



Motion-induced blindness does not affect the formation of negative afterimages

Constanze Hofstoetter^{a,*}, Christof Koch^b, Daniel C. Kiper^a

^a *Institute of Neuroinformatics, Swiss Federal Institute of Technology and University of Zurich, Zurich, Switzerland*

^b *Division of Biology, 139-74, California Institute of Technology, Pasadena, USA*

Received 22 December 2003

Available online 23 August 2004

Abstract

Aftereffects induced by invisible stimuli constitute a powerful tool to investigate what type of neural information processing can occur in the absence of visual awareness. This approach has been successfully used to demonstrate that awareness of oriented gratings or translating stimuli is not necessary to obtain a robust orientation-specific or motion-specific aftereffect. We exploit motion-induced blindness (MIB, [Bonneh, Cooperman, & Sagi, 2001](#)) to investigate the related question of the influence of visual awareness on the formation of negative afterimages. Our results show that MIB does not affect the persistence and intensity of afterimages. Thus, there is no significant contribution to the formation of afterimages beyond the sites mediating MIB.

© 2004 Elsevier Inc. All rights reserved.

Keywords: Perceptual suppression; Color; Aftereffect; Visual awareness; Consciousness

1. Introduction

In the last two decades, the scientific community's attitude toward consciousness has undergone a profound change. From a topic that was previously considered beyond the boundaries of rigor-

* Corresponding author. Fax: +41-1-635-30-53.

E-mail address: connie@ini.phys.ethz.ch (C. Hofstoetter).

ous, quantitative investigation, it is now often referred to as one of the last frontiers of science (Metzinger, 1995). Thus, neuroscientists set out to study the neural basis for consciousness (for review, see e.g., Zeman, 2001). A first step towards this goal is to determine minimal sets of neural events that are necessary and sufficient to evoke certain conscious experiences (Crick & Koch, 2003; Metzinger, 2000). In our research we focus on the visual system, the most extensively studied mammalian sensory system, to begin our quest for such neuronal correlates of consciousness (NCC; Koch, 2004).

A common strategy to investigate the neuronal correlates of visual awareness is to study conditions when a stimulus is present, but observers fail to perceive it (see, e.g., Blake, 1997; Logothetis, 1998; Merikle, Smilek, & Eastwood, 2001; Rajimehr, 2004). In healthy human observers, different spatio-temporal manipulations can induce a dissociation between the physical and the perceived exposure to a visual stimulus. A stimulus can be rendered completely invisible for a short period of time by forward or backward masking (Bachmann, 2000; Breitmeyer, 1984; Breitmeyer & Ögmen, 2000; Lamme, Zipser, & Spekreijse, 2002; Macknik & Livingstone, 1998; Rolls, Tovee, & Panzeri, 1999), binocular rivalry (Blake & Fox, 1974; Blake & Logothetis, 2002; Engel, Fries, König, Brecht, & Singer, 1999; Lee & Blake, 2002; Leopold & Logothetis, 1996; Logothetis, Leopold, & Sheinberg, 1996), flash suppression (Wilke, Logothetis, & Leopold, 2003; Wolfe, 1984), or by motion-induced blindness (Bonneh, Cooperman, & Sagi, 2001). Other manipulations prevent the perception of a specific stimulus attribute for a prolonged period of time. For example, observers may fail to perceive the orientation of a grating when its spatial frequency exceeds the resolution of the eye (He & MacLeod, 2001) or when it is surrounded by similar gratings (He, Cavanagh, & Intriligator, 1996; Toet & Levi, 1992). By exploiting such stimulus configurations, experimenters aim to compare the neural events underlying normal visual perception to those when the physical stimulus is identical but is not consciously perceived.

Visual aftereffects, transient percepts induced by prior exposure to an adapting stimulus, constitute a useful, non-invasive tool to study the neuronal basis of perception in humans. The formation of aftereffects to a particular stimulus attribute is generally thought to reflect adaptation of neurons processing this attribute. For example, after adapting to a grating whose dominant orientation is slightly tilted toward the left, away from the vertical, a perfectly vertical grating presented thereafter will appear to be slightly tilted to the right. In addition, the detection threshold for a grating with the same orientation increases significantly. These aftereffects are generally assigned to the adaptation of orientation-selective cells in visual cortex (Blakemore & Campbell, 1969; Maffei, Fiorentini, & Bisti, 1973; Movshon & Lennie, 1979).

By studying aftereffects induced by perceptually invisible stimuli, we can infer the type of neural information processing occurring in the absence of visual awareness. For example, orientation-specific adaptation has been studied to gratings that were perceptually invisible because they underwent binocular rivalry (Blake & Fox, 1974), were masked due to crowding (He et al., 1996), were too fine to resolve (He & MacLeod, 2001; Rajimehr, 2004) or were suppressed by motion-induced blindness (Kouhsari, Moradi, Zand-Vakili, & Esteky, 2004). These studies consistently demonstrate that orientation-selective adaptation occurs despite the lack of conscious visual awareness of the adapting stimulus. Thus, activation of cells mediating this type of adaptation is not sufficient to generate a conscious percept. Collectively, these findings imply that orientation-selective cells in V1 are not part of the NCC (Crick & Koch, 1995; Koch, 2004).

The purpose of this study was to examine how visual awareness affects the formation of negative afterimages (AIs). Negative AIs are visual traces whose color and contrast polarity are reversed relative to that of the adaptor. For example, prolonged exposure to a dark green adaptor results in the subsequent perception of a light, reddish negative AI. We measured negative AIs induced by stimuli that were temporarily invisible due to motion-induced blindness (MIB). In MIB, salient stimuli are intermittently suppressed when surrounded by a moving stimulus (Bonneh et al., 2001). There is evidence that perceptual suppression caused by MIB occurs in the cortex (Bonneh et al., 2001; Hsu, Yeh, & Kramer, 2004; Kouhsari et al., 2004; Rajimehr, 2004; see discussion). While the origin of AIs is often assigned solely to retinal adaptation (Craik, 1940; Lack, 1978; Wilson, 1997), recent studies suggest an additional cortical contribution to the formation of AIs (Hadjikhani, Liu, Dale, Cavanagh, & Tootell, 1998; Shimojo, Kamitani, & Nishida, 2001; Suzuki & Grabowecky, 2003). Thus, if MIB modulated AI strength, this would be clear evidence for an extra-retinal contribution to AI formation.

In our experiments, we measured AIs induced by two bright yellow squares, the adaptors. Under MIB conditions, these yellow squares were temporarily suppressed from visual awareness by a cloud of moving blue random dots evoking the percept of a rotating sphere. In 50% of the trials, the moving dots were present and induced MIB (MIB trials). In the remaining 50% of the trials (playback, or PB, trials), the blue dots were absent and the adaptors were physically removed in a way that precisely mimicked the subject's percept recorded in the preceding MIB trial.

For PB trials, we compared AI persistence and intensity induced by an adaptor that had just been physically removed to one that had been present throughout. It is well established that the longer an adaptor is presented, the more persistent and the more intense the AI it induces (see, e.g., Kelly & Martinez-Uriegaz, 1993). The PB data we present clearly reflect this dependence of the strength of the perceived AI on the physical exposure time to the adaptor. Therefore, our method is sensitive enough to detect an effect in our stimulation conditions.

For MIB trials, we compared AI persistence and intensity induced by an adaptor that had just been temporarily suppressed by MIB to one induced by an adaptor that had been visible throughout. The data we present demonstrate that MIB does not disrupt AI persistence and intensity. Thus, there is no significant contribution to the formation of negative afterimages beyond the site(s) mediating MIB.

2. Methods

2.1. Subjects

Thirty-two volunteers participated in this study. As described below, the analysis is based on data of 16 subjects (aged from 25 to 48; mean 29.6 years, 13 male/3 female) with normal or corrected to normal vision.

All experiments were undertaken with the understanding and written consent of each subject. Experiments conform to the institutional and national guidelines for experiments with human subjects and with the Declaration of Helsinki.

2.2. Apparatus

Stimuli were generated on a Macintosh G4/800 computer (Apple, Cupertino, CA) using MatLab (Mathworks, Natick, MA) including the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997). The stimuli were presented at 1024×768 pixel resolution and 120 Hz refresh rate on a 19 in. computer screen (Hitachi CM 772, Tokyo, Japan). The CIE coordinates of the screen were: red: 0.625, 0.34; green: 0.285, 0.605; blue: 0.150, 0.065. The screen was located at 57 cm viewing distance. To minimize head movements, subjects placed their forehead and chin against mechanical restraints.

2.3. Stimulus

The two adaptors were static yellow squares with a width of 0.35° visual angle and a luminance of 61 cd/m^2 presented on a black background (0.132 cd/m^2) with a white central fixation cross (width 0.4° , luminance 70 cd/m^2). The two adaptors were centered at 1.8° to the left and the right of the fixation cross.

In the MIB stimulus (see Fig. 1), adaptors were juxtaposed next to 250 blue moving dots (diameter: 0.08° visual angle, luminance: 6.4 cd/m^2). The dots 2D-position resembled a random distribution on the surface of a 3D-sphere (10.5° in diameter) and their displacement from frame to frame resulted in the impression of a rotation around a diagonal axis ($1/3$ rotations/s). No dots were displayed in the immediate vicinity of the adapting stimuli, giving rise to square-shaped “protection zones” that extended 0.5° around the yellow squares to prevent any local interaction between the adaptors and the small blue dots.

Although the adaptors were physically present throughout the adaptation phase of an MIB trial, they were occasionally suppressed from conscious awareness (see Fig. 2). The adapting phase was terminated when one and only one of the adaptors was suppressed for more than 0.4 s and one of the following criteria was met: the suppressed adaptor reappeared or the other adaptor dis-

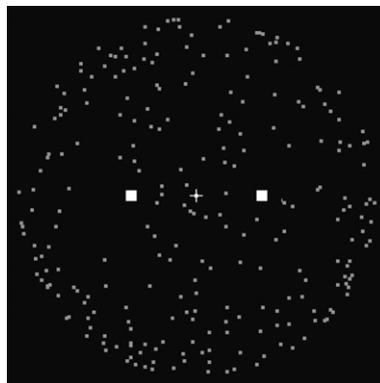


Fig. 1. The MIB inducing stimulus. Two static squares, the adaptors, are juxtaposed on a background of small moving dots that give the impression of a rotating sphere. (The stimulus used in our experiments was comprised of yellow adaptors and blue moving dots, for a demo see www.ini.unizh.ch/~connie/demos.html.)

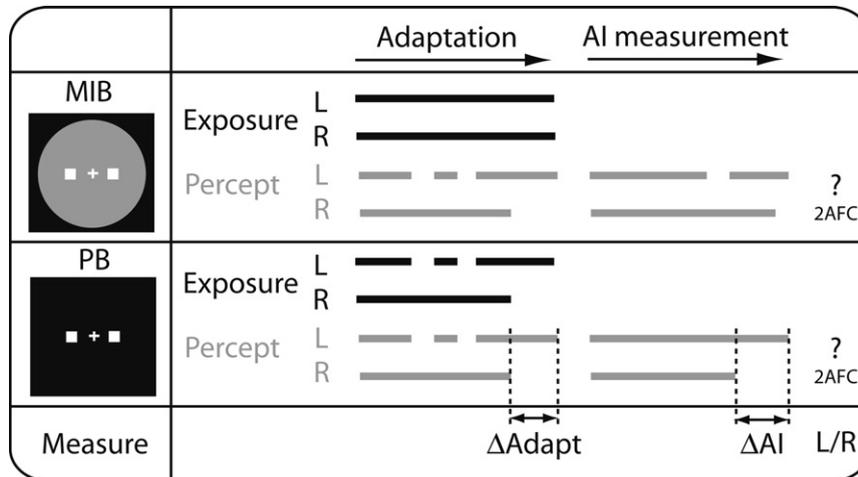


Fig. 2. Trial design. A trial consists of two phases: The adaptation phase and the phase to measure the AIs, in which the subject is asked to report on his/her perception of the left and right yellow adaptors or AIs, respectively. At the end of the trial, the subject is asked in a two alternative forced-choice manner (2AFC) whether the left or the right AI was more intense. Motion-induced-blindness (MIB) trials and playback (PB) trials alternate during the experiment. In an MIB trial the physical exposure to the left and right adaptor is equal and continuous. However, due to MIB, the duration for which the two adaptors are perceived will be different. In a PB trial the adaptors are removed from the display so that the physical exposure to the left and right adaptor mirrors the subject's perception in the preceding MIB trial. Since there are no moving blue dots in the PB condition, no perceptual suppression occurs. Thus, the perceived exposure to the left and right adaptor will be the same as in the MIB trial, but the physical exposure will be different. The parameters used for the data analysis are the relative difference in perceived adaptation time triggering the AI measuring period, ΔAdapt , the relative difference in the duration of AI perception, ΔAI , and the intensity rating, L/R.

appeared. Therefore, the duration of the MIB stimulus presentation varied from trial to trial, depending on the subject's perception.

In the PB condition, the yellow adaptors were presented alone. They were physically removed from the screen, mimicking exactly the suppression periods recorded in the preceding MIB trial. An adapting phase was thus terminated after one and only one of the adaptors had been presented alone for at least 0.4s. The duration of the adaptation phase was thereby equal to that of the perceptually matched MIB trial, but the physical exposure to the adaptors was reduced.

The AIs induced by the two adaptors were observed on a grey square (size 11° , luminance 22cd/m^2). The properties of the background that AIs are viewed upon are known to affect AI appearance (Anstis, Rogers, & Henry, 1978). Pilot studies indicated that subjects felt most comfortable describing AIs on a grey background with this luminance. With our stimulation parameters, subjects experienced negative AIs, i.e., bluish AIs appearing as dark squares on the grey background. The duration of the AI measurement was dependent on the previous adaptation phase, being twice as long as the adaptation phase but no shorter than 10s and no longer than 60s. The central fixation cross turned black for 1s during the AI measurement period. As explained below, this marked the period during which the subject was asked to make his/her AI intensity judgment. All stimuli were viewed binocularly in a darkened room.

2.4. Task

During the adaptation phase of MIB and PB trials, the subject's task was to report on his/her percept of the two adaptors by means of button presses. Subjects were told to use the left and/or right arrow key to independently monitor the disappearance of the left and/or right adaptor. The key had to be held down while the corresponding adaptor was not visible.

Similarly, during the AI measurement, the subject had to indicate when the left and/or right AI was not perceived. After termination of each AI measurement, the subject was asked to indicate with a single button press "Which of the two afterimages was more intense?" It has been observed that the intensity of AIs increases over the first few seconds after the removal of an adaptor, and then begins to wax and wane. Pilot studies had revealed that subjects felt more comfortable making a relative intensity judgment based on a brief period in time rather than having to judge the AIs average intensity over a long period. Thus, subjects were asked to base their judgment on their perception between the third and fourth second in the AI measurement period (indicated by a short beep and a change in fixation cross brightness). This interval was chosen because AIs had generally already reached their maximum intensity but had not yet vanished, even for short adaptation times. Thus, AI strength was quantified in terms of duration and relative intensity.

2.5. Experimental design

An experimental session lasted approximately 40 min and consisted of the following stages:

- Instructions and training (approx. 5 min).
- MIB-baseline recording (2.5 min).
- Alternating MIB and PB trials (2×15 min).

After adjusting head and chin rest, subjects were given written instructions on the screen. At the beginning of the experiment, each subject completed several training trials, during which we made sure that the instructions were understood correctly. At least two simulated MIB (in which one of the adaptors was also physically removed after a couple of seconds) and PB trials had to be completed. After that, the subject could continue training until feeling comfortable with the procedure.

Next, the MIB stimulus was presented continuously for 2.5 min to measure the range of suppression intervals occurring for each subject. This adaptation to the stimulus was necessary because naïve subjects tend to experience MIB only after some exposure to the stimulus.

The rest of the experiment consisted of alternating MIB and perceptually matched PB trials, as explained above (see Fig. 2). Subjects could rest shortly before initiating a new trial and were given the chance to take a longer break after 15 min.

Two assumptions underlie our experimental design: First, we assume that the stronger the difference in adaptation, the stronger the difference in AI formation. Second, we assume that the most recent event has the strongest effect on AI strength. Thus, to maximize the chances to detect a perceptual effect on their formation, AIs were measured immediately after a prolonged perceptual suppression of one and only one adaptor. As discussed in Section 3, there is a tradeoff be-

tween maximizing the perceptual adaptation difference triggering an AI measurement and the amount of data that can be recorded during an experiment. We chose suppression intervals longer than 0.4 s to trigger an AI measurement. We present our data using ΔAdapt , the length of the longest and most recent suppression period, as index for the difference in adaptation to the left and right adaptor (see Fig. 2). However, data analysis was also performed using the total difference in perceptual suppression of the left and the right adaptor. The results obtained by these two types of analyses are highly similar.

2.6. Data analysis

2.6.1. Recording of data

The computation of ΔAdapt and ΔAI was based on the timestamps marking changes in the status of the left and right arrow key, which were polled with an average sampling rate of 120 Hz and a maximal error of 11 ms.

2.6.2. Exclusion of single trials and data sets

Both MIB and AIs are known to show high within- and between-subject variability (see, e.g., Carter & Pattigrew, 2003; Lack, 1978; Loomis, 1972; Shimojo et al., 2001). To study the effect of MIB on AI formation, we had to use subjects and trials in which both phenomena were present. To maximize the chances of detecting an effect of perceptual suppression and physical removal of the adaptor on AI formation, we applied strict exclusion criteria for single trials and complete data sets. The precise choice of exclusion criteria, however, does not affect the findings we report.

The experiment was terminated after the initial training session if subjects reported no percept or only very faint and brief AIs, or after the 2.5 min baseline recording if subjects did not experience MIB at all. The data of the remaining subjects was processed offline.

Single trials were excluded from analysis if the recorded key presses revealed that:

- no or only very short AIs were perceived (i.e., AI persistence < 0.5 s);
- an AI was still perceived at or close to the termination of the AI measurement phase (i.e., AI persistence > measuring phase – 0.5 s);
- it was likely that a motor error occurred when reporting on AIs (i.e., persistence of one AI > persistence of other AI + 7 s);
- the induced AI was expected to have saturated (i.e., adaptation phase > 45 s).

A subject was completely excluded from the study if he/she:

- showed little suppression, i.e., less than 10 valid MIB trials;
- had a strong suppression bias, i.e., if the left or right adaptor triggered the AI measurement in less than three MIB trials.

Based on these criteria, the data of 16 of the 32 subjects were excluded from the analysis: 2 subjects did not perceive strong enough AIs, 4 did not experience MIB, 6 showed too little suppression, and another 4 had a strong suppression bias.

2.6.3. Analysis across vs. within subjects

The analysis across subjects is based on the data of the 16 subjects that satisfied all criteria. A within-subject analysis was performed with one subject (mk), who completed five experimental sessions on different days. This subject was very sensitive to MIB (i.e., the adaptors disappeared soon after presentation of the MIB stimulus and for quite long intervals) and was naïve to the purpose of the experiment. For the data analysis across all 16 subjects, we used only the data of the first experimental session of mk.

2.6.4. Type of data analysis—within-trial comparison

The analysis of the data is based on an initial within-trial comparison of the AIs induced by the left and the right adaptor. We chose this type of analysis because of the high variability observed when measuring AIs. This variability can partly be explained by certain factors that change over the course of the experiment. For example, the degree of fixation, blinking, fatigue or attention are known to affect AI formation. These factors will affect the two AIs measured within the same trial equally, but their effect on AIs recorded in different trials may differ greatly. Thus, the first step in our data analysis was to compare the left and right AIs measured within a trial. Combined data from different trials were analyzed later.

3. Results

Using single subject data, we first show results to characterize the distribution of MIB suppression intervals and the dependence of the AI duration on the physical exposure time to an adaptor. Then, we present data describing the effect of perceptual suppression on relative AI persistence. Finally, we present data on relative AI intensity.

3.1. MIB suppression periods

The subjective quality of MIB can be evaluated by studying the periods of perceptual suppression recorded during the MIB baseline measurement. This is illustrated by mk's data recorded in five experimental sessions (Fig. 3). The left adaptor was suppressed for 18% of the stimulation time (Fig. 3: top, left), the right for 13% (Fig. 3: top, right). Thus, mk had a small bias for suppression of the left adaptor. The histograms show that this is due to the left adaptor being suppressed often for very short periods. Longer suppression intervals occur equally often for the left and right adaptor. This is reflected by higher mean and median values for the suppression of the right than of the left adaptor.

As discussed above, the periods when one and only one of the adaptors is perceptually suppressed are relevant. Only these periods could possibly induce a difference in AI formation. One and only one of the adaptors was suppressed for only 20% of the presentation time (Fig. 3: bottom, right). The difference in the total suppression time and the time that only one of the adaptors was suppressed is due to those instances when both adaptors were suppressed. The distribution of suppression intervals for one and only one adaptor is significantly shifted towards shorter intervals when compared to that of suppression intervals recorded independently for the left and the right adaptor (Fig. 3: bottom, left). The

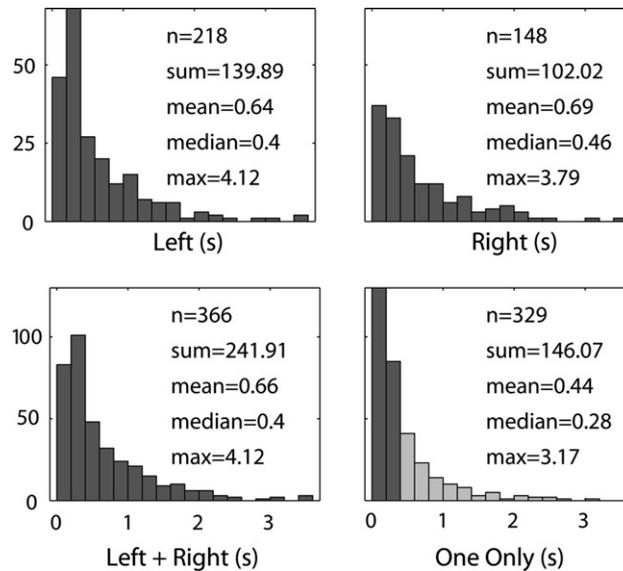


Fig. 3. Histograms of suppression intervals for one subject (mk). Histograms are based on MIB baseline recordings of five experimental sessions (total presentation time = 750s, bin width: 0.2s). Upper left: suppression intervals of left adaptor. Upper right: suppression intervals of right adaptor. Lower left: independently summed suppression intervals of left and right adaptor. Lower right: suppression intervals of one and only one adaptor (*exclusive or*). Light grey bars mark those intervals that would have triggered an AI measurement in MIB trials (>0.4 s). The number of suppression instances and the summed, mean, median, and maximal suppression interval (in seconds) are given in each panel.

shape of the distribution illustrates the trade-off mentioned above: although it would be best to measure AIs after long differences in suppression of the left vs. right adaptor, such events occur only very rarely. To collect a useful amount of data per experiment, a relatively low threshold for suppression difference (i.e., $\Delta\text{Adapt} > 0.4$ s, indicated by light grey bars) had to be used, despite the fact that the data are expected to be more noisy for shorter differences.

3.2. Effect of physical adaptation on AI persistence

The strength of AIs is known to increase with prolonged physical exposure to an adapting stimulus (see, e.g., Kelly & Martinez-Uriegaz, 1993). We verified this for our stimulation protocol using AI persistence. Fig. 4 shows this with the data from a single subject (mk) for both the PB and the MIB condition. Despite high variability in AI duration, data show that the longer mk was physically exposed to an adaptor, the longer the AI he perceived. Fitting an exponential constrained to pass through the origin to both sets of data results in very similar time constants and asymptotic values for these two conditions (time constant of 3.84 and 3.55s and asymptotic values of 7.52 and 7.53s for PB and MIB condition, respectively). Thus, despite the absence of blue dots during the PB condition, the AIs measured under the two conditions behaved very similarly.

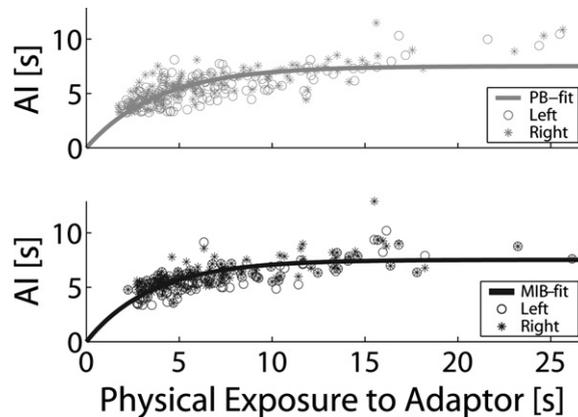


Fig. 4. The longer the physical exposure to the adapting stimulus, the stronger the AI persistence (illustrated with data from subject mk). Top panel: Absolute AI persistence as a function of the physical exposure to the adapting stimulus (less or equal to trial duration) in the PB condition. Bottom panel: Absolute AI persistence as a function of the physical exposure to the adapting stimulus (equal to trial duration) in MIB condition. Circles and stars represent AIs induced by the left and right adaptor ($n = 276$), respectively.

3.3. Relative AI persistence

3.3.1. Single subject

The effect of a physical (PB) and perceptual (MIB) difference in adaptation on the difference in AI persistence is illustrated in Fig. 5 for subject mk. $\Delta AI(s)$, the difference in AI persistence induced by the left and right adaptor, is plotted as a function of $\Delta Adapt(s)$, the difference in adaptation triggering the AI measurement, for all PB (light grey stars) and MIB (dark grey circles) trials ($n = 139$). Positive (negative) values denote that the left (right) adaptor or AI was perceived longer. The mean ΔAI plotted over the mean $\Delta Adapt$ is indicated as a medium grey star (circle) for the PB (MIB) condition. The mean value of $\Delta Adapt$ indicates that the subject did not have a strong suppression bias. The sum of the AI measurements triggering suppression is approximately the same for the left and right adaptor. However, the mean value of ΔAI indicates that the subject did have a bias with respect to AI perception. The mean ΔAI for both PB and MIB is about $-0.4s$, which means that right AIs were generally perceived longer.

Further analysis of AI persistence is based on computing a linear regression for ΔAI expressed as function of $\Delta Adapt$. The Null hypothesis to be evaluated is that ΔAI is independent of $\Delta Adapt$, i.e., that the slope of a linear regression through the data is not different from zero. The PB data should reveal a positive slope, reflecting the effect of physical adaptation time onto AI formation. Thus, for the PB condition, the Null hypothesis would have to be rejected. If perceptual visibility had a significant effect on AI persistence, the same trend expected for the PB data should be observed for the MIB data.

The best linear fit (least mean squares) through the PB data has a slope of 0.79 with 95% confidence intervals lying between 0.67 and 0.90. Thus, the Null hypothesis can be rejected for the PB condition ($r^2 = 0.56$, $F = 172.97$, $p < 0.5 \times 10^{-3}$). The best linear fit through the MIB data has a slope of -0.06 with 95% confidence intervals lying between -0.16 and 0.03. The Null hypothesis can thus not be rejected for the MIB condition ($r^2 = 0.012$, $F = 1.66$, $p = 0.2$). Thus, our method is

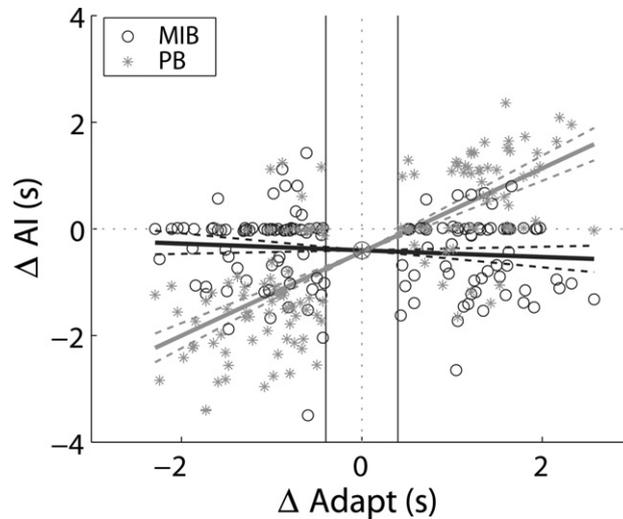


Fig. 5. Effect of physical (PB, light grey stars) and perceptual difference (MIB, dark grey circles) in adaptation on AI persistence for one subject (mk). $\Delta AI(s)$, the difference in duration of the AI between left and right adaptors is plotted as a function of $\Delta Adapt(s)$, the difference in visibility of the left and right adaptors. The solid lines represent the best linear fits through the data sets ($n = 139$ per condition). The dashed lines mark the 95% confidence intervals. Vertical lines through the x values of -0.4 and $0.4s$ indicate that an AI measurement was only triggered if the left or the right adaptor was suppressed (MIB) or removed (PB) for more than $0.4s$. The medium grey circle and star indicate the mean ΔAI over the mean $\Delta Adapt$ for the MIB and PB condition, respectively. As expected, the slope of the PB regression has a positive value, reflecting the effect of physical exposure time on AI formation. The slope of the MIB regression is not significantly different from 0.

sensitive enough to detect the expected effect of the physical adaptation time difference on AI persistence in the PB condition. However, no such effect is detected for the MIB condition.

3.3.2. Across subjects

The data obtained for an additional 15 subjects are presented in Fig. 6. Each panel contains a single subject's data recorded in one experimental session. As explained for Fig. 5, ΔAI is plotted as a function of $\Delta Adapt$ for PB and MIB trials, where the linear regressions through PB and MIB data sets are represented by solid lines. The individual slopes for the PB data are positive for all subjects. This reflects the known dependency of the physical adaptation time on AI persistence. However, the slopes for the MIB data are only positive for about half (7/15) of the subjects. This is expected if MIB does not influence AI persistence: by chance, individual slopes are expected to vary around the value of 0, yielding approximately 50% positive and 50% negative regression slopes. When computing the average of the 16 regression slopes for the PB data, the value obtained is—as expected—positive (0.691, std. of mean ± 0.106) and significantly different from 0 ($p < 1 \times 10^{-5}$). However, the average slope for the MIB data (0.038, std. of mean: ± 0.083) is not significantly different from 0 ($p = 0.65$). In addition, we found no correlation between the slopes obtained for MIB and PB data.

When computing the average slopes for the PB and MIB data across subjects, each subject's data and thus each experimental session, is weighted equally, although the amount of data

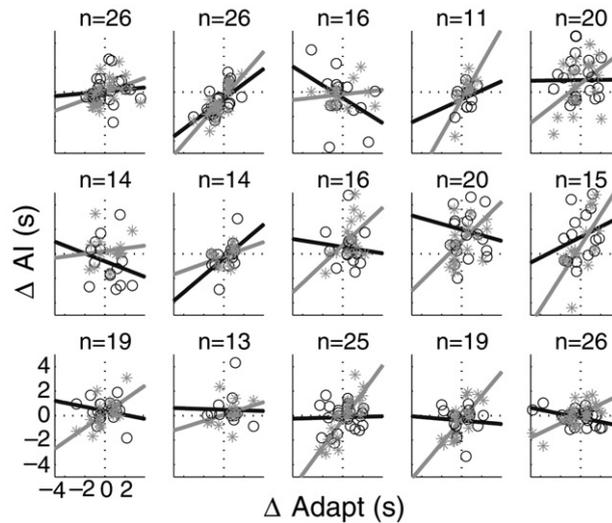


Fig. 6. Effect of physical difference (PB, light grey stars) and perceptual difference (MIB, dark grey circles) in adaptation on AI persistence for 15 additional subjects. Each panel displays ΔAI as a function of $\Delta Adapt$ for the valid trials obtained from a single subject in a single experimental session (see also Fig. 5). The number of valid trials per condition is given in the title of each panel. All axes have the same scale (numerical values are indicated in the panel in the lower left corner). The best linear fit through the data of the MIB and PB condition is indicated by the dark and light solid grey lines, respectively. As expected, the regression slope of all subjects has a positive value for the PB condition. For the MIB condition, only about half (7/15) of the regression slopes have a positive, while the rest have a negative value. This is to be expected if MIB does not affect AI persistence (see text).

obtained per subject differs (as indicated in Fig. 6). Because the computation of the slopes for a single subject is based on little data, the 16 underlying values themselves are not very reliable. An alternative approach is to first combine all data recorded for all subjects and then compute a linear regression. To do so, the data of each subject must first be transformed, to account for biases a subject may have. To eliminate these biases we subtracted the mean of $\Delta Adapt$ and ΔAI of each subject from each value of $\Delta Adapt$ and ΔAI , respectively. Fig. 7 displays the transformed data ($n = 305$) of all 16 subjects. The linear regression for the PB data (left, light grey stars) and the MIB data (right, dark grey circles) is indicated by solid lines, 95% confidence intervals by dashed lines. The best linear fit (least mean squares) through the PB data has a slope of 0.64 with 95% confidence intervals lying between 0.52 and 0.76. The best linear fit (least mean squares) through the MIB data has a slope of -0.03 with 95% confidence intervals lying between -0.15 and 0.09 . The Null hypothesis can be rejected for the PB condition ($r^2 = 0.27$, $F = 109.57$, $p < 0.5 \times 10^{-3}$), but not the MIB condition ($r^2 = 0.001$, $F = 0.29$, $p = 0.59$).

3.4. AI intensity

So far our data analysis focused on AI persistence as a measure for AI strength. The same type of result was obtained when analyzing the AI intensity data. AI intensity was measured using a two alternative forced-choice (2AFC) paradigm, leading to a binary measure for AI intensity (left

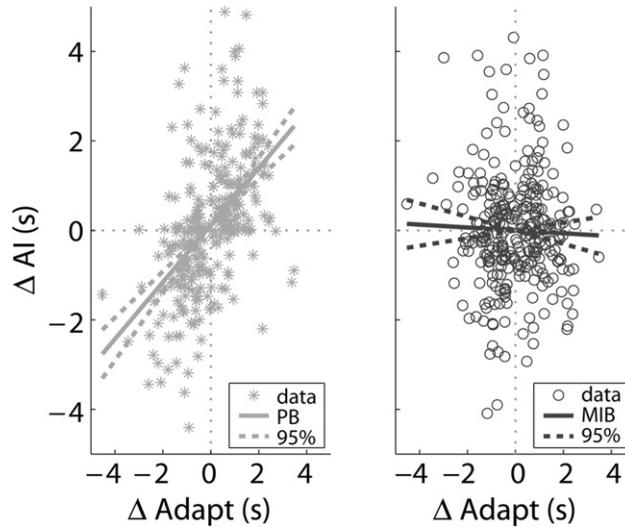


Fig. 7. Transformed data of all 16 subjects. For each subject, ΔAI minus mean ΔAI is plotted over $\Delta Adapt$ minus mean $\Delta Adapt$. Left: PB data (light grey stars). Right: MIB data (dark grey circles). The solid lines show the linear regression, the dashed lines mark the 95% confidence intervals. The slope of the PB regression has a positive value, reflecting the effect of physical exposure time on AI formation. The slope of the MIB regression is not significantly different from 0.

AI more intense vs. right AI more intense). The analysis of AI intensity is based on studying the distribution of $\Delta Adapt$ intervals for which the AI evoked by the left or right adaptor was rated as most intense.

In the top row of Fig. 8 the left (right) panel displays the distribution of $\Delta Adapt$ in PB (MIB) trials for which subjects reported to have perceived a more intense right AI (dark grey, inverted bars) or left AI (light grey, upright bars). The bins closest to $\pm 0.4s$ are represented strongest, because trials with short $\Delta Adapt$ occurred more often than those with long $\Delta Adapt$. Subjects did not show a significant suppression bias. They did show a bias for perceiving a more intense right AI in the MIB condition, but not the PB condition. The bottom panels show normalized data from the panel above: for each bin the number of trials when the right AI was perceived to be more intense was subtracted from that when the left AI was perceived to be more intense and then divided by their sum. Thus, a y value of +1 (–1) means that for all trials in this $\Delta Adapt$ interval the left (right) AI was perceived as more intense, while a value of 0 indicates that the left and right AI were rated most intense equally often. The PB data show clearly the expected trend: the longer one adaptor was presented compared to the other, the more likely it was that it induced the more intense AI. Such a trend does not occur for the MIB condition. When the left adaptor was suppressed, subjects were about equally likely to perceive the right or left AI to be more intense. Surprisingly, when the right adaptor was suppressed, subjects were more likely to perceive the right AI to be more intense.

The analysis of AI intensity data recorded for subject mk revealed similar results: a clear effect of the physical difference in adaptation on AI intensity, and no effect of a perceptual difference in adaptation (data not shown). In summary, the results obtained by averaging across 16 individual subjects are in perfect accord with those described in detail for one subject.

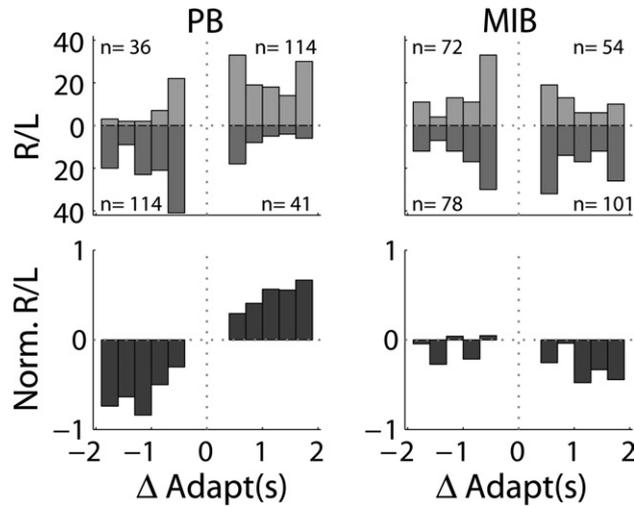


Fig. 8. Histograms analyzing the 2AFC intensity rating over ΔAdapt intervals. The left and the right plots display the results for the PB and MIB condition, respectively. Bins of 0.3 s width range from -1.9 s (left triggering suppression) to $+1.9$ s (right triggering suppression). Top panels: distributions when the left (light grey, upright bars) and right (dark grey, inverted bars) AI was rated as most intense. The values in the corner of the four quadrants are the sum of their histograms, i.e., indicate how often the left/right AI was rated as more intense when the left/right adaptor was perceived longer. Bottom panels: Normalized histograms $[(L - R)/(L + R)]$. As expected, the adaptor presented for the longest time tended to yield the more intense afterimage. There is no such trend when the adaptor was suppressed by MIB.

4. Discussion

Aftereffects have been referred to as the “psychologist’s microelectrode” (Frisby, 1979) because they permit the noninvasive study of the relationship between perception and its underlying neural processes. Crick and Koch suggested that aftereffects induced by invisible stimuli constitute a useful tool in the search of the NCC (Crick & Koch, 1995; see also Blake, 1997). Here, we present experiments that study the effect of MIB of the inducing stimuli on the persistence and intensity of negative AIs.

We measured AIs immediately after one of two moderate-intensity, colored adaptors had been temporarily perceptually suppressed due to MIB. Playback (PB) trials, in which the adaptors were physically removed in a way mimicking the perception in MIB trials served as controls. As expected, the intensity and duration of AIs were reduced by the physical removal of the adaptor in the PB condition. However, equivalent perceptual suppression of the adaptor in the MIB condition did not reduce intensity or persistence of AIs.

As discussed in Section 2, we assumed that the more recent the physical or perceptual removal of the adaptor, the stronger its impact on the strength of the AI. Accordingly, we analyzed ΔAI as a function of ΔAdapt , the duration of the suppression interval that triggered the AI measurement (see Fig. 2). It could be argued, however, that the total difference in suppression between the left and the right adaptor during the adaptation phase constitutes a more appropriate measure for the difference in adaptation. We performed this analysis and obtained highly similar results for both AI persistence and AI intensity.

In summary, MIB does not disrupt the formation of afterimages. To discuss the implications of this finding, we will briefly review what is known about the formation of AIs and the sites involved in MIB.

4.1. *The formation of AIs*

Although AIs have been investigated for decades (see, e.g., Goethe, 1810/1970; Helmholtz, 1867/1962; Hering, 1874/1964), the processes underlying their formation are still not completely understood. The term “afterimage” is used to describe a family of phenomena which are likely to have different origins (Gerrits, van Erning, & Eijkman, 1988; Sakitt, 1976). At least two types of AIs must be differentiated: positive and negative AIs. A positive AI has a color and contrast polarity similar to that of the adapting stimulus. Positive AIs are induced by brief exposure to high intensity stimuli and are mostly attributed to photochemical bleaching in rods and cones (Brindley, 1962; Sakitt, 1976). We study negative AIs induced by prolonged exposure to non- or only weakly bleaching adaptors. The origin of negative AIs is generally assigned to neural adaptation (see, e.g., Wilson, 1997). Electrophysiological and psychophysical studies clearly demonstrate a contribution of retinal and LGN cells to the formation of negative AIs (Virsu & Laurinen, 1977; Virsu, Lee, & Creutzfeld, 1977). Assigning the formation of AIs purely to such low-level processes is consistent with the general finding that AIs do not transfer from the adapted eye to the non-adapted one. However, it has been suggested that cortical neurons also participate in and modulate the formation of negative AIs (Chan, Crutch, & Warrington, 2001; Davis, 1973; Hayhoe & Williams, 1984; Loomis, 1972, 1978; Schiller & Dolan, 1994; Weiskrantz, 2002; Weiskrantz, Cowey, & Hodinott-Hill, 2002). For example, Hadjikhani et al. (1998) presented functional magnetic resonance imaging data indicating that perception of negative AIs correlates with an increase in the blood oxygen level-dependent signal in an extrastriate area (which they call V8), much more than it does in earlier cortical areas. Suzuki and Grabowecy (2003) demonstrated that attention during adaptation modulates consecutive AI perception and argue that this effect is most likely caused by attentional modulation of polarity sensitive cells in IT and V4. Shimojo et al. (2001) showed that negative AIs can be induced in the absence of a local adaptor, i.e., without the direct stimulation of retinal cells. The authors studied AIs induced by perceptually filled-in surfaces. The strength of these AIs correlated with the subjects’ ability to perceptually fill-in the illusory contour generated by several inducers in the adaptation phase. It did not correlate—but instead rived—with the AIs produced by the inducers themselves. The authors concluded that this type of negative AI must be caused by the adaptation of a cortical representation of surface. Such findings imply that low-level processes alone are insufficient to explain the generation of AIs and that these phenomena are more complex than often assumed.

4.2. *What do we know about MIB?*

Even in the absence of a moving background, perceptual fading of stimuli eventually occurs under stabilized viewing conditions (Troxler, 1804). MIB, however, occurs suddenly, may be experienced after very brief observation times, and is stronger when the contrast of the target is high (Bonneh et al., 2001). Because slowly moving or flickering target stimuli can also be perceptually suppressed by MIB, this phenomenon is unlikely to be caused by retinal, local adaptation (Bon-

neh et al., 2001). Furthermore, MIB is not a form of local masking, as suppression still occurs if a protection zone around the stationary stimuli prevents spatial interactions between the moving elements and the target. Recent studies revealed that MIB has no detectable effect on orientation-specific adaptation (Kouhsari et al., 2004) and that MIB is sensitive to cues that affect grouping (Bonneh et al., 2001; Hsu et al., 2004) or surface completion (Graf, Adams, & Lages, 2002), processes usually assigned to striate and even extrastriate visual areas. Thus, there is evidence that MIB is a cortical phenomenon, and there are indications that it may originate in extrastriate cortex.

This study demonstrates that MIB does not significantly affect the strength of negative AIs induced by moderate intensity, colored adaptors. Therefore, in the presence of a physical adaptor, there is no significant contribution to the formation of negative AIs beyond the sites mediating MIB. These findings are consistent with an earlier study suggesting that binocular rivalry does not reduce afterimage formation (Lack, 1978) and the common hypothesis that the generation of negative AIs is mostly due to low-level processes. Our study is restricted in that it only allows us to infer the relative location of sites governing AI formation and MIB. A better understanding of the precise mechanisms underlying MIB will enhance the power of this study. Our results show that AI perception is not modulated by the perceived exposure time. It has been suggested that afterimages constitute a tool for defining the visual NCC (Kirschfeld, 1999). Our findings imply that the neurons mediating the formation of AIs are not part of the visual NCC.

Acknowledgments

This research was supported by grants from NSF, NIMH, the Mind Science Foundation, and the Keck Foundation to C.K., and Swiss National Science Foundation Grant 31-67980.02 to D.C.K. We thank Peter König for providing laboratory equipment and Wolfgang Einhäuser, Jörg Hipp, and Farshad Moradi for valuable feedback on the manuscript.

References

- Anstis, S., Rogers, B., & Henry, J. (1978). Interactions between simultaneous contrast and coloured afterimages. *Vision Research*, *18*, 899–911.
- Bachmann, T. (2000). *Microgenetic approach to the conscious mind*. Amsterdam: Johns Benjamins.
- Blake, R. (1997). What can be perceived in the absence of visual awareness?. *Current Directions in Psychological Science*, *6*, 157–162.
- Blake, R., & Fox, R. (1974). Adaptation to invisible gratings and the site of binocular rivalry suppression. *Nature*, *249*, 488–490.
- Blake, R., & Logothetis, N. K. (2002). Visual competition. *Nature Reviews Neuroscience*, *3*, 13–21.
- Blakemore, C. B., & Campbell, F. W. (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.
- Bonneh, Y. S., Cooperman, A., & Sagi, D. (2001). Motion-induced blindness in normal observers. *Nature*, *411*, 798–801.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.
- Breitmeyer, B. G. (1984). *Visual masking: An integrative approach*. Oxford: Oxford University Press.
- Breitmeyer, B. G., & Ögmen, H. (2000). Recent models and findings in visual backward masking: A comparison, review, and update. *Perception & Psychophysics*, *62*, 1572–1595.

- Brindley, G. S. (1962). Two new properties of foveal after-images and a photochemical hypothesis to explain them. *Journal of Physiology*, *164*, 168–179.
- Carter, O. L., & Pattigrew, J. D. (2003). A common oscillator for perceptual rivalries. *Perception*, *32*, 295–305.
- Chan, D., Crutch, S. J., & Warrington, E. K. (2001). A disorder of colour perception associated with abnormal colour after-images: A defect of the primary visual cortex. *Journal of Neurology, Neurosurgery and Psychiatry*, *71*, 515–517.
- Craik, W. J. K. (1940). Origin of afterimages. *Nature*, *148*, 512.
- Crick, F., & Koch, C. (1995). Are we aware of neural activity in the visual cortex?. *Nature*, *375*, 121–123.
- Crick, F., & Koch, C. (2003). A framework for consciousness. *Nature Neuroscience*, *6*, 119–126.
- Davis, P. (1973). The role of central processes in the perception of visual after-images. *British Journal of Psychology*, *64*, 325–338.
- Engel, A. K., Fries, P., König, P., Brecht, M., & Singer, W. (1999). Temporal binding, binocular rivalry, and consciousness. *Consciousness and Cognition*, *8*, 128–151.
- Frisby, J. P. (1979). *Seeing: Illusion, brain, and mind*. Oxford: Oxford University Press.
- Gerrits, H. J., van Erning, L. J., & Eijkman, E. G. (1988). Afterimages: A collective term for percepts of different origin. *Experimental Brain Research*, *72*, 279–286.
- von Goethe, J. W. (1810/1970). *Theory of colours* (English translation by Eastlake, C. L.). Cambridge, MA: The MIT Press.
- Graf, E. W., Adams, W. J., & Lages, M. (2002). Modulating motion-induced blindness with depth ordering and surface completion. *Vision Research*, *42*, 2731–2735.
- Hadjikhani, N., Liu, A. K., Dale, A. M., Cavanagh, P., & Tootell, R. B. (1998). Retinotopy and color sensitivity in human visual cortical area V8. *Nature Neuroscience*, *1*, 235–241.
- Hayhoe, M. M., & Williams, D. R. (1984). Disappearance of afterimages at ‘impossible’ locations. *Perception*, *13*, 455–459.
- He, S., Cavanagh, P., & Intriligator, J. (1996). Attentional resolution and the locus of visual awareness. *Nature*, *383*, 334–337.
- He, S., & MacLeod, D. I. (2001). Orientation-selective adaptation and tilt after-effect from invisible patterns. *Nature*, *411*, 473–476.
- von Helmholtz, H. (1867/1962). *Treatise on physiological optics* (Vol. 3, English translation by Southall, J. P. C.). New York: Dover Press.
- Hering, E. (1874/1964). *Outlines of a theory of the light sense* (English translation by Hurvich, L. M.). Cambridge, MA: Harvard University Press.
- Hsu, L. C., Yeh, S. L., & Kramer, P. (2004). Linking motion-induced blindness to perceptual filling-in. *Vision Research* (in press).
- Kelly, D. H., & Martinez-Uriegaz, E. (1993). Measurement of chromatic and achromatic afterimages. *Journal of the Optical Society of America*, *10*, 29–37.
- Kirschfeld, K. (1999). Afterimages: A tool for defining the neural correlate of visual consciousness. *Consciousness and Cognition*, *8*, 462–483.
- Koch, C. (2004). *The quest for consciousness: A neurobiological approach*. Denver, Colorado: Roberts Publishers.
- Kouhsari, L. M., Moradi, F., Zand-Vakili, A., & Esteky, H. (2004). Orientation-selective adaptation in motion-induced blindness. *Perception*, *33*, 249–254.
- Lack, L. C. (1978). Selective attention and the control of binocular rivalry. In Barendregt, et al. (Eds.). *Psychological studies* (Vol. 11). The Hague, The Netherlands: Mouton Publishers.
- Lamme, V. A., Zipser, K., & Spekreijse, H. (2002). Masking interrupts figure-ground signals in V1. *Journal of Cognitive Neuroscience*, *14*, 1044–1053.
- Lee, S. H., & Blake, R. (2002). V1 activity is reduced during binocular rivalry. *Journal of Vision*, *2*, 618–626.
- Leopold, D. A., & Logothetis, N. K. (1996). Activity changes in early visual cortex reflect monkey’s percepts during binocular rivalry. *Nature*, *379*, 549–553.
- Logothetis, N. K., Leopold, D. A., & Sheinberg, D. L. (1996). What is rivaling during binocular rivalry?. *Nature*, *380*, 621–624.
- Logothetis, N. K. (1998). Single units and conscious vision. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *353*, 1801–1818.

- Loomis, J. (1972). The photopigment bleaching hypothesis of complementary afterimages: A psychophysical test. *Vision Research*, *12*, 1587–1594.
- Loomis, J. (1978). Complementary afterimages and the unequal adapting effects of steady and flickering light. *Journal of the Optical Society of America*, *68*, 411–416.
- Macknik, S. L., & Livingstone, M. S. (1998). Neuronal correlates of visibility and invisibility in the primate visual system. *Nature Neuroscience*, *1*, 144–149.
- Maffei, L., Fiorentini, A., & Bisti, S. (1973). Neural correlate of perceptual adaptation to gratings. *Science*, *182*, 1036–1038.
- Merikle, P. M., Smilek, D., & Eastwood, J. D. (2001). Perception without awareness: Perspectives from cognitive psychology. *Cognition*, *79*, 115–134.
- Metzinger, T. (1995). The problem of consciousness. In T. Metzinger (Ed.), *Conscious experience*. Imprint Academic, Schöningh.
- Metzinger, T. (2000). In T. Metzinger (Ed.), *Neural correlates of consciousness*. The MIT Press, Cambridge, Massachusetts.
- Movshon, J. A., & Lennie, P. (1979). Pattern-selective adaptation in visual cortical neurones. *Nature*, *278*, 850–852.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Rajimehr, R. (2004). Unconscious orientation processing. *Neuron*, *41*, 663–673.
- Rolls, E. T., Tovee, M. J., & Panzeri, S. (1999). The neurophysiology of backward visual masking: Information analysis. *Journal of Cognitive Neuroscience*, *11*, 300–311.
- Sakitt, B. (1976). Psychophysical correlates of photoreceptor activity. *Vision Research*, *16*, 129–140.
- Schiller, P. H., & Dolan, R. P. (1994). Visual aftereffects and the consequences of visual system lesions on their perception in the rhesus monkey. *Visual Neuroscience*, *11*, 643–665.
- Shimojo, S., Kamitani, Y., & Nishida, S. (2001). Afterimage of perceptually filled-in surface. *Science*, *293*, 1677–1680.
- Suzuki, S., & Grabowecky, M. (2003). Attention during adaptation weakens negative afterimages. *Journal of Experimental Psychology: Human Perception and Performance*, *29*, 793–807.
- Toet, A., & Levi, D. M. (1992). The two-dimensional shape of spatial interaction zones in the parafovea. *Vision Research*, *32*, 1349–1357.
- Troxler, D. (1804). Über das Verschwinden gegebener Gegenstände innerhalb unsers Gesichtskreises. In K. Himley & J. A. Schmidt (Eds.), *Ophthalmologische Bibliothek* (Vol. 2). Jena: Fromann.
- Virsu, V., & Laurinen, P. (1977). Long-lasting afterimages caused by neural adaptation. *Vision Research*, *17*, 853–860.
- Virsu, V., Lee, B. B., & Creutzfeld, O. D. (1977). Dark adaptation and receptive field organization of cells in the cat lateral geniculate nucleus. *Experimental Brain Research*, *27*, 35–50.
- Weiskrantz, L. (2002). Prime-sight and blindsight. *Consciousness and Cognition*, *11*, 568–581.
- Weiskrantz, L., Cowey, A., & Hodinott-Hill (2002). Prime-sight in a blindsight subject. *Nature Neuroscience*, *5*, 101–102.
- Wilke, M., Logothetis, N. K., & Leopold, D. A. (2003). Generalized flash suppression of salient visual targets. *Neuron*, *39*, 1043–1052.
- Wilson, H. R. (1997). A neural model of foveal light adaptation and afterimage formation. *Visual Neuroscience*, *14*, 403–423.
- Wolfe, J. M. (1984). Reversing ocular dominance and suppression in a single flash. *Vision Research*, *24*, 471–478.
- Zeman, A. (2001). Consciousness. *Brain*, *124*, 1263–1289.