ATTENTION AND AWARENESS:

VISUAL PSYCHOPHYSICS AND AVERSIVE CONDITIONING IN HUMANS

Thesis by

Naotsugu Tsuchiya

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2006

(Defended July 18, 2005)

Acknowledgement

As I write this acknowledgement, I remember all the exciting events and troubles that happened to me one after another since I came to Caltech. I started as a volunteer researcher in April 2000 and joined the CNS program as a graduate student during that summer. Over these five years, I've been fortunate to be surrounded by my colleagues and friends, whom I want to acknowledge here.

Mark Konishi and Shin Shimojo have been my role models of successful Japanese researchers in the American academic community. I owe a lot to their advice and encouragement throughout my graduate life at Caltech. Especially, Shin's recommendation to Christof Koch in April 2000 made it possible for me to work as a volunteer. My Caltech life started from there.

I introduced myself in a five-minute meeting to Christof, trying to convince him that I was desperately interested in the problem of consciousness and that I wanted to work in his lab. It was the most intense five minutes. He took the risk of taking me. I only hope I have brought more fruit than troubles to him. I cannot imagine such a great (and colorful) advisor like him if I were in Japan.

Jochen Braun was the advisor for my first project on visual motion and attention (Chapter 3). I thank him for his generous invitation to England and Germany, which makes me think that I should be generous to him even if our manuscript will never be published. Working with Geraint Rees at University College London on this motion and attention project with functional magnetic resonance imaging (fMRI) was very exciting. Through this project, I learned a lot about visual psychophysics and fMRI.

McKell Carter taught me how to conduct electrical shock experiments with Caltech students (and sometimes with our mentor!) in the name of science. Throughout the collaboration with him, we published one paper and submitted another one now in the second round of the review (in Chapter 4), at the cost of a few friends who really hated the electrical shocks. I thank more than 200 subjects who participated in our 'uncomfortable but not painful' experiments (See Method 4.4.1).

Ryota Kanai at the Universiteit Utrecht, the Netherlands, has been my best friend, colleague, and rival since we shared our office in Kyoto University, Japan, in 1999-2000. With Ryota, I translated Christof's "The Quest for Consciousness" into Japanese as a private project. We will continue our collaboration.

I thank all of my colleagues in the Koch lab, past and present. Connie Hofstoetter and Farshad Moradi were very helpful with my project on afterimage using continuous flash suppression. I enjoyed the collaboration with Melissa Saenz as well as two of my SURF students, Bobby Rohrkemper and Dania Adamuszek. Kamran Diba invited me to give a talk at Rutgers, The State University of New Jersey, which was a very good experience for me. I organized several parties in the Koch lab, which went very well usually. I was pleased when people liked the sushi, which Dirk Walther and I prepared at the farewell party for the Reddy sisters. There is one thing I do regret, however: it was my mistake to introduce karaoke to Patrick Wilken and Wei Ji Ma at the farewell party for Chunhui Mo.

I am grateful to my committee for two candidacy exams and the defense: Yaser Abu-Mostafa, Pietro Perona, David Dubowitz, Joe Bogen, Ralph Adolphs, John O'Doherty, Steve Quartz, Richard Andersen.

My wife, Noriko, has supported and encouraged my research throughout. She has been keeping an eye on my health, both mental and physical. Nori's comments on my practice talks before conferences have been always critical and useful. We shared one joy when I received the Best Student Presentation Award in the ASSC 9 conference in June of 2005. Nori and I often let out our stress at parties with our Japanese friends. These Japanese friends were my secret pool for many of my psychophysics experiments.

Finally, I want to thank my family, Hisao, Junko, and Rie for their generous support and encouragement.

Abstract

We studied the neuronal correlates of consciousness by characterizing the role of attention and awareness in three psychophysical experiments. First, we investigated the role of visual awareness in the formation of afterimages, phenomena believed to occur in the retina. Visibility of the afterimage-inducing stimuli was manipulated by a powerful dichoptic suppression technique, *continuous flash suppression*, which allows us to project visual stimuli onto the retina without subjects noticing them at all, sometimes longer than three minutes. We found that reliably suppressing the inducer weakens afterimage strength. Paradoxically, trial-to-trial variability in visibility did not correlate with the intensity of afterimage. As afterimages are enhanced when attention is withdrawn from the adaptor, the opposite effects between awareness and attention were demonstrated. Second, we examined visual motion processing outside the focus of spatial, top-down attention using a dual-task paradigm. Attentional effects in motion processing were characterized by our novel *wavelet* motion stimuli. Our stimuli effectively activate neurons in the first stage of motion processing, while they are poor stimuli for higher motion processing. Using a contrast-masking paradigm, we found that attention mainly affected the strength of inhibition for high-contrast motion stimuli in an orientation-specific, but not direction-specific manner, presumably reflecting the physiological properties for divisive inhibition within the primary visual cortex. Third, we characterized the role of awareness in classical aversive conditioning. Subjects associated previously neutral auditory stimuli (CS) with aversive mild electric shocks (US). We used skin conductance response, an index for autonomic arousal, as implicit measure for

the conditioned response. In *delay* conditioning, CS was paired with *delayed* but *overlapping* US, while in *trace* conditioning CS was followed by US after a three-second *temporal gap*. We intermixed these two CSs with another control CS that never predicted US to examine whether awareness plays different roles depending on the temporal relationships between CS and US. Subjects expressed their shock expectancy using their gaze direction, from which we inferred the onset of awareness of CS-US contingency. By aligning the skin conductance response with the onset of awareness, we found that trace, but not delay, conditioning coincided with the onset of awareness.

Table of contents

List of Figures

List of Tables

List of abbreviations

Chapter 1

NCC	Neuronal correlates of consciousness

Chapter 2

BR	Binocular rivalry
FS	Flash suppression
CFS	Continuous flash suppression
s.e.m.	standard error of the mean
cpd	cycles per degree

Chapter 3

V1	Primary visual cortex
MT	Middle temporal area (sensitive to visual motion)

Chapter 4

CS	Conditioned stimulus
CR	Conditioned response
US	Unconditioned stimulus
UR	Unconditioned response

1. General Introduction

1.1. The neuronal correlates of consciousness

The problem of consciousness is one of the biggest mysteries left unsolved in modern science. *How* and *why* can our brains, being physical systems, produce subjective phenomenon, conscious experience, or qualia? This is the central question, both in neuroscience and philosophy, and finding the answer to this question is the ultimate goal of my scientific career. Why can some neuronal structures, i.e., the cerebral cortex, produce consciousness while others, i.e., the cerebellum, cannot? Why does the cerebral cortex sometimes, i.e., after waking up and before going to sleep, produce consciousness, while it doesn't at other times, i.e., during dreamless sleep. These questions are sometimes referred to as the Hard Problem of consciousness because there appears to be no way to directly answer or even attack these questions rigorously using scientific reductionism (Chalmers, 1996).

Instead of tackling this extremely difficult puzzle of 'why' physical brains produce any subjective phenomena, we can start from the Easy Problem, trying to characterize the properties of consciousness and its correlates in the nervous system (Crick and Koch, 1998). The search for the neuronal correlates of consciousness, the NCC, is the first step towards solving the deep puzzle. Here, the NCC is defined as "the minimal set of neuronal mechanisms or events jointly sufficient for a specific conscious percept or experience" (Koch, 2004). (For problems associated with defining the term, NCC, see (Chalmers, 2000).)

To infer the mechanisms of the NCC, in this thesis, we obtained verbal or behavioral reports from human subjects and measured behavioral or autonomic response,

rather than directly measuring neuronal activity. For studies of the NCC, it is most effective to focus on experimental situations where conscious percepts probed by behavioral reports[1] are dissociated from sensory inputs or behavioral outputs (Frith et al., 1999). The focus of this thesis is to identify such experimental paradigms that are ideal for the search of the NCC.

In the course of our search for the NCC, we expect to obtain significant insight about the mechanisms that are key to understanding consciousness. For example, attention is one such mechanism (Posner, 1994). Selective attention amplifies attended information and suppresses ignored information, serving as a gate for consciousness. When we are occupied with something and do not pay attention, we may fail to notice an important event. How does attention modulate neuronal activity and prevent incoming information from entering into consciousness? Are we aware of only what we attend to, or do we have a subjective experience outside the focus of attention? Is there an instance where attending to a stimulus and becoming aware of the stimulus result in opposite results? What kinds of tasks require attention or awareness for successful performance? These questions constitute the main theme of the three experiments described in this thesis. The relationship between attention and awareness is further discussed in the final chapter.

---

[1] The term "consciousness" possesses broad and multiple meanings; however, throughout this thesis we use 'consciousness' to mean the specific content of consciousness, distinguishing from consciousness as such (i.e., whether a person is conscious at all or not) or conscious states (i.e., awake, dreaming, hallucinating, etc). Specifically, the content of consciousness we refer to in the following chapters includes "visibility of stimulus projected to the eye" (in Chapter 2), "sense of visual motion" (in Chapter 3), and "contingency awareness that one stimulus predicts an aversive event in near future" (in Chapter 4).

1.2. Our three approaches to the search for the NCC

To study the NCC, the *content of consciousness (reports)* has to be dissociated from *sensory inputs* (*stimulation*) or *behavioral outputs* (*responses*) (Frith et al., 1999). (Figure 1.1). In each of the following three chapters, we focus on a situation where one of these three components is held constant while the other two components are experimentally controlled.

One aspect of the study in Chapter **2** deals with *unconscious* visual processing. Visual input to one eye does not give rise to conscious percepts because continuously flashed images to the other eye strongly suppress it. The degree of unconscious visual processing of the invisible input is measured in terms of the intensity of the associated afterimages[2]. This represents one approach to help isolate the NCC, where *the content of consciousness is kept constant* (subjectively invisible, unconscious) during adaptation. While subjectively it is always invisible, sensory input has some parameters that are varied, and their effects are measured by the extent of subsequent aftereffects (Blake and Fox, 1974; Lehmkuhle and Fox, 1975; Wade and Wenderoth, 1978; O'Shea and Crassini, 1981; He et al., 1996; He and MacLeod, 2001; Hofstoetter et al., 2004; Rajimehr, 2004; Moradi et al., 2005). The method of constant unconsciousness with varied inputs is commonly employed in the studies of unconscious processing, such as blindsight (Cowey and Stoerig, 1991, 1995), masking (Macknik and Livingstone, 1998; Morris et al., 1998;

[2] Although after-*images* from invisible stimuli sounds like an oxymoron, it is not. Neuronal activity in the retina does not correlate with consciousness but it is sufficient to produce afterimages, to some extent. In Chapter 3, however, we show that retinal activity is not sufficient for the formation of afterimages with full-blown intensity.

Ohman and Soares, 1998; Macknik and Martinez-Conde, 2004) and priming (Bar and Biederman, 1998; Dehaene et al., 1998; Dehaene et al., 2001; Naccache et al., 2002; Melcher et al., 2005).

In Chapter **3**, we study vision outside the focus of spatial attention. In this experiment, *identical visual stimuli are presented (constant stimulation)* while subjects' attentional states are manipulated; subjects attend only to the stimuli presented at the periphery (single-task situation), or both at the periphery and at the fixation (dual-task situation). Task performance on the peripheral stimuli is compared under these two different attentional states (Braun, 1994; Lee et al., 1997; Lee et al., 1999b; Lee et al., 1999a; Zenger et al., 2000; Li et al., 2002). In the past decades, the neural responses to identical sensory inputs have been characterized when it is task-relevant (attended) or task-irrelevant (ignored) in monkeys with single-unit recording (Moran and Desimone, 1985; Spitzer et al., 1988; Maunsell, 1995; Treue and Maunsell, 1996; McAdams and Maunsell, 1999b, a; Seidemann and Newsome, 1999; Treue and Martinez Trujillo, 1999) and in humans with fMRI (Beauchamp et al., 1997; Watanabe et al., 1998a; Gandhi et al., 1999; Huk et al., 2001; Saenz et al., 2002). The constant input paradigm is very powerful, and it includes a hallmark of neurophysiological studies during bistable percepts (Logothetis and Schall, 1989; Leopold and Logothetis, 1996; Sheinberg and Logothetis, 1997). In this situation, identical input results in alternating percepts or interpretations over time. The NCC is straightforwardly defined as the neuronal activity that most closely follows in time the perceptual fluctuations reported by subjects.

Chapter **4** describes an effort to find a paradigm, where *performance in two tasks is held constant*, while the relationship between inputs and consciousness is different;

successful performance in one task, but not in the other, is critically dependent on consciousness, although the performance in two tasks and inputs in two conditions are very similar. Previous studies implied that consciousness plays a critical role in performing a task that requires holding information over a few seconds, while a similar task can be carried out automatically, independent of consciousness, if there is no temporal gap between stimulus presentation and behavioral response. In this chapter, we report that classical aversive conditioning depends on consciousness when there is a temporal gap between CS and US (trace conditioning), while it is independent of awareness if CS and US overlaps (delay conditioning). Our study is based on findings in the eyeblink conditioning literature (Clark and Squire, 1998; Clark et al., 2002). Notably, a similar relationship between consciousness and temporal gap is reported in clinical studies of blindsight (Marcel, 1993) and visual form agnosia (Milner and Goodale, 1995) (For a review (Rossetti, 1998).) When tested, these patients show detection and orienting performance well above chance for a target presented in their affected visual field without having phenomenal experience of the presence of the target. However, their residual ability is severely compromised if a brief gap is inserted between the presentation of the stimuli and the reaction to it. (Here, subjects are kept unconscious of the stimuli in both cases, while residual visuo-behavioral abilities are altered.)

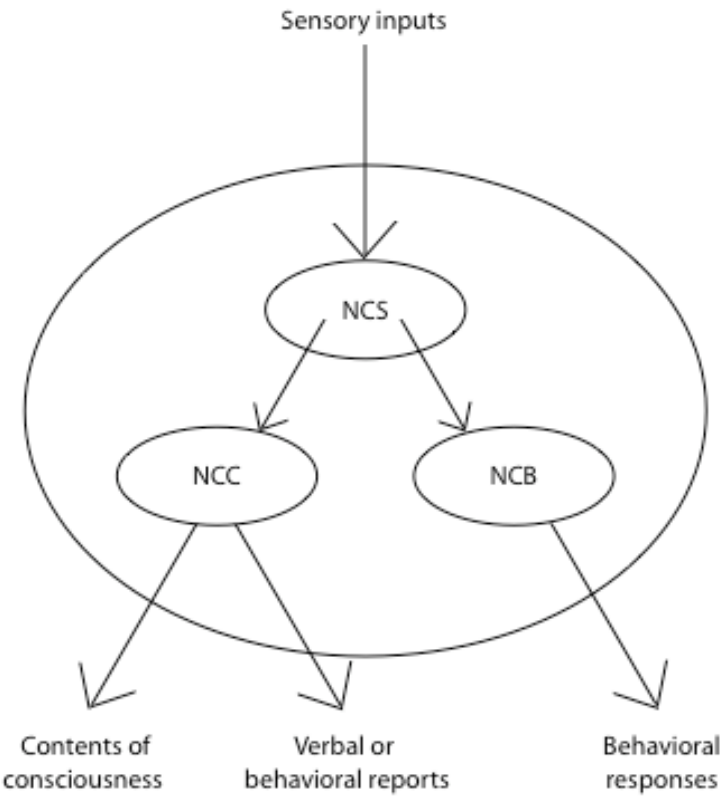

Figure 1.1 Scheme for the empirical studies for the NCC

For the study of the neuronal correlates of consciousness (NCC), the neuronal correlates of sensory input (NCS) and the neuronal correlates of behavioral response (NCB) have to be controlled experimentally. Ideally, two factors are kept exactly the same while the remaining one is varied[3]. If one of the three factors is held constant while the other two

[3] In binocular rivalry, NCS and NCB are held constant while NCC is varied. To do this, we need to compare the NCC across trials by equating behaviors that are used to probe subjects' perceptual flips (e.g., in one trial, pull the lever to indicate perception of a vertical grating, and in another trial, push the lever to indicate seeing a horizontal grating, and so on). Holding NCS and NCC while varying NCB is the classical example of implicit learning and implicit motor behavior. Even when NCS is varied, NCC and NCB can be held constant in the case of aftereffects from invisible stimuli and false memories (Frith et al., 1999).

are manipulated strategically, the NCC can be studied adequately as well. The contents of consciousness are probed by verbal or voluntary behavioral reports. Implicit behavioral responses can be used to assess the degree of unconscious processing.

1.3. Structure of the thesis

In the following chapters, we begin by reviewing some background information and the motivation for the study. After the summary of our findings, the research paradigms are described, and the results are presented.

2. Role of awareness and attention in the formation of negative afterimages

2.1. Overview

What are the differences between the neuronal activities that give rise to consciousness and those that are buried in the unconscious? The study of unconscious visual processing is one of the most powerful approaches to characterizing the neuronal correlates of consciousness. Although much of the computation in the brain is performed without reaching consciousness, the exact nature of the difference in neuronal activity for conscious and unconscious percepts remains unclear. Understanding the nature of unconscious processing would provide us with insights as to what kinds of neuronal processing are *insufficient* for conscious perception. This kind of eliminative approach is a big step towards elucidating *necessary* and *sufficient* conditions for the NCC (Crick and Koch, 1995; Milner and Goodale, 1995).

In this chapter, we present and characterize a novel and powerful technique for unconscious visual presentation, Continuous Flash Suppression (CFS) (See Section 2.4). With CFS, the duration of perceptual invisibility is greatly extended, from a few seconds to several minutes, at least 10-fold longer compared to standard binocular rivalry suppression. Further, using a probe detection technique, we characterized the depth of continuous flash suppression to be much deeper than the standard binocular rivalry suppression. As CFS allows us to present stimuli unconsciously in a reliable and sustained manner, it is very useful for the study of unconscious processing.

Using this tool, we revisited the question of the neuronal origin of negative afterimages. Though negative afterimages have been traditionally thought to arise from

the adaptation process in the retina, our results, as well as other recent studies, show that the formation of afterimages involves processing that accesses information from both eyes. Although the reliability of suppression correlated with the reduction of afterimage, trial-to-trial variations in the duration the adaptor was visible did not correlate with the intensity of the afterimage. Thus, awareness seems only indirectly responsible for the formation of afterimages.

One intriguing aspect of our findings is the fact that the effects of attention and awareness in the formation of afterimages are opposite (See Section 2.5.5). While withdrawing attention enhances the intensity of afterimages, reliable unawareness weakens it. Recently, there has been an emerging body of studies that support the view that neuronal mechanisms of attention and awareness can be dissociated under some conditions. Some psychophysical studies show that attention can be attracted by invisible cues in normal subjects (McCormick, 1997; Lambert et al., 1999; Rajimehr, 2004) and blindsight patients (Kentridge et al., 1999a, b, 2004) (but the exogenous cues made invisible by binocular rivalry suppression do not attract attention (Schall et al., 1993)). Further evidence from psychophysics shows that attention is necessary for invisible stimuli to modulate sensory processing (Naccache et al., 2002; Montaser-Kouhsari and Rajimehr, 2005) and that feature-based attention spreads to the invisible stimuli (Kanai et al., 2004; Melcher et al., 2005). Finally, recent studies using backward masking while measuring event-related potentials show that an initial 130-200 msec response component can be modulated by attention and awareness independently (Koivisto et al., 2005b; Koivisto et al., 2005a).

Our and other studies imply post-retinal components for the formation of afterimages. On the other hand, recent electrophysiological studies show that the retinal network adapts to complex spatiotemporal patterns, which were thought to arise only post-retinally (Smirnakis et al., 1997; Hosoya et al., 2005). We dissociated retinal components, which should not be affected by CFS, and post-retinal components, which should be suppressed by CFS, by holding retinal adaptation constant while *manipulating the duration of suppression*. We found that post-retinal components of afterimages decay and buildup much more rapidly than the perceptual afterimages and retinal adaptation process, indicating previous estimates mainly reflect the retinal adaptation (Section 2.5.6).

Because afterimages have traditionally been considered as retinal phenomena, nearly nothing is known about the extent to which neuronal mechanisms in cortex and elsewhere are involved in the formation, expression, and modulation of afterimages. In the context of the search for the NCC, the study of afterimages may advance our understanding of the neuronal basis of attention and awareness.

2.2. Summary

Illusions that produce perceptual suppression despite constant retinal input are used to manipulate visual consciousness. Here we report on a powerful variant of existing techniques, continuous flash suppression. Distinct images flashed successively at ~10 Hz into one eye reliably suppress an image presented to the other eye. The duration of perceptual suppression is at least ten times greater than that produced by binocular rivalry. Using this tool, we show that the strength of the negative afterimage of an adaptor was reduced by half when it was perceptually suppressed by input from the other eye. The more likely the adaptor was completely suppressed, the larger the reduction of the afterimage intensity. Paradoxically, trial-to-trial visibility of the adaptor did not correlate with the degree of reduction. Our results imply that formation of afterimages involves neuronal structures that access input from both eyes, but that do not correspond directly to the neuronal correlates of perceptual awareness.

2.3. Introduction

The question of the neuronal correlates of conscious perception has seen renewed interest over the last decades (Koch, 2004). One powerful tool in this area are illusions that give rise to effects that are measurable, yet are not, or only occasionally, consciously seen (Blake and Fox, 1974; He et al., 1996; He and MacLeod, 2001; Hofstoetter et al., 2004; Rajimehr, 2004). In backward masking (Macknik and Livingstone, 1998; Macknik and Martinez-Conde, 2004), inattentional blindness (Mack and Rock, 1998; Rees et al., 1999), motion-induced blindness (Bonneh et al., 2001), binocular rivalry (Logothetis and Schall, 1989; Leopold and Logothetis, 1996; Sheinberg and Logothetis, 1997; Lumer et al., 1998; Tong et al., 1998; Polonsky et al., 2000; Tong and Engel, 2001; Pasley et al., 2004; Williams et al., 2004), and flash suppression (Wolfe, 1984; Sheinberg and Logothetis, 1997; Kreiman et al., 2002; Wilke et al., 2003), an image is presented to one or both eyes of the observer, yet is not seen.

Binocular rivalry (BR) is a popular method used to determine if a visual aftereffect occurs after or before the neuronal site for the suppression of rivalry for those stimuli (Blake and Fox, 1974; Lehmkuhle and Fox, 1975; Lack, 1978; O'Shea and Crassini, 1981; Wiesenfelder and Blake, 1990; Blake, 1995; Moradi et al., 2005). In BR, two different images are shown to the two eyes, and the subject's percept alternates between the two images (Blake and Logothetis, 2002). The strength of the aftereffect when the adaptor is presented to one eye and remains plainly visible throughout the adaptation period is compared to the aftereffect when the adaptor is suppressed by the input to the other eye.

However, the duration and timing of perceptual suppression are difficult to control because of the stochastic nature of rivalry. Flash suppression (FS) (Wolfe, 1984; Sheinberg and Logothetis, 1997; Kreiman et al., 2002; Wilke et al., 2003) provides better control over the timing of suppression, but at the price of shorter periods of suppression, too brief to produce strong aftereffects. Furthermore, FS requires a pre-adapting period, preventing complete unawareness of the adaptor. Here we combine aspects of both BR and FS into a potent procedure we term continuous flash suppression (CFS). We continuously flash different images rapidly into one eye, while the input to the corresponding location in the other eye remains the same (see demonstration at http://www.klab.caltech.edu/~naotsu/CFSdemo.html). Most observers fail to see the image in one eye, even though it is present for a long time, sometimes for several minutes.

We used CFS to examine the neuronal site for negative afterimages. These are vivid percepts that demonstrate the tenuous link between physical stimuli and their associated subjective percepts. A variety of evidence supports their origin among neurons in the retina (Alpern and Barr, 1962; Brindley, 1962; Loomis, 1972; Sakitt, 1976; Virsu and Laurinen, 1977; Loomis, 1978; Wilson, 1997) or lateral geniculate nucleus (LGN) (Kelly and Martinez-Uriegas, 1993). In particular, negative afterimages do not transfer across eyes, nor is their strength reduced by suppression of the inducing image by pressure blinding (Craik, 1940; Lack, 1978) (but see (Virsu and Laurinen, 1977)). Neither BR (Lack, 1978) nor motion-induced blindness (MIB) (Hofstoetter et al., 2004) reduces either the duration or the strength of afterimages. All of these observations suggest that afterimages are retinal phenomena.

However, both BR as well as MIB only suppress the adaptor intermittently. By using CFS, we asked what happens when the adapting stimulus is completely suppressed from awareness. We found that when an adaptor was reliably suppressed by CFS, the intensity of the negative afterimage of the adaptor was reduced by half. Our results imply that formation of afterimages involves neuronal structures that access input from both eyes, but that do not correspond directly to the neuronal correlates of perceptual awareness.

2.4. Characterizing Continuous Flash Suppression (CFS)

2.4.1. Experiment 1: Prolonged invisibility by CFS

In Experiment 1, we compared the initial duration of stimulus suppression in CFS and BR without pre-exposure to the suppressed image (Methods for all experiments in chapter 2 are described in Section 2.7 Methods). While a constant, gray image—here a picture of an angry face—is presented to one eye, CFS stimuli, composed of different Mondrian patterns, were presented at the corresponding location in the other eye (Figure 2.1). Each Mondrian was replaced by a different pattern every 100 msec. Seventeen naïve subjects pressed a button as soon as any part of the gray figure became visible. The mean initial suppression time in 16 trials was 4.3 sec for BR, and 56.0 sec for CFS, which is 13 times longer (paired *t*-test, *t*-score = 4.81, d.f. = 16, $P < 0.001$). In 40 out of 272 CFS trials, subjects reported seeing no part of the gray image at all for the full three-minute trial! As we treated those trials' suppression times as 180 sec, we are likely to underestimate the true duration of the initial period during which the adaptor remains invisible.

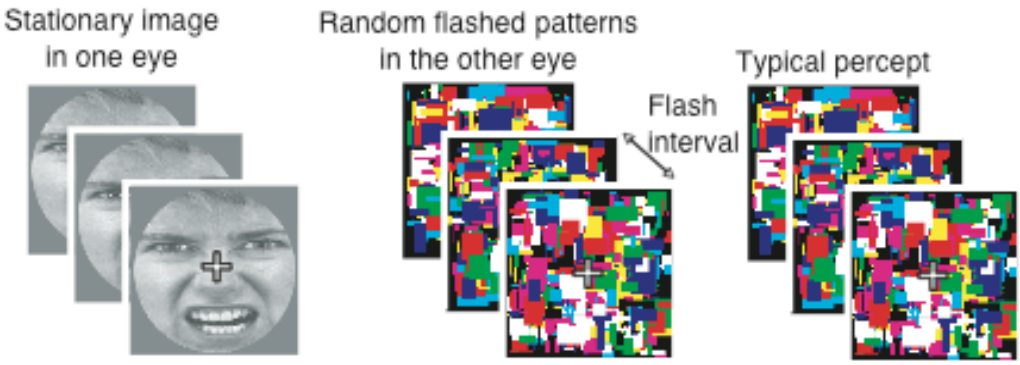

Figure 2.1 A basic setup for Continuous flash suppression

A stationary gray stimulus was presented to one eye, while different, colored Mondrian patterns were flashed into the other eye every 100 msec. Subjects fixated on the central cross and pressed a button to report when the gray figure started to become visible. Initial suppression duration in continuous flash suppression (CFS) was more than ten times longer than in binocular rivalry (BR), using the same stimulus but with a stationary Mondrian pattern.

2.4.2. Experiment 2: Optimal flash interval for CFS

An important parameter for successful CFS is the flash interval between successive presentations of distinct Mondrian patterns. In Experiment 2, we studied dominance during a one-minute observation period as a function of different flash intervals (Figure 2.2 a and b). The most effective flash interval for long suppression was between 80 and 320 msec (~3-12 Hz flash rate). We used a 10 Hz flash rate for Experiments 3-11.

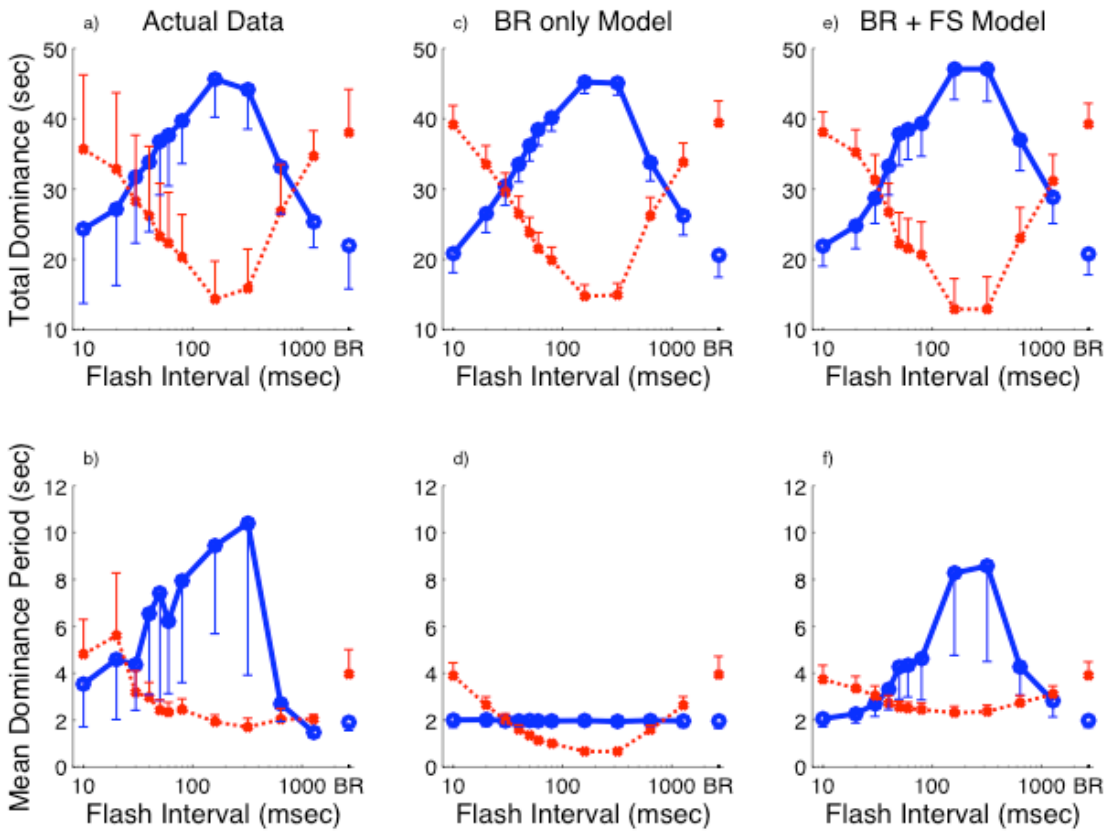

Figure 2.2 Optimal flash interval for continuous flash suppression.

(**a,c**, and **e**) Total dominance duration (TD) and (**b, d**, and **f**) mean of each dominance period (MD) were plotted as a function of flash intervals. Error bars correspond to standard error. (**a** and **b**) Data from the actual experiment. Four subjects tracked the visibility of the Gabor patch during one-minute continuous viewing when any part of the Gabor pattern was visible (thin dotted lines) or the whole Gabor pattern was completely invisible (thick solid lines) while Mondrian patterns were continuously flashed into the other eye with a inter-flash intervals indicated on the x-axis. The right-most two points are for the case of binocular rivalry (BR). (**c** and **d**) 200 computer-simulated one-minute

trials at each flash interval without a flash suppression component. The strength of the Mondrian was modulated to fit TD. (**e** and **f**) With the FS component included in our phenomenological model (Sections 2.4.2 to 2.4.5), the peak of TD and MD is located around 80-320 msec flashed intervals (200 simulated trials). The strength of the Mondrian stays constant but the probability of successful FS was fit to the TD data (Figure 2.3). Without the FS component, these TD and MD curves cannot both be fitted simultaneously.

2.4.3. Is CFS just a stronger version of binocular rivalry?

Levelt's second proposition of binocular rivalry (BR) states that the stimulus strength of 'A' primarily determines the mean dominance duration (MD) of the stimulus 'B' (MD_B) presented to the other eye, with little effect on the MD of the stimulus A (MD_A) (Levelt, 1965; Fox and Rasche, 1969; Blake, 1977; Blake and Logothetis, 2002). Let the number of perceptual switches be *n*, the total experimental duration be *T* (here it was one minute), the total dominance duration of A and B be TD_A and TD_B respectively, then,

$$T = n * (MD_A + MD_B) = TD_A + TD_B$$

Levelt's second proposition predicts that when the strength of stimulus A increases into A', $MD_{B'}$ decreases (i.e., stronger suppression) while $MD_{A'}$ remains constant (i.e., $MD_{A'} = MD_A$). Accordingly the frequency of perceptual reversals *n increases* into *n'* (i.e., $\frac{n'}{n} = \frac{MD_A + MD_B}{MD_{A'} + MD_{B'}} > 1$) and,

$$T = n' * (MD_{A'} + MD_{B'}) = TD_{A'} + TD_{B'}$$

,where $TD_{A'} = n' * MD_{A'} = n' * MD_A > n * MD_A = TD_A$, that is, the total dominance of A *increases*. For the total dominance of B, $TD_{B'} = T - TD_{A'} < T - TD_A = TD_B$, that is, the total dominance of B *decreases*).

If CFS is a straightforward extension of BR, one would expect the result of Figure 2.2 a and b, with the dependence of the total dominance (TD) on the flash interval reflecting the 'effective stimulus strength' of Mondrian flashes. In other words, the increase of TD_A would be accompanied by the decrease of MD_B and with little effect on MD_A. Our analysis of TD (Figure 2.2 a) and MD (Figure 2.2 b) as a function of flash

intervals shows that the extended TD of Mondrians is mainly due to the extended MD of Mondrian percepts. This is not what would be expected from a simple extrapolation of BR in which a “strong” stimulus primarily reduces the MD of the rival stimulus (here the Gabor patch).

We propose a simple model assuming that the prolonged MD of Mondrians, which depends on the flash intervals, can be explained by the combined effect of BR and repetitive FS (Figure 2.2 e and f). Adding the FS component, whose sensitivity depends on the flash interval (Figure 2.3), explains most of the variance of the data (Figure 2.2 e and f).

2.4.4. CFS is a combined effect of binocular rivalry and flash suppression

Can CFS be modeled in a quantitative manner as a combination of BR and FS? We started from a simple phenomenological model of BR (Levelt, 1965). During BR, the percept flips randomly between two interpretations with two statistical regularities: 1) Each period of dominance for a stimulus presented to one eye is well fit by a Gamma distribution and 2) each dominance period is dependent on the strength of the stimulus presented to the other eye. Levelt assumes a fourth-order Gamma distribution for the probability density functions:

$$f_{4,G}(t) = \frac{\lambda_M}{3!}(\lambda_M t)^3 \exp(-\lambda_M t), \quad (1)$$

$$f_{4,M}(t) = \frac{\lambda_G}{3!}(\lambda_G t)^3 \exp(-\lambda_G t), \quad (2)$$

where the subscripts G and M stand for the Gabor and Mondrian percept, respectively, and λ represents the *stimulus strength*[4]. The mean duration of the Gamma distribution is inversely related to λ, such that $MD_G = \frac{4}{\lambda_M}$. As the MD for Gabor and Mondrian under BR are about 4 and 2 sec (Figure 2.2 b), we chose $\lambda_G = 2, \lambda_M = 1$.

As in many rivalry situations (Blake and Logothetis, 2002), this model produces longer TD of Mondrian when the strength of Mondrian increases (Figure 2.2 c).

[4] *Stimulus strength* is an abstract concept proposed by Levelt. Levelt showed that the effect of luminance contrast on the dominance duration follows the relationship described in equations 1 and 2. Motion energy and spatial frequency roughly follow these rules (Blake and Logothetis, 2002, but see Bossink et al., 1993). In Section 2.4.4, we tested if the frequency of Mondrian flashes can be considered as stimulus strength. Note that the dominance period of the Gabor depends on the strength of Mondrian, λ_M (equation 1), and vice versa (equation 2).

However, the stronger the Mondrian stimulus (larger λ_M values), the shorter the MD of Gabor (equation 1), without any effect on the MD of Mondrian (equation 2). As a result, we cannot reproduce TD and MD at the same time just by changing the strength of Mondrian. For example, we can approximate TD (Figure 2.2 c) but fail miserably for MD (Figure 2.2 d) with this minimal model. Clearly, the MD of Mondrian is independent of flash intervals (Figure 2.2 d), which is quite different from the actual inverse U-shape function we observe (Figure 2.2 b). For Figure 2.2 c and d, we set the strength of Mondrian as shown in Table 2.1.

Flash interval (msec)	10	20	30	40	50	60	80	160	320	640	1280	BR
λ_M	1	1.5	2	2.5	3	3.5	4	6	6	2.5	1.5	1

Table 2.1 Fitted strength of Mondrian for each flash interval

The strength of Mondrian was adapted to reproduce the total dominance data (TD) for each of the 12 points in Figure 2.2 to obtain the TD in Figure 2.2c. Therefore, we had 12 degrees of freedom to fit the data.

2.4.5. A phenomenal model of CFS

In standard FS, the success of suppression crucially depends on the length of pre-adaptation. Wolfe showed a monotonically increasing sigmoidal relationship between the pre-adaptation duration and successful FS (see Fig. 6 in (Wolfe, 1984)).

We added an FS component to the above BR model. When the current percept is dominated by Mondrian patterns, each Mondrian flash 'refreshes' its percept with probability of $p(t_{FI})$ (t_{FI} is the flash interval). If the flash is successful, it resets the dominance period according to equation (2), maintaining the Mondrian percept. When the percept flips to that of a Gabor, the model asserts that Mondrian flashes are unable to reverse the percept to the Mondrian for some duration. We refer to this duration as the "refractory period," which we set to 2 sec. After the refractory period, each Mondrian flash to the other eye has the potential to flip the percept instantaneously (with no time delay) to a Mondrian. This flip occurs with probability, $p(t_{FI})$, which is dependent on flash interval (thus, the degree of freedom of this model is 12 (11 levels of flash intervals (excluding the BR condition) and the duration of the refractory period)). This minimal addition explains most of the variance of TD and MD (Figure 2.2 e and f). $p(t_{FI})$ (Figure 2.3) was derived to fit TD (Figure 2.2 e).

As long as the $p(t_{FI})$ function shows a similar shape, such as a sigmoidal on a log time scale, both TD and MD show an inverted-U shape simultaneously. In psychological terms, this means that our model can explain that the extended TD of Mondrian is caused by the extended MD of Mondrian, but not by the reduced MD of Gabor. The BR model introduced in Section 2.4.4 does not show such behavior.

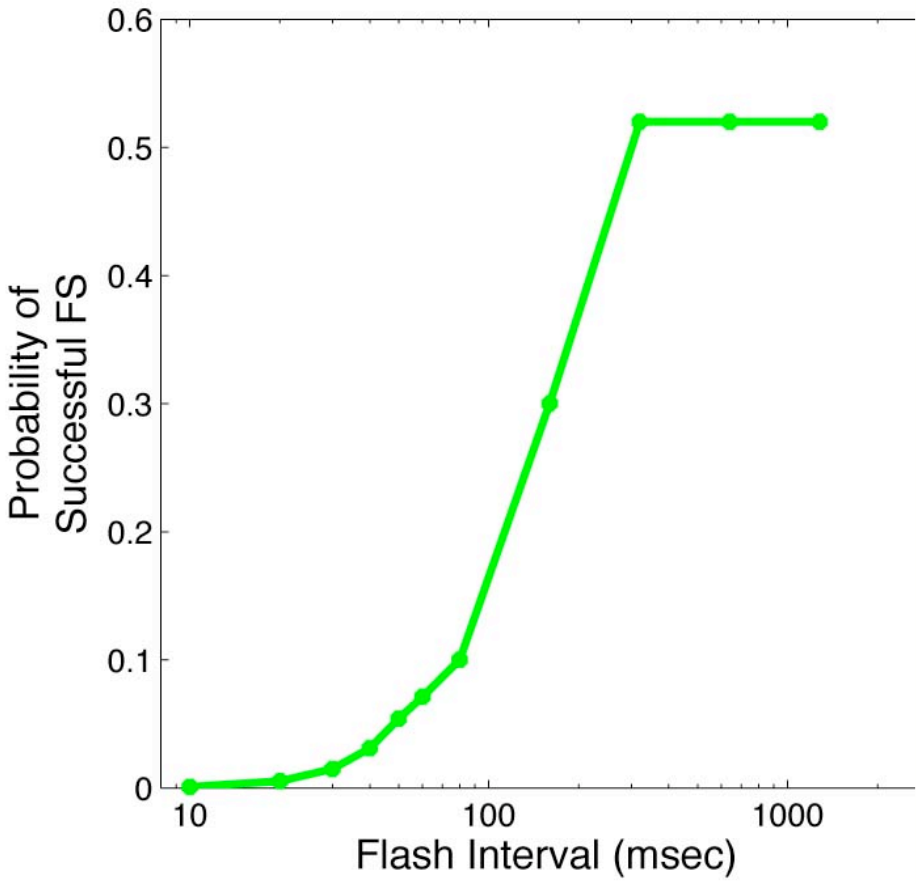

Figure 2.3 Probabilistic model of flash suppression

This monotonically increasing function was obtained by fitting the TD (Figure 2.2 e) for the actual data in Figure 2.2 a. Although the probability of suppression for each flash is low for small intervals, it accumulates.

2.4.6. Objective measurement of strength of CFS and binocular rivalry

In Sections 2.4.3 to 2.4.5, we argued that we could not explain the more than 10-fold longer perceptual suppression in CFS on the basis of the 'stimulus strength' theory of binocular rivalry. A model that incorporated flash suppression into binocular rivalry did, however, explain the relationship between the dominance duration and flash intervals at a quantitative level.

So far, we characterized CFS using subjective reports on whether one stimulus is consciously visible or not. While this method allows us to study the phenomenal visibility of CFS, it is not suited for quantitative examination of how 'strong' continuous flash suppression really is. For that purpose, measures of sensitivity, such as detection thresholds, obtained by an objective forced-choice procedure are ideal. Through Experiments 3 to 5, we used a probe detection task to infer the mechanisms of CFS. In terms of the depth of suppression, we characterized CFS in relation to single- and multiple- flash suppression, as well as binocular rivalry.

2.4.7. Experiment 3: Depth of suppression in CFS

In Experiment 3, we measured the depth of continuous flash suppression relative to binocular rivalry suppression using a probe technique (Fox and Check, 1972; Norman et al., 2000; Nguyen et al., 2001; Nguyen et al., 2003; Watanabe et al., 2004). In the CFS condition, a 10% stationary grey-scale baseline grating was projected to one eye and a 10Hz Mondrian stream, which suppressed the grating, to the other eye. After making sure no part of the grating was visible, subjects pressed a button, which triggered a 500 msec Gaussian contrast increment pulse superimposed onto the upper or lower half of the baseline grating (Figure 2.4 a). Subjects reported which location the probe appeared in a two-alterative forced choice manner. Threshold contrast increment was established using a two-down-one-up staircase method.

In Figure 2.4, we show the results comparing the three conditions. In the monocular-viewing condition without any suppression, the threshold was very low (1.9±0.2%, error is s.e.m.). For the binocular rivalry (BR) condition, in each trial, subjects waited until a stationary Mondrian pattern was perceptually dominant exclusively (with no hint of the grating) and, at that point, triggered presentation of the contrast increment. Under these conditions, the increment threshold (6.2±0.9%) was elevated only moderately for the BR condition consistent with previous results (Nguyen et al., 2001; Watanabe et al., 2004). In contrast, we measured a 20-hold increase in increment thresholds (44.7±10.3%) for the CFS condition.

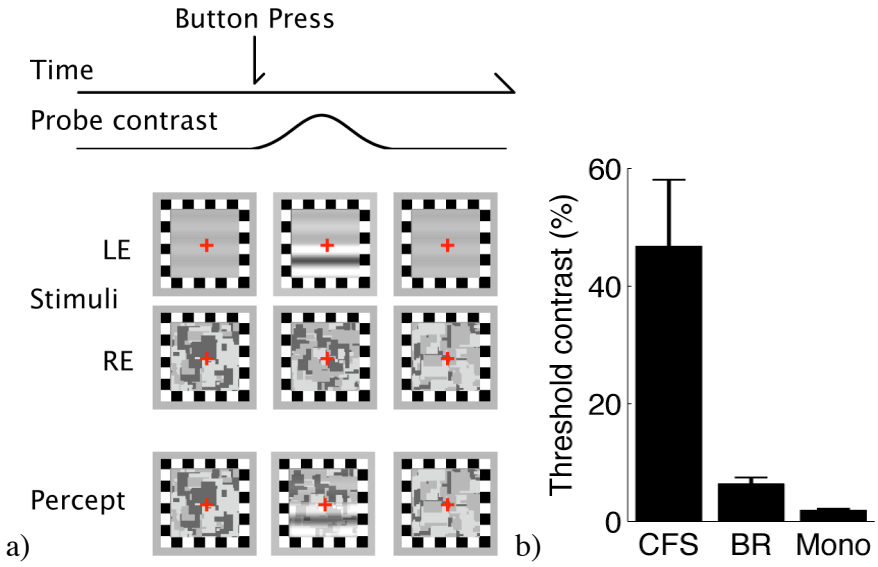

Figure 2.4 Depth of suppression during CFS

a) Physical stimuli to two eyes and typical percept in the CFS condition. Subjects viewed a 10% grating in one eye (in this example, left eye (LE)) and a 10Hz stream of grey-scale Mondrians in the other eye (right eye (RE)). Upon a button press, either upper or lower half of the grating increased its contrast gradually over 500 msec and subjects discriminated which field contained the probe (2AFC). In the binocular rivalry (BR) condition, stimulus to the other eye was a stationary Mondrian. In the monocular-viewing (Mono) condition, a blank field was presented to the other eye and contrast increment started 500 msec after the button press. **b**) Contrast increment thresholds for the three conditions.

These results clearly show that CFS does indeed produce much deeper suppression than does BR, consistent with the fact that durations of exclusive dominance with CFS are an order of magnitude longer than average dominance durations produced by BR (Experiment 1). However, this conclusion might appear inconsistent with Section 2.4.3 to 2.4.5, where we argued that CFS is not merely a sstrong version of BR. This apparent inconsistency can be resolved when we consider the role of 'neuronal adaptation', in bistable perception (Leopold et al., 2002; Blake et al., 2003). In Discussion (2.6.2), we consider the possibility that CFS avoids adaptation of neurons for the dominant percept, which triggers perceptual switches in standard binocular rivalry and other bistable percepts.

In Experiments 4 and 5, we dissected CFS into several components using the probe technique. In Experiment 4, we measured the depth of suppression in a single flash suppression episode as a function of flash timing, and we characterized its interaction with the binocular rivalry suppression. In Experiment 5, we asked whether multiple flashes sum their effects to amplify the depth of suppression. If so, how many flashes are necessary to reach the depth of continuous flash suppression?

2.4.8. Experiment 4: Depth of single flash suppression

In Experiment 4, we first dissected 'continuous' flashes into single flashes. Does a single flash elevate the detection threshold more than binocular rivalry? If so, how is it characterized by the timing between the onsets of a probe and a Mondrian flash? Previous studies of flash suppression (Wolfe, 1984; Wilke et al., 2003) showed a tight relation between the effectiveness of perceptual suppression and flash timing, but they did not study the relation between the depth of suppression and flash timing. Second, we asked if suppression following a single flash becomes stronger when rivalry suppression is already operating. Do flash suppression and rivalry suppression sum and, therefore, amplify the depth of suppression?

Using a paradigm similar to that used in Experiment 3 (See Methods) we measured the thresholds while manipulating the stimulus onset asynchrony (SOA) between the onset of a probe and a Mondrian flash. SOA ranged from -1000 to +500 msec, where negative SOA means that the Mondrian flash occurred first followed by the probe. Figure 2.5 shows the thresholds for the single flash suppression (FS-only) condition with filled triangles (error bars are s.e.m.). Single flash suppression elevated the thresholds only within a narrow time window, with the peak threshold elevation (12.7±1.4%) occurring at about SOA = +100 msec, which was still lower by 64% than CFS (34.8% for the three subjects who participated in this experiment) but higher by 132% than BR (5.5% for the three subjects who participated in this experiment). Note that the contrast of the probe was modulated in a Gaussian manner with its peak at SOA of 250 msec. In other words, a single flash suppressed the probe most strongly when it

was presented to the other eye 150 msec before the contrast of the probe becomes maximal. For SOA longer than 150 msec, the effects of single flash diminished dramatically, and with SOA > 300 msec, they reached almost the level of the monocular-viewing condition in Experiment 3 (a thin broken line and empty triangle). For SOA < -100 msec, the effects were comparable to that of binocular rivalry in Experiment 3 (a thick broken line and empty circle), suggesting that with this long negative SOA, the effects of flash suppression had already dissipated but binocular rivalry was still operating.

Next, we tested the effect of flash suppression after binocular rivalry was already induced. As a single flash was combined with binocular rivalry, the situation might be closer to CFS. Figure 2.5 shows the thresholds for this condition (FS + BR, filled circles). The main difference from the FS-only condition was found mainly at SOA > 100 msec. Up to SOA = 100 msec, no difference was observed between the two.

Observing Figure 2.5, we noticed three characteristics. First, when a Mondrian flash occurred too early (SOA < -100 msec), the thresholds were comparable with the BR condition (5.5%) in both the FS-only and the FS + BR conditions. At SOA = 500 msec, the threshold was also similar to the BR condition in the FS + BR condition. These thresholds reflected binocular rivalry because at these SOA, timings of a flash were too far from the optimal suppression. Second, the thresholds in the FS-only condition seem to agree relatively well with those in the FS + BR condition for SOA < 100 msec. This implies that the rising part of thresholds for both the FS-only and FS + BR conditions were due to the effects of FS and not due to the interactive effects between FS and BR. Third, for SOA > 150 msec, however, the thresholds for the FS-only condition dropped

quickly. Note that, at SOAs of 150 and 200 msec, the thresholds for the FS-only condition were comparable to the BR condition, while those for the FS + BR condition almost reached their peak. This implies that a single flash presented later than 150 msec on its own did not impair detection of the probe (no suppression effect), but combined with rivalry suppression, that single flash was highly effective.

To characterize the data more quantitatively, we fitted two Gaussian curves to the data with three assumptions. First, in the FS + BR condition, the thresholds would be fitted with a Gaussian function with asymptote value being the threshold in the BR condition in Experiment 3. Second, the thresholds for the rising part in the FS-only condition (SOA < 50 ms) would be captured by the above Gaussian fit. Thus, we made no attempt to fit the data here. Third, the thresholds for the falling part of the FS-only condition would be fitted by a second Gaussian curve with a different asymptote value being the monocular-viewing condition in Experiment 3.

Based on these assumptions, we first fitted a Gaussian function of the form $y = a \times \exp\left(-\left(\frac{x-b}{c}\right)^2\right) + 5.5$ to the thresholds in the FS+BR condition, where y is the threshold, x is SOA, a is the peak threshold, b is the SOA at the peak, c is the width of the tuning, and 5.5% is the threshold in the BR condition. Using 13 points, we obtained the best fit, shown as a thick line in Figure 2.5 (degree of freedom = 10, a = 7.9%, b = 85.5 msec, c = 201.0 msec, and R^2 = 0.85). Next, we fit another Gaussian to the 7 data points for SOA ≥ 50 msec in the FS-only condition, with fixing a and b to be 7.9% and 85.5 msec and changing the asymptote value to the monocular-viewing threshold (1.8%). We obtained the best fit for c to be 74.9 msec with R^2 = 0.97 (degree of freedom = 6). Thus,

when FS was combined with BR, the temporal tuning of suppression was broadened from 74.9 to 201.0 msec, implying a nonlinear interaction between FS and BR.

Our conclusions in this section are threefold. First, flash suppression impaired the probe detection much more than did binocular rivalry, but only within a relatively narrow time window, peaking at SOA = 85.5 msec. But even at its maximum effectiveness, flash suppression was weaker than CFS. This time dependency contrasted starkly with the time-course of binocular rivalry, where the depth of suppression is reported to be constant at all points throughout a dominance phase (Fox and Check, 1972; Norman et al., 2000). This difference provides further evidence that flash suppression and binocular rivalry suppression are mediated by different mechanisms (Section 2.4.3 - 2.4.5). Second, flash suppression did not summate with binocular rivalry suppression before the peak of the flash suppression. Third, threshold elevation in the FS-only condition dropped to the monocular-viewing level with a tuning width of 74.9 msec, yet it synergistically interacted with binocular rivalry suppression to broaden the temporal tuning curve (201.0 msec) to deepen suppression. These second and third conclusions suggest that the very robust strength of suppression accompanying CFS is due to the cooperative inhibition among multiple flashes, which we studied in Experiment 4.

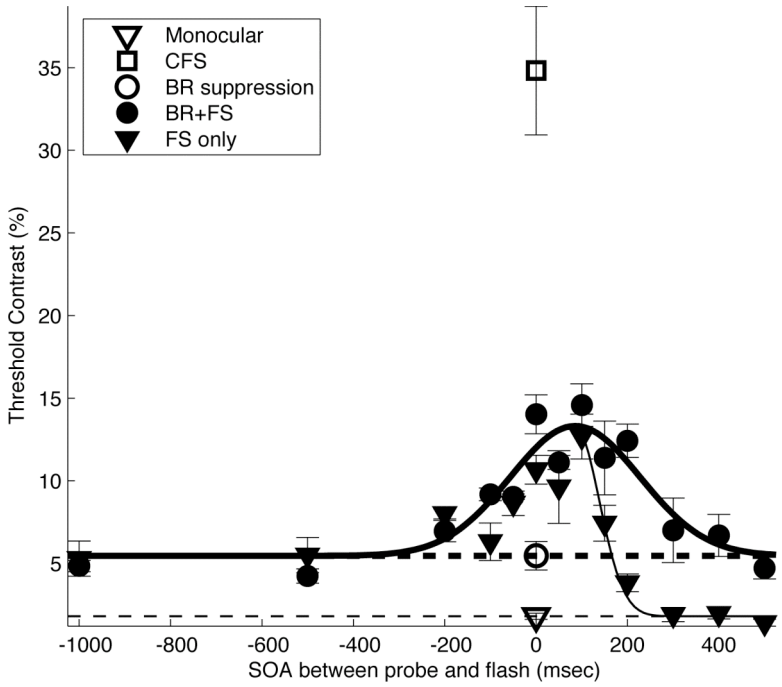

Figure 2.5 Does a single flash suppression summate with rivalry suppression?

Contrast increment thresholds as a function of SOA between the onset of a probe and a Mondrian flash. For comparison, the results in Experiment 3 are re-plotted here with empty symbols (Monocular, CFS at 10Hz, and BR suppression). Filled symbols show the results in the single flash suppression (FS-only, circles) and the FS combined with binocular rivalry suppression (FS+BR, triangles) conditions, respectively (error bars are s.e.m.). The solid curves were the best-fit Gaussian functions to the FS+BR (thick) and FS-only (thin) conditions, respectively. The dotted lines are the asymptotic values for the Gaussian, taken from the thresholds for the BR suppression (thick) and the monocular-viewing (thin) conditions in Experiment 3, respectively. When the SOA was negative, a Mondrian flash came first, then a probe was presented. For both the FS-only and FS+BR conditions, the peak suppression occurred around at SOA = 85.5 msec, where the

thresholds were higher than for BR suppression but lower than CFS in Experiment 3. The FS-only and the FS+BR conditions did not show a clear difference when SOA was below +100 msec, but showed clear differences in tuning width when SOA was above +150 msec.

2.4.9. Experiment 5: How many flashes are necessary to achieve CFS?

In Experiment 4, we found that a single flash never elevated the thresholds as much as CFS does, regardless when the single flash is presented. Further, we found that when combined with binocular rivalry suppression, the suppression produced by a single flash, while unchanging in magnitude, was extended in time. To reconcile the results in Experiment 3 and 4, we hypothesize that a single flash *cooperates* with neighboring flashes to enhance inhibition, as might be happening in CFS. In Experiment 5, we tested this hypothesis by varying the number of flashes from one to five and, at the same time, varying the SOA between the onsets of the first flash and the probe systematically (Figure 2.6 a). Our goal was to determine how many flashes are necessary to achieve the threshold elevation of the full CFS condition. Are there critical flashes embedded among the continuous flashes? Although we know that many flashes should eventually amplify suppression cooperatively as in CFS, it is difficult to predict what exactly happens for intermediate number of flashes. In the visual masking literature, it is known that when a target is inhibited by a first masker, a second masker can inhibit the first masker to recover the visibility of the target (*disinhibition*) (Robinson, 1966; Dember and Purcell, 1967; Breitmeyer et al., 1981), and in other cases they can enhance inhibition cooperatively (Macknik and Livingstone, 1998; Macknik et al., 2000; Macknik and Martinez-Conde, 2004). Which effects are observed depends on the precise spatio-temporal arrangement of experiments. We know *a priori* what the answer to this question is because CFS is so much more powerful than a single flash; the flashes obviously have to interact cooperatively.

In Figure 2.6 b, the thresholds at the best SOA are shown as a function of the number of flashes, including the CFS condition in Experiment 3. Notably, the threshold for the five-flash suppression was comparable to the threshold for CFS. However, the thresholds in the conditions involving two to four flashes resulted in almost the same magnitude observed in the single-flash, far from the CFS condition.

In Figure 2.6 c, all the measured thresholds are shown as a function of SOA. When SOA was less than 200 msec, the thresholds for the five-flash suppression (filled squares) were higher than those for the four-flash suppression (empty triangles), which in turn were higher than those for the rest of conditions. When SOA was 200 msec, the thresholds for all conditions dropped below the level of the BR condition in Experiment 3. Thus, at least one flash must be presented no more than 200 msec prior to the onset of the probe for effective multiple-flash suppression to occur. The thresholds in the five-flash suppression condition was much higher from those in the four-flash suppression at the SOA of 0 msec, but not at the other SOAs. Such strong SOA dependency was not found in the conditions that involved one to four flashes.

Detailed examination of Figure 2.6 c reveals that the threshold was 44.4±8.8% for the five-flash suppression at SOA = 0 msec, very close to the threshold for the CFS condition (44.7%). The thresholds were considerably lowered by removing the first flash (four-flash, SOA = 100 msec) or the last flash (four-flash, SOA = 0 msec) to 19.2±4.1% or 17.4±1.5%, respectively. Adding one flash to the four-flash suppression at the 'wrong' time also did not help; the thresholds were 20.2±2.3% for five-flash at SOA = -100 msec and 30.7±6.5% for five-flash at SOA=100 msec. This kind of highly non-linear interaction can be observed in other data points as well. For example, though the

three-flash suppression starting at SOA = 200 msec did not elevate thresholds at all on their own (4.1±0.9%), combined with the two-flash suppression starting at SOA = 0 msec, whose threshold was 9.3±1.0%, they amounted to the strongest five-flash suppression.

The depth of suppression produced by CFS was reached when at least five flashes were successively presented. The effects of each of the five flashes were cooperative, but highly non-linear; it was difficult to rank order the importance of each flash because the suppression became weaker substantially by removing a single flash from five flashes or adding a single flash to four flashes at the wrong time. As for the intermediate number of flashes, they neither facilitated suppression nor disinhibited the suppressed target. It is possible that each additional flash acts as a masker that acts on the immediately following or the preceding flashes and, as a result, it cancels its suppressive effect (disinhibition (Robinson, 1966; Dember and Purcell, 1967; Breitmeyer et al., 1981)). Why, then, does the inhibition become suddenly strong when there are five flashes? One possibility is that disinhibition may be strong only for the flash at the onset or the offset of a set of multiple flashes, that is, the first or last flash. When there are five flashes, disinhibition from the first or last flash may not be able to disinhibit the third flash, and threshold elevation might have been observed (Macknik and Livingstone, 1998; Macknik et al., 2000; Macknik and Martinez-Conde, 2004).

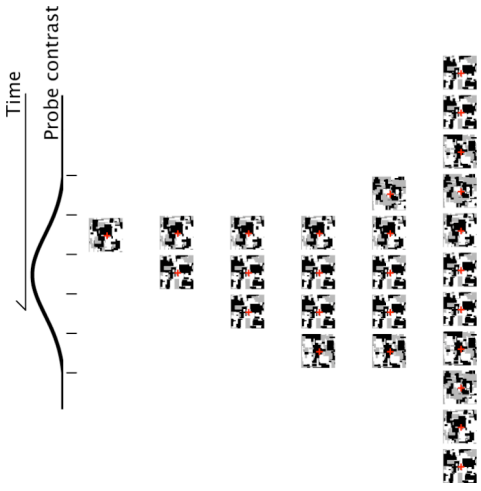
Time
Probe contrast

a)

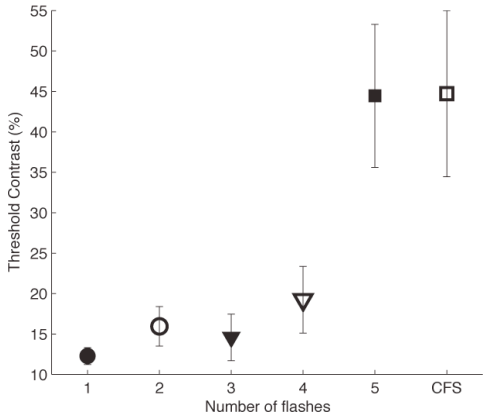
Threshold Contrast (%)
Number of flashes
1
2
3
4
5
CFS

b)

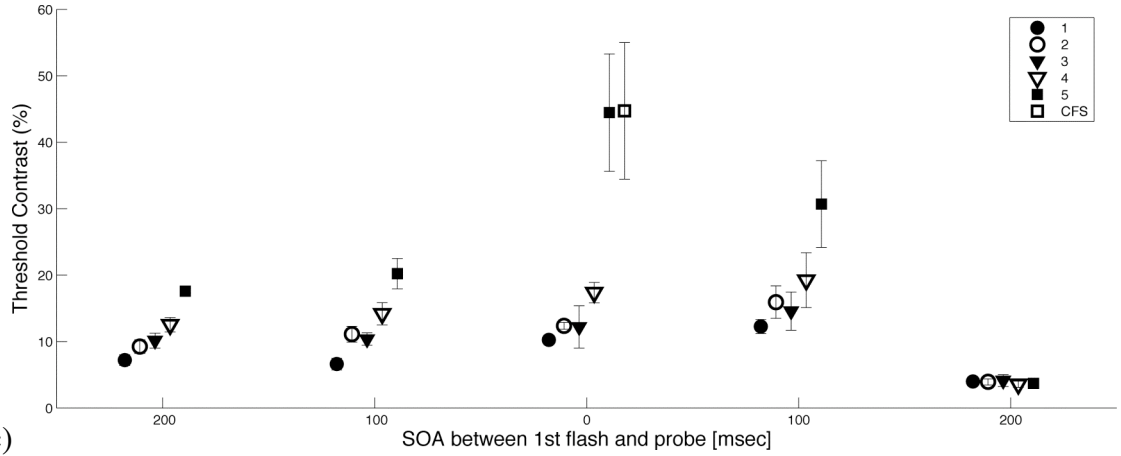
Threshold Contrast (%)
SOA between 1st flash and probe [msec]
1
2
3
4
5
CFS

c)

Figure 2.6 How many flashes are necessary to attain the depth of CFS?

a) Probe detection thresholds in the multiple-flash suppression. Each flash was followed by another flash with a 100-msec interval. The timing of the first flash was varied from -200 to +200 msec relative to the onset of the probe (in the schematic figure here, vertical positions of the Mondrians indicate the timing of flashes). The contrast of the 500 msec probe was modulated in a Gaussian manner, with its peak at 250 msec. **b**) Thresholds as a function of the number of flashes (error bars for s.e.m.). For comparison, the threshold for the CFS condition from Experiment 3 is re-plotted. For other conditions, the highest thresholds among the five SOAs were plotted (For 1-4 flashes, SOA was +100 msec, and for 5 flashes 0 msec. Shown schematically at the corresponding columns in **a**). Note the much higher thresholds in the five-flash suppression and CFS. A single flash suppression did not differ much from 2-, 3-, and 4-flash suppression. **c**) Thresholds as a function of all the tested SOAs. A sharp temporal tuning was observed only for the five-flash suppression. Symbols: filled circles, 1 flash; empty circles, 2 flashes; filled triangles, 3 flashes; empty triangles, 4 flashes; filled triangles, 5 flashes; an empty square, CFS.

2.5. CFS reduces the intensity of negative afterimages

2.5.1. Experiment 6: Pilot experiments on afterimage reduction

We next examined if CFS interferes with the formation of negative afterimages. These experiments were partly motivated by the observation that only a single subject in one out of 40 trials—where subjects did not see the suppressed figure for 3 minutes—reported a negative afterimage of the gray figure. We did not expect this, given that the image was present for three minutes on the retina.

In Experiment 6, we presented two isoluminant Gabor patches in one eye to the left and right of fixation for 5 sec (Figure 2.7, left). At the same time, suppressing CFS stimuli were continuously flashed only to one side of the other eye (Figure 2.7, center). CFS in one eye effectively renders the Gabor patch at the corresponding location in the other eye invisible (Figure 2.7, right).

Sixteen naïve subjects verbally described their percepts after 5-sec adaptation in two trials (Table 2.2). Subjects usually reported that the adaptor suppressed by CFS produced a weaker afterimage (87% in the 2-trial experiment and 83% in the 30-trial experiment). No subjects reported seeing an afterimage of the Mondrians. This consistency was notable, given the known variability in the strength of afterimages across trials, subjects, and hemifields (Loomis, 1972; Georgeson and Turner, 1985; Shimojo et al., 2001; Hofstoetter et al., 2004). Under the retinal origin hypothesis, input from the other eye should not influence afterimage formation. As adaptation at the retina is the

same for both visible and suppressed locations, the weakened afterimage must be due to interference from sites at or beyond binocular convergence.

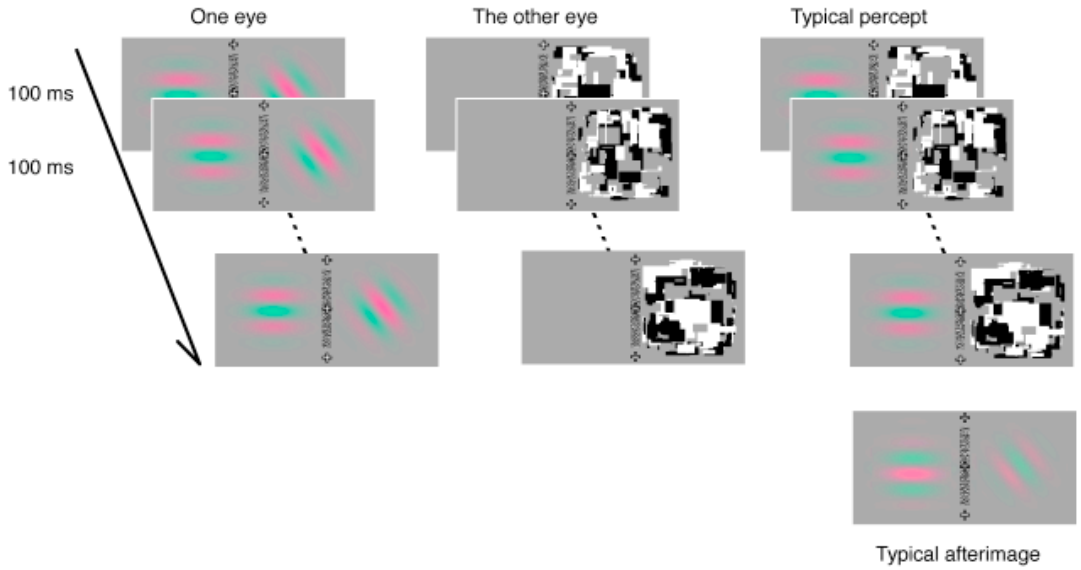

Figure 2.7 CFS suppresses a Gabor patch and reduces its afterimage

(Left) Two isoluminant Gabor patches were presented to the left and right of fixation in one eye for 5 sec (for Experiment 6 and 9) or 3 sec (for Experiment 8) during adaptation. The Gabor patches had 30% contrast and 0.6 cpd spatial frequency but with a spatial phase and orientation randomly drawn for each stimulus. **(Center)** Different Mondrian patterns that changed every 100 msec were projected into one half of the visual field (right side in this case) of the other eye. **(Right)** Typically, subjects saw a Gabor patch on one side (in this example, on the left visual field) and changing Mondrians on the other (the right visual field), failing to perceive the Gabor patch on the right. Subjects verbally described their percepts at the end of the adaptation period.

	16 naïve subjects		8 subjects	
	Total number		CFS trials	Catch trials
Visible only	7	22%	14%	95%
Visible > Suppressed	13	41%	65%	0%
Same	2	6%	16%	1%
Suppressed > Visible	1	3%	1%	0%
Suppressed only	0	0%	0%	0%
No afterimage	9	28%	5%	4%

Table 2.2 Detailed results for Experiment 6

Sixteen subjects described their percepts immediately after a tone sounded at the end of a 5 sec adaptation period. The Mondrians were placed to the left of fixation in one trial and to the right in the other. Subjects did not know there would be Gabor patches on both sides during the adaptation. Verbal reports of two trials by 16 naïve subjects (32 trials total) were classified into six categories: an afterimage was perceived only from the visible (or suppressed) adaptor, two afterimages were perceived but the afterimage corresponding to the visible (or suppressed) adaptor was stronger, and the two afterimages were of the same strength or were not perceived at all.

Eight of the subjects carried out another 20 trials with Gabor patches on both sides (CFS trials) and 10 trials with a Gabor patch only on the side that was not suppressed by the CFS (catch trials). Catch trials were included to ensure the subjects did not report imaginary afterimages or afterimages from Mondrian patterns. 10 catch trials

were randomly interleaved with 20 CFS trials. They described their percepts by choosing one of the six categories as explained above (six-alternative forced choice). The subjects were told there would be Gabor patches on both sides in some of these trials.

The numbers of trials and rounded percentages of the report (in the six-alternative forced choice) are shown. Excluding those trials where the afterimage was invisible on both sides, subjects usually reported that the plainly visible adaptor produced a stronger afterimage than the adaptor suppressed by CFS (87% in the 2-trial experiment and 83% in the 30-trial experiment).

2.5.2. Experiment 7: Ruling out the possibility of non-specific effects of the Mondrian flashes

Although no subjects reported seeing afterimages to the ever-changing Mondrians, such dynamic and luminance-equated patterns could have created afterimages (Virsu and Laurinen, 1977). Though their contrast may have been too low to perceive, they may nevertheless have interfered with the afterimage from the Gabor adaptor (Breese, 1899; Blake et al., 1971). In Experiment 7, we tested for this possibility. We compared the subjective ratings of the afterimage intensity from three intermittently presented adaptors (2-sec 'on' and 2-sec 'off' for 30 sec, Figure 2.8). One eye was stimulated by three separate Gabor patches while the other eye was stimulated at two of these three locations with Mondrians in such a way that this pattern synchronously coincided for 2 sec with one of the Gabor patches but was asynchronously delayed by 2 sec from the other Gabor pattern. A third location was never suppressed by Mondrians but received a Gabor patch intermittently for 2 sec and served as a control for the strength of the afterimage (pegged at a subjective rating of 10). Any putative afterimage of the Mondrian would interfere with the afterimage of the Gabor in both synchronous and asynchronous CFS locations. If CFS has to be presented simultaneously with the adaptors to weaken afterimage amplitude, the afterimages should be equally strong for the control and the asynchronous CFS locations and should be weaker for the synchronous CFS location.

Each of 6 naïve subjects performed 20 trials, rating the subjective intensity of the afterimages induced by the synchronously and asynchronously suppressed Gabor adaptors relative to the control, which received a rating of 10 (Figure 2.8, bottom). The

mean afterimage rating from the asynchronous CFS location was 11.1 ± 1.13 (standard error of the mean, or s.e.m.), which was not significantly different from 10 (two-tailed *t*-test, $P > 0.35$, *t*-score = 0.98, d.f. = 5). The mean rating from the synchronous CFS location was 5.85 ± 1.63, a reduction of 47% (one-tailed paired *t*-test on the rating between synchronous and asynchronous, $P < 0.02$). We conclude that the Mondrians themselves did not reduce the afterimage, and that coincidence of the adaptor with CFS was key to the observed reduction of the afterimage.

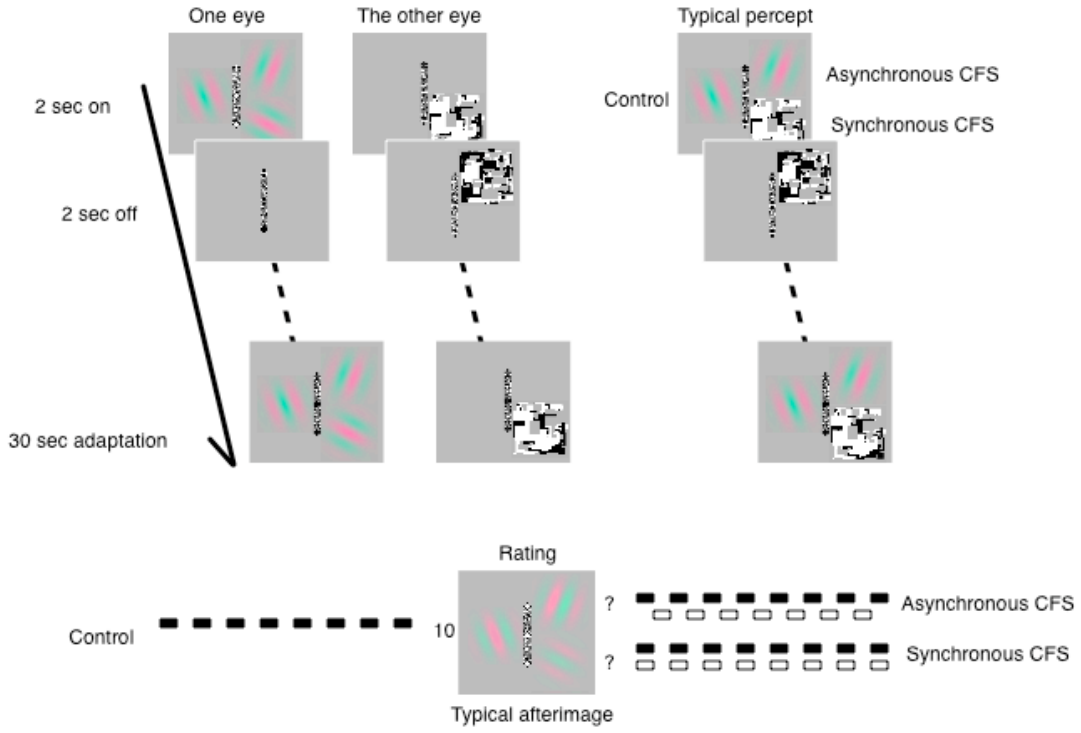

Figure 2.8 Mondrian flashes themselves do not reduce the afterimage of the Gabor

(**Left**) Three adaptors were presented to the left, top right, and bottom right of fixation during 2-sec 'on' periods and removed during 2-sec 'off' periods. The contrast of the adaptors was 50%. The position of the adaptors and Mondrians was balanced across top/bottom and left/right across 20 trials (5 trials for each configuration). (**Center**) Mondrian flashes were presented *synchronously* with the adaptors during 2-sec 'on' periods at the bottom right and *asynchronously* during 2-sec 'off' periods at the top right. (**Right**) During 2-sec 'on' periods, 6 naïve subjects perceived two adaptors and a stream of Mondrian patterns at the bottom right, while during 2-sec 'off' periods they saw only a

stream of asynchronous Mondrians at the top right. (**Bottom**) After 30 seconds of adaptation, subjects rated the intensity of the afterimage relative to that at the left visible location (control), which was pegged at ’10.’ Next to the expected afterimage, the time course of adaptors and Mondrians is shown for each location. Small black squares indicate the eight, 2-sec adaptor-on periods while empty squares denote the 2-sec CFS-on periods.

2.5.3. Experiment 8: Reliability of suppression and afterimage reduction

Why does CFS reduce the intensity of the afterimage, while previous studies showed that perceptual suppression did not influence the duration or the intensity of the afterimage using BR (Lack, 1978) or MIB (Hofstoetter et al., 2004)? One notable difference is that CFS suppresses adaptors in a more consistent and complete manner than either techniques; most subjects do not see the suppressed stimuli at all throughout the adaptation period.

In Experiment 8, we tested the extent to which complete invisibility is necessary for the observed reduction in afterimage strength to occur. We measured the reduction in the afterimage while manipulating the reliability of suppression of the Gabor patches by changing the stimulus properties of both adaptors and Mondrians. In preliminary experiments, we found that complete suppression occurred less frequently as the adaptors contained more and more high spatial frequencies. Furthermore, as the contrast of adaptors (respectively Mondrians) increased, the suppression became less (respectively more) reliable. We used 5 different spatial frequencies for the adaptors and 3 combinations of adaptor/Mondrian contrasts to manipulate the reliability of suppression. With the same setup as in Figure 2.7 (except for a 3-sec adaptation period), subjects indicated which side had the stronger afterimage and whether they saw the Gabor adaptor at the CFS location during adaptation (e.g., the right side in Figure 2.7). We also used BR to suppress the inducing image to compare with the efficiency of CFS methods.

The results (Figure 2.9) are quite unambiguous: the less the Gabor patch is visible during adaptation, the weaker the associated afterimage (for CFS, $r^2 = 0.82$, $p<1e\text{-}5$; for

BR, $r^2 = 0.76$, $p<1e\text{-}4$). Or, put differently, the more reliable the suppression (in a statistical sense), the weaker the afterimage. If suppression is sufficiently reliable, the intensity of the afterimage is reduced.

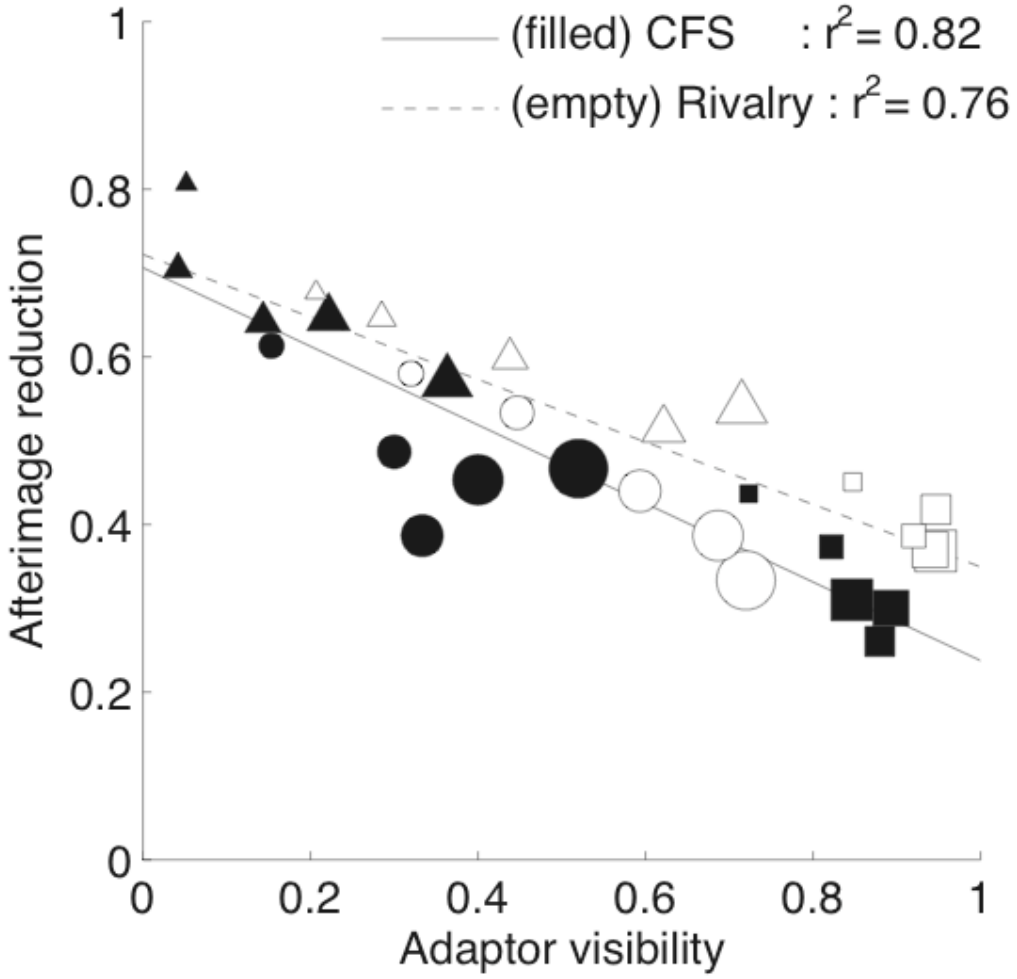

Figure 2.9 Afterimage reduction and the reliability of suppression

We used two different dichoptic suppression protocols: CFS (as in Figure 2.7, except that the adaptation period was shortened to 3 sec) and BR with moving stimuli. 5 subjects reported which of the afterimages were stronger and whether or not they saw the suppressed Gabor patch during adaptation. To modulate the reliability of suppression, three combinations of contrasts for adaptors and Mondrians were used: 30% and 100% (triangles), 100% and 100% (circles), or 100% and 5% (squares). The Gabor had one of

five different spatial frequencies. Subjects compared the intensity of the afterimage from two adaptors with the same contrast and spatial frequency. Increasing symbol size represents increasing spatial frequency of the adaptors: 0.60 through 0.84, 1.2, 1.7, to 2.5 cpd. In total, 30 different experimental conditions were evaluated. Filled symbols represent data obtained from CFS, and open symbols represent data from BR. Each data point represents the average across five subjects. The y-axis is the proportion of trials in which the afterimage from the suppressed adaptor was weaker than the afterimage from the plainly visible Gabor patch, representing the degree of afterimage reduction. The x-axis is the fraction of trials during which any part of the adaptor was visible, representing the reliability of complete suppression in a statistical sense. The data clearly show that for both CFS and BR, the less frequently the adapting stimulus is seen, the weaker its associated afterimage.

2.5.4. Experiment 9: Trial-by-trial visibility and the degree of afterimage reduction

CFS-induced suppression may reduce the afterimage either by eliminating the afterimage entirely on some fraction of trials or by lowering the afterimage intensity uniformly on all trials.

In Experiment 9, we tried to distinguish between these hypotheses. We repeated three of the conditions from Experiment 8 more extensively to estimate the matching contrast. In Figure 2.10 a, we show the results in the condition with the low-spatial frequency Gabor patch and high contrast Mondrian patterns (the small filled triangle in Figure 2.9). Here, the contrast of a test Gabor patch that matched the 60% contrast Gabor patch suppressed by CFS was 42.7% ± 7.1% (n = 5, t-score = 2.45, $P < 0.05$, one-tailed t-test against 60%), replicating our previous finding (in Sections 2.5.1 to 2.5.3). The average adaptor visibility duration was 0.51 ± 0.24 sec during the 5-sec adaptation period (Figure 2.10 b). Subjects did not report seeing any part of the Gabor in 61.8 ± 15.2% of trials (Figure 2.10 c). In the high-spatial frequency condition (the large filled triangle in Figure 2.9), the matching contrast was 62.8 ± 6.1%; that is, there was no reduction of afterimage intensity (n = 4, t-score = 0.46, $P > 0.6$). The mean adaptor visibility duration was 1.73 ± 0.38 sec, and complete suppression occurred in 18.6 ± 10.9% of trials.

To evaluate the effect of adaptor visibility, we sorted the 60 trials with the high-spatial frequency patches at each test contrast into ten bins according to the duration of the visibility of the adaptor. Figure 2.10 d shows each data point averaged across four subjects and psychometric curves fitted for each of ten bins. Matching contrast was independent of adaptor visibility duration (Figure 2.10 e, $r^2 = 0.18$, $P = 0.22$). Although it

did not reach significance, the slope is slightly negative (-2.9% s^{-1}), contrary to the prediction that only invisible trials contribute to afterimage reduction. We obtained similar results using low-spatial frequency Gabor patches and low-contrast Mondrians (Figure 2.11 a and b; n = 5; the matching contrast was 61.2 ± 1.8%, the mean adaptor visible duration was 2.03 ± 0.38 sec and complete suppression occurred in 4.12 ± 10.9% of trials). Again, we did not find any correlation between matching contrast and the adaptor visibility duration (r^2 = 0.23, P = 0.16; slope of the regression line was, again, slightly negative: -3.2% s^{-1}). Dividing trials into two or three bins did not change the results. Although there was only a little variation in adaptor visibility, we repeated the same analysis on the data set in the condition with low-spatial frequency Gabor patches and low-contrast Mondrians (Figure 2.11 c). There was no correlation between matching contrast and the adaptor visibility duration (P = 0.52) and the slope of the regression line was, again, slightly negative: -3.7% s^{-1}.

For a given stimulus setting, trial-by-trial variability in the visibility of the adaptor did not change the intensity of the afterimage. This is consistent with previous studies of afterimages (Lack, 1978; Hofstoetter et al., 2004) but was contrary to the suppression of high-level aftereffects (Wiesenfelder and Blake, 1990; Moradi et al., 2005). Since the statistical reliability of suppression is correlated with the reduction of the afterimage, the visibility of the adaptor seems only indirectly related to the percepts of the associated afterimage. Thus, the answer to the question that triggered this experiment, “does CFS-induced suppression eliminate the afterimage entirely on some trials or does it lower the afterimage intensity uniformly on all trials,” is the latter.

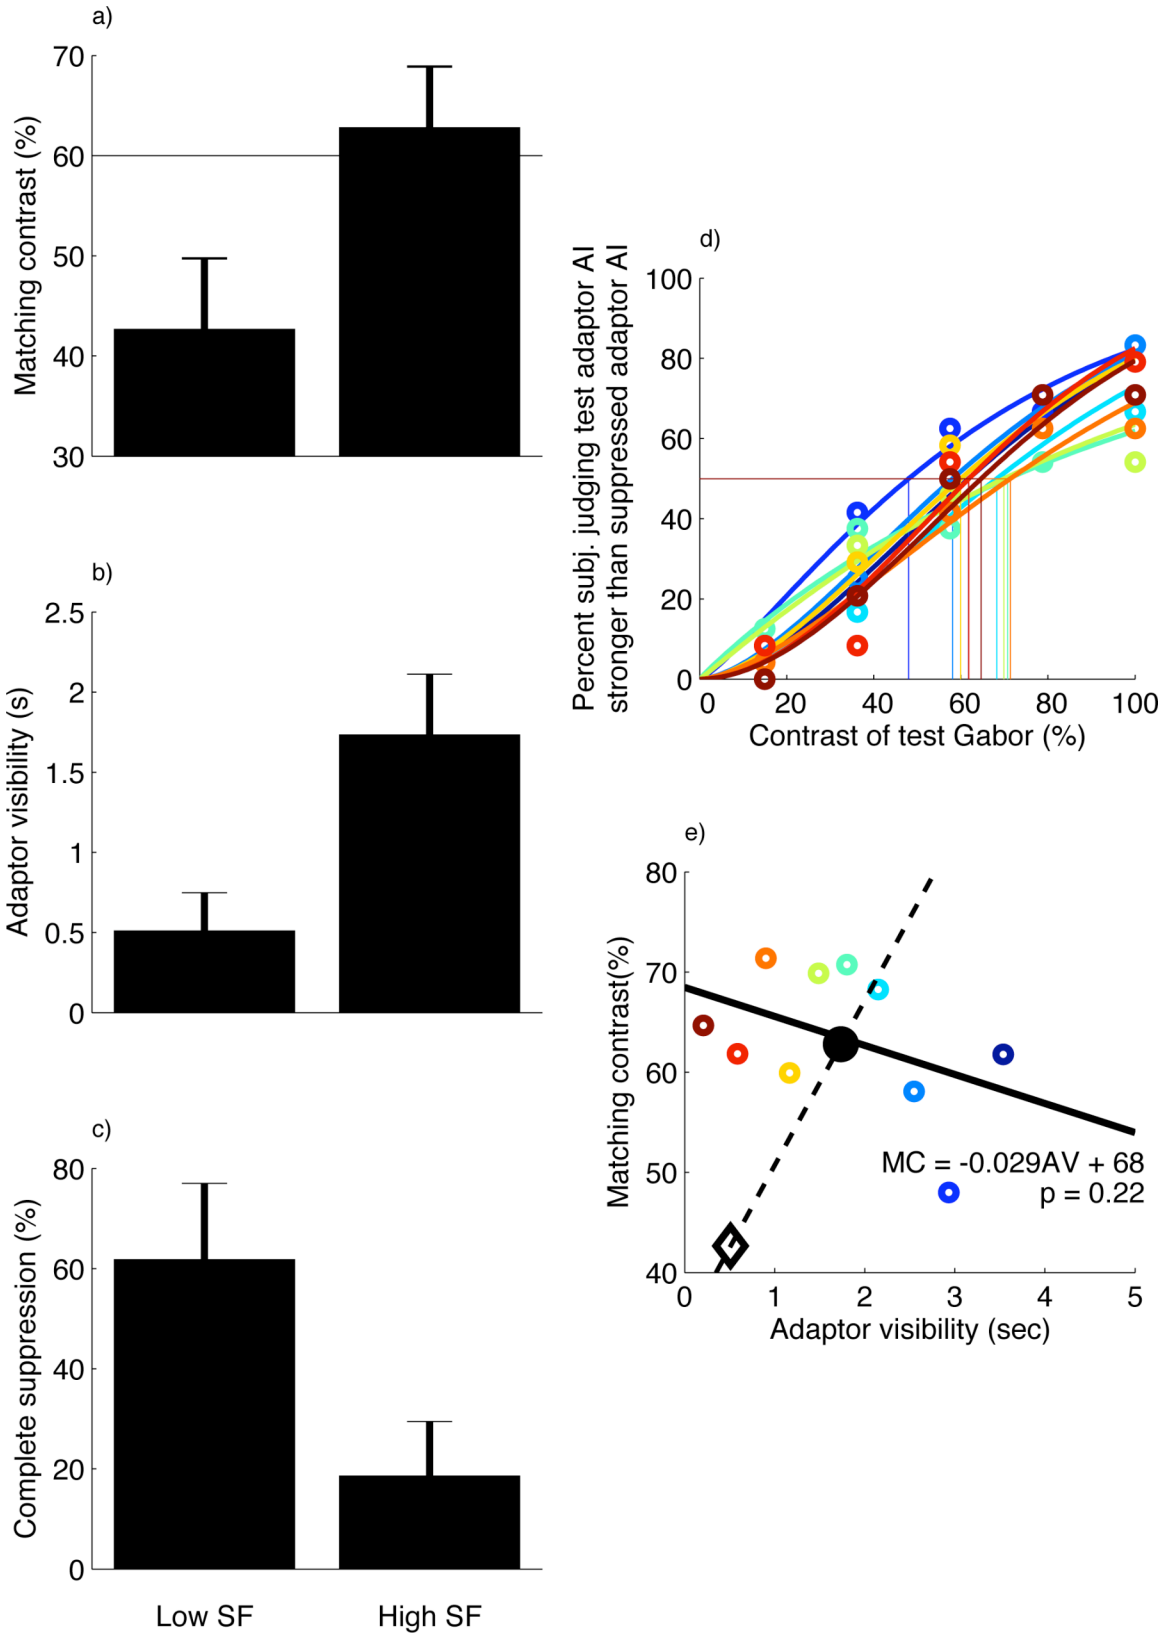
a)
Matching contrast (%)
70
60
50
40
30
b)
Adaptor visibility (s)
2.5
2
1.5
1
0.5
0
c)
Complete suppression (%)
80
60
40
20
0
Low SF
High SF
d)
Percent subj. judging test adaptor AI stronger than suppressed adaptor AI
100
80
60
40
20
0
0
20
40
60
80
100
Contrast of test Gabor (%)
e)
Matching contrast(%)
80
70
60
50
40
0
1
2
3
4
5
Adaptor visibility (sec)
MC = -0.029AV + 68
p = 0.22

Figure 2.10 Visibility and afterimage reduction

We repeated three conditions of Experiment 8. 4 subjects reported the visibility of the suppressed Gabor by holding a key during the 5-sec adaptation period and then compared the afterimage intensity. (**a**-**c**) We used low- (0.6 cpd, left) and high-spatial frequency (2.0 cpd, right) Gabor patches. (**a**) Matching contrast (estimated by the method of constant stimuli), showing a significant reduction of matching contrast only in the low-spatial frequency condition (error bars represent s.e.m.). (**b**) Mean duration that the adaptor was visible. (**c**) Proportion of trials where subjects did not see the adaptor at all. (**d**) We divided 60 trials at each test contrast for each subject into ten bins based on the adaptor visibility duration in the high-spatial frequency condition (a-c right). Six trials from each of four subjects were pooled to fit a Weibull function, which was used to estimate the matching contrast for which the afterimage (AI) induced by the test adaptor was stronger than the afterimage induced by a suppressed Gabor adaptor in 50% of trials (vertical lines in d). (**e**) The duration for which the adaptor was visible was weakly but negatively correlated (one-tailed t-test; $P = 0.22$) with the matching contrast (thick solid line). Different colors in d and e represent different durations of mean adaptor visibility. The large filled black circle and the empty square in e represent the average matching contrast in a and adaptor visibility in b for the high- and the low-spatial frequency condition, respectively. If the adaptor visibility would correlate with the intensity of afterimages in a trial-to-trial manner, there should be a positive correlation (thin broken line).

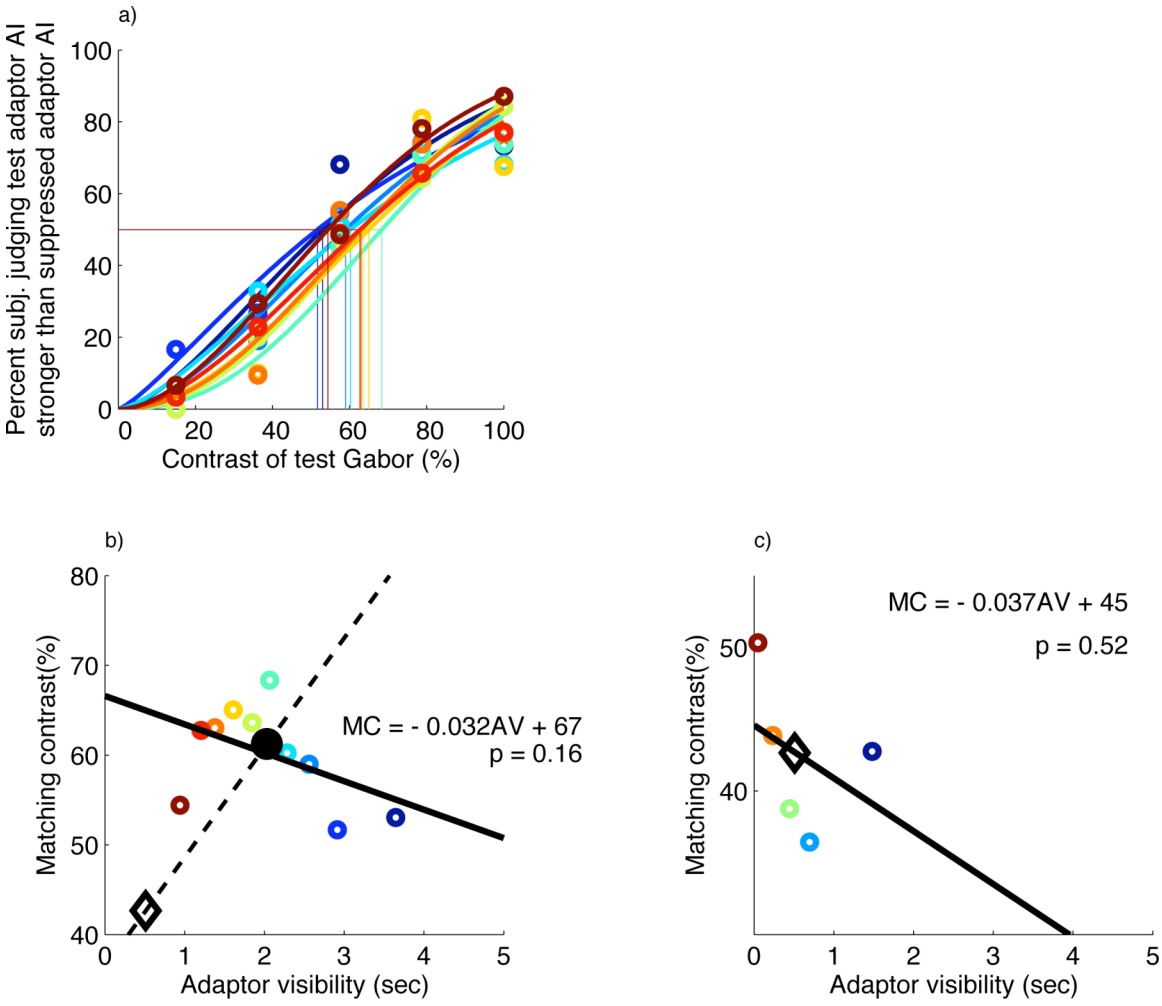

Figure 2.11 Visibility and afterimage reduction (control)

(**a-b**) Data from 5 subjects in the condition with a Gabor patch of low spatial frequency (0.6 cpd) and low-contrast (2-4%) Mondrian patterns. The formats are the same as those in Figure 2.10 d and e. We divided 60-70 trials at each test contrast for each subject into ten bins based on the adaptor visibility. Six to seven trials from each of five subjects were pooled to fit a Weibull function, which was used to estimate the matching contrast. If the adaptor visibility would correlate with the intensity of afterimages in a trial-to-trial

manner, there should be a positive correlation (thin broken line). (**c**) Data from 5 subjects in the baseline condition with a Gabor patch of low spatial frequency (0.6 cpd) and high-contrast (100%) Mondrian patterns, analyzed as in Figure 2.10 d and e. We divided 20-30 trials at each test contrast for each subject into five bins based on the adaptor visibility. Four to six trials from each of five subjects were pooled to fit a Weibull function, which was used to estimate the matching contrast. Because the suppression was very strong, there was not much variability in the adaptor visibility. Although there is no significant correlation ($P = 0.52$), again, there is no trend of positive correlation, which would be expected if the adaptor visibility would correlate with the intensity of afterimages in a trial-to-trial manner.

2.5.5. Experiment 10: Withdrawing attention enhances the intensity of afterimages

The effects of attention and awareness are similar in most experimental conditions, and they are usually difficult to disentangle. Although some recent research claims a dissociation of attention and awareness (McCormick, 1997; Kentridge et al., 1999a, b; Lambert et al., 1999; Naccache et al., 2002; Kanai et al., 2004; Kentridge et al., 2004; Rajimehr, 2004; Koivisto et al., 2005b; Koivisto et al., 2005a; Melcher et al., 2005; Montaser-Kouhsari and Rajimehr, 2005), none of them show the *opposite* effects of attention and awareness. We wanted to replicate the original finding that withdrawing attention enhances afterimages (Lou, 2001; Suzuki and Grabowecky, 2003) in our setting.

In Experiment 10, we compared the apparent contrast of the afterimage when the adaptor was ignored because the spatial attention was withdrawn to a task-relevant rapid digit stream presented at the fixation (Figure 2.12 a), and when the adaptor could be attended because they did not perform the task. The central task required subjects to count the appearance of digit '4's among a stream of 9 other distractor digits. The difficulty of the task was kept around 70% using a 2-down-1-up staircase procedure to control the duration of each digit presentation. The monitoring of digit stream engaged spatial attention for a sustained duration, sufficient to modulate the intensity of afterimages (Suzuki and Grabowecky, 2003). A task on rapid visual serial presentation is shown to withdraw focal attention to the same degree as other attention demanding tasks (Braun, 1998).

The adaptor was either suppressed by CFS or plainly visible without Mondrians (visible vs. invisible x attended vs. unattended). These four experimental conditions was

blocked and the order was randomized across subjects. In each trial, a Gabor patch appeared only on one side for 5 sec, while the digit stream was presented at the fixation (the same visual inputs for attended vs. unattended condition, Figure 2.12 a). At the end of the adaptation, a Gabor patch with the same appearance as the negative afterimage (*i.e.*, phase-reversed and same orientation) was presented on the other side. Subjects compared the apparent contrast of the afterimage with that of the physical Gabor just after the end of adaptation period, cued by a tone signal. Matching contrast was estimated with a 1-up-1-down staircase method.

Making adaptor *invisible* via CFS *reduced* the intensity of afterimage from 17.5% to 11.9% (across two attentional conditions), whereas *withdrawing attention* from the adaptor by a central task *increased* the intensity of afterimages from 13.2% to 16.1% (across two visibility conditions, Figure 2.12 b). Using ANOVA, we confirmed significant main effects of CFS (F = 18.30 P = 0.00058) and attention (F = 4.86, P = 0.0425). The visibility effect was significant separately for each attentional condition, confirming our results in Experiments 6-9; a *post-hoc* paired one-tailed t-test showed a significant effect of visibility when subjects performed the attention-demanding task (t-score = 3.67, P = 0.011) and when they passively viewed the display without any task (t-score = 3.55, P = 0.012). Interestingly, the effect of visibility did not interact with the effect of attention (F = 0.15, P = 0.702). In other words, the magnitude of the attentional effect was similar when the target was visible (increase from 15.8 to 19.2%) and when the adaptor was invisible (increase from 10.7 to 13.1%). A *post-hoc* paired one-tailed t-test confirmed the attentional effect separately when the adaptor was visible (t-score = 2.07, P = 0.054), as expected from the previous study of (Suzuki and Grabowecky,

2003). When the adaptor was invisible, we did not have a prior hypothesis, and a *post-hoc* two-tailed test for attention did not reach significance (*t*-score = 2.24, *P* = 0.089)[5]. However, we note the magnitude of the attentional effect was similar in both visibility conditions and the direction of attentional effect was consistent with the previous study of (Suzuki and Grabowecky, 2003).

We showed performing the central task and thereby 'drawing attention' away from the Gabor increased the intensity of the afterimage. The lack of interaction indicates that the attentional effect was independent of awareness of the adaptor. We conclude that the lack of attention during CFS-suppressed adaptation does not explain our main findings. Indeed, this experiment supports the view that selective visual attention and awareness can be dissociated (Lamme, 2003; Koch, 2004).

[5] Rather small size of the attentional effect seems comparable with the study of Suzuki and Grabowecsky, who hired > 20 naïve subjects to confirm the significance of the effect. The effect might be enhanced if the subjects actively attend to the adaptor by performing a task on the adaptor, rather than passively viewing the display without doing any task. Here, we did not opt for such an experimental design because we were concerned about the role of fixational eye movements. It is very likely that attending to a target outside the fixation would induce (micro-) saccades towards the adaptor, which would reduce the intensity of afterimage in a trivial manner.

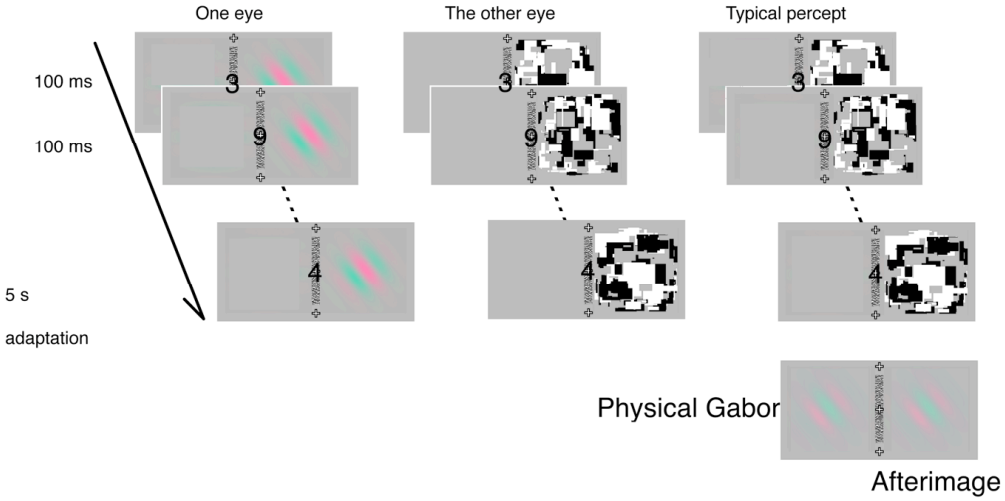

a)

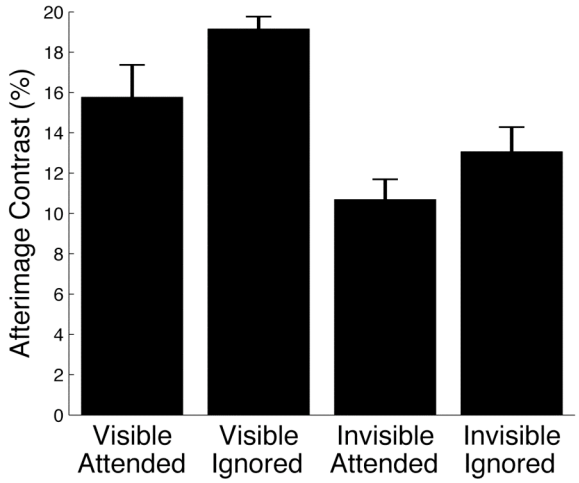

b)

Figure 2.12 Apparent contrast of the afterimage in the presence and the near-absence of focal attention

(**a**) In each trial, a Gabor patch of 60% contrast with random orientation and a spatial phase was presented on one side of the fixation to one eye. Visibility of the Gabor patch

was manipulated by a 10Hz stream of Mondrian patterns on the same side of the other eye. The focal attention was manipulated by having subjects ignored or engaged in the central digit stream, appearing in a fixation area of 0.5 x 0.5° (the size of the digit in the figure is exaggerated). After 5-sec adaptation, subjects compared the apparent contrast of the afterimage and the physical Gabor, and in addition, reported the occurrence of the digit '4's in an "adaptor-ignored" condition. (**b**) Matching contrasts of the physical Gabor are shown for four experimental conditions. Error bars are s.e.m.

2.5.6. Do we need to give up localizing visual aftereffects?

Recent psychophysical studies, including our own (Tsuchiya and Koch, 2005), show that negative afterimages originate not only from retinal adaptation, as traditionally thought, but that post-retinal components also significantly contribute to the formation of negative afterimages (Shimojo et al., 2001; Gilroy and Blake, 2005). Recent electrophysiological studies show that neuronal network in mammalian retina possesses substantial adaptability to the contrast variation of the input images (Smirnakis et al., 1997), or even to highly complex spatiotemporal patterns, such as orientation and motion (Hosoya et al., 2005).

The findings discussed so-far question two traditional, well-established assumptions about localizing neural functions using visual aftereffects. First, when an aftereffect does not transfer across eyes, as is the case for afterimages, it is assumed that monocular channels relatively early in the visual processing are solely responsible for mediating the aftereffect (Blake et al., 1981). Thus, no inter-ocular transfer of afterimages has been taken as an evidence for the retinal origin of afterimages (but see Discussion 2.6.6). Second, when an aftereffect requires selectivity to complex visual features, such as orientation, direction of motion, or facial identity, it is assumed that the aftereffect is mediated by neurons with complex receptive properties, generally found in later, *i.e*. cortical, stages of visual processing. For example, circular receptive fields in the retina and LGN imply that these structures cannot be the neural basis of the orientation-selective adaptation (Blakemore and Campbell, 1969) (but see (Hosoya et al., 2005)). The recent findings discussed above (Smirnakis et al., 1997; Shimojo et al.,

2001; Gilroy and Blake, 2005; Hosoya et al., 2005; Tsuchiya and Koch, 2005) call into question these assumptions for localizing the neuronal sites for a given aftereffect within the visual system.

Do we need to abandon hope to track down the neuronal correlates of visual adaptation to one or more highly localized neuronal populations? We believe not. In Experiment 11, we provide a new technique to dissociate the microgenesis of retinal and post-retinal adaptation, using continuous flash suppression (CFS). While we induced afterimages presented by an adaptor to one eye for a constant duration, we manipulated the duration of CFS presented to the other eye. CFS should prevent post-retinal components from adapting, but should not affect retinal adaptation.

2.5.7. Experiment 11: Dissociating the microgenesis of retinal and post-retinal adaptation

We showed in Experiment 9 that variability in trial-to-trial visibility did not correlate with the intensity of afterimages. In that experiment, we needed to make CFS less reliable to obtain sufficient variability in visibility. Here, using an optimal CFS, we manipulated the duration of suppression by physically starting or removing CFS during the course of adaptation. When reliable CFS coincides with the adaptor (Experiments 8-10), the intensity of afterimages was reduced. By changing the duration and timing of overlap between CFS and the adaptor, we can characterize the temporal characteristic (microgenesis) of the post-retinal components of afterimages. If the adaptor is initially visible (or suppressed) and subsequently becomes suppressed (or visible) until the end of adaptation, would the afterimage be stronger than the afterimage induced by the completely suppressed adaptor? If so, what is the relationship between the duration of suppression and the reduction of afterimage intensity?

In Experiment 11, we measured the intensity of afterimages with a contrast matching technique using a staircase method (Kelly and Martinez-Uriegas, 1993) (see Methods). We held the duration of retinal adaptation constant for 5 sec, while manipulating the duration of CFS, which should affect adaptation of post-retinal, but not retinal, components. For example, when the matching adaptor was fully visible throughout (Figure 2.13 a), post-retinal components would adapt maximally and afterimages would be most intense. When the adaptor was suppressed all the time during adaptation (Figure 2.13 b), post-retinal components would adapt minimally and

afterimages would be least intense. Critically, when the adaptor was suppressed (via CFS) only during the last fraction of the 5-sec adaptation period (Figure 2.13 c), post-retinal components would adapt until the start of CFS but would then stop adapting and decay. Likewise, to study the buildup of adaptation, we suppressed the adaptor only after a variable period of visibility. In this case, post-retinal components would not adapt until the release from CFS, and from then on would start adapting until the end of the 5-sec adaptation period. However, retinal exposure of the adaptor would be exactly the same in all conditions, as the CFS stimulus was only presented in the other eye. Thus, we can study the microgenesis of post-retinal components of afterimages, that is, how fast they buildup or decay.

In order to minimize covert shifts of attention and eye movements induced by the transient apparent motion of the Mondrian patterns from an upper to a lower quadrant, subjects performed an attentionally demanding task at the fixation. The same digit-counting task was used in Experiment 10 and it did not interfere the formation of afterimage (rather, it enhanced the intensity of afterimage, if any).

Surprisingly, the estimated intensity of afterimages was reduced below 60% of the strength of the normal afterimage, even when the adaptor was suppressed for as short as the last 0.5 sec, that is, only five Mondrian flashes were presented during the last 500 msec of adaptation[6] (Figure 2.13 d). In other words, there was little or no dependency of how long the adaptor was visible initially, prior to the last period of suppression. This was confirmed by one-way ANOVA, showing a significant difference of the intensity

[6] We did not use suppression durations shorter than 0.5 sec because we thought that too few flashes presented at the end were likely to create afterimages of their own and they might make the interpretation of results more difficult.

($F(5) = 12.8$, $P < 1e\text{-}6$). A *post-hoc* Tukey's least significant difference test confirmed a significant difference between the no-suppression (decay time = 0sec) and the other conditions ($P < 0.05$), but not for any other comparisons, including a comparison between complete suppression (decay time = 5 sec, Figure 2.13 b) and partial suppression at 0.5 sec ($P > 0.05$). Similarly, when we made the adaptor became visible only for the last 0.5 sec, the estimated intensity quickly rose above 75% (Figure 2.13 e).

Our results cannot be explained by differential pupil diameters across conditions, because transient CFS flashes were always present in the display for 5 sec (Figure 2.13 a-c). We excluded the possibility of the effects of task set and/or attention by having subjects engaged in an attentionally demanding task at fixation throughout adaptation and by interleaving all conditions randomly within each block to avoid any shift of criterion in intensity judgment (see Methods).

For more quantitative analyses, we fitted these data with exponential functions of the form, $AI = a * \exp(-T/\tau) + b$, where AI is the estimated intensity of afterimages, a is a scaling constant (%), b is the asymptotic afterimage intensity (%), T is the duration of decay or buildup (sec), and τ is the time constant for 1/e decay or buildup (sec). The best-fit curves are shown in Figure 2.13 d and e. For decay, the best fit resulted in a = 37.4% (with the 95% confidence interval of [21.0, 55.83]), b = 57.0 [49.4, 64.4] %, and τ = 0.17 [-0.32, 0.65] sec, with $R^2 = 0.95$. For buildup, the best fit was obtained with a = -26.1 [-32.6, -19.6] %, b = 85.7 [81.2, 90.3] %, and τ = 0.51 [0.19, 0.84] sec, with R^2 = 0.97.

These data suggest that post-retinal components of afterimages rapidly decay and buildup, possibly reflecting the rapid adaptation and recovery of polarity-selective

neurons (See Discussion 2.6.5). Alternatively, the visibility of afterimages may be rapidly modulated by cortical mechanisms that have access to the information about the states of perceptual suppression just before the expression of afterimages. When the adaptor is strongly suppressed at the end, the putative mechanisms may turn off afterimages. Such modulation is implicated in afterimages created at the periphery, which can be turned off when subjects move their eyes so that these locations fall outside of the normal visual field (Hayhoe and Williams, 1984). Another possible explanation might be dichoptic forward masking onto the initial part of afterimages, which resulted in the reduction of afterimage. However, this is unlikely for several reasons. First, similar suppression was also obtained by reliable binocular rivalry suppression (Experiment 8), where no transients for dichoptic forward masking were present. Second, there was no suppression when unreliable CFS was used (Experiment 8 and 9), where transients for dichoptic forward masking were present.

We note the suppression obtained by five flashes in Experiment 5 was comparable with the suppression in CFS. This is consistent with both the rapid microgenesis and modulatory cortical mechanisms.

The rapid time constants allow us infer the neuronal mechanisms of afterimages. Even when we estimate the time constants conservatively at the longer end (0.65 and 0.84 sec for decay and buildup, respectively), these are still much more rapid than previous estimates of microgenesis of afterimage induced by fully visible stimuli, which typically shows exponential decay/buildup on the order of 4-8 sec (Kelly and Martinez-Uriegas, 1993; Hofstoetter et al., 2004). Our results imply the genuine cortical components of afterimages adapt and recover much faster than the retina.

Other adaptation phenomena, such as contrast adaptation (Blakemore and Campbell, 1969) and flicker adaptation (Schieting and Spillmann, 1987), are also characterized by slow time constant. Recent studies (Smirnakis et al., 1997; Hosoya et al., 2005) indicate that these types of adaptation may also be found within the retinal circuits. It is plausible that the slow time course in contrast adaptation may as well mainly reflect the retinal adaptation processes.

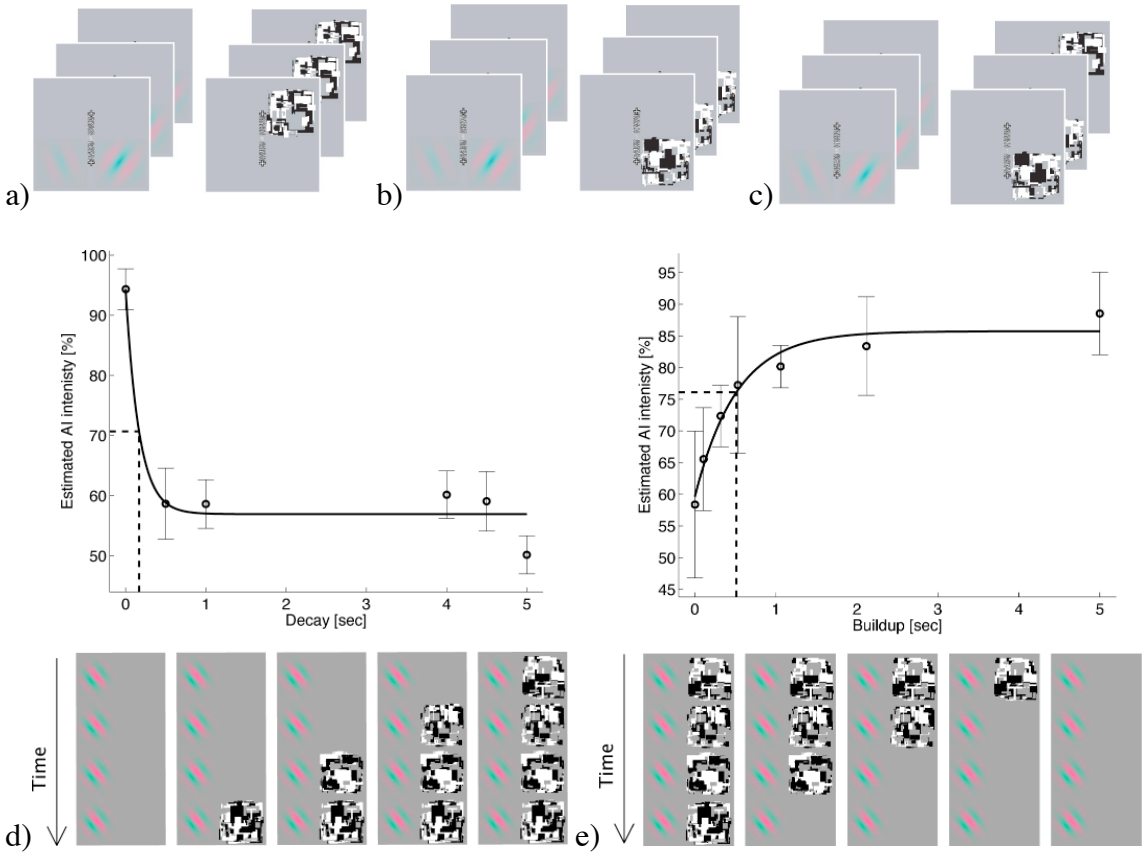

Figure 2.13 Microgenesis of post-retinal components of afterimage

(**a-c**) The constant adaptor of a 30% contrast (bottom left of the left column) and the matching adaptor (bottom right) were presented to the non-dominant eye. A 10Hz stream of Mondrians was presented in one of the four quadrants of the dominant eye. Stimuli were presented for 5 sec. **(a)** Mondrian patterns were presented at the top right corner to equate the display luminance and transients. In this condition, post-retinal components of afterimages were not suppressed at all, corresponding to decay time = 0 sec in d and buildup time = 5sec in **e**. (**b**) Mondrians coincided with the location of the matching adaptor for 5 sec, corresponding to decay time = 5 sec in d and buildup time = 0 sec in e.

(**c**) To characterize the decay and buildup time constant of post-retinal components of afterimage, we started CFS only after some time (here, from the 2nd frame). (**d)** Decay and (**e**) buildup of post-retinal components of afterimage. X-axis is the duration of decay or buildup and y-axis is the estimated intensity of afterimage, i.e., normalized matching contrast (Error bars are s.e.m.). The curves were obtained by the best exponential fit. The broken horizontal and vertical lines indicate the time constant for 1/*e* decay and buildup and are obtained using a best exponential fit.

2.6. Discussion

2.6.1. Prolonged phenomenal invisibility in continuous flash suppression (CFS)

We identified dichoptic visual stimuli that, for at least ten times longer than existing techniques, reliably suppress from conscious perception salient figures presented to one eye (Experiment 1). With CFS as a tool, vivid images can be rendered invisible for long periods with excellent control of timing. This suppression of a continuously presented stimulus at the fovea dissociates physical stimuli from their associated subjective percepts. CFS does not require pre-adaptation, a key aspect of flash suppression (Wolfe, 1984; Wilke et al., 2003; Pasley et al., 2004; Moradi et al., 2005) to achieve reliable disappearance. This property makes CFS attractive for studies that require complete unawareness.

CFS extends the total duration for which Mondrians are perceived by prolonging their period of dominance without shortening their period of suppression (Experiment 2). In binocular rivalry, strong stimuli have the effect of shortening the period of dominance of the other, weaker, stimuli (that is, effectively shortening the period during which they are suppressed), with little effect on their period of dominance (Levelt, 1965; Fox and Rasche, 1969; Blake, 1977; Blake and Logothetis, 2002) (but see (Mueller and Blake, 1989; Bossink et al., 1993; Sobel and Blake, 2002)). Thus, CFS is not simply a stronger version of binocular rivalry. This observation is compatible with the hypothesis that CFS involves a repetitive flash suppression component, in addition to binocular rivalry. Indeed, we present a simple model (Sections 2.4.3 to 2.4.5) that combines aspects of flash

suppression and binocular rivalry and describes the measured periods of CFS dominance and suppression in a quantitative manner.

2.6.2. CFS suppresses the target strongly but avoids adaptation to prolong dominance

Using a probe detection task, we found that the detection threshold was 20-times higher in CFS than under monocular-viewing conditions (Experiment 3), but only three-times higher in binocular rivalry. The prolonged dominance in CFS (Experiment 1) may be due to the stronger suppression. This conclusion seems in conflict with the evidence that 'strong' stimuli in binocular rivalry do not prolong its dominance (Levelt, 1965; Fox and Rasche, 1969; Blake, 1977; Blake and Logothetis, 2002) as we argued in (Sections 2.4.3 to 2.4.5).

The apparent conflict is resolved when we consider the effects of adaptation for perceptual switches in bistable percepts. We argue that CFS effectively minimize neuronal adaptation, which is one possible cause of terminating dominance in binocular rivalry (See a review by (Leopold and Logothetis, 1999) for other possible causes of perceptual switches). In this framework, two competing neuronal populations reciprocally inhibit each other via feedback, and the winning population dominates the percept (Matsuoka, 1984; Lehky, 1988; Mueller, 1990; Wilson, 2003). When the dominating population fatigues, adapts or recalibrates, its inhibition becomes weaker and a perceptual switch occurs. In this model, any methodology that alleviates adaptation of neuronal population for the dominant percept would prolong the dominance. Indeed, two techniques to prolong dominance in rivalry has been recently reported; inserting blanks

between bistable percepts (*i.e.*, freezing (Leopold et al., 2002)) and changing retinal location during rivalry (Blake et al., 2003) both prolong the dominance duration. In CFS, randomly generated Mondrians were replaced every ~100 msec. This short presentation of each pattern minimize adaptation of neurons that respond to patterns as a whole or to local, high contrast edges, because each Mondrian pattern is unique and its edges are located at random positions. However, in conventional 'strong' rivalry stimuli, especially high contrast ones, most strong rivalry stimuli are stationary or periodic if they are dynamic, thus, they strongly stimulate the *same* population of neurons during perceptual dominance. Compared to weak stimuli, neuronal response to strong stimuli would be initially more vigorous, but reduced faster by adaptation, fatigue, and/or recalibration, leading to the similar duration of perceptual dominance (Matsuoka, 1984; Lehky, 1988; Mueller, 1990) (During perceptual suppression, these models posit that neuronal responses to the preferred stimuli recover from adaptation. The speed of recovery depends on the strength of stimuli, explaining the major effects of stimulus strength on the duration of perceptual suppression (See also Section 2.4.3. to 2.4.5)).

2.6.3. Reducing the intensity of afterimage via dichoptic suppression

We applied CFS to ascertain the extent to which dichoptic inhibition interferes with the formation of negative afterimages. Though it is widely believed that afterimages originate among retinal neurons (Craik, 1940; Alpern and Barr, 1962; Brindley, 1962; Loomis, 1972; Sakitt, 1976; Virsu and Laurinen, 1977; Lack, 1978; Loomis, 1978; Wilson, 1997), and their formation therefore immune from influences from the other eye, some experiments imply that post-retinal processing can modulate (Anstis et al., 1978; Hayhoe

and Williams, 1984; Lou, 2001; Suzuki and Grabowecky, 2003) or possibly even create (Weiskrantz, 1950; Shimojo et al., 2001) negative afterimages. Our results are consistent with these latter studies. Notably, they imply that such post-retinal components may be necessary for the formation of afterimages. Dichoptic inhibition that underlies the afterimage reduction has been found as early as the LGN (Sengpiel et al., 1995) in anesthetized cats. In contrast, dichoptic inhibitory effects in a masking paradigm were not found in the LGN of awake monkeys (Macknik and Martinez-Conde, 2004). For binocular rivalry, electrophysiological evidence for the suppression in LGN is controversial conflicting between the results from anesthetized cats (Varela and Singer, 1987) and awake monkeys (Lehky and Maunsell, 1996). Recent human fMRI studies provide evidence for the hemodynamic modulation of the LGN during binocular rivalry (Haynes et al., 2005; Wunderlich et al., 2005).

The possibility that latent afterimages from the Mondrians interfered with the afterimage of the Gabor pattern was ruled out by Experiment 7. Asynchronously presented Mondrian patterns did not reduce the intensity of the afterimage, whereas synchronously presented Mondrians reduced it by about 50%. This suggests that one peculiarity of CFS, continuously present transient signals, is not sufficient for the reduction of the afterimage. Rather, the adaptor has to be suppressed strongly by stimuli presented simultaneously to the other eye.

2.6.4. Reliability of suppression and trial-to-trial visibility in a two-stage hierarchical model

We found that the degree of afterimage reduction correlated with how reliably adaptors are suppressed in a statistical sense. This relationship exists for both binocular rivalry and CFS (Experiment 8), implying that the inconsistency between our results and previous studies (Lack, 1978) arises partly from the reliable suppression induced by CFS. Although cortical neurons may adapt under partial suppression, they seem to adapt less under reliable suppression. In an analogy with lesion studies, partial suppression techniques can be compared with 'unilateral lesions' and CFS with 'bilateral lesions'; often, bilateral, but not unilateral, lesions result in behavioral deficits. Partial suppression by binocular rivalry and MIB may leave sufficient residual activity to produce full-blown adaptation and an afterimage.

Further, we confirmed (Lack, 1978; Hofstoetter et al., 2004) that afterimage intensity was not influenced by trial-by-trial variability in adaptor visibility, indicating that post-retinal components for the afterimage are not directly related to the neuronal correlates of awareness.

Although reliability of suppression correlated with the reduction of afterimage intensity, trial-to-trial visibility did not. Because these two measures were obtained by independent methods[7], it is not so surprising to find that the degrees of correlation were different. Still, it is important to point out that completely invisible CFS trials did lead to reduced afterimage intensities, while equally invisible trials that are a subset of unreliable

[7] Indeed, difference in reliability of suppression comes from the difference in the experimental protocols including suppression technique (e.g., CFS vs. BR vs. MIB), contrast and spatial frequency of stimuli.

suppression did not reduce afterimage intensity (Figure 2.10). This is reminiscent of the finding on the depth of suppression in Experiment 3. Perceptual suppression induced by reliable CFS and binocular rivalry were phenomenally indistinguishable; however, the depth of suppression were very different.

We can explain the results of Experiments 8 and 9 with the help of a two-stage hierarchical model, where reliability of suppression and trial-to-trial visibility reflect outputs from the first and the second stage, respectively. In the first stage, inputs from two eyes converge and the strength of each input pattern is compared. In the second stage, based on the analysis at the first stage, only one pattern survives because of a winner-take-all computation. Because of internal noise, adaptation and other factors (Discussion 2.6.2), the winner is not necessarily the stronger one and who the winner is may fluctuate unpredictably on a moment-by-moment basis. Trial-to-trial visibility corresponds to the output from this second stage. Reliability of suppression reflects the output from the first stage because it corresponds to the probability that one pattern wins over many trials. This two-stage model is consistent with electrophysiological studies of binocular rivalry, where the proportion of neurons whose firing correlates with monkey's percept is low in the early visual cortex (20% in V1 and V2) and higher as the processing goes deeper (40% in MT, MST and V4, and 90% in IT and STS (Logothetis, 1998). V1 and V2 would correspond to the first stage, and IT and STS the second stage. It is also consistent with a computational model based on the two layers of spiking neurons, which reconciles contradicting ideas between eye-specific vs. pattern-specific rivalry hypotheses (Wilson, 2003).

According to the two-stage model, post-retinal components of afterimages are mainly formed in the neural structures corresponding to the first stage (Experiment 8), but not the second stage (Experiment 9). Difference in reliability of suppression at the first stage resolves the apparent contradiction regarding the effects of perceptual suppression on the formation of afterimages; the intensity of afterimages is reduced by strong reliable suppression (Gilroy and Blake, 2005; Tsuchiya and Koch, 2005), but not by unreliable suppression (Lack, 1978; Hofstoetter et al., 2004) at the first stage. Trial-to-trial visibility, the output from the second stage, does not influence the intensity of afterimages (Lack, 1978; Hofstoetter et al., 2004; Tsuchiya and Koch, 2005).

There are other types of aftereffects that appear to be formed in the first stage. Such candidate aftereffects include aftereffects of (orientation-specific) contrast, translational motion and tilt. Although original studies show that visibility does not influence the magnitudes of these aftereffects (Blake and Fox, 1974; Lehmkuhle and Fox, 1975; Wade and Wenderoth, 1978; O'Shea and Crassini, 1981; Moradi et al., 2005), recent studies show that strong suppression or weak adaptors reduce these aftereffects (Lehky and Blake, 1991; Kanai et al., 2004; Sobel et al., 2004). It remains to be tested if the magnitudes of these aftereffects do not correlate with trial-to-trial visibility, as we here demonstrate for afterimages ((Moradi et al., 2005) reports no correlation between the orientation-specific contrast aftereffects and trial-to-trial visibility).

Some aftereffects appear to be formed in the second stage. The magnitude of the spiral motion aftereffects (Wiesenfelder and Blake, 1990) and the face-identity specific aftereffects (Moradi et al., 2005) correlates with trial-to-trial visibility. Both aftereffects were reduced even with binocular rivalry.

2.6.5. The opposing effects of awareness and attention in the formation of afterimages

In Experiment 10, we showed that our results in Experiments 6-9 could not be explained by a lack of attention to the adaptor owing to complete suppression. Attending to adaptors during adaptation weakens the afterimage (Lou, 2001; Suzuki and Grabowecky, 2003) as we replicated in Experiment 10. If one assumes that lack of attention to an object is equivalent to not being aware of it (O'Regan and Noe, 2001), one would expect both would cause the same effects on the afterimage. However, lack of attention enhances afterimages, while reliable perceptual suppression reduces afterimages. These results support the view that attention and awareness involve different mechanisms (Lamme, 2003; Koch, 2004).

Why are the effects of attention and awareness on the formation of afterimages opposing? In terms of neuronal activity, attending to and becoming aware of a stimulus is usually recorded as an increase of firing (for a review of the neurophysiology of attention see (Desimone and Duncan, 1995; Maunsell, 1995) and for awareness see (Logothetis, 1998)). There are some exceptions, though, for example, some neurons in MT, MST, and V4 decrease their firing rate when monkeys becomes aware of the preferred stimuli of the recorded neuron (Logothetis and Schall, 1989; Leopold and Logothetis, 1996). FMRI studies in humans likewise report that BOLD signals increase when subjects attend to (Watanabe et al., 1998a; Gandhi et al., 1999; O'Craven et al., 1999; Saenz et al., 2002) or become aware of the stimuli during binocular rivalry (Lumer et al., 1998; Tong et al., 1998; Polonsky et al., 2000).

It is known from psychophysics that the aftereffects of tilt, motion and face are all enhanced by attending to adaptors during adaptation (Chaudhuri, 1990; Spivey and Spirn, 2000; Kanai et al., 2004; Moradi et al., 2005) and reduced when adaptors are made invisible by inter-ocular suppression, such as binocular rivalry or continuous flash suppression (Kanai et al., 2004; Sobel et al., 2004; Moradi et al., 2005). Therefore, we cannot explain why awareness and attention cause opposing effects on the strength of afterimages in terms of any common mechanism for aftereffects in general[8]. We have to consider something specific to afterimages.

Perceived afterimages are polarity-reversed version of the adapted pattern, hence, they are called 'negative' afterimages (in our experiments, red-green polarity was reversed). Negative afterimages are formed due to adaptation of polarity-sensitive neurons, which are abundant in the retina, LGN and among simple cells in V1. Once induced, the intensity of afterimages also depends on the state of other processes, namely, polarity-independent (or contrast sensitive) cells, such as complex cells in V1. The involvement of such polarity-independent process is shown by experiments showing that adaptation to counter-phase or drifting grating, which do not produce interfering afterimages on their own, reduces the intensity and duration of the subsequently induced afterimages (Leguire and Blake, 1982; Burbeck and Kelly, 1984; Georgeson and Turner, 1985). (Suzuki and Grabowecky, 2003) suggested that *attention enhances* adaptation of

[8] For example, one might think of 'rebound' of attention as a possible explanation; attending to the adaptor during the *formation* of afterimages might repel attention from the adaptor location during *perception* of afterimages, much like inhibition-of-return of attention. If this were true, attention would be directed to the previously ignored adaptor, therefore, the ignored adaptor would produce more intense afterimages than the attended one. This logic is not satisfactory because reduction of aftereffects should be also observed for aftereffects of tilt, motion and facial identity.

polarity-independent processes (e.g., contrast adaptation) more than polarity-sensitive processes (e.g., formation of afterimages). This prediction receives some support by experiments showing attending to the adaptor during adaptation enhances contrast adaptation (Festman and Ahissar, 2004), but its effect may be rather small (Moradi et al., 2005). To be conclusive, the magnitude of contrast aftereffects and its attentional modulation should be compared directly with the magnitude of attentional reduction in afterimage intensity in a within-subject design[9].

2.6.6. Reconciling with no inter-ocular transfer of afterimages

A reduction of 50% in the strength of the afterimage when the inducing image is present on the retina but not seen by the observer seems to be at odds with the fact that afterimages do not transfer across eyes; when the subject closes the adapted eye, no afterimage is seen. It is known that the binocular components of afterimages have access to the direction of gaze (Hayhoe and Williams, 1984). Likewise, it may be possible that these mechanisms have access to the overall brightness from the eye and, if it is closed, may reduce or even eliminate afterimages, resulting in no transfer of afterimages. The involvement of a cortical suppression mechanism is supported by patients with cortical lesions who report abnormally long afterimages that transfer across eyes (Chan et al.,

[9] According to this framework, reliable suppression prevents adaptation of the polarity-sensitive processes more so than adaptation of the polarity-independent processes (Experiment 8) and trial-to-trial visibility either does not influence adaptation of the polarity-selective process or influences adaptation of both processes but each one cancels the effects of each other. This framework is consistent with our two-stage model (Discussion 2.6.4) if we assume the proportion of polarity-sensitive neurons is higher for the first stage than the second stage.

2001; Weiskrantz, 2002). This may reflect the disruption of cortical mechanisms for afterimage reduction. If normal observers open both eyes during the test period, the inter-ocular transfer effect, albeit weak, can be measured psychophysically in detection or discrimination procedures (Schiller and Dolan, 1994).

2.6.7. Rapid microgenesis of post-retinal components of afterimages: compensatory rapid mechanisms for slow retinal adaptation?

Are there any functional reasons why post-retinal components of afterimages have faster microgenesis? A more fundamental question would be, is there any functional role at all for negative afterimages?

Traditionally, vision scientists considered 'retinal' negative afterimages as a by-product or an epiphenomenon of luminance adaptation (Virsu, 1978). By adapting to the ambient light level, the visual system becomes more sensitive around the mean luminance level and optimizes information transfer. Luminance adaptation is highly functional as the ambient light level varies daily over more than nine orders of magnitude, while optic nerves can encode information with neuronal spikes with less than two orders of magnitude of bandwidth (Barlow, 1981) However, because this adaptive mechanism is mainly based on local mechanisms (i.e., photoreceptors and retinal neurons with small receptive fields), it can be mal-adaptive when it manifests itself as negative afterimages, which lower sensitivity (Leguire and Blake, 1982; Burbeck and Kelly, 1984; Georgeson and Turner, 1985).

If local negative afterimages are non-adaptive, it makes sense if the cortical mechanisms, which have access to much richer information, compensate for the

maladaptive afterimages. The post-retinal processing of afterimages we characterized in this chapter are largely consistent with this idea. As we have seen, post-retinal components of afterimages can be modulated depending on the efferent signals for gaze direction (Hayhoe and Williams, 1984) and eye closure (See Discussion 2.6.6), attention (Experiment 10) and the interference from the other eye (Experiments 6-9, 11). None of the information above is available at the level of retinal processing. Thus, the efferent signals and attention may be utilized to alleviate the mal-adaptive aspects of afterimages. These mechanisms may operate rapidly as revealed by dichoptic suppression (Experiment 11), which may explain why we are largely oblivious to afterimages induced in everyday life.

2.7. Methods

Subjects were recruited from the Caltech campus, except for Experiment 3-5, which were performed at Vanderbilt University. Subjects gave consent before participating in the experiments. They had normal or corrected eyesight and normal stereo vision. Subjects observed the display through a set of mirrors. The distance between the eyes and the display was 92 cm. To stabilize fixation, a head and chin rest was employed. We used Matlab 6.5 under Windows 98 (Experiment 1 and 2) and Matlab 5.2.1 under Mac OS and the Psychophysics Toolbox (Brainard, 1997). The Mondrian images consisted of randomly generated squares of random colors (Experiment 1 and 2) or white, black, and gray squares superimposed onto each other. Twenty to forty distinct Mondrians were generated prior to each session.

2.7.1. Experiment 1: Prolonged invisibility by CFS

Seventeen naïve subjects participated. They were instructed to hit a space bar when any part of a gray image became visible and to describe it verbally to the experimenter. The time to key press was taken as the duration of initial suppression. One of four types of gray images was used in each trial: A 45° left-tilted Gabor patch of 30% contrast, a 45° right-tilted Gabor patch of 60% contrast, an angry face, and a blurred angry face. Spatial frequency and standard deviation of the Gabor patches were 0.5 cpd and 1°, respectively. Each type of image appeared once in a block of 4 trials. In total, 4 blocks of 16 trials were run. The images were presented at the fovea and extended 6 x 6°.

2.7.2. Experiment 2: Optimal flash interval for CFS

Four naïve subjects participated. During a one-minute observation period, a gray Gabor patch was presented to one eye while color Mondrian patterns at flash intervals ranging from 10 to 1280 msec or a stationary Mondrian (binocular rivalry) were presented to the other eye in the same set-up as in Figure 2.1. Subjects pressed and held one of three keys to indicate their current percept: Mondrian only, Gabor only, and a mixture of the two. The flash interval was randomized within sessions. The measurement was repeated four times. Subjects took at least a one-minute break between trials. In the analysis, 'Gabor only' and 'mix percept' were treated the same (Gabor visible) and were contrasted against 'Mondrian only' percepts (Gabor invisible). The mean dominance period was calculated by excluding periods that were terminated by the end of the one-minute observation interval.

2.7.3. Experiment 3: Depth of suppression in CFS

The author of this thesis and three experienced observers in the Blake lab at Vanderbilt University participated in Experiment 3-5. The primary objectives of Experiment 3-5 were to compare the depth of suppression in CFS with the conventional binocular rivalry. When binocular rivalry is induced using small targets, only one pattern, but not the other, is almost always exclusively visible at a given moment (for exceptions to this rule, see (Wolfe, 1983; Liu et al., 1992)). As the rivalry targets become larger, the global percept becomes patchy or piecemeal, while retaining local exclusive visibility (Blake et al.,

1992). Accordingly, in Experiment 3-5, we used smaller suppressed targets and Mondrians (1.6 x 1.6°) to facilitate the comparison between binocular rivalry and CFS. Forty different black/gray/white Mondrian patterns were created before each block using the same algorithm as in the other experiments, but they were shrunk to fit to the square (Figure 2.4 a). As a result, the spatial frequency of the Mondrians was higher than the other experiments. A red fixation cross and black/white stripes surrounding targets helped binocular fusion. The monitor refresh rate was 100Hz. The monitor was carefully calibrated using a color-bit stealing technique that provided an effective resolution of 10-bits after linearization (Tyler, 1997).

Subjects initiated a trial by pressing a button. In the CFS condition, subjects saw a 10Hz stream of Mondrian patterns presented to one eye and did not see a horizontal sinusoidal grating of 10 % contrast and 1.8 cpd presented to the other eye. After making sure that no part of the grating was visible, they pressed a button, which triggered a contrast increment pulse (The first frame for the contrast increment was synchronized with the next closest Mondrian flash). Either upper or lower half of the grating increased its contrast. The phase of the grating was 0° at the horizontal center to avoid an artifactual edge. The contrast was smoothly increased in a Gaussian manner (standard deviation of 100 msec, Figure 2.4 a) for an extent of 500 msec to avoid abrupt onset/offset signals. At the end of probe presentation, a brief beep sound alerted subjects to press an up or a down arrow key to report which field contained the probe (2AFC). CFS and the baseline grating were kept on the display until subjects responded. No feedback was given to subjects after each response. After two consecutive successful responses contrast increment was reduced by 30% of the current increment, and after

each error it was set by 30% higher (a two-down-one-up staircase converging to the level for 71% correct performance (Levitt, 1971)). After four reversals, the contrast adjustment was reduced to 15%. One block terminated after 12 reversals and a threshold was calculated by taking the geometric mean of the contrast increment values of last 10 reversals.

In the binocular rivalry condition, a single stationary Mondrian pattern was randomly chosen from 40 patterns for each trial. After initiating a trial, subjects waited until the stationary Mondrian achieved complete dominance. Upon complete dominance, they pressed a key, immediately followed by a contrast increment. In the monocular-viewing condition, we presented a blank field to the suppressing eye. In this case, a contrast increment was induced 500 msec after the key press, until then subjects saw the 10% contrast grating. The other aspects of these conditions were the same as the CFS condition.

During one block of staircase, the same eye was used for probe detection and the suppression protocol was held constant. The author of this thesis used one eye as the probe eye throughout and the other subjects used both eyes in a balanced manner. Measurement was repeated twice to six times for each condition.

2.7.4. Experiment 4: Depth of single flash suppression

Three experienced subjects in Experiment 3, including the author of this thesis, participated in the experiment. The equipments, the stimuli (Mondrians and probes), and the basic procedure to estimate thresholds (2AFC and staircase) were identical to Experiment 3.

In the single flash suppression (FS-only) condition, subjects initially viewed a blank field on the suppressing eye and the baseline grating on the probe eye. The blank field was swapped into a Mondrian pattern at time T_M, which remained on until a response. A contrast increment probe was superimposed to the grating from time T_C and extended for 500 msec. Across blocks, the stimulus onset asynchrony (SOA) between the onsets of a probe and a Mondrian flash was manipulated (SOA = $T_C - T_M$). When SOA < 0 (a Mondrian flash occurred first, followed by a probe), the actual sequence of a trial was, 1) subjects pressed a button, which triggered a baseline grating of 10% contrast for 500 msec, 2) the blank field was swapped into the Mondrian pattern at T_M, then 3) the probe was presented. When SOA ≥ 0, the order of 2) and 3) were swapped. When the single flash was applied after binocular rivalry completed (the FS+BR condition), subjects initiated a trial by a key press and saw a stationary Mondrian on the suppressing eye and the baseline grating on the probe eye. When the stationary Mondrian completely dominated over the baseline grating, they pressed a button again, upon which the Mondrian swapped into another Mondrian pattern, then the probe was presented (when SOA < 0). When SOA ≥ 0, the order was reversed. For each trial, Mondrian patterns were chosen randomly from 40 patterns that were created before each block.

2.7.5. Experiment 5: How many flashes are necessary to achieve CFS?

The four subjects in Experiment 4 participated in this experiment. The equipments, stimuli, and basic procedure were the same as for the single flash suppression only condition in Experiment 4. The number of flashes was varied from one to five and, at the

same time, the SOA between the onsets of the probe and the first Mondrian flash from -200 to 200 msec was varied across staircases.

Subjects initiated a trial by a button press, which triggered 500 msec presentation of 10% contrast baseline grating to the probe eye and blank field to the suppressing eye. With an appropriate SOA (see Method for Experiment 4), the probe and the first of multiple Mondrian patterns were presented. Until subjects made a response, the last Mondrian pattern and the baseline grating remained on the display to avoid any transient offset signal. Except for the last Mondrian pattern, each Mondrian pattern was replaced with another one in an interval of 100 msec.

2.7.6. Experiment 6: Pilot experiments on afterimage reduction

Sixteen naïve subjects participated. Isoluminant Gabor patches (spatial frequency, 0.6 cpd; standard deviation, 0.83°) were used as afterimage inducers. The isoluminant green level was calibrated to equate with pink (CIE [x, y]=[0.389,0.205], luminance 18.7 cd/mm^2) using a full-field flicker-minimizing technique for each subject. The average green level was [x,y]=[0.201,0.278]. The contrast of the isoluminant Gabor patch was defined as the contrast modulation of the red or the green intensity, $Contrast = \frac{\max(red, green) - \min(red, green)}{\max(red, green) + \min(red, green)}$. Peaks of the red intensity coincided with troughs of the green to keep the luminance level roughly at 20 cd/mm^2. The luminance level for black and white was 0.028 and 67.6 cd/mm^2. Three crosses on a rectangle with random texture (0.48° x 4.8°) served to stabilize binocular fusion. Subjects fixated the middle cross. Each of two Gabor patches of 30% contrast were presented within an

imaginary 4.8° x 4.8° square, with the center of the square 2.6° away from the fixation (Figure 2.7, left). The phase and orientation of the adaptor were randomized for each trial. The black and white Mondrians flashed at 10 Hz were presented in the corresponding imaginary square (Figure 2.7, center). After a 5 sec adaptation period, the adaptors and CFS stimuli were replaced with a uniform gray background to induce negative afterimages. We asked subjects to describe their percepts immediately after a tone sounded at the end of the adaptation period. In the two-trial experiment, the CFS stimuli were placed to the left of fixation in one trial and to the right in the other. The verbal reports from 16 naïve observers were categorized into six classes. Eight subjects were further tested in a 30-trial experiment in a six-alternative forced-choice manner.

2.7.7. Experiment 7: Ruling out the possibility of non-specific effects of the Mondrian flashes

Six naïve subjects performed 20 trials. Three 50% contrast Gabor adaptors were presented spaced apart (Figure 2.8). One visible control adaptor was placed to the left (or right) of fixation (3.6° square, with its center location 2° from fixation; two other adaptors were placed at the top or bottom right (or left) from the fixation (3.6° square, 1.8° above or below fixation). Three adaptors appeared simultaneously for 2 sec and turned off for 2 sec, repeating over 7 cycles and ending with a 2-sec on period. After 30 sec of adaptation, subjects rated the intensity of the afterimage on a linear scale, relative to the control, which was pegged at ’10.’ If no afterimage was visible, the rating was ‘0.’ A rating of ‘5’ (or ‘20’) was given when the intensity was half (or twice) as strong as the afterimage from the control adaptor.

2.7.8. Experiment 8: Reliability of suppression and afterimage reduction

Four naïve subjects and the author of this thesis participated. Subjects compared the intensity of afterimages produced by a pair of adaptors with the same contrast and spatial frequency. Different spatial frequencies (0.60, 0.84, 1.2, 1.7, and 2.5 cpd) and two types of dichoptic suppression (CFS and BR) were randomly interleaved within a block, ruling out possible accumulative effects. The contrast of adaptors and Mondrians was held constant in one block of 100 trials. Each subject completed at least 30 trials for each combination of the suppression protocol, the spatial frequency of adaptors, and contrasts of adaptors and Mondrians (30% vs. 100%, 100% vs. 100%, 100% vs. 5%). We created motion binocular rivalry (BR) stimuli by sliding Mondrian patterns horizontally. The right and left half of the moving texture converged at the midline of the pattern (see the QuickTime movie at http://www.klab.caltech.edu/~naotsu/CFSdemo.html). The speed of horizontal motion was 0.71 deg/sec. The texture of the motion BR stimulus was randomly created before each trial.

2.7.9. Experiment 9: Trial-by-trial visibility and the degree of afterimage reduction

Four experienced (but naïve to the hypothesis of this experiment) subjects and the author of this thesis participated. Subjects pressed a key to indicate if the suppressed Gabor became visible during 5-sec adaptation and then reported on which side the afterimage was stronger. Only CFS was used as a suppressing protocol. The spatial frequency of Gabor was either 0.6 or 2.0 cpd, and the contrast of Mondrian was either 100% or 2-4%.

In the low spatial frequency and high Mondrian contrast condition, the test contrast was adjusted between either 10% to 50% or 15% to 100% in five linear steps, depending on the matching contrast for each subject. In one block of 50 trials (10 trials at each of 5 contrast levels), test contrasts were randomized. Each subject completed at least two blocks. In other conditions, the test contrast was always adjusted 15% to 100% in five linear steps, and 6 blocks were conducted.

To estimate the matching contrast for the test adaptor, we fitted a Weibull function to the data. For the correlation analysis, 60 trials at each test contrast were sorted according to the adaptor visibility duration; the 6 trials with the shortest visibility duration were categorized in the first bin, another 6 trials with the next shortest visibility duration were categorized in the second bin, and so on. For each bin, responses from 4 subjects were pooled, thus each point in Figure 2.10 represents the average of 24 trials.

2.7.10. Experiment 10: Withdrawing attention enhances the intensity of afterimages

Four subjects who were experienced as subjects but naïve to the hypothesis of this experiment and the author of this thesis participated. The central task was a continuous digit-counting task at fixation (Suzuki and Grabowecky, 2003), i.e., counting the occurrence of '4' embedded in a stream of numbers between 0 to 9. '4' appeared 0 to 9 times. The digit was presented in a fixation are of 0.5 x 0.5°. This task was shown to withdraw spatial attention from the adaptor and delay the onset of the afterimage (Suzuki and Grabowecky, 2003). The task difficulty was maintained by a 'three-down-one-up' staircase procedure; after each correct answer the presentation time of each digit was decreased by 10 msec, while after each mistake it was increased by 30 msec. This

procedure converges to a 79.4% correct performance level (Levitt, 1971). Subjects performed the task at approximately 150 msec for each digit presentation.

The apparent contrast of the afterimage was measured by controlling the contrast of a 'physical' isoluminant Gabor patch with a one-up-one-down staircase procedure. The physical Gabor had the same orientation but the opposite phase of the adaptor, that is, the same orientation and phase as the negative afterimage. The orientation and phase were randomized for each trial. The spatial frequency was set to 0.6 cpd and the adaptor contrast to 60%. The contrast of the physical Gabor decayed exponentially with a time constant of 5.0 sec (Kelly and Martinez-Uriegas, 1993). As the appearance of the physical Gabor was comparable to that of the negative afterimage, all subjects felt comfortable with comparing the apparent contrast of the two. We emphasized the importance of comparing the contrasts just after the end of the adaptor presentation (in the beginning of the physical Gabor presentation), signaled by a tone. Step size of the staircase for the physical Gabor was set approximately to 1/10 of the threshold, estimated during practice blocks.

Each block contained four independent staircases interleaved randomly. Two staircases controlled the contrast of the physical Gabor that appeared on the left side of fixation and the other two controlled on the right side. Two staircases started from high contrast (~60%) and the other two started from 0% contrast. Each staircase was terminated after 6 reversals, that is, one block terminated after 24 reversals. The arithmetic mean of the last 5 reversal contrasts from 4 staircases was taken as the apparent contrast of the afterimage.

2.7.11. Experiment 11: Dissociating the microgenesis of retinal and post-retinal adaptation

Six naïve subjects and the author of this thesis participated. All subjects participated in the experiment measuring the decay time, and three naïve subjects and the author of this thesis measured the buildup time.

In each trial, we presented a 30% contrast adaptor and another matching adaptor whose contrast was adjusted by the staircase method. These two isoluminant Gabors were presented at the bottom left and right quadrants for 5 sec. The 30% contrast adaptor was never suppressed and the matching adaptor was suppressed for a variable duration. A 10Hz stream of Mondrian patterns was presented for 5 sec, but it could change its location among four quadrants (Figure 2.13 a-c). By presenting Mondrians for the entire 5 sec in all conditions, we equated the average luminance level and image contrast to equate the pupil diameter. After the trial subjects made two responses. First, they reported how many times the digit '4' was embedded in the digit stream presented at the fixation (see Experiment 10). This task helped fixation and forced them to ignore the adaptors and the Mondrians. Second, they reported which side contained the stronger afterimages. The contrast of the matching adaptor was raised or lowered by 10% according to the response.

When measuring the decay time, 12 independent staircases were run simultaneously in one session: two Mondrian locations (bottom left or right) x six decay durations (0, 0.5, 1, 4, 4.5, or 5 sec). By intermixing these conditions we avoided shifts of criterion and attention. Each session terminated when all staircases reversed twice. Eight to twelve reversal contrasts were averaged to estimate the matching contrast.

When measuring the buildup time, we varied buildup durations from 0, 0.1, 0.3, 0.5, 1, 2, to 5 sec (simultaneous 14 staircases), and twelve reversals were averaged to estimate the matching contrast. The intensity of afterimages was estimated by the following formula,

afterimage intensity = 30% / measured matching contrast (%).

2.8. Appendix

2.8.1. Appendix 1: Prolonged first dominance period in CFS

In addition to our arguments in Sections 2.4.2 to 2.4.5, there is additional evidence that suggests that CFS involves a repetitive FS component that maintains the dominant percept, besides contributions from BR. This additional FS component is probably responsible for the extended dominance. Thus, CFS is not simply an extension of BR.

We analyzed the duration of the first, second, and third dominance periods in Experiment 2 (Figure 2.2 a and b). In BR, the duration of successive dominance is independent and the mean of both is constant, indicating that rivalry does not accelerate nor slow down and that it does not employ a periodic mechanism (Fox and Hermann, 1967; Blake et al., 1990).

This is different in CFS. We tested if the successive dominance durations are constant in a subset of 176 CFS trials of Experiment 2 (four subjects x four trials x eleven levels of flash intervals (10, 20, 30, 40, 50, 60, 80, 160, 320, 640, 1280 msec)). Some trials were excluded from the analysis if the trial terminated before the second (in Figure 2.14 a) or third (Figure 2.14 b) dominance was completed. For CFS (blue crosses), we found that the Mondrian dominance (D_M) differs between the first ($D_{M,1}$) and the second ($D_{M,2}$) intervals (Figure 2.14 a: two-tailed paired *t*-test, $P < 0.001$, *t*-score = 3.89, d.f. = 147), but not between the second ($D_{M,2}$) and the third ($D_{M,3}$) intervals (Figure 2.14 b: $P > 0.9$, *t*-score = -0.034, d.f. = 137). For 16 BR trials (red circles), there was no difference

between $D_{M,1}$ and $D_{M,2}$ (Figure 2.14 a: $P > 0.7$, t-score = -0.39, d.f. = 15) and $D_{M,2}$ and $D_{M,3}$ (Figure 2.14 b: $P > 0.4$, t-score = -0.84, d.f. = 14).

In Figure 2.15, we plot the data in Figure 2.14 in a different representation; the difference between $D_{M,1}$ and $D_{M,2}$ is plotted as a function of flash interval as blue points and that between $D_{M,2}$ and $D_{M,3}$ is plotted as red points. The number beside each point denotes the number of trials included in the analysis.

$D_{M,1}$ was significantly longer than $D_{M,2}$ and $D_{M,3}$ in CFS, perhaps due to fatigue or adaptation of the mechanism that is sensitive to repetitive FS. This suggests that CFS involves an additional flash-sensitive component that is not present for BR.

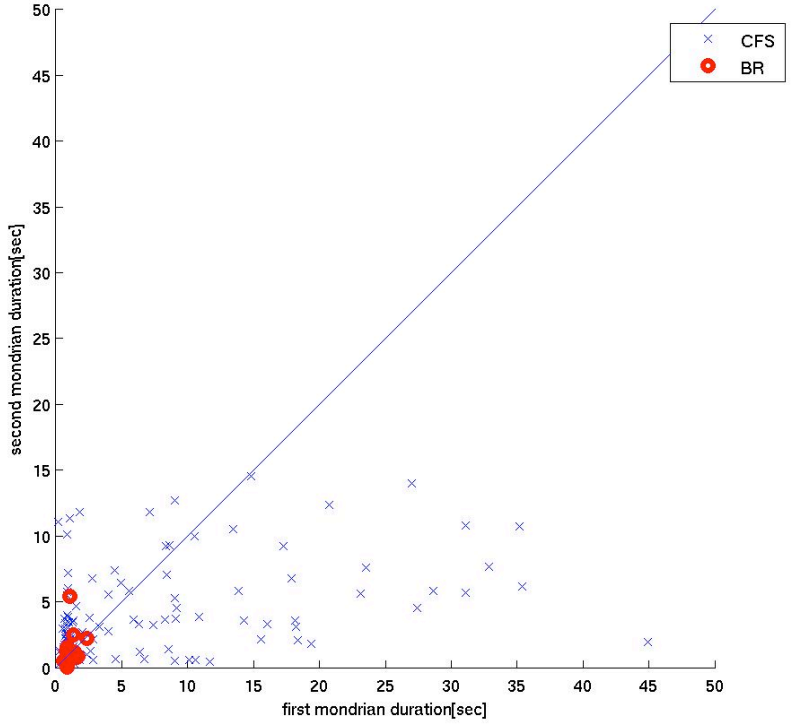
CFS
BR
second mondrian duration[sec]
first mondrian duration[sec]
0
5
10
15
20
25
30
35
40
45
50

(a)

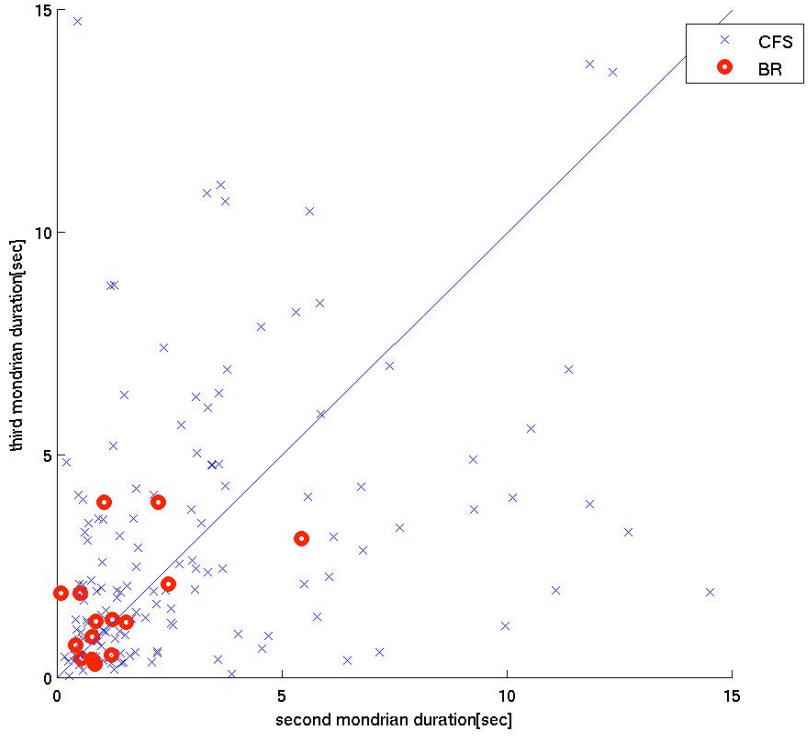
CFS
BR
third mondrian duration[sec]
second mondrian duration[sec]
0
5
10
15

(b)

Figure 2.14 Scatter plot for the successive dominance durations

We analyzed the dominance durations in Experiment 2. Red circles and blue crosses are data for BR and CFS, respectively. Data for CFS consists of a subset of 176 trials. (**a**) x- and y-axes are durations of first and second dominance, respectively. We excluded the trials if the second dominance did not finish before the trial ended (60 sec). Clearly first-dominance duration is much longer than the second dominance only in CFS but not in BR. (**b**) x- and y-axes are durations of second and third dominance, respectively. We excluded the trials if the third dominance did not finish before the trial ended (60 sec). Note that the duration is much shorter than in a and that there is no apparent difference between second- and third-dominance durations for both CFS and BR.

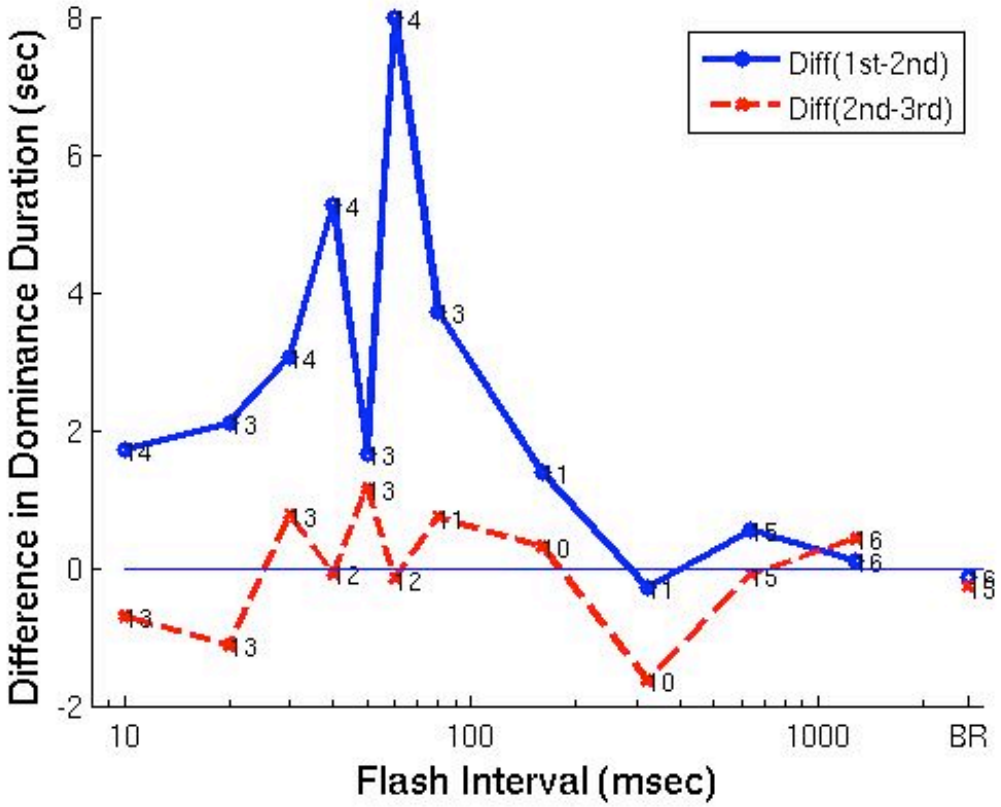

Figure 2.15 Difference between the successive duration of dominance as a function of flash interval in CFS

The same data in Figure 2.14 are plotted in a different representation; data in Figure 2.14 a and b correspond to blue and red points here, respectively. Y-axis is the difference in dominance duration between the first and second (solid blue line) or second and third (broken red line) dominance phases. The numbers along each data point show the number of trials that contributed to that data point. Note that the number of trials that contributed to the point for flash intervals from 60-320ms is rather small because CFS is too effective

and many trials ended before second (blue) or third (red) dominance finished before the 60 sec trials ended.

2.8.2. Appendix 2: The effects of novelty in CFS

In preliminary experiments we investigated the question of how important novelty is for CFS. The basic set-up is the same as for Experiment 2, but we fixed the flash interval and varied the number of Mondrian patterns, which was fixed to 20 in Experiment 2. If novelty plays an important role, we would expect to see the total dominance duration (TD) and the mean duration (MD) increase as the number of Mondrian patterns increases.

One naïve subject and the author of this thesis participated in this experiment (Appendix 2). To avoid the saturation of the CFS effects, we set the flash interval to 200 msec and the contrast of the Gabor to 100%. The number of Mondrian patterns was varied (1, 2, 4, 8, 16, 32). The order of trials was randomized within a block of 6 trials. Each subject completed 4 blocks. The data was analyzed and represented as in Figure 2.2. The results are shown in Figure 2.16. Although we used rather abstract and meaningless Mondrian stimuli, Mondrian dominance measured as both TD and MD increased as the number of Mondrian patterns increased, while Gabor dominance stayed rather constant. This pattern is again difficult to reconcile with a notion that novel patterns are considered as 'strong' stimuli (see Section 2.4.3). We conclude that novelty plays some role in extending CFS. Note that we did not use meaningful naturalistic stimuli, usually deployed to examine the effect of novelty. Such experiments remain to be carried out.

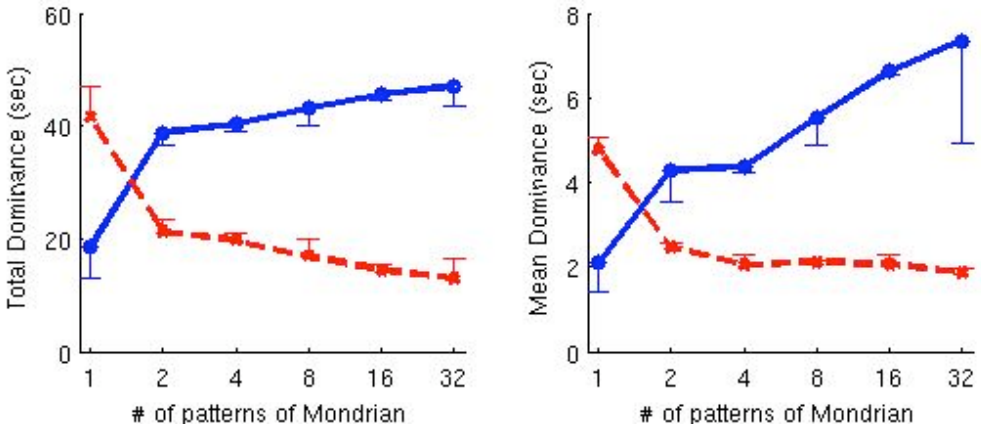

Figure 2.16 The effects of novelty in CFS

Total dominance (left) and mean dominance (right) as a function of the number of Mondrian patterns for 2 subjects. The flash interval was 200 msec. The contrast of Gabor was 100%. The error bars represent s.e.m. The solid lines correspond to the times for which the Gabor pattern remains invisible, the broken lines for when the Gabor is visible.

2.9. Acknowledgement

Connie Hofstoetter did the work that inspired me to study whether the negative afterimage can be reduced by continuous flash suppression. I appreciated discussions with her as well as Ryota Kanai and Farshad Moradi during the development of this project. Daw-An Wu suggested an experiment, similar to Experiment 7. Rufin VanRullen and Shinsuke Shimojo gave us useful comments on the manuscript that appeared in Nature Neuroscience, which contains Experiment 1, 2, and 6-9.

Randolph Blake invited me to Vanderbilt University over the summer of 2005. Experiment 3-5 were performed there with Randolph and his postdoc, Lee Gilroy. I appreciate Randolph and Lee in stimulating and helpful discussion.

All the experiments reported in this chapter were designed, programmed, and analyzed by me.

This research was funded by grants from NIMH, NSF, the Keck Foundation, and the Moore Foundation.

3. Component motion processing outside the focus of attention

3.1. Overview

Attention is thought to be key to the scientific understanding of consciousness (Posner, 1994). In everyday life, sometimes we do not notice some salient objects or events consciously if we are occupied by some thoughts or concentrating on a task at hand. Indeed, such situations have been extensively studied by psychologists, showing the importance of attention for conscious vision. This point is made very clear in some dramatic demonstrations, including inattentional blindness (Mack and Rock, 1998) and change blindness (O'Regan and Noe, 2001). In inattentional blindness, subjects do not notice a salient object appearing at fixation and cannot report on its presence during interview immediately after the trial. In change blindness subjects cannot detect unsubtle, big changes in alternating two pictures unless attention is directed to an appropriate location. Based on these phenomena, these authors as well as other researchers claim that conscious vision without any attention is impossible[10]. With appropriate expectation and familiarity of the stimuli, however, we do experience phenomenal vision at the periphery in certain aspects, even when significant attentional resources are dedicated to a demanding task at the fixation (Braun, 1998).

[10] There is an on-going debate as to how to interpret inattentional blindness and change blindness. Braun argues that inattentional blindness occurs when subjects *fail to expect* stimuli and that "inattentional" blindness is a misnomer (Braun J (2001) It's Great But Not Necessarily About Attention. In: Psyche.). Wolfe argues that inattentional blindness and change blindness occur not because subjects fail to see the objects but because they *immediately forget* them after seeing them (inattentional amnesia hypothesis).

In this chapter, we studied visual motion processing outside the focus of attention using a *dual-task* paradigm. Here, the inputs are held constant while the attentional states are varied in two conditions. In a single-task situation, only the peripheral stimulus is task-relevant and it is fully attended. In a dual-task situation, subjects perform a demanding T/L discrimination task at the fixation in addition to the peripheral task, leaving the peripheral target poorly or not at all attended. Since the display is identical across the two conditions, any difference in the peripheral task performance between the two conditions is due to whether the peripheral target is fully or poorly attended.

The design of our study in visual motion is based on the previous efforts (Lee et al., 1999b; Itti et al., 2000), where performance in an array of the static visual tasks is modeled in a quantitative and unified manner. There, the effect of visual attention is explained by an increased non-linear interaction among populations of neurons tuned to a particular orientation and spatial frequency. Here, we showed that attention strongly modulated lateral masking and local direction-selective inhibition for the component motion processing, whereas it did not affect local facilitation or local orientation-selective inhibition (See Section 3.6.7).

Here, we applied this strategy to the visual motion domain and further improved our approach by utilizing the hierarchical structure of motion processing and tailoring our novel "targeted" motion stimuli—*wavelet* motion—to a particular stage of the visual motion system. Visual motion of an object is processed in at least two stages; motion of the luminance-defined edge is detected by spatio-temporal energy filters (the component motion stage) (Adelson and Bergen, 1985). The outputs from the component stage, which suffers from ambiguity (the aperture problem), are integrated to compute the

unique direction of motion of the object (the pattern motion stage) (Movshon et al., 1985; Welch, 1989; Simoncelli and Heeger, 1998; Weiss et al., 2002). We designed each wavelet with a specific spatio-temporal tuning so that it effectively activates detectors in the first, but not the second, stage. This is different from the commonly employed random dot motion stimulus; random dot motion has a broad spatio-temporal spectrum, and its luminance step can be controlled only by a pixel-level luminance resolution, both of these factors making it difficult to control the activity of the neurons, especially in V1. Within a circular patch composed of multiple wavelets, however, each wavelet has a counterpart wavelet that moves in the opposite direction. Thereby, units for the pattern motion processing are inhibited due to motion opponency (Qian and Andersen, 1994; Heeger et al., 1999). Further, unlike intersections of the plaid motion, our wavelet composite does not possess any traceable features so that second-order motion processing is minimally involved. As a result, our results are expected to reflect attentional modulation of the component motion stage without secondary attentional effects from higher visual motion processing stages.

3.2. Summary

We describe a novel motion display, which combines advantages of luminance gratings (frequency selectivity) and of dynamic random dots (coherence manipulation) and which distinguishes between representations of component and of pattern motion. In particular, we characterize visual filters selective for component motion with contrast masking experiments. To ascertain the extent to which the representation of component motion is altered by attention, we establish contrast masking thresholds for fully and for poorly attended stimuli, controlling the availability of attention with a demanding concurrent task. Our results show a prominent threshold reduction (facilitation of detection) for low mask contrast (4%), revealing an accelerating contrast response in the low contrast regime. This contrast response of component motion mechanisms is not changed by attention, establishing a clear difference to static visual mechanisms. Second, we find that thresholds for component motion are impaired by presence of nearby components, but that this impairment is attenuated by attention. In short, attention ameliorates lateral masking for component motion. With displays that stimulate pattern motion mechanisms, we observe a direction-selective threshold elevation, which again is ameliorated by attention. In summary, we find specific attention effects not on the representation of component motion, but on representation of pattern motion.

3.3. Introduction

Contrast thresholds measured in the presence of a masking pattern reveal the tuning of, and interaction between, visual filters. The paradigm of "contrast masking" has been developed and refined in the course of extensive efforts to understand spatial vision, that is, the perception of contrast, orientation, and spatial frequency of static visual patterns (Legge and Foley, 1980; Wilson, 1980; Foley, 1994; Lee et al., 1999b; Itti et al., 2000). The perceptual account that resulted from these efforts agrees well, and in some regards quantitatively, with the tuning of, and interaction between, cortical neurons selective to contrast, orientation and spatial frequency (Geisler and Albrecht, 1997; Itti et al., 2000).

When combined with manipulations of visual attention, contrast masking experiments can also uncover how visual representations are altered by attention (Lee et al., 1999b; Carrasco et al., 2000; Zenger et al., 2000; Freeman et al., 2001; Morrone et al., 2002). For example, we found previously that attention intensifies competitive interactions among visual filters, resulting in a higher effective gain and a sharper effective tuning (Lee et al., 1999b).

Here we ask whether contrast masking experiments can uncover comparable attention effects on the visual representation of moving patterns. Traditionally, the effect of attention on the perception of visual motion was thought to be slight. Manipulations of attention with cueing and visual search paradigms typically produce little or no effect on the perception of visual motion (Raymond, 2000). However, more recent psychophysical work (Chaudhuri, 1990; Raymond et al., 1998) as well as a series of neuroimaging (Watanabe et al., 1998b; Gandhi et al., 1999; Huk and Heeger, 2000; Saenz et al., 2002)

and neurophysiological studies (Treue and Maunsell, 1996; Seidemann and Newsome, 1999; Martinez-Trujillo and Treue, 2002) have established robust attention effects on the neural representation of visual motion. We hoped that psychophysical experiments with contrast masking, which have proven so successful in spatial vision, would uncover sizeable attention effects also in the perception of moving visual patterns.

One obstacle to achieving this goal is that visual motion is represented at multiple levels in the visual system. Particularly relevant here are the levels of "component motion" and "pattern motion" (Adelson and Movshon, 1982; Welch, 1989; Wilson and Kim, 1994; Simoncelli and Heeger, 1998). Filters or neurons selective for local spatio-temporal energy represent "component motion" which is inherently ambiguous (Adelson and Bergen, 1985). Filters or neurons that integrate selected component motions over a larger region of visual space resolve this ambiguity and represent "pattern motion" (Adelson and Movshon, 1982; Welch, 1989). Several areas of visual cortex, including area V1, exhibit selectivity for component motion, whereas selectivity for pattern motion appears restricted to middle temporal cortex (area MT or V5)(Movshon et al., 1985; Movshon and Newsome, 1996; Huk and Heeger, 2002) (but also see (Guo et al., 2004)). The neural circuits that underlie this transformation are under active study (Movshon and Newsome, 1996; Heuer and Britten, 2002).

The distinct representations of component and pattern motion were first described with the help of 'moving plaid' stimuli, which superimpose two moving sinusoidal gratings (Adelson and Movshon, 1982). However, the conspicuous 'nodes' at which the gratings intersect are a drawback of this stimulus (Stoner et al., 1990; Stoner and Albright, 1992; Wilson and Kim, 1994). Schrater and colleagues avoided the problem by

filtering dynamic noise such as to achieve a comparable distribution of motion energy without introducing conspicuous features (Schrater et al., 2000). We adopted a similar approach and combined discrete moving 'wavelets' to create spatially uniform dynamic textures. In addition, we matched the motion energy carried by each 'wavelet' to the neural tuning for component motion.

To isolate perceptions based on component mechanism and to distinguish them from perceptions derived from pattern mechanisms, we took advantage of known properties of pattern selective neurons in middle temporal area MT. The response of such neurons to a preferred motion is reduced, and in some cases even suppressed, by the simultaneous presence of motion in the opposite direction. This non-linear interaction between different motion components is known as 'motion opponency' (Snowden et al., 1991; Qian and Andersen, 1994; Heeger et al., 1999). In fact, area MT as a whole responds only minimally to multiple motion components in random directions (Britten et al., 1993; Rees et al., 2000). Presumably, the response to any one component is inhibited by the simultaneous presence of the other components (Simoncelli and Heeger, 1998). Accordingly, dynamic patterns with equal energy in all directions of motion are expected to stimulate component mechanisms comparatively well and pattern mechanisms comparatively poorly.

With such a stimulus, we carried out contrast masking experiments to characterize facilitatory and inhibitory interaction among visual filters selective for component motion. Our experiments confirmed and extended the results of a number of earlier studies on motion masking (Anderson and Burr, 1985; Ferrera and Wilson, 1987; Anderson and Burr, 1989; Anderson et al., 1991; Lu and Sperling, 1995, 1996). To

ascertain the impact of attention, we used an established dual-task paradigm (Braun, 1994; Lee et al., 1997; Braun, 1998; Braun and Julesz, 1998; Lee et al., 1999b; Lee et al., 1999a; Zenger et al., 2000; Li et al., 2002) to compare masking thresholds when moving patterns are either fully or poorly attended.

3.4. Methods

3.4.1. Psychophysical task

3.4.1.1. Subjects and apparatus

Six naïve subjects participated in daily 1-hour sessions lasting for a period of 1 to 3 months (20 to 40 hours total) and were paid $13 per hour. The study was divided into three parts, with at least two subjects contributing to each part.

Stimuli were generated by a Silicon Graphics O2 workstation and displayed on a 19" raster monitor (1280 x 1024 pixels RGB) with 74 Hz refresh (13 msec/frame). Viewing was binocular from a distance of 120 cm (80 pixels per 1° visual angle). The average display luminance was 40 cd/m^2, and gamma correction combined with color bit stealing (Tyler, 1997) provided linear luminance step size of 0.07 cd/m^2. Ambient luminance was 3 cd/m^2.

Subjects initiated each trial by fixating a cross at display center and by pressing a key. Central and peripheral stimuli appeared and subjects produced one or two responses pertaining to either or both stimuli, depending on instructions (see below). The short presentation time and random polar angle of peripheral stimuli renders eye-movements counterproductive[11]. As a result, subjects maintain central fixation even when reporting only on the peripheral stimulus.

---

[11] Because the phase of individual wavelets moves alternating its polarity rapidly, visual persistence of the stimuli should be negligible. For the contrast discrimination task in our study, therefore, the optimal strategy is to integrate contrast information over the time period, especially around the 8th frame where the contrast of individual wavelets becomes maximal (Figure **3.3** and Figure 3.4). The 8th frame occurs ~104 msec after the onset of

3.4.1.2. Central task (letter discrimination)

To draw attention away from the visual periphery, we had subjects perform a demanding letter discrimination task with an array of seven letters (Figure 3.1 a) as a central task in the dual-task situation. Seven letters (T's and/or L's of size 0.5°) appeared at nine possible locations (one at fixation and the other eight horizontal, vertical, and diagonal to the fixation), rotated randomly and independently, and appeared within 1.4° eccentricity. Subjects were instructed to press the 'S' key for seven T's or seven L's and the 'F' key for six T's and one L or six L's and one T. They received auditory feedback immediately after each mistake. Each letter was followed by an independently rotated masking letter (F) to limit visual persistence. Stimulus onset asynchrony (SOA) was adjusted for each observer to maintain performance of approximately 80% (164–250 ms).

This letter task engages spatial attention to such an extent that a concurrent, attention-demanding task cannot be performed significantly above chance (Braun, 1994; Lee et al., 1997; Braun, 1998; Braun and Julesz, 1998; Lee et al., 1999b; Lee et al., 1999a; Zenger et al., 2000; Li et al., 2002).

---

the stimuli. Even if the wavelets are detected at the first frame (when their contrast is very low), a saccade to the target is too late to reach there when the contrast is maximal (in humans, the latency of saccade is approximately 200 to 250 msec). Therefore, saccades are very counter-productive in our task.

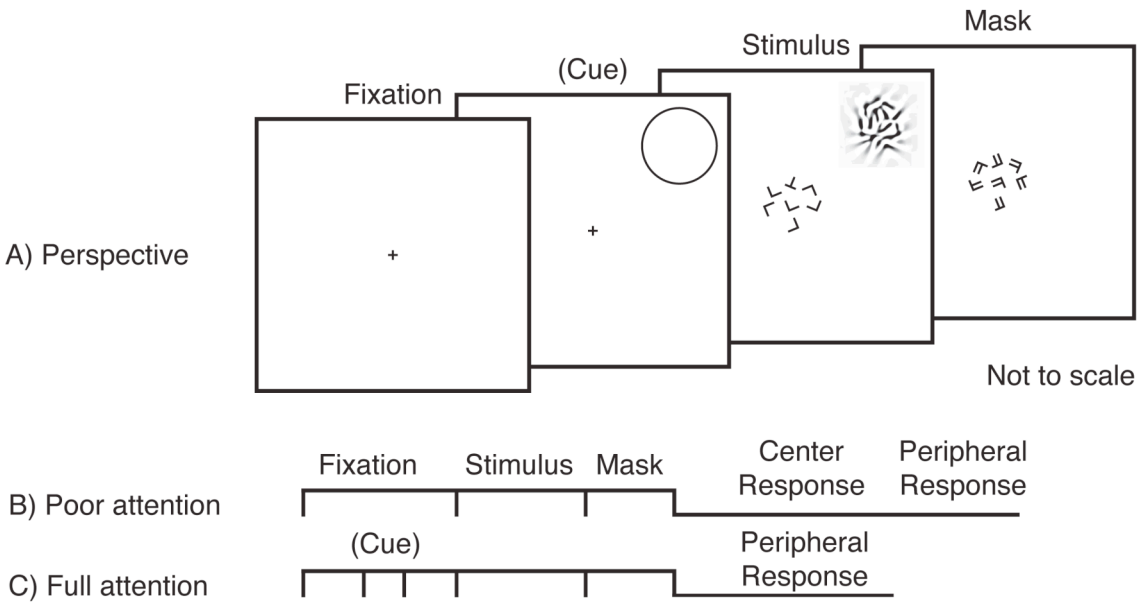

Figure 3.1 Psychophysical procedure

a) Subjects fixated at the centre of the display sequence throughout. The stimulus display consisted of two parts: a central letter cluster (masked, varying positions < 1.4° eccentricity) and a peripheral array of moving wavelets (not masked, varying position at a constant eccentricity of 3.7°). In a control experiment, a circular cue indicated the future position of the peripheral array. **b**) In the dual-task condition, subjects reported whether all letters were identical or whether one letter differed from the others and, independently, whether moving wavelets comprised only 'mask' wavelets or both 'target' and 'mask' wavelets (see Figure 3.2 and Figure 3.3). **c**) In the single-task condition, subjects ignored the central letters and reported only on the peripheral array.

3.4.1.3. Peripheral task (contrast masking)

The peripheral stimulus appeared at a random polar angle at 3.7° eccentricity and filled a circular region of up to 1.5° in diameter (Figure 3.1 a). It was composed of 1 or 23 pairs of 'moving wavelets' (see below), each pair comprising a target wavelet and a superimposed masker wavelet. A standard adaptive staircase method was used to establish contrast thresholds for target detection (2-down-1-up, converging to 70.7% correct performance (Levitt, 1971)). Subjects were instructed to press the 'J' key to report presence of target wavelets and the 'L' key to report absence (contrast masking) (Figure 3.1 b). Contrast of masking wavelets remained fixed throughout each block, which terminated after 25 reversals (typically 80–120 trials for a staircase). Subjects received immediate auditory feedback after each mistake. The mean target contrast during the last 40 trials of each block was taken as a threshold estimate. After one to three practice blocks at each condition (see below), the contrast step size was set at approximately one-tenth of the threshold value achieved during practice. The initial contrast value was set at approximately twice the threshold value. When a staircase failed to converge, the results were discarded and stair parameters adjusted. For each subject, data from 3-12 converged staircases were combined to estimate the individual threshold and standard error. More staircases were collected for high mask contrasts to compensate for larger variance.

For each observer, a mean threshold and a standard error were computed by averaging over all staircases collected at each condition. A representative threshold and standard error were obtained by averaging the means and square standard errors of all

subjects. Note that the resulting error bars reflect the variance observed within subjects, not between subjects.

In a subset of experiments, the peripheral stimulus was preceded by a cue (a thin white circle of 30% contrast and 1.8° in diameter, Figure 3.1 a). When present, the cue was flashed briefly (26 ms) and preceded the stimulus onset by 104 msec. The relatively short cue lead time was chosen to ensure a facilitatory effect of cueing and to eliminate the possibility of 'inhibition of return' (Klein, 2000).

3.4.1.4. Dual task

When performing both tasks concurrently, subjects responded to the central task first and to the peripheral task second (Figure 3.1 c). Subjects were instructed firmly to give priority to the central task and to perform as well as they possibly could on the peripheral task. The dual-task blocks with central performance lower than 75% were discarded, to ensure enough attention was paid to the central task. Four subjects completed 536 blocks of dual task in total.

Note that subjects are aware of the peripheral wavelets and are able to judge their contrast even though their attention focuses on the central task. This allows thresholds to be established for 'poorly attended' stimuli (Lee et al., 1997, 1999a; Zenger et al., 2000; Morrone et al., 2002).

We compared thresholds with poor attention (peripheral-task performance in the dual task) to thresholds with full attention (the peripheral task is carried out alone and the central task was ignored). However, the central and peripheral stimuli were always both

present, and subjects always fixated at display center, so that the visual inputs were identical under single and dual task conditions.

3.4.2. Log Gabor wavelets

3.4.2.1. Single Log Gabor wavelet

We used a set of self-similar wavelets to generate visual motion. To facilitate comparison with the sensitivity of V1 neurons, we chose to use log Gabor wavelets (Field, 1987) instead of the more conventional linear Gabors (Figure 3.2. See the demos at http://www.klab.caltech.edu/~naotsu/DemoComponentMotion/index.html). Log and linear Gabors are similar in that both present a drifting contrast phase within a stationary spatiotemporal envelope and in that Fourier energy is concentrated around one particular spatiotemporal frequency $(\Omega_x,\Omega_y,\Omega_t)$. However, in the case of log Gabors, the Fourier energy conforms to Gaussian distributions with respect to the logarithm of spatial frequency and the logarithm of temporal frequency (as well as with respect to linear spatial direction), as is also the case for neurons in V1 (Geisler and Albrecht, 1997). Accordingly, the Fourier amplitude of a log Gabor wavelet is given by

$$E(\omega_x,\omega_y,\omega_t) = E(\omega_r,\theta,\omega_t) = \frac{\sqrt{\pi^3}}{\Omega_r^2\,\Omega_t\,\Lambda_r^2\,\Lambda_t\,\Lambda_\theta \ln^2 2}\exp\left(-\frac{\ln^2\frac{w_r}{\Omega_r}}{2\Lambda_r^2\ln^2 2}\right)\left(A^+_{(\theta)}\times B^+_{(\omega_t)}e^{i\Phi} + A^-_{(\theta)}\times B^-_{(\omega_t)}e^{-i\Phi}\right)$$

$$\omega_r = \sqrt{\omega_x^2+\omega_y^2}$$

$$\theta = \arctan(\omega_x/\omega_y)$$

$$\Omega^+_\theta = \Omega_\theta - \pi$$

$$\Omega^-_\theta = \Omega_\theta$$

where the Cartesian coordinates of Fourier space (ω_x, ω_y, ω_t) are replaced by polar coordinates (ω_r, θ, ω_t). Ω_r [cpd], Ω_t [Hz], and Ω_θ [°] are the peak spatial and temporal

frequencies and directions, respectively; Λ_r [octaves], Λ_t [octaves], and Λ_θ [°] are the standard deviations or bandwidth; and Φ is the phase of the wavelet.

A spatiotemporal wavelet W(x, y, t) was obtained as the inverse Fourier transform of $E(\omega_x,\omega_y,\omega_t)$ (Schrater et al., 2000). The normalization of $E(\omega_x,\omega_y,\omega_t)$ was chosen such that |W(x,y,t)| takes maximal values on the order of unity. The same normalization factor was used for all the 144 wavelets.

The functions A^+, A^-, B^+, and B^- denote positive and negative lobes of the Fourier amplitude, which jointly determine the wavelet motion in space-time.

For example, a horizontally oriented and vertically upward-moving wavelet (going in the 90° direction) has A^+ with Ω_θ = -90° and A^- with Ω_θ = +90°, whereas a downward moving wavelet has A^+ with Ω_θ = +90° and A^- with Ω_θ = -90°.

$A^\pm$ gives the direction dependency θ,

$$A^{\pm}_{(\theta)} = \exp\left(-\frac{\left(\theta-\Omega^{\pm}_{\theta}\right)^2}{2\Lambda^2_\theta}\right)$$

and, $B^\pm$ gives the Gaussian dependency on the logarithm of temporal frequency ω_t

$$B^{\pm}_{(\omega_t)} = \exp\left(-\frac{\ln^2\frac{\left|\pm\omega_t\right|_+}{\Omega_t}}{2\Lambda^2_t\ln^2 2}\right) \quad where \quad \left|x\right|_+ = \begin{cases} x & if \quad x \geq 0 \\ 0 & if \quad x < 0 \end{cases}.$$

The dimensions of each log Gabor wavelet were 128 pixels x 128 pixels x 16 video frames. The peaks and standard deviations of Fourier amplitude were Ω_r = 2.5 cpd, Λ_r = 0.6 octave, Λ_θ = 13°, Ω_t = 6.0 Hz, and Λ_t = 0.6 octaves. For comparison, the median values for V1 neurons of macaque are Ω_r = 4.2 cpd, Λ_r = 0.72 octave, Λ_θ = 15°, Ω_t = 7.2

Hz, and Λ_t = 1.2 octave (Geisler and Albrecht, 1997). Note, however, that V1 neurons in a given locus show a continuum of spatial frequency peaks over a very broad range of more than 3 octaves in the parafoveal (3 – 5°) region (De Valois et al., 1982).

Using 3D inverse Fourier transform, we computed 144 wavelets covering 36 directions (0, 10, 20°…) and 4 phases (0, 90, 180, and 270°), each in the form of a 3D real-valued matrix (128 x 128 pixels x 16 frames). The contrast of individual wavelets was defined as

$$C_{set} = \frac{\max_{set} W_{(x,y,t)} - \min_{set} W_{(x,y,t)}}{\max_{set} W_{(x,y,t)} + \min_{set} W_{(x,y,t)}}$$

where maxima and minima are taken over the entire set of 144 wavelets. Additional contrast values were obtained by linear scaling.

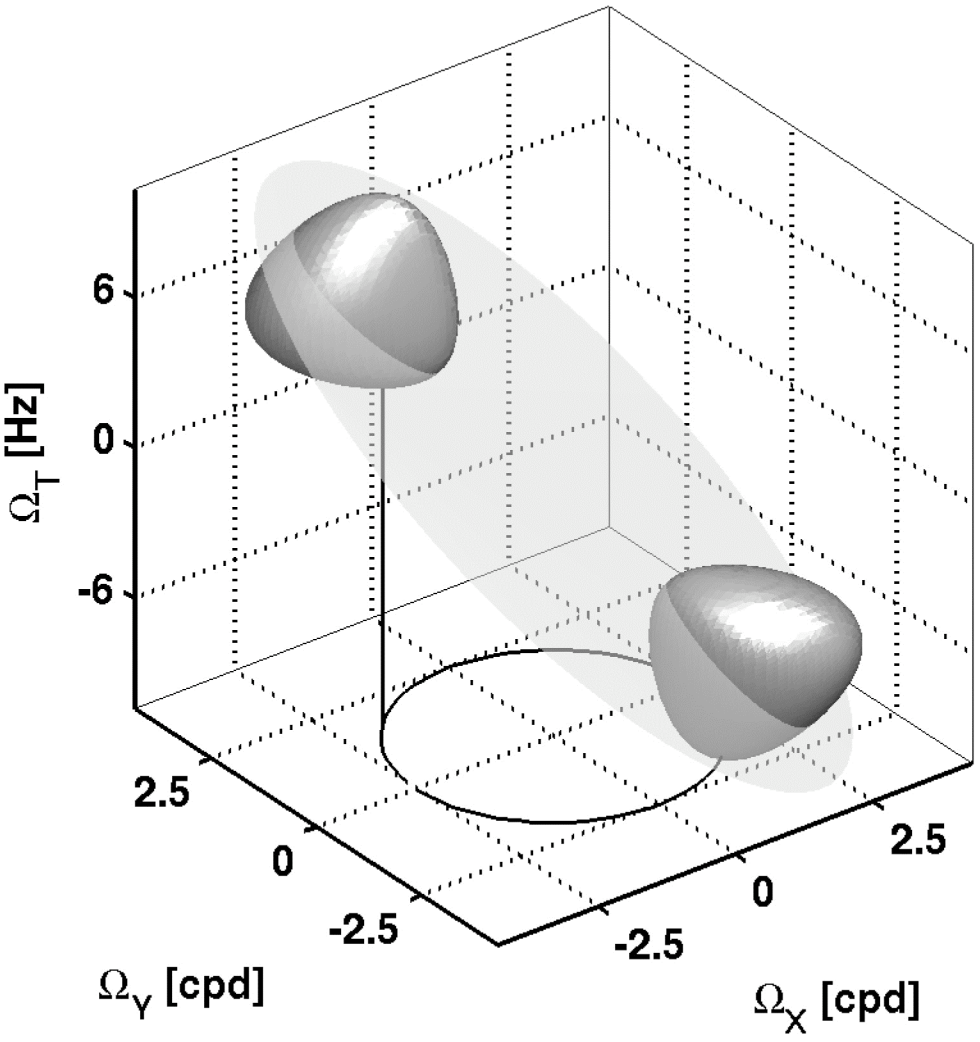

Figure 3.2 Tuning characteristics of a log Gabor wavelet in Fourier space

Here represented by an iso-power surface (7% of peak power) and its symmetry plane. Tuning is separable in log spatial frequency, log temporal frequency, and direction, and conforms to a Gaussian distribution in each of these dimensions. The peak sensitivities and standard deviations of the wavelets used were Ω_r = 2.5 cpd, Λ_r = 0.6 octaves, Ω_t = 6.0 Hz, Λ_t = 0.6 octaves, and $\Omega_\theta = 45°$, $\Lambda_\theta = 13°$. For comparison, the median values for

neurons in area V1 of macaque are Ω_r = 4.2 cpd, Λ_r = 0.72 octaves, Ω_t = 7.2 Hz, Λ_t = 1.2 octaves, and Λ_θ = 15° (Geisler and Albrecht, 1997). The circle at the bottom indicates the projection of the peak spatial frequency (Ω_r = 2.5 cpd) for wavelets of all possible directions with identical spatio-temporal tuning.

3.4.2.2. Pairs of wavelets (target & masker)

We formed pairs of wavelets by superimposing target and masker wavelets that coextended in space and time (Figure 3.3. See the demos at http://www.klab.caltech.edu/~naotsu/DemoComponentMotion/Single.html). The phase difference (relative phase) between target and masker wavelets was fixed at 0° to maximize the interaction. The direction difference (relative direction) between target and masker wavelets (0, 30, 90, 150, and 180°) and masker contrast (0 to 64%) was held constant during each staircase (while target contrast varied from trial to trial, see above).

With a relative direction of 0°, a pair of target and masker wavelets effectively forms a single wavelet (of higher contrast). For other relative directions, a target and masker pair forms various 'interference patterns.' For example, an intersection may be visible at 90°, and a contrast flash or pulse may be seen at 180° (a 'counterphase' situation). Well above threshold contrast, subjects tend to use the presence or absence of these 'interference patterns' as a cue to the presence or absence of the target wavelet. At near threshold contrast, however, 'interference patterns' are not visible, obliging subjects to base their judgment solely on contrast information.

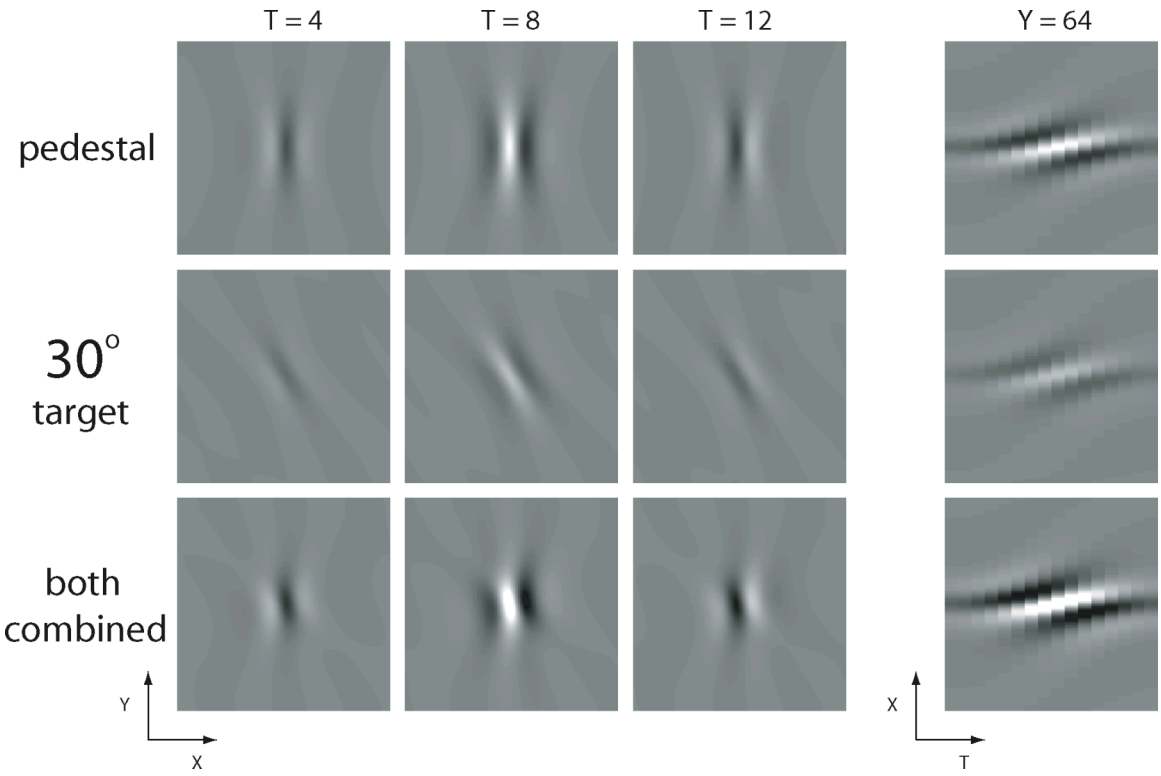

Figure 3.3 Superposition of target and mask wavelets

Wavelets occupy a volume in space-time (X=1…128, Y=1…128, T=1…16) and are represented by slices through this volume. (**Top**) Mask wavelet of a given contrast (exaggerated for clarity). Instantaneous appearance is illustrated by three X-Y slices (at times T=4,8,12) and temporal evolution by one X-T slice (at position Y=64). (**Middle**) Target wavelet of half the masking contrast, here differing by 30° in its direction of motion from mask wavelet. Other directional differences used were 0, 90, 150, and 180° (not shown). (**Bottom**) Target and mask wavelet superimposed. Both the instantaneous appearance and temporal evolution are affected by the superposition. See the demos at http://www.klab.caltech.edu/~naotsu/DemoComponentMotion/Single.html

3.4.2.3. Wavelet composites

To create wavelet composites, 23 wavelet pairs were placed randomly (but with a minimal center-to-center spacing of 0.25°) in a circular area of 1.5° diameters, centered at an eccentricity ε = 3.7°. For comparison, the average diameter of central receptive fields at this eccentricity has been estimated as 0.22° in area V1 (0.224° + 0.000617 × ε; (Dow et al., 1981)) and as 3.3° in area MT (1.04° + 0.61 × ε; (Albright and Desimone, 1987)).

Three conditions were compared (Figure 3.4): (i) a single wavelet pair of random direction placed randomly in the 1.5° area, (ii) 23 parallel wavelet pairs of identical orientation and direction of motion, chosen randomly for each trial (See the demo at http://www.klab.caltech.edu/~naotsu/DemoComponentMotion/Parallel.html), and (iii) 23 random wavelets pairs with different orientations and directions of motion, randomly assigned to each pair (See the demo at http://www.klab.caltech.edu/~naotsu/DemoComponentMotion/Random.html). Iso-orientation-specific lateral interactions between wavelet pairs are expected to be maximized by the parallel and minimized by the random configuration. Lateral interaction, which is insensitive to the relative orientation, should be present for both cases. In the absence of lateral interactions, thresholds should depend only on local interactions between target and masker wavelets of each pair.

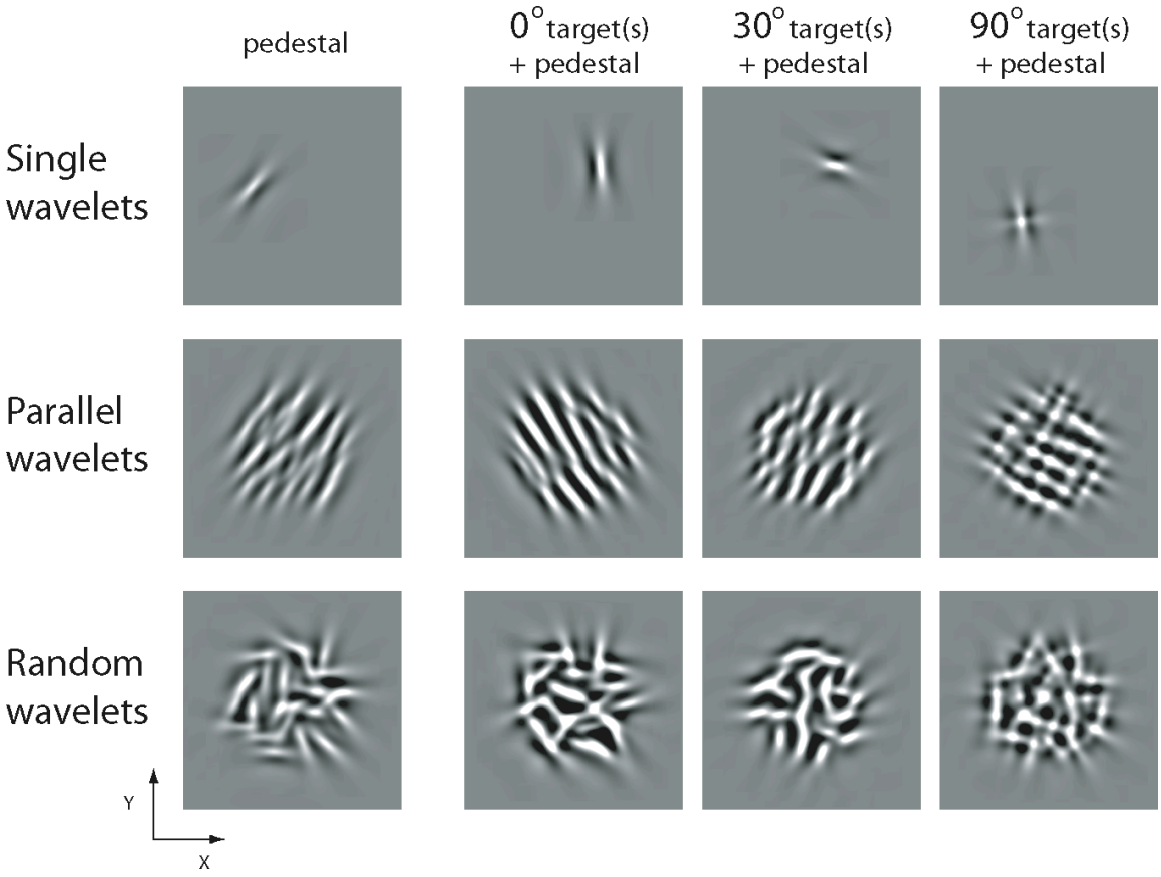

Figure 3.4 Instantaneous appearance of wavelet arrays (at the time of maximal contrast)

(**Top**) Single wavelets. A single wavelet pair (target and mask) appeared at a random position within a circular array of diameter 1.5°. Observers discriminated between a mask wavelet (leftmost frame) and the superposition of mask and target wavelets (other frames). The respective directions of motion of mask and target wavelets differed by 0, 30, 90, 150, or 180° (only the first three possibilities are shown. For 150 and 180°, see http://www.klab.caltech.edu/~naotsu/DemoComponentMotion/Parallel.html and http://www.klab.caltech.edu/~naotsu/DemoComponentMotion/Random.html). Target

contrast was adjusted to determine the observer's discrimination thresholds. In the 0° case, the threshold was in effect a 'contrast increment' threshold. In the other cases, the measurement determined a 'contrast masking' threshold. (**Middle**) Parallel wavelets. Twenty-three wavelet pairs appeared within the circular array. Observers discriminated between mask wavelets only (leftmost frame) and the superposition of an equal number of mask and target wavelets (other frames). Within each pair, the relative directions of target and mask were fixed (0, 30, 90, 150, or 180°), in some cases creating the appearance of a 'plaid' (rightmost frame). (**Bottom**) Random wavelets. Similar to above, except that mask (pedestal) wavelets assumed random directions. The relative direction of target and mask wavelets remained fixed within each pair. See the demos at http://www.klab.caltech.edu/~naotsu/DemoComponentMotion/index.html

3.4.3. Contingency analysis

In any dual-task situation, the observer may on some trials favor one task and on other trials favor the other task. If this should occur, there would be a positive correlation between success on one task and failure on the other. Let the numbers of trials with particular dual-task outcomes be n_{00}, n_{10}, n_{01}, n_{11}, where the subscripts 1 and 0 stand for correct and incorrect responses in the central and peripheral task, respectively. If the tasks are performed independently (i.e., in the absence of any correlation), the expected numbers of trials of each type are

$$\bar{n}_{00} = \frac{(n_{01}+n_{00})(n_{10}+n_{00})}{n_{00}+n_{01}+n_{10}+n_{11}} \qquad \bar{n}_{01} = \frac{(n_{01}+n_{00})(n_{11}+n_{01})}{n_{00}+n_{01}+n_{10}+n_{11}} \qquad \bar{n}_{10} = \frac{(n_{11}+n_{10})(n_{10}+n_{00})}{n_{00}+n_{01}+n_{10}+n_{11}} \qquad \bar{n}_{11} = \frac{(n_{11}+n_{10})(n_{11}+n_{01})}{n_{00}+n_{01}+n_{10}+n_{11}}$$

where the statistical significance of any difference between expected and actual values is given by the χ^2 measure of association:

$$\chi^2 = \sum_{i,j} \frac{\left(n_{ij} - \bar{n}_{ij}\right)^2}{\bar{n}_{ij}} \quad .$$

As peripheral performance changes systematically during each staircase block, we restricted the contingency analysis to the subset of trials in which target contrast was near threshold contrast. This ensured that average peripheral performance in the analyzed subset was close to the limiting performance of the staircase (i.e., 70.7%) (Levitt, 1971). In case of a negative correlation, there would be fewer-than-expected trials with identical outcomes on both tasks ($n_{00} < \bar{n}_{00}, n_{11} < \bar{n}_{11}$) and more-than-expected trials with different outcomes on both tasks ($n_{10} > \bar{n}_{10}, n_{01} > \bar{n}_{01}$).

3.4.4. Statistical Analysis

As a criterion of significance, we used $P < 0.05$ unless noted otherwise. When we used ANOVA to assess contrast-increment thresholds, we submitted log contrast thresholds in order to achieve the equal variance assumption between lower and higher mask contrast regimes. When comparing across stimulus configuration using *different* subjects, data from a single subject is collapsed, and variability across subjects served as residual. When comparing other factors (i.e., effects of mask contrast and direction, withdrawing attention, and so on) using the *same* subjects, variability among multiple blocks served as residual. Only interesting main effects and interactions are reported (i.e., except for main effects of subject and so on).

3.5. ***Results***

We studied interactions between components of visual motion and how such interactions are altered by attention. Our stimuli combined several 'moving wavelets,' which approximately matched the known sensitivity of component motion mechanisms in the visual cortex (Figure 3.2). We were interested both in *local interactions*, such as between wavelets superimposed in space and time, and in *lateral interactions,* such as those that may arise from the presence of nearby wavelets in close proximity. To characterize both interactions, we measured the contrast necessary for detecting target wavelets in the presence of masker wavelets. Target and masker wavelets were superimposed one-on-one, so that each target wavelet was paired with a masker wavelet at the same spatio-temporal location (Figure 3.3). To characterize *local interactions*, we varied the configuration *within* each wavelet pair, varying contrast and direction of a masker wavelet relative to the associated target wavelet. To assess *lateral interactions*, we manipulated the configuration *between* wavelet pairs, comparing multiple *random* wavelets, multiple *parallel* wavelets, and *single* wavelets (Figure 3.4). We expected lateral interactions to be moderate for random wavelets and maximal for parallel wavelets. Single wavelets cannot, of course, generate lateral interactions.

Our peripheral stimulus was a roughly circular array of wavelets measuring 1.5° in diameter (see Methods 3.4.1.3). Subjects reported the presence or absence of target wavelets, which were paired with superimposed masker wavelets. As target and masker wavelets of each pair shared the same phase (relative phase 0°), a special situation arose when they also shared the same direction (relative direction 0°). In this case, target and

masker wavelets had identical waveforms and the observer's task was reduced to detecting a contrast increment, contrast-increment threshold measurements. For other relative directions (30, 90, 150, and 180°), the observer's task involved detecting one complex pattern (array of target wavelets) in the presence of another (array of masker wavelets), contrast-masking threshold measurements.

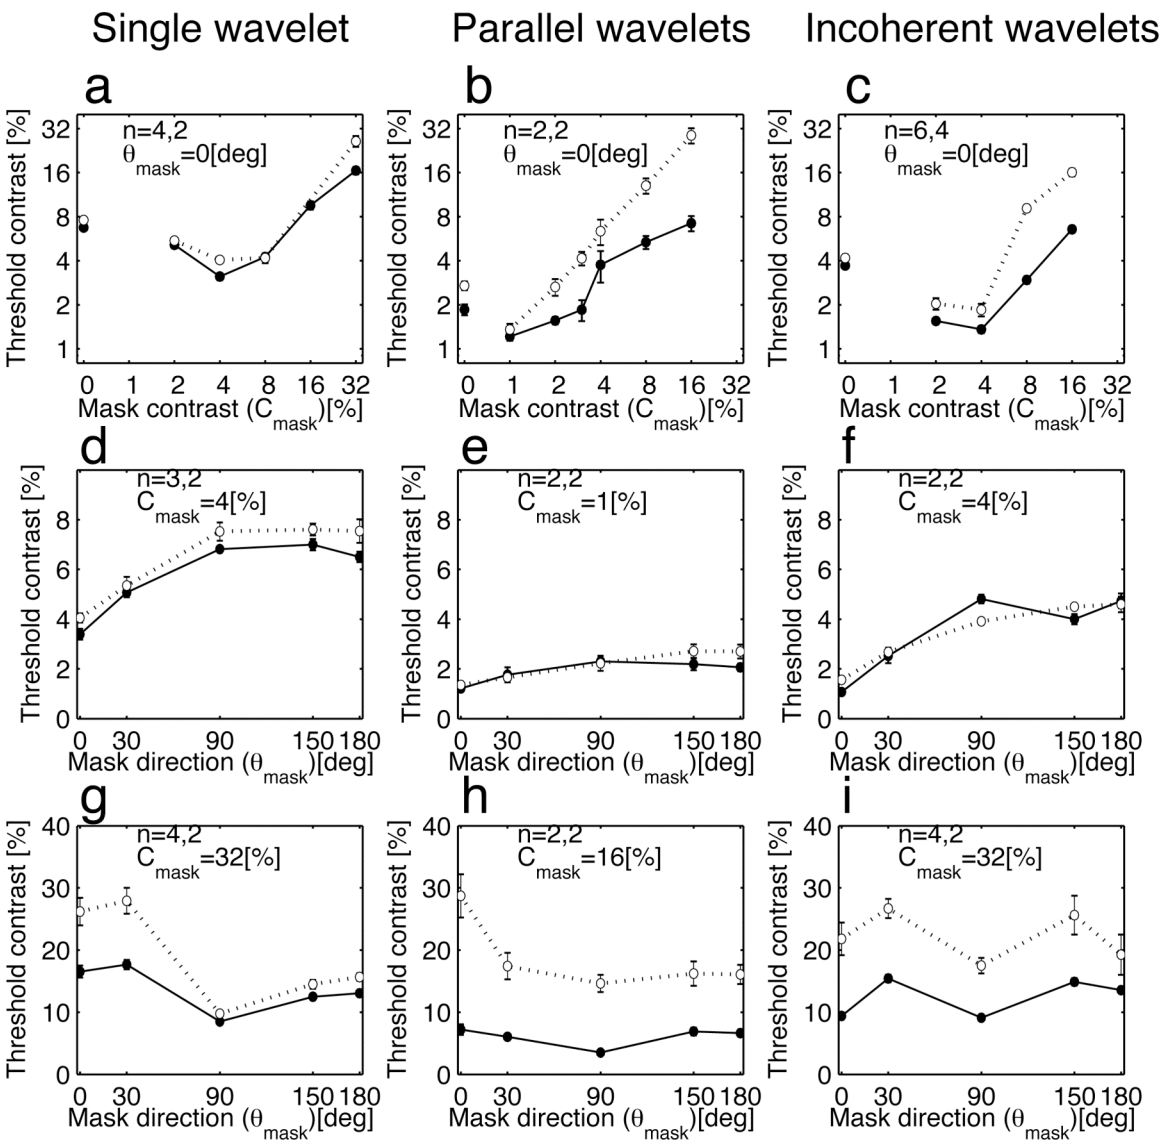

Figure 3.5 Contrast increment and contrast masking thresholds measured with full and poor attention

Thresholds are shown with closed and open symbols for full and poor attention, respectively. (**Top, a-c**) Contrast increment thresholds as a function of mask contrast in log scale for both x- and y-axes. (**Middle d-f**) Thresholds in linear scale as a function of

the relative direction of target and mask wavelets at low masking contrast (1% or 4%), which tend to facilitate target detection. (**Bottom g-i**) Thresholds in linear scale as a function of the relative direction of target and mask wavelets at high masking contrast (16% or 32%), which tend to inhibit target detection. Results for single wavelets, parallel wavelets, and random wavelets are shown in the **left (a, d, g), middle (b, e, h), and right (c, f, i)** columns, respectively.

3.5.1. Contrast-increment thresholds with full attention (Figure 3.5 a-c)

The average absolute detection threshold (masker contrast 0%, relative direction 0°) was 6.7 ± 0.2% for single, 1.9 ± 0.2% for parallel, and 3.7 ± 0.1% for random wavelets (mean and standard error from 4, 2, and 6 observers, respectively). This rank order of wavelet configurations (parallel < random < single) held consistently for all observers. The difference between configurations was significant (one-tailed *t*-test on single vs. random: *t*-score = 4.73, d.f. = 8; random vs. parallel, *t*-score = 2.04, d.f. = 6). Multiple wavelets are expected to exhibit lower thresholds than single wavelets, due to signal summation. The additional, two-fold difference between parallel and random wavelets may reflect sensitive higher-level mechanisms for coherent motion (see Discussion 3.6.3).

At relative direction 0°, detecting target wavelets is equivalent to detecting a contrast increment. Contrast-increment thresholds measured for a series of masker contrasts (0, 1, 2, 4, 16, and 32%) are depicted in Figure 3.5 a-c. As is typical, thresholds improve as masker contrast increases from zero to approximately the level of detection threshold ('dipper'), before rising as masker contrast increases beyond the detection level. The former regime (reduced thresholds) is thought to reflect facilitatory local interactions between target and masker wavelets, whereas the latter regime (elevated thresholds) is likely to reflect inhibitory local interactions (Zenger and Sagi, 1996; Itti et al., 2000). The 'dipper' was pronounced for single wavelets (46% reduction; from 6.7 ± 0.2% at 0% mask to 3.1 ± 0.2% at 4% mask) and for random wavelets (37% reduction; 3.7 ± 0.1% to 1.4 ± 0.1%), but comparatively slight for parallel wavelets (65% reduction; 1.9 ± 0.2% to 1.2 ± 0.1% at 1% mask). Two-way ANOVA (subject x contrast

[at the dipper and at the absolute detection]) on log contrast thresholds revealed a very strong main effect of contrast ($P < 1e\text{-}6$), with $F = 66.76, 42.07, 179.55$ for the single, parallel, and random configurations, respectively. In the high-contrast regime (above 4%), all three wavelet configurations produced similar results. This is confirmed by a two-way ANOVA (stimulus configuration x contrast at 8% and 16% mask contrast) without significant configuration x contrast interaction ($F = 1.03$), whereas at low contrast mask (0, 2, and 4%) significant interaction was found ($F = 5.24$). Previous studies of contrast increment thresholds for moving stimuli report little or no 'dipper' (see Discussion 3.6.5).

The location of the peripheral stimulus varied from trial to trial, creating 'positional uncertainty' when the total stimulus (target + mask) contrast was at or near the detection threshold (Solomon et al., 1997; Foley and Schwarz, 1998). To determine how much of our results (especially at low masker contrast) were contaminated by positional uncertainty, we re-measured some thresholds with positional cueing for the peripheral stimulus to remove uncertainty (see Methods 3.4.1.3). The absolute detection thresholds for both single and random wavelets were reduced by 82% from 6.7 ± 0.2% to 5.5 ± 0.2% ($F = 17.76$) and by 68% from 3.8 ± 0.1% to 2.6 ± 0.2% ($F = 26.41$). Though cueing reduced the absolute detection thresholds, they nevertheless remained above the minimal thresholds at the bottom of the dipper with full attention (3.2 ± 0.2% for single and 1.3 ± 0.1% for random).

To make sure that the dipper is present with cueing, we measured the thresholds at low mask contrast (2, 4, and 8%) in two subjects. Reduction of thresholds was evident for both stimulus types. For random wavelets, the minimal threshold was 1.6 ± 0.1% at

mask contrast of 2% (vs. detection threshold of 2.6 ± 0.2%), whereas for single wavelets the minimal threshold was 2.7 ± 0.1% at mask contrast 4% (vs. detection threshold of 5.5 ± 0.2%). Accordingly, the dipper was obtained for both stimuli even after removing positional uncertainty.

3.5.2. Contrast-masking thresholds with full attention (Figure 3.5 d-i)

At relative directions other than 0°, the nature of the task changes from detecting a contrast increment to discriminating a complex, moving pattern (target wavelets) that is superimposed over another complex, moving pattern (masker wavelets). We studied the dependence of thresholds on the relative direction of motion within each wavelet pair in order to characterize the direction tuning of local interactions between target and masker wavelets. In addition, we used two levels of masker contrast in order to characterize separately the facilitatory and inhibitory regimes of masker-target interactions. A masker contrast at the lowest point of the 'dipper' (4% for single and random wavelets and 1% for parallel wavelets) resulted in facilitatory local interactions. The highest masker contrast available (32% for single and random and 16% for parallel wavelets) revealed inhibitory local interactions. The relative direction of target and masker wavelets in each pair was set to 0, 30, 90, 150, or 180°. A relative direction of 0° produced 'in-phase' motion, whereas one of 180° resulted in 'counter-phase' (or opponent) motion.

The effect of *low-contrast* maskers is depicted in Figure 3.5 d-f (filled symbols). As expected, the overall effect was facilitatory in that thresholds were lower than or equal to absolute detection threshold for each wavelet configuration. In general, thresholds increased as relative direction increased from 0 to 30°, but for relative directions of 90°

and above reached a plateau at the level of the absolute detection threshold. The particular results for each wavelet configuration follow closely the contrast increment thresholds in Figure 3.5 a-c. In each case, the lowest point (single 3.4 ± 0.2%, parallel 1.2 ± 0.1%, random 1.1 ± 0.1%) and the plateau level, [average across 90, 150, and 180°] (6.8 ± 0.2%, 2.2 ± 0.2%, 4.5 ± 0.2%) in Figure 3.5 d-f correspond to, respectively, the lowest point of the 'dipper' (3.1 ± 0.2%, 1.2 ± 0.1%, 1.4 ± 0.1%) and the absolute detection threshold (6.7 ± 0.2%, 1.9 ± 0.2%, 3.7 ± 0.1%) in Figure 3.5 a-c. In short, the threshold increase is pronounced for wavelet configurations with a pronounced 'dipper' (single and random wavelets) but slight for the configuration with a shallow 'dipper' (parallel wavelets).

Interestingly, no significant facilitation occurred at 180° (opponent or 'counterphase' motion), even though target and masker wavelets shared the same spatial orientation. This shows that the local facilitation is selective for direction of motion (rather than merely for spatial orientation) and thus that it is mediated by direction-selective motion-specific mechanisms.

Overall, the differences between the results for three wavelet configurations appeared quantitative rather than qualitative and mirrored the rank order of absolute detection thresholds (see above): Thresholds were highest for single wavelets and reduced for multiple wavelets, presumably reflecting summation of multiple signals. In addition, thresholds for the parallel configuration were up to 2 times lower than for the random configuration, perhaps due to facilitatory lateral interaction (see Discussion 3.6.5).

The effect of *high-contrast* maskers (Figure 3.5 g-i, filled symbols) tended to be inhibitory, in that thresholds were generally higher than absolute detection thresholds. For all wavelet configurations, thresholds were lowest at relative direction 90°. In the case of single wavelets, thresholds with 90° maskers were only a little above absolute detection thresholds (8.5 ± 0.6% and 6.7 ± 0.5%, respectively). In the case of parallel and random wavelets, 90° maskers raise thresholds approximately two-fold above absolute detection levels (from 1.9 ± 0.2% to 3.5 ± 0.4% and 3.7 ± 0.1% to 9.1 ± 0.4%, respectively). For relative directions either greater or smaller than 90°, thresholds rise to even higher levels. The details of this rise suggest that there may be qualitative differences between wavelet configurations.

In the case of single wavelets, there was a pronounced asymmetry in the effect of maskers at 0 and 30° (16.5 ± 1.0% and 17.7 ± 0.7%) on the one hand, and at 150 and 180° (12.5 ± 0.5% and 13.1 ± 0.4%, respectively), on the other, with in-phase motion providing a more effective 'mask' than counter-phase motion. This asymmetry, which was also observed with low-contrast maskers (see above), reveals the operation of a local inhibition that is specific for *direction* of motion. In contrast, the random wavelet configuration produced a more *symmetric* pattern of thresholds, with comparable values at 0 and 30° (9.4 ± 0.5% and 15.4 ± 0.6%) and at 150 and 180° (14.9 ± 0.6% and 13.6 ± 0.5%, respectively). This symmetric pattern hints at the presence of another type of local inhibition, which is specific for spatial *orientation* rather than for direction of motion (Figure 3.5). In the case of parallel wavelets, no clear pattern emerged with full attention, and the results at 0 and 30° (7.2 ± 0.8% and 6.0 ± 0.5%) and at 150 and 180° (6.9 ± 0.6% and 6.6 ± 0.5%, respectively) may well reflect a combination of orientation-selective

(symmetric) and direction-selective (asymmetric) local inhibition. Taken together, these results suggest that there are both *direction*-specific and *orientation*-specific types of local inhibition and that different wavelet configurations are differentially susceptible to each type of inhibition. As both symmetric and asymmetric threshold patterns become more pronounced with poor attention, we will return to these points further below.

With 3 subjects, we conducted additional contrast masking experiments for single and random wavelets (data not shown). The resulting pattern of masking thresholds was consistently asymmetric for low-contrast maskers (2 and 4%) and symmetric for high-contrast maskers (16, 32, and 64%), consistent with direction-selective facilitation and orientation-selective inhibition, respectively. For maskers of intermediate contrast (8%), the results varied and were not consistent across subjects.

3.5.3. Central task performance and attentional strategy

Performance of the central task was nearly constant across conditions. Mean performance and mean standard error of central performance were 81.1, 83.4, and 81.1% in 175, 156, and 205 blocks, when combined with detecting single, parallel, and random wavelets, respectively. This was due to the fact that instructions emphasized the central task and that the analysis excluded blocks of trials where central task performance was compromised (see Methods 3.4.1.4).

Although central task performance was consistently high on average, attention might have swerved to the peripheral array in a subset of trials, e.g., in the trials conducted closest to threshold. To test for this possibility, we analyzed correlations between central and peripheral responses (see Methods 3.4.3) for the critical subset of

trials in each experiment, that is, the trials in which target contrast equaled threshold contrast**.** Among a total 91 contingency analyses (pooled across different masking contrast and direction in 2-4 subjects in 3 stimulus configurations), χ^2 measures of association revealed no significant correlation ($P > 0.05$) between central and peripheral responses in 87 cases. A significant positive correlation was observed in four cases; no significant negative correlation was observed. In short, we find no evidence that attention 'swerves' away from the central task, confirming that dual-task thresholds were indeed established under conditions of poor attention.

3.5.4. Contrast-increment thresholds with poor attention (Figure 3.5 a-c)

Under conditions of poor attention, absolute detection thresholds (masker contrast 0%, relative direction 0°) were 7.6 ± 0.2%, 2.7 ± 0.2%, and 4.2 ± 0.2% (12, 46, and 13% higher than under conditions of full attention for 2, 2, and 4 subjects), significant for single and parallel (F = 11.03, 10.44), and almost significant ($P < 0.07$, F = 3.73) for random wavelets. For masker contrasts above zero (1, 2, 4, 8, 16, and 32%), the effect of attention depended strongly on array configuration: for single wavelets, poorly attended thresholds averaged 22% above fully attended thresholds whereas, for multiple wavelets, the average difference was 109 and 94% for parallel and random, reaching 300% in the most extreme case (i.e., parallel wavelets, 16% masker contrast). Three-way ANOVA (subject x mask contrast x attention) revealed a significant main effect of contrast (F = 127.59, 82.37, and 195.91) and attention (F = 15.09, 82.56, 111.82), as well as a significant interaction between contrast and attention (F = 4.48, 4.28, and 11.24) for single, parallel, and random stimuli. This confirms the significant effect of attention

overall; especially strong effects were seen at high contrast revealed by significant interaction.

Thus, thresholds measured with multiple wavelets more than doubled under conditions of poor attention, overturning the advantage (i.e., lower thresholds) enjoyed by multiple wavelets under conditions of full attention. As lateral masking is one of the factors that disadvantage multiple wavelets compared to single wavelets, these observations suggest that attention substantially alters the lateral interactions that mediate masking (see Discussion 3.6.4 and 3.6.7). For all wavelet configurations, the difference between poorly and fully attended thresholds increased disproportionately with masker contrast (significant interaction), so that the largest proportional difference was always obtained with higher masker contrasts.

3.5.5. Contrast-masking thresholds with poor attention (Figure 3.5 d-i)

The combined effect of poor attention and maskers of low contrast (4% for single and random wavelets, 1% for parallel wavelets) is depicted in Figure 3.5 d-f (open symbols). In general, thresholds followed the asymmetric same pattern as with full attention, increasing from parallel to random to single wavelets and exhibiting facilitation when masker wavelets moved with target wavelets (relative direction 0 and 30°), but not when maskers moved against target wavelets (relative direction 90, 150, and 180°), consistent with direction-specific facilitation. Compared to full attention, thresholds tended to be slightly higher (12% for single wavelets, 12% for parallel wavelets, and 9% for random wavelets). However, withdrawing attention raised thresholds significantly only in 3 out of 15 experiments (5 directions and 3 configurations). In one case, it decreased thresholds,

which suggests weak and inconsistent effects of attention. Three-way ANOVA (subject x mask direction x attention) revealed significant effect of mask direction (F = 45.15, 10.65, 105.07) but attention only in single (F =8.93). Interaction between attention and direction was significant only in random stimulus (F = 3.92). This confirms strong threshold dependency on local mask direction with rather weak effect of attention at low mask contrast.

Taken together with the findings on contrast-increment thresholds and contrast-masking at higher thresholds (see below), this demonstrates that attention is of little consequence as long as the interaction between masker and target wavelets remains facilitatory, that is, as long as masker contrast remains in the 'dipper' regime.

Poor attention had a rather more dramatic effect when combined with maskers of high contrast (32% for single and random wavelets and 16% for parallel wavelets, Figure 3.5 g-i, open symbols). Compared to full attention, poorly attended thresholds were 33% higher with single wavelets, 216% higher with parallel wavelets, and 82% higher with random wavelets.

Three-way ANOVA (subject x mask direction x attention) revealed significant effect of mask direction (F = 39.28, 6.97, 5.73) and attention (F = 69.28, 123.52, 80.56), but the two factors (attention and mask direction) interacted significantly only with *single* and *parallel wavelets* (F = 10.54, 4.31), not *random wavelets* (F = 1.22). This confirms that thresholds with all wavelet configurations depend on both mask direction and attention. Note, however, that for *random wavelets* the attention effect was the same at all mask directions.

The comparatively large effect of attention seems to reflect a combination of several factors. Firstly, poor attention impaired thresholds for multiple wavelets far more than those for single wavelets, as had been true already for contrast-increment thresholds (see above), presumably because poor attention exacerbates lateral masking. Due to this, a substantial gap separated multiple wavelet thresholds with full and poor attention at all relative directions (0, 30, 90, 150, 180°). Secondly, for *single* and *parallel* wavelets, we obtained an *asymmetric* pattern of masking thresholds, so that poorly attended thresholds were higher when maskers moved with (0 and 30°) rather than against (150 and 180°) target wavelets. This pattern, which was noted already with full attention (see above) but proved even stronger with poor attention (as is revealed by the significant interaction between attention and direction), suggests a direction-specific local inhibition by masker wavelets. Thirdly, for *random* wavelets, the pattern of masking thresholds was symmetric (as is shown by the lack of interaction between attention and direction), so that poorly attended thresholds were elevated to the same extent when maskers moved either with (0 and 30°) or against (150 and 180°) target wavelets, compared to when maskers moved orthogonally to target wavelets (90°). We take this pattern to be evidence for another type of local inhibition, namely, an inhibition that is specific to static *orientation* rather than motion *direction*. Both types of inhibition (*direction*- and *orientation*-specific) appeared to also modulate thresholds for parallel wavelets. Lastly, maskers with axial motion (0 and 180°) tended to elevate thresholds less than maskers with slightly off-axial motion (30 and 150°). This general rule was violated substantially only in the case of parallel wavelets, where one would expect strong grouping effects that are absent in

the other two cases. It is not clear why there should be this difference between axial and off-axial masking.

Our results corroborate one incremental effect of attention, specifically, poor attention exacerbating direction-selective inhibition (producing an asymmetric pattern of incremental threshold elevation). They do not support any incremental attention effect on orientation-selective inhibition (as there seems to be no symmetric pattern of incremental threshold elevation). In sum, our results show an attentional modulation of lateral masking and local direction-selective inhibition, but not of local facilitation or local orientation-selective inhibition.

3.6. Discussion

3.6.1. Purpose of experiment

Our aim was to compare the representation of visual motion under conditions of full and poor attention. We distinguished the representations of component and of pattern motion, which appear to involve different levels of visual cortex. Specifically, we sought to maximally stimulate component-sensitive mechanisms while simultaneously minimizing stimulation of pattern-sensitive mechanisms by using random (or incoherent) patterns of motion (Snowden et al., 1991; Britten et al., 1993; Qian and Andersen, 1994; Heeger et al., 1999; Rees et al., 2000). To probe interactions between component-sensitive mechanisms, we used the psychophysical paradigm of contrast masking (Legge and Foley, 1980; Wilson, 1980; Foley, 1994; Lee et al., 1999b; Itti et al., 2000). The observer's attention was controlled with the same dual-task technique that we have used previously for studying spatial vision (Lee et al., 1999b).

3.6.2. Simoncelli and Heeger model

A widely accepted model of the representation of visual motion is the two-stage model by Simoncelli and Heeger (Simoncelli and Heeger, 1998). The first stage consists of visual filters selective for component motions, similar to direction-selective neurons in area V1 or to filters that measure energy at a particular spatio-temporal frequency (Adelson and Bergen, 1985). The sensitivity of each first-stage filter can be represented by a pair of spheres, symmetrically arranged about the origin, in the Fourier domain (Figure 3.2). The second stage is composed of filters selective for pattern motion, that is, for a moving pattern with a particular overall direction and velocity. The sensitivity of second-stage filters can be represented by a zero-crossing plane in the Fourier domain and resembles that of pattern selective cells in area MT. At both stages, the output of each filter is half-rectified and normalized (divided) by the linear sum of a pool of filters (Heeger, 1993). Figure 3.6 illustrates this divisive normalization for the first-stage filters (i.e., the stage that matter for contrast masking thresholds). Among other advantages, divisive normalization renders tuning properties largely independent of stimulus contrasts. In addition to divisive normalization, second-stage filters are subject to a subtractive inhibition termed 'motion opponency,' (Snowden et al., 1991; Qian and Andersen, 1994; Heeger et al., 1999).

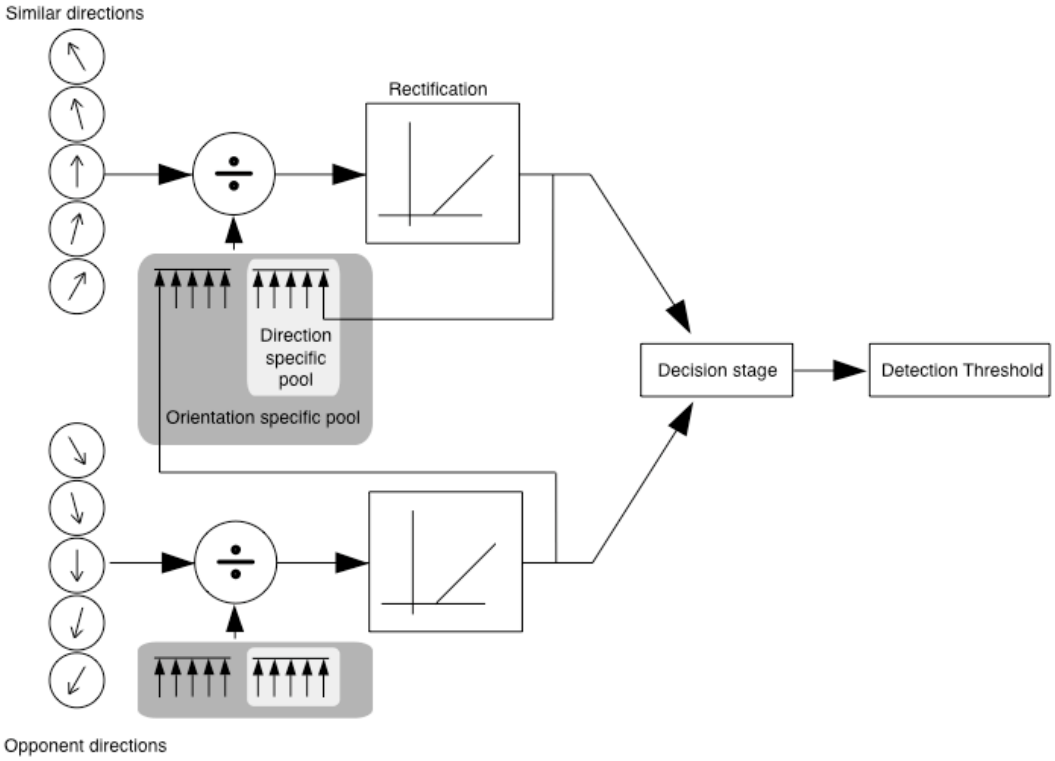

Figure 3.6 Divisive normalization by direction- and orientation-specific filters

Contrast masking thresholds reflect the contrast response of first-stage filters, which are selective for the direction of component motion. They can be modeled by comparing the activity of filters representing directions similar to, and opposite to, the probe stimulus direction. Each filter response results from divisive normalization (shown by a symbol for division) of the linear response, followed by rectification. In general, two filter pools are expected to contribute to the normalization: a "direction-specific pool" with filters of

similar direction and an “orientation-specific pool” with filters of both similar and opposite directions.

3.6.3. Absolute detection threshold

Different stimulus configurations show a consistent rank order of detection thresholds. Some difference is due to the summation of signals (i.e., single vs. 23 wavelets). Theoretical consideration of signal-to-noise ratio based on ideal signal summation predicts reduction of the detection thresholds by a factor of 4.8. However, previous empirical studies consistently show a less-than-ideal degree of summation (Quick, 1974; Bonneh and Sagi, 1998; Meese and Williams, 2000; Tyler and Chen, 2000). The reason for this is unclear, in particular as less-than-ideal summation is observed even with well-separated signals (Bonneh and Sagi, 1998; Meese and Williams, 2000).

Table 3.1 compares the measured and predicted thresholds for 23 wavelets, with the latter calculated from the following formula, $\frac{1}{C_{MULTI}} = \left(\sum^{23} \left(\frac{1}{C_{SINGLE}} \right)^{q} \right)^{1/q}$. Here, C_{SINGLE} is the measured detection threshold for single wavelets and C_{MULTI} is the calculated threshold for multiple wavelets. As in other studies, a value of q=4 produces good agreement between predicted thresholds and the average measured thresholds. However, we also observe an additional effect of wavelet configuration that is not predicted.

Detection thresholds for multiple random wavelets are poorer than predicted, perhaps due to subtractive inhibition at the pattern motion stage (Schrater et al., 2000). Thresholds for multiple parallel wavelets are better than predicted, perhaps reflecting collinear facilitation at the level of component motion.

	Measured single	Average multiple	Predicted multiple	Measured multiple random	Measured multiple parallel
Poor attention	7.6 ± 0.2	3.4	3.5 ± 0.1	4.2 ± 0.2	2.6 ± 0.2
Full attention	6.7 ± 0.2	2.8	3.1 ± 0.1	3.7 ± 0.1	1.9 ± 0.2
Cueing	5.5 ± 0.2	2.6	2.5 ± 0.1		

Table 3.1 Probability summation

Comparison of absolute detection thresholds (in % contrast and s.e.m.) for patterns with one and with multiple (23) wavelets. The fourth-root power-law was used to predict the effect of probability summation. Thresholds for cued detection were established only for single and random wavelets. The predictions underestimate the observed thresholds for multiple random wavelets and overestimate those for multiple parallel wavelets.

3.6.4. Inhibitory lateral interaction (relative mask direction 90°)

Lateral interactions between wavelets are expected to become inhibitory at high contrast (Polat and Sagi, 1993). Accordingly, any such interactions should be evident in the comparison of increment thresholds between single and multiple wavelets. The lowest increment thresholds (at high mask contrast) are obtained with orthogonal target and mask wavelets (relative direction 90°), where *local* interactions are expected to be minimal. Accordingly, situations with orthogonal target and mask wavelets are the most suited for observing threshold elevation due to *lateral* interactions in the presence of multiple wavelets (Table 3.2). With full attention, the presence of a single masking wavelet elevates detection thresholds by 126% (presumably due to local inhibition), whereas the presence of multiple masking wavelets elevates thresholds by 250% (random wavelets) and 180% (parallel wavelets). Note that the latter value was obtained for a lower mask contrast. This additional increase presumably reflects lateral inhibitory interactions.

	Single (32% mask)	Random (32% mask)	Parallel (16% mask)
Detection (0% mask)	6.7 ± 0.2	3.7 ± 0.1	1.9 ± 0.2
Full attention	8.5 ± 0.6 (126%)	9.1 ± 0.4 (250%)	3.5 ± 0.4 (180%)
Poor attention	9.8 ± 0.4 (145%)	17.5 ± 1.3 (470%)	14.6 ± 1.4 (770%)

Table 3.2 Threshold elevation due to lateral interactions

Comparison of detection and increment thresholds for single and multiple wavelets, with orthogonal target and mask wavelets. Threshold value (% contrast and s.e.m.) and threshold elevation (percentage points). Note that values for random and parallel wavelets were obtained with different mask contrasts.

3.6.5. Facilitation and inhibition (relative mask direction 0°)

The "dipper" regime in the contrast increment threshold function reveals facilitatory effects of mask wavelets that move in parallel to probe wavelets (relative direction 0°). (That the "dipper" did not result from positional uncertainty was shown by a control experiment with cueing.) We find facilitatory mask effects in the contrast regime of 2-4%. For larger mask contrasts, the mask wavelets become inhibitory and the contrast increment threshold function saturates. This inhibition may reflect local divisive normalization and/or lateral inhibition between nearby wavelets.

We observe a comparatively shallow "dipper" with multiple parallel wavelets, echoing previous results with sinusoidal gratings (Bowne, 1990; Lu and Sperling, 1995, 1996). For single and multiple random wavelets, however, we obtain a prominent "dipper" regime (Figure 3.5 a-c). This strongly suggests that component motion mechanisms exhibit a markedly sigmoidal contrast dependency. At the level of pattern motion mechanism, this sigmoidal contrast response may be lost due to summation. Tootell and colleagues (Tootell et al., 1995) compared contrast response function with gratings using fMRI in V1 and MT, finding an early saturation of BOLD signal in MT but not in V1.

3.6.6. Direction dependence of contrast masking (relative mask direction 30°, 150°)

We observe facilitation at low contrast (Figure 3.5 d-f) when mask wavelets move in directions similar to target wavelets. This is consistent with earlier results that sub-threshold summation occurs between similar, but not opposite, directions of motion (Levinson and Sekuler, 1975; Stromeyer et al., 1984; Wilson, 1985). In contrast, the observed inhibition at high contrast (Figure 3.5 g-i) depends on the configuration of wavelets (random or parallel). For multiple random wavelets, the threshold increment reveals an orientation-specific inhibition (symmetric 'M' shape in Figure 3.5 i). This selectivity matches the organization of direction-columns in area V1 organization, where adjacent columns represent opposite directions of motion (Shmuel and Grinvald, 1996; Weliky et al., 1996; Ohki et al., 2005). Accordingly, indiscriminate connectivity between nearby columns would be expected to result in a normalizing pool that includes columns selective for both similar and opposite directions. In other words, a local inhibitory pool is expected to be *orientation*-selective without being *direction*-selective.

Using low-pass filtered one-dimensional noise motion as a masking stimulus, Anderson and Burr (Anderson and Burr, 1985) found weak inhibitory effects with the masker moving in the opposite direction. The magnitude of inhibitory effects in their experiments resembles our asymmetrical inhibition with single and parallel wavelets at high contrast masking. A common property of these stimuli is that they activate MT efficiently in mask-only trials, but the target moving in the opponent direction suppresses MT activity. The reduction of MT activity can be utilized for target detection. Thus, the

reduction of inhibition at the opposite direction indicates motion opponency at the level of MT.

An unexpected feature of our results is the reduced inhibition for maskers of identical (0°) and exactly opponent (180°) direction ('M' symmetry rather than inverse 'U' symmetry in Figure 3.5 i). The reason is unclear. One possibility is that the inhibitory pool "spares the center (0 and 180°)" and is shaped like an annulus rather than like a Gaussian blob. There are some reports implicating stronger divisive inhibition from relative directions of 30-45° than of 0° (Foley, 1994; Zenger and Sagi, 1996) (however, see also (Anderson and Burr, 1985)).

3.6.7. Attentional effects

We found no evidence that attention alters the facilitating effect of low-contrast masks. Neither the depth of the "dipper" function (Figure 3.5 a-c) nor its dependence on relative mask direction (Figure 3.5 d-f) was affected by attention. Apparently, the initial contrast response (i.e., response at low contrast) of component motion mechanisms is not affected by the state of attention.

This finding differs markedly from our earlier results for static visual patterns. In that case, mask facilitation at low contrast was significantly enhanced by attention (Lee et al., 1999b). Specifically, attention enhanced both the effective gain and sharpened the effective tuning of visual filters selective for static patterns. The effects observed for static patterns were quantitatively consistent with the possibility that visual attention intensifies a competitive interaction among overlapping visual filters.

Taken together, our findings for static and dynamic visual patterns imply that any change in competitive interactions be restricted to static filters. The presumed neural correlates of static and dynamic visual filters are, respectively, orientation and direction columns in primary visual cortex. Our observations imply, therefore, that attentional feedback should differentiate between the two classes of filters.

Besides this negative finding, our results demonstrate two clear effects of attention. The first effect is a *uniform* reduction of inhibition by high-contrast masks (Figure 3.5 hi). This effect is observed for all relative mask directions, but only for multiple wavelets (parallel or random). Accordingly, the interaction modulated by attention appears to be a lateral inhibition by nearby wavelets of high contrast. This

would be consistent with previous reports that attention modulates lateral interactions between high-contrast stimuli (Zenger et al., 2000; Freeman et al., 2001, 2004). Inhibitory interactions (i.e., lateral masking) are particularly affected, as attention may decrease their effect by a factor of four or more (Zenger et al., 2000).

The second effect of attention is a reduction of *direction*-selective inhibition (Figure 3.5 gh). This inhibition in question is direction-selective as it elevates thresholds for relative mask directions of 0 and 30°, but not 150 and 180°. It is observed only for single wavelets and parallel multiple wavelets, not for random multiple wavelets. This restriction to stimuli that activate pattern motion mechanisms suggests that the inhibition in question operates at that level. It might reflect either motion opponency or divisive normalization or both. In either case, this direction-selective inhibition in the representation of pattern motion is reduced significantly by attention.

In summary, we find no evidence for an attentional modulation of the representation of component motion. We do find that attention reduces a direction-selective inhibition in the representation of pattern motion and, consistent with previous studies, that attention reduces lateral masking by high-contrast stimuli (both moving and stationary).

3.7. Acknowledgement

I started this project under the guidance of Jochen (Achim) Braun in April 2000. Achim created the display for the basic multiple random wavelets. After he left for England in June 2000, Geraint Rees helped me do the programming. Achim, Geraint, and I used the motion stimuli to look at the neuronal activity with fMRI in University College London from July to September 2001, which is yet to be studied more rigorously and extensively. The fMRI work is likely to be published within 10 years from now and was not included in this thesis.

4. Role of awareness in trace and delay aversive conditioning

4.1. Overview

For the study of consciousness, it would be desirable to have a battery of tests for the presence of consciousness, which can be applied to any non-speaking humans, including infants and neurological patients, as well as animals, from mammals to even invertebrates[12]. Such tests would be useful in objectively assessing the presence of consciousness in these organisms, in much the same way as a *Turing-test* has been used to assess artificial intelligence (Koch, 2004). In this chapter, we explore the possibility that a simple learning paradigm may serve as a Turing-test for consciousness.

A hallmark study that alluded to such a possibility is reported by Clarke and Squire in 1998 (Clark and Squire, 1998). Successful performance in simple eyeblink conditioning relies on attention and awareness, depending on the temporal relationships between the two stimuli; in *delay* conditioning, a previously neutral tone (CS) is paired with a temporally *overlapping* airpuff stimulus to the eyelid (US), while in *trace* conditioning, a temporal gap is inserted between CS and US, so that a memory *trace* has to be stored in working memory to bridge the gap between the offset of the CS and the onset of the US. Clark and Squire report that trace eyeblink conditioning depends on contingency awareness of the relationship of CS and US, while delay eyeblink conditioning occurs independently of the awareness. Further, while trace conditioning is

[12] Even if babies, patients during epileptic seizures, or animals show apparently complicated behaviors, these behaviors by themselves do not guarantee the presence of consciousness, as there are many examples of so-called "zombie behaviors" that are highly elaborated and operate without giving rise to consciousness.

disrupted when attention is distracted by a concurrent task, delay conditioning is resilient to distraction. Previously, we extended their findings on attention to aversive conditioning, using mild electrical shocks as US and skin-conductance response as a measure of conditioned responses (Carter et al., 2003). Here, we show that trace conditioning develops at that point in time at which subjects become aware of the relationship between CS and US, while such a temporal synchrony is not found in delay conditioning. The observed difference in dependency on consciousness between the two protocols is remarkable; the inputs are very similar (the only difference being the presence or absence of a 3-sec temporal gap), and the overall performances are comparable (See Section 4.5).

Our results suggest that we might be able to use the trace/delay-conditioning paradigm as a Turing-test for the presence of awareness and/or attention in non-verbal organisms. During trace but not delay conditioning, candidate neural systems for the NCC should be actively involved in the maintenance of the memory trace. These neural systems are likely to be disturbed when attention is distracted, which would eliminate trace but not delay conditioning. The above inference would be extremely powerful when it is applied to flies and mice, whose genomes can be rather easily manipulated using ever more powerful molecular tools. In addition to the exquisite spatial resolution, molecular markers allow us to target a specific class of neurons, which may be distributed across space. Probing a particular population of neurons required to carry out a Turing-test for consciousness might turn out to be a critical step in the study of the NCC, because only a subset of neurons, that are distinguished by peculiar molecular

markers yet distributed across different areas, might be critical for producing consciousness.

4.2. Summary

Previous studies showed that human associative conditioning can depend on subjects becoming aware of the exact, temporal relationship between neutral and aversive stimuli. Such awareness seems necessary for eyeblink trace conditioning (where the airpuff follows the tone after some temporal gap) but not for eyeblink delay conditioning (where the airpuff overlaps with the tone). We investigated if conditioning and awareness develop synchronously during trace and delay aversive conditioning. Trace and delay conditioning were performed in a single session to reduce inter-subject variability. The degree of conditioning was assessed by skin-conductance response, a measure of autonomic arousal. Each subject expressed their shock expectancy throughout the experiment, allowing us to infer the onset of awareness of the relationship between tones and shock. We found that the level of differential skin-conductance response did not differ between trace and delay conditioning when averaged across all the trials. However, when the response was aligned at the onset of awareness, trace, but not delay, conditioning showed significant discrimination between paired and unpaired stimuli. Our result is consistent with the idea that trace conditioning is more tightly related to awareness than delay conditioning and that different mechanisms mediate the two types of learning.

4.3. Introduction

In classical conditioning, subjects learn to associate a previously meaningless conditioned stimulus (CS) with a meaningful unconditioned stimulus (US). Whether they need to be aware of the CS-US contingency in order to be conditioned remains controversial (Dawson and Furedy, 1976; Dawson and Schell, 1985; Clark and Squire, 1998; Clark et al., 2002; Lovibond and Shanks, 2002). A tight relationship between conditioning and contingency awareness is supported by aversive autonomic conditioning studies (Dawson and Furedy, 1976; Dawson and Schell, 1985; Lovibond and Shanks, 2002). Successfully conditioned subjects verbalize the correct CS-US contingency, but knowledge of the contingency itself is insufficient for conditioning. When subjective shock expectancy is concurrently assessed on a trial-by-trial basis, conditioning occurs when subjects realize the correct CS-US contingency (Dawson and Biferno, 1973; Biferno and Dawson, 1977). No evidence of conditioning is seen before they express this awareness. These results are explained within a 'necessary-gate' framework, in which successful conditioning requires awareness and conditioning can occur only after awareness has developed; however, awareness by itself is not sufficient for conditioning to occur (Dawson and Furedy, 1976; Dawson and Schell, 1985; Lovibond and Shanks, 2002).

On the other hand, classical conditioning is regarded as a non-declarative learning process, which does not require conscious association and can occur automatically (Milner et al., 1998). This notion has been refined by Clark and Squire, who claim that the link between conditioning and awareness depends on the temporal relationship between CS and US; awareness is necessary when CS and US are separated in time (trace conditioning), while it is not required when CS and US overlap (delay conditioning)

(Clark and Squire, 1998; Clark et al., 2002). In an eyeblink conditioning paradigm, where an airpuff to the eyelid serves as US, delay conditioning occurs even though subjects are unaware of the contingency, either due to a concurrent distracting task or due to hippocampal amnesia (Gabrieli et al., 1995; Clark and Squire, 1998). Attentional distraction interferes, however, with the acquisition of trace eyeblink conditioning (Clark and Squire, 1998). We previously showed that trace aversive conditioning was easily disrupted by a concurrent working-memory task while delay aversive conditioning was more resilient (Carter et al., 2003).

We here revisit the phenomenon of synchronous development of conditioning and awareness by monitoring shock expectancy throughout a learning session (Dawson and Biferno, 1973; Biferno and Dawson, 1977) and contrast differences between delay and trace protocols. To reduce inter-subject variability, we directly compare trace and delay by intermixing three CSs in a single session (Knight et al., 2004). Based upon the claim that trace conditioning is more dependent on consciousness than delay conditioning (Clark and Squire, 1998; Clark et al., 2002; Carter et al., 2003), we predict that trace conditioning develops with the onset of contingency awareness, while delay conditioning does not.

4.4. Methods

4.4.1. Subjects and equipment

To record skin conductance, we attached electrodes to the subjects' index and middle fingers at the second phalanx of the non-dominant hand. Skin conductance was sampled at 50Hz. Electrical shocks were delivered through electrodes attached to the ring and little fingers at the second phalanx of the dominant hand. The delivery of shocks and recording of skin conductance was performed by a system from Contact Precision Instruments controlled by PSYLAB software. We used Silver/Silver Chloride electrodes filled with Med Associates paste TD-246.

Subjects expressed their immediate shock expectancy by gaze position, which was recorded at a sampling rate of 120 Hz with an infrared video-based eyetracker (ISCAN, Cambridge, MA). A head and chin rest minimized head and body motion. The display was placed 85 cm from the subjects. Presentation® software (Version 0.70) controlled the experimental procedure, played recorded instructions, and presented visual stimuli for the distracter task.

We used three different auditory, 2-sec-long CSs to conduct trace and delay conditioning in a single session for each subject: a 2kHz tone (101dB), a white noise sound (83dB), and a bubbling water sound (98.5dB). The volume of each was determined in a pilot study with 20 naïve subjects so that each sound produced approximately the same SCR amplitude. Each auditory CS was assigned to either CS+D (delay), CS+T (trace), or CS-, and was balanced across subjects. The CS+D co-

terminated with a 0.25-sec electrical shock (US), the CS+T was followed by the US after a 3-sec trace interval, and the CS- was never accompanied by the US (Figure 4.1 a).

Thirty-three subjects (13 female) employed at Caltech gave consent and were paid for participating in the experiment. The protocol was approved by Caltech's IRB. Subjects' ages ranged from 18 to 48 years (mean 24.2). After signing the consent form and before the experiment, each subject adjusted the intensity of the shock level so that it was 'uncomfortable but not painful.' Three subjects were excluded because they reached the maximal shock amplitude before rating it as uncomfortable. An additional six subjects were excluded because they did not respond to shocks in more than 8 out of 40 trials.

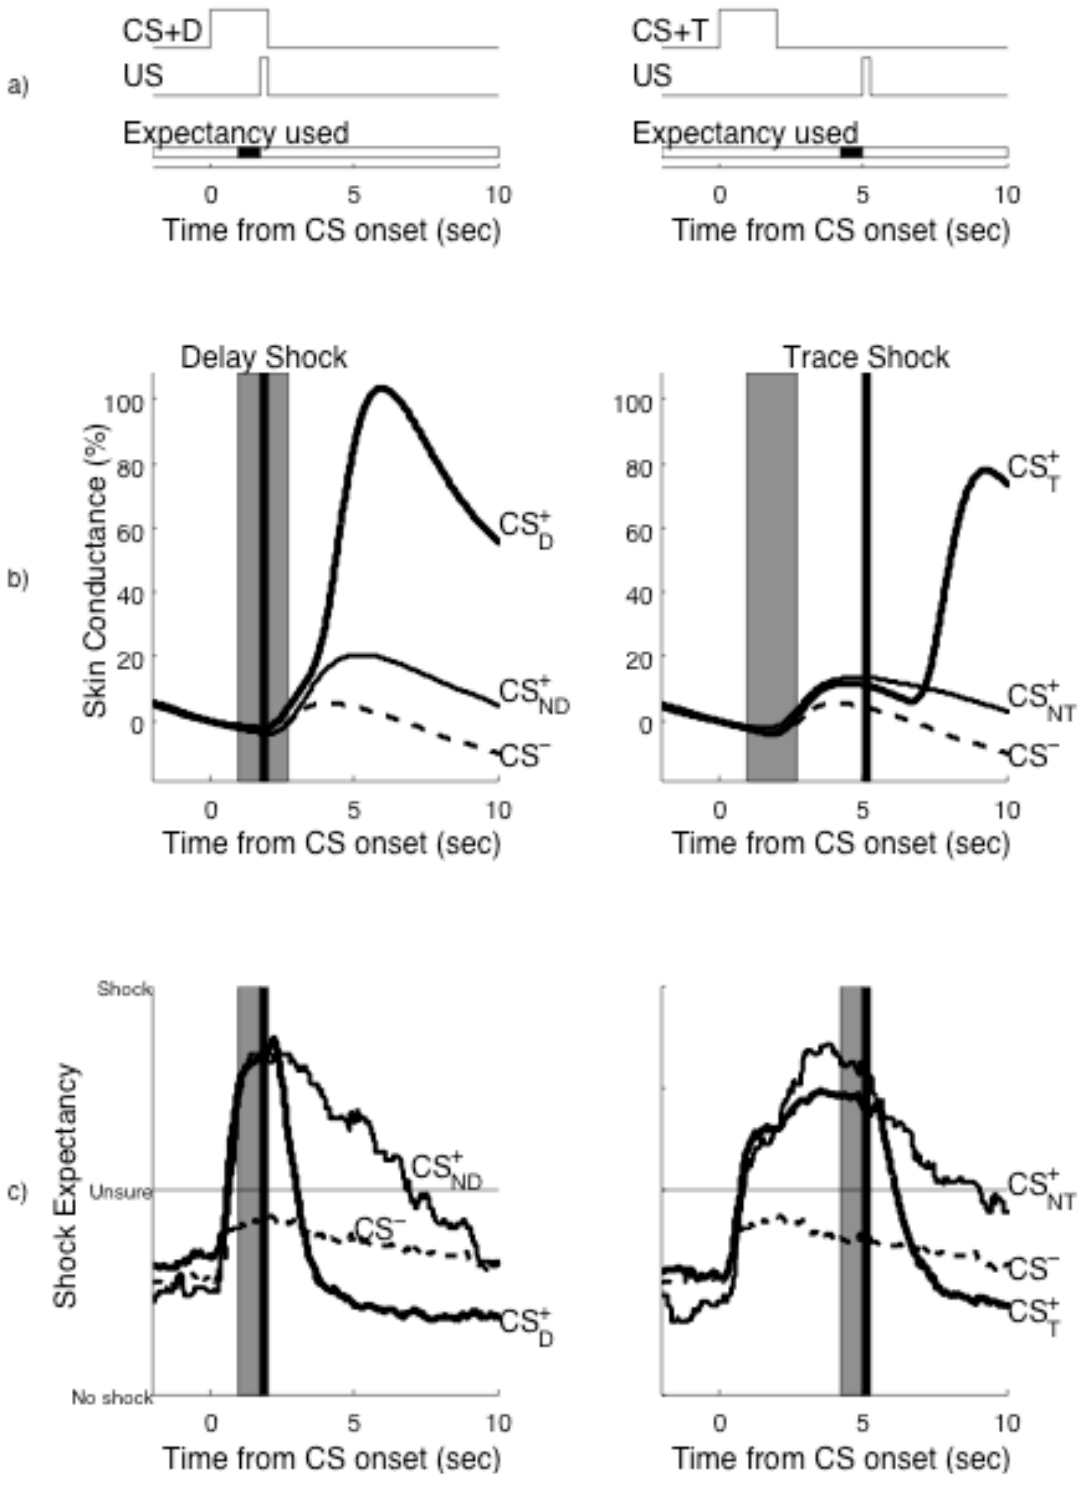
a)
CS+D
US
Expectancy used
CS+T
US
Expectancy used
0
5
10
Time from CS onset (sec)
b)
Delay Shock
Trace Shock
Skin Conductance (%)
100
80
60
40
20
0
CS^+_D
CS^+_{ND}
CS^-
CS^+_T
CS^+_{NT}
c)
Shock
Unsure
No shock
Shock Expectancy

Figure 4.1 Stimuli characteristics and event-related time course of SCR and shock expectancy

Left column is for delay and right column is for trace conditioning. **a**) Delay and trace conditioning were performed in a single session by intermixing three auditory 2-sec-long CSs (CS+D, CS+T, CS-). Shock expectancy expressed by direction of gaze was continuously monitored. The average expectancy during the 0.75 sec before the US onset was used as an expectancy score for each trial (black squares). CS- expectancy from the corresponding period was compared with CS+T or CS+D appropriately. **b**) The average time course of the skin conductance in all acquisition trials for the "no task" group (n = 12, see Section 4.4.3) was first normalized by the mean response to US and then averaged across subjects. Shocks for delay and trace conditioning are represented as vertical lines at 1.75 sec and 5 sec, respectively. We defined a SCR for each trial using a time window from 1 sec to 2.75 sec (gray zone). For trace conditioning, the thick solid line for the 83% reinforced trials (CS+T) diverges at ~7 sec (~2 sec after shocks) from the thin solid line for the non-reinforced trials (CS+NT). Similarly, the line for the reinforced delay trials (CS+D) diverges from the line for the non-reinforced delay trials (CS+ND) around ~3.5 sec. **c**) The average shock expectancy during acquisition. All aware subjects could reliably move their eyes to signal high shock expectancy well before 1 sec after CS onset while performing a visual task. The expectancy score is the average expectancy during 0.75 sec before the shock (gray zone).

4.4.2. Experimental design

The conditioning session consisted of three phases (habituation, acquisition, and extinction), with no explicit cue for the transition between phases (Table 4.1). In habituation and extinction, no shock reinforcements occurred. There were 2, 24, and 8 blocks for habituation, acquisition, and extinction, respectively. One block consisted of three trials whose order was randomized so that the same trial type only occurs, at most, twice in a row. During acquisition, CS+T and CS+D were reinforced with shocks in 83% of the trials (20 reinforced and 4 non-reinforced trials). To facilitate conditioning, the first 3 blocks in acquisition were always reinforced. Inter-trial intervals were chosen from time periods of 10, 12, 14, 16, 18, or 20 sec.

Block	H1			H2			A1			A2		
Trial	1	2	3	4	5	6	7	8	9	10	11	12
CS type	D	-	T	-	T	D	-	T	D	D	T	-

Block	A3			A4			A5			A6		
Trial	13	14	15	16	17	18	19	20	21	22	23	24
CS type	-	D	T	T	-	D	D	NT	-	T	D	-

Block	A7			A8			A9			A10		
Trial	25	26	27	28	29	30	31	32	33	34	35	36
CS type	-	T	ND	-	T	D	-	D	T	D	-	T

Block	A11			A12			A13			A14		
Trial	37	38	39	40	41	42	43	44	45	46	47	48
CS type	T	ND	-	D	-	T	NT	-	D	ND	T	-
OTP							$T\text{-}3^{+}$	$T\text{-}3^{-}$			$T\text{-}2^{+}$	$T\text{-}2^{-}$

Block	A15			A16			A17			A18		
Trial	49	50	51	52	53	54	55	56	57	58	59	60
CS type	T	D	-	T	-	D	D	-	T	-	NT	D
OTP	$T\text{-}1^{+}$		$T\text{-}1^{-}$	$T0^{+}$	$T0^{-}$			$T1^{-}$	$T1^{+}$	$T2^{-}$	$T2^{+}$	

Block	A19			A20			A21			A22		
Trial	61	62	63	64	65	66	67	68	69	70	71	72
CS type	ND	T	-	T	-	D	T	-	D	NT	D	-

Block	A23			A24			E1			E2		
Trial	73	74	75	76	77	78	79	80	81	82	83	84
CS type	T	D	-	-	D	T	-	D	T	D	-	T

Block	E3			E4			E5			E6		
Trial	85	86	87	88	89	90	91	92	93	94	95	96
CS type	-	D	T	D	T	-	T	-	D	D	-	T

Block	E7			E8		
Trial	97	98	99	100	101	102
CS type	-	D	T	D	-	T

Table 4.1 Exemplar trial sequence for depiction of experimental design

This table shows the actual trial sequence used for an exemplar subject, whose shock expectancy data is shown in Figure 4.2 (See also Section 4.5.1). The conditioning session consisted of three phases (2 blocks of habituation, 24 blocks of acquisition, and 8 blocks of extinction). There was no explicit cue for the transition between phases and blocks. Three types of CS (CS+T, CS+D, and CS-) were randomized within a block so

that the same trial type only occurred, at most, twice in a row across blocks (i.e., here, both trial 9 and 10 were CS+D). During acquisition, 4 non-reinforced trials, CS+NT and CS+ND, occurred (i.e., CS+NT for block 5 in acquisition) with a restriction that two non-reinforced trials did not occur within a single block. To facilitate conditioning, the first 3 blocks in acquisition (A1-A3) were always reinforced. Inter-trial intervals were on average 15 sec. This particular subject became aware of the CS-US contingency for trace conditioning at the 52^{nd} trial according to the generalized onset of awareness (See Section 4.5.1 and Figure 4.2). The three trials for CS+ (in the table for the onset trial pair, OTP, depicted as $T0^{+}$, $T1^{+}$, and $T2^{+}$) and the other three trials for CS- ($T0^{-}$, $T1^{-}$, and $T2^{-}$) are paired to form three trial pairs to calculate the difference of SCR between CS+ and CS- (ΔSCR) (See Section 4.5.1 and 4.5.2, Figure 4.3). For the preceding 6 trials, three trial pairs ($T\text{-}3^{\pm}$, $T\text{-}2^{\pm}$, and $T\text{-}1^{\pm}$) are likewise formed to calculate ΔSCR.

4.4.3. Distraction task and expectancy monitoring task

As a distracter, we used a visual discrimination task known to function in the near absence of focal visual attention (Li et al., 2002). We assigned subjects blindly to the "distraction task" group (n = 12, 4 female) or to the "no task" group (n = 12, 4 female).

During practice of the discrimination task, we presented a central white fixation cross (0.4 x 0.4°) that was visible throughout the session and a natural scene picture (4.6 x 7.9°). Each picture was shown for 0.1 sec and replaced with a gray field for 1.9 sec until the next picture was presented. Subjects pressed a mouse button with the dominant hand (the one with the shock electrodes) whenever the natural scene contained a means of transportation (e.g., cars, trucks, trains, boats, aircraft, and hot-air balloons). The sixty practice stimuli were different from those used in the main experiment (437 targets, 911 distracters). During conditioning, pictures were presented once every 2 sec. When the CS is presented, pictures were shown 0.3 sec after the CS onset, so that pictures do not interfere with the processing of CS.

During conditioning, 3 sets of identical stimuli, consisting of an image and fixation cross, were projected onto a horizontal line with a separation of 8.3°. Subjects expressed their shock expectancy by fixating one of the three crosses. The recorded instructions indicated that they should either: 1. look at the left fixation cross when they 'think it is likely that there will be a shock within the next few seconds,' 2. look at the center point when they 'do not know whether or not there will be a shock within the next few seconds,' 3. fixate the right cross when they 'think it is unlikely that there will be a shock within the next few seconds.' Whether the left or right fixation cross indicated

high shock expectancy was randomized across subjects. Subjects practiced the expectancy task for 40 sec, where text on the screen indicated which cross should be fixated.

Finally, subjects practiced reporting their shock expectancy by moving their eyes while performing the distraction task. The recorded instructions emphasized the primacy of reporting shock expectancy over performance on the distraction task. None of the subjects had problems performing both tasks at the same time (mean reaction time = 451 ± 56 msec (s.d.), 97.1% correct detection, 3.5% false alarm for the distraction group (n = 12) in the main experiment). After the final instruction, subjects were assigned to the distraction or the “no task” group. We told the subjects in the "no task" group that they would only perform the expectancy monitoring task.

4.4.4. Data Analysis

Figure 4.1 a shows the time course of the CS and US for each protocol, and Figure 4.1 b shows the average time courses of skin conductance. Skin conductance was first normalized by the mean SCR amplitude of the responses to shocks within the 1-10 sec interval after CS onset for that subject, and then was averaged across subjects (Lykken, 1972; Ben-Shakhar, 1985). An event-related time course was obtained by aligning skin conductance to 0% at CS onset (time = 0 sec). We defined the skin conductance response (SCR) for each trial as the increase in normalized conductance from the trough to the peak in the 1-2.75-sec interval after CS onset (gray zone in Figure 4.1 b). The tail of the time window, 2.75 sec, occurs 1 sec after the onset of the 0.25-sec long shocks for CS+D trials, which is well before the average SCR onset of 1.8 sec (Venables and Christie,

1980). We confirmed the SCR latency in each trial for each subject and found no subjects with a SCR latency faster than 1 sec. This allows us to include the SCR from reinforced trials as well as those from non-reinforced CS+ presentations. SCR was corrected by a square root transformation for normality.

Shock expectancy was obtained by converting the gaze position into three discrete expectancy scores, –1 (no shock expected), 0 (unsure), and +1 (shock expected). The expectancy score for each trial was the average expectancy during 0.75 sec before the shock (black squares in Figure 4.1 a and gray zone in Figure 4.1 c), that is, the average from 1 to 1.75 sec in CS+D (delay) trials and from 4.25 to 5 sec in CS+T (trace) trials. Shock expectancy in CS- trials was measured within the 1 to 1.75 sec interval to compare with CS+D and from 4.25 to 5 sec to compare with CS+T.

4.5. Results

4.5.1. Generalized definition for the onset of awareness

Dawson and Biferno (Dawson and Biferno, 1973; Biferno and Dawson, 1977) introduced two ways to define the onset of awareness. When subjects expect shocks for CS+ and no shocks for CS-, they have 'valence' awareness. When their shock expectancy for CS+ is higher than for CS-, they are said to have 'discrimination' awareness. The original definition required correct expectancy to be 'sustained' until the end of acquisition.

The requirement for 'sustained' expression may be too strict to allow for subjects who were somewhat aware of the contingency. Since CS-US contingency awareness for the entire experiment is likely to be a graded phenomenon rather than black-and-white, vague awareness is also a valid measure (Merikle and Reingold, 1992). Subjects' confidence in their judgment concerning CS-US contingency can wax and wane over the course of the experiment. This is particularly true if partial reinforcement was used (as in our study) since not all CS+'s are followed by a shock. In the post-experimental interview, subjects volunteered several possible reasons why they changed their expectancy after they became aware of the correct contingency. Some commented that once they realized the correct contingency, they further noticed non-reinforced trials and tried to predict the absence of a shock in a CS+ trial. Others noted that they thought their expectancy responses might trigger shocks via feedback, and that they tried to confirm this by expressing the opposite expectancy. Furthermore, subjects may forget to move their eyes on some trials. Therefore, we need measures of awareness that capture a graded level of awareness.

We generalize the original definition of the onset of awareness by introducing two variables: 'consecutiveness,' N, defines how many consecutive trials are necessary for the onset of awareness to be detected, while 'violations tolerated,' V, quantifies how often, i.e., for how many trials, the expectancy can be violated after the onset of awareness. The original definition by Dawson & Biferno is equivalent to our generalized definition with N=6 and V=0; in this case, the onset trial is defined as the i^{th} trial; from then on, the correct shock expectancy is expressed in the following N=6 consecutive trials, equivalent to three CS+/CS- pairs. The three (= N/2) trials for CS+ and the other three trials for CS- are paired to form three trial pairs (trial pair 0, the onset pair; +1; and +2) to calculate the difference of SCR between CS+ and CS- (ΔSCR) (Table 4.1). For the preceding N = 6 trials, three trial pairs (-3, -2, and -1) are likewise formed to calculate ΔSCR. After the i^{th} trial, there are no trials (V = 0) where the expression of the correct shock expectancy can be violated.

In the following analysis, we varied N from 4 to 6 and 8 and V from 0, 1, 3, 7, 15, to 31 (The data for $V \geq 31$ were identical). The onset of awareness was detected separately for trace and delay shock expectancy. We rejected the onset of awareness that occurred too early ($i < N$) or too late ($i + N > 78$, the end of the acquisition phase). Though we performed the analyses with both discrimination and valence criteria, we report only the results with the discrimination criterion because significant effects were not obtained using the valence criterion.

Figure 4.2 a and b show the shock expectancy of an exemplar subject as a function of trials, to illustrate how the onset trial is detected. In Figure 4.2 c, we show the number of aware subjects in a color-coded matrix whose columns and rows represent the

consecutiveness (N) and the violations tolerated (V), respectively. Because our main concern is the comparison between trace and delay conditioning in a within-subject manner, we only included subjects when their awareness was detected for both trace and delay protocols. For example, the exemplar subject was considered aware and was included in the analysis with the discrimination criterion only when $V \geq 3$ for any N ('o' in Figure 4.2 c), as he failed to express the correct shock expectancy twice in trace conditioning (in the 76^{th} and 78^{th} trials). The mean number of aware subjects across all N's and V's was 11.28 out of 24 subjects, ranging from 7 to 19. As expected, more aware subjects were detected as the definition became looser, that is, as N decreases or as V increases (towards the upper left corner).

To estimate how many of these aware subjects could have been detected by chance, we performed a Monte Carlo simulation. We randomly created 100,000 sets of a simulated shock expectancy. Each set of simulated data consisted of 78 values, randomly taken from -1, 0, or 1. We treated three sets of 26 values as shock expectancy for CS+T, CS+D, and CS- and submitted each set to the same awareness detection algorithm as the real data. Figure 4.2 d shows the number of aware subjects detected by chance alone, multiplied by a ratio between the participants (n=24) and the simulated trials (n=100,000). The mean was 1.51, ranging from 0.037 to 12.32. The differences between the data and the simulation were well separated everywhere (Figure 4.2 e); the mean difference was 9.77, ranging from 6.68 to 13.14. In other words, our generalized awareness criterion detects more aware subjects in the data than the random data in any tested combination of N and V.

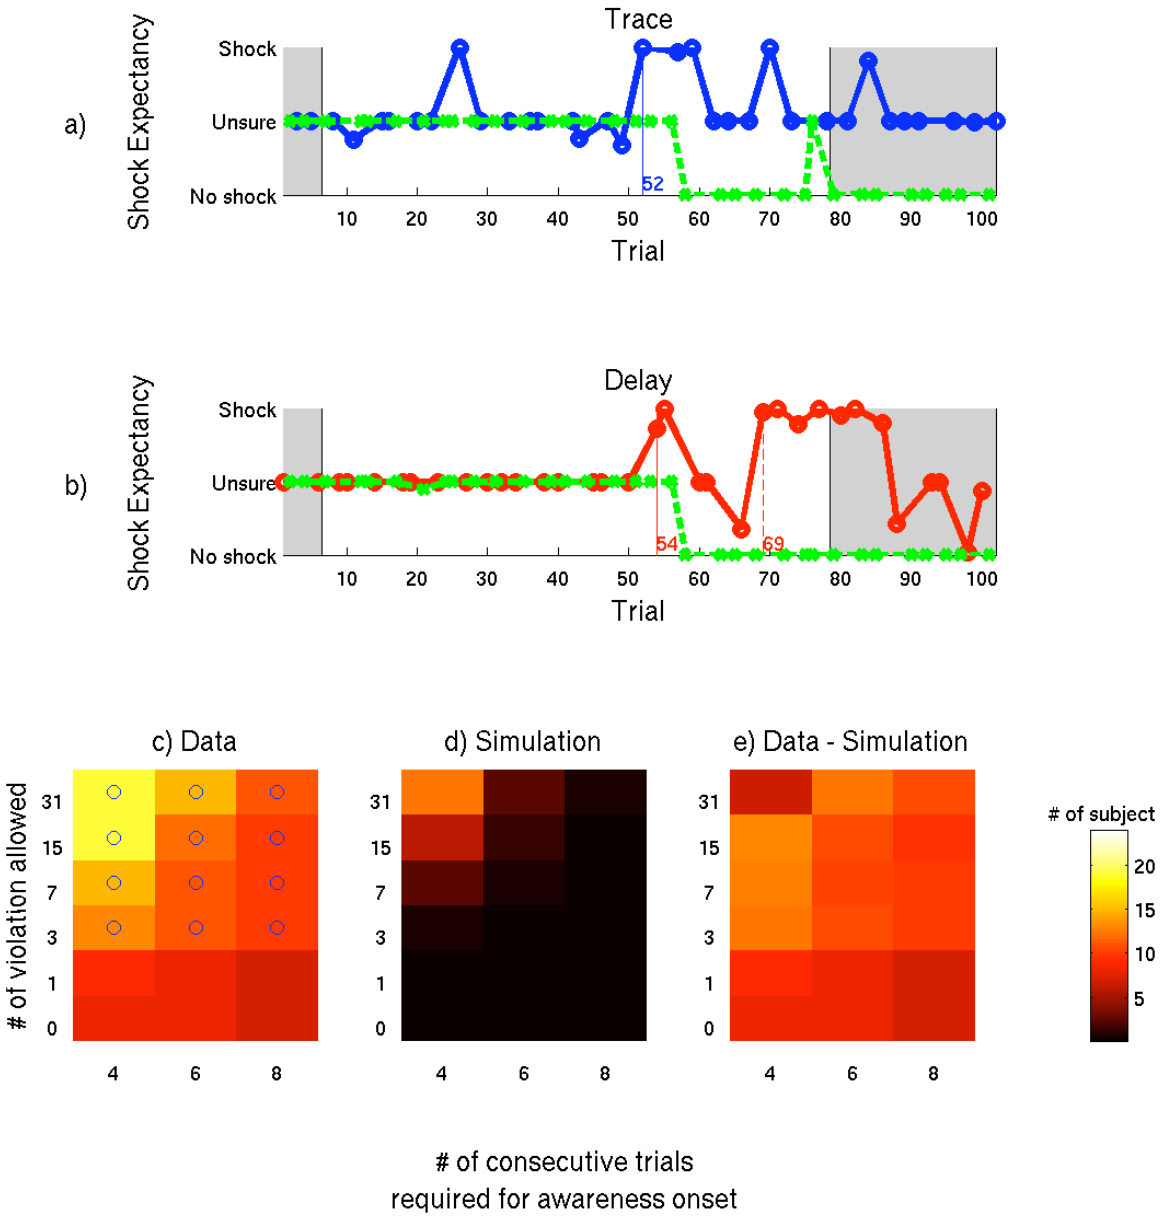
a)
Trace
Shock Expectancy
Shock
Unsure
No shock
52
10 20 30 40 50 60 70 80 90 100
Trial
b)
Delay
Shock Expectancy
Shock
Unsure
No shock
54
69
10 20 30 40 50 60 70 80 90 100
Trial
c) Data
d) Simulation
e) Data - Simulation
# of violation allowed
31
15
7
3
1
0
4 6 8
# of subject
20
15
10
5
# of consecutive trials
required for awareness onset

Figure 4.2 Definition of the onset of awareness

Shock expectancy from an exemplar subject in **a)** delay and **b)** trace conditioning, as a function of the trial number. The acquisition phase started at the 7th trial and ended at the 78th trial (the two gray regions correspond to habituation and extinction). Blue, red, and green corresponds to the shock expectancy scores for CS+T, CS+D, and CS- trials, respectively. **a**) For trace conditioning, this subject started expressing correct discrimination awareness from the 52nd trial onward (vertical blue), but violated it at the 76th and 78th trial. He did not demonstrate valence awareness. **b**) For delay conditioning, he demonstrated discrimination and valence awareness, starting at the 54th trial (solid vertical red) and 69th trial (broken vertical red), respectively. No violations occurred. **c**) The number of detected "aware" subjects (out of a total of 24) changes in a graded manner depending on both how many consecutive trials are required to detect the onset of awareness (N, x-axis), and how many violations are allowed (V, y-axis). The color scale on the right indicates the number of aware subjects for a given set of criteria. The 'o' symbols in **c**) denote the combinations of N and V where the exemplar subject is considered aware. **d**) The same analysis was repeated on simulated shock expectancy to estimate how many subjects could have been detected by chance. **e**) The difference between data and simulation. The number of aware subjects increases as the criterion becomes looser and is always larger than simulated data.

4.5.2. Trace but not delay conditioning coincides with the onset of awareness

Following the convention of the matrix notation in Figure 4.3 c-e, we show ΔSCR in four trial pairs (i.e., two trial pairs prior to the onset of awareness, at the onset of awareness, and one trial pair following the onset) in color-coded matrices in Figure 4.3 for trace (top row) and delay (middle row) conditioning, respectively; the same subjects contributed to both matrices at corresponding locations. To assess the degree of conditioning, we tested whether ΔSCR was significantly above 0 at each entry of the matrices; we used a one-tailed paired *t*-test (the degree of freedom was the number of subjects -1 at each entry in Figure 4.2 c; a Bonferroni correction was applied for multiple comparisons across the 4 trials). No correction for multiple comparisons was applied across N's and V's, as each entry was highly correlated with other entries (i.e., in Figure 4.2 c, the trial pairs of the exemplar subject were identical across all the entries with 'o' marks, that is, $V \geq 3$ for all N). ΔSCR was significantly above 0 ($P < 0.05$), only at trial pair 0, the onset of awareness, and only for trace conditioning. No other entries were significant following the Bonferroni correction.

Utilizing our advantage of within-subject comparison between trace and delay conditioning, we calculated the difference between ΔSCR for trace and delay for each trial pair and applied the same one-tailed *t*-test as above because we had a prior hypothesis that trace conditioning is more tightly related to awareness than is delay conditioning. In the bottom row of Figure 4.3, we show the difference as *t*-score (i.e., the difference of ΔSCR divided by the square root of the number of subjects in each entry of the matrix in Figure 4.2 c). For presentation purposes, we thresholded the *t*-score at $P <$

0.05. When the *P*-value did not survive the Bonferroni correction for the four trials considered, we marked the pixel with 'o.' The difference between trace and delay protocols was significant only for the trial pair 0, that is, at the onset of discrimination awareness. Notably, this was true when at least six consecutive trials were required for the detection of awareness. When the requirement was more transient (N = 4), only V = 0 or 3 violations could be allowed to find a significant difference between trace and delay conditioning.

The significant difference at trial pair 0 cannot be explained by the better overall conditionability in trace than in delay conditioning. At each entry of the matrix, we tested for a difference in the mean ΔSCR during acquisition (the first trial pair was excluded) between trace and delay (two-tailed *t*-test). There was no difference in any of the entries ($P > 0.17$ for all). Thus, the difference between trace and delay conditioning was seen specifically at the onset of discrimination awareness.

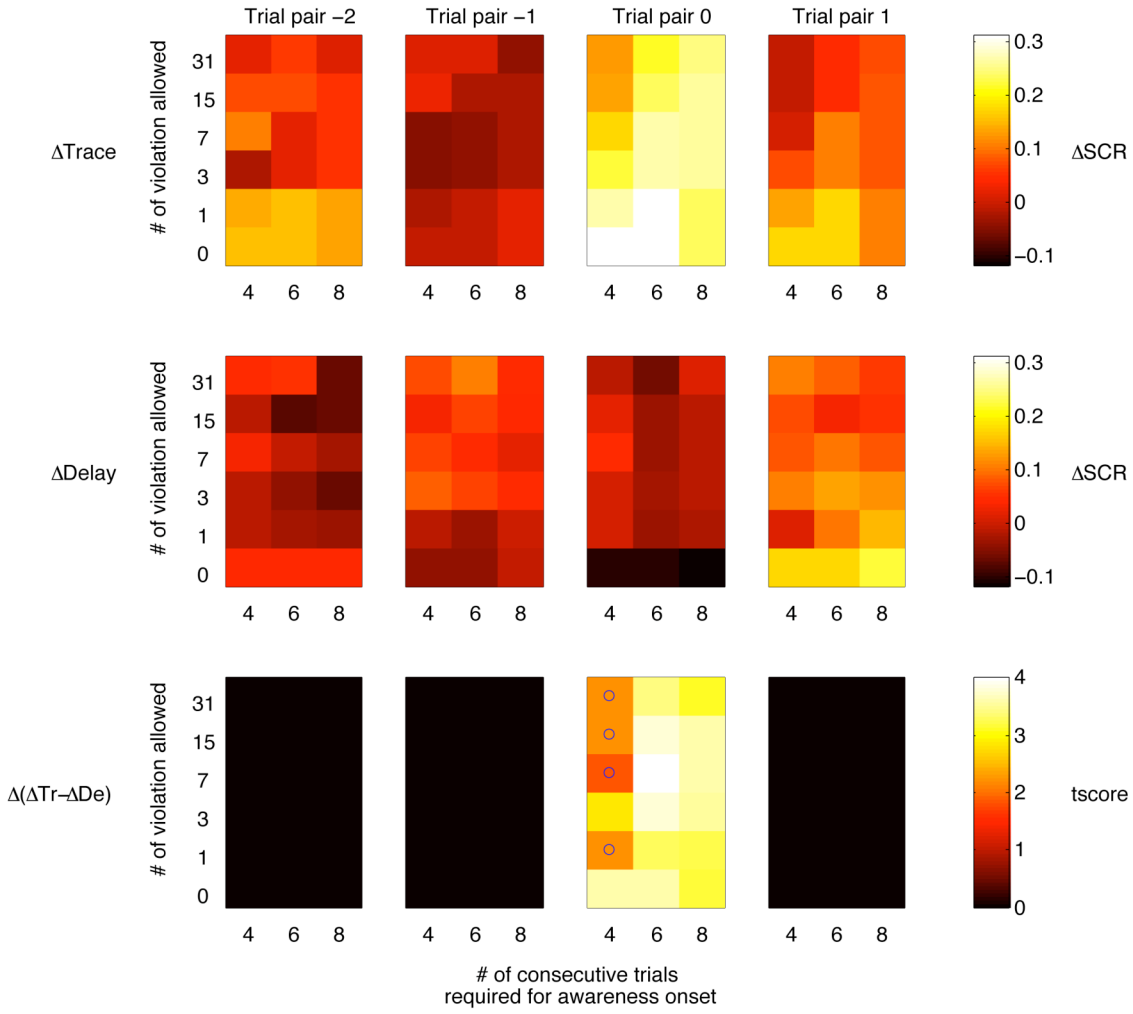

Figure 4.3 Differential skin conductance response (ΔSCR) around the onset of awareness

The convention for x- and y-axes in each color matrix follows Figure 4.2 c-e. ΔSCR is color-coded from trial pair -2 to +1 for trace (**top**) and delay (**middle**) conditioning around the onset of awareness. Note that the same subjects contributed to both trace and delay conditioning. (**bottom**) The difference between trace and delay ΔSCR is divided by the standard error to calculate a *t*-score (thresholded at $P < 0.05$). Those pixels that did

not survive the Bonferroni correction across 4 trials were marked by ‘o.’ Trace, but not delay conditioning, coincides with the onset of discrimination awareness.

4.5.3. Trace conditioning correlates to shock expectancy score

We performed linear regression analyses between the average ΔSCR and the mean difference between shock expectancy for CS+ and CS- (Δexpectancy). We found a significant correlation in 23 acquisition trials (excluding the first trial pair) within the time interval 1-2.75 sec in trace (Figure 4.4 left, blue, R = 0.476, *F* score = 6.45, *P* = 0.0187) but not delay conditioning (Figure 4.4 left, red, R = 0.294, *F* score =2.09, *P* = 0.162). Significant correlations were also found when we correlated the mean Δexpectancy with the difference between the mean SCR of 4 non-reinforced CS+ trials and the mean SCR of 23 CS- trials with the 1-6 sec interval in trace (Figure 4.4 middle, blue, R = 0.561, *F* score =10.1, *P* = 0.00433), but not delay conditioning (Figure 4.4 middle, red, R = 0.298, *F* score = 2.15, *P* = 0.157), and also when we correlated the mean Δexpectancy with the mean ΔSCR in all 23 acquisition trials with the 1-6 sec interval for trace conditioning (Figure 4.4 right, blue, R = 0.643, *F* score = 15.5, *P* = 0.000695).

To test if the correlation coefficient for trace conditioning is significantly better than delay conditioning, we used a one-tailed *z*-test (Because (Clark et al., 2002) reported a stronger correlation for trace conditioning, we had a prior hypothesis that trace correlates better than delay). Transforming the correlation coefficient R into Fisher's Z scores, we computed *P*-values on the difference between the two Fisher's Z scores. The difference was not significant when we used either the 1-2.75-sec interval (*P* = 0.24358, Figure 4.4 left) or the 1-6-sec interval and non-reinforced trials (*P* = 0.14467, Figure 4.4 middle).

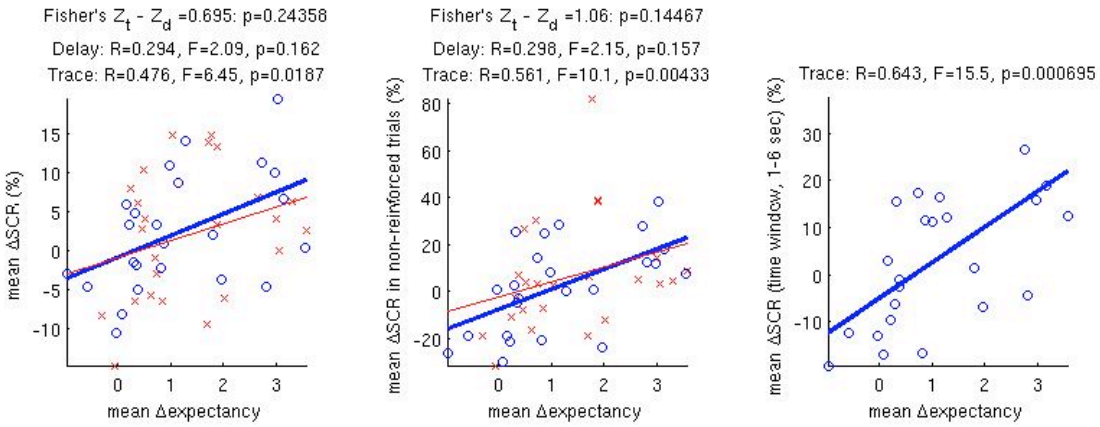

Figure 4.4 Trace conditioning correlates with shock expectancy score

The x-axis corresponds to the mean difference between the shock expectancy for CS+ and CS- during acquisition. The y-axis corresponds to the mean ΔSCR in the 1-2.75-sec interval in all acquisition trials (**left**), the mean ΔSCR in four non-reinforced trials with the 1-6-sec interval (**middle**) or, the mean ΔSCR in all 23 acquisition trials with the 1-6-sec interval (**right**). Trace (blue circles, thick blue lines) and delay (red cross, thin red lines) conditioning can be directly compared in the left and middle plots.

4.5.4. Negative evidence for unconscious delay conditioning

We did not find any evidence of conditioning for unaware subjects in any entry in the matrix (one-tailed *t*-test on the mean ΔSCR, $P > 0.38$ for all delay-unaware subjects and $P > 0.44$ for all trace-unaware subjects). We did not find any difference in the degree of conditioning among those unaware subjects between trace and delay conditioning in any entry of the matrix (paired two-tailed *t*-test on the mean ΔSCR, $P > 0.53$ for all).

In Figure 4.5, we show the number of unaware subjects (out of $n = 24$) in the same format as in Figure 4.2 c. The mean number of unaware subjects is 5.61, ranging from 2 to 8. The numbers in the corresponding entries of the two matrices (unaware subjects for Figure 4.5 and aware subjects for Figure 4.2 c) do not add up to 24 because some subjects were classified as aware in either delay or trace conditioning and as unaware in the alternative protocol.

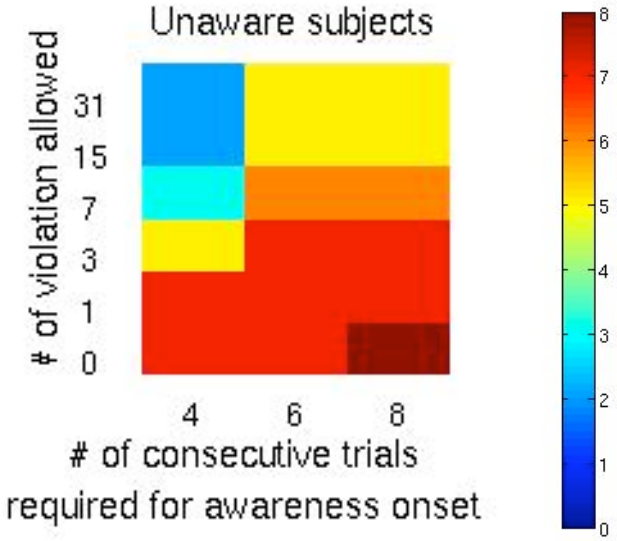

Figure 4.5 The number of unaware subjects

The convention follows Figure 4.2 c.

4.6. Discussion

We performed trace and delay conditioning in a single session, while subjects monitored their immediate shock expectancy throughout the session. By generalizing the definition of the onset of discrimination awareness (Dawson and Biferno, 1973; Biferno and Dawson, 1977), we detected awareness with graded levels of 'consecutiveness (N)' and 'violations tolerated (V).' Our results showed that trace, but not delay, conditioning coincided with the onset of discrimination awareness when we required at least 6 consecutive trials ($N \geq 6$) for detecting awareness. We reached the same conclusion when we included subjects who did not consistently express the correct shock expectancy after the onset of awareness (i.e., $V > 0$). This confirmed the reports by some subjects who commented that they did not express the correct expectancy throughout, and sometimes even intentionally violated it. Our generalized definition of the onset of awareness circumvented problems due to these uncertain strategies or behaviors and showed that including these subjects did not add noise. Our approach may be useful even when the reinforcement rate is 100% because some subjects may believe that their act of expressing expectancy (monitored by gaze direction in this study) influences the probability of shocking or because they may forget and fail to express the correct shock expectancy after they become aware of the CS-US relationship.

The contrasting relationship between trace and delay conditioning is consistent with the idea that trace and delay conditioning are mediated by two distinctive systems; trace conditioning is mediated by mechanisms with access to declarative knowledge, while delay conditioning is supported by non-declarative mechanisms. Our results are

consistent with previous aversive (Carter et al., 2003) and eyeblink studies (Clark et al., 2002), which showed that trace conditioning was more dependent on awareness than delay.

Comparison between trace and long delay conditioning (Ivkovich et al., 2000; Herbert et al., 2003) would be interesting and deserves extensive study in the future, though such a study would require a very large group of subjects due to the difficulty in balancing 4 CSs in a single session or in recruiting two separate groups with 2 CSs.

Delay aversive conditioning, presumably mediated by the amygdala (Fanselow and LeDoux, 1999), was not affected by the onset of awareness. The reason for this lack of modulation is unclear. In eyeblink conditioning, similar results are observed (Clark et al., 2001): subjects expected an airpuff with a higher probability when there had been no recent airpuff. This effect was observed in both trace and delay conditioning. Clark and colleagues found that the conditioned response increased for trace but decreased for delay as US expectancy increased, similar to our findings. An inhibitory interaction between the prefrontal cortex and the amygdala has been shown in rats (Rosenkranz et al., 2003). In human studies, the prefrontal cortex has been implicated in awareness during learning (McIntosh et al., 1999), (Carter et al., submitted). It is possible that a similar inhibitory interaction between the prefrontal cortex and the amygdala is at work in our subjects.

Similar to the conditioning studies where a brief temporal gap between CS and US eliminates successful trace conditioning under attentional distraction (Clark et al., 2002; Carter et al., 2003), implicit perception and behaviors are remarkably retarded when subjects have to hold their response for a brief duration (on the order of 1 to 10 sec) (Marcel, 1993; Milner and Goodale, 1995; Rossetti, 1998). This raises the possibility

that our brains may switch between conscious and unconscious operating modes depending on the temporal structures of the required task performance. However, previous studies of implicit processing have largely ignored the importance of the temporal structure between stimulus input and behavioral response. Because the implicit perception and behaviors themselves are still controversial (Morgan et al., 1997; Lovibond and Shanks, 2002), reports on implicit processing tend to emphasize only the positive existence of the subjects' ability to perform tasks without conscious percepts, and therefore the relationship between temporal factors and implicit processing has not been extensively studied. Implicit processing may degrade after some temporal delay, perhaps because the memory for implicit process may degrade very fast. However, it is also possible that neuronal systems for unconscious processing may be inhibited by that for conscious processing (Rossetti, 1998); in this study, we saw that delay conditioning did not increase at the onset of awareness, which cannot be explained by the short-memory-for-unconsciousness hypothesis but can be explained by the conscious-unconscious-competitive hypothesis.

In summary, our results support the 'necessary-gate' framework (Dawson and Furedy, 1976; Dawson and Schell, 1985; Lovibond and Shanks, 2002) and are consistent with (Dawson and Biferno, 1973; Biferno and Dawson, 1977) assuming that their 8-sec long CS delay protocol approximates trace conditioning at the neuronal level (Ivkovich et al., 2000; Herbert et al., 2003). The difference between trace and delay conditioning at the onset of awareness supports the view that these two types of learning are related to awareness in distinctive ways (Clark and Squire, 1998; Clark et al., 2002; Carter et al., 2003).

4.7. Acknowledgement

This work is a collaboration with McKell Carter. McKell developed the skin-conductance recording and electrical-shocking systems with PSYLAB and recorded the voice instructions. He developed a basic algorithm for detecting skin-conductance response, and I further improved the analytical methods. Rob Peters set up the eye tracker and gave me some tips for its usage. I appreciate Rufin VanRullen for providing the natural scene images used to distract subjects. Gabriel Kreiman provided me with detailed feedback on the original manuscript and suggested doing a computer simulation to assess the false-alarm rate for the onset of awareness. I thank John O'Doherty for his advice on the statistical analysis for Figure 4.3. David Anderson, Michael Franslow, Henry Lester, and their lab members provided us with useful comments throughout the development of this project in the series called the Conscious Mouse Project. This work was funded by grants from NIMH, NSF, the Keck Foundation, and the Moore Foundation.

5. References

Adelson EH, Movshon JA (1982) Phenomenal coherence of moving visual patterns. Nature 300:523-525.

Adelson EH, Bergen JR (1985) Spatiotemporal energy models for the perception of motion. J Opt Soc Am A 2:284-299.

Albright TD, Desimone R (1987) Local precision of visuotopic organization in the middle temporal area (MT) of the macaque. Exp Brain Res 65:582-592.

Alpern M, Barr L (1962) Durations of the after-images of brief light flashes and the theory of the Broca and Sulzer phenomenon. J Opt Soc Am 52:219-221.

Anderson SJ, Burr DC (1985) Spatial and temporal selectivity of the human motion detection system. Vision Res 25:1147-1154.

Anderson SJ, Burr DC (1989) Receptive field properties of human motion detector units inferred from spatial frequency masking. Vision Res 29:1343-1358.

Anderson SJ, Burr DC, Morrone MC (1991) Two-dimensional spatial and spatial-frequency selectivity of motion-sensitive mechanisms in human vision. J Opt Soc Am A 8:1340-1351.

Anstis S, Rogers B, Henry J (1978) Interactions between simultaneous contrast and coloured afterimages. Vision Res 18:899-911.

Bar M, Biederman I (1998) Subliminal visual priming. Psychol Sci 9:464-469.

Barlow HB (1981) The Ferrier Lecture, 1980. Critical limiting factors in the design of the eye and visual cortex. Proc R Soc Lond B Biol Sci 212:1-34.

Beauchamp MS, Cox RW, DeYoe EA (1997) Graded effects of spatial and featural attention on human area MT and associated motion processing areas. J Neurophysiol 78:516-520.

Ben-Shakhar G (1985) Standardization within individuals: a simple method to neutralize individual differences in skin conductance. Psychophysiology 22:292-299.

Biferno MA, Dawson ME (1977) Onset of Contingency Awareness and Electrodermal Classical-Conditioning - Analysis of Temporal Relationships During Acquisition and Extinction. Psychophysiology 14:164-171.

Blake R (1977) Threshold Conditions for Binocular Rivalry. J Exp Psychol 3:251-257.

Blake R (1995) Psychoanatomical strategies for studying human visual perception. In: Early vision and beyond (T.V. P, ed), pp 17-25. Cambridge, MA, USA: Massachusetts Institute of Technology.

Blake R, Fox R (1974) Adaptation to invisible gratings and the site of binocular rivalry suppression. Nature 249:488-490.

Blake R, Logothetis NK (2002) Visual competition. Nat Rev Neurosci 3:13-21.

Blake R, Overton R, Lema-Stern S (1981) Interocular transfer of visual aftereffects. J Exp Psychol Hum Percept Perform 7:367-381.

Blake R, Westendorf D, Fox R (1990) Temporal perturbations of binocular rivalry. Percept Psychophys 48:593-602.

Blake R, O'Shea RP, Mueller TJ (1992) Spatial zones of binocular rivalry in central and peripheral vision. Vis Neurosci 8:469-478.

Blake R, Sobel KV, Gilroy LA (2003) Visual motion retards alternations between conflicting perceptual interpretations. Neuron 39:869-878.

Blake RR, Fox R, McIntyre C (1971) Stochastic properties of stabilized-image binocular rivalry alternations. J Exp Psychol 88:327-332.

Blakemore C, Campbell FW (1969) On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. J Physiol 203:237-260.

Bonneh Y, Sagi D (1998) Effects of spatial configuration on contrast detection. Vision Res 38:3541-3553.

Bonneh YS, Cooperman A, Sagi D (2001) Motion-induced blindness in normal observers. Nature 411:798-801.

Bossink CJ, Stalmeier PF, De Weert CM (1993) A test of Levelt's second proposition for binocular rivalry. Vision Res 33:1413-1419.

Bowne SF (1990) Contrast discrimination cannot explain spatial frequency, orientation or temporal frequency discrimination. Vision Res 30:449-461.

Brainard DH (1997) The Psychophysics Toolbox. Spat Vis 10:433-436.

Braun J (1994) Visual search among items of different salience: removal of visual attention mimics a lesion in extrastriate area V4. J Neurosci 14:554-567.

Braun J (1998) Vision and attention: the role of training. Nature 393:424-425.

Braun J (2001) It's Great But Not Necessarily About Attention. In: Psyche.

Braun J, Julesz B (1998) Withdrawing attention at little or no cost: detection and discrimination tasks. Percept Psychophys 60:1-23.

Breese BB (1899) On Inhibition. Psychological Monographs 3:1-65.

Breitmeyer BG, Rudd M, Dunn K (1981) Metacontrast investigations of sustained-transient channel inhibitory interactions. J Exp Psychol Hum Percept Perform 7:770-779.

Brindley GS (1962) Two new properties of foveal after-images and a photochemical hypothesis to explain them. J Physiol 164:168-179.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. Vis Neurosci 10:1157-1169.

Burbeck CA, Kelly DH (1984) Role of local adaptation in the fading of stabilized images. J Opt Soc Am A 1:216-220.

Carrasco M, Penpeci-Talgar C, Eckstein M (2000) Spatial covert attention increases contrast sensitivity across the CSF: support for signal enhancement. Vision Res 40:1203-1215.

Carter RM, Hofstotter C, Tsuchiya N, Koch C (2003) Working memory and fear conditioning. Proc Natl Acad Sci U S A 100:1399-1404.

Chalmers DJ (1996) The conscious mind : in search of a fundamental theory. New York: Oxford University Press.

Chalmers DJ (2000) What is a neural correlate of consciousness? In: Neural Correlates of Consciousness: Empirical and Conceptual Questions (Metzinger Ts, ed), pp 17-40. Cambridge: MIT Press.

Chan D, Crutch SJ, Warrington EK (2001) A disorder of colour perception associated with abnormal colour after-images: a defect of the primary visual cortex. J Neurol Neurosurg Psychiatry 71:515-517.

Chaudhuri A (1990) Modulation of the motion aftereffect by selective attention. Nature 344:60-62.

Clark RE, Squire LR (1998) Classical conditioning and brain systems: the role of awareness. Science 280:77-81.

Clark RE, Manns JR, Squire LR (2001) Trace and delay eyeblink conditioning: contrasting phenomena of declarative and nondeclarative memory. Psychol Sci 12:304-308.

Clark RE, Manns JR, Squire LR (2002) Classical conditioning, awareness, and brain systems. Trends Cogn Sci 6:524-531.

Cowey A, Stoerig P (1991) The neurobiology of blindsight. Trends Neurosci 14:140-145.

Cowey A, Stoerig P (1995) Blindsight in monkeys. Nature 373:247-249.

Craik KJW (1940) Origin of Visual After-Images. Nature 145:512.

Crick F, Koch C (1995) Are we aware of neural activity in primary visual cortex? Nature 375:121-123.

Crick F, Koch C (1998) Consciousness and neuroscience. Cereb Cortex 8:97-107.

Dawson ME, Biferno MA (1973) Concurrent Measurement of Awareness and Electrodermal Classical-Conditioning. Journal of Experimental Psychology 101:55-62.

Dawson ME, Furedy JJ (1976) The role of awareness in human differential autonomic classical conditioning: the necessary-gate hypothesis. Psychophysiology 13:50-53.

Dawson ME, Schell AM (1985) Information processing and human autonomic classical conditioning. In: Advances in psychophysiology (Ackles PK, Jennings JR, Coles MGH, eds), pp 89-165. Greenwich, CT: JAI Press.

De Valois RL, Albrecht DG, Thorell LG (1982) Spatial frequency selectivity of cells in macaque visual cortex. Vision Res 22:545-559.

Dehaene S, Naccache L, Cohen L, Bihan DL, Mangin JF, Poline JB, Riviere D (2001) Cerebral mechanisms of word masking and unconscious repetition priming. Nat Neurosci 4:752-758.

Dehaene S, Naccache L, Le Clec HG, Koechlin E, Mueller M, Dehaene-Lambertz G, van de Moortele PF, Le Bihan D (1998) Imaging unconscious semantic priming. Nature 395:597-600.

Dember WN, Purcell DG (1967) Recovery of masked visual targets by inhibition of the masking stimulus. Science 157:1335-1336.

Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. Annu Rev Neurosci 18:193-222.

Dow BM, Snyder AZ, Vautin RG, Bauer R (1981) Magnification factor and receptive field size in foveal striate cortex of the monkey. Exp Brain Res 44:213-228.

Fanselow MS, LeDoux JE (1999) Why we think plasticity underlying Pavlovian fear conditioning occurs in the basolateral amygdala. Neuron 23:229-232.

Ferrera VP, Wilson HR (1987) Direction specific masking and the analysis of motion in two dimensions. Vision Res 27:1783-1796.

Festman Y, Ahissar M (2004) Attentional states and the degree of visual adaptation to gratings. Neural Netw 17:849-860.

Field DJ (1987) Relations between the statistics of natural images and the response properties of cortical cells. J Opt Soc Am A 4:2379-2394.

Foley JM (1994) Human luminance pattern-vision mechanisms: masking experiments require a new model. J Opt Soc Am A Opt Image Sci Vis 11:1710-1719.

Foley JM, Schwarz W (1998) Spatial attention: effect of position uncertainty and number of distractor patterns on the threshold-versus-contrast function for contrast discrimination. J Opt Soc Am A 15:1036-1047.

Fox R, Hermann J (1967) Stochastic properties of binocular rivalry alternations. Percept Psychophys 2:432-436.

Fox R, Rasche F (1969) Binocular rivalry and reciprocal inhibition. Percept Psychophys 5:215-217.

Fox R, Check R (1972) Independence between binocular rivalry suppression duration and magnitude of suppression. J Exp Psychol 93:283-289.

Freeman E, Sagi D, Driver J (2001) Lateral interactions between targets and flankers in low-level vision depend on attention to the flankers. Nat Neurosci 4:1032-1036.

Freeman E, Sagi D, Driver J (2004) Configuration-specific attentional modulation of flanker- -target lateral interactions. Perception 33:181-194.

Frith C, Perry R, Lumer E (1999) The neural correlates of conscious experience: an experimental framework. Trends Cogn Sci 3:105-114.

Gabrieli JD, McGlinchey-Berroth R, Carrillo MC, Gluck MA, Cermak LS, Disterhoft JF (1995) Intact delay-eyeblink classical conditioning in amnesia. Behav Neurosci 109:819-827.

Gandhi SP, Heeger DJ, Boynton GM (1999) Spatial attention affects brain activity in human primary visual cortex. Proc Natl Acad Sci U S A 96:3314-3319.

Geisler WS, Albrecht DG (1997) Visual cortex neurons in monkeys and cats: detection, discrimination, and identification. Vis Neurosci 14:897-919.

Georgeson MA, Turner RS (1985) Afterimages of sinusoidal, square-wave and compound gratings. Vision Res 25:1709-1720.

Gilroy LA, Blake R (2005) The interaction between binocular rivalry and negative afterimages. Curr Biol 15:1740-1744.

Guo K, Benson PJ, Blakemore C (2004) Pattern motion is present in V1 of awake but not anaesthetized monkeys. Eur J Neurosci 19:1055-1066.

Hayhoe MM, Williams DR (1984) Disappearance of afterimages at 'impossible' locations in space. Perception 13:455-459.

Haynes JD, Deichmann R, Rees G (2005) Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus. Nature.

He S, MacLeod DI (2001) Orientation-selective adaptation and tilt after-effect from invisible patterns. Nature 411:473-476.

He S, Cavanagh P, Intriligator J (1996) Attentional resolution and the locus of visual awareness. Nature 383:334-337.

Heeger DJ (1993) Modeling simple-cell direction selectivity with normalized, half-squared, linear operators. J Neurophysiol 70:1885-1898.

Heeger DJ, Boynton GM, Demb JB, Seidemann E, Newsome WT (1999) Motion opponency in visual cortex. J Neurosci 19:7162-7174.

Herbert JS, Eckerman CO, Stanton ME (2003) The ontogeny of human learning in delay, long-delay, and trace eyeblink conditioning. Behav Neurosci 117:1196-1210.

Heuer HW, Britten KH (2002) Contrast dependence of response normalization in area MT of the rhesus macaque. J Neurophysiol 88:3398-3408.

Hofstoetter C, Koch C, Kiper DC (2004) Motion-induced blindness does not affect the formation of negative afterimages. Consciousness and Cognition.

Hosoya T, Baccus SA, Meister M (2005) Dynamic predictive coding by the retina. Nature 436:71-77.

Huk AC, Heeger DJ (2000) Task-related modulation of visual cortex. J Neurophysiol 83:3525-3536.

Huk AC, Heeger DJ (2002) Pattern-motion responses in human visual cortex. Nat Neurosci 5:72-75.

Huk AC, Ress D, Heeger DJ (2001) Neuronal basis of the motion aftereffect reconsidered. Neuron 32:161-172.

Itti L, Koch C, Braun J (2000) Revisiting spatial vision: toward a unifying model. J Opt Soc Am A Opt Image Sci Vis 17:1899-1917.

Ivkovich D, Paczkowski CM, Stanton ME (2000) Ontogeny of delay versus trace eyeblink conditioning in the rat. Dev Psychobiol 36:148-160.

Kanai R, Tsuchiya N, Verstraten FAJ (2004) Featural, but not spatial, attention modulates unconscious processing of visual stimuli. Perception 33:8-8.

Kelly DH, Martinez-Uriegas E (1993) Measurements of chromatic and achromatic afterimages. J Opt Soc Am A 10:29-37.

Kentridge RW, Heywood CA, Weiskrantz L (1999a) Effects of temporal cueing on residual visual discrimination in blindsight. Neuropsychologia 37:479-483.

Kentridge RW, Heywood CA, Weiskrantz L (1999b) Attention without awareness in blindsight. Proc R Soc Lond B Biol Sci 266:1805-1811.

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

Klein RM (2000) Inhibition of return. Trends Cogn Sci 4:138-147.

Knight DC, Cheng DT, Smith CN, Stein EA, Helmstetter FJ (2004) Neural substrates mediating human delay and trace fear conditioning. J Neurosci 24:218-228.

Koch C (2004) The Quest for Consciousness: A neurobiological Approach. CO: Roberts and Publishers.

Koivisto M, Revonsuo A, Salminen N (2005a) Independence of visual awareness from attention at early processing stages. Neuroreport 16:817-821.

Koivisto M, Revonsuo A, Lehtonen M (2005b) Independence of Visual Awareness from the Scope of Attention: an Electrophysiological Study. Cereb Cortex.

Kreiman G, Fried I, Koch C (2002) Single-neuron correlates of subjective vision in the human medial temporal lobe. Proc Natl Acad Sci U S A 99:8378-8383.

Lack LC (1978) Selective attention and the control of binocular rivalry. The Hague, The Netherlands: Mouton Publishers.

Lambert a, Naikar N, McLachlan K, Aitken V (1999) A new component of visual orienting: Implicit effects of peripheral information and subthreshold cues on covert attention. Journal of Experimental Psychology-Human Perception and Performance 25:321-340.

Lamme VA (2003) Why visual attention and awareness are different. Trends Cogn Sci 7:12-18.

Lee DK, Koch C, Braun J (1997) Spatial vision thresholds in the near absence of attention. Vision Res 37:2409-2418.

Lee DK, Koch C, Braun J (1999a) Attentional capacity is undifferentiated: concurrent discrimination of form, color, and motion. Percept Psychophys 61:1241-1255.

Lee DK, Itti L, Koch C, Braun J (1999b) Attention activates winner-take-all competition among visual filters. Nat Neurosci 2:375-381.

Legge GE, Foley JM (1980) Contrast masking in human vision. J Opt Soc Am 70:1458-1471.

Leguire LE, Blake R (1982) Role of threshold in afterimage visibility. J Opt Soc Am 72:1232-1237.

Lehky SR (1988) An astable multivibrator model of binocular rivalry. Perception 17:215-228.

Lehky SR, Blake R (1991) Organization of binocular pathways: modeling and data related to rivalry Neural Comput 3:44-53.

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

Lehmkuhle SW, Fox R (1975) Effect of binocular rivalry suppression on the motion aftereffect. Vision Res 15:855-859.

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

Leopold DA, Logothetis NK (1999) Multistable phenomena: changing views in perception. Trends Cogn Sci 3:254-264.

Leopold DA, Wilke M, Maier A, Logothetis NK (2002) Stable perception of visually ambiguous patterns. Nat Neurosci 5:605-609.

Levelt WJM (1965) On binocular rivalry. Soesterberg, The Netherlands: Institute for perception RVO-TNO.

Levinson E, Sekuler R (1975) The independence of channels in human vision selective for direction of movement. J Physiol 250:347-366.

Levitt H (1971) Transformed up-down methods in psychoacoustics. J Acoust Soc Am 49:Suppl 2:467+.

Li FF, VanRullen R, Koch C, Perona P (2002) Rapid natural scene categorization in the near absence of attention. Proc Natl Acad Sci U S A 99:9596-9601.

Liu L, Tyler CW, Schor CM (1992) Failure of rivalry at low contrast: evidence of a suprathreshold binocular summation process. Vision Res 32:1471-1479.

Logothetis NK (1998) Single units and conscious vision. Philos Trans R Soc Lond B Biol Sci 353:1801-1818.

Logothetis NK, Schall JD (1989) Neuronal correlates of subjective visual perception. Science 245:761-763.

Loomis JM (1972) The photopigment bleaching hypothesis of complementary after-images: a psychophysical test. Vision Res 12:1587-1594.

Loomis JM (1978) Complementary afterimages and the unequal adapting effects of steady and flickering light. J Opt Soc Am 68:411-416.

Lou L (2001) Effects of voluntary attention on structured afterimages. Perception 30:1439-1448.

Lovibond PF, Shanks DR (2002) The role of awareness in Pavlovian conditioning: empirical evidence and theoretical implications. J Exp Psychol Anim Behav Process 28:3-26.

Lu ZL, Sperling G (1995) The functional architecture of human visual motion perception. Vision Res 35:2697-2722.

Lu ZL, Sperling G (1996) Contrast gain control in first- and second-order motion perception. J Opt Soc Am A Opt Image Sci Vis 13:2305-2318.

Lumer ED, Friston KJ, Rees G (1998) Neural correlates of perceptual rivalry in the human brain. Science 280:1930-1934.

Lykken DT (1972) Range correction applied to heart rate and to GSR data. Psychophysiology 9:373-379.

Mack A, Rock I (1998) Inattentional blindness. Cambridge, Mass.: MIT Press.

Macknik SL, Livingstone MS (1998) Neuronal correlates of visibility and invisibility in the primate visual system. Nat Neurosci 1:144-149.

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.

Macknik SL, Martinez-Conde S, Haglund MM (2000) The role of spatiotemporal edges in visibility and visual masking. Proc Natl Acad Sci U S A 97:7556-7560.

Marcel AJ (1993) Slippage in the unity of consciousness. Ciba Found Symp 174:168-180; discussion 180-166.

Martinez-Trujillo J, Treue S (2002) Attentional modulation strength in cortical area MT depends on stimulus contrast. Neuron 35:365-370.

Matsuoka K (1984) The dynamic model of binocular rivalry. Biol Cybern 49:201-208.

Maunsell JH (1995) The brain's visual world: representation of visual targets in cerebral cortex. Science 270:764-769.

McAdams CJ, Maunsell JH (1999a) Effects of attention on the reliability of individual neurons in monkey visual cortex. Neuron 23:765-773.

McAdams CJ, Maunsell JH (1999b) Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. J Neurosci 19:431-441.

McCormick PA (1997) Orienting attention without awareness. Journal of Experimental Psychology-Human Perception and Performance 23:168-180.

McIntosh AR, Rajah MN, Lobaugh NJ (1999) Interactions of prefrontal cortex in relation to awareness in sensory learning. Science 284:1531-1533.

Meese TS, Williams CB (2000) Probability summation for multiple patches of luminance modulation. Vision Res 40:2101-2113.

Melcher D, Papathomas TV, Vidnyanszky Z (2005) Implicit attentional selection of bound visual features. Neuron 46:723-729.

Merikle JS, Reingold EM (1992) Measuring unconscious perceptual processes. In: Perception without awareness: Cognitive, clinical, and social perspectives (Bornstein RF, Pitman TS, eds), pp 55-80. New York: Guilford Press.

Milner B, Squire LR, Kandel ER (1998) Cognitive neuroscience and the study of memory. Neuron 20:445-468.

Milner DA, Goodale MA (1995) The visual brain in action. Oxford: Oxford University Press.

Montaser-Kouhsari L, Rajimehr R (2005) Subliminal attentional modulation in crowding condition. Vision Res 45:839-844.

Moradi F, Koch C, Shimojo S (2005) Face adaptation depends on seeing the face. Neuron 45:169-175.

Moran J, Desimone R (1985) Selective attention gates visual processing in the extrastriate cortex. Science 229:782-784.

Morgan MJ, Mason AJ, Solomon JA (1997) Blindsight in normal subjects? Nature 385:401-402.

Morris JS, Ohman A, Dolan RJ (1998) Conscious and unconscious emotional learning in the human amygdala. Nature 393:467-470.

Morrone MC, Denti V, Spinelli D (2002) Color and luminance contrasts attract independent attention. Curr Biol 12:1134-1137.

Movshon JA, Newsome WT (1996) Visual response properties of striate cortical neurons projecting to area MT in macaque monkeys. J Neurosci 16:7733-7741.

Movshon JA, Adelson EH, Gizzi MS, Newsome WT (1985) The analysis of moving visual patterns. In: Study Group on Pattern Recognition Mechanisms (Chagas C, and al e, eds), pp 117-151. Vatican City: Pontifica Academia Scientiarium.

Mueller TJ (1990) A physiological model of binocular rivalry. Vis Neurosci 4:63-73.

Mueller TJ, Blake R (1989) A fresh look at the temporal dynamics of binocular rivalry. Biol Cybern 61:223-232.

Naccache L, Blandin E, Dehaene S (2002) Unconscious masked priming depends on temporal attention. Psychol Sci 13:416-424.

Nguyen VA, Freeman AW, Wenderoth P (2001) The depth and selectivity of suppression in binocular rivalry. Percept Psychophys 63:348-360.

Nguyen VA, Freeman AW, Alais D (2003) Increasing depth of binocular rivalry suppression along two visual pathways. Vision Res 43:2003-2008.

Norman HF, Norman JF, Bilotta J (2000) The temporal course of suppression during binocular rivalry. Perception 29:831-841.

O'Craven KM, Downing PE, Kanwisher N (1999) fMRI evidence for objects as the units of attentional selection. Nature 401:584-587.

O'Regan JK, Noe A (2001) A sensorimotor account of vision and visual consciousness. Behav Brain Sci 24:939-973; discussion 973-1031.

O'Shea RP, Crassini B (1981) Interocular transfer of the motion after-effect is not reduced by binocular rivalry. Vision Res 21:801-804.

Ohki K, Chung S, Ch'ng YH, Kara P, Reid RC (2005) Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. Nature 433:597-603.

Ohman A, Soares JJ (1998) Emotional conditioning to masked stimuli: expectancies for aversive outcomes following nonrecognized fear-relevant stimuli. J Exp Psychol Gen 127:69-82.

Pasley BN, Mayes LC, Schultz RT (2004) Subcortical discrimination of unperceived objects during binocular rivalry. Neuron 42:163-172.

Polat U, Sagi D (1993) Lateral interactions between spatial channels: suppression and facilitation revealed by lateral masking experiments. Vision Res 33:993-999.

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.

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

Qian N, Andersen RA (1994) Transparent motion perception as detection of unbalanced motion signals. II. Physiology. J Neurosci 14:7367-7380.

Quick RF, Jr. (1974) A vector-magnitude model of contrast detection. Kybernetik 16:65-67.

Rajimehr R (2004) Unconscious orientation processing. Neuron 41:663-673.

Raymond JE (2000) Attentional modulation of visual motion perception. Trends Cogn Sci 4:42-50.

Raymond JE, O'Donnell HL, Tipper SP (1998) Priming reveals attentional modulation of human motion sensitivity. Vision Res 38:2863-2867.

Rees G, Friston K, Koch C (2000) A direct quantitative relationship between the functional properties of human and macaque V5. Nat Neurosci 3:716-723.

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

Robinson DN (1966) Disinhibition of visually masked stimuli. Science 154:157-158.

Rosenkranz JA, Moore H, Grace AA (2003) The prefrontal cortex regulates lateral amygdala neuronal plasticity and responses to previously conditioned stimuli. J Neurosci 23:11054-11064.

Rossetti Y (1998) Implicit short-lived motor representations of space in brain damaged and healthy subjects. Conscious Cogn 7:520-558.

Saenz M, Buracas GT, Boynton GM (2002) Global effects of feature-based attention in human visual cortex. Nat Neurosci 5:631-632.

Sakitt B (1976) Psychophysical correlates of photoreceptor activity. Vision Res 16:129-140.

Schall JD, Nawrot M, Blake R, Yu K (1993) Visually guided attention is neutralized when informative cues are visible but unperceived. Vision Res 33:2057-2064.

Schieting S, Spillmann L (1987) Flicker adaptation in the peripheral retina. Vision Res 27:277-284.

Schiller PH, Dolan RP (1994) Visual aftereffects and the consequences of visual system lesions on their perception in the rhesus monkey. Vis Neurosci 11:643-665.

Schrater PR, Knill DC, Simoncelli EP (2000) Mechanisms of visual motion detection. Nat Neurosci 3:64-68.

Seidemann E, Newsome WT (1999) Effect of spatial attention on the responses of area MT neurons. J Neurophysiol 81:1783-1794.

Sengpiel F, Blakemore C, Harrad R (1995) Interocular suppression in the primary visual cortex: a possible neural basis of binocular rivalry. Vision Res 35:179-195.

Sheinberg DL, Logothetis NK (1997) The role of temporal cortical areas in perceptual organization. Proc Natl Acad Sci U S A 94:3408-3413.

Shimojo S, Kamitani Y, Nishida S (2001) Afterimage of perceptually filled-in surface. Science 293:1677-1680.

Shmuel A, Grinvald A (1996) Functional organization for direction of motion and its relationship to orientation maps in cat area 18. J Neurosci 16:6945-6964.

Simoncelli EP, Heeger DJ (1998) A model of neuronal responses in visual area MT. Vision Res 38:743-761.

Smirnakis SM, Berry MJ, Warland DK, Bialek W, Meister M (1997) Adaptation of retinal processing to image contrast and spatial scale. Nature 386:69-73.

Snowden RJ, Treue S, Erickson RG, Andersen RA (1991) The response of area MT and V1 neurons to transparent motion. J Neurosci 11:2768-2785.

Sobel KV, Blake R (2002) How context influences predominance during binocular rivalry. Perception 31:813-824.

Sobel KV, Blake R, Raissian TA (2004) Binocular rivalry suppression does impede buildup of the motion aftereffect. Journal of Vision 4:243a.

Solomon JA, Lavie N, Morgan MJ (1997) Contrast discrimination function: spatial cuing effects. J Opt Soc Am A Opt Image Sci Vis 14:2443-2448.

Spitzer H, Desimone R, Moran J (1988) Increased attention enhances both behavioral and neuronal performance. Science 240:338-340.

Spivey MJ, Spirn MJ (2000) Selective visual attention modulates the direct tilt aftereffect. Percept Psychophys 62:1525-1533.

Stoner GR, Albright TD (1992) Neural correlates of perceptual motion coherence. Nature 358:412-414.

Stoner GR, Albright TD, Ramachandran VS (1990) Transparency and coherence in human motion perception. Nature 344:153-155.

Stromeyer CF, 3rd, Kronauer RE, Madsen JC, Klein SA (1984) Opponent-movement mechanisms in human vision. J Opt Soc Am A 1:876-884.

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

Tong F, Engel SA (2001) Interocular rivalry revealed in the human cortical blind-spot representation. Nature 411:195-199.

Tong F, Nakayama K, Vaughan JT, Kanwisher N (1998) Binocular rivalry and visual awareness in human extrastriate cortex. Neuron 21:753-759.

Tootell RB, Reppas JB, Kwong KK, Malach R, Born RT, Brady TJ, Rosen BR, Belliveau JW (1995) Functional analysis of human MT and related visual cortical areas using magnetic resonance imaging. J Neurosci 15:3215-3230.

Treue S, Maunsell JH (1996) Attentional modulation of visual motion processing in cortical areas MT and MST. Nature 382:539-541.

Treue S, Martinez Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. Nature 399:575-579.

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

Tyler CW (1997) Colour bit-stealing to enhance the luminance resolution of digital displays on a single pixel basis. Spat Vis 10:369-377.

Tyler CW, Chen CC (2000) Signal detection theory in the 2AFC paradigm: attention, channel uncertainty and probability summation. Vision Res 40:3121-3144.

Varela FJ, Singer W (1987) Neuronal dynamics in the visual corticothalamic pathway revealed through binocular rivalry. Exp Brain Res 66:10-20.

Venables PH, Christie MJ (1980) Electrodermal Activity. In: Techniques in Psychophysiology (Martin I, Venables PH, eds), pp 3-67. New York: Wiley & Sons.

Virsu V (1978) Retinal mechanisms of visual adaptation and afterimages. Med Biol 56:84-96.

Virsu V, Laurinen P (1977) Long-lasting afterimages caused by neural adaptation. Vision Res 17:853-860.

Wade NJ, Wenderoth P (1978) The influence of colour and contour rivalry on the magnitude of the tilt after-effect. Vision Res 18:827-835.

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

Watanabe T, Harner AM, Miyauchi S, Sasaki Y, Nielsen M, Palomo D, Mukai I (1998a) Task-dependent influences of attention on the activation of human primary visual cortex. Proc Natl Acad Sci U S A 95:11489-11492.

Watanabe T, Sasaki Y, Miyauchi S, Putz B, Fujimaki N, Nielsen M, Takino R, Miyakawa S (1998b) Attention-regulated activity in human primary visual cortex. J Neurophysiol 79:2218-2221.

Weiskrantz L (1950) An unusual case of after-imagery following fixation of an "imaginary" visual pattern. Quarterly Journal of Experimental Psychology:170-175.

Weiskrantz L (2002) Prime-sight and blindsight. Conscious Cogn 11:568-581.

Weiss Y, Simoncelli EP, Adelson EH (2002) Motion illusions as optimal percepts. Nat Neurosci 5:598-604.

Welch L (1989) The perception of moving plaids reveals two motion-processing stages. Nature 337:734-736.

Weliky M, Bosking WH, Fitzpatrick D (1996) A systematic map of direction preference in primary visual cortex. Nature 379:725-728.

Wiesenfelder H, Blake R (1990) The neural site of binocular rivalry relative to the analysis of motion in the human visual system. J Neurosci 10:3880-3888.

Wilke M, Logothetis NK, Leopold DA (2003) Generalized flash suppression of salient visual targets. Neuron 39:1043-1052.

Williams MA, Morris AP, McGlone F, Abbott DF, Mattingley JB (2004) Amygdala responses to fearful and happy facial expressions under conditions of binocular suppression. J Neurosci 24:2898-2904.

Wilson HR (1980) A transducer function for threshold and suprathreshold human vision. Biol Cybern 38:171-178.

Wilson HR (1985) A model for direction selectivity in threshold motion perception. Biol Cybern 51:213-222.

Wilson HR (1997) A neural model of foveal light adaptation and afterimage formation. Vis Neurosci 14:403-423.

Wilson HR (2003) Computational evidence for a rivalry hierarchy in vision. Proc Natl Acad Sci U S A 100:14499-14503.

Wilson HR, Kim J (1994) A model for motion coherence and transparency. Vis Neurosci 11:1205-1220.

Wolfe JM (1983) Influence of spatial frequency, luminance, and duration on binocular rivalry and abnormal fusion of briefly presented dichoptic stimuli. Perception 12:447-456.

Wolfe JM (1984) Reversing ocular dominance and suppression in a single flash. Vision Res 24:471-478.

Wunderlich K, Schneider KA, Kastner S (2005) Neural correlates of binocular rivalry in the human lateral geniculate nucleus. Nat Neurosci 8:1595-1602.

Zenger B, Sagi D (1996) Isolating excitatory and inhibitory nonlinear spatial interactions involved in contrast detection. Vision Res 36:2497-2513.

Zenger B, Braun J, Koch C (2000) Attentional effects on contrast detection in the presence of surround masks. Vision Res 40:3717-3724.