 47.58 ± 0.34 (Visible–High Attention), 32.74 ± 1.76 (Visible–Low Attention), 28.16 ± 3.65 (Invisible–High Attention), 32.79 ± 2.16 (Invisible–Low Attention). The lowest mean number of trials of any of the contrast conditions was as follows: 5.67 ± 0.33 (Visible–Low Attention, contrast = 0.06) for the low SF condition and 4.0 ± 0.68 (Invisible–Low Attention, contrast = 0.5) for the high SF condition.

Calculating psychophysical gain functions
The afterimage durations in the condition with invisible trials that were not attended (Att⁻/Vis⁻) were taken as a baseline measure, as this condition has a maximally reduced influence of attention and conscious visibility in our design. The attention-induced gain function was then calculated by subtracting, per participant, this baseline from the attended, but invisible conditions (Att⁺/Vis⁻): $\Delta AI = Att^+/Vis^- - Att^-/Vis^-$. The visibility-induced gain function was similarly calculated by subtracting Att⁻/Vis⁻ from the visible but less-attended condition: $\Delta AI = Att^-/Vis^+ - Att^-/Vis^-$. ΔAI-measures were averaged over participants, and calculated per contrast value.

Contrast-gain index (CGI) and response gain index (RGI)
CGI and RGI were calculated per individual. The CGI was calculated as the difference between the average response at the middle two contrasts minus the average response at the lowest and highest contrast, as follows:

$0.5 \times ((\Delta AI_{0.13} + \Delta AI_{0.25}) - (\Delta AI_{0.03} + \Delta AI_1))$, where ΔAI_c refers to ΔAI at contrast c. The RGI was calculated as the average response at the highest two contrasts minus the average response at the lowest two contrasts. It was then multiplied by -1 to yield a positive index, as follows: $-0.5 \times ((\Delta AI_{0.50} + \Delta AI_1) - (\Delta AI_{0.03} + \Delta AI_{0.06}))$. Two subjects missed data for some of the ΔAI_c values; these participants were removed for this analysis. CGI and RGI values were analyzed with one-sample t tests versus 0. Comparisons between experiments for both CGI and RGI were made using two-tailed paired t tests.

Z-transformed partial correlations and descriptive model fitting
Both the attention and visibility ΔAI data were fit with predictions from a contrast gain and response gain model. For the contrast gain model, the contrast-response curve change is characterized by a shift to the left or right. For a response gain model, the change is characterized by a multiplicative increase in response. Therefore, using the NonLinearModel class in MATLAB (2014b, 2015b) with default parameters, the average ΔAI data were fitted with the following functions:

$$\text{Contrast gain model}: \frac{R_{max} * \gamma C^n}{\gamma C^n + C_{50}^n} - \frac{R_{max} * C^n}{C^n + C_{50}^n} \quad (1)$$

$$\text{Response gain model}: \frac{\gamma * R_{max} * C^n}{C^n + C_{50}^n} \quad (2)$$

Where R_{max} is the maximum obtainable response, C is the contrast of the inducer, n is an exponent that determines the steepness of the curve, and C_{50} is the contrast at which half-maximum response is reached. The parameter γ is the gain of the process under scrutiny; that is, attention or visibility. For both high and low SF experiments, we fixed all parameters, except γ. R_{max} was set to 150% of the maximum AI duration at 100% contrast, C_{50} to 0.25, and the exponent n was set to 2 (Herrmann et al., 2010).

To remove correlations between the contrast gain and response gain models themselves, partial correlations were computed (Movshon et al., 1985; Smith et al., 2005). The partial correlations for contrast-gain (R_c) and response-gain (R_r) models are as follows:

$$R_c = \frac{r_c - r_c r_{cr}}{\sqrt{(1 - r_c^2)(1 - r_{cr}^2)}} \quad (3)$$

and

$$R_r = \frac{r_r - r_r r_{cr}}{\sqrt{(1 - r_r^2)(1 - r_{cr}^2)}} \quad (4)$$

where r_c and r_r are the correlations between the ΔAI data and the contrast-gain and response-gain models, respectively, and r_{cr} is the correlation between the two fitted gain models.

Pearson's r correlations are not normally distributed, so the Fisher r-to-Z transformation (shown just for R_c) was used as follows:

$$Z_c = \frac{0.5 \frac{(1 + R_c)}{(1 - R_c)}}{\sqrt{\frac{1}{df}}} \quad (5)$$

where df is the degrees of freedom, which is equal to the number of contrast values measured minus 3. A Z larger than 1.65 (equal to $p = 0.05$, one-tailed) was taken as the threshold for significance.

Model fitting
Several models were considered as potentially being able to explain our data. The models that were considered had a very similar two-level architecture, with a level 1 (L_1) activity feeding into a L_2. Both levels represent the neural activity of that level, with L_1 representing a contrast-polarity sensitive level and L_2 representing a contrast-polarity-insensitive level. Both levels include adaptation. The layout of all models was as follows:

$$\text{Resp}_{L1}(C) = F_{L1} \quad (6)$$

$$\text{Resp}_{L2}(C) = F_{L2} \quad (7)$$

$$S_p = 1 - m \times \text{Resp}_{L2}(\text{Resp}_{L1}) \quad (8)$$

$$\text{Resp}_{Adapt} = s \times S_p \times \text{Resp}_{L2}(\text{Resp}_{L1}) + O \quad (9)$$

where $\text{Resp}L_1$ is the response at L_1; $\text{Resp}L_2$ is the response at L_2 and takes the output of L_1 as input; FL_1 and FL_2 describe the contrast-response functions of L_1 and L_2, which are explained below for each of the models separately; and S_p describes the sensitivity postadaptation at L_2, which is dependent on the preadaptation activity of this level. The parameter m determines the strength of adaptation at L_2. Note that a stronger adaptation at the second level lowers the sensitivity S_p which is analogous to increasing the contrast detection threshold (cf. Brascamp et al., 2010). The final response of the model is Resp_{Adapt}, which is scaled (with s) and offset (with O) to best fit the data. It was assumed that the afterimage duration is directly related to this final response (i.e., a larger response means a longer time to decay).

Note that the activity in L_1 is also dependent on adaptation, but this is only implicitly present in the model. Specifically, it was assumed that there exist two populations of neurons at L_1, a population L_1^+ that is responding to the presented contrast polarity of the inducer and a population L_1^- that is not. Without adaptation, the responses of both populations in L_1 would lead to an unbiased response to a gray screen: $L_1^- - L_1^+ = 0$. However, adaptation at L_1 leads to a decreased sensitivity for cells in L_1^+, but not in cells that are nonresponsive to the stimulus (L_1^-). Therefore, after adaptation, there is now a biased response ($L_1^- - L_1^+ = \text{Resp}L_1 > 0$), causing the afterimage.

Models were fitted using the fmincon function in MATLAB. The models were simultaneously fit to the four contrast-response functions, as well as to the two ΔAI curves (the consciousness and attention gain functions). Both upper and lower bounds were provided on all parameters. Lower and upper bounds for the various parameters were as follows: n [lower = 0, upper = 3], C_{50} [0, 1], m [0, 1], s [0, 100], O [0, 2]. The various multiplicative attention terms (A) were bounded as follows [1, 100], with 1 being no attentional effect. When the parameter C_{cfs} was free to vary, it was limited to [0, 1].

In all models except the full model, C_{cfs} was set to 0 when no CFS stimulus was presented and 1 when a CFS stimulus was presented. In the full model, C_{cfs} was a free parameter when the CFS stimulus was presented (independently at L_1 and at L_2). In all models, multiplicative attention (A) parameters were set to 1 in low-attention conditions and were free to vary in the high-attention condition.

Contrast-response functions for the different models
Monocular normalization only model. The monocular normalization-only model was the base model against which other models were compared. This model consists of two stages, each with a normalization operation but without attention and consciousness manipulation. Figure 4B shows the model layout. The CRFs, which include the normalization step, for L_1 and L_2 are as follows:

$$F_{L1} = C^n/(C^n + C_{50}^n) \quad (10)$$

$$F_{L2} = C^n/(C^n + C_{50}^n) \quad (11)$$

This model has five free parameters: n, C_{50}, m, s, and O. Parameters m, s, and O are as explained above. The parameter n controls the steepness of the tuning curves and it is the same for L_1 and L_2. C_{50} determines the point where the tuning curve reaches half the maximum height; it is also the same for L_1 and L_2. The parameters n, C_{50}, s, m, and O are free parameters in all models.

Attention-first model. The attention-first (and consciousness second) model includes the influence of attention at L_1 and the influence of the CFS stimulus at L_2 (see model architecture in Fig. 4B). The CRFs for L_1 and L_2 are as follows:

$$F_{L1} = A_1 C^n/(A_1 C^n + C_{50}^n) \quad (12)$$

$$F_{L2} = C^n/(C^n + C_{CFS}^n + C_{50}^n) \quad (13)$$

A Contrast response functions **C**

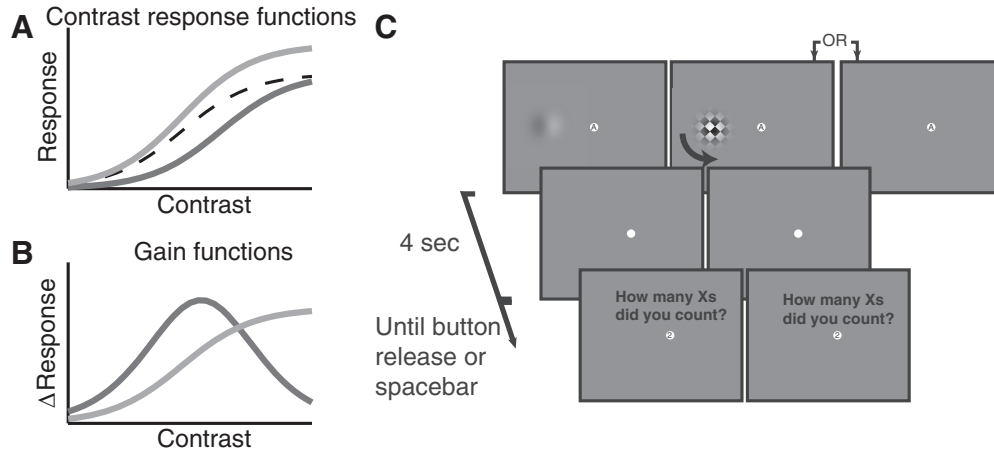**B** Gain functions

4 sec

Until button release or spacebar

Figure 1. **A**, Schematic illustration of contrast-gain (dark solid curve) and response-gain (light solid curve) effects on a CRF (dashed black). A contrast gain is evident as a wholesale rightward shift of the curve, whereas a response gain is evident as a multiplication of the response. **B**, Gain functions. Taking the difference between the dashed line and gain-modulated responses will result in a signature profile for both contrast and response gain, with contrast gain (dark solid curve) identifiable by a bell-shaped curve with the largest effects at intermediate contrasts and response gain (light solid curve) identifiable by a sigmoidal curve with the largest effects at high contrasts. **C**, Schematic of our experimental paradigm in which attention was modulated with an RSVP task at fixation (white dot) and consciousness was modulated by presenting (or not presenting) an interocular mask (see first row, middle, and right).

where A_1 is the multiplicative influence of attention at L_1 and C_{CFS} is the contrast of the CFS stimulus (set to 0 when the CFS stimulus was absent and to 1 when the CFS stimulus was present).

Consciousness-first model. The consciousness-first (and attention second) model includes the influence of CFS at L_1 and the influence of attention at L_2 (see model architecture in Fig. 4B). The CRFs for L_1 and L_2 are as follows:

$$F_{L1} = C^n/(C^n + C_{CFS}^n + C_{50}^n) \qquad (14)$$

$$F_{L2} = A_2 C^n/(A_2 C^n + C_{50}^n) \qquad (15)$$

where A_2 is the multiplicative influence of attention at L_2.

Attention-first model with response gain for attention. In this model, the CRFs for L_1 and L_2 are identical to the attention-first model apart from excluding the attentional effects in the denominator (thus resulting in an attention-induced response-gain) as follows:

$$F_{L1} = A_1 C^n/(C^n + C_{50}^n) \qquad (16)$$

$$F_{L2} = C^n/(C^n + C_{CFS}^n + C_{50}^n) \qquad (17)$$

Parameter definitions are as in the attention-first model.

Consciousness-first model with response gain for attention. In this model, the CRFs for L_1 and L_2 are identical to the consciousness-first model apart from excluding the attentional effects in the denominator (thus resulting in an attention-induced response-gain) as follows:

$$F_{L1} = C^n/(C^n + C_{CFS}^n + C_{50}^n) \qquad (18)$$

$$F_{L2} = A_2 C^n/(C^n + C_{50}^n) \qquad (19)$$

Parameter definitions are as in the consciousness-first model.

Full model. The full model includes the influence of CFS and attention at L_1 and at L_2 (see model architecture in Fig. 4B). Attention is implemented as spatial attention, and therefore boosts both the activity to the afterimage inducer (C), and the CFS stimulus (CFS). The CRFs for L_1 and L_2 are as follows:

$$F_{L1} = A_1 C^n/(A_1 C^n + A_1 C_{CFS1}^n + C_{50}^n) \qquad (20)$$

$$F_{L2} = A_2 C^n/(A_2 C^n + A_2 C_{CFS2}^n + C_{50}^n) \qquad (21)$$

Model comparison
The fits were compared using the Bayesian information criterion (BIC), in which a low BIC is better (a perfect fit leads to a BIC of $-\infty$). The BIC depends on the likelihood, which was calculated as follows:

$$L = -0.5 \cdot (df_e + N \cdot (\log(2\pi) + \log(MSE))) \qquad (22)$$

Where L is the likelihood, N is the number of data points, df_e is the degrees of freedom of the error ($= N$ minus the number of parameters that are estimated), and MSE equals the sum of squared errors of all the fit residuals divided by df_e.

As a baseline, we took the monocular normalization-only model (N) that has only monocular normalization, but no influence of attention and interocular interactions. Models were compared by calculating the difference in the BIC score (ΔBIC) relative to the monocular normalization-only model (N). According to Kass and Raftery (1995), ΔBIC values not worth more than a bare mention are $0 < abs(\Delta BIC) < 2$; those that show positive evidence are $2 < abs(\Delta BIC) < 6$; those that show strong evidence are $6 < abs(\Delta BIC) < 10$, and those that show very strong evidence are $abs(\Delta BIC) > 10$. Note that the ΔBIC thresholds of 2, 6, and 10 convert to the Bayes factors (BF_{10}) of 2.7, 20, and 148, respectively. Apart from the models discussed here, we also considered various other models, which are not discussed below, but our conclusions remain the same.

Results
Psychophysical data
We measured the complete CRFs under various levels of focused attention and consciousness (operationalized as visibility) while using a full-factorial design. This approach allowed us to investigate the independent influences of attention and consciousness on visual processing (Fig. 1C). Previous research has shown that, when the influences of attention and consciousness are not strictly controlled, it may be difficult to attribute experimental effects to either of these processes (Koch and Tsuchiya, 2007; van Boxtel et al., 2010a). Masking is a powerful technique to modulate conscious perception (Kim and Blake, 2005) and it allows one to separate the influences of attention from those of consciousness (Dehaene et al., 2006; Kanai et al., 2010; van Boxtel et al., 2010b). We therefore used the most versatile method to render stimuli invisible through the process of masking, a method called continuous flash suppression (CFS) (Tsuchiya and Koch, 2005). We paired the CFS technique with a stringent control of the participant's attention allocation. We minimized attention to the inducer stimulus in half of the trials by means of a distracting RSVP task (Rees et al., 1999). On other trials, participants reported the visibility of the inducer stimulus, which resulted in high levels of attention to the inducer stimulus. By using this
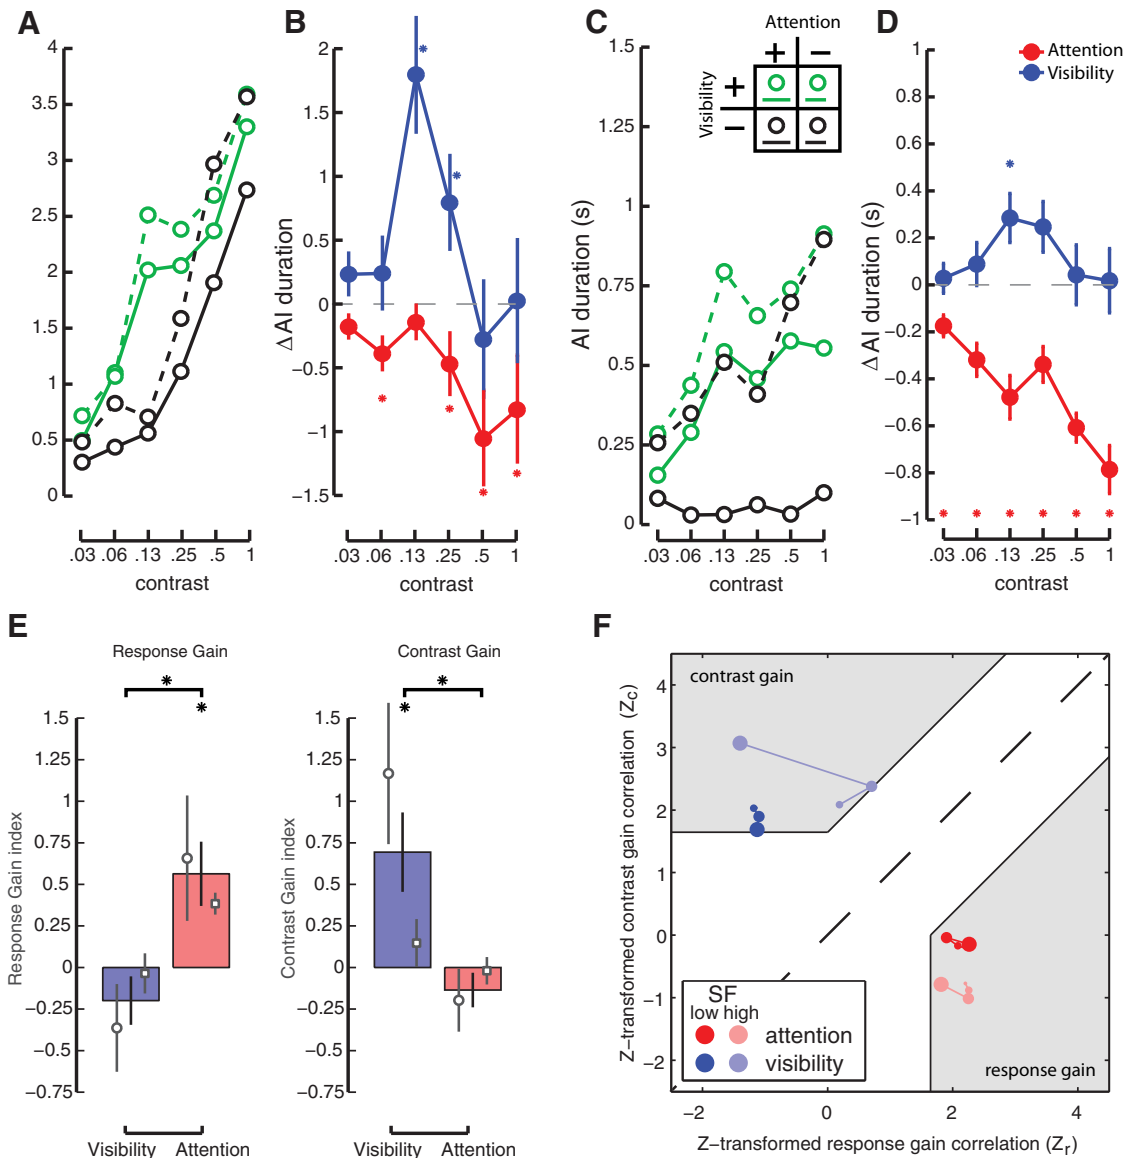
Figure 2. Afterimage durations in the 2 × 2 design. ***A***, Afterimage durations for low SF depend on the contrast of the afterimage inducer. Visible = green lines, invisible = black lines, high-attention = solid lines, and low-attention = dashed lines (see legend in ***C***). ***B***, Effects of attention (red curve), and visibility (blue curve) for low SF. These curves are the difference between high-attention/invisible and low-attention/invisible data (for attention), and low-attention/visible and low-attention/invisible data (for visibility). Asterisks indicate significant deviations from zero (one-tailed *t*-tests). ***C***, As in ***A*** but for high SF. ***D***, As in ***B*** for high SF. Note that the *y*-axes have a different scales in ***A–D***. ***E***, RGI and CGI. Attention shows a significant RGI, whereas visibility shows a significant CGI. Bars represent the average over all participants; circles are low SF data; and squares are high SF data. Asterisks indicate significant deviations from zero (two-tailed *t*-tests). ***F***, Z-transformed correlations between ΔAI duration data and contrast and response gain functions. The shaded area represents significance. The different symbols are data with different performance thresholds in the RSVP task. Larger dots correspond to stricter RSVP thresholds (and consequently fewer trials); large to small: 1, 2, 3, 4 items off correct. In all cases, the attention curve is best fit with a response gain function, whereas the visibility curve is best fit with a contrast gain function. Error bars indicate SEM.

design, attention and consciousness were independently modulated, resulting in a 2 × 2 matrix that consisted of four conditions: low-attention/invisible (i.e., Att⁻/Vis⁻), low-attention/visible (i.e., Att⁻/Vis⁺), high-attention/invisible (i.e., Att⁺/Vis⁻), and high-attention/visible (i.e., Att⁺/Vis⁺). This design allowed us to derive the individual contributions of attention and consciousness.

We determined psychometric CRFs, as measured by afterimage durations, for all four conditions (Fig. 2*A*,*C*). We looked at the influence of attention and consciousness separately relative to a baseline condition that lacked both of these influences. Therefore, to calculate the gain functions for attention and consciousness (Fig. 2*B*,*D*), we took the low-attention/invisible condition (i.e., black dashed lines, Att⁻/Vis⁻, Fig. 2*A*,*C*) as a baseline because this condition lacked influences of both attention and vis-

ibility. The gain function for attention was then calculated as the difference in AI duration (ΔAI duration) between the high-attention/invisible condition and this baseline (i.e., ΔAI = Att⁺/Vis⁻ − Att⁻/Vis⁻; Fig. 2*B*,*D*, red curve). Similarly, the gain function for visibility was then calculated as the ΔAI duration between this baseline and the low-attention/visible condition (i.e., ΔAI = Att⁻/Vis⁺ − Att⁻/Vis⁻; Fig. 2*B*,*D*, blue curve).

These results show strikingly different gain functions for attention and consciousness. Attention decreased afterimage durations (Fig. 2*B*,*D*, red curve), whereas visibility increased afterimage durations (Fig. 2*B*,*D*, blue curve) (Suzuki and Grabowecky, 2003; Tsuchiya and Koch, 2005; Brascamp et al., 2010; van Boxtel et al., 2010b). Furthermore, attention follows a response-gain function with small effects at low inducer contrasts and large significant

effects at high inducer contrasts (Fig. 2 B, D, red stars), whereas visibility follows a contrast-gain function, with significant effects only at intermediate inducer contrasts (Fig. 2 B, D, blue stars). We quantified these effects further by calculating for each subject a CGI and an RGI (see Materials and Methods; Fig. 2E). These data revealed that, over all subjects (and both experiments), attention had a significant RGI ($t_{(29)} = 2.92$, $p = 0.007$, Cohen's $d = 0.53$), but no significant CGI ($t_{(29)} = -1.30$, $p = 0.20$, Cohen's $d = -0.24$). Conversely, visibility showed a significant CGI ($t_{(29)} = 2.91$, $p = 0.007$, Cohen's $d = 0.53$), but no significant RGI ($t_{(29)} = -1.36$, $p = 0.18$, Cohen's $d = -0.25$). Similar findings are obtained for the low and high SF experiments separately in that there are positive RGIs for attention (low SF: $t_{(14)} = 1.7$, $p = 0.10$, Cohen's $d = 0.45$, ns; high SF: $t_{(13)} = 5.83$, $p < 0.0001$, Cohen's $d = 1.56$), and a positive CGI for visibility (low SF: $t_{(14)} = 2.75$, $p = 0.016$, Cohen's $d = 0.71$; high SF: $t_{(13)} = 1.0235$, $p = 0.32$, Cohen's $d = 0.27$, ns), although significance was not always reached. There was no significant difference between the RGI for attention for the low and high SF data ($t_{(27)} = 0.69$, $p = 0.50$, Cohen's $d = 0.26$). There was a significant difference between the CGI for attention for the low and high SF data ($t_{(27)} = 2.20$, $p = 0.036$, Cohen's $d = 0.83$). Nevertheless, the overall data suggest that attention and consciousness operate through different gain functions.

Interestingly, we find that attention and consciousness also have opposite effects at 3 cycles/°, which is inconsistent with a previous report (Brascamp et al., 2010). That report showed that attention decreased afterimage durations at low and high SFs, but that the effect of visibility depended on SF: visibility increased afterimage duration at low SFs (consistent with our findings), whereas it decreased the afterimage durations at higher (3 cycles/°) SFs (inconsistent with our findings). This discrepancy could potentially be explained by the use of different paradigms (i.e., afterimage nulling versus duration paradigms). However, an alternative explanation is related to the fact the previous study did not use a 2 × 2 design. For example, the attention effect was only measured in the visible conditions and the visibility effect was measured without having a demanding secondary task. Therefore, attention was not strictly controlled for in the measurements of the visibility effect. Without this task, participants potentially paid more attention to the stimulus when it was visible then when it was invisible.

Based on our data, we can estimate the results of this methodological difference. To approximate the effects of the previous study, we looked at the condition that combined the effects of attention and consciousness. Therefore, we computed the gain function based on the difference between our standard baseline condition (i.e., the invisible/low-attention condition, Att⁻/Vis⁻) and the condition with both high visibility and high attention (i.e., the visible/high-attention condition): Att⁺/Vis⁺ − Att⁻/Vis⁻. These data are plotted with open circles in Figure 3. We compared this with the gain function without attention confound (open circles in Fig. 3, replotted from Fig. 2).

We found that, at low SFs, the lack of a demanding secondary task (and the consequent potential for a confounding influence of attention) has no or little influence of the gain function for visibility. However, at higher frequencies, an increase of attention (due to poor attentional control) shifts the visibility curve downward, resulting in apparent negative effects of visibility at many contrasts (being significantly negative at the highest contrast; $t_{(15)} = -2.19$, $p = 0.045$, Cohen's $d = -0.55$). Therefore, this analysis suggests that the negative effects of visibility reported previously (Brascamp et al., 2010), measured at a contrast of 0.62, are potentially due to the lack of a stringent control of attention.
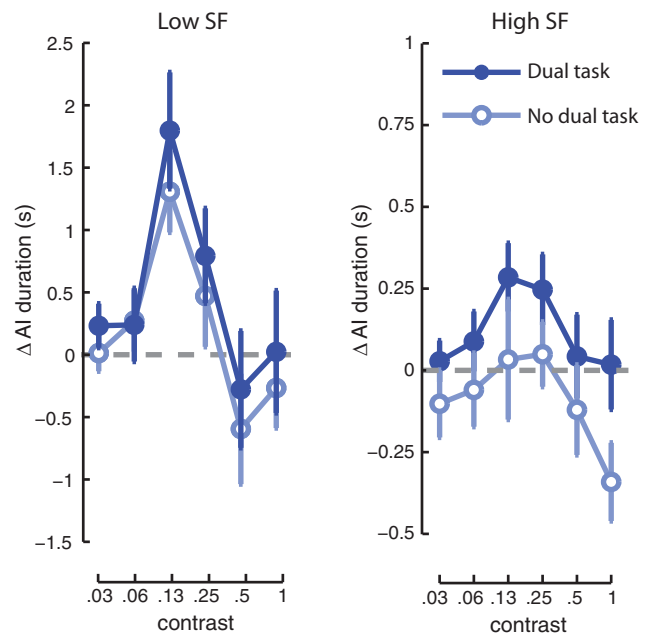

Figure 3. Gain functions for visibility in conditions without attention (filled symbols) and those with attention (open symbols). Left, In low SF conditions, the presence of attention in addition to visibility does not appear to alter the gain function. Right, In the high SF condition, the addition of attention shortens the afterimage durations overall, pushing the ΔAI measure below zero (significantly so at inducer contrasts of 100%). Note that the y-axes have different scales in the left and right panel. Error bars indicate SEM over participants.

Descriptive model fitting

To investigate whether our measured gain functions can be described by pure contrast-gain and response-gain functions, we determined whether single parameter modulations from a baseline CRF could fit our data (Ling and Carrasco, 2006). We fitted both response gain and contrast gain models to the visibility and attention ΔAI data with one free parameter (see Materials and Methods). To control for correlations between response-gain and contrast-gain models, we computed Z-transformed partial correlations (Movshon et al., 1985; Smith et al., 2005) (see Materials and Methods for details). The resulting test-statistics, Z_c and Z_r, refer to the z-transformed partial correlations for the contrast-gain and response-gain fit, respectively. A Z of zero means that there is no correlation between the data and the model. We found that the attention data was best fit with a response-gain model (for the low SF data: $Z_r = 2.26$, $p = 0.012$, one-tailed, whereas $Z_c = -0.162$, $p = 0.56$; Fig. 2F). Conversely, visibility was better fit with a contrast gain model ($Z_r = -1.09$, $p = 0.86$, $Z_c = 1.67$, $p = 0.047$; Fig. 2F). These results are further supported when the Z-values are converted back into r-values (data not shown). Finally, using more lenient trial selection criteria on RSVP performance (e.g., correct number of items ±2 or ±3) did not change these results (Fig. 2F, different symbol sizes), showing that our results are not due to the specific trial selection criterion that we used. The high SF data are similar to the low SF data (Fig. 2F) and the same conclusions will be reached. These data show that gain functions of attention and consciousness can be modeled by changes of single parameters of the baseline response function and can be approximated by pure response and contrast gain functions, respectively.

Model comparisons with the BIC

The previous analysis showed that descriptive models of contrast gain and response gain can describe our data, but to provide a

better mechanistic account, we constructed several computational models based on a normalization framework to investigate which model architecture could best explain our data.

In a normalization framework, the response to a target stimulus in one eye is divided by the activity of a normalization pool (Heeger, 1992; Reynolds and Heeger, 2009; Carandini and Heeger, 2011). This normalization pool combines the neural activity in response to both the target stimulus and the competing stimulus in the other eye (Baker and Meese, 2007; Ling and Blake, 2012). In this framework, an interocular mask will result in a contrast-gain effect.

Using a normalization framework, the modeled influence of attention on perception could result in both response gain and contrast gain effects (Reynolds and Heeger, 2009; Ling and Blake, 2012). However, an increase in attention cannot produce a negative effect on afterimage duration in this model. To explain a negative effect of attention, other researchers have proposed a two-level process (Suzuki and Grabowecky, 2003; Wede and Francis, 2007; Brascamp et al., 2010; van Boxtel et al., 2010b). The first level (L_1) is sensitive to the particular configuration of light and dark patches; that is, it is contrast-polarity sensitive. After a prolonged adaptation period, the polarity sensitivity of this level causes it to produce a negative afterimage when a blank screen is shown. This process is generally assumed to underlie afterimages. On top of this first level, the two-stage model assumes that the output of L_1 is the input of a second level. This second level is polarity insensitive, meaning that it is sensitive to the contrast of the stimulus independently of the light/dark arrangement of luminance. A change in activity at this level will change the perceived contrast of the stimulus, but will not generate an afterimage when no input is received. Here, we implement this conceptual two-level model into a simple computational model, incorporating the idea of response normalization (Heeger, 1992; Reynolds and Heeger, 2009; Carandini and Heeger, 2011) at both stages.

We modeled each of the two levels, L_1 and L_2, by a normalization process (see Materials and Methods; Fig. 4A) in which the output response of each level was modeled as follows:

$$L_{\{1,2\}}(C, a, C_{CFS}) = \frac{aC^n}{aC^n + aC_{CFS}^n + C_{50}^n} \qquad (23)$$

At L_1, C is the contrast of the inducer and, at L_2, C is the output of L_1. The parameter a is the influence of attention (where $a = 1$ means unattended and $a > 1$ means attended). Here, we let attention influence both the processing of the afterimage inducer and the CFS stimulus because our design mainly manipulated spatial attention (in fact, feature-based attention would be difficult to deploy in our stimuli because the orientation and contrast was randomized over trials). C_{CFS} is the contrast of the mask and C_{50} is the contrast at which half-maximum response is reached. The parameter n controls the steepness of the curves. Note that most of the models that we considered incorporate only contrast-gain changes because our stimuli were relatively small compared with the presumed size of the attentional window (Reynolds and Heeger, 2009). See Materials and Methods for full model descriptions.

To incorporate the differences between L_1 and L_2 in terms of their sensitivity to contrast polarity, the effects of adaptation were modeled differently in both levels. In L_1, adaptation is polarity sensitive, so, after adaptation, when presented with a blank screen, unadapted neurons will be more active then adapted neurons, resulting in a negative afterimage. Therefore, we modeled the activity of L_1 after adaptation as equal to the activity of the neurons sensitive to the opposite polarity before adaptation (as

explained in the Materials and Methods). Adaptation at L_2 was incorporated by assuming that L_2 was less sensitive to the input from L_1 after adaptation (i.e., similar to threshold elevation) by an amount proportional to its activity before adaptation (see Eq. 8 and 9 in the Materials and Methods).

Different model architectures were fitted to the data, differing in the level at which attention and consciousness took effect (i.e., at L_1 and/or L_2; Figure 4B). As a baseline, we constructed a model that included monocular normalization only, but no influences of attention or consciousness through interocular suppression (the "monocular normalization-only" (N) model). Other models were an "attention-first" model in which attentional modulated L_1 whereas consciousness modulated L_2; a "consciousness-first" model in which consciousness modulated L_1 and attention modulated L_2; and a "full model" in which both attention and consciousness could influence L_1 and L_2. We also included two models in which attention was explicitly modeled as a response gain. The models were compared using the BIC (Schwarz, 1978), which weighs the maximum value of the likelihood function against the number of free parameters in the model. A low BIC is preferred. All models were compared with the N model because this model did not incorporate the effects of our attention and visibility manipulations and therefore serves as a good baseline. Models were fit concurrently on both the afterimage durations (Fig. 2A,C) and attention and consciousness effects (Fig. 2B,D) with equal weights. Fitted model parameters for the low SF data are given in Table 1.

The BIC analysis revealed that the consciousness-first model best fits the data for low SF (Fig. 4C) and high SF (Fig. 4G) data. The attention-first model did not fit the data very well because the direction of modulation of attention and consciousness were always in the same direction. The full model fitted the data very well (Fig. 4D,H), but did not obtain a very low BIC score because of the added complexity of the model. Interestingly, the optimal fit parameters (Table 1) of the full model showed that the CFS parameter was 0.61 at L_1, and 0.05 at L_2, suggesting that CFS should only influence L_1. In addition, it showed that attention was 2.2 at L_1 (where 1 means no attention), whereas it was 100 at L_2, showing that attention mostly influenced L_2. Both of these findings conform the simpler model architecture of the consciousness-first model, suggesting that, indeed, the consciousness-first model is an optimal model. The models with attention modeled as a response gain, labeled Cr and Ar in Figure 4, C, D, G, and H, produced worse fits than the respective models with attention modeled as a contrast-gain (i.e., models C and A). Of these two, the Cr model was significantly better, but it was unable to show strong increases of attentional effects at medium contrasts and it did not evidence a plateau effect at high contrasts, thus its underperformance relative to the consciousness-first (C) model (ΔBIC = 3.84 between the two models). Posterior model comparisons (Wasserman, 2000) showed that the consciousness-first (C) model has a posterior probability of 0.83 of being the correct model, whereas the second-best model was the Cr model with a posterior probability of 0.12.

These same analyses were also performed for the high SF data, resulting in the same conclusions (Fig. 4G–J), with the consciousness-first model fitting the data the best (ΔBIC = 5.45 between the C and Cr model). Posterior model comparisons showed that the C-model has a posterior probability of 0.93 of being the correct model, whereas the second-best model was the Cr model with a posterior probability of 0.06.

This model comparison revealed that our data are best explained by a two-level normalization model in which consciousness operates at the first level, whereas attention operates at the second level.
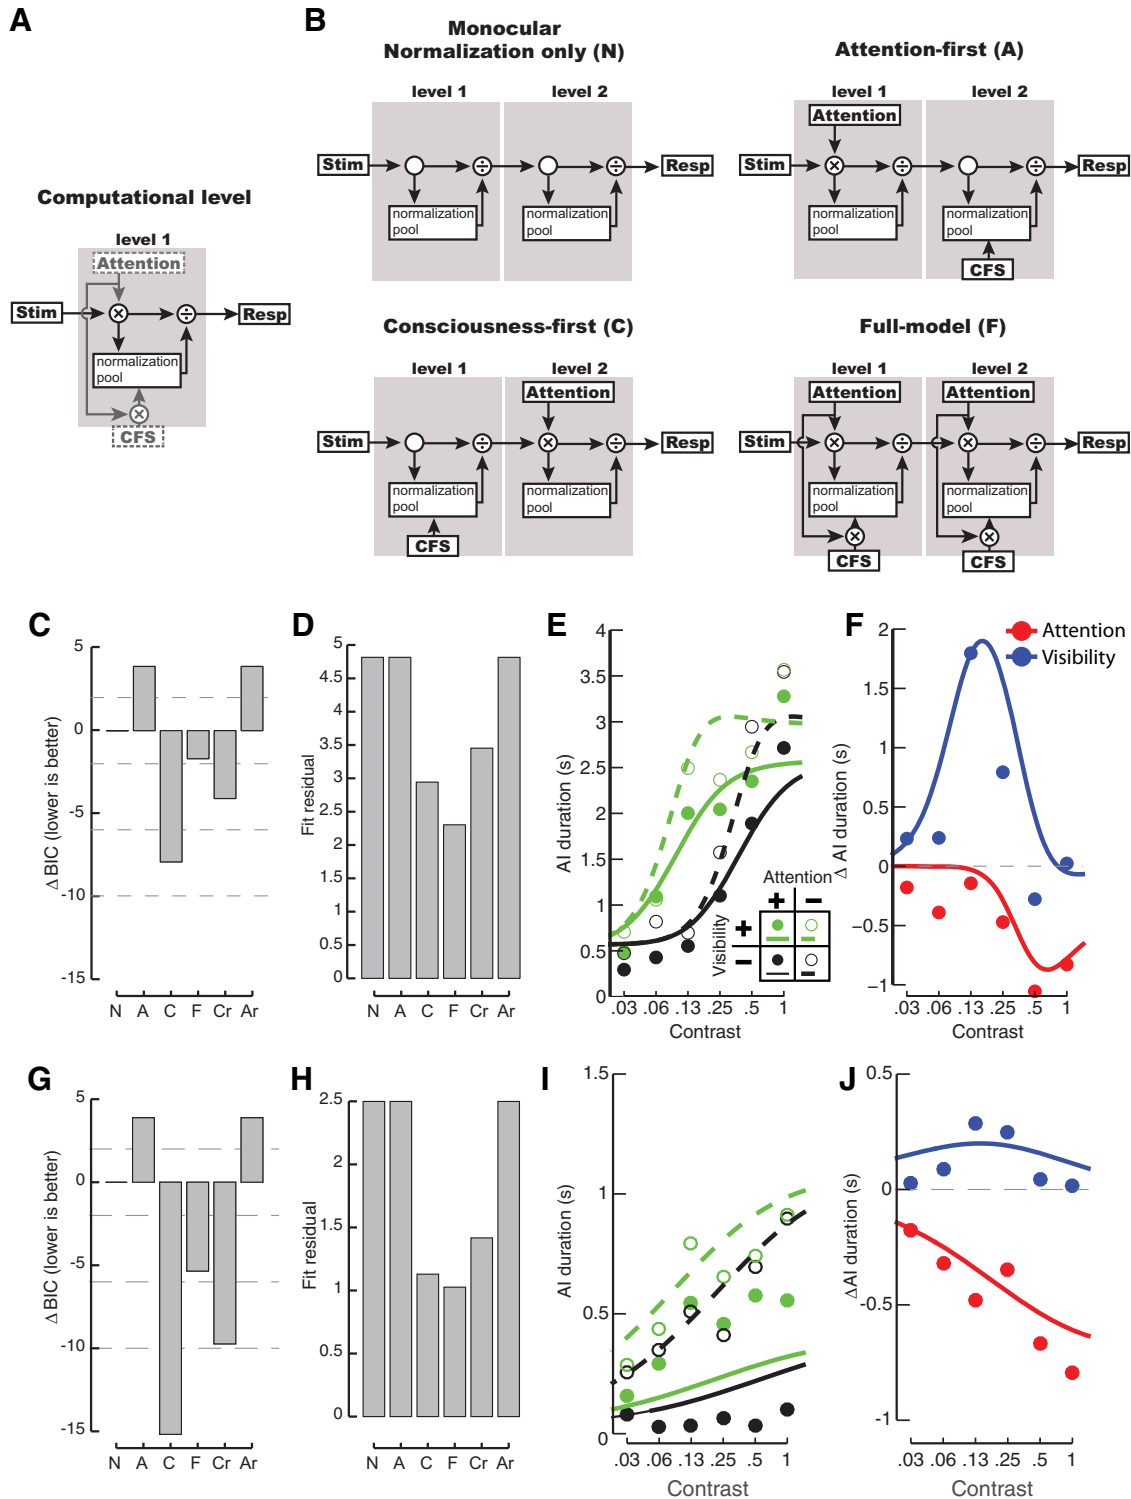
Figure 4. Model comparisons. **A**, Main computational unit that includes normalization, as well as potential influences of attention and consciousness (through CFS). These units were stacked on top of one another to create a two-level model architecture. **B**, Four main model architectures that were considered. **C**, ΔBIC scores from the BIC (lower is better) for the low SF data. Horizontal lines represent different levels of evidence (see Materials and Methods). **D**, Fit residuals from each of the models. **E**, **F**, Model fit of the consciousness-first model, the best-fitting model, through the afterimage data (**E**; data replotted from Fig. 2A), and through the attention and consciousness data (**F**; data replotted from Fig. 2B). Curves are the fits, circles are the data, color assignments as before. **G–J**, Same as **C–F**, but for the high SF data (experimental data replotted from Fig. 2C,D).

Influence of the mask

One potential explanation of our visibility findings is that they are due to the presence/absence of the mask, not the visibility/invisibility of the inducer. Past research has shown that the presence/absence of the mask does not fully determine the strength of the

aftereffects (Tsuchiya and Koch, 2005; Blake et al., 2006; van Boxtel et al., 2010b), but the mask may have an influence, through contrast adaptation, at low inducer SFs (Brascamp et al., 2010). To test for the influence of the mask, we compared linear mixed models through likelihood ratio tests (Bates et al., 2014).

Table 1. Model fit for the free parameters for the four models

	Monocular normalization only model	Attention-first model	Consciousness-first model	Full model
n	0.920	0.920	1.567	1.785
C_{50}	1	1	0.276	0.267
O	0.102	0.102	0.567	0.472
s	8.982	8.962	6.686	6.865
m	0	0	0.671	0.665
$A_1{}^a$		1		2.024
$A_2{}^a$			7.650	99.999
$C_{cfs1}{}^b$				0.610
$C_{cfs2}{}^b$				0.050

[a]Attention was set to 1 in the low-attention conditions, the reported value is the fit parameter for the attended high-attention condition.

[b]C_{cfs} was fixed to 0 or 1 when the CFS was absent or present in most models; however, it was free to vary in the "present" condition in the full model; these values are reported here. Subscripts 1 and 2 refer to level 1 and level 2.

Afterimage duration was taken as the dependent variable and fixed effects were the inducer contrast, the CFS presence/absence, and the visibility (as indicated by the participants button presses). Random intercepts per participant were included. The unrestricted model included all fixed effects and their interactions, which we compared with restricted models that excluded either the factor CFS (and its interaction terms) or visibility (and its interaction terms). For low SF conditions, visibility had a significant effect ($\chi^2_{(4)} = 33.30$, $p < 0.0001$), and CFS had a significant effect ($\chi^2_{(4)} = 9.94$, $p = 0.041$). For the high SF condition, visibility had a significant effect ($\chi^2_{(4)} = 353.87$, $p < 0.0001$), but CFS did not ($\chi^2_{(4)} = 2.57$, $p = 0.63$). The model with the visibility influence (but not CFS) had a lower BIC than the model with the CFS influence (but not visibility): ΔBIC = 23.36 for low SF and ΔBIC = 351.31 for high SF). Therefore, the main influence appears to be visibility, but CFS does influence afterimage durations for low SF inducers.

Discussion

Modulations of CRFs are one of the most valuable descriptors of visual function (Reynolds and Heeger, 2009; Carrasco, 2011), allowing one to investigate the computational mechanisms underlying visual processing (Reynolds and Heeger, 2009). We measured the modulations of afterimage duration CRFs, or gain functions, induced by attention and consciousness. We show that attention and consciousness can have opposite effects on afterimage durations and, more importantly, that attention displays a response gain function, whereas consciousness displays a contrast gain function (which is possibly stronger for low SFs). These results show that attention and consciousness, arguably the two most important cognitive functions, operate through different underlying mechanisms.

These results are consistent with previous demonstrations of a dissociation between attention and consciousness (Kentridge et al., 2004; Wyart and Tallon-Baudry, 2008; Brascamp et al., 2010; van Boxtel et al., 2010b; Watanabe et al., 2011), but additionally provide an explanation of the processes that underlie this dissociation. Showing such clear dissociations between attention and consciousness is not always possible. For example, using a different experimental setup and parameters, attention produces contrast-gain functions at a behavioral level (Reynolds and Heeger, 2009; Herrmann et al., 2010), just like consciousness often does. Attention and consciousness would then appear to work synergistically. However, even though attention and consciousness show synergistic effects in those experiments, our

findings suggest their computational underpinnings can be dissociated using the current paradigm.

Through extensive model comparisons, we show that the dissociation is best explained by a hierarchical two-level normalization model, in which attention acts on a late contrast-polarity-insensitive level, whereas consciousness manipulations, by means of interocular suppression, act on an early contrast-polarity-sensitive level. Because contrast-polarity sensitivity is predominantly found subcortically in LGN and in layer 4 in V1 (Hubel and Wiesel, 1968; Levitt et al., 2001), our findings are consistent with the idea that interocular suppression operates at early visual stages (Blake, 1989; Tong et al., 2006), potentially as early as LGN or V1 (Sengpiel et al., 1998; Polonsky et al., 2000; Meese and Hess, 2004; Watanabe et al., 2004; Haynes et al., 2005; Li et al., 2005; Bahrami et al., 2008; Yuval-Greenberg and Heeger, 2013). However, other research suggests that interocular masking at LGN or V1 is weak or nonexistent (Lehky and Maunsell, 1996; Leopold and Logothetis, 1996; Macknik and Martinez-Conde, 2004; Keliris et al., 2010; Watanabe et al., 2011). Our results cannot resolve this issue because, although we show sensitivity to contrast polarity, cells with such sensitivity (Zhou et al., 2000), as well as (rare) ocularly biased cells (Zeki, 1978; Hubel and Livingstone, 1987), exist in both striate and extrastriate visual areas. However, most results (including ours) are consistent with neurophysiological (Leopold and Logothetis, 1996) and computational (Freeman, 2005) findings that suggest that interocular inhibition at early stages is small and noisy and complete perceptual suppression builds up over consecutive stages in the visual hierarchy.

Our data further show that attention mainly operates on a higher contrast-polarity-insensitive level. This is largely consistent with the idea that attention allows for vision with scrutiny through top-down modulation of neural activity (Hochstein and Ahissar, 2002; Buffalo et al., 2010), acting more strongly higher up in the visual hierarchy (Rainer et al., 1998).

Different from our two-stage model, two recent models proposed that attention operates at the same level as interocular suppression (Ling and Blake, 2012; Li et al., 2015). Although seemingly at odds, this approach is not inconsistent with ours. For example, in Li et al. (2015), attention is binocular, consistent with the higher-level (second) stage in our model. Attention then influences (presumably through feedback) interocular competition. Importantly, the attentional influences on interocular suppression were strong only when competing stimuli were large, but not when they were small (Ling and Blake, 2012; Li et al., 2015). Because our stimuli were small, this may have obviated the need to model the first-level attentional influences for our experiments.

Other data suggest that attention is not merely a modulator of, but actually a prerequisite for binocular rivalry (Zhang et al., 2011; Brascamp and Blake, 2012). Although attention generally influences binocular rivalry minimally (Meng and Tong, 2004; van Ee et al., 2005), the finding that binocular rivalry may require attention could suggest that there exists an interaction between our attention manipulation and our consciousness manipulation (see also Ling and Blake, 2012; and our Fig. 3), thus potentially making them not completely independent in our experimental design. Although this situation is not ideal, the existence of such an interaction means that, if we were able to better separate attention and consciousness, then the results would probably be still clearer.

Interestingly, whereas attention may be required for binocular rivalry, conscious awareness of the visual conflict is not (Brascamp et al., 2015; Zou et al., 2016). This finding supports the proposed dissociation of attention and consciousness. It must be

mentioned that recent neurophysiological measurements suggest that binocular rivalry may occur in the absence of both attention and consciousness (Xu et al., 2016), meaning that the above concerns about our paradigm are moot.

From a computational perspective, it is interesting that the best-performing consciousness-first model only included contrast-gain computations, even outperforming models that explicitly included response-gains for attention. It nevertheless produced a response gain for attention because the modeled contrast-response function at the second level was not yet at saturation at high contrasts, so a contrast-gain change induced by attention translated into increased activity, even at high contrasts. Arguably, this unsaturated response is not consistent with neurophysiological data. However, this "true" contrast response behavior of a neuron cannot be measured in the brain. One can only measure the contrast response curve at a certain level determined by the contrast-dependent input from preceding neuronal levels. In other words, the output of level X is dependent on the contrast-dependent input it receives from level X − 1. This "observed" contrast-response curve is the one that is measured in neurophysiological experiments. Consistent with these observations, our model also shows response saturation of the contrast-response function of L_2 given the input of L_1 (Fig. 4E). In addition to this effect, the response gain is then further strengthened by adaptation (which is effectively a response gain change) that is dependent on the attention-modulated activity.

Our modeling suggests that both the effects of attention and consciousness on afterimage duration are the result of a signal enhancement at the neuronal level. It has been suggested that these effects could therefore operate through identical or similar mechanisms (Brascamp et al., 2010; Blake et al., 2014), thus implying a weak separation between attention and consciousness. This interpretation, however, was based on measurements at just a single inducer contrast. Here, we have measured the entire contrast response curve and analyzed the results with different two-level models. Our analysis reveals that, even though attention and consciousness both cause increases in activity, they achieve that through processes that give rise to different (gain) functions and by influencing different levels within the visual processing hierarchy. Our overall set of findings is not consistent with alternative interpretations based on the assumption that the effects of suppression are purely due to decreasing attention to a stimulus that is rendered invisible (Blake et al., 2014; Li et al., 2015).

A separate concern is that the effects that we attribute to consciousness may instead be due to a difference in contrast adaptation when the interocular mask is present versus when it is absent (i.e., it is stimulus dependent) and not due to changes in conscious perception per se (Cohen et al., 2012). Consistent with past findings (Brascamp et al., 2010), we find that, for low SF stimuli, there is an influence of the mask, but not for high SF stimuli. It is thus worthwhile to investigate and control for this potential influence (Brascamp et al., 2010; Watanabe et al., 2011; Yuval-Greenberg and Heeger, 2013). However, conscious visibility had a significantly stronger effect in our data, consistent with previous indications that aftereffect strength is largely determined by the conscious percept and that the presence/absence of the mask is not (fully) determining the strength of the aftereffects (Tsuchiya and Koch, 2005; Blake et al., 2006; van Boxtel et al., 2010b).

It is currently not possible to determine how general our findings are. However, apart from afterimages, attention can decrease perception or performance in motion-induced blindness (Geng et al., 2007; Schölvinck and Rees, 2009) and the motion aftereffect (Murd and Bachmann, 2011), for Troxler/peripheral fad-

ing (Lou, 1999; De Weerd et al., 2006), visual memory (Voss and Paller, 2009), SF (Yeshurun and Carrasco, 1998), the attentional blink (Olivers and Nieuwenhuis, 2005), and visual search (Smilek et al., 2006). These findings suggest that our conclusions may be valid beyond the context of afterimages.

In conclusion, our results indicate that signal enhancement functions differently through attention and consciousness. Because it is easy to conflate attention and consciousness effects, future research would gain from carefully controlling parameters that influence attention and consciousness.

References

Albrecht DG, Hamilton DB (1982) Striate cortex of monkey and cat: contrast response function. J Neurophysiol 48:217–237. Medline

Bahrami B, Carmel D, Walsh V, Rees G, Lavie N (2008) Spatial attention can modulate unconscious orientation processing. Perception 37:1520–1528. CrossRef Medline

Baker DH, Meese TS (2007) Binocular contrast interactions: dichoptic masking is not a single process. Vision Res 47:3096–3107. CrossRef Medline

Bates D, Mächler M, Bolker B, Walker S (2014) Fitting linear mixed-effects models using lme4. Journal of Statistical Software 67:1–48.

Blake R (1989) A neural theory of binocular rivalry. Psychol Rev 96:145–167. CrossRef Medline

Blake R, Tadin D, Sobel KV, Raissian TA, Chong SC (2006) Strength of early visual adaptation depends on visual awareness. Proc Natl Acad Sci U S A 103:4783–4788. CrossRef Medline

Blake R, Brascamp J, Heeger DJ (2014) Can binocular rivalry reveal neural correlates of consciousness? Philos Trans R Soc Lond B Biol Sci 369: 20130211. CrossRef Medline

Brascamp JW, Blake R (2012) Inattention abolishes binocular rivalry: perceptual evidence. Psychol Sci 23:1159–1167. CrossRef Medline

Brascamp JW, van Boxtel JJ, Knapen TH, Blake R (2010) A dissociation of attention and awareness in phase-sensitive but not phase-insensitive visual channels. J Cogn Neurosci 22:2326–2344. CrossRef Medline

Brascamp J, Blake R, Knapen T (2015) Negligible fronto-parietal BOLD activity accompanying unreportable switches in bistable perception. Nat Neurosci 18:1672–1678. CrossRef Medline

Buffalo EA, Fries P, Landman R, Liang H, Desimone R (2010) A backward progression of attentional effects in the ventral stream. Proc Natl Acad Sci U S A 107:361–365. CrossRef Medline

Carandini M, Heeger DJ (2011) Normalization as a canonical neural computation. Nat Rev Neurosci 13:51–62. CrossRef Medline

Carrasco M (2011) Visual attention: the past 25 years. Vision Res 51:1484–1525. CrossRef Medline

Chun MM, Wolfe JM (2000) Visual attention. In: Blackwell's Handbook of Perception (Goldstein EB, ed), pp 272–310. Oxford, UK: Blackwell.

Cohen MA, Cavanagh P, Chun MM, Nakayama K (2012) The attentional requirements of consciousness. Trends Cogn Sci 16:411–417. CrossRef Medline

De Brigard F, Prinz J (2010) Attention and consciousness. In: Wiley interdisciplinary reviews: cognitive science (Nadel L, ed), pp 51–59. New York: Wiley.

Dehaene S, Changeux JP, Naccache L, Sackur J, Sergent C (2006) Conscious, preconscious, and subliminal processing: a testable taxonomy. Trends Cogn Sci 10:204–211. CrossRef Medline

De Weerd P, Smith E, Greenberg P (2006) Effects of selective attention on perceptual filling-in. J Cogn Neurosci 18:335–347. CrossRef Medline

Freeman AW (2005) Multistage model for binocular rivalry. J Neurophysiol 94:4412–4420. CrossRef Medline

Geng HY, Song QL, Li YF, Xu S, Zhu Y (2007) Attentional modulation of motion-induced blindness. Chinese Science Bulletin 52:1063–1070. CrossRef

Haynes JD, Deichmann R, Rees G (2005) Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus. Nature 438:496–499. CrossRef Medline

Heeger DJ (1992) Normalization of cell responses in cat striate cortex. Vis Neurosci 9:181–197. CrossRef Medline

Herrmann K, Montaser-Kouhsari L, Carrasco M, Heeger DJ (2010) When size matters: attention affects performance by contrast or response gain. Nat Neurosci 13:1554–1559. CrossRef Medline

Hochstein S, Ahissar M (2002) View from the top: hierarchies and reverse hierarchies in the visual system. Neuron 36:791–804. CrossRef Medline

Hubel DH, Livingstone MS (1987) Segregation of form, color, and stereopsis in primate area 18. J Neurosci 7:3378–3415. Medline

Hubel DH, Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. J Physiol 195:215–243. CrossRef Medline

Iwasaki S (1993) Spatial attention and two modes of visual consciousness. Cognition 49:211–233. CrossRef Medline

Kanai R, Walsh V, Tseng CH (2010) Subjective discriminability of invisibility: a framework for distinguishing perceptual and attentional failures of awareness. Conscious Cogn 19:1045–1057. CrossRef Medline

Kass RE, Raftery AE (1995) Bayes factors. Journal of the American Statistical Association 90:773–795. CrossRef

Keliris GA, Logothetis NK, Tolias AS (2010) The role of the primary visual cortex in perceptual suppression of salient visual stimuli. J Neurosci 30: 12353–12365. CrossRef Medline

Kentridge RW, Heywood CA, Weiskrantz L (2004) Spatial attention speeds discrimination without awareness in blindsight. Neuropsychologia 42: 831–835. CrossRef Medline

Kim CY, Blake R (2005) Psychophysical magic: rendering the visible 'invisible'. Trends Cogn Sci 9:381–388. CrossRef Medline

Koch C, Tsuchiya N (2007) Attention and consciousness: two distinct brain processes. Trends Cogn Sci 11:16–22. CrossRef Medline

Koch C, Tsuchiya N (2012) Attention and consciousness: related yet different. Trends Cogn Sci 16:103–105. CrossRef Medline

Lehky SR, Maunsell JH (1996) No binocular rivalry in the LGN of alert macaque monkeys. Vision Res 36:1225–1234. CrossRef Medline

Leopold DA, Logothetis NK (1996) Activity changes in early visual cortex reflects monkeys' percepts during binocular rivalry. Nature 379:549–553. CrossRef Medline

Levitt JB, Schumer RA, Sherman SM, Spear PD, Movshon JA (2001) Visual response properties of neurons in the LGN of normally reared and visually deprived macaque monkeys. J Neurophysiol 85:2111–2129. Medline

Li B, Peterson MR, Thompson JK, Duong T, Freeman RD (2005) Cross-orientation suppression: monoptic and dichoptic mechanisms are different. J Neurophysiol 94:1645–1650. CrossRef Medline

Li HH, Carrasco M, Heeger DJ (2015) Deconstructing interocular suppression: attention and divisive normalization. PLoS Comput Biol 11:e1004510. CrossRef Medline

Ling S, Blake R (2012) Normalization regulates competition for visual awareness. Neuron 75:531–540. CrossRef Medline

Ling S, Carrasco M (2006) Sustained and transient covert attention enhance the signal via different contrast response functions. Vision Res 46:1210–1220. CrossRef Medline

Lou L (1999) Selective peripheral fading: evidence for inhibitory sensory effect of attention. Perception 28:519–526. CrossRef Medline

Macknik SL, Martinez-Conde S (2004) Dichoptic visual masking reveals that early binocular neurons exhibit weak interocular suppression: implications for binocular vision and visual awareness. J Cogn Neurosci 16: 1049–1059. CrossRef Medline

Meese TS, Hess RF (2004) Low spatial frequencies are suppressively masked across spatial scale, orientation, field position, and eye of origin. J Vis 4:843–859. CrossRef Medline

Meng M, Tong F (2004) Can attention selectively bias bistable perception? Differences between binocular rivalry and ambiguous figures. J Vis 4:539–551. CrossRef Medline

Mole C (2008) Attention and consciousness. Journal of Consciousness Studies 15:86–104.

Movshon JA, Adelson EH, Gizzi MS, Newsome WT (1985) The analysis of moving visual patterns. Pontificiae Academiae Scientiarum Scripta Varia 54:117–151.

Murd C, Bachmann T (2011) Spatially localized motion aftereffect disappears faster from awareness when selectively attended to according to its direction. Vision Res 51:1157–1162. CrossRef Medline

Olivers CN, Nieuwenhuis S (2005) The beneficial effect of concurrent task-irrelevant mental activity on temporal attention. Psychol Sci 16:265–269. CrossRef Medline

Polonsky A, Blake R, Braun J, Heeger DJ (2000) Neuronal activity in human primary visual cortex correlates with perception during binocular rivalry. Nat Neurosci 3:1153–1159. CrossRef Medline

Posner MI (1994) Attention: the mechanisms of consciousness. Proc Natl Acad Sci U S A 91:7398–7403. CrossRef Medline

Prinz J (2012) The conscious brain. Oxford: OUP.

Rainer G, Asaad WF, Miller EK (1998) Selective representation of relevant information by neurons in the primate prefrontal cortex. Nature 393: 577–579. CrossRef Medline

Rees G, Russell C, Frith CD, Driver J (1999) Inattentional blindness versus inattentional amnesia for fixated but ignored words. Science 286:2504–2507. CrossRef Medline

Reynolds JH, Heeger DJ (2009) The normalization model of attention. Neuron 61:168–185. CrossRef Medline

Schölvinck ML, Rees G (2009) Attentional influences on the dynamics of motion-induced blindness. J Vis 9:38.1–38.9. CrossRef Medline

Schwarz G (1978) Estimating the dimension of a model. Annals of Statistics 6:461–464. CrossRef

Sengpiel F, Baddeley RJ, Freeman TC, Harrad R, Blakemore C (1998) Different mechanisms underlie three inhibitory phenomena in cat area 17. Vision Res 38:2067–2080. CrossRef Medline

Smilek D, Enns JT, Eastwood JD, Merikle PM (2006) Relax! Cognitive strategy influences visual search. Vis Cogn 14:543–564. CrossRef

Smith MA, Majaj NJ, Movshon JA (2005) Dynamics of motion signaling by neurons in macaque area MT. Nat Neurosci 8:220–228. CrossRef Medline

Suzuki S, Grabowecky M (2003) Attention during adaptation weakens negative afterimages. J Exp Psychol Hum Percept Perform 29:793–807. CrossRef Medline

Tong F, Meng M, Blake R (2006) Neural bases of binocular rivalry. Trends Cogn Sci 10:502–511. CrossRef Medline

Tsuchiya N, Koch C (2005) Continuous flash suppression reduces negative afterimages. Nat Neurosci 8:1096–1101. CrossRef Medline

van Boxtel JJ, Tsuchiya N, Koch C (2010a) Consciousness and attention: on sufficiency and necessity. Front Psychol 1:217. CrossRef Medline

van Boxtel JJ, Tsuchiya N, Koch C (2010b) Opposing effects of attention and consciousness on afterimages. Proc Natl Acad Sci U S A 107:8883–8888. CrossRef Medline

van Ee R, van Dam LC, Brouwer GJ (2005) Voluntary control and the dynamics of perceptual bi-stability. Vision Res 45:41–55. CrossRef Medline

Voss JL, Paller KA (2009) An electrophysiological signature of unconscious recognition memory. Nat Neurosci 12:349–355. CrossRef Medline

Wasserman L (2000) Bayesian model selection and model averaging. J Math Psychol 44:92–107. CrossRef Medline

Watanabe K, Paik Y, Blake R (2004) Preserved gain control for luminance contrast during binocular rivalry suppression. Vision Res 44:3065–3071. CrossRef Medline

Watanabe M, Cheng K, Murayama Y, Ueno K, Asamizuya T, Tanaka K, Logothetis N (2011) Attention but not awareness modulates the BOLD signal in the human V1 during binocular suppression. Science 334:829–831. CrossRef Medline

Wede J, Francis G (2007) Attentional effects on afterimages: theory and data. Vision Res 47:2249–2258. CrossRef Medline

Wundt W (1874) Grundzüge der physiologischen Psychologie. Leipzig: Engelmann.

Wyart V, Tallon-Baudry C (2008) Neural dissociation between visual awareness and spatial attention. J Neurosci 28:2667–2679. CrossRef Medline

Xu H, Han C, Chen M, Li P, Zhu S, Fang Y, Hu J, Ma H, Lu HD (2016) Rivalry-like neural activity in primary visual cortex in anesthetized monkeys. J Neurosci 36:3231–3242. CrossRef Medline

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

Yuval-Greenberg S, Heeger DJ (2013) Continuous flash suppression modulates cortical activity in early visual cortex. J Neurosci 33:9635–9643. CrossRef Medline

Zeki SM (1978) Uniformity and diversity of structure and function in rhesus monkey prestriate cortex. J Physiol 277:273–290. CrossRef Medline

Zhang P, Jamison K, Engel S, He B, He S (2011) Binocular rivalry requires visual attention. Neuron 71:362–369. CrossRef Medline

Zhou H, Friedman HS, von der Heydt R (2000) Coding of border ownership in monkey visual cortex. J Neurosci 20:6594–6611. Medline

Zou J, He S, Zhang P (2016) Binocular rivalry from invisible patterns. Proc Natl Acad Sci U S A 113:8408–8413. CrossRef Medline