

Perception of ambiguous images within temporal context

Perceptie van ambigue afbeeldingen in temporele context
(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht op gezag van de rector magnificus, prof.dr. J.C. Stoof, ingevolge het besluit van het college voor promoties in het openbaar te verdedigen op maandag 29 september 2008 des middags te 12.45 uur

door

Jan Willem Brascamp

geboren op 24 december 1977
te Ede

Promotor: Prof.dr. A.V. van den Berg
Co-promotor: Dr. R. van Ee

Contents

Introduction	2
General introduction	4
Reasons for studying ambiguous stimuli	5
Perceptual alternations	6
Onset dominance	7
1 Distributions of alternation rates in various forms of bistable perception	8
1.1 Introduction	11
1.2 Data collection	15
1.3 Gamma distributed percept durations?	16
1.4 Inter-stimulus comparison	20
1.5 Discussion	22
1.6 Conclusion	26
2 The time course of binocular rivalry implies a crucial role of noise	27
2.1 Introduction	30
2.2 Methods	31
2.3 Results	32
2.4 Discussion	37
2.5 Conclusion	41
Appendix	42
3 Sensory memory for ambiguous vision	47
3.1 Stabilizing unstable vision	50
3.2 Information held in memory	50
3.3 Structure from motion	50
3.4 Binocular rivalry	54
3.5 Temporal characteristics of memory	55
3.6 Unambiguous prior stimuli	56
3.7 Modeling	57
3.8 Related memory phenomena	58
3.9 Conclusion	60

4	Flash suppression and flash facilitation in binocular rivalry	61
4.1	Introduction	64
4.2	Methods	65
4.3	Results	67
4.4	Discussion	74
	Appendix	78
5	Multi-timescale perceptual history resolves visual ambiguity	83
5.1	Introduction	86
5.2	Results	88
5.3	Discussion	94
5.4	Materials and Methods	97
	Appendix	98
6	Intermittent ambiguous stimuli cause periodic perceptual alternations	106
6.1	Introduction	109
6.2	Methods	110
6.3	Results	111
6.4	Discussion	118
6.5	Conclusion	121
	Appendix	122
7	Perceptual memory: toward neural measurement	128
7.1	Introduction	131
7.2	Results	133
7.3	Discussion	144
7.4	Conclusion	146
7.5	Methods	146
	Appendix	148
	Summary and discussion	150
	Summary per chapter	150
	General discussion	152
	Samenvatting	157
	Acknowledgements	171
	Publications	173
	Curriculum vitae	175

Introduction

I learned a lot of new things as a PhD student, and only a fraction of them are represented in the following chapters. For one thing, the way scientific papers are written they present results in a logical context; they usually do not describe the often haphazard sequence of events that led to the results. Also, a large part of becoming a scientist is not about the subject matter itself, but about things like learning what questions to ask, how to design experiments, when to abandon a plan that initially looked promising, and how to communicate your findings to fellow scientists. At the beginning of every chapter I therefore inserted a short text in italics, like this one, to provide some context to the science. Sometimes I will explain how we came across a particular experimental result, sometimes I will describe how I think the work fits into the literature, sometimes I will suggest what I feel we could have done better, and sometimes I will just recount an anecdote that I find amusing.

General introduction

Whenever we open our eyes the mind conjures up a vivid impression of the world around us. We expect this impression to be accurate and reliable, and we usually do not spend much thought on how it came into being. In fact, the movie in our mind's eye is so naturally associated with the outside world that it is easy to forget that the two are not one and the same, and that a process of active interpretation is required to construct one from the other. For me personally, I remember one of the first times I realized how thoroughly misleading our familiar personal impressions can be regarding the workings of the visual brain. Someone told me of a condition called motion blindness, where a patient has perfectly good vision, apart from the fact that they cannot perceive motion (usually due to a stroke). Instead, these patients perceive the world as a series of stills, presumably somewhat similar to the effect of a stroboscope. The great surprise to me was that one component of visual experience, motion in this case, could simply be removed while leaving the remainder intact. By revealing modularity in the apparently indivisible movie in the mind's eye, this finding drew my attention to the hidden machinery underlying visual perception.

We do not need to turn to rare disorders to realize that intricate processing takes place on the way between light falling on the retina and visual experience. Many everyday observations would lead to similar conclusions. For instance, why does the moon look larger just above the horizon than up in the sky? How can separate eyewitnesses give highly diverging accounts of the same event? Why doesn't the world look half as dark when I close one eye, like a photo does when the shutter size of the camera is reduced by half? Such observations allow a peek behind the scenes of visual processing, permitting inferences regarding the construction of a visual percept. This is one of the main goals of vision science: to understand how visual experience is constructed from the pattern of light that hits our retinas.

This thesis contributes to this understanding by investigating perception of so-called ambiguous visual stimuli. These are images that are equally consistent with each of several perceptual interpretations. They thus challenge the visual system's ability to make sense out of visual input, and as such are well suited to illuminate mechanisms involved in the construction of a visual percept. Figure 1 shows a simple example of ambiguous stimuli in a screenshot taken from a standard e-mail client. The grey arrows to the left of the inbox icon and the sent box icon can each be seen as pointing in any of three directions (see Attneave, 1968). These interpretations are perceptual ones: it is not a matter of acknowledging the existence of three possible solutions, but of alternately perceiving distinct kinds of arrows while being exposed to unchanging visual input.

Typically, when faced with an ambiguous stimulus, perception alternates over time between alternative interpretations. This arguably reflects the brain's continuous effort to re-evaluate its organization of incoming information. The first two chapters of this thesis are aimed at studying perceptual alternations in response to ambiguous input, and thereby gaining insight into this process of perceptual organization.

When constructing a perceptual interpretation, the visual system shows remarkable resourcefulness in assembling and integrating various sources of contextual information. For instance, in Figure 1, one reason why the arrow beside the sent box is most readily seen as pointing rightward, may be that this direction is aligned with the direction of the word 'Sent'. Also, my impression of this arrow is that it is more easily seen as pointing to the bottom left than as pointing to the top left. This may well be due to a repulsive interaction with the paper



Figure 1: *One regularly encounters ambiguous images in everyday life. The triangular arrows in this screenshot of a popular e-mail client have multiple perceptual interpretations. Each arrow can be seen as pointing in either of three directions. Upon prolonged inspection most people will see all three alternatives in turn.*

plane next to it, which points to the top right.

These are examples of the use of spatial context in percept formation. Similarly, any visual stimulus is embedded within a temporal context. Because sequences of visual input do not progress randomly but contain a degree of predictability, the visual system could use temporal context to interpret current input. The greater part of this thesis, starting from the third chapter, deals with the question of how perception of ambiguous stimuli depends on the preceding history of visual input and perception.

Reasons for studying ambiguous stimuli

Ambiguous stimuli have been the subject of scientific research for nearly two centuries (Necker, 1832; Wheatstone, 1838). Why are people drawn to this subject?

It is not common to see an object we know to be static spontaneously flip between configurations right before our eyes. Undoubtedly this counterintuitive quality has helped draw many investigators to the study of ambiguous stimuli. In fact, the first paper on the well-known Necker cube (Figure 2) was not written by a vision scientist working to understand the brain, but by a geologist called Louis Albert Necker, reporting a striking phenomenon he observed while studying drawings of crystals. Similarly, the frequent appearance of ambiguous stimuli in visual art (e.g. M.C. Escher, Salvador Dalí, various eighteenth century artists) testifies to their inherent appeal to the human mind.

Another reason for studying ambiguous images is the idea that the observed perceptual alternations, although exceptional in their saliency, reflect selection processes that are involved in visual perception in general. Indeed, some degree of ambiguity is common to all visual input, in the sense that the problem of reconstructing a physical scene from a given pattern of retinal stimulation is generally underdetermined. Moreover, eye movements and visual attention in normal vision constitute selection processes akin to perceptual dominance and suppression in ambiguous vision (Von Helmholtz, 1867; Leopold and Logothetis, 1999; Mitchell et al., 2004). Studying the perception of ambiguous stimuli can therefore further our understanding of organizational principles inherent to all vision.

A more specific aspect of perception whose study has recently caused a surge in ambiguity research, is the distinction between consciousness and unconsciousness. What does it mean to be conscious of visual input, and how does this relate to brain activity? In attempts to understand the nature of human conscious experience within a scientific context, researchers ask

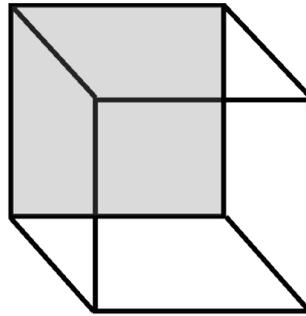


Figure 2: *A Necker cube. The shaded surface can be perceived either as being in front or as being at the back. This well-known ambiguous image was first reported by a geologist inspired by drawings of crystals.*

what brain processes are active while an observer consciously perceives a given stimulus, but inactive when the same stimulus does not reach awareness (Logothetis and Schall, 1989; Crick and Koch, 1998). Similarly, one can ask whether certain psychophysical phenomena, such as aftereffects, selectively occur when a stimulus is consciously perceived, thus inferring what neural events are specifically associated with conscious experience (Tsuchiya and Koch, 2005; Blake et al., 2006; Blake and Fox, 1974). Ambiguous stimuli provide a useful tool for such studies, as the perceptual alternations prompted by ambiguous stimuli constitute a change in an observer's conscious state, without accompanying confounding changes in any other experimental variable.

One final aspect that renders ambiguous stimuli interesting as a subject of scientific study, is the fact that the perceptual competition process that ambiguous stimuli bring about, is quite sensitive to small imbalances, whereas the resulting manifestation of these imbalances (that is, perception of one interpretation over the other) is highly salient. In other words, ambiguous stimuli can be viewed as a type of amplifier that translates subtle shifts in neural activity into major perceptual events. Arguable examples of such use of ambiguous stimuli as a psychophysical amplifier include Alais and Blake (1999); Kanai and Verstraten (2005); Paffen et al. (2005) and Chong and Blake (2005).

Perceptual alternations

A patent characteristic of perception of ambiguous stimuli is the ongoing cycle of perceptual alternations between alternative interpretations. A major fraction of the ambiguity literature focuses on this alternation cycle. Main questions include what factors influence the relative amount of time spent in either percept (e.g. Levelt, 1966; Paffen et al., 2005); what we can infer from naturally occurring variability in the duration between consecutive alternations (e.g. Levelt, 1967; Fox and Herrmann, 1967; Borsellino et al., 1972; Murata et al., 2003), and how transitions between alternative interpretations are organized in the brain (e.g. Wilson et al., 2001; Lee et al., 2004; Knapen et al., 2007).

Chapters 1 and 2 focus on the properties of the alternation cycle to infer characteristics of the neural machinery responsible for alternations, and to compare these between distinct types of ambiguous stimuli.

Onset dominance

A second informative variable is the identity of the first interpretation to dominate after a stimulus is switched on. This variable turns out to be particularly useful for disclosing effects of temporal context, one of the main topics of this thesis (e.g. Nawrot and Blake, 1989; Long et al., 1992; Long and Olszweski, 1999; Orbach et al., 1963; Leopold et al., 2002). Apparently, the process that determines perception of a newly appearing ambiguous stimulus is exceptionally sensitive to slight imbalances, including those due to stimulus history. Onset dominance, in other words, is the most obvious demonstration of the power of ambiguous stimuli to act as a 'psychophysical amplifier'. This helps reveal effects of temporal context, but many other effects as well. For instance, perception of an appearing stimulus also exhibits a strong influence of stationary percept bias (Carter and Cavanagh, 2007), attention (Chong and Blake, 2005; Mitchell et al., 2004) and stimulus imbalances (e.g. interocular contrast differences in binocular rivalry; Kamphuisen et al., 2007).

Chapters 3 through 7 of this thesis use onset dominance to investigate how the visual system uses traces left by prior visual events to modify its processing of present visual input.

Chapter 1

Distributions of alternation rates in various forms of bistable perception

Published as: Brascamp, J. W., R. van Ee, W. R. Pestman and A. V. van den Berg (2005) Distributions of alternation rates in various forms of bistable perception. *J Vis*, 5(4): 287-98

When I first started work in Utrecht, Bert and Raymond decided a good way for me to get acquainted with things would be for me to do some analyses on data that Raymond had collected previously. In the meantime I would have an opportunity to read literature, learn to use software, and so on. I think the basic idea was to try and compare distributions of percept durations in Raymond's perceptual ambiguity data with distributions of saccade latencies collected elsewhere. Analyzing the perceptual ambiguity data we ran into some findings about gamma distributions that we thought were interesting. We decided to write a paper about those, and saccades never really became part of my project. One of the things that struck me when I read this text over as a thesis chapter, is perhaps somewhat of an imbalance between novelty and technical detail. This kind of loss of perspective seems to be a common factor in many beginning PhD students. I recall it was one of our referees who suggested we broaden the appeal of the work by expanding on the comparison between different kinds of ambiguous stimuli. It is funny that in hindsight, I do not care very much about the gamma distributions anymore, but I do think the inter-stimulus comparison is still interesting.

Abstract

Studying the temporal dynamics of bistable perception can be useful for understanding neural mechanisms underlying the phenomenon. We take a closer look at those temporal dynamics, using data from four different ambiguous stimuli. We focus our analyses on two recurrent themes in bistable perception literature. First, we address the question whether percept durations follow a gamma distribution, as is commonly assumed. We conclude that this assumption is not justified by the gamma distribution's approximate resemblance to distributions of percept durations. We instead present two straightforward distributions of reciprocal percept durations (i.e. rates), that both surpass the classic gamma distribution in terms of resemblance to empirical data. Second, we compare the distributions arising from binocular rivalry with those from other forms of bistable perception. Parallels in temporal dynamics between those classes of stimuli can be an indication of a similar neural basis, but have not yet been studied in detail. Our results demonstrate that the distributions arising from binocular rivalry and other forms of bistable perception are indeed similar up to a high level of detail.

1.1 Introduction

When subjects are shown a visual stimulus that has two distinct interpretations, those interpretations will take turns reaching awareness; a phenomenon called bistable perception (for examples, see Figure 1.1). Bistable perception has over the years been studied by numerous researchers, being an interesting phenomenon in its own right, but also because the partial decoupling of stimulus and percept might help to gain insight into the relation between awareness and brain function. In spite of these efforts, the neural mechanisms causing bistability are still under debate. One way to gain insight into these mechanisms is by studying the temporal dynamics of perceptual alternations. For a subject experiencing bistable perception, the duration of a percept commonly varies widely from one to the next, so that speaking of a mean percept duration is of limited use. In contrast to this unpredictability, one of the most reproducible aspects of bistable perception is the shape of percept duration distributions. Invariably, these distributions are unimodal with a skew toward high durations (like the top curves in Figure 1.2). These distinctly shaped distributions may provide clues toward understanding the neural processes involved.

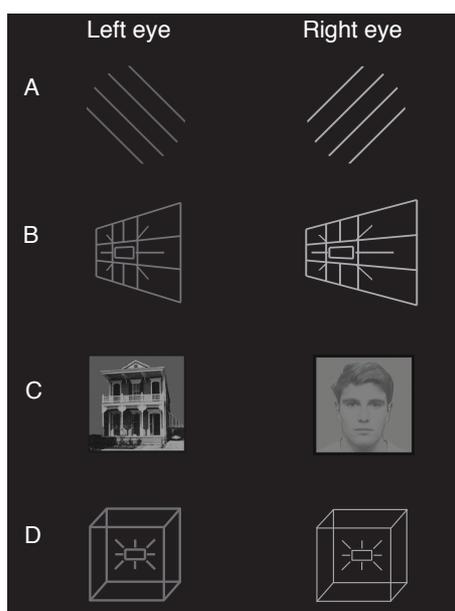


Figure 1.1: Left and right eye components of the four stimuli of which bistable perception was examined: orthogonal gratings (A), a bistable slant stimulus (B), a house-face stimulus (C), and a Necker cube (D). Images are rendered in black and white, but the actual stimuli were green and red. See text for details. Adapted from Van Ee (2005).

In this paper, we investigate the temporal dynamics of bistable perception, using data that was previously collected using four different ambiguous stimuli (1.1). We focus our analyses on two recurrent themes in the literature on the subject. First, we address the question whether distributions of percept durations follow a gamma distribution. Second, we compare distributions over different stimuli, since parallels and differences in temporal dynamics between stimuli can serve as an indication for parallels and differences in underlying neural mechanisms.

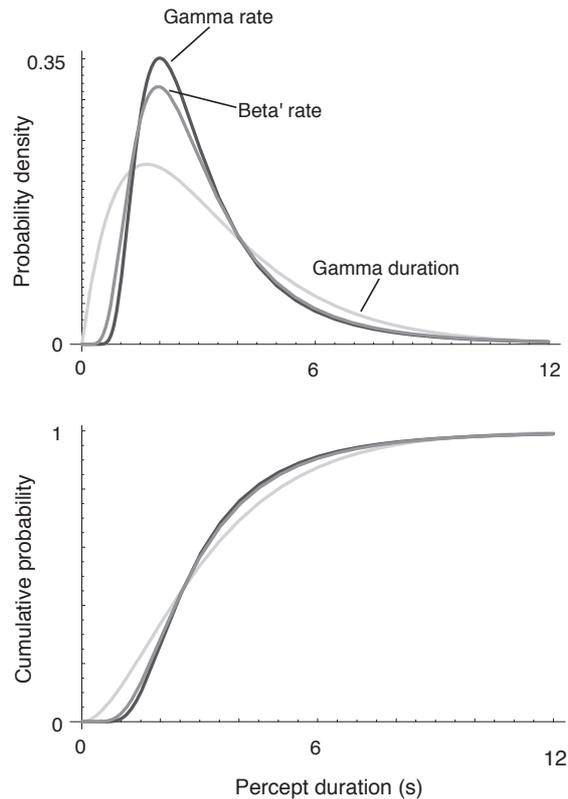


Figure 1.2: Examples of percept duration distributions predicted by the three models we compare in this paper. Top: PDFs of percept durations, under the assumption of gamma distributed alternation rates (darkest shade), beta' distributed alternation rates (intermediate shade), and gamma distributed percept durations (lightest shade). To facilitate comparison, the parameters were chosen such that both mean and variance are the same for the three curves. Bottom: the corresponding CDFs. The PDFs show that all three theoretical distributions satisfy the basic criteria of unimodality and rightward skew, but that the classic gamma duration distribution differs from the other two in having a shorter initial phase of low slope, and a lower peak. The CDFs on the other hand, emphasise that any difference between the three distributions is modest in light of the overall similarity.

1.1.1 Gamma distributed percept durations?

An interesting aspect in the study of distributions of percept durations, is their resemblance to known statistical distributions, since these are often associated with clearly defined mechanisms. One such theoretical distribution with the desired unimodal, right-skewed shape is the gamma distribution, introduced in this context by Levelt (1967). Its probability density function is given by

$$f(t | k, \lambda) = \frac{1}{\lambda^k \Gamma(k)} t^{k-1} e^{-\frac{t}{\lambda}} \quad (1.1)$$

By definition, $\Gamma(n)$ is the canonical continuous extension of $(n - 1)!$, which itself is only defined for natural n . The parameters k and λ in the equation are called the shape parameter and the scale parameter, respectively.

Among other similarly shaped distributions, the gamma distribution is noteworthy because it has over the last decades become the standard for performing fits to empirical distributions of percept durations (e.g. Borsellino et al., 1972; Kovács et al., 1996; Leopold and Logothetis, 1996; Murata et al., 2003; Walker, 1975).

Four papers that statistically analyzed the gamma distribution's fit performance, did not find a particularly good fit to empirical data. The authors of two such papers (Borsellino et al., 1972; De Marco et al., 1977) judged gamma distributions to fit their data acceptably well for their purposes, but their analyses leave considerable room for doubt (Borsellino and coworkers stated that around 15 to 30% of their gamma fits had a chi² probability lower than 1%, while De Marco and colleagues mentioned two alternative theoretical distributions to fit equally well as the gamma distribution, although less favorable in the light of parsimony.) The two remaining papers (Cogan, 1973; Zhou et al., 2004) showed an unacceptable fit quality for the gamma distribution: in both cases more than half of the fitted distributions should be rejected at the 5% significance level. As a point in favor of the gamma distribution, it should be mentioned that on the basis of the above studies, one can not identify an alternative distribution with better fit performance. Although Zhou et al. (2004) did show the lognormal distribution to fit their data better than the gamma distribution, Cogan (1973) rejected the lognormal distribution as an acceptable fit to her data (note that lognormal distributed percept durations were also proposed by Lehky (1995)).

What makes the question whether the gamma distribution is appropriate particularly interesting, is the neural model that is associated with it. As Levelt (1967) pointed out, one process known to lead to gamma distributed latencies between events, is one in which every event occurs after a fixed number of consecutive ticks from a 'Poisson clock', i.e. randomly occurring ticks. Therefore, he proposed that "the summative effect of [...] successive spikes from the recessive [percept] is necessary and sufficient to re-establish dominance for that [percept]", without defining the exact nature of these Poisson 'spikes' or ticks. In the case of a Poisson clock, λ in the above equation is the basic duration between two ticks, while the value of the shape parameter k is the number of ticks causing a perceptual alternation.

One of the goals of this paper is to use our data to compare the fit quality of the gamma distribution and alternative distributions. The question is: what alternative distributions are viable candidates? We believe a step towards finding such distributions could be to switch the focus of attention away from distributions of percept durations, towards distributions of reciprocal percept durations, or rates. Since distributions of alternation rates are just as characteristic for bistable perception as those of percept durations, this provides no a priori reason

to make fits to one or the other. We believe, however, that distributions of alternation rates can increase our understanding of underlying processes. This idea stems from saccadic (i.e. rapid eye movement) search literature, in which there is some tradition of using such an approach. There are several parallels between saccadic search and bistable perception, one of them being the timing of saccades: the duration distribution of periods that a subject fixates between saccades, has a shape that is very similar to that of distributions of percept durations. Much research has been aimed at linking saccadic latency distributions to the neural mechanisms involved (for an overview, see Smith and Ratcliff, 2004). However, Carpenter (1981) proposed that these mechanisms might be more directly represented in distributions of reciprocal latencies in stead of latencies, and since then various studies have made use of this idea (Carpenter and Williams, 1995; Van Loon et al., 2002; Reddi and Carpenter, 2000) The key element underlying these authors' point of view is the notion of a decision network, that in the presence of an extrafoveal target accumulates information about that target, and causes a saccade to be made toward it as soon as this rising decision signal exceeds a certain threshold level. If one now assumes this threshold level to be fixed over trials, then variations in the rate of information accumulation are directly reflected in the reciprocal latency, or rate, of saccades. What we propose in analogy regarding bistable perception, is that distributions of *perceptual alternation rates* (i.e. reciprocal percept durations) may be more directly related to underlying neural processes than distributions of percept durations, and could therefore give insight into neural processes. This especially makes sense in light of the idea that the numerous phenomenological similarities between saccadic search and bistable perception are not merely coincidental, but instead reflect a more fundamental link between the two (Leopold and Logothetis, 1999; Sheinberg and Logothetis, 1997). Completing the analogy between rate models in saccadic search and in bistable perception, we suggest a decision signal that starts rising at the beginning of a dominance period of percept *A*, and causes an alternation to percept *B* as soon as it reaches its threshold, so that characteristics of the rate of information accumulation may be reflected in distributions of perceptual alternation rates.

In this paper we pit two different rate distributions against the standard gamma fit. First, we explore the possibility that alternation rates have a gamma distribution. We will hereafter refer to these distributions as 'gamma rate' distributions, not to be confused with the classical 'gamma duration' distributions. Since Van Loon et al. (2002) showed gamma rate fits to be applicable to saccade data, there is some reason to expect an acceptable fit to our data too.

The second distribution we fit to our rate data is the beta' (pronounce: "beta prime") distribution, a two-parameter distribution related to the gamma distribution¹, and similar in shape. Its probability density function is given by:

$$f(r | S, F) = \frac{\Gamma(S + F)}{\Gamma(S) + \Gamma(F)} \frac{r^{S+1}}{(1 + r)^{S+F}} \quad (1.2)$$

This distribution arises from a model that assumes two rising decision signals in stead of one: one in favour of maintaining the present percept; the other in favour of switching. These signals both increase towards a threshold over time, but compete for the same resource, so that at any given time only one of them can gain an increment (i.e. the model is based on a Bernoulli

¹The gamma distribution and the beta' distribution are linked by the fact that dividing two gamma distributions of equal scale parameter produces a beta' distribution. The beta distribution (a scaled version of the beta' distribution) previously made its appearance in bistable perception literature when Borsellino et al. (1972) made use of this feature in their analysis of supposedly gamma distributed percept durations.

process). In terms of this model, the parameters S and F of the beta' distribution are determined by the thresholds for switching percepts and maintaining the current one, respectively. The parameter r is proportional to the observed alternation rate via a constant R : $\text{rate} = Rr$. In the present paper, we fix R at 1 s^{-1} , performing our fits with a two-parameter distribution. The main reason for this is that the full three-parameter version caused divergence in our fitting algorithm, while much of the distribution's flexibility can already be achieved by varying only the two remaining parameters. It should be kept in mind however, that removing this constraint on the third parameter would provide room for an improved fit quality compared to what we present here. A complete derivation of the beta' distribution from the model assumptions is given by Van den Berg and van Loon (2005).

Figure 2.1 gives an impression of the overall shape of the probability density functions (PDFs) and the cumulative distribution functions (CDFs) associated with the gamma duration, the gamma rate and the beta' rate model. It shows both the overall similarity between them, and the more subtle aspects in which they differ.

1.1.2 Inter-stimulus comparison

There are many stimuli that lead to bistable perception. One particularly well-studied class of stimuli are the ones causing binocular rivalry. Binocular rivalry is elicited by presenting dissimilar images to corresponding areas of the two retinas, resulting in a percept that alternates between the two images (e.g. stimulus A and C in Figure 1.1). In recent literature on bistable perception, much attention has been aimed towards answering the question whether binocular rivalry on the one hand, and perceptual rivalry of ambiguous figures (e.g. stimulus B and D in Figure 1.1) on the other, reflect distinct or similar neural mechanisms (Blake and Logothetis, 2002; Blake, 2001; Tong, 2001). Advocates of the former point of view argue that in binocular rivalry, alternations arise as the two monocular channels converge in the primary visual cortex, as a direct result of the incongruence between the retinal images. Researchers favouring the latter position on the other hand, claim that competition in binocular rivalry is between stimulus representations rather than eyes, generally locating it in the higher visual areas and placing it in the same category as reversals in ambiguous figure perception. One argument in this discussion is that percept durations have similarly shaped distributions in both binocular rivalry and ambiguous figure rivalry (e.g. Carter and Pettigrew, 2003; Logothetis et al., 1996; Logothetis, 1998; Lumer et al., 1998). However, at present this argument is based on a broad resemblance and a fair fit to the gamma duration distribution. Our second goal in this paper therefore is to test whether the claim holds in the face of a detailed comparison of distributions across stimuli.

1.2 Data collection

A detailed account of the conditions in which data were gathered, was given by Van Ee et al. (2004). In short, a total of six subjects were tested observing four different stimuli (Figure 1.1):

1. Orthogonal gratings. One way to elicit bistable perception is by presenting dissimilar images to corresponding areas of the two retinas, resulting in a percept that alternates between the two images. Orthogonal grating rivalry is a classical form of such binocular

rivalry, where orthogonal gratings are presented to the two eyes. In our case the stimulus consisted of four parallel lines at 45 degrees with the vertical, that had orthogonal orientations in the two eyes.

2. Bistable slant stimulus. This is a stimulus first described by Van Ee et al. (2002), consisting of a trapezoid that is viewed binocularly. Bistability arises from the fact that linear perspective and binocular disparity specify opposite slants (e.g. in Figure 1.1B perspective information corresponds to a rectangle seen with the right side in front, whereas uncrossed fusing of the two images would cause disparity information to signal a trapezoid seen with the left side in front).
3. House-face stimulus. This is a stimulus developed by Tong et al. (1998). Like the orthogonal gratings described above, it gives rise to a form of binocular rivalry. In house-face rivalry however, the conflicting images are not orthogonal lines but pictures of a house and a face.
4. Necker cube. This is arguably the best known ambiguous stimulus. It is an image that can be interpreted as a cube seen from either of two viewpoints.

All stimuli were depicted in red and green and viewed through red and green anaglyph glasses to separate the left and right eye image where necessary. Even though there is no need for such separation in Necker cube rivalry, anaglyph glasses were used here as well for consistency. During three-minute trials, a stimulus was viewed on a computer screen, leading to an unstable percept alternating back and forth between the two alternatives. Alternations were reported by the subjects using button presses. Stimuli covered an angular width of 1.2 degrees, and were surrounded by a background of fusible squares to maintain alignment of the eyes. Subjects were instructed to maintain central fixation during all trials, making use of a sunburst figure in the cases of the slant and the Necker cube stimulus.

It was established that drift in the alternation rates was restricted to the first 30 seconds of a session, during which it never exceeded a few percents of the mean rate, so our data was reasonably stable both across small data chunks and across experimental sessions.

Because the data were originally gathered, among other reasons, for investigating the influence of voluntary control on the dynamics of bistable perception, subjects were given four different instructions: either to view a stimulus in a natural way without attempting to control the alternation rate, or to try to keep one or the other of the two percepts for as great a fraction of the time as possible, or to speed up the alternation rate as much as they could. For each experiment, we treated the durations of the one percept (e.g. the house) separately from the durations of the other percept (e.g. the face). All these variables taken together, led to $6 \text{ (subjects)} \times 4 \text{ (stimuli)} \times 4 \text{ (instructions)} \times 2 \text{ (percepts)} = 192$ distributions. After removal of the 2% largest and 2% smallest values, the distributions had an average size of 243 data points.

1.3 Gamma distributed percept durations?

1.3.1 Methods

We determined the best fitting gamma and beta' distributions to our empirical data in two different ways. First, we employed the classical least squares method, minimizing the sum of squared residuals SSE . For the fits we obtained in this manner, we calculated the fit quality as

$SSE/(n - 2)$ (where n is the number of data points in a given distribution), but also using the Kolmogorov-Smirnov test for goodness of fit. The Kolmogorov-Smirnov test is based on the largest overall deviation between empirical and fitted distribution, and the associated probability pKS . In addition to these least squares fits, we performed maximum likelihood fits on the same data, providing an alternative estimate of the best fitting parameter values, as well as of the fit quality: the likelihood L . Note that likelihood fits by definition involve PDFs, whereas for the least squares fits we used continuity corrected CDFs, for reasons of robustness and objectivity (contrary to the PDF, the CDF does not involve an arbitrary bin size). All fitting algorithms were implemented in the software package Scilab (<http://scilabsoft.inria.fr/>).

1.3.2 Results

Figure 1.3 shows examples of the three distributions fitted to an empirical distribution. As expected, at first glance all seem in fair agreement with the data, but there are some deviations. For the gamma duration fit, the fitting algorithm's tendency to align model and experimental data on the CDF's steeply ascending flank, results in the fit overshooting the empirical CDF at both inflection points. The two fits to rate data on the other hand, frequently undershoot the empirical distributions at the inflection points, however these deviations seem smaller and less consistent.

This impression of the rate distributions fitting our data better, is confirmed by quantitative analysis. Figure 1.4 displays the fit quality of the three distributions, in terms of three different measures. It can be seen that, irrespective of the specific fitting algorithm or goodness measure used, both rate fits perform better than the gamma duration fit, and the gamma rate fit slightly outperforms the beta' rate fit.

Figure 1.4 clearly shows how the fit qualities of the three distributions relate to each other, but it does not tell us whether particular fits should be rejected or not. To assess that question, we visualised the fit qualities in a different way in Figure 1.5, plotting the fraction of fits that would be accepted, as a function of the critical Kolmogorov-Smirnov p level that one chooses. An advantage of this presentation method is that it summarizes the acceptance of fits, without choosing the critical pKS level in advance. This shows that both rate distributions have a higher acceptance than the gamma duration distribution at any critical p level. Should one for instance choose a level of 0.1 (vertical line in Figure 1.5), fractions of 0.92, 0.86 and 0.69 of the gamma rate, beta' rate and gamma duration fits pass, respectively.

We can conclude that this analysis demonstrates a better fit for the two rate distributions than for the gamma duration distribution.

1.3.3 Control experiment for motor bias

The shown difference in fit quality, although consistent, is based on fairly subtle differences between distributions (Figure 1.3). We should therefore be wary of any bias introduced by our experimental design. Particularly, since we relied on key presses to infer our subjects' perceptual state, it is important to rule out the possibility that our results reflect characteristics of the motor system rather than the perceptual system. It is conceivable that subjects do experience percepts that have gamma distributed durations, but that systematic variations in their manual response time cause the two rate distributions to fit the reported distributions better after all.

To investigate this possibility, we performed a control experiment determining the relation

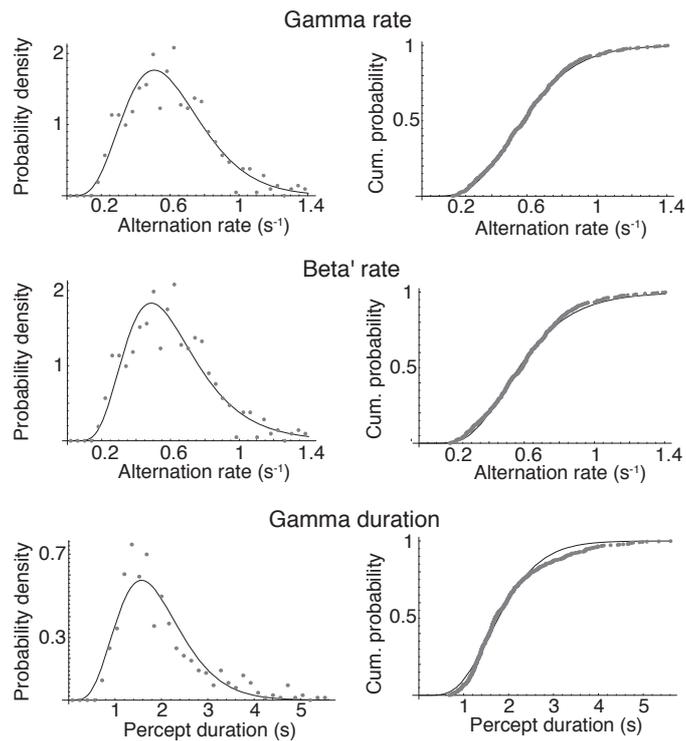


Figure 1.3: *Gamma rate, beta' rate and gamma duration distributions (drawn lines), fitted to an example data set (dots). Probability density functions (left) are shown for illustration, but cumulative distribution functions (right) were used for fitting. Fitted curves' deviations from the empirical cumulative distributions, especially evident for the gamma duration fit, are discussed in the text. The number of points of this data set is 541; the PDF bin size was chosen such that 35 bins span the entire data range. The pKS values associated with these fits, purposefully low for illustrating their weaknesses, are 0.06, 0.009 and 0.0002 for the gamma rate, the beta' rate and the gamma duration distribution, respectively. Note that the top four graphs have 'alternation rate' on the horizontal axis, where the bottom two have 'percept duration'.*

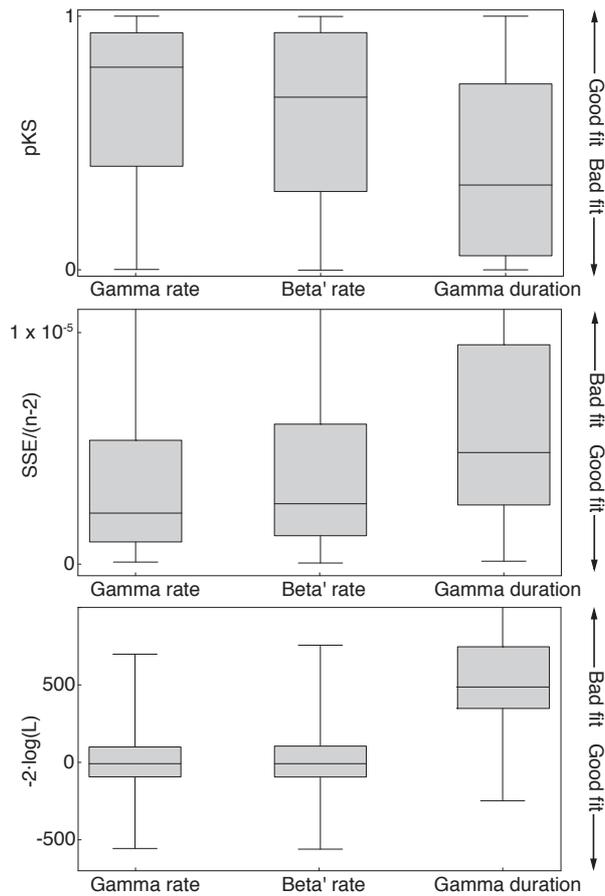


Figure 1.4: Quality of the fits to our 192 empirical distributions, measured using the Kolmogorov-Smirnov test (top), the sum of squared residuals SSE (middle) and the likelihood L (bottom). Each box with whiskers summarizes 192 values, displaying median, 25% and 75% quantile and extreme values (but note that these fall outside the plot range in some cases). All three figures indicate a low fit quality for the gamma duration distribution as compared to the rate distributions, and a slightly better fit for the gamma rate than for the beta' rate distribution.

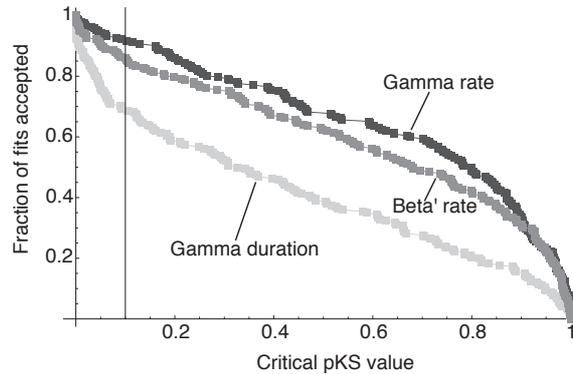


Figure 1.5: *Fit quality as measured by the Kolmogorov-Smirnov p value. The fraction of fits that would be accepted is plotted as a function of the critical pKS level. Irrespective of the particular critical p level that is chosen, the gamma duration distribution, represented here by the lightest shade, clearly fits the data less well than the beta' rate distribution (intermediate shade) and the gamma rate distribution (darkest shade), which fits best. The vertical line at $p=0.1$ crosses the three curves at 0.69, 0.86 and 0.92: the fractions of accepted fits in case one tests at the 10% level.*

between perceived and reported durations. In this experiment we showed subjects a stimulus that resembled our orthogonal grating stimulus, but that did not cause bistable perception. Instead, 'perceptual alternations' were induced by presenting the green and the red set of lines in an alternating fashion on the computer screen, and subjects were asked to press a key when they saw the stimulus change. Importantly, we drew the durations of these successive presentations from a gamma distribution, so if all was well it would be the gamma duration distribution that fitted the subjects' reported distributions best.

We did find some deviations between presented and reported distributions, but these were too small to explain our findings. This is illustrated in Figure 1.6, which displays the results of one experiment in which we presented six subjects with durations drawn from a gamma distribution (200 durations with a mean of 1.2 s and a standard deviation of 0.7 s). The left panel shows, in the same way as Figure 1.5, the quality of fits to the distributions we presented to our subjects. Not surprisingly, the gamma duration distribution fits best. The right panel displays fit quality to the reported distributions. Even though there seems to be some improvement in the rate fits, the gamma duration fit has retained its superior quality, and the gamma rate fit (the best fit to our actual data) fits worst by far. Our results can therefore not be explained by a bias introduced by the motor response, but should instead be interpreted in terms of the perceptual process itself.

1.4 Inter-stimulus comparison

1.4.1 Methods

As a first step in comparing the distributions between stimuli, we re-examined the previously determined acceptance levels of our fits, now treating the results of the four stimuli separately. We were particularly alert to differences in the results for binocular rivalry versus other forms

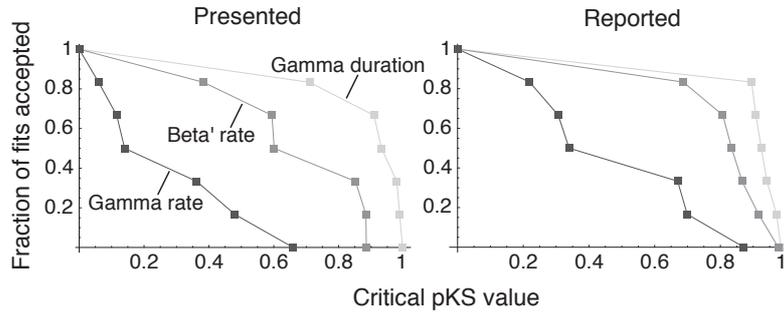


Figure 1.6: In this control experiment we presented our subjects with a sequence of two alternating images that resembled the two percepts of binocular grating rivalry, and asked them to report every alternation by means of a button press. By drawing the durations between successive alternations from a gamma distribution, we mimicked the time course of a Poisson driven bistable perception experiment. The left panel shows, in the manner introduced in Figure 1.5, the quality of the gamma rate, the beta' rate and the gamma duration fits to the distributions we presented to our subjects in this way. The gamma duration distribution fits best, as was to be expected because we drew the presentation durations from a gamma distribution. In the panel on the right, showing fit quality to the distributions of alternations that our subjects reported, this superior fit quality is preserved. This experiment indicates that the inferior gamma duration fit in our real experiments, is probably not due to a bias introduced by the motor response.

of rivalry, that might reflect differences in neural basis.

We performed an additional analysis in which we compared the distributions' shapes in a more direct way, by looking at the nature of the deviations between empirical and fitted distributions. For this purpose, we calculated for all three fits (gamma duration, gamma rate, beta' rate) the mean fit residual as a function of position along the CDF, averaging over all 48 distributions associated with a particular stimulus. Our procedure is illustrated in Figure 1.7.

1.4.2 Results

Figure 1.8 depicts the acceptance of fits in the same way as Figure 1.5, except that now the results of the four stimuli were treated separately. This procedure reveals that the isolated results for the bistable slant stimulus (top right) deviate from the general pattern, showing hardly any difference in acceptance between the three distributions. The results for the remaining stimuli however, do conform to the overall finding that the rate distributions outperform the gamma duration distribution. Therefore, contrary to what one might predict, no dichotomy between the binocular rivalry stimuli (grating and house face) and the other stimuli can be made on the basis of this figure. On the contrary, any separation would be between the isolated results for the bistable slant stimulus, and those for the other stimuli.

The results of our second analysis, comparing fit residuals between stimuli, are depicted in Figure 1.9. This figure agrees with our preliminary impression that gamma duration fits deviate from the empirical distributions in a highly consistent way (bottom left panel; see also figure caption). Since the gamma duration distribution has a fairly poor fit quality, it is not so surprising to see that this pattern of deviations is similar over all four stimuli: this simply means that any subtle differences that may exist between the four groups of distributions, are overshadowed by the gamma duration distribution's large structural fit error.

In terms of the inter-stimulus comparison, the more interesting panels are the top two,

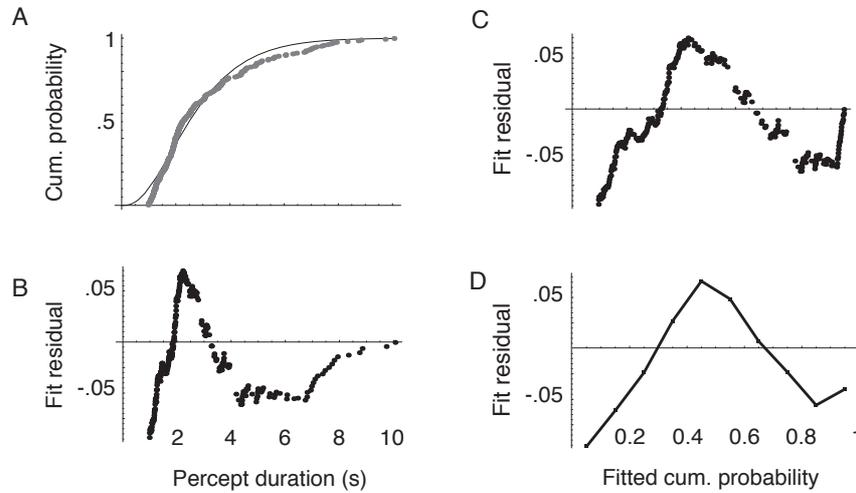


Figure 1.7: Analyzing fit residuals. In order to compare the distributions' shapes between stimuli, we analysed the deviations between our empirical and fitted distributions. A. An example of an empirical CDF of percept durations (gray dots), and the corresponding gamma duration fit (drawn line). B. Residuals of the same fit, as a function of percept duration. The values plotted here are simply the differences between the empirical and the fitted distribution depicted in A. C. For each distribution, panel B may span a different domain of percept durations, preventing us from averaging residuals over distributions. Therefore, in panel C we standardized the x-axis in panel B by replacing each x-value with the corresponding value of the fitted CDF. Note that in statistics, plots such as these, depicting residuals as a function of cumulative probability, are referred to as 'detrended Q-Q plots'. D. The same data as in panel C, after averaging over bins spanning 0.1 unit on the x-axis. After this final step, residuals could be combined over fits.

depicting residuals for the two (better fitting) rate distributions. The pattern of deviations is comparable for all but the slant stimulus (second lightest shade). The remaining three stimuli, therefore, are not only similar in their fit quality to the tested distributions, but also in the pattern of fit residuals. This again provides no evidence for a dichotomy between the binocular rivalry stimuli and the remaining stimuli, but rather between the slant stimulus and the other three stimuli.

1.5 Discussion

Previous work has shown the gamma duration distribution to be surpassed in fit quality by a lognormal distribution (Zhou et al., 2004). To find out how these findings relate to our results, we performed a lognormal fit to our duration distributions as well, and we can confirm an improved fit quality compared to the gamma duration distribution. However, this improvement was modest, and not as large as what we have shown for the two rate distributions. Using, as we did earlier, the fraction of accepted fits at a 10% significance level as a brief indication of fit quality, the lognormal distribution, with a fraction of 0.78, falls right in between the values of the gamma duration distribution (0.69) and the two rate distributions (0.86 and 0.92). With regard to our data set these rate distributions should therefore be preferred to the lognormal

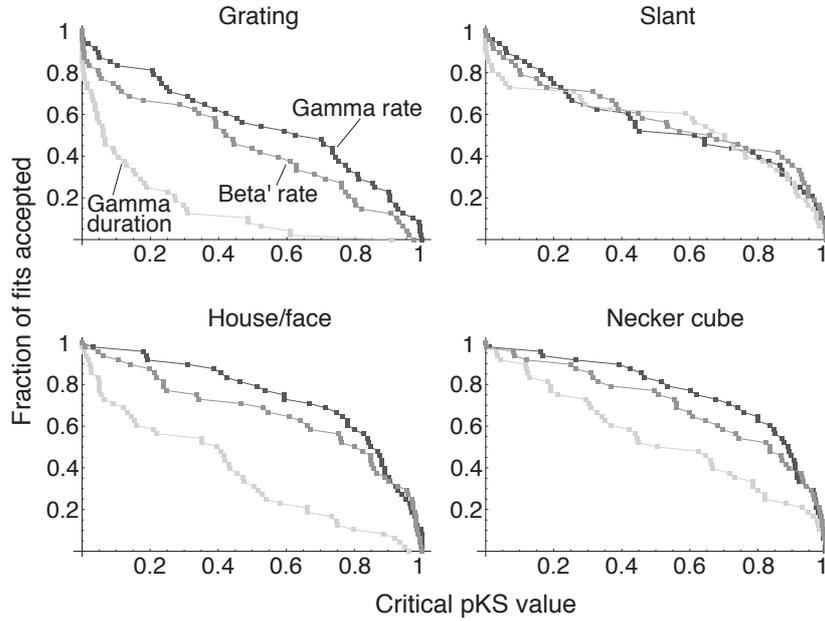


Figure 1.8: Fit quality as measured by the Kolmogorov-Smirnov p value, separated by stimulus. The fraction of fits that would be accepted is plotted as a function of the critical pKS level. For grating, house/face and Necker cube rivalry, the results conform to the overall pattern, with the gamma duration distribution (lightest shade) fitting our data less well than the beta' rate distribution (intermediate shade) and the gamma rate distribution (darkest shade). The results for the bistable slant stimulus form an exception, with no clear difference in quality between any of the three fits. There is however no indication of a dichotomy between the binocular rivalry stimuli (grating and house/face) and the other two. All curves are built up of 48 data points.

distribution.

We consider it beyond the scope of this paper to present an exhaustive comparison of fit qualities of the numerous distributions that have over the years been proposed in this context. As mentioned previously, many can be discarded, based on the fact that other studies showed their fit qualities not to exceed that of the gamma duration distribution. One last distribution that we do wish to discuss here however, is the 'dual Gaussian rate distribution'. Since this is a rate distribution of some renown in the field of saccadic search, it seems appropriate to include it here. When using this distribution, which is actually a composite of two separate distributions, in the conventional way, one assumes that a part A of an empirical rate distribution, the part containing low rates, can be described by one truncated Gaussian distribution Φ_A , and the remaining part B by another one Φ_B (Carpenter and Williams, 1995; Reddi and Carpenter, 2000). The raw data is plotted on probability paper, so that the distinction between data sections A and B can be made by eye, after which the parameters of Φ_A and Φ_B can be estimated (see Figure 1.10, left panel). If we want to compare the fit quality of the dual Gaussian distribution to that of the ones we investigated, this approach will not do, mainly because it involves visual inspection. Since it does seem that the combination of two Gaussians describes rates of saccadic eye movements well, we adjusted the dual Gaussian distribution in order to be able to investigate its fit quality to our data on alternation rates. According to this adaptation, the

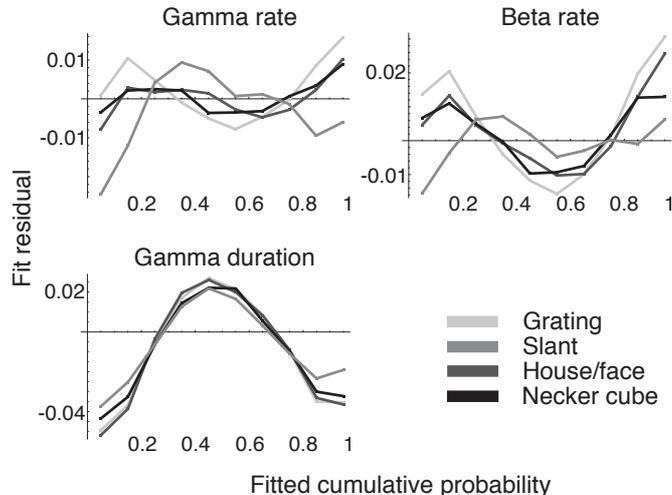


Figure 1.9: Average residuals of our three fits, separated by stimulus. Each curve represents the binned residuals averaged over 48 distributions, as a function of fitted CDF value (see 1.7). A positive residual means that the empirical data overshoots the fit; a negative value means it undershoots it. Standard deviations on the points in these graphs, not shown for the sake of clarity, are on the order of 0.02. The clearest panel is the one depicting the results for the gamma duration fits (bottom left). For all four stimuli, the curves show the same pattern of the empirical data initially undershooting the fit, then overshooting it near the mean, and again undershooting it towards the end. The two rate distributions (top panels) give rise to more noisy residuals, as expected in light of their superior fit quality. Still, for both rate fits the pattern of residuals is fairly consistent for all but the slant stimulus. We can therefore state that, in agreement with our previous analysis, the distributions associated with the bistable slant stimulus are slightly different from those associated with the other stimuli.

probability density function at rate r is given by:

$$f(r \mid \mu_1, \sigma_1, \mu_2, \sigma_2, a) = a \cdot \phi(r \mid \mu_1, \sigma_1) + (1 - a) \cdot \phi(r \mid \mu_2, \sigma_2) \quad (1.3)$$

Here, ϕ is a Gaussian distribution. Note that, instead of fitting two truncated Gaussians to separate parts of the empirical distribution, we fit a mix of two Gaussians to the entire distribution, and introduce a free scaling parameter a to ensure a total cumulative probability of 1. This mixed Gaussian distribution was introduced over a century ago (Pearson, 1894), and similar mixed distributions are presently used in many fields of research (McLachlan and Peel, 2000). This dual Gaussian model produced excellent fits to our data (all were accepted at a critical pKS level of 0.1, after removal of non-convergent fits), but a direct comparison to the other models is not possible since these contain only two free parameters, instead of five (the dual Gaussian distribution is over-parameterised, as witnessed by the convergence problems we experienced). In addition, it is worth mentioning that our adapted dual Gaussian fits often reach quite different results than the conventional ones (see Figure 1.10). A final conceptual problem one might have with these fits is that, both in the original form and in our adaptation, they allow for the occurrence of negative alternation rates.

As discussed in the introduction section, the gamma duration distribution is associated with a so-called ‘Poisson clock’ model. Since we showed a relatively poor fit quality for this distribution, it is important to address a recent paper by Murata et al. (2003), presenting data that seem to provide strong support for this model. According to the Poisson model, the gamma duration

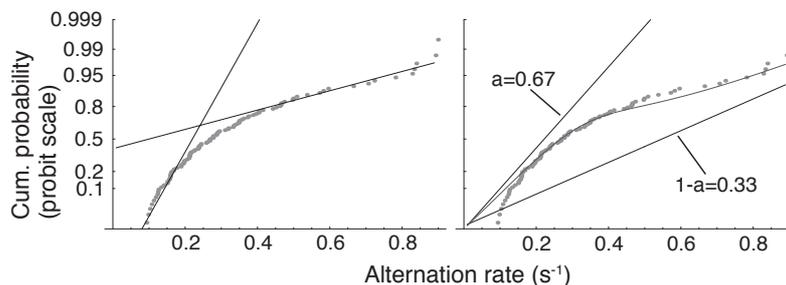


Figure 1.10: Dual Gaussian fits to one of our distributions. Both plots depict cumulative probability as a function of alternation rate, on probit scale. Left: the conventional method. The data (dots) are divided into two sections, each of which can be described by a Gaussian distribution: the two straight lines. This method does not readily lend itself for a comparison of fit quality. Right: an adapted method, applied to the same data. The two straight lines depict two Gaussian distributions that together comprise the fitted distribution (curved line). The relative contributions of the two Gaussians, determined by weighing factor a , are displayed next to the corresponding lines.

distribution's shape parameter k reflects the number of Poisson ticks after which a perceptual transition takes place (see Introduction). On the basis of the model, one might therefore predict that this shape parameter should be a natural number. What Murata and coworkers showed, is that the gamma duration distributions they fitted to their data indeed had shape parameters that grouped around natural numbers. It would seem that this result can not be explained unless by accepting the idea of a Poisson clock and the associated gamma duration distribution, but Murata and coworkers did not support their claim with a statistical analysis. One way of statistically testing for grouping around natural numbers, using estimated shape parameters from gamma duration fits, would be to subtract the nearest natural number from each of these estimates. This operation would produce a distribution of residual shape parameters ranging from -0.5 to 0.5 , that should be peaked around 0 in case of natural k -values. One can statistically test for the presence of such a peak using a standard test. We have performed such an analysis on our data without finding any evidence for natural shape parameters, however Monte Carlo simulations show that a data set as large as ours (192 distributions of 243 points on average), does not provide enough statistical power to demonstrate natural k -values in this way, even if they are present. Similar simulations show that Murata and colleagues' data set (227 distributions of 350 data points), might be just large enough to successfully perform such an analysis, but it should be noted that these simulations were performed without adding any noise.

The distributions associated with the bistable slant stimulus, seem to differ from the ones produced using other stimuli. They are exceptional in that they fit the gamma duration distribution just as well as the two rate distributions, and also in their pattern of fit residuals. Since another notable characteristic of the bistable slant stimulus is the relatively low rate of perceptual alternations it elicits, it is tempting to associate both features. As previously mentioned, and illustrated by Figure 1.2, the gamma duration distribution mainly differs from the other two in the sense that there is less probability mass near short percept durations (or near high rates). Our findings might therefore be in agreement with some extension of the classic Poisson model, in which the addition of some low-pass filter results in the Poisson-generated gamma distributions being deformed at shorter percept durations. This idea would explain the improved gamma duration fit for the slower alternating slant stimulus, and seems somewhat in

agreement with Figure 1.9, which shows that the difference in fit residuals between the slant stimulus and the other stimuli, is partly concentrated in the region of short durations (high rates), although also at the other extreme.

Our analyses did not reveal any difference between binocular rivalry and other forms of rivalry. In fact, differences among non binocular rivalry stimuli (Necker cube and slant) were larger than between binocular rivalry and non binocular rivalry stimuli. This does not necessarily indicate a shared neural basis for binocular rivalry and other forms of bistable perception. Differences, other than our analyses would reveal, have previously been identified. For instance, the extent to which observers can consciously influence the alternation process differs between the two kinds of bistable perception (Van Ee et al., 2004; Meng and Tong, 2004). In addition, quoting Blake (2001): '[...] comparable temporal fluctuations may reflect a fundamental property of neural dynamics, but not necessarily a common neural mechanism'. Nevertheless, there are two conclusions that we think are justified. First, the claim of similar temporal dynamics of binocular rivalry and other rivalry, hitherto based on a broad resemblance of the distributions, still holds after detailed analysis. Our results therefore render it more convincing as one of the pieces of converging evidence linking the two phenomena. Second, any model placing binocular and other rivalry in different brain 'modules', should be able to account for the highly similar distributions of alternation rates.

1.6 Conclusion

We show two rate distributions, the beta' rate distribution and the gamma rate distribution, to conform to our bistability data better than the more traditional gamma duration distribution. This, together with the idea that they form a more direct representation of the neural decision process, suggests that rate distributions may be more appropriate than duration distributions for studying the time course of bistable perception. Furthermore, our results demonstrate that the temporal dynamics of binocular rivalry and ambiguous figure perception are similar up to a high level of detail.

Chapter 2

The time course of binocular rivalry implies a crucial role of noise

Published as: Brascamp, J. W., R. van Ee, A. J. Noest, R. H. A. H. Jacobs and A. V. van den Berg
(2006) The time course of binocular rivalry reveals a fundamental role of noise. *J Vis*, 6(11):
1244-56

Among the more widely cited older work on binocular rivalry, is the work that Willem Levelt published in the nineteen sixties on the relation between stimulus contrast and percept dominance durations. I noticed at some point that no such characterization existed of the relation between stimulus contrast and the durations of the other part of the alternation cycle – transition periods when both images are perceived. I set out to provide such a characterization. Because transition periods cannot be measured in isolation, we ended up describing various features of the alternation cycle beside transition durations. One aspect of this work that strikes me as typical of our approach in Utrecht, is that we chose not to center our paper on those empirical observations, but instead on their neural implications. I find this particularly remarkable because I recall spending ages trying to get a model to reproduce our observations, which we felt was necessary for a paper about neural mechanisms. The model work did not seem to be going anywhere, until it occurred to me that if our model would not reproduce the data, I could try to get published models to do it. That way it would be newsworthy, both if they could do it and if they could not. In the end, we could not get things to work with existing models either, so we wrote a paper pointing out that existing models might need revision. Tomas aptly described this process as ‘instead of stubbornly continuing banging your heads into a wall, you took a step back and wrote a paper saying: “Look! There’s a wall there.”’

Abstract

When our two eyes view incongruent images we experience binocular rivalry: an ongoing cycle of dominance periods of either image and transition periods when both are visible. Arguably, two key forces underlying this process are adaptation of, and inhibition between the images' neural representations. Models based on these factors meet the constraints posed by data on dominance periods but these are not very stringent. We extensively studied contrast dependence of dominance and transition durations and of the occurrence of return transitions: occasions when an eye loses and regains dominance without intervening dominance of the other eye. We found that dominance durations and the incidence of return transitions depend similarly on contrast; transition durations show a different dependence. Regarding dominance durations we show that the widely accepted rule known as Levelt's 2nd proposition is only valid in a limited contrast range; outside this range the opposite of the proposition is true. Our data refute current models, based solely on adaptation and inhibition, as these cannot explain the long and reversible transitions that we find. These features indicate that noise is a crucial force in rivalry, frequently dominating the deterministic forces.

2.1 Introduction

When we present two incompatible images to our two eyes binocular fusion gives way to binocular rivalry (e.g. Figure 2.1), a phenomenon characterized by ongoing perceptual alternations between the two images (Blake, 2001; Alais and Blake, 2005), similar to those occurring when viewing ambiguous figures such as the Necker cube. Understanding such bistable phenomena is important for the study of visual awareness. Arguably, they are an evident demonstration of fundamental processes underlying awareness, reflecting an ongoing effort of the visual system to select and reorganize sensory input to form consistent interpretations (Andrews and Purves, 1997; Leopold and Logothetis, 1999). In addition, because of the partial decoupling of stimulus and percept, rivalry is a valuable tool in studies into neural correlates of awareness in isolation from stimulation.

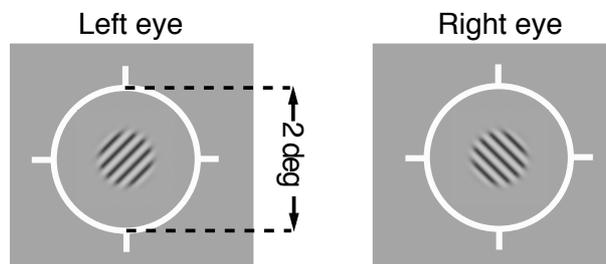


Figure 2.1: Our binocular rivalry stimulus.

According to current ideas binocular rivalry revolves around *cross-inhibition* and slow *self-adaptation* (e.g. Wilson, 2005). That is, both interpretations of an ambiguous stimulus have a neural representation, and these representations inhibit each other's activity. When one representation becomes stronger than the other, its inhibitory force on the other also increases and the situation develops to full activity (dominance) for one, and minimal activity (suppression) for the other representation. Then slow self-adaptation gradually causes the dominant representation's activity to decline until a point is reached at which the balance tips in the opposite direction, et cetera. Models based on these ideas can account for existing data on dominance durations (Kalarickal and Marshall, 2000; Stollenwerk and Bode, 2003; Wilson, 2005; Lehky, 1988; Mueller, 1990). Notably, they agree with Levelt's 2nd proposition. In its original form this proposition stated that changing the strength (contrast) of one eye's stimulus while fixing the other one, affects dominance durations only of the eye with the fixed contrast (Fox and Rasche, 1969; Levelt, 1966). More recent work (Bossink et al., 1993; Mueller and Blake, 1989; Shiraishi, 1977) demonstrated small changes in the eye with the variable contrast as well, so the currently accepted interpretation is that unilateral contrast changes primarily affect dominance durations in the fixed contrast eye.

The good agreement between models and psychophysics is encouraging but also somewhat misleading, because the constraints posed by current psychophysical data are not very rigorous, as witnessed by substantial differences between the models cited above. It would seem, then, that current data do not allow all too specific inferences on the rivalry mechanism. Fortunately, there are indications that rivalry's alternation cycle is more diverse in its features than is usually taken into consideration, and therefore carries more information on underlying mech-

anisms than the observations that constrain current models.

Specifically, each alternation cycle comprises not only periods with complete dominance of either percept, but also substantial transition periods in which a compound of both is perceived (Blake et al., 1992; Bossink et al., 1993; Hollins, 1980; Wilson et al., 2001; Yang et al., 1992). A further informative complication lies in the fact that, although most transitions mediate a dominance change from one eye to the other, some end up with dominance returning to the previously dominant eye (Mueller and Blake, 1989). These latter events, which we call return transitions, have not yet been studied systematically. Finally, what we do know of rivalry's dynamics all comes from studies that measured at a limited set of specific contrast settings (e.g. fixing one eye's contrast at a chosen value and varying the other one) but not throughout the range of contrasts that we may present the eyes with.

In this study we aim to get a more complete view of the dynamics of binocular rivalry, including dominance durations, transition durations and the occurrence of return transitions. We study these for a matrix of left-eye and right-eye contrast combinations, spanning the entire range from near the detection threshold to the theoretical maximum. We subsequently test whether current rivalry models can account for these dynamics.

Using a conventional orthogonal grating stimulus (Figure 2.1) we demonstrate that dominance durations, transition durations and the frequency of occurrence of return transitions (FRT) all systematically depend on stimulus contrast. This dependence is similar for dominance durations and the FRT but different for transition durations.

Regarding dominance durations we show that Levelt's 2nd proposition (in its currently accepted form) is accurate only in a limited contrast range, and may be replaced more generally by the statement that unilateral contrast changes mainly affect dominance durations of the *higher contrast eye*.

These findings show that the data underlying current ideas represent only a part of the rich behavior of this system. The more complete characterization presented here allows new inferences on the mechanism of rivalry and poses more stringent model constraints. Indeed, we demonstrate using simulations that existing models cannot reproduce our findings. Specifically, systems based exclusively on cross-inhibition and slow self-adaptation cannot account for the many return transitions and long transition durations that we find. These features of the alternation cycle point toward neural noise as an essential driving force underlying rivalry.

2.2 Methods

Four subjects, one author and three naive, with normal or corrected to normal vision, participated. The stimulus (Figure 2.1) consisted of sine wave gratings (6.5 cyc/deg) filling a circular patch ($r=0.31$ deg) with constant contrast plus a surrounding region in which intensity fell off following a Gaussian profile (half-width 0.06 deg), i.e. a 'soft stimulus edge'. Average luminance of both stimulus and background was 15 cd/m². For fusion we used an alignment ring ($r=1$ deg; 50 cd/m²) with four lines extending 0.27 deg outward in the cardinal directions, and a binocular pattern of open squares (side 0.34 deg; 50 cd/m²) sparsely scattered across the screen from 13 deg above and below the centre onward. Subjects reported percepts by pressing and holding either of two buttons corresponding to exclusive visibility of either eye's image or releasing all keys in case of a transition. Because releasing the keys is also a natural response when one is uncertain of the percept (e.g. due to temporally unaligned eyes), subjects were instructed to press a third button to indicate such episodes. We sampled a 4×4

matrix spanning the full domain of left-eye/right-eye contrast combinations. The four contrast values were customized per subject on the basis of their detection threshold. We therefore determined for each subject the contrast yielding 75% correct for gratings presented monocularly (an otherwise identical stimulus) in a two-interval two-alternative forced choice Quest procedure (Watson and Pelli, 1983). We verified that this value was similar for both eyes and orientations, and used the average to determine the contrast range as follows: The lowest one was 0.75^{10} log-units Michelson above 75% detection, which meant 15% Michelson on average for these subjects; the highest one was the maximum of 100% Michelson, and we interpolated the other two in equal log steps. Contrast conditions were distributed randomly over trials. Sessions consisted of four 5-min experimental trials (see Appendix for remarks on this duration) and a 2-min control trial each. The first minute of each trial was discarded. Control trials, in which contrasts were always 100%/100% Michelson, were compared over sessions to verify subjects' constant performance, leading a fifth subject's data to be excluded from analysis. Each subject produced an average of about 180 dominance durations per condition per eye, and the accompanying transition durations.

2.3 Results

Figure 2.2 shows the results for a typical subject (panel A) and averaged over all four (panel B). The top and middle charts show dominance and transition durations; the bottom ones show the fraction of return transitions (FRT), i.e. the fraction of transition periods after which dominance returned to the previously dominant eye instead of crossing over. Return transitions were not included in the calculation of transition durations, a choice that did not notably affect any of our conclusions. None of the subjects showed a significant eye preference so data could be pooled over eyes. Consequently, regarding dominance durations, contrasts are given for the ipsilateral and contralateral eye, i.e. for the eye of which durations are plotted and for the other eye, respectively. Concerning transitions we use the terms 'departure' and 'destination' contrast: in the eye that was dominant before the transition and in the other one, respectively.

Figure 2.2 shows that both dominance and transition durations are on the order of seconds, dominance phases taking up the most time. Furthermore, return transitions are almost absent in some conditions but in other conditions make up as much as about half of all transitions starting from a given eye. Dominance durations (top) and the FRT (bottom) show similar patterns of contrast dependence, both being enhanced by a high ipsilateral/departure contrast and attenuated by a high contralateral/destination contrast. In addition, both values increase slightly when contrast is lowered symmetrically in both eyes, but this rise is small compared to the off-diagonal effects. Transition durations (middle) show a different pattern of contrast dependence, with the roles of departure and destination contrast being largely equivalent. As a result, transition durations are largest when both contrasts are minimal (cf. Hollins, 1980).

2.3.1 A restriction to Levelt's 2nd proposition

A hallmark of current ideas on rivalry, known as Levelt's 2nd proposition, is the notion that a change in one eye's contrast mainly ('solely' in the original formulation) affects dominance durations of the other eye and only to a lesser extent those of the eye itself. This has been shown in experiments in which one eye's contrast was fixed while the other one was varied. To test if our data support this notion, we reconstructed four such experimental regimes from our data,

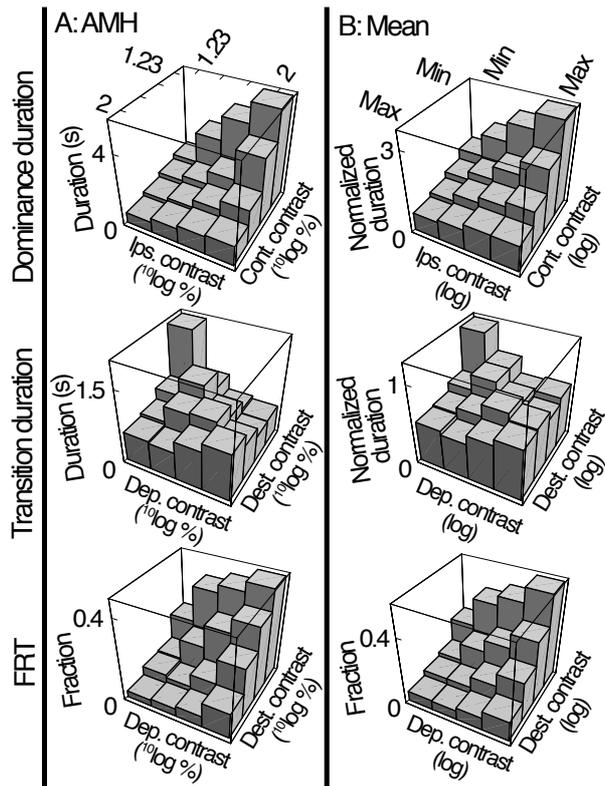


Figure 2.2: Dominance durations, transition durations and the fraction of return transitions (FRT) as a function of the two eyes' contrasts, for one subject (A) and averaged over all four (B). Ipsilateral (ips.) and contralateral (cont.) refer to the eye whose dominance durations are plotted and to the other eye, respectively. Departure (dep.) and destination (dest.) refer to the eye that was dominant before a transition started and the other eye, respectively. Contrasts were customized for each subject, 'Min' meaning near detection threshold, and 'Max' meaning 100% Michelson. Durations in panel B were normalized per subject relative to the dominance duration in the 100%/100% contrast condition (1.2, 0.9, 1.6 and 0.9 s for these subjects), so proportional relations between dominance and transition durations are preserved. The figure shows that dominance and transition durations are both on the order of seconds and that the FRT varies between about 0 and as much as 0.5, depending on condition. All three quantities show a systematic dependence on contrast, which is similar for dominance durations and the FRT but different for transition durations.

as shown in Figure 2.3. The bottom right inset, depicting a schematic top view of a chart such as in Figure 2.2, illustrates this. Taking a section through the matrix at a given ipsilateral contrast (dashed line; e.g. darkest shade) and combining it with the section at the same contralateral contrast (solid line; same shade), yields the data for one such regime. The four reconstructions differ in the level of the fixed contrast (dotted arrows in panels A-D). Panels A-D depict the reconstructions from our across-subject averages, shades corresponding to those in the inset. Panels A and B show the classic result: as the contrast in one eye is altered, dominance durations change mainly in the other eye. However, panels C and especially D show that the pattern subsides and actually reverses for lower values of the fixed contrast: Here the main effect is in the eye in which the contrast is changed. The validity of the proposition therefore depends on the subsection of the data one considers. It is valid at high values of the fixed contrast but invalid at lower ones.

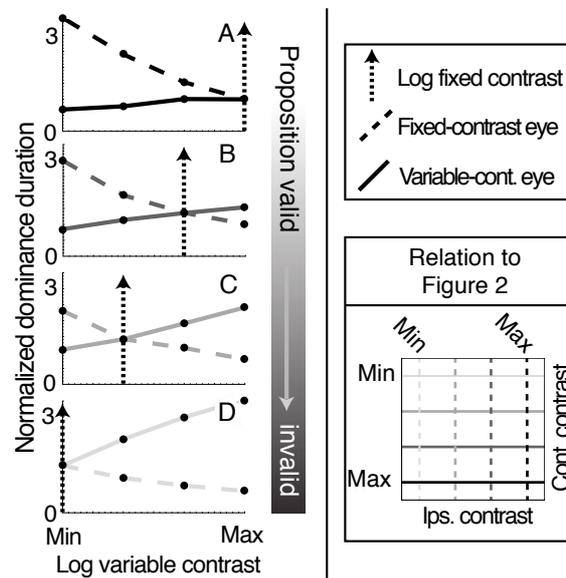


Figure 2.3: Levelt's 2nd proposition. In its current form this proposition states that a change in one eye's contrast primarily affects the other eye's dominance durations rather than those of the eye itself. It is based on experiments where one eye's contrast was fixed and the other one was systematically varied. From our data we took four subsets corresponding to such experimental regimes, as shown in the bottom-right inset, schematically depicting our 4×4 ipsilateral/contralateral contrast matrix. Each shade denotes one subset, the solid lines indicating the data points for the variable-contrast eye (i.e. contralateral contrast fixed) and the dashed lines those for the fixed-contrast eye (i.e. ipsilateral contrast fixed). Panels A-D show across-subject averages, using the same shades as in the inset. A and B. At high values of the fixed contrast (dotted arrows) we see the classic pattern. C and D. At lower values, however, this pattern subsides and reverses. The proposition therefore applies only in a restricted portion of the contrast domain, namely where the fixed contrast is high. Outside this domain, the opposite of the proposition becomes true.

2.3.2 Return transitions and dominance durations

Figure 2.2 showed similar contrast dependencies for dominance durations and the FRT. Here we examine this similarity more closely. Figure 2.4 displays across-subject averages of the FRT as a function of mean dominance duration in the departure eye, i.e. it shows how transitions starting from dominance for a given eye are related to dominance durations of that eye. The

horizontal and vertical extent of the 16 rhombi symbolize ipsilateral and contralateral contrast, smaller corresponding to lower. There is a clear positive correlation between dominance durations and the FRT, as quantified by the regression line.

It turns out there is no direct causal link between the occurrence of a return transition and the occurrence of a long-lived percept (see Appendix), i.e. one does not cause the other. Instead, the correlation stems from a common dependence on an underlying variable, strongly influenced by contrast. In addition, irrespective of contrast, our data showed a tendency for subjects with long dominance durations to experience many return transitions (not shown), strengthening the notion that both quantities are connected.

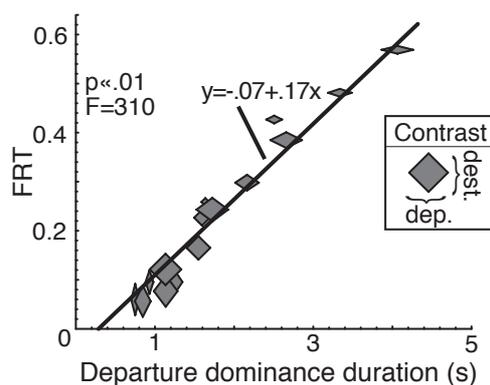


Figure 2.4: Relation between the FRT and mean dominance duration of the departure eye, in various contrast conditions. The width and height of the rhombi symbolize ipsilateral and contralateral contrast. Dominance durations and the FRT show a positive linear correlation, as indicated by the regression line. p is ANOVA p -value; F is ANOVA F -ratio. These data point toward a common underlying variable, responsible for trends in both departure dominance durations and the FRT.

2.3.3 Comparison with existing models

We performed simulations (see Appendix) with three models based on adaptation and inhibition, that agree with experiments to date (Kalarickal and Marshall, 2000; Stollenwerk and Bode, 2003; Wilson, 2005). To verify if these models can reproduce our findings, we tested their predictions on dominance durations, transition durations, and the FRT for a matrix of ipsilateral/contralateral input strengths.

To adequately cover transitions we chose two distinct classes of models, each addressing one of the two types of transitions that may occur (see also Discussion). The Wilson and Kalarickal & Marshall models cover superposition transitions, during which both images are seen superimposed (Burke et al., 1999; Liu et al., 1992), while the Stollenwerk & Bode model covers piecemeal transitions, during which parts of both images are seen in complementary regions of the stimulus (Blake et al., 1992; Lee et al., 2004; O'Shea et al., 1997; Silver and Logothetis, 2004; Wilson et al., 2001).

The models by Wilson and Kalarickal & Marshall comprise a single oscillator formed by the two percepts' representations interacting via adaptation and inhibition. Piecemeal transitions are beyond their scope as they have no spatial dimension. We defined superposition phases as those when neither of the representations strongly dominates, but instead both are intermedi-

ately active. The model by Stollenwerk & Bode involves oscillators similar to those described above, but several of them linked together in a 2D network, corresponding to neighbouring zones in visual space. Neighbouring oscillators are coupled such that they tend to follow each other's dominance states, so that once a dominance change emerges in one location a piecemeal transition may occur. We defined transition periods as those during which less than 100% of all oscillators were in the same dominance state.

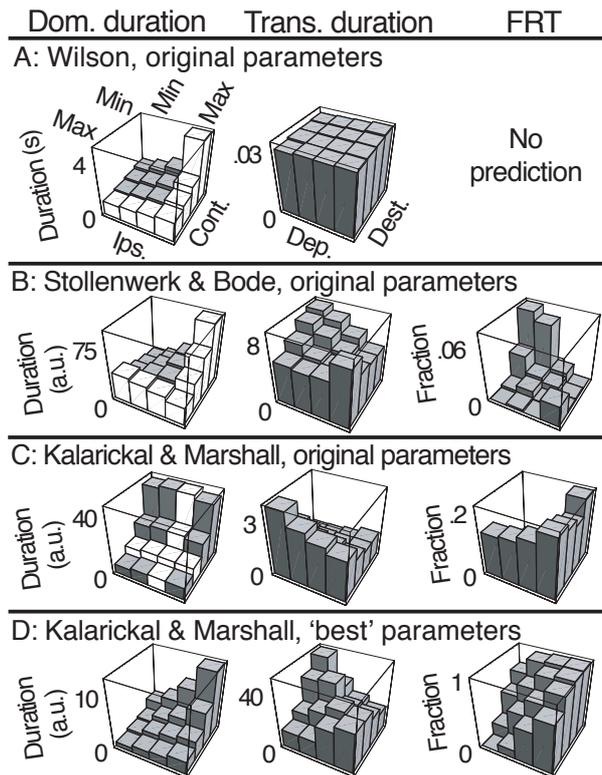


Figure 2.5: Simulations with three existing models. The left charts of each panel show simulated dominance durations as a function of ipsilateral (*ips.*) and contralateral (*cont.*) input strength; the middle and right ones show transition durations and the FRT as a function of departure (*dep.*) and destination (*dest.*) input strength (the Wilson model produces no return transitions as it is noise-free). Durations are given in seconds for the Wilson model, and in arbitrary units (a.u.) for the remaining ones. Panels A-C show results at the original parameter settings, with the white bars reproducing the original demonstrations of agreement with Levelt's 2nd proposition. Although all three models support the proposition (for a limited range of inputs) their behavior diverges at most other points. In addition, there are marked deviations from our data. First, the models incorrectly predict at least an order difference between dominance and transition durations, and also underestimate the FRT. Second, none of the models reproduces the found patterns of contrast dependence of transition durations and the FRT. Panel D. In simulations at other parameter settings we found one parameter region for the Kalarickal & Marshall model where the results were reminiscent of our data. Here contrast dependencies of all three variables were qualitatively correct, but were accompanied by an overestimation of both transition durations and the FRT. Further analysis (see Appendix) shows that in this region model dynamics are essentially stochastic. Deterministic forces keep the system in a state intermediate between both dominance states, and it is noise that causes incidental excursions into either dominance percept.

Figure 2.5, panels A-C display simulation results at the original papers' parameter settings. From left to right, the charts in each panel show dominance durations, transition durations and the FRT as a function of input strength. The white bars in the left charts reproduce the original

papers' demonstrations of agreement with Levelt's 2nd proposition. They also illustrate how we complemented the original input strength combinations to form 4×4 matrices. Aside from the white bars the models show a striking lack of concordance, underscoring the added constraints posed by the present characterization.

The charts show marked deviations from our data. First, all models predict transitions to be much shorter than dominance phases rather than of the same order. Second, they underestimate the FRT. Finally, contrast dependence of particularly transition durations and the FRT are not correctly reproduced.

To see if model predictions would improve at other parameter settings we performed simulations at an extensive range of values, confirming the shortcomings discussed above. There was one parameter region for the Kalarickal & Marshall model, however, (panel D) where the contrast dependence of all three variables was in qualitative agreement with our data, although both transition durations and FRT were overestimated. Interestingly, at these parameter settings the system is in a mode entirely different from the one that was originally intended (see Appendix). Its deterministic dynamics have settled at intermediate activities for both representations, and it is purely noise that causes the oscillations manifested in the figure.

2.4 Discussion

2.4.1 Noise and rivalry dynamics

The most interesting aspect of the present findings may be that current ideas on the system underlying rivalry do not agree with the observed dynamics. Broadly speaking, we show that there is more to the alternation cycle than current models can explain, and no single adjustment will likely bridge this gap. More specifically, however, there is one aspect regarding which the deviation between models and data is particularly strong and which to us indicates one necessary model addition. This aspect is the behavior of transitions. These are presently treated as brief and irreversible switches between two dominance periods, whereas the data show that they take considerable time and that return transitions are common. We will argue that the features of transitions in current models follow from an overemphasis on the role of slow adaptation, and that our results point to a system in which part of this role is taken over by stochastic variations (i.e. noise) in the system components.

This is illustrated in Figure 2.6, schematically showing perceptual dynamics of rivalry as motions across energy landscapes (e.g. Billock and Tsou, 2003). The system state (ball) always develops toward lower energy so minima in the landscapes are fixed points (attractors), corresponding to left-eye and right-eye dominance. Changes in the landscapes are due to changes in adaptation state, temporarily modifying the relative strength of the attractors. Sequence $A \rightarrow B \rightarrow C-I$ shows the situation as many current models treat it. A: the ball occupies the left attractor, indicating e.g. left-eye dominance. $A \rightarrow B \rightarrow C-I$: slow adaptation destabilizes the occupied attractor (black arrow) while recovery from adaptation deepens the other one (white arrow), until the left attractor disappears (C-I) causing the ball to move (dashed arrow) through the transition region (gray) to the remaining attractor. This transition takes place quickly and not until considerable time has passed, allowing the left attractor to reappear, may the system return to its previous state. This separation of timescales between landscape changes and ball movement is necessary for these models to oscillate, yet it is clear that the long transition durations and particularly the existence of return transitions, corresponding to the ball moving

halfway between both attractors and then returning, argue against this type of system.

In panels C-II and C-III we show two (not mutually exclusive) alternatives to panel C-I that do agree with our data. C-II: in this scenario transitions are not initiated by destruction of the occupied attractor, but by noise (curved arrow) tossing the system into the transition region near the maximum separating the left and right domain of attraction. Here, the slope is shallow, i.e. deterministic forces are weak, and may be positive or negative depending on the actual system state. This may explain both the long duration of transitions and the frequent occurrence of returns. C-III: here the slope in the transition region is so low, i.e. the deterministic dynamics here are so slow, that the attractor may reappear before the transition is over. Again, this is in line with long transitions and provides an opportunity for noise to tip the system over in either direction. We think this latter scenario is particularly likely at lower contrasts (see below). Both scenarios point to a strong stochastic term as a key ingredient missing from present thinking.

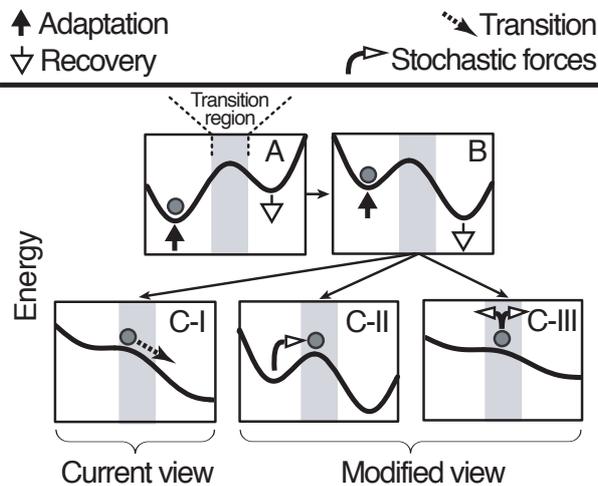


Figure 2.6: Schematic representation of rivalry's dynamics as currently understood ($A \rightarrow B \rightarrow C-I$) and as our data imply ($A \rightarrow B \rightarrow C-II$ and $A \rightarrow B \rightarrow C-III$). Dynamical modes of the system are shown as energy landscapes. The two minima are stable states (attractors), corresponding to left-eye and right-eye dominance. A. The ball occupies the left attractor, indicating left-eye dominance. $A \rightarrow B \rightarrow C-I$. Adaptation (black arrow) destabilizes the occupied attractor as recovery from adaptation (white arrow) deepens the other one. C-I shows transitions as treated by current models. As the occupied attractor disappears due to adaptation the system moves to the remaining one (dashed arrow). It cannot return to the left attractor until (recovery from) adaptation has changed the landscape to the mirror image of panel C-I. This scenario is incompatible with the slow and reversible transitions we observe, which point toward scenarios such as sketched in panels C-II and C-III. C-II. Transitions are initiated by noise (curved arrow) driving the system out of the still present attractor to a location in the transition region (gray) near the separation between both attractor domains. Here deterministic forces are small and the system may develop in either direction. C-III. The system remains in the attractor until it is gone, but the deterministic forces in the transition region are so small that the attractor may reappear before the transition is over. Again, noise may tip the system to the left or to the right. Both scenarios imply a crucial role for noise.

As such, the notion that noise plays a role in rivalry is not new, and indeed it is obvious that no biological system is noise-free. Particularly, stochastic variations in dominance durations are well known (Levelt, 1967; Brascamp et al., 2005; Fox and Herrmann, 1967) and can be reproduced by current models (Kalarickal and Marshall, 2000; Lehky, 1988; Stollenwerk and Bode, 2003; Wilson, 2005). However, to account for the stochastics of dominance durations, on

the level of detail that has hitherto been considered, it suffices to add jitter to essentially deterministic systems (e.g. Lehky, 1988). The exact nature of this jitter is of little importance and even if it is left out entirely the dynamics are not notably affected (Wilson, 2005). Consequently, the view of noise that has dominated the literature so far is of a detail that introduces some unpredictability to a course of events that is itself governed by deterministic forces. The role that we propose for noise is of a different order. Our data imply stochastic forces that are on a par with deterministic ones, frequently dominating them. Relevant models therefore require careful consideration of the nature of this noise and will display dynamics that differ fundamentally from those of current models. To put it another way, the point is not that noise plays a role in rivalry; the point is that it is a crucial factor for the system to function the way it does.

One recent study (Kim et al., 2005) extensively dealt with noise in rivalry, and it is important to discuss their findings in relation to ours. Like those authors, we are convinced of the importance of noise and we also endorse their assertion that more empirical constraints, like the ones here presented, are required to infer the internal workings of this system. In addition, Kim and co-workers took some important steps toward characterizing the relevant noise components, by calculating the amount of external perturbation that is equivalent to the system noise, and showing that this amount scales linearly with stimulus contrast. We significantly add to their conclusions by providing compelling evidence that noise is indeed a dominant factor underlying oscillations in unperturbed rivalry. In addition, whereas their data did not allow them to distinguish between various models (although the data did constrain the nature of the noise component) our data indicate shortcomings in all models we tested.

Given the importance of noise we may wonder where in the system these random fluctuations originate. There is presently no definitive answer to this question but let us consider the options. From studies dealing with noise in the context of visual detection and discrimination we know that signal loss may arise at any level of visual processing: from retinal photoreceptors and ganglion cells to the cortex itself, and both prior to and beyond binocular combination (Raghavan, 1989; Pelli, 1990). In addition, eye movements (both blinks and saccades) can influence the rivalry percept (Van Dam and van Ee, 2005; Wade, 1975), providing a possible noise source on the input side (although note that Kim et al. (2005) report similar results in the absence and presence of blinks). Finally, there is the possibility of deterministic chaos (Laing and Chow, 2002; Wilson, 2005) which, although technically deterministic, is equivalent to noise on the present level of analysis.

One intriguing possibility to note on the side is that, regardless of its origins, here noise may represent more than simply a physical limit on the system's accuracy. In the context of rivalry it may instead be of functional use. Arguably, rivalry's alternations reflect a general property of vision to continuously reorganize sensory input to reach a perceptual solution (Andrews and Purves, 1997; Leopold and Logothetis, 1999). In such a framework, noise may act to destabilize the present organization and prevent the brain from getting trapped in a single interpretation while others may have more survival value (Kim et al., 2005).

Although we emphasize the role of noise in rivalry we do not deny the importance of deterministic forces like adaptation and inhibition. Figure 2.5D is instructive in this respect, showing results from the Kalarickal & Marshall model in a parameter regime (see Appendix) where the system has only one, permanent, attractor that is itself moved around by noise. One definitive argument against this system, and against most purely stochastic systems, is that it produces exponential-like distributions of dominance durations, which is incorrect. Rather, all evidence points to a system characterized by an interaction between deterministic and stochastic forces.

Specifically, deterministic dynamics likely dominate the part of the alternation cycle when the occupied attractor is deep, while noise drives the system in the temporal vicinity of a transition. This view is further supported by recent findings (Lankheet, 2006) showing that noise added externally (i.e. to the stimulus) primarily has an effect in a small time window near the transition.

2.4.2 Levelt's 2nd proposition

A surprising finding of our study is the fact that Levelt's 2nd proposition, a widespread notion considered crucial for assessing the validity of models, is not generally applicable. It may be emphasized that the present findings are qualitatively different from other demonstrations of the limits of the proposition. Previous work has already led to the attenuated view that the effect of a unilateral contrast change is not *exclusively* but *mainly* in the fixed contrast eye (Bossink et al., 1993; Mueller and Blake, 1989). Our data on the other hand show that under certain conditions the main effect is in the variable contrast eye, so that the opposite of the proposition becomes true. The reason that this has so far been overlooked seems to be simply that the proposition had not yet been tested for the full contrast range that we considered. Levelt based the proposition on measurements in which one target was fixed at 89% contrast while the other was adjusted down to 8% (Levelt, 1966), and subsequent studies used similar designs. To our knowledge no study has fixed one contrast at a *low* value and then increased the other one to near maximum, the conditions crucial to our conclusions. In other words, although the proposition is correct in the contrast range in which Levelt and subsequent authors measured, it does not generalize beyond this range. It is therefore not a general feature of rivalry and one should act with caution in attaching any particular significance to it. In terms of model constraints the most appropriate description of the relation between contrast and dominance durations is probably not any single proposition but the three-dimensional surface implied by Figure 2.2. As a rule of thumb for experimenters, on the other hand, the most accurate alternative for the proposition may be the statement that *changes to one of the two contrasts mainly affect dominance durations in the higher contrast eye*.

2.4.3 The nature of transitions

As previously noted, transitions may be either local (superposition) or spatial (piecemeal) (Blake et al., 1985; Hollins, 1980; Liu et al., 1992; Yang et al., 1992). For the specific conclusion that noise is of great importance, it is not necessary to distinguish these two types: although slow transitions may partly be explained by a series of fast local transitions adding up to a longer spatial one, the return transitions in our data form a compelling argument for noise regardless of transition type. Ultimately, however, binocular rivalry cannot be understood without a clear view of the nature of transitions, because the two types correspond to entirely different system states. Specifically, the occurrence of piecemeal transitions points toward parallel rivalry in a number of separate (but linked) modules as in the Stollenwerk & Bode model (Blake et al., 1992; Stollenwerk and Bode, 2003; Wilson et al., 2001). A powerful paradigm to study such transitions in isolation has recently been developed (Lee et al., 2004; Wilson et al., 2001). The occurrence of superposition, on the other hand, indicates that the system is reaching the limits of bistability and approaching fusion, as noted by (Liu et al., 1992), who were able to study this perceptual state in the period directly following stimulus onset.

Clearly then, it is of value to know which type dominated in our study. Since the nature of our main experiment, in which subjects tracked their percepts in real time, did not allow a distinction between the two types, which often occur simultaneously or in quick succession, we performed a control experiment (see Appendix) to address this issue. Subjects observed 10-second binocular rivalry trials and assigned to each trial as a whole a number from 1 to 5, 1 meaning that all transition percepts in this trial were superposition percepts, 5 meaning they were all piecemeal. This control involved the same stimuli and subjects as the main experiment, but we used only the four symmetric contrast conditions. The outcome was clear-cut: going from high to low contrast all subjects had a monotonic decrease in their scores, from 'mostly piecemeal' (4.1 on average) at high contrast to 'mostly superposition' (2.1 on average) at low contrast. These results underscore that both local and spatial transitions play a role in rivalry, with piecemeal percepts dominating at high contrast and superposition becoming more prominent as contrast decreases. Our findings dovetail nicely with existing literature, as our transition durations at high contrast (about 0.5 to 1 s) agree well with the lower limit predicted in case of pure piecemeal transitions (the time required for a border between two regions of opposite dominance to sweep over our stimulus, calculated on the basis of (Horton and Hoyt, 1991) and (Wilson et al., 2001)), whereas superposition periods at stimulus onset have been shown to be particularly prominent at low contrast (Liu et al., 1992).

Note that as contrast decreases, changes in the transition phase are gradual, with the first effects, both regarding their durations and regarding subjects' assessment of their nature, being apparent at contrasts as high as 50%. This argues against the notion (Liu et al., 1992) that a separate neural mechanism is responsible for binocular summation at low contrast, which would be supported, for instance, by a discontinuity in the transitions' features near the low contrast end. A more parsimonious interpretation of the present data is that what we observe is a gradual evolution of the binocular rivalry process, from one producing abrupt perceptual flips to one that has more gradual transitions. Stable binocular summation would then be a limiting case at near-threshold contrasts. Such a gradual evolution is a natural property of many oscillator models, as illustrated by panel C-III in Figure 2.6, which shows that shallowing of the attractors at low contrast can be accompanied by increased stability of the region in between. So, although the possibility of a separate summation process can not presently be excluded, the data do not force us to invoke any such additional mechanism.

2.5 Conclusion

We have investigated binocular rivalry's alternation cycle in terms of both mean dominance and transition duration and the fraction of return transitions (FRT), in relation to stimulus contrast. Both dominance and transition durations were on the order of seconds; the FRT varied between about 0 and as much as 0.5. We found systematic patterns of contrast dependence for all three variables, similar for dominance durations and the FRT but different for transition durations. Regarding dominance durations we show that Levelt's 2nd proposition is valid only in a limited portion of the contrast domain, and may be replaced more generally by the proposition that changes in one eye's contrast mainly affect dominance durations in the *higher contrast* eye. We found a strong correlation between the FRT and mean dominance duration of the departure eye, showing the presence of a common variable underlying both quantities.

Our data allow inferences regarding rivalry's mechanism not allowed by previous data, and provide more stringent model constraints. They refute prevailing models driven by adaptation

and inhibition as these underestimate both transition durations and the FRT, and predict incorrect patterns of contrast dependence. The data imply a crucial influence of stochastic variations in the neural circuitry mediating rivalry, which these models overlook.

Appendix

Trial duration and fatigue

Our experimental trials were longer than usual in the literature (5 minutes) and it is important to rule out the possibility that subject fatigue played any role in our results. Figure A2.1 shows what Figure 2.2 would have looked like, had we analysed only the first 2 minutes of every trial rather than the last 4. Clearly, there are no important differences between this figure and Figure 2.2, indicating that trial duration is not a crucial factor.

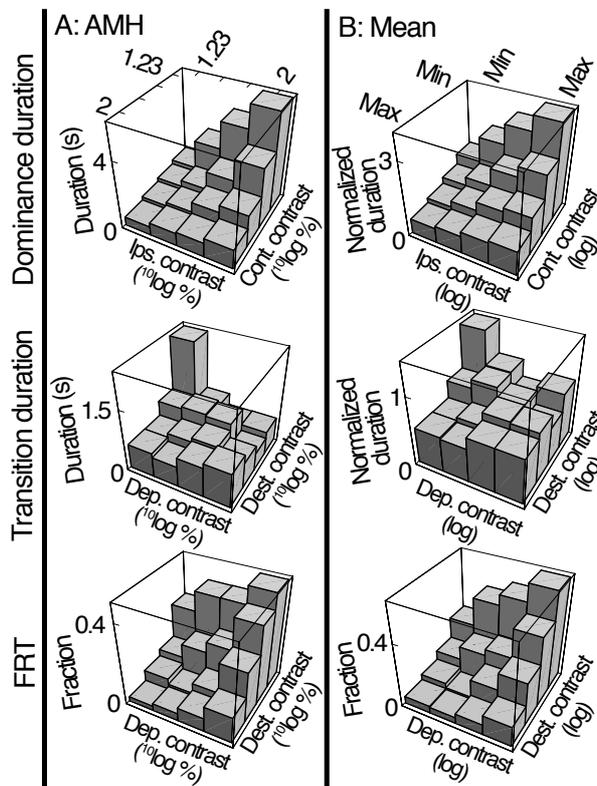


Figure A2.1: Same as Figure 2.2, using data from only the first 2 minutes of every trial. This figure and Figure 2.2 are almost identical, demonstrating that trial duration did not crucially influence our results.

Dominance durations and return transitions

The correlation between departure dominance duration and FRT shown in Figure 2.4 of the main text can have either of two explanations. First, both may depend on a common underlying variable (an indirect causal link). Second, the occurrence of a return transition may cause a long

duration or vice versa (a direct causal link). Figure A2.2 dissociates these options. For every return transition away from, and back to a given eye, we considered the 10 last dominance durations (normalized per trial and eye separately) of that eye preceding the return transition, and the 10 first ones occurring afterward. We then rank numbered these in time from -10 to -1 and from 1 to 10, respectively. This enabled us to calculate the average dominance duration per rank number. A direct causal link would show up as a positive deflection near the return transition in this event related average. Figure A2.2 shows no such deflection, so the correlation is instead due to a common factor causing both phenomena. Strikingly, there is even a negative deflection at rank numbers -1 and 1. This may be explained as follows. Subjects employ some decision criterion as to the amount of suppressed-eye contamination in a dominant percept that will cause them to report the onset of a transition. If, as we propose, a transition is a gradual, errant process, then there are occasions on which a dominant percept is temporarily disturbed by some *near-criterion* amount of contamination, and then re-stabilizes. On some of these occasions subjects will, and on some occasions they will not, report a return transition. When they do, this cuts the present dominance duration into two shorter ones, which would explain why the last percept preceding a return transition, and the first one following it, are relatively brief on average.

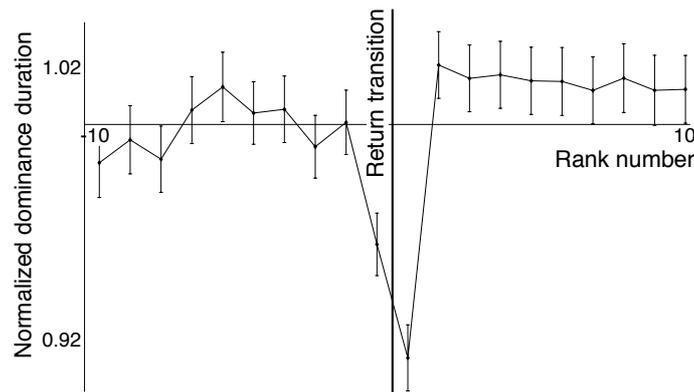


Figure A2.2: Event related average dominance duration for the departure eye, relative to the occurrence of a return transition. The absence of a positive deflection for dominance durations near the return transition excludes the option that the positive correlation shown in Figure 2.4 arises because return transitions cause long dominance durations or vice versa. Instead, the correlation must be due to a common underlying variable. The negative deflection that is evident instead is not unexpected if we consider a transition as a gradual, errant process (see text). Error bars indicate standard errors.

Simulations

The results in Figure 2.5A-C of the main text are for the parameter settings used by the original authors. For the input strengths we took four values equidistant in log-space, the smallest and largest ones based on the original papers: 0.9 and 1.5 for Stollenwerk & Bode, 0.853 and 1 for Wilson, and 0.3 and 0.7 for Kalarickal & Marshall. For the single-oscillator models we defined transitions as those periods during which both pools' activities lay within a factor 4 from each other; for the multiple-oscillator model the criterion was that less than 100% of the oscillators should be in the same dominance state, i.e. have the strongest activity for the same percept.

For the single-oscillator models we explored the surrounding parameter space by varying all parameters from half to twice their original values in three (Kalarickal & Marshall) or four (Wilson) equal log-steps. For the multiple-oscillator model, simulation time did not allow this so instead we fixed the two parameters defining the sigmoid non-linearity involved in the model, and varied the remaining six variables from $\frac{2}{3}$ to $\frac{3}{2}$ times their base values in two equal log-steps. For these latter simulations we reduced the number of coupled oscillators to 10x10 to fit our stimulus, based on the assumption that one oscillator covers about 0.1 deg of visual angle. This is roughly the maximum stimulus size for which rivalry is non-piecemeal (Blake et al., 1992). Along with the lattice size we reduced the base value for ρ_0 , the spatial extent of noise, by a factor two, since Stollenwerk & Bode tuned that value for a 20x20 lattice. For all simulations we used an explicit Euler iteration scheme, with step sizes of 0.1 (Stollenwerk & Bode), 0.02 (Kalarickal & Marshall) and 0.001 (Wilson). Note that in the Kalarickal & Marshall model, noise is applied to the adaptation *state*, rather than to its *rate* as implied by the original equations (G.J. Kalarickal, pers. comm.).

Besides the above simulations we performed simulations at the original settings to verify our correct reconstruction of the models. Figure A2.3 shows the outcomes along with those read from the original figures. The close agreement confirms that our reconstructions were correct.

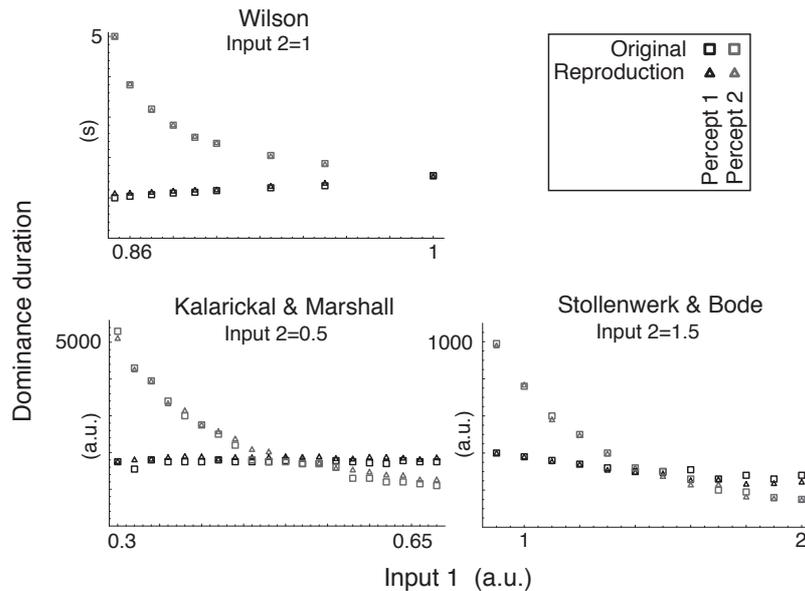


Figure A2.3: The original papers' simulation results (squares) and reproductions using our reconstructed models (triangles). The close agreement confirms the correct reconstruction of the models.

Kalarickal & Marshall model, 'best' parameters

As shown in Figure 2.5,D there exist parameter settings at which the Kalarickal & Marshall model is in qualitative agreement with our data. It turns out that at these settings ($W_1^+ = W_2^+ = 0.315$; $W_{12}^- = W_{21}^- = 125$; $c_1 = 0.005$; $c_2 = 0.016$; $c_3 = 0.105$; $s = 0.005$) the model has not two but one attractor,

and the dynamics are entirely governed by noise. This situation is illustrated in Figure A2.4A by a phase plane plot, displaying the states that the system may occupy in terms of both pools' activities (x_1 and x_2). For any system state, or x_1, x_2 combination, x_1 and x_2 develop as indicated by the flow arrows. The two curves (null-clines), however, indicate the locations at which either x_1 (black) or x_2 (gray) does not change, so at an intersection between the null-clines neither changes. At these parameter settings, contrary to the original ones, the lines have only one intersection, an attractor (plus sign), which all flow arrows lead toward. The only reason the system 'oscillates' is that the attractor itself is displaced by noise. It is unlikely that in reality binocular rivalry is entirely stochastic, as stochastic systems are generally associated with monotonically decreasing distributions of dominance durations, instead of the well-known unimodal ones observed in rivalry. Panel B illustrates this, showing a distribution of dominance durations obtained at these settings.

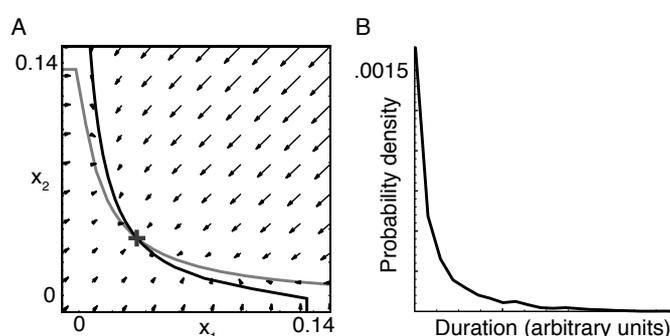


Figure A2.4: The Kalarickal & Marshall model in a mono-stable regime. At these settings the system has not two but one attractor (panel A), and the dynamics are entirely governed by noise. B. This system produces exponential-like distributions of dominance durations, rendering it an unlikely candidate for explaining rivalry, along with most other purely stochastic systems.

Superposition versus piecemeal transitions

In our main experiment we were unable to distinguish superposition and piecemeal transitions. In this control experiment we investigated their relative importance, and how this changes with contrast. Our four subjects viewed the grating stimuli for 10 s at a time, and afterward scored the subjective nature of the transition percepts observed during the period, on a scale ranging from 1 (only superposition) to 5 (only piecemeal). Each report was followed by another 10 s before the start of the next trial, preventing after-images from interfering with the stimulus on the next trial. This was repeated five times per subject per condition, randomly interleaving conditions and applying only the four symmetric contrast combinations. As shown in Figure A2.5 there is a clear shift from mainly superposition at low contrast to mainly piecemeal percepts at high contrast, although note that both types of percepts were perceived throughout the contrast range, as neither 1 nor 5 was scored very often.

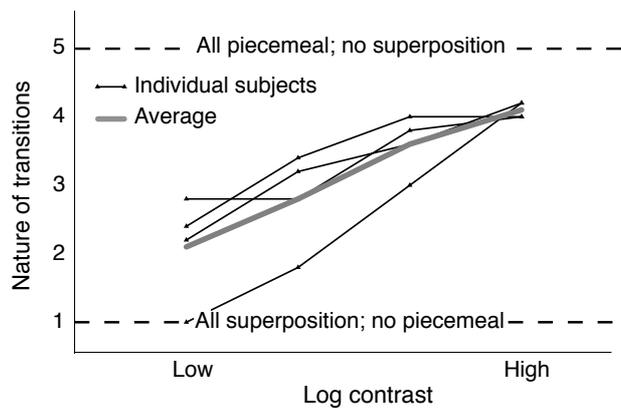


Figure A2.5: Relative importance of piecemeal and superposition percepts as a function of stimulus contrast. Both for individual subjects (black) and averaged over all four (gray) there is a clear shift from mainly superposition at low contrast to mainly piecemeal percepts at high contrast, although both types of percepts seem to occur throughout the contrast range.

Chapter 3

Sensory memory for ambiguous vision

When you spend a lot of time working on the same subject with the same people, you run the risk of losing contact with alternative ways of looking at the subject. After having worked on perceptual stabilization in Utrecht for a couple of years, it was a valuable change to go over to Nashville and work on stabilization there. This chapter is based on a review paper that Joel wrote with me during my time in Nashville. Both in terms of content and in terms of style the text is a blend of Joel's input and mine, and I think that increased the quality beyond what either of us could have reached individually. On a practical note, because this chapter was written later than some of the remaining ones, the stance taken in some later chapters is based on the state of the literature a while before this chapter was written. This may be a bit confusing, and another option would have been for me to place this chapter at the very end of the thesis. I chose not to do that, because I think the general background that the review paper provides, is a good introduction to the more technical chapters that follow.

Abstract

In recent years the overlap between visual perception and memory has shed light on our understanding of both. When ambiguous images that normally cause perception to waver unpredictably are presented briefly with intervening blank periods, perception tends to freeze, locking into one interpretation. This suggests a form of memory storage across the blank interval. This memory trace codes low-level characteristics of the stored stimulus. While a transient trace is evident after a single perceptual instance, a more persistent trace builds over many separate stimulus presentations, indicating storage on multiple timescales. This memory shares important characteristics with priming by non-ambiguous stimuli. Computational models now provide a framework to interpret many empirical observations in this paradigm.

3.1 Stabilizing unstable vision

While most visual input produces stable vision, sometimes when visual information is ambiguous awareness tends to waver continuously between alternative interpretations (Blake and Logothetis, 2002; Tong et al., 2006). For one class of ambiguous sensory stimuli, known as bistable stimuli, an observer's perception will alternate almost exclusively between only two interpretations (Figure 3.1A-C). During continuous viewing conditions alternations are unavoidable, and it is impossible to predict what someone will see in the near future (e.g. Fox and Hermann, 1967) (Figure 3.1D).

Surprisingly, perception of a bistable stimulus can be made stable and predictable by the simple manipulation of periodically removing the stimulus from view (Figure 3.1E). In 1963 Orbach and colleagues reported that intermittent presentation of a bistable pattern – making it appear and disappear every few seconds – reduced the overall number of perceptual alternations (Orbach et al., 1963). More recently, a more general account by Leopold and colleagues demonstrated that such presentation techniques can effectively make perception freeze on one interpretation of a bistable stimulus (Leopold et al., 2002) (Figure 3.2; compare left and right panels). This perceptual stabilization contrasts dramatically with the continual perceptual changes experienced during uninterrupted viewing. This remarkable phenomenon implies a form of memory in which the visual system maintains information from past perception across blank intervals.

This is an exciting time for this relatively new phenomenon. We are beginning to understand its many characteristics and components, from the type of information the brain actually stores, to where and how the brain achieves this. In this review we will first consider which aspects of perception are held in memory and which ones seem irrelevant. Second, we will discuss the mechanism of the memory, its temporal dynamics, and how recent modeling efforts have changed the way we think about it.

3.2 Information held in memory

What types of information are held across these blank intervals? One way to answer this question is to physically change different characteristics of the stimulus in synchrony with the on/off presentation cycle (Chen and He, 2004; Pearson and Clifford, 2004; O'Shea, 2004). For instance: if we change the color of the stimulus from one presentation to the next (across the blank interruption), from blue to red, and as a result an observer experiences an alternation (the stimulus interpretation changes) this suggests that the memory trace from the blue stimulus did not stabilize the red stimulus. This implies that the stimulus color is stored in the memory trace. Over multiple presentations this scenario would result in lower perceptual stability. If, instead, stability is unaffected because the observer sees the same stimulus interpretation despite the color change, memory can be said to be blind to this feature (see Figure 3.3 for a summary of findings).

3.3 Structure from motion

When rotating three-dimensional stimuli, such as a cylinder or sphere, are presented on a monitor in two dimensions their direction of rotation is inherently ambiguous. This results in per-

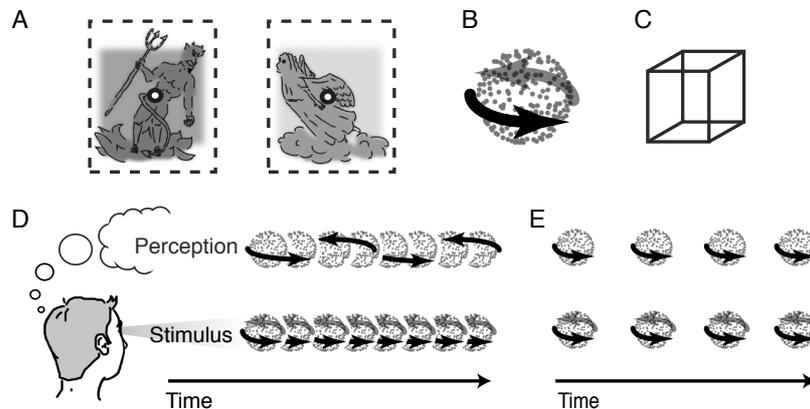


Figure 3.1: *Ambiguous stimuli and stabilization timeline. A. A binocular rivalry stimulus. Each eye is shown a different image, causing the observer to perceive both images in alternation. If you can cross-fuse these images you can experience binocular rivalry right here on the page. B. A static snap-shot of a bistable structure from motion stimulus, see glossary for definition. C. A wire-frame cube, known as a Necker cube. The depth relation between the lines is ambiguous. Perception alternates between two configurations of a cube. D. Continuous viewing of an ambiguous image causes unpredictable and inescapable perceptual alternations between interpretations. E. Periodically removing the stimulus from view can cause one percept to dominate on each presentation. This perceptual stabilization implies some sort of automatic memory that stores perceptual information across periods of stimulus absence.*

ception that oscillates between two possible directions of rotation (Figure 3.1B). Like many other bistable patterns, such structure from motion stimuli are perceptually stabilized by intermittent presentation. If the color, size or rotation speed of the stimulus is changed on each presentation, perception remains stabilized (Chen and He, 2004), suggesting these features are not contained within the memory trace. Similar findings were obtained in a related paradigm where two different structure from motion stimuli were alternately presented one after the other (Maier et al., 2003). If the two stimuli differed in color, size, rotation speed, shape, or (to a lesser extent) eye of origin, alternations, although infrequent, were highly correlated between the two interleaved stimuli. A reversal of the perceived direction of motion of one stimulus was likely to be followed by a reversal in the perceived motion direction of the other. This indicates that in terms of memory two stimuli that differ along these feature dimensions are effectively treated as the same.

One factor that is stored in memory is the orientation of the axis around which a structure from motion stimulus revolves (e.g. horizontal or vertical). When two stimuli that differ sufficiently in their angle of rotation are presented in an interleaved cycle, they become independently stabilized (Maier et al., 2003) (see Figure 3.4A). Here, the rotational angle of the pattern is so important to the memory that the visual system treats these two stimuli as if they were categorically different. Each pattern becomes independently stabilized, as though the other pattern was not there. This indicates the coexistence of two independent memory traces, each specific to a rotation angle.

A second factor represented in the memory trace is the visual location of the stimulus. When the retinotopic location of structure from motion stimuli is varied between consecutive presen-

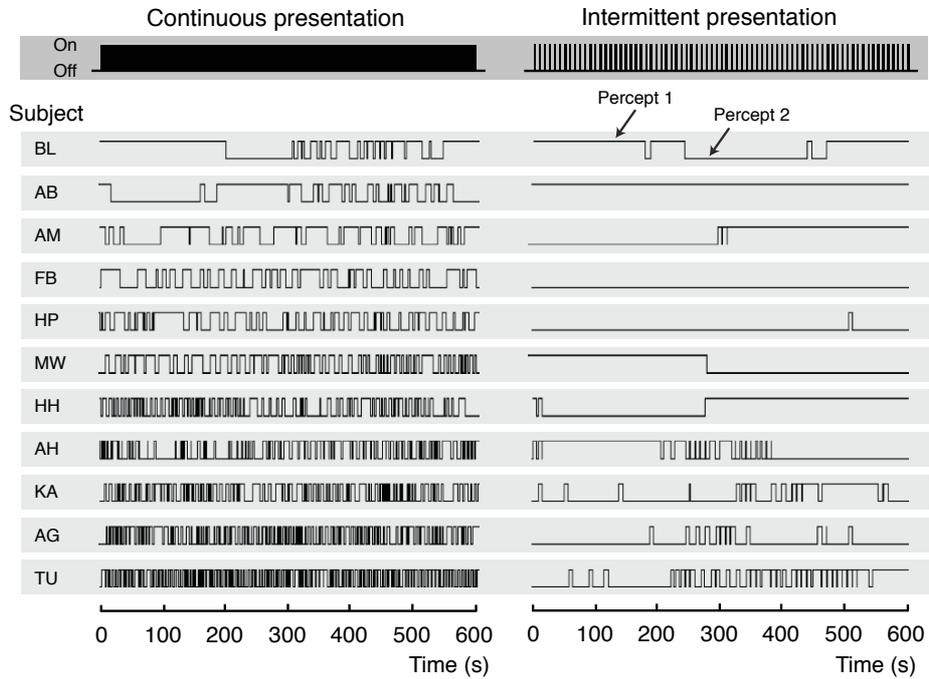


Figure 3.2: *Perceptual stabilization. Eleven individual subjects tracked perception of a structure from motion stimulus presented either continuously (left) or intermittently (right). Each subject demonstrates a dramatic reduction in the number of perceptual alternations during intermittent viewing compared to continuous viewing. Adapted from Leopold et al. (2002).*

Feature \ Stimulus	Eye of origin	Color	Orientation	Position on retina	Rotation direction	Size	Shape	Rotation speed	Depth
					X	?	?	X	X
									

 Not stored in memory
  Stored in memory

Figure 3.3: *What type of information is stored in perceptual memory? This table summarizes which attributes of a percept are stored across a blank period and which ones are seemingly forgotten. The top row shows data for binocular rivalry; the bottom row for a structure from motion stimulus. Dark brains indicate that the particular information is not stored in the memory trace, light brains denote the information is stored in memory.*

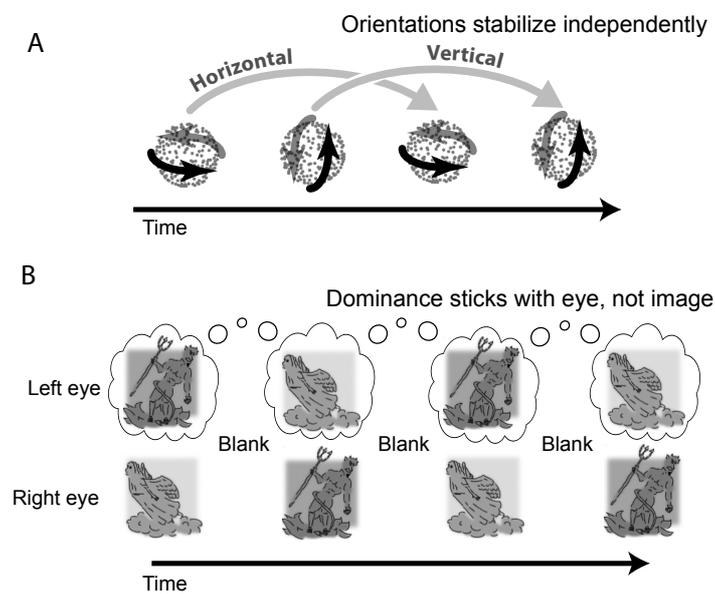


Figure 3.4: Different methods for investigating memory content. A. When two structure from motion stimuli with different angles of rotation are intermittently cycled at the same location in retinotopic space in an interleaved fashion, both patterns are independently stabilized. The two memory traces can co-exist for the same area of retinotopic space. Hence, the memory trace is specific to the angle of rotation (Maier et al., 2003). B. By swapping the two patterns between the eyes in synchrony with the intermittent cycle of binocular rivalry, researchers have shown that it is primarily the eye of origin that is stored in memory (Chen and He, 2004; Pearson and Clifford, 2004).

tations, subjects see more alternations (Chen and He, 2004). This indicates perceptual memory is specific to the location in retinotopic space. The spatial layout of visual information on the retina is strongly topographical in early visual cortex. Hence, the memory exhibits known characteristics of early visual areas.

In sum, for structure from motion stimuli perceptual memory entails storage of the angle of rotation within a limited region of the visual field, and carries little information on features such as size, shape, speed or color.

3.4 Binocular rivalry

Binocular rivalry is a form of bistability that occurs when two dissimilar stimuli are presented one to each eye, in corresponding locations of visual space (see Figure 3.1A). Binocular rivalry memory, like memory for structure from motion stimuli, is local in terms of retinotopic space (Chen and He, 2004). Another important factor is the eye of origin (the eye which sources each stimulus) of the dominant stimulus. When two competing static images are exchanged between the eyes in synchrony with the intermittent presentation cycle, dominance will stick with the same eye rather than the same image (see Figure 3.4B) (Chen and He, 2004; Pearson and Clifford, 2004). The memory retained between presentations effectively stabilizes the dominance relation between the eyes, even when different patterns are presented. An individual experiencing binocular rivalry is unaware of which eye is sourcing the dominant pattern (Ono and Barbeito, 1985), suggesting that what is stabilized can be dissociated from the conscious percept. Hence, the content of memory is not necessarily conscious perceptual information, rendering ‘perceptual memory’ somewhat of a misnomer.

Swapping the eye of origin across presentations is less disruptive to perceptual stability if moving, three-dimensional shapes are used as binocular rivalry images, instead of simple static patterns (Grossmann and Dobbins, 2005). This suggests that the memory can be specific to the attributes in the stimulus. For example, if colored and oriented patterns are used for the rivalry stimuli, the memory not only consists of eye of origin information, but also to a lesser extent color and orientation information (Pearson and Clifford, 2004).

Recent results have revealed that subjects exhibit a retinally specific preference for one percept over the other during binocular rivalry (Carter and Cavanagh, 2007). For example, an individual may have an intrinsic preference for a green pattern in the top right region of visual space. This location preference could partially drive changes in perception across presentations at different locations. Hence, care is needed when interpreting memory measured across different retinotopic locations.

Perceptual bistability similar to conventional binocular rivalry can be observed when the two patterns undergoing rivalry are flickered at about 18 Hz and exchanged between the eyes around 3 times per second (Logothetis et al., 1996). This form of ‘stimulus rivalry’ exhibits memory very similar to classical binocular rivalry (Pearson and Clifford, 2004). However, memory for stimulus rivalry lacks the eye of origin component prevalent in classical binocular rivalry memory (Pearson and Clifford, 2004). Despite this difference the memory for these two types of rivalry is effectively interchangeable (Pearson and Clifford, 2005). Perceptual dominance in one form of rivalry will carry over to the other, stabilizing perception across the two. This is surprising because evidence suggests that these two types of rivalry are contingent on neural activity at different levels of visual processing (Pearson et al., 2007).

In sum, the emerging view from these experiments is that perceptual memory elicited by

a given stimulus has characteristics very similar to those of the early sensory neurons that respond to the stimulus during continuous presentation. For binocular rivalry these would be neurons sensitive to information like eye of origin, color and orientation, and specific to a region of retinotopic space. For structure from motion stimuli one would expect neurons tuned to motion direction as well as visual location. We interpret this as evidence that the perceptual memory arises across the specific sensory areas that deal with the perception of a given ambiguous stimulus. This idea is consistent with results based on event related brain potentials (ERPs) (Kornmeier and Bach, 2004, 2005). Recent functional imaging work also points to sensory areas, while also suggesting the involvement of frontal areas (Sterzer and Rees, 2008).

3.5 Temporal characteristics of memory

Many studies of memory for ambiguous perception demonstrate correlations between current and past dominance. However, the exact causal relationship in this situation is unclear. (Pearson and Clifford, 2005) used an experimental paradigm known as flash suppression to control perceptual dominance during binocular rivalry. During flash suppression the two rivalry images are switched on asynchronously, and as the second image appears it becomes dominant. This presentation technique allows perceptual dominance to be reset to the other eye at any time throughout a cycle of stabilized rivalry presentations. When perceptual dominance was reset like this, dominance in the next normal rivalry presentation in the cycle tended to follow the percept imposed by flash suppression. This demonstrates a causal relation between perceptual dominance before and after a blank interval, as opposed to only a correlational relationship.

Recent experiments show that perceptual memory is not limited to the effect of the single last percept, but instead accumulates across many perceptual events (Maloney et al., 2005; Brascamp et al., 2008). To investigate this issue, Brascamp et al. (2008) first stabilized perception using an intermittent presentation cycle. Perception was then made to oscillate either by leaving the stimulus on long enough to prompt perceptual alternations, or by forcing an alternation using flash suppression. After dominance had switched to the opposite percept, intermittent presentations were resumed and perception re-stabilized (Figure 3.5A). Strikingly, perception typically did not re-stabilize to the new 'reset' percept, but resumed the old stabilization sequence from before any alternations. The single last percept only had a transient influence on subsequent perception (Figure 3.5A). The driving force for this sustained memory was the overall ratio of dominance of either percept during the prior minute of perception. Hence, memory cannot only have a transient effect across a single interval, but has the potential to grow stronger and become more persistent over time, across multiple brief presentations. These findings indicate that memory is built up of several components which each operate on a different timescale. The sustained components survived at least four consecutive alternations during the intervening continuous presentation. However, the probability of stabilization surviving these alternations did decline as the number of continuous alternations increased (Figure 3.5B). These results show that, rather than providing a snapshot of the most recent perceptual situation prior to a blank, perceptual memory can reflect a balance of forces accumulated over the course of many presentations.

One implication of this flexible long-term storage is that the longevity of this memory cannot be fully estimated by comparing perception on two consecutive presentations separated by various durations. Instead, the maximum blank duration that perceptual memory can sur-

vive depends on the specific sequence of preceding perception. Another complication is that the lifetime of this memory depends on the specific conditions during memory retention. For instance, attentional manipulations can modulate the longevity of memory for some stimuli (Kanai and Verstraten, 2006).

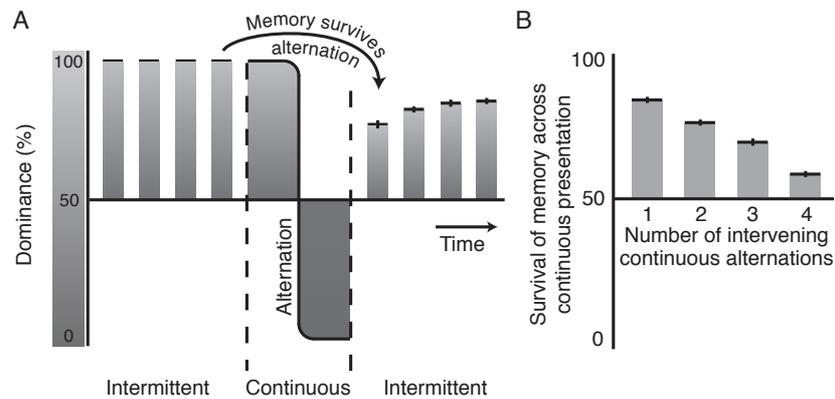


Figure 3.5: *Perceptual memory spans back further than the most recent percept. (A). A period of continuous viewing was inserted halfway into an intermittent presentation sequence, prompting a perceptual alternation. In spite of this alternation, perceptual stabilization was little affected when intermittent presentation was resumed. (B). In further experiments the intervening continuous viewing period contained more than one alternation. Perceptual memory showed above-chance survival across up to four alternations. This suggests that perceptual memory is not simply memory of the latest percept, but of a more elaborate history of prior perception. Adapted from Brascamp et al. (2008).*

When the blank interruption during intermittent presentation is exceedingly brief the stabilizing effect reverses. Following a blank interval shorter than about half a second, the observer is more likely to see the interpretation that was previously suppressed than the one that was previously dominant, a process of perceptual destabilization (Orbach et al., 1963; Noest et al., 2007; Klink et al., 2008; Orbach et al., 1966). This negative relation still indicates an interaction across a period of stimulus absence, but this interaction does not involve a priming effect. Instead, it is reminiscent of the suppressive effect often associated with prior unambiguous stimulation on ambiguous perception (see next paragraph). At such brief blank intervals the rate of switching increases to a point well above the rate during continuous viewing (Orbach et al., 1963; Noest et al., 2007; Klink et al., 2008; Orbach et al., 1966). Here, instead of slowing alternations, switching the stimulus on and off has the opposite effect: it speeds up the alternation rate. Not until the blank interruptions are contracted (<100ms) so that the presentation is almost continuous, do alternations tend to slow again, approaching the rate during continuous presentation. Recent models can account for both the stabilizing and the destabilizing effects (see below).

3.6 Unambiguous prior stimuli

One way to think of perceptual stabilization is as a priming effect of past dominance on percept choice at the next stimulus presentation. In most conditions when an ambiguous stimulus follows an *unambiguous* stimulus corresponding to one of the percepts, this prompts perception

of the opposite percept upon ambiguous presentation (e.g Long and Olszweski, 1999; Nawrot and Blake, 1989; Wolfe, 1984; Pearson and Clifford, 2005). This tendency toward the opposite percept contrasts with the overall tendency for ambiguous stimuli to prime dominance of the same percept (Pearson and Clifford, 2005). Still, a good case can be made that ambiguous and unambiguous prior stimuli do not have categorically distinct effects on subsequent processing, but are on different ends of one continuum.

On the one hand many conditions have been identified in which an unambiguous prior stimulus (like an ambiguous one) primes perception for ambiguous viewing (Wilton, 1985; Long et al., 1992; Jiang et al., 1998; Pinkus and Pantle, 1997; Kanai and Verstraten, 2005; Long and Moran, 2007; Brascamp et al., 2007; Kanai et al., 2007). Conversely, an ambiguous prior stimulus (like an unambiguous one) can cause dominance of the opposite percept if the interval between both presentations is sufficiently brief, as discussed in the previous section (Orbach et al., 1963; Noest et al., 2007; Klink et al., 2008; Orbach et al., 1966).

These results suggest that the persisting neural effects evoked by a prior stimulus – be it ambiguous or unambiguous – are characterized by both positive components (priming) and negative components (suppression). In this context, an appearing ambiguous stimulus could be seen as a particularly sensitive psychophysical tool to gauge the balance between these components. In fact, binocular rivalry perception has recently been used in this way to measure the persistent effects caused by mental imagery, in the absence of incoming visual signals (Pearson et al., 2008).

3.7 Modeling

Recent modeling efforts now unite many of the psychophysical findings within a plausible account of this memory for ambiguous vision (Wilson, 2007; Noest et al., 2007). In fact, perception during continuous and intermittent viewing can now be explained within a single framework.

Initial theories of perceptual stabilization were strongly influenced by the historical emphasis on models of *continuous* ambiguous perception. These models center on the concept that competing perceptual interpretations are represented in the brain by opposing neural populations. For continuous viewing, typically these models posit that perceptual alternations occur because a form of self-adaptation or neural fatigue weakens the neural representation of the dominant stimulus. When the neural representation of the dominant stimulus becomes critically weak an alternation takes place (Blake et al., 2003; Wilson, 2003; Freeman, 2005). The evidence from neuroimaging and neurophysiology studies is ambiguous, with some studies supporting the idea that neural adaptation drives perceptual changes (Tong et al., 1998; Tong and Engel, 2001; Leopold and Logothetis, 1996), whereas other studies suggest that a distributed high level process, rather than adaptation, is responsible for driving perceptual alternations (Lumer et al., 1998; Lumer and Rees, 1999; Sterzer et al., 2002). Inspired by these earlier models, perceptual stabilization was thought to occur because periodic stimulus removal slows the build-up of adaptation, causing adaptation to take longer to reach a critical level (Orbach et al., 1963; Leopold et al., 2002; Chen and He, 2004). It is now clear that this reasoning is incongruent with empirical findings. A slower increase in adaptation as the cause of stabilization is hard to reconcile with the finding that, given the right timing, interruptions can also promote alternations (Noest et al., 2007; Klink et al., 2008; Orbach et al., 1966). Furthermore, adaptation is freely allowed to reach its critical level when a period of continuous presentation is inserted halfway through a stabilized cycle of intermittent presentations, however this does not reset

the stabilization cycle (Brascamp et al., 2008) (see Figure 3.5).

Classical adaptation-inhibition models as they stand are insufficient to explain perceptual stabilization. Mutual inhibition between two competing neural populations is a crucial element of models of continuous ambiguous perception (Noest et al., 2007; Wilson, 2007; Freeman, 2005; Laing and Chow, 2002; Mueller, 1990). In these models the dominant representation strongly inhibits the suppressed representation, which reinforces the present dominance relation. Only when adaptation of the dominant population reaches a critical value at which it outweighs this stabilizing force does perception switch. After a seconds-long blank interval, however, the spiking response to a stimulus (and hence mutual inhibition) has likely died out (Wilson, 2005). Adaptation, on the other hand, decays much more slowly. Therefore, when the stimulus reappears there is still residual adaptation, but no inhibition. In other words, without the counteracting force of inhibition there is no threshold left for adaptation to overcome, and any residual adaptation will cause instant suppression. Indeed, no classic adaptation-inhibition model can explain perceptual stabilization.

One way to expand these models to account for perceptual stabilization is to add a component that is persistent, like adaptation, but, unlike classical adaptation, has a facilitating effect on future perception. Two recent studies account for stabilization using this approach. One model (Noest et al., 2007) put forward by Noest and colleagues posits that being in an adapted state has two distinct effects on a neural population. First, it causes a reduced response gain; a classic interpretation of adaptation. The second effect is different from classic accounts, and can be interpreted as a sub-threshold elevation in the baseline activation of the adapted neurons. This second, positive effect gives the more adapted representation a head-start when a stimulus reappears (see Figure 3.6B; right panel). Because this is generally the recently dominant representation, this allows for the occurrence of perceptual stabilization. In addition, the model parsimoniously explains why perceptual stabilization turns into destabilization at short blank intervals (Noest et al., 2007). A second model, proposed by Wilson (Wilson, 2007), expands an adaptation-inhibition framework with a form of long-term synaptic potentiation. Here, activity brings a neural population into a persistent state of elevated excitability (Figure 3.6B; left panel). It is this potentiation that provides the competitive advantage that is needed to explain perceptual stabilization. Figure 3.6 shows the time courses for the two hypothetical memory traces implemented by the models for a given intermittent presentation cycle. The light shaded plots show activity and traces for the dominant interpretation, while the darker plots show the suppressed interpretation.

The models differ in their various predictions and simulated data. The longer a stimulus interpretation has been perceived before offset, the more likely subsequent perception is to follow the same interpretation (Leopold et al., 2002; Brascamp et al., 2008). The model proposed by Noest and colleagues (Noest et al., 2007) predicts this empirical observation, whereas Wilson's account predicts the opposite (Wilson, 2007). The model by Noest and colleagues has recently been expanded in a straightforward manner to capture the multi-timescale nature of perceptual memory observed empirically (Brascamp et al., 2008).

3.8 Related memory phenomena

What makes this memory phenomenon unique is its manifestation – as dominance during perceptual ambiguity. We can, however, look beyond the specific expression of this memory, and

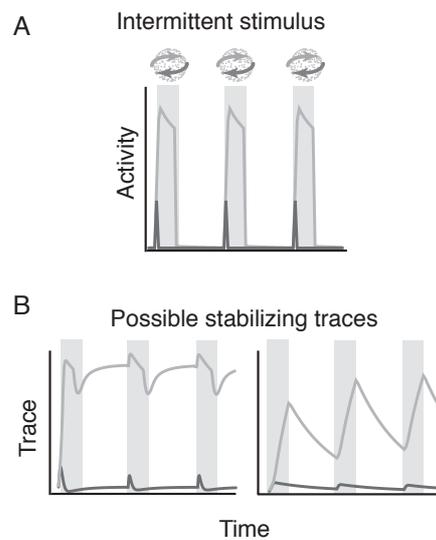


Figure 3.6: Model activity and traces during stabilized perception. A. Activity of two competing perceptual representations during stabilized perception. During stimulus presence the representation of the dominant percept is highly active (light shade), while that of the suppressed interpretation is only weakly active (dark shade). During the blank interval activity for both representations falls to near zero. B. Proposed memory traces that can account for stabilized perception. Left panel: traces of past perception proposed by Wilson and on the right those by Noest et al. Left: a facilitatory trace of the dominant percept persists across the blank intervals. This maintains the stabilized representation in a state of heightened excitability, allowing it to gain dominance at stimulus reappearance. Right panel: a trace of the stabilized percept builds during dominance and passively decays during stimulus absence. This trace slightly elevates the baseline activation of the stabilized representation, which can allow it to gain dominance at stimulus reappearance.

tentatively identify a more general internal mechanism underlying it. Perception during brief presentations of an ambiguous stimulus reflects the almost instantaneous choice between percepts that occurs at each stimulus appearance, and this competitive choice is very sensitive to small imbalances (Wilson, 2007; Noest et al., 2007). Due to the binary nature of this process and its wholesale effect on awareness this memory is a cogent and powerful phenomenon in terms of its influence on an observer's experience. However, internally it might be more accurately described as a modest shift in readiness. The question that now arises is: how would such shifts be manifested in paradigms not involving perceptual ambiguity?

Similar shifts in readiness might underlie the response time reduction that occurs over repetitions of a recurring oddball target (Kristjánsson, 2006; Campana et al., 2002; Maljkovic and Nakayama, 1994), or automatic acceleration over repetitions of a motor action (Dorris et al., 2000; Soetens et al., 2004). We suggest these various priming phenomena, despite obvious differences, may be similar in terms of the neural events that underly them.

In another paradigm low contrast primes were shown to reduce subsequent detection thresholds (Tanaka and Sagi, 1998). Interestingly, the same kind of primes prompt dominance in binocular rivalry (Brascamp et al., 2007; Pearson et al., 2008), suggesting a possible relation between history effects in these paradigms.

Researchers have also studied active-manual forms of conscious perceptual memory (Magnussen and Greenlee, 1999). This typically involves participants having to actively remember a particular sensory feature, then making a comparative judgement. This is a clear dissociation from the automatic forms of perceptual memory evident during ambiguous vision.

3.9 Conclusion

Memory for ambiguous vision is largely composed of 'early level' sensory information. While a single instance of ambiguous vision can directly bias subsequent ambiguous perception, the memory trace does not only carry information from the single last perceptual event, but can build over many perceptual events spread across time. The memory trace shares characteristics with various types of priming. Computational models can currently account for a broad spectrum of empirical observations.

This is an exciting time to watch the rapidly growing body of empirical findings complemented by computational frameworks forging our understanding of this memory phenomenon. We now have some understanding of what information is remembered and how and where the brain might hold this information. We look forward to future growth and new directions, as research into this phenomenon continues to excite and surprise.

Chapter 4

Flash suppression and flash facilitation in binocular rivalry

Published as: Brascamp, J. W., T. H. J. Knapen, R. Kanai, R. van Ee and A. V. van den Berg (2006) Flash suppression and flash facilitation in binocular rivalry. *J Vis*, 7(12): 1-12

More than any other finding in this thesis, we stumbled upon the main one in this chapter by accident. We were investigating perceptual stabilization of binocular rivalry. We wondered whether the memory trace that is responsible for perceptual stabilization can also alter detection thresholds of patterns that have previously dominated in rivalry. In our setup, we alternately presented a pair of binocular rivalry images (so that the dominant image could leave a trace) and a single one of those images at a very low contrast (for the observer to detect). As I was performing the task, I noticed that if the low-contrast monocular image was, say, a left-tilted grating, the dominant pattern in the immediately following rivalry presentation was usually also the left-tilted grating. It seemed that the monocular presentation was facilitating rivalry dominance in the subsequent trial. This surprised me, because I knew of only literature that reported the opposite effect. This opposite effect, where a monocular presentation suppresses subsequent rivalry dominance, is called 'flash suppression'. I went to Ryota's office and asked him if he had ever heard of anyone reporting this. He thought for a second and said: 'Hmm... flash facilitation'. We set out to explore the characteristics of the phenomenon, and found that they fit surprisingly well with the theory that André had been developing for perceptual stabilization. Tomas later commented on my serendipitous encounter with this phenomenon by concluding: 'Het belangrijkste van visueel onderzoek is goed kijken'. This almost translates to: 'The most important aspect of vision science is looking good'. But not quite.

Abstract

We show that prior exposure to a monocular image can make that image gain initial dominance in subsequent binocular rivalry. We term this novel phenomenon *flash facilitation*. It is the converse of a known effect called flash suppression, where the previewed image becomes suppressed upon rivalrous presentation. The exact effect of previewing an image depends on both the duration and the contrast of the prior stimulus. Brief, low-contrast prior stimuli facilitate, whereas long, high-contrast ones suppress. These effects have both an eye-based component and a pattern-based component. Our results suggest that, instead of reflecting two unrelated mechanisms, both facilitation and suppression are manifestations of a single process that occurs progressively during presentation of the prior stimulus. The distinction between the two phenomena would then lie in the extent to which the process has developed during prior stimulation. This view is consistent with a neural model previously proposed to account for perceptual stabilization of ambiguous stimuli, suggesting a relation between perceptual stabilization and the present phenomena.

4.1 Introduction

Ambiguous images are images that have multiple, mutually exclusive interpretations. Notable examples are the Necker cube and Rubin's face-vase illusion. Binocular rivalry is a form of ambiguous perception that arises when we present two incompatible images to the two eyes, resulting in a percept that wavers between both images, involving the complete perceptual disappearance of the temporarily suppressed pattern. Binocular rivalry can be used as a tool for bringing to light subtle shifts in neural activity that accompany various experimental manipulations. For instance, changes in predominance of one percept over the other have been used to disclose effects of center-surround interactions on visual processing (Paffen et al., 2005), and to demonstrate the action of so-called Gestalt grouping cues (Alais and Blake, 1999). A particularly sensitive use of rivalry as an indicator of neural state, is to simultaneously switch on both competing images and observe which one gains dominance first. Here both neural representations start racing for dominance simultaneously, and even a slight imbalance in the two images' processing may considerably bias initial dominance. For instance, drawing attention to one of two, otherwise balanced, images can cause a threefold to fourfold shift in their initial dominance ratio (Chong and Blake, 2005; Kamphuisen et al., 2007; Mitchell et al., 2004). Another well known example is flash suppression (Kreiman, 2001; Wolfe, 1984): briefly showing one of the competing patterns in isolation can virtually eliminate the possibility of it winning the race upon subsequent rivalrous presentation. Flash suppression is a valuable experimental tool that allows investigators of visual perception and awareness to dictate their subjects' perception of an ambiguous stimulus (Kreiman et al., 2002; Sengpiel et al., 1995; Sheinberg and Logothetis, 1997).

Here we use initial dominance in binocular rivalry to systematically study how prior exposure to a pattern affects subsequent visual processing. We do this (Figure 4.1) by briefly presenting an image (prior stimulus), and then after a blank interval letting that image compete for dominance against one in the other eye (test). Surprisingly, prior exposure to one of the rival patterns in isolation can not only impair that pattern's initial dominance (flash suppression), but can also facilitate it, a novel phenomenon that we call *flash facilitation*. The exact effect of a prior stimulus depends on both its duration and its contrast. These two are largely interchangeable, so that (after correcting for a contrast transfer nonlinearity) the effect of a prior stimulus shows a systematic relation with their arithmetic product, here denoted as the 'energy' of the prior stimulus. Initial dominance is suppressed by high-energy prior stimuli, but facilitated by low-energy ones.

In additional conditions we vary parameters such as the eye of origin of the prior stimulus relative to the test stimulus. These manipulations indicate that both the facilitatory effect and the suppressive effect at least partly originate in lower visual areas, and that both are built up of an eye-based component and a pattern-based component. In the standard condition, these components add up to jointly produce a prior stimulus' effect.

We discuss our findings in relation to a growing body of literature demonstrating a dual effect of prior exposure in a variety of cognitive and perceptual tasks, as well as possibly related instances of facilitation in non-rivalrous vision. We argue that the most parsimonious interpretation of our findings is not one in which facilitation and suppression reflect entirely distinct mechanisms. Instead, our results are consistent with the idea of a continuous neural process that underlies both facilitation and suppression.

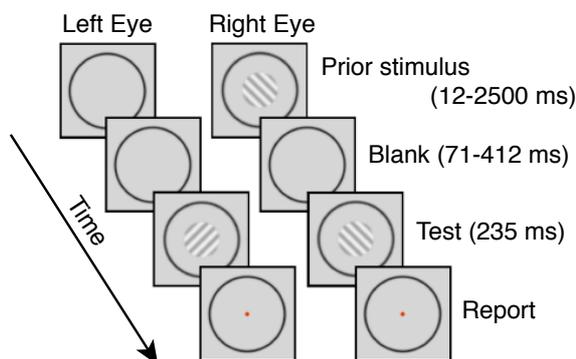


Figure 4.1: Time course of a trial in the main condition. Subjects consecutively viewed one half of a dichoptic orthogonal grating stimulus (prior stimulus), a blank interval, and both stimulus halves together (test). They then reported which of the two rivaling orientations was perceptually dominant during the test, revealing effects of the prior stimulus. We systematically varied the durations of the prior stimulus and the blank, as shown, as well as the contrast of the prior stimulus. In additional conditions we varied the nature of the prior stimulus, in order to specifically address eye-based and pattern-based effects, and to investigate the importance of the spatial correspondence between prior stimulus and test stimulus. We also varied the luminance of the background, to investigate if the relative luminance of the stimulus with respect to its surround affects the observed effects.

4.2 Methods

Each trial (Figure 4.1) consisted of the sequential presentation of a prior stimulus, a blank interval, a test, and a fixation dot prompting subjects to report. In most conditions the test was a pair of dichoptic orthogonal sine wave gratings, oriented ± 45 deg from vertical. In one condition (see below) we used square wave gratings. Subjects fixated the stimuli through a stereoscope at a viewing distance of 47 cm, within a black alignment ring ($r=1.1$ deg, or $r=2.3$ deg for the larger stimuli) on a gray background (30 cd/m^2 ; the stimuli's mean luminance) unless otherwise stated.

Seven subjects participated in the sessions of Figures 4.2 and 4.3, four took part in those of Figures 4.4 and 4.5. Two authors were subjects in all of these experiments, the others were naive.

Subjects reported which of the two test gratings was perceived more strongly. This instruction allowed them to make a choice even if dominance was incomplete. Subjects could discard a trial if they felt they were unable to make a choice, which happened on 2% of the trials. To verify that our subjects experienced strong rivalry suppression in our experiments we performed a control experiment (not shown) where normal trials were randomly interleaved with ones where the orthogonal test gratings were presented superimposed to one eye. In those conditions subjects generally perceived a fairly balanced plaid and they discarded 76% of the trials. This confirms that there was clear rivalry suppression during the test in our experiments.

In the main condition (Figure 4.1) we used one half of the pair of test gratings as a prior stimulus, to differentially address one of the two competing neural representations. Then, there were three conditions that were designed to tease apart the effects of prior stimulation of one eye, and prior exposure to a particular pattern, respectively.

1. In the *eye* condition, the prior stimulus was a monocular pattern that was unlike either test

grating, namely concentric rings. Because this prior stimulus did not specifically coincide with either test pattern, it allowed us to single out eye-of-origin effects.

2. The *pattern* condition had the complementary objective of isolating pattern related effects on initial dominance. Here, like in the main condition, the prior stimulus consisted of one of the two test patterns, but we presented it to both eyes. Because this prior stimulus did not specifically target either eye, it formed a probe into pattern related effects.
3. In the *swap* condition the prior stimulus consisted of one eye's test pattern, but presented to the opposite eye (i.e. the pattern was swapped between eyes in between prior stimulus and test). This condition bypassed any effects of the prior stimulus on monocular orientation channels, and allowed us to further constrain the source of the observed effects.

Two further conditions were designed to assess the importance of the spatial correspondence between the prior stimulus and the test stimulus.

1. The *annulus* condition was similar to the main condition, but here the prior stimulus was an annulus that surrounded the location of the test stimulus. Any retinotopic effects should disappear here.
2. The *phase shift* condition was similar to the *pattern* condition, but it involved a phase difference between the prior stimulus and test stimulus, allowing further inferences on the neural location of the effects.

In a final condition we investigated if the results from a classic flash suppression study (Wolfe, 1984) would be affected by changing the luminance of the stimulus' background. For that condition we used a square wave grating similar to the one used in the original study.

For all conditions the test duration was 235 ms, a duration sufficiently long to allow rivalry to develop (e.g. Chong and Blake, 2005; Wolfe, 1984, 1983) but too short for perceptual alternations to occur. Test contrast was 50% Michelson unless otherwise stated. For the main condition we parametrically varied the contrast and duration of the prior stimulus from 2.7 to 100% Michelson and from 12 to 2500 ms, respectively, and the duration of the blank from 71 to 412 ms. In the other conditions we applied only a subset of durations and contrasts, and used only a 71 ms blank unless otherwise stated. The sine wave gratings (3.6 cyc/deg) filled a circular patch of radius 0.55 deg, except in the phase shift condition, where we used a larger stimulus ($r=1.4$ deg; 1.4 cyc/deg) to minimize the effect of eye movements. The concentric rings used for the eye condition (5.2 cyc/deg) were designed to equate the gratings in total contour length. The ring pattern's contrast fell off along a Gaussian profile ($\sigma=0.14$ deg) at the edge so that even the outlines did not coincide between prior stimulus and test. In the annulus condition the annulus had an inner and outer radius of 0.61 deg and 1.1 deg, respectively. In the experiment replicating Wolfe's data and investigating the role of background luminance we copied Wolfe's original stimulus with some slight modifications. That is, we used diagonal square wave gratings ($r=1.35$ deg; 3.8 cyc/deg) and 90% Michelson contrast for the test stimulus.

We randomly interleaved the main condition, eye condition, pattern condition and swap condition within sessions. We also randomized the relation between orientation and eye during the test, as well as the eye receiving the prior stimulus. Randomization of conditions minimizes the interaction between consecutive trials, such as observed in 'perceptual stabilization'

(Leopold et al., 2002). Randomization of the eye and orientation has the same effect, as well as allowing us to control for systematic biases toward one orientation or eye. All data figures show pooled data over all eye-orientation combinations, with about 50 observations underlying each data point for each subject.

4.3 Results

4.3.1 Flash suppression and flash facilitation

Figure 4.2 shows the subject-averaged results from our main condition, in which the prior stimulus was one half of a pair of dichoptic orthogonal gratings (top right inset). The key finding here is that prior exposure can both suppress and facilitate dominance of the previewed pattern. The top panel shows results using a 71 ms blank between prior stimulus and test. Facilitation occurs when the prior stimulus is weak, being either of short duration (leftmost sections of the curves) or low contrast (lighter shaded curves). Suppression, on the other hand, follows long, high-contrast prior stimuli. Contrast and duration are to some extent interchangeable, as the curves shift rightward with decreasing contrast.

All facilitatory effects dissipate within hundreds of milliseconds after the offset of the prior stimulus, as seen in the bottom panels, which show the outcome at larger blank durations.

Note that these plots, as well as the remaining ones, represent pooled data from trials where the prior stimulus could be presented to either eye and could have either orientation, and the two test patterns were distributed to the eyes accordingly. Any systematic bias in rivalry dominance toward one eye or orientation will therefore cancel out in these representations, and one may safely interpret any deviations from 0.5 as effects of the prior stimulus.

4.3.2 Eye effects vs. pattern effects

Because in the main condition we used one of the two test patterns as a prior stimulus, our results could in principle be the result of prior stimulation of one eye, prior exposure to one pattern, or both. Figure 4.3 shows the outcome of experiments that dissociate these options. Here we define ‘pattern facilitation’ and ‘pattern suppression’ as the tendency for the previewed pattern to become dominant or suppressed during the test, respectively. The terms ‘eye facilitation’ and ‘eye suppression’ refer to dominance or suppression, respectively, of the eye that received the prior stimulus.

Isolating eye effects (Panel A) by stimulating one eye with a pattern unlike either test pattern (concentric rings), yields results similar to those in the main condition. Eye dominance is facilitated by brief or low contrast prior stimuli, and suppressed when using a longer duration or a higher contrast. We take this to be a general eye effect, not an orientation specific effect that is due to the minor orientation correspondence that exists between the rings and the gratings. This interpretation is strengthened by the fact that we get similar results when using a vertical grating or a luminance patch as a prior stimulus (data not shown).

Surprisingly, the complementary condition (Panel B), wherein we single out pattern specific channels by presenting one of the two test patterns to both eyes during prior stimulation, again produces comparable results. Thus, both eye-based effects and pattern-based effects show a similar dependence on the parameters of the prior stimulus, and they jointly underlie the outcome of our main condition.

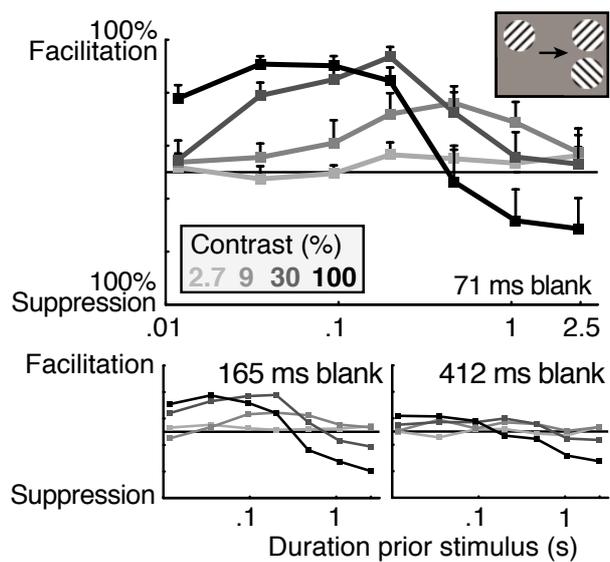


Figure 4.2: Dominance as a function of the duration and contrast of the prior stimulus, in the main condition. The top panel shows the results for the shortest blank duration (71 ms). Considering first the black curve (100% prior stimulus contrast) we see a biphasic effect of the duration of the prior stimulus. Brief exposure to one stimulus-half enhances its subsequent dominance (facilitation), but exposure durations over about .5 s have the opposite effect (suppression). The other curves in this panel indicate that decreasing the contrast of the prior stimulus causes this evolution to slow down, as well as reducing its amplitude. These differences are significant: Taking, for all subjects separately, the location and height of each curve's highest point, contrast correlates negatively with log peak time ($p < 0.05$), and positively with peak amplitude ($p < 0.05$). The two bottom panels show a rapid decay of facilitation at increasing blank durations, and much less so of suppression. Error bars (not shown in the bottom panels) are standard errors of the mean ($n=7$; see Appendix for individual subjects' data).

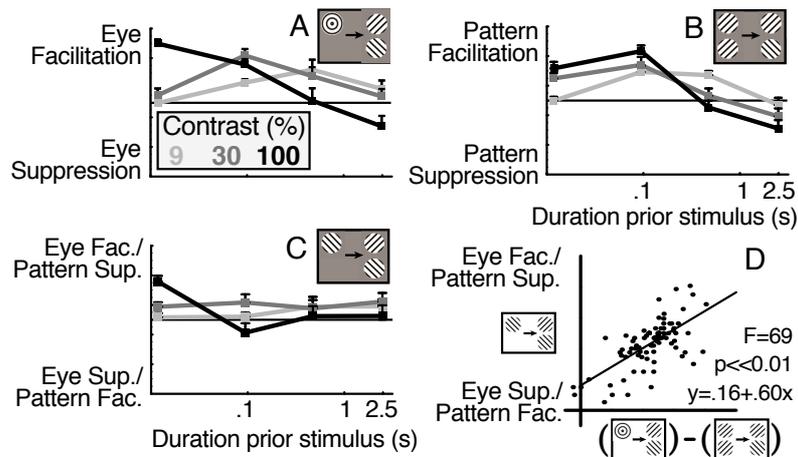


Figure 4.3: *Dissociating eye-based effects and pattern-based effects. A. Eye condition. Prior stimulation of one eye using a pattern that is unrelated to either test pattern causes facilitation and suppression much like in the main condition. The converse condition of previewing one test pattern dioptically (B, pattern condition) again has a comparable effect. Clearly, eye-based effects and pattern-based effects obey similar laws, and both contribute to the outcome of our main condition. If as a prior stimulus we let one eye view the opposite eye's test pattern (C, swap condition), the pattern-based effects and eye-based effects largely cancel out. In panel D the y-axis depicts the outcome of the swap condition for each combination of subject, contrast and duration. The x-axis depicts the difference between the outcome of the eye condition and the pattern condition, for the corresponding combinations. This difference turns out to form a fair prediction for the outcome of the swap condition. This underscores that eye-based effects and pattern-based effects act against each other in the swap condition. Panels A-C show subject-averaged data; error bars are standard errors of the mean ($n=7$). Like in the main condition, contrast correlates negatively with log peak time ($p < 0.05$) and positively with peak amplitude ($p < 0.05$) for the curves of the eye and pattern condition. For the swap condition the correlation with peak amplitude remains but the one with log peak time is no longer significant.*

The effects shown in Figure 4.3A (eye condition) necessarily originate at processing levels where both monocular streams have not yet fully converged. What channels underlie the data in panel B (pattern condition), however, is less clear-cut. It seems clear that pattern selective channels are involved, but are these monocular or binocular? Panel C addresses this question, showing what happens when we let one eye preview the other eye's test grating (swap condition). This bypasses effects in monocular orientation detectors, so if these are responsible for the outcome of the pattern condition, we expect only a general eye-based effect here, and predict results similar to those of panel A (eye condition). If, in contrast, the outcome of the pattern condition has a binocular origin, it is pitted against eye-based effects here, presumably yielding an outcome that is a trade-off between both forces. Panel C strongly points to the latter option, showing a result that is clearly different from that of the eye condition in panel A, and that could well be due to the opposed action of eye-based effects and pattern-based effects. In fact, the curves in panel C are similar to the arithmetic difference between those found for the eye condition (panel A) and those found for the pattern condition (panel B). This is quantified in panel D. The y-axis here depicts the outcome of the swap condition, for each combination of subject, contrast and duration separately. The x-axis depicts a prediction for the corresponding combinations, calculated by subtracting the outcome of the pattern condition from that of the eye condition (as indicated by the icons on the axis). Although this linear subtraction is arbitrary, the positive trend is clearly consistent with a trade-off between both forces. These findings therefore indicate that the pattern-based effects at least partly arise at stages beyond binocular convergence, and exist independent of the eye based-effects.

Incidentally, the findings in Figure 4.3 also rule out the option that facilitation is due to a cognitive bias to report the previewed pattern. First, none of the rival patterns was previewed in the eye condition (Figure 4.3A). Second, the conditions of Figures 4.2 and 4.3 were interleaved randomly and could not be distinguished by the subjects, yet the tendency to report the previewed pattern is absent in the condition of Figure 4.3C (see Appendix for further evidence against this concern).

4.3.3 *Retinal location and phase*

To further elucidate the neural basis of the observed phenomena, we tested their dependence on the spatial correspondence between the prior stimulus and the test stimulus: their relative retinal location and relative phase. Figure 4.4A shows data from a condition where the location of the prior stimulus did not coincide with that of the test patch, but instead surrounded it. In this condition the facilitatory effect is largely abolished, indicating a high degree of retinal specificity. Regarding suppression the data are less conclusive, as suppression was weak even in the baseline condition. Figure 4.4B shows that phase differences have no influence if the prior stimulus has a low contrast, but that at a high contrast they do cause a shift toward suppression. This combination of retinal specificity, partial (but not complete) phase dependence and both monocular and binocular components (Figure 4.3) suggests a distributed neural origin with an emphasis on lower visual areas.

4.3.4 *Comparison with flash suppression studies*

We wondered how the current experiments could give rise to such pronounced facilitatory effects, considering that previous studies found only suppressive effects of showing one of the

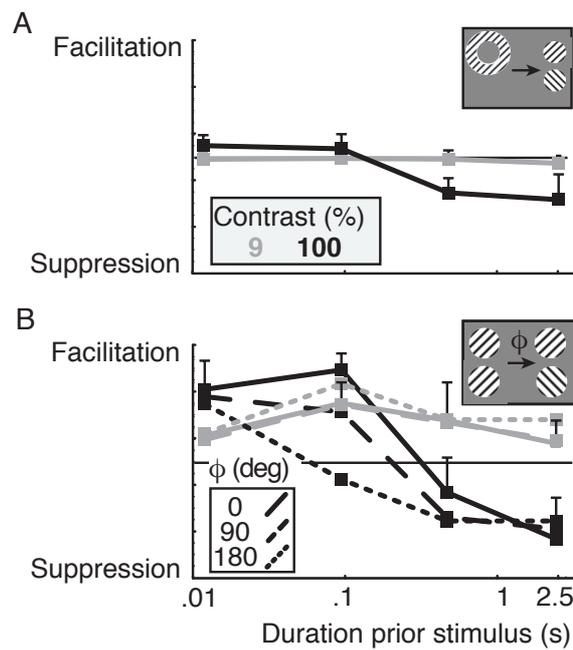


Figure 4.4: The influence of the spatial relation between the prior stimulus and the test stimulus. A. When the retinal location of the prior stimulus does not coincide with that of the test stimulus, its effect is negligible (None of the points here differs significantly from chance; two-sided t-test; $p > 0.05$). B. A phase difference (ϕ) between the prior stimulus and the test stimulus preserves the basic effect, but with an increased tendency toward suppression when using a high contrast prior stimulus. This difference is significant for the 90 ms prime, when comparing $\phi=0$ deg with $\phi=180$ deg (two-sided paired t-test; $p < 0.01$). These features, together with those shown in Figure 4.3, are consistent with a distributed neural basis that includes lower visual areas. Error bars, only shown for $\phi=0$ deg in panel B, are standard errors of the mean ($n=4$).

rival patterns as a prior stimulus. Although this can partly be attributed to the long, high contrast prior stimuli that are often used, we explored two additional options.

First we examined whether the fact that many studies excluded a blank interval between the prior stimulus and the test could be a relevant factor. Figure 4.5A shows results from an experiment where we systematically varied the blank duration, using the same stimulus as in our main condition. The plot shows that without a blank interval (blank duration 0 s) suppression is somewhat stronger than at the blank duration we used (71 ms; dashed vertical line). More specifically, it appears that at blank durations above about 50 ms there is a continuous evolution of the effect of the prior stimulus. In some cases (12.5 and 25% contrast) this is simply a gradual decline, and in others (50 and 100%) this involves an excursion toward suppression before the effect is extinguished. At shorter blank durations (0 and 12 ms, leftmost points) the curves appear to depart from this gradual progression, and show a shift toward suppression (this discontinuity is marked by dashing of the curves).

A second factor we examined is the luminance of the background on which the stimuli are presented. It is common to use a black background in experiments such as these, whereas in the present experiments the background was gray, equal to the mean luminance of the stimuli. To make a more direct comparison with existing flash suppression studies we copied the stimulus from a classic study that used a black background (Wolfe, 1984). This is a square wave grating somewhat larger than our basic stimulus. In one condition we attempted to replicate that study's results by using the original black background, and in a second condition we raised the background luminance to equal the mean luminance of the stimulus. Figure 4.5B depicts the outcome of these experiments. Using a black background (top) we found robust suppression, in agreement with the original study (the star indicates the settings of that study's main experiment). With this background luminance, facilitation remained weak or absent even at a lower contrast and duration of the prior stimulus. However, when the background luminance was increased (bottom) we again observed the familiar pattern of strong facilitation and slight suppression seen in our main experiments. Clearly, background luminance is another important factor affecting the balance between flash facilitation and flash suppression. Note that in Wolfe's original experiments suppression may have been even stronger, due to a much higher luminance difference between background and stimulus. We could not apply this difference because the mean stimulus luminance on Wolfe's tachistoscope (398 cd/m^2) is well beyond our CRT monitor's capacity (our mean luminance was 30 cd/m^2).

4.3.5 *Energy of the prior stimulus determines facilitation and suppression*

We have shown that the contrast and duration of the prior stimulus are to some extent interchangeable: lower-contrast prior stimuli need to be presented longer to reach their maximal effect. Therefore, we may be able to unify the outcomes from all contrasts by using some measure of stimulus energy, i.e. 'power' \times duration. This is shown in Figure 4.6A. Here we replot the data from Figure 4.2, as a function of the prior stimulus' contrast^{0.75} \times duration. Plotted in this way, the course of the curves practically coincides for all contrasts (For instance, compare the locations of the maxima).

A parsimonious explanation of this invariant 'energy' is that facilitation and suppression depend on a continuous neural process that occurs progressively during the presentation of the prior stimulus, and at a rate proportional to the contrast of that stimulus. This is illustrated schematically in Figure 4.6B. Here a hypothetical signal rises linearly during the presentation

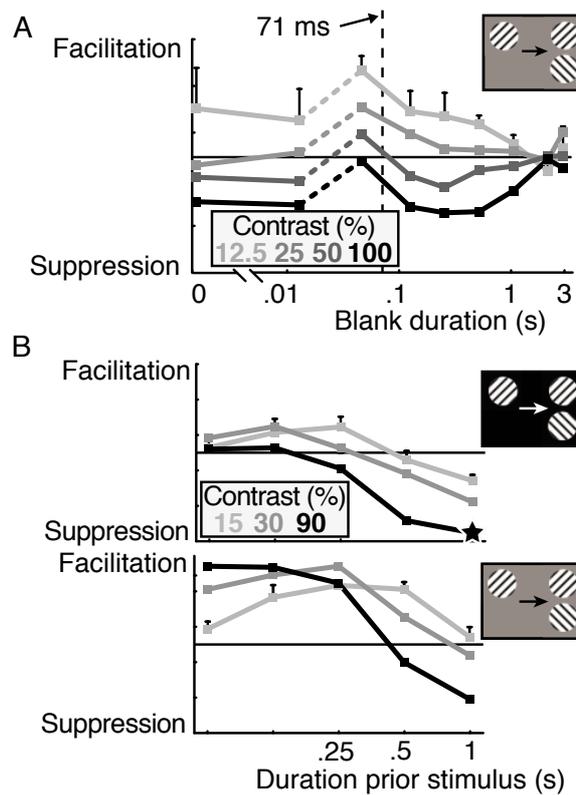


Figure 4.5: Factors that shift the balance between facilitation and suppression. *A. Removing the blank interval.* At blank durations above about 50 ms (we mainly used 71 ms, dashed vertical line) there appears to be a continuous evolution of the effect of the prior stimulus. Shorter blanks, however, (0 and 12 ms, leftmost points) yield an increased suppression that seems discontinuous with the remainder of the curves (marked by dashed). This shows that leaving out the blank interval weakens facilitation. The duration of the prior stimulus was 800 ms here. *B: Background luminance.* Top: using a stimulus from a classic flash suppression study (Wolfe, 1984) we found mainly suppression, even if the contrast and duration of the prior stimulus were relatively low. After raising the background luminance from black to the stimulus' mean luminance, however, the same stimulus produced mainly facilitation (bottom). Background luminance, therefore, is a key parameter. Suppression is significantly stronger in the black background condition for all points in these plots, apart from the two briefest prior stimuli at minimal contrast and the longest one at maximal contrast (one-sided paired t-test on individual subjects' data; $p < 0.05$). Error bars (only shown for lowest contrasts) are standard errors of the mean ($n=4$).

of the prior stimulus (Figure 4.6B, top panel), with a rate of rise that is a function of its contrast (light versus dark curve). If we further assume that a low end-level of this signal produces facilitation, whereas a high end-level yields suppression, we obtain an outcome (bottom panel) very similar to our experimental results (see Discussion for a further remark on this assumption).

Within the context of this account, the deviations between the curves' amplitudes, evident both in our empirical data (Figures 4.2-4.5) and the bottom panel of Figure 4.6B, are a natural consequence of adding noise to the accumulating signal. That is, a lower rate of rise causes a greater spread in the total accumulated value, as indicated by the shaded areas in Figure 4.6B, top. Consequently, there is a lower peak probability for the accumulated value to lie within the 'facilitation region' at any given time.

The fact that our data line up nicely when using an exponent of 0.75 for calculating stimulus 'energy', suggests in the context of this explanation that the rate of rise of the signal is proportional to contrast^{0.75}. If we take this rate to be proportional to the instantaneous neural response, an exponent of 0.75 compares favorably with the compressive contrast responses found in lower visual areas (e.g. Carandini et al., 1997; Sclar et al., 1990).

Clearly, the illustration of a single signal and linear rise to threshold in Figure 4.6B is a simplification. Nevertheless, the observed invariant 'energy' and the increase in curve amplitude with contrast (Figure 4.6A) are consistent with the view that facilitation and suppression do not depend on two unrelated neural processes, but instead on the level of advancement of a progressive neural process. We elaborate on this idea in the Discussion section.

4.4 Discussion

4.4.1 Facilitation and suppression

We have shown that prior exposure to a pattern can both facilitate and suppress its initial dominance in binocular rivalry. Facilitation occurs with prior stimuli that are short and/or of low contrast, whereas suppression is observed using high-contrast, long-duration prior stimuli. These effects are in part eye-based and in part pattern-based. They are also to a considerable extent retinotopic, and they show a partial dependence on spatial phase. Our findings stand among a number of examples where processing is inhibited by prior stimuli of long duration, but enhanced by brief ones. An overview of such instances is given by Huber and O'Reilly (2003), who focused mainly on higher-level cognition, but also within the realm of ambiguous perception there is accumulating evidence that this is a general feature. For instance, in structure-from-motion rivalry prolonged prior exposure to a disparity-disambiguated stimulus impairs initial dominance (Nawrot and Blake, 1989), but brief exposure has the converse effect (Brascamp et al., 2008). Such antagonism between long and short prior stimuli has further been found in ambiguous motion perception (Kanai and Verstraten, 2005) and Necker cube rivalry (Long et al., 1992). In binocular rivalry, however, the facilitatory property has not previously been reported.

Our data suggest that the reason that flash facilitation has hitherto been overlooked in binocular rivalry, is that experimenters commonly use relatively long, high-contrast prior stimuli, no blank interval prior to the test, and a black background. All these factors shift the balance between facilitation and suppression toward the latter. In hindsight, influences of the blank interval and background luminance seem reasonable. Without a blank interval the added

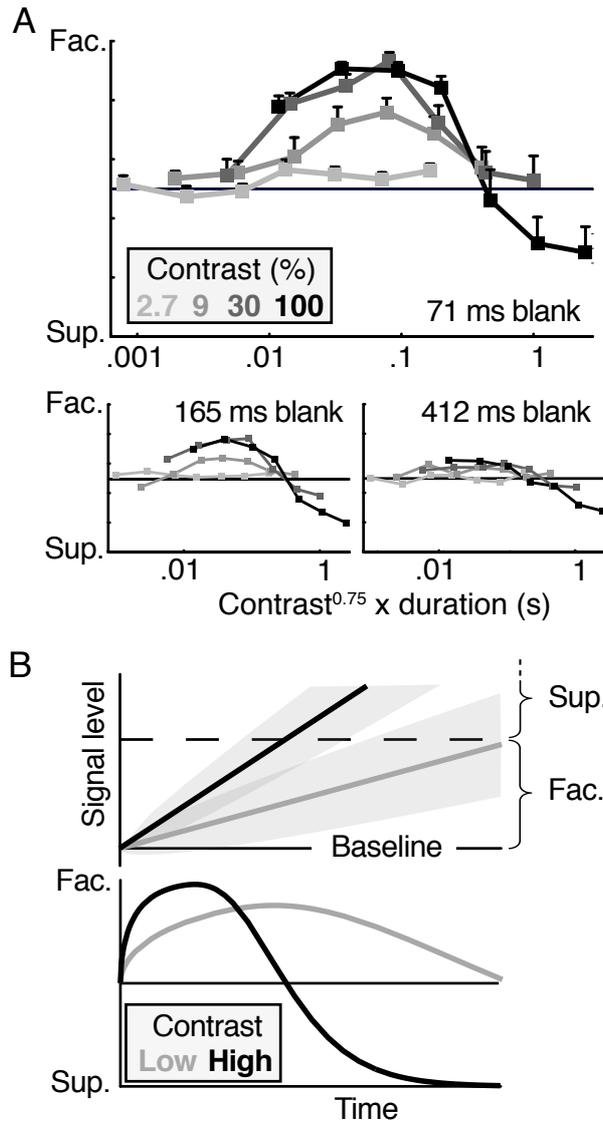


Figure 4.6: A single progressive process may underlie both facilitation and suppression. A. The data from Figure 4.2 replotted as a function of the ‘energy’ of the prior stimulus: $\text{contrast}^{0.75} \times \text{duration}$. In this representation the course of all curves corresponds closely (the exponent of 0.75 was chosen by eye). This suggests a mechanism such as schematically illustrated in panel B. A signal builds up linearly during the presentation of the prior stimulus (top panel), at a rate proportional to $\text{contrast}^{0.75}$. If by the end of the prior stimulus’ presentation the signal lies between baseline and a certain threshold (dashed line) this yields facilitation, whereas if it rises beyond this threshold or ends up below baseline we get suppression. Such a scenario would also explain the different curve amplitudes as due to different noise levels in the accumulated signal (shaded areas). That is, a signal that is twice as weak needs twice as much time to reach the same end level, but, assuming it is integrated with additive noise, four times as much to reach the same signal-to-noise ratio. The bottom panel shows facilitation and suppression corresponding to the top panel signals.

stimulus half may benefit from the strong transient response at its onset (Keyser and Perrett, 2002). In fact, a nearby stimulus onset can cause a percept to fade even without interocular conflict (Kanai and Kamitani, 2003; Wilke et al., 2003). Regarding background luminance, a black background could aid suppression because it increases the strength of the prior stimulus in several ways. First, the observed eye-based effects may be partly luminance based; second, there is a large contrast step between the stimulus and a black background; and third, a black background means a higher *effective* stimulus contrast in case luminance gain control integrates over an area larger than the stimulus itself.

Most previous studies that reported combined facilitatory and suppressive effects of prior stimulation, ascribed these two effects to two entirely distinct mechanisms. Suppression is often thought to be the result of satiation or ‘fatigue’ type processes in neurons that code the pre-viewed stimulus (e.g. Huber and O’Reilly, 2003; Long et al., 1992), whereas facilitation has been attributed to a number of other factors, including cognitive expectancy (Long et al., 1992) and integration of the response to the test stimulus with a persisting response to the prior stimulus (Huber and O’Reilly (2003); Georgeson and Georgeson (1987); Pinkus and Pantle (1997)). We do not think cognitive expectancy can explain the present data because these bear the stamp of a fairly low-level phenomenon, depending in part on retinal location, eye of origin and stimulus phase. Similarly, temporal integration of the prime and test response does not seem adequate because it can only apply to our situation if it takes place at relatively high processing stages, where cells’ temporal properties fit the data (Huk and Shadlen, 2005; Keyser and Perrett, 2002). The faster neural responses in lower visual areas (Breitmeyer, 1984) might just be persistent enough to allow facilitation to survive a 165 ms blank interval, but it is unlikely they could explain the buildup of facilitation over up to 500 ms (Figure 4.2). If not expectancy or neural persistence, what can explain the facilitation we observe? The analysis of Figure 4.6 indicates that, instead of assuming entirely separate mechanisms for facilitation and suppression, it may be useful to think of both as distinct products of one continuous neural mechanism.

4.4.2 A possible link with perceptual stabilization

In light of this idea it is relevant to discuss a recent publication by Noest et al. (2007) that addresses the effects of *ambiguous* prior stimuli on initial dominance. An established example of such effects is so-called perceptual stabilization of ambiguous stimuli. This refers to the tendency, when an ambiguous stimulus is periodically removed from view, for the same percept to gain dominance on many consecutive reappearances (Leopold et al., 2002; Orbach et al., 1963; Pearson and Clifford, 2004). Arguably, this stabilization reflects the repeated action of a facilitatory effect of dominance during one (prior) ambiguous presentation, on dominance during the following (test) presentation. Noest and colleagues systematically varied the timing of such an intermittent presentation cycle, and found that in certain temporal regimes this facilitatory effect changes into an opposite, suppressive effect (see also (Orbach et al., 1966)). That is, in those timing regimes subjects tend to see the percept opposite to the previous percept on each presentation.

Noest and colleagues account for these findings using a model that treats both the facilitatory and suppressive effect as distinct manifestations of one continuous neural process. Specifically, in this model the rivalrous prior stimulus exerts its effect via progressive sensitivity changes, or adaptation, of neurons coding the dominant percept. At low adaptation levels the model predicts facilitation, but at higher adaptation levels it predicts suppression. This

mechanism is very similar to the one suggested in Figure 4.6, where suppression and facilitation are also proposed to depend on the degree of progression of a cumulative process that occurs during the presentation of the prior stimulus. Indeed, in simulations with this model we can reproduce our present data (see Appendix), when using a non-rivalrous prior stimulus as input instead of the rivalrous input for which the model was designed. This result indicates that the effects of unambiguous prior stimuli observed in the present work may reflect similar mechanisms as the effects of ambiguous prior stimuli observed elsewhere.

Having said this, it is important to point out several nuances. First, in Figure 4.6B facilitation is simply postulated to change into suppression at some arbitrary threshold level of the accumulating signal. In the model by Noest and colleagues, on the other hand, this transition between both behaviors is more natural. There, it is an automatic consequence of the assumption that adaptation not only has a conventional divisive effect on activity, but also has a slight additive effect. Which of the two opposed forces decides dominance at stimulus onset is determined, in part, by the level of adaptation, leading to the observed transition from facilitation to suppression (Noest et al., 2007). Second, although we believe that ambiguous and unambiguous prior stimuli may both engage similar neural mechanisms, there is at least one clear difference. An unambiguous prior stimulus differentially adapts the corresponding neural processing stream from the very lowest level upward. An ambiguous prior stimulus, on the other hand, contains equal evidence for both percepts, and therefore arguably adapts both processing streams equally up to some level where the conflict is resolved. Only from this level onward, there is greater adaptation of the dominant representation. Differences such as these may contribute to known phenomenological differences between the effects of ambiguous and unambiguous prior stimuli. For instance, contrary to the situation with unambiguous prior stimuli, the facilitating effect of ambiguous prior stimuli does not require the prior stimulus to be of low contrast. Also, the facilitating effect of ambiguous prior stimuli can survive longer blank intervals than the effects observed in the present work (Leopold et al., 2002).

4.4.3 *Other related phenomena*

A phenomenon that seems related to flash facilitation is the facilitatory influence of attention on rivalry dominance. Cueing attention to a pattern enhances its probability of gaining initial dominance in subsequent rivalrous viewing (Chong and Blake, 2005; Kamphuisen et al., 2007; Mitchell et al., 2004). It is likely that with flash facilitation we have at least partly probed the same mechanisms as researchers who studied attention effects. For one thing, the onset of the prior stimulus in our experiments likely draws attention to the previewed pattern (e.g. Corbetta and Shulman, 2002), so in that sense existing attention studies may provide a partial explanation for the present findings. Conversely, we may also consider that the present findings provide a new perspective on existing attention studies. On a neural level attending to a pattern is in some respects similar to a slight increase in its contrast (e.g. Martínez-Trujillo and Treue, 2002; Reynolds et al., 2000), so our conclusions regarding facilitation may in part also apply to facilitation due to attention. Our results suggest that a kind of gradual gain control underlies the present observations, raising the possibility that a similar mechanism contributes to the attention effects on initial dominance observed in previous studies. We should note, however, that there is at least one qualitative difference between flash facilitation and attention effects, namely that the latter seem to lack an eye-based component. When attention is drawn to a monocular stimulus that is then presented to the other eye during the test, it stays associ-

ated with the primed pattern (Kamphuisen et al., 2007), with no sign of an opposed eye-based effect (cf. Figure 4.3C).

Prior exposure can also have facilitatory and inhibitory effects in situations that do not involve perceptual ambiguity. A well-known example is masking, altered detection or judgment of a pattern due to an immediately preceding ‘mask’ stimulus. Like investigators of flash suppression (Kreiman et al., 2005; Wolfe, 1984) we argue that masking is distinct from the present findings, because it shows very different parameter dependencies. Among other features, facilitatory masking effects are restricted to mask-stimulus intervals below about 50 ms (Georgeson and Georgeson, 1987), whereas we find facilitation even at 165 ms blank intervals (Figure 4.2).

There are several other examples where a visual pattern is detected more easily or elicits a faster response, if it has previously been presented. Like the present effects, those phenomena often have characteristics that indicate the involvement of sensory cortical areas (Campana et al., 2002; Tanaka and Sagi, 1998), and their temporal properties do not exclude an overlap with the present phenomena either. It seems unlikely that the mechanisms probed in the present study are entirely restricted to situations of perceptual ambiguity. We think it would be an interesting subject of future research to search for a direct association between effects of prior stimuli on initial dominance on the one hand, and on detection or reaction times on the other.

Appendix

Individual subjects

Figure A4.1 shows the individual subjects’ data that underlie the top panel of Figure 4.2. All subjects show the same trends, although there are clear differences both in the temporal evolution of the patterns, and in the relative importance of facilitation versus suppression.

The effect of a prior stimulus on orientation discrimination

Because our main measure consisted of observers’ subjective perceptual judgment, we wished to make sure that our findings were not due to a cognitive bias. The results from Figures 4.2 and Figures 4.3 were obtained within the same sessions by randomly interleaving conditions, and the fact that these figures show very different outcomes for subjectively identical prior stimuli (depending on eye of origin), took away much of our concern. Nevertheless, two authors performed an additional control experiment where we quantified the effect of a prior stimulus objectively, as change in orientation discrimination performance. We used the experimental sequence of our main condition but with the instruction to report whether the right-tilted test grating was rotated clockwise or counterclockwise from 45 deg. In agreement with our previous results, for brief prior stimuli performance was best when the prior stimulus was the right-tilted grating (flash facilitation, left column in Figure A4.2), and for long prior stimuli performance was best when it was the opposite grating (flash suppression, right column in Figure A4.2). The effect of prior stimulation does therefore not depend on a subjective dominance judgment.

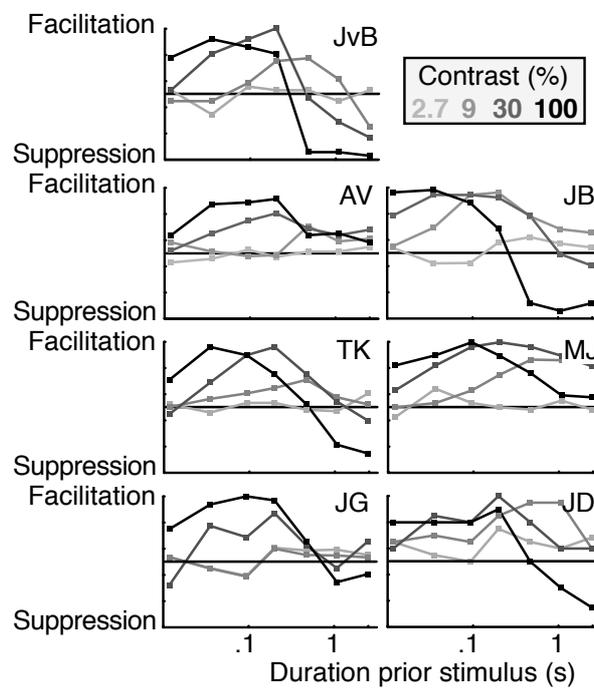


Figure A4.1: Individual subjects' data from our main condition, at a 71 ms blank duration. All subjects show the same trends.

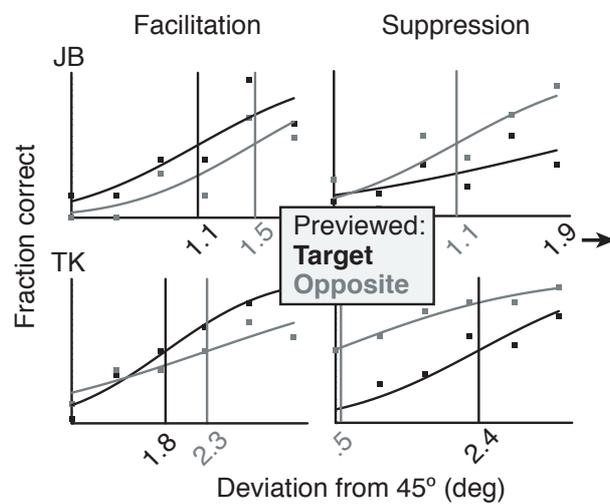


Figure A4.2: Orientation discrimination performance as a measure of facilitation and suppression. In the case of a weak prior stimulus (100% contrast, 94 ms), performance was better if the prior stimulus was identical to the grating on which an orientation judgment was required. With a strong prior stimulus (100% contrast, 2.5 s) this pattern reversed. This indicates that flash suppression and flash facilitation are not due to a cognitive bias in subjects' subjective dominance reports. Vertical lines mark the means of the cumulative Gaussians fitted to the data. We determined subjects' baseline thresholds beforehand in a staircase procedure, and made sure we used prior stimuli that in the main condition yielded robust facilitation and suppression, respectively. All points represent 40 measurements. The effects of facilitation and suppression are significant both within subjects (bootstrap, $p < 0.05$) and combined over both (one-sided paired t-test, $p < 0.05$).

A model for the effects of rivalrous and non-rivalrous prior stimuli

Like many rivalry models (Wilson, 2005), the model by (Noest et al., 2007) comprises, for each of the competing neural representations, two differential equations: one for the representation's activity, and one for its adaptation. The time derivative of activity H_i is given by

$$\tau \partial_t H_i = X_i - (1 + A_i)H_i + \beta A_i - \gamma S[H_j] \quad (\text{A4.1})$$

That of adaptation A is given by

$$\partial_t A_i = -A_i + \alpha S[H_i] \quad (\text{A4.2})$$

Here $\tau \ll 1$ is a time constant, H is activity, X is input strength and A is adaptation. γ and α are constants governing the strengths of, respectively, cross inhibition between the representations and adaptation. A notable difference between this model and other ones is the term βA_i in equation A4.1, which describes an additive effect of adaptation on activity. It is this term that allows adaptation A to facilitate dominance in some conditions. Subscripts i and j refer to the two percepts, and $S[X]$ is a sigmoidal function of X , here we used $S[X] = \frac{X^2}{1+X^2}$ for $X > 0$, and $S[X] = 0$ otherwise. Simulations were performed at the following parameter settings. During the test phase $X_i=X_j=0.6$ and during the blank $X_i=X_j=0$. For the prior stimulus $X_i=0$ and $X_j=0.4$ or 1.0 to simulate a low or high contrast, respectively. The constants were set to $\tau=0.02$, $\gamma=3.3$, $\alpha=4$, and $\beta=0.2$. The duration of the prior stimulus was varied from $10^{0.1}$ to $10^{2.1}$, and the blank interval was 8. The one modification we made to the original model is that we included a noise term in the adaptation equation, namely additive Gaussian noise of $\mu=0$ and $\sigma=2.5$, at an integration step size of 0.1. Figure A4.3 shows that at these settings the model qualitatively reproduces our main result.

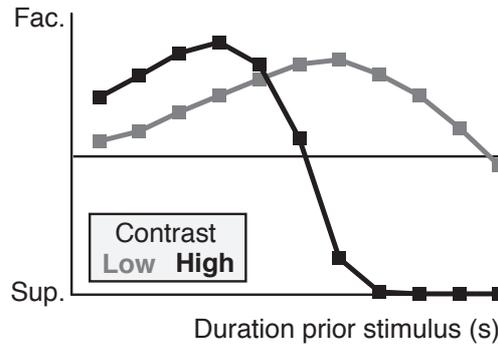


Figure A4.3: Simulation results using the model by Noest and colleagues. Although developed to describe the effect of a rivalrous prior stimulus on initial dominance, the model qualitatively accounts for our data as well.

As a final remark, note that we focus on the general principle that the mechanism employed by this model can explain our data. A quantitative application of the model to our data is complicated by the fact that in our experiments differential adaptation occurs, not only at and beyond the level where rivalry is resolved (A), but also at all stages leading up to it. This causes an asymmetry in the input strengths X (Figure A2 in Noest et al. (2007), see also Discussion section). Such lower stages have successfully been added to the model to explain other effects that

presumably originate at those lower stages (Klink et al., 2008). However, in our case expanding the model with lower stages would allow for too much freedom (for instance, adaptation time constants and contrast gain functions could be chosen for each stage independently), leaving it underconstrained by our data.

Chapter 5

Multi-timescale perceptual history resolves visual ambiguity

Published as: Brascamp, J. W., T. H. J. Knapen, R. Kanai, A. J. Noest, R. van Ee and A. V. van den Berg (2008) Multi-timescale perceptual history resolves visual ambiguity. *PLoS ONE*, 3(1): e1497

I think this is an important chapter. The results in this chapter genuinely made me think of a subject – in this case perceptual stabilization – in a different way. Building on the basis that André and coworkers had laid in a 2007 paper (Noest et al., 2007), I think this work brought us a lot closer to understanding how perceptual stabilization really works. Although I like this chapter, I am not entirely happy about the writing. I remember how difficult I found the writing process this time. There was an exotic experimental design that had to be explained. There was the wish to write for a broader audience while also conveying the details. But perhaps most importantly, by the time we started writing, our view of the subject had drifted substantially away from ideas that were current at the time. This meant our text had to start off with an introduction rooted in ideas and concepts that we were not comfortable using anymore, end with a conclusion in the vein of our own thinking, and still have a flowing story in between. To put this dilemma another way, we found it hard to explain even why our experiment made any sense at all, without first showing its conclusion. By the time we submitted the manuscript, we had gone through so many versions that I could not really judge it objectively anymore. When recently I read it over for the thesis, the first half of the text still seemed rather forced to me. On a positive note, by the time I reached the results of Figure 5.3 I did remember why I had been so excited about the work, and I think that anyone who has ever made it that far into the text should have at least somewhat of a different perspective on perceptual stabilization than they had before they started reading.

Abstract

When visual input is inconclusive, does previous experience aid the visual system in attaining an accurate perceptual interpretation? Prolonged viewing of a visually ambiguous stimulus causes perception to alternate between conflicting interpretations. When viewed intermittently, however, ambiguous stimuli tend to evoke the same percept on many consecutive presentations. This perceptual stabilization has been suggested to reflect persistence of the most recent percept throughout the blank that separates two presentations. Here we show that the memory trace that causes stabilization reflects not just the latest percept, but perception during a much longer period. That is, the choice between competing percepts at stimulus reappearance is determined by an elaborate history of prior perception. Specifically, we demonstrate a seconds-long influence of the latest percept, as well as a more persistent influence based on the relative proportion of dominance during a preceding period of at least one minute. In case short-term perceptual history and long-term perceptual history are opposed (because perception has recently switched after prolonged stabilization), the long-term influence recovers after the effect of the latest percept has worn off, indicating independence between time scales. We accommodate these results by adding two adaptation terms, one with a short time constant and one with a long time constant, to a standard model of perceptual switching.

5.1 Introduction

The visual system adjusts its processing of current input on the basis of past experience. Such dynamic adjustment allows, for instance, faster responses to recurrent stimuli (Maljkovic and Nakayama, 2000) and tuned weighting of visual cues depending on their previous validity (Ernst et al., 1999). A fundamental question for such adaptive systems is *how long* a history to incorporate in current processing.

An opportunity to examine the role of history in vision within a controlled experimental setting is provided by ambiguous stimuli (Figure 5.1A). These images convey conflicting information to the eyes, causing perception to waver randomly between alternative interpretations, or percepts (Figure 5.1B, top). For instance, if a movie of a transparent revolving sphere with dots on its surface is stripped of all unambiguous depth information, such as perspective and occlusion, it causes alternating perception of either possible rotation direction (Figure 5.1A, left). Alternatively, presenting incongruent images to the two eyes simultaneously causes alternating perception of either image in isolation (binocular rivalry; Figure 5.1A, right). Strikingly, prior experience can allow the state of perceptual indecision brought about by ambiguous stimuli to be overcome. That is, when observers are presented with an ambiguous stimulus they have viewed before, they often instantly perceive the same interpretation as they did on the prior encounter, even though the immediate visual input remains inconclusive. This can lead to prolonged periods of perceptual stabilization if an ambiguous stimulus is periodically removed from view and the same percept keeps reappearing on consecutive presentations (Figure 5.1B, bottom) (Pearson and Clifford, 2005; Maloney et al., 2005; Chen and He, 2004; Noest et al., 2007; Orbach et al., 1963; Maier et al., 2003; Leopold et al., 2002). This salient expression of visual memory provides a convenient measure to study how traces of past perception interact with current input in shaping what we see.

Perceptual memory for ambiguous stimuli has a persistence of at least minutes, in the sense that even if an ambiguous stimulus does not reappear until several minutes after disappearing, the previous percept often still recurs (Leopold et al., 2002). Does this imply that a minute-scale perceptual history is incorporated in processing current visual stimuli? On the contrary, considering that the percept that is experienced (or *dominant*) at reappearance is generally simply the one that also dominated during the most recent encounter, one is tempted to conclude that only a single percept is stored, and that memory is ‘overwritten’ whenever perception changes (*switches*) to the alternative interpretation. Memory would then reach back no further than the moment of the latest switch. This is indeed implied by the common view that perceptual stabilization of ambiguous images reflects persistence after stimulus removal of the present state of perceptual organization (Pearson and Clifford, 2005; Chen and He, 2004; Orbach et al., 1963; Maier et al., 2003; Leopold et al., 2002). It would appear that a system centered on persistence of the present dominance state could store only a single percept at a time.

Within the broader context of history dependence in vision this view is remarkable. If the goal of a visual memory system is to optimize processing based on past experience, storage of a single percept or event is of limited use. Processing would benefit from incorporating a more elaborate record of past events. Indeed, the literature does contain indications that persistence of the current perceptual state may be insufficient to explain history effects in ambiguous vision. For instance, if an ambiguous stimulus is removed from view only shortly after a perceptual switch occurred (under about 2 s), the percept that dominated before the switch – not the most recent percept – often regains dominance at stimulus reappearance (Leopold

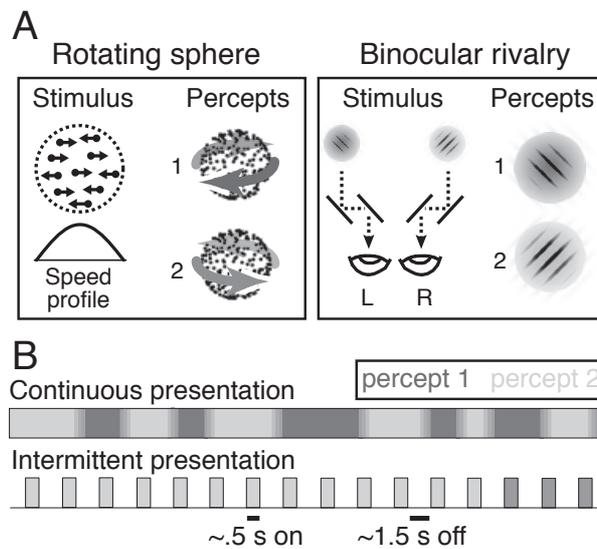


Figure 5.1: Stimuli and presentation sequences. *A.* Each of our stimuli has two distinct perceptual interpretations. Only one percept is experienced at any given moment. *Left:* an ambiguous rotating sphere is a two-dimensional projection of dots covering the surface of a transparent sphere that rotates around a central axis. Because no cue indicates which dots are in front, the rotation direction is ambiguous and subjects perceive either direction in turn, as indicated by the arrows. *Right:* in binocular rivalry ambiguity arises because two incompatible images are projected into the two eyes (designated as ‘L’ and ‘R’). Subjects perceive the left eye’s image or the right eye’s image in turn. *B. Top:* viewing an ambiguous stimulus continuously, observers experience random alternations between both percepts every few seconds. *Bottom:* periodically removing the stimulus from view (here: on-time ≈ 0.5 s; off-time ≈ 1.5 s) causes perception to stabilize in one interpretation for sometimes minutes, with only incidental switches between alternative interpretations: Perceptual stabilization.

et al., 2002). Second, when an ambiguous stimulus is preceded by a sequence of stimuli that are similar but contain unequal evidence for either interpretation (biased stimuli), influences of several preceding stimuli on current perception can be measured (Maloney et al., 2005).

We study how prior perception of an ambiguous stimulus influences how it is perceived at reappearance, specifically aiming at distinguishing persistence of the most recent percept from more intricate influences of past perception. We interleave episodes of intermittent viewing with episodes of continuous viewing (Figure 5.2A). During intermittent viewing perception stabilizes into one interpretation, whereas continuous viewing prompts spontaneous switches between percepts. The switches are essential, as they permit a dissociation between the most recent percept (following a switch) and preceding perception, enabling us to pit the effects of immediate and more remote perceptual history against each other.

5.2 Results

5.2.1 *The influence of a spontaneous perceptual switch on stabilization*

Two distinct ambiguous stimuli were used in our experiments (Figure 5.1). Results presented in the main text are for the ambiguous rotating sphere stimulus. Data from binocular rivalry are very similar (Appendix, Figures A5.1 and A5.3). The general layout of the sessions was the same for all experiments. It is illustrated in Figure 5.2A using experiment 1 as an example. Sessions consisted of blocks of intermittent presentation of an ambiguous stimulus, interleaved with periods of continuous presentation where perception was allowed to switch spontaneously. Subjects reported their perceptual state when a stimulus reappeared and whenever the percept switched. Transitions between intermittent presentation and continuous presentation were interactively initiated on the basis of observers' perceptual reports (Figure 5.2A, bottom). Intermittent presentation sequences proceeded until an observer reported the same percept on eight consecutive presentations, signaling robust stabilization. When this occurred, the stabilized percept was termed the *winner* of that intermittent presentation sequence, and continuous presentation was started. Continuous presentation periods, in turn, were terminated a fixed period after an observer reported a perceptual switch, and then intermittent presentation started again.

In the first experiment we allowed a single switch during continuous viewing. The delay (t_2' in 5.2A) between this switch and stimulus offset was drawn randomly from the values 0.5, 1.5 and 3 s. This mimicked the aforementioned situation in (Leopold et al., 2002) where the effect of an incidental switch on stabilization was shown to depend on this delay. Perceptual switches naturally occur at random intervals, so the percept preceding the switch (t_1' in Figure 5.2A) also had a variable duration. This allowed us to study how perception at stimulus reappearance depended on a well-controlled fraction of perceptual history; that is, on the duration of both perceptual episodes t_1 and t_2 .

As a main measure we will use the probability that the winner of a given intermittent presentation sequence equals the winner of the preceding intermittent presentation sequence. If it does not, the perceptual switch successfully disrupted perceptual stabilization; if it does, perception during this intermittent presentation sequence reverted to the winner of the previous intermittent presentation sequence in spite of the intervening switch to the opposite percept.

Figure 5.2B plots the subject-averaged probability that the winner of the current intermittent presentation sequence equals the previous winner, as a function of the intervening percept

durations t_1 and t_2 ; that is, the durations preceding and following the perceptual switch, respectively. The durations are categorized as 'brief', 'moderate' and 'long', as indicated on the x-axis. In the case of t_2 these three categories correspond simply to the three delay durations applied, whereas for t_1 (whose duration varied in a continuous fashion) we divided the data up into three percentiles to form the three categories. The light curve shows that the current winner is less likely to equal the previous winner if the period following the switch (t_2) is longer. In other words, the longer the final percept of the continuous presentation episode dominated before stimulus offset, the more likely it was to remain stabilized during the subsequent intermittent presentation sequence, replicating (Leopold et al., 2002). A new finding here is that the duration of the percept preceding the switch (t_1 ; dark curve) has the opposite effect. That is, the longer the previous winner remained dominant before the perceptual switch occurred, the more likely it was to regain dominance during the subsequent intermittent presentation sequence, in spite of the intervening switch. Note that the durations t_1 and t_2 were not correlated in this design, as the delay between the perceptual switch and stimulus offset was varied independently of the spontaneous percept duration preceding the switch. Both curves therefore reflect different, orthogonal, subdivisions of the same data set.

In isolation, the influence of the final percept duration (t_2) would be consistent with an explanation that during a blank the visual system retains its latest perceptual organization, provided this organization has had sufficient time to establish. The finding of a comparable influence of the preceding percept duration (t_1), however, argues against such an exceptional role for the final percept. Instead, it suggests that ambiguous figure memory is determined by a more global perceptual history. Specifically, Figure 5.2B leads to the tentative interpretation that the longer a percept has dominated in the past – be it during the final dominance period or earlier – the more readily it will regain dominance when the stimulus reappears.

Apart from the effects of durations t_1 and t_2 , Figure 5.2B also shows a general tendency for the winner of the current intermittent presentation sequence to equal the previous winner; that is, all points lie above 0.5. Figure 5.2C addresses whether this could be due to a systematic tendency for subjects to report one particular percept during intermittent presentation, regardless of history (Carter and Cavanagh, 2007). For individual subjects this panel shows the probability that the current winner equals the previous one, both for the cases where percept 1 won the previous intermittent presentation sequence (dark shade) and those where percept 2 did (light shade). Both the light and the dark bars consistently reach beyond 0.5. This demonstrates that the current winner generally equaled the previous winner regardless of which percept it was, ruling out systematic bias as an explanation. As an extra measure we calculated each subject's individual bias, as the overall fraction of intermittent presentation sequences won by that subject's more predominant percept. This value was 0.57 on average, which is insufficient to explain the values in Figure 5.2B. Further analyses suggest that our comparatively long sessions (40 min) may have reduced the influence of a systematic bias in our experiments (Appendix, Figure A5.1).

Considering the above findings, one reason why the winner of the previous intermittent presentation sequence often recurs during the current intermittent presentation sequence could be that it has enjoyed much dominance in the past. That is, if it is true that prior dominance in general increases the probability a percept will recur after a blank, then the extensive dominance during the previous intermittent presentation sequence (Figure 5.2A) could underlie the previous winner's elevated probability of regaining dominance during the current intermittent presentation sequence.

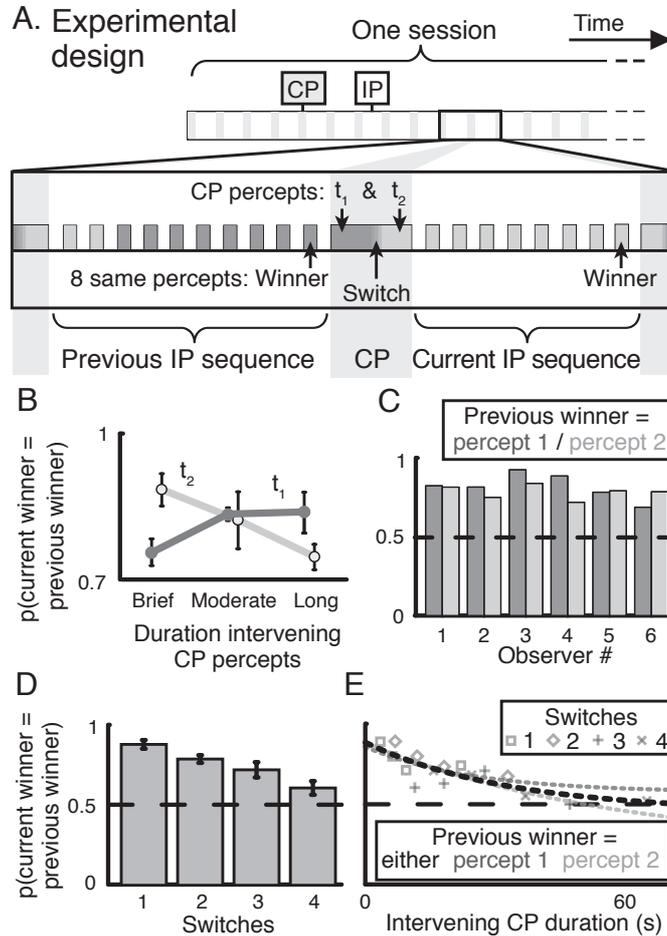


Figure 5.2: The effect of spontaneous perceptual switches on perceptual stabilization. *A. Experimental design.* Sequences of intermittent presentation (IP) were interleaved with periods of continuous presentation (CP). The first percept to reach eight consecutive dominance periods during intermittent viewing (termed the winner) triggered the start of continuous presentation. A continuous presentation period was terminated after a randomly varied delay (t_2) following the first perceptual switch. The duration (t_1) that preceded the switch varied naturally. *B.* The probability that the current intermittent presentation sequence is won by the same percept as the previous intermittent presentation sequence rises with t_1 but falls with t_2 (both $p < 0.01$; Spearman on individual subjects' data; $n=6$; $\rho=0.68$ for t_1 and -0.56 for t_2). The three data points for t_1 were based on the three $33\frac{1}{3}\%$ quantiles (average durations: 1.7, 4.1 and 11.6 s). *C.* The same probability plotted separately for cases where the previous winner was percept 1 (dark shade) or percept 2 (light shade). All bars reach above chance level, ruling out systematic bias as a key factor. *D.* When allowing multiple perceptual switches (x-axis) during continuous viewing, the probability that the current winner equals the previous winner decreases as the number of intervening switches increases ($p < 0.01$; Spearman on individual subjects' data; $n=7$, $\rho=-0.69$). *E.* For one representative subject (others shown in the Appendix, Figure A5.2) the probability decreases gradually with increasing duration of the intervening continuous presentation period (x-axis), reaching chance after about a minute. Again, the identity of the previous winner (light vs. dark thinner curves) does not play a large role. Error bars indicate standard errors.

5.2.2 *The influence of longer continuous presentation episodes on stabilization*

Would the winner of the previous intermittent presentation sequence still predominate during the current intermittent presentation sequence if a more extensive period of spontaneous switching intervened, thereby moving the previous intermittent presentation sequence further into the past? Our second experiment was similar in design to the first one, but continuous presentation episodes were allowed to include up to four perceptual switches instead of just one. The delay between the final switch and the end of continuous presentation was no longer varied but fixed at the same duration as one intermittent presentation (≈ 0.5 s).

Figure 5.2D shows the subject-averaged probability that the current intermittent presentation sequence was won by the same percept as the previous intermittent presentation sequence, as a function of the number of intervening switches. Indeed, the probability decreased with an increasing number of switches, with almost no memory of the previous winner left after four perceptual switches. Figure 5.2E (black curve) quantifies how this memory decay progressed over time, by depicting the same probability as a function of the duration of the intervening continuous presentation episode, for one representative subject. It reveals a gradual reduction with increasing duration, and the probability reaches chance level after a minute or so. This indicates that the influence of the previous winner fades during continuous viewing and is completely gone after about a minute.

In Figure 5.2E we again controlled for a potential role of a systematic preference for one of the two percepts, by reanalyzing the data separately for occasions where percept 1 won the previous intermittent presentation sequence (dark thin curve) and those where percept 2 did (light thin curve). The dark curve runs slightly above the light curve, indicating an overall tendency for this subject to perceive percept 1 more than percept 2. Nevertheless, the similarity between these curves confirms that our results do not depend on a systematic percept bias. Indeed, the systematic bias in this experiment was only 0.52 for this subject, and 0.56 on average.

5.2.3 *The influence of dominance during continuous presentation on stabilization*

The above findings support the idea that dominance of a given percept in the recent past facilitates its regained dominance at stimulus reappearance. Consistent with this idea, the preference toward the previous winner decreased as a longer period of alternating dominance separated the moment of stimulus reappearance from the winner's dominance streak during the previous intermittent presentation sequence. A more specific prediction from the hypothesis, however, is that the preference toward the previous winner should not decay passively during continuous viewing. Instead, the evolution of the preference during continuous viewing should depend on what is being perceived. Hence, perception during the current intermittent presentation sequence should depend on the previous winner as well as on perception during the intervening period of continuous presentation.

Figure 5.3A depicts additional analyses of the second experiment, which confirm this prediction and refine it. Like Figure 5.2E, Figure 5.3A shows the probability that the current intermittent presentation sequence is won by the winner of the previous intermittent presentation sequence, as a function of the intervening continuous presentation duration. To assess how perception during continuous viewing affects this probability, we now separated our data according to the fraction of the intervening continuous viewing period that was taken up by dominance of the percept opposite to the previous winner. This fraction could vary because spontaneous switches occur at random intervals during continuous viewing, and also because

the number of switches varied. During continuous presentation periods where the opposite percept dominated a large fraction of the time (black curve), the bias toward the previous winner decayed rapidly over time, and eventually even turned into an opposite bias. If the continuous presentation period contained little opposite dominance, in contrast (lightest shaded curve), the bias toward the previous winner remained strong even after prolonged continuous viewing. The curve for an intermediate fraction of opposite dominance (intermediate shade) falls in between these two extremes. These results confirm that the preference for the previous winner does not decay passively during continuous viewing. Instead, it decays rapidly during dominance of the opposite percept but stays high during dominance of the previous winner percept itself. This is consistent with the notion that perception at stimulus reappearance reflects a balance that continuously evolves while viewing an ambiguous stimulus, ever shifting toward the currently dominant percept.

The final analysis goes beyond a comparison of the winners of consecutive intermittent presentation sequences, to include perception throughout an intermittent presentation sequence (Figure 5.3B). We now investigated perception during the first eight presentations of the current intermittent presentation sequence, as well as during the winning presentation (which could be either the eighth or a later one, depending on perception). Figure 5.3C shows the probability that the previous winner dominated on these nine presentations, with the number of intervening perceptual switches during continuous viewing depicted in diagrams on the left. The plots confirm that the tendency toward the previous winner decays with an increasing number of intervening switches. More importantly, however, they show an influence of the final percept during continuous viewing on perception during the subsequent intermittent presentation sequence. This is visible in the conditions where the intervening continuous presentation period contained either one or three perceptual switches (left plots), and therefore ended in the percept opposite to the previous winner (see diagrams). In those conditions, subjects reported the previous winner less often during the initial presentations of an intermittent presentation sequence than during later presentations. The resulting trend is marked by curved arrows. In the conditions involving two or four switches, in contrast, where the last percept during continuous viewing was identical to the previous winner, no such trend is visible.

An interesting aspect of the influence of the final percept of continuous viewing (Figure 5.3C) is its short duration in comparison to the minute-scale trace left by the winner of the previous intermittent presentation sequence (Figure 5.2E). The influence of the final percept of a continuous presentation period dissipates within a few intermittent presentations, or about 10 seconds. In our design, the winning percept dominated on eight consecutive intermittent presentations, whereas the final percept of a continuous viewing period dominated only briefly. This suggests that the longer a percept has dominated, the more persistent a bias it leaves. Moreover, it appears that a transient bias toward the most recent percept and a persistent bias toward the previous winner can exist simultaneously and independently, in the sense that the bias toward the previous winner is not erased by the transient bias toward the most recent percept. Instead, as soon as the transient bias wears off during the course of an intermittent presentation sequence, the bias toward the previous winner turns out to be unaffected (Figure 5.3C). This rebound toward the previous winner after a temporary tendency toward the opposite percept is even more pronounced if one externally forces perception to the opposite interpretation by means of a disambiguated stimulus, instead of waiting for a spontaneous switch to occur during continuous presentation (Appendix, Figure A5.3). These data suggest the existence of multiple parallel biases, each reflecting a different timescale of perceptual history.

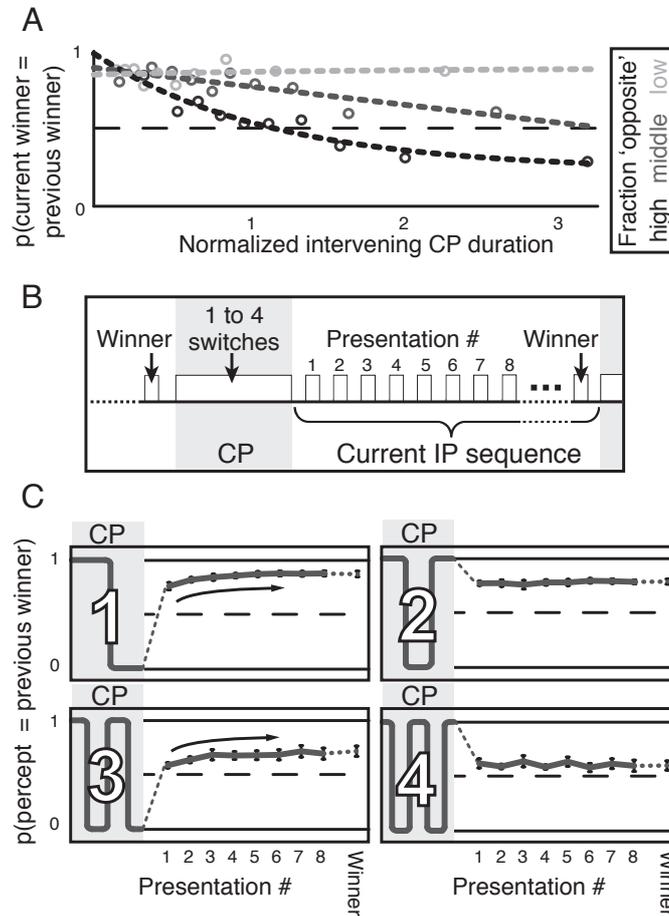


Figure 5.3: The influence of dominance throughout a continuous presentation episode on stabilization. A. The data of the second experiment were regrouped according to the fraction of a continuous viewing period during which the percept opposite to the previous winner dominated. The probability for the previous winner to also win the current intermittent presentation sequence decays faster if the opposite percept dominates much (black curve) than if dominates little (intermediate and light shade). For this panel the durations of continuous presentation periods were normalized per subject by dividing by their mean continuous presentation duration (17 s on average), and then pooled over subjects. Data were split into three quantiles of the fraction of opposite dominance. Normalized continuous presentation durations for each of the three groups were then divided into ten quantiles to yield ten data points. The black and dark gray data points show a negative trend ($p < 0.01$, Spearman; $\rho = -0.98$ and -0.68 , respectively), but the light gray ones do not ($p > .25$; $\rho = -0.07$). B. The data of the second experiment were reanalyzed to study perception during the first eight presentations of an intermittent presentation sequence, as well as during the winning presentation. C. Probability that the previous winner dominates on individual presentations of the current intermittent presentation sequence, for up to four intervening perceptual switches (digits and diagrams on the left). In those conditions where continuous viewing ended in the percept opposite to the previous winner (after one or three switches), the probability of perceiving the previous winner is lower during the initial intermittent presentations following continuous viewing than during later presentations (trend marked by arrows; Spearman $p < 0.01$; $\rho = 0.88$ for both one and three switches). Error bars indicate standard errors ($n = 7$).

5.2.4 A model account based on multi-timescale adaptation

We interpret our findings as follows. (i) During perceptual dominance the visual system accumulates a bias toward the currently dominant percept. (ii) This accumulation takes place on several timescales, such that prolonged dominance (e.g. during intermittent presentation) leaves a persistent biasing trace, whereas brief dominance (e.g. just before the end of a continuous presentation period) leaves a more transient trace. (iii) Separate timescales work independently, such that the system can briefly be biased toward one percept without losing its longer-term bias toward the other percept. (iv) These biases become evident in perception when an ambiguous stimulus reappears after an interruption (rather than during ongoing viewing). Then, the visual system's choice between both percepts reflects the balance between various biasing traces that have so far accumulated. It is worth emphasizing that this approach treats perceptual stabilization as a repeated choice for the same percept on many stimulus onsets; not as persistence of a single perceptual state (during stimulus absence the system is in neither perceptual state). What is to be explained, therefore, is how perceptual history can make the system choose one percept over the other at stimulus onset.

We have constructed a computational model (Appendix, Figure A5.4) that implements the above four concepts. The model, an extension of Noest et al. (2007), attributes perceptual stabilization to a history-driven bias in percept choice at stimulus onset. It is a natural property of the model that this bias takes effect specifically at stimulus appearance, and not during continuous viewing. An indication that this is an appropriate property is the experimental finding that factors that drive dominance at stimulus onset need not have a similar effect during prolonged viewing (Carter and Cavanagh, 2007).

In the model the bias gradually accumulates during perceptual dominance, due to gradual sensitivity changes, or adaptation (Ohzawa et al., 1985; Bonds, 1991), of neurons that code the currently dominant percept. The resulting imbalance in adaptation states persists for some time after the dominance period itself has ended, and therefore carries information on past perception – that is, acts as a memory store. At stimulus reappearance the difference in adaptation state between model neurons that code a recently dominant percept and those coding the other percept causes the recently dominant percept to win the competition. Interestingly, using only a single adaptation term with a single timescale of persistence, such as used by Noest et al. (2007), we were unable to account for our findings. In contrast, when incorporating two adaptation terms, one with fast decay and the other more persistent, our data could be replicated in considerable detail (Figure 5.4).

5.3 Discussion

Our results indicate that when a visual conflict reoccurs, traces of past perception increase the probability that the visual system will assume a previous state of interpretation rather than a different one. Despite its apparently simple nature, such a mechanism could have great merit for visual function. Visual input quite generally contains ambiguities, and in normal conditions only one perceptual interpretation is veridical. Selecting the one correct interpretation often requires the combination of multiple information sources (Ernst and Bühlhoff, 2004) and engages extensive regions of the brain (Leopold and Logothetis, 1999). The current observations suggest that by biasing the system toward previous perceptual interpretations – be it the most recent one or one that consistently dominated longer ago – visual memory automatically

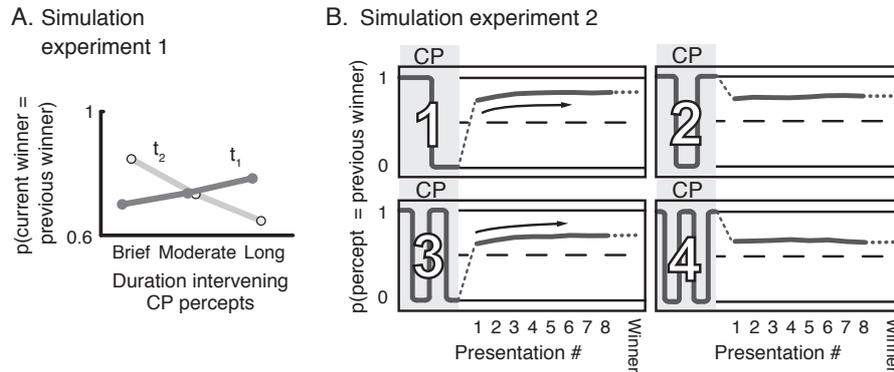


Figure 5.4: Model simulations. Our key findings are replicated by a model in which information on past perception is contained in the adaptation states of sensory neurons. A. Influence of one intervening perceptual switch on perceptual stabilization (cf. Figure 5.2B). The relation between percept durations t_1 and t_2 (see Figure 5.2A) and perceptual stabilization is replicated by the model. B. Influence of multiple intervening perceptual switches on perceptual stabilization (cf. Figure 5.3C). The probability that the previous winner is perceived during the current intermittent presentation sequence decreases with an increasing number of intervening switches, in line with the experimental data. In case the final percept during continuous viewing was opposite to the previous winner (one or three switches) the probability rises during the current intermittent presentation sequence (Spearman $p < 0.01$; $\rho = 0.95$ and 0.92 for one and three switches, respectively), consistent with our experimental findings.

enforces the outcome of previous perceptual conflicts, and thereby eliminates the need for the same conflicts to be resolved repeatedly.

Our model work shows that in neural terms this memory could be carried by the adaptation state of sensory neurons, provided adaptation occurs on more than one timescale. Indeed, adaptation is known to occur on a wide range of timescales in sensory cortex (Ohzawa et al., 1985; Bonds, 1991; Müller et al., 1999; Ulanovsky et al., 2004; Albrecht et al., 1984; Rose and Lowe, 1982), and it has been argued on theoretical grounds that this holds promise for functional forms of history-dependence in neural systems (Drew and Abbott, 2006; Gilboa et al., 2005). The present work thus suggests that the perceptual memory observed here constitutes one such functional correlate.

Our observations seem to conflict with an earlier study (Maloney et al., 2005) where percept choice at stimulus reappearance was interpreted as an attempted continuation of the preceding percept sequence. For instance, a sequence of percepts ABA would cause percept B at reappearance, forming the regular sequence ABAB. This is inconsistent with our data mainly because it involves suppressive effects of past dominance on subsequent percept choice (for instance when ABA causes B whereas BBA causes A, the initial percept stimulates opposite perception at the end of the sequence), whereas we find only facilitation. The discrepancy may be due to the fact that the stimuli used by Maloney et al. (2005) were not fully ambiguous but biased toward one interpretation. The use of ambiguous stimuli throughout our experiments has allowed us to characterize memory of the internally constructed interpretation of a stimulus in isolation. Effects of unbalanced stimuli are probably a combination of the present effects and differential gain control at stages prior to where the percept arises (Noest et al., 2007). This is known to cause different effects on subsequent perception (Pearson and Clifford, 2005; Brascamp et al., 2007; Kanai and Verstraten, 2005).

We modeled our findings by adding a longer adaptation timescale to an existing model of

perceptual stabilization (Noest et al., 2007). An alternative model of stabilization has also been proposed (Wilson, 2007). Both models are essentially standard oscillator models expanded with an additional interaction (Noest et al., 2007) or storage mechanism (Wilson, 2007) to allow a trace of previous perception to bias the next percept choice at stimulus reappearance. Both models in their original form have the limitation of lacking multiple timescales of storage. Regardless of the number of timescales, a drawback of Wilson (2007) is that it predicts stabilization of a percept that has dominated briefly before stimulus offset but no stabilization of a percept that has dominated longer (Figure 7 in Wilson (2007)) This is opposite to the experimental finding that brief dominance will prevent stabilization, and longer dominance is required for a percept to recur (Figure 5.2B in the present work, Figure 3C in Leopold et al. (2002)). Models of the type of Noest et al. (2007) do reproduce this feature. Arguably a second objection to Wilson (2007) is that it entails a binary memory, where the system is in one of two states of ‘remembering’ either percept. Experiments indicate that, instead, the system’s bias toward one or the other percept varies over time in a continuous fashion (Figures 5.2B, 5.2E and 5.3A in the present work, also Pastukhov and Braun (2007)), consistent with the model we used.

Our view of ambiguous figure memory suggests a relation to visual memory in other situations. Previous notions that attributed perceptual stabilization to prolongation of a perceptual state during stimulus absence seemed to imply that it is specifically an ambiguity-related phenomenon. The present view of perceptual stabilization as a bias in a decision network – in this case regarding a perceptual decision at stimulus onset – allows more room for extensions beyond ambiguous perception. Specifically, the accumulation of a bias during perceptual dominance that we observe here is reminiscent of the progressive decrease in response time that is observed when subjects direct their attention (Maljkovic and Nakayama, 2000; Kristjánsson, 2006) or eye fixation (McPeck et al., 1999; Dorris et al., 2000) to a similar search target appearing on several consecutive trials. This type of attention priming occurs automatically, independent of conscious recollection. It has been attributed to progressive use-related changes that build up in the neural structures activated when the target is attended (Kristjánsson, 2006; Dorris et al., 2000), so that every allocation of attention or gaze to an item simultaneously acts to stimulate reorientation to that item in the future. This is analogous to the accumulating bias that facilitates repeated perceptual dominance in our paradigm, a similarity that is particularly remarkable considering the numerous other parallels between attentional selection and perceptual dominance (Leopold and Logothetis, 1999; Mitchell et al., 2004).

Our findings bear directly on the question asked at the outset, how long a history to incorporate into current processing. Functionally, the answer depends on the liability for the conditions to change. If they change every few seconds it is useless to incorporate a minute-scale history because what happened a minute ago bears little relation to the present situation. If, in contrast, the conditions remain relatively stable for minutes, incorporating a longer-term history prevents unfavorable sensitivity to seconds-long (noisy) excursions. Our findings suggest how just such a strategy is implemented in vision, by use of parallel biasing traces on several timescales. In the case of ambiguity resolution, if recent perception was highly stable, slow biases have built up sufficiently to outweigh the fast bias due to the most recent percept. If perception was variable, however, no slow biases have accumulated and the most recent percept becomes the main driving factor. This organization therefore ensures automatic adjustment of the effective memory timescale, dependent on the changeability of the situation at hand.

5.4 Materials and Methods

5.4.1 Subjects and task

Subjects were two authors and sixteen naive observers. All had normal or corrected-to-normal acuity. After showing subjects the stimuli and explaining their task, we orally obtained an informed consent statement before proceeding with the experiment. All experiments were conducted in agreement with Utrecht University ethics and safety guidelines. Three subjects showed a strong preference for one of the percepts during pilot experiments, and were not included in further testing. The remaining subjects had an average preference, as measured by the fraction of all intermittent presentation sequences won by their preferred percept, of 0.57 ($\sigma=0.05$) and 0.61 ($\sigma=0.06$) for the sphere and rivalry, respectively. Subjects were instructed to fixate the center of the display passively, and report their percepts via key presses. Experimental sessions took 40 minutes.

5.4.2 Apparatus and stimuli

Ambiguous stimuli were an ambiguous rotating sphere ($r=0.65$ deg; $\omega=2.23$ rad/s; 90 black dots of $r=0.02$ deg; dot lifetime=1 s) and dichoptic ± 45 deg grayscale Gabor patches ($\sigma=0.37$ deg; 100% contrast; $\text{spf}=2.7$ c/deg). Stimuli were presented on a gray background (35 cd/m²) within a white alignment ring ($r=1.7$ deg) and with a red plus sign (side=0.2 deg) marking fixation for the sphere. They were presented via a mirror stereoscope, on a CRT monitor (1600x1200 dpi) at a visual distance of 47 cm.

5.4.3 Intermittent presentation sequences

The timing of intermittent presentation was optimized for each subject beforehand, to find a regime with robust perceptual stabilization. We therefore designed an adaptive procedure that dynamically adjusted stimulus timing according to a subject's perceptual reports, until no alternation was reported during 60 s of intermittent presentation. The average presentation duration was 0.5 s for both stimuli; the average blank duration was 1.4 s for binocular rivalry and 1.2 s for the ambiguous sphere. In all experiments we terminated and discarded an intermittent presentation sequence if a subject did not reach a stable percept within 24 intermittent presentations. This happened on 1.3% of the occasions.

5.4.4 Continuous viewing periods

The blank interval between the end of continuous viewing and the initial intermittent presentation was equal to the interval between consecutive presentations during intermittent viewing. In the first experiment the delay between the single perceptual switch and the end of continuous viewing was either 0.5, 1.5 or 3 s. A continuous viewing period was discarded if a second switch occurred during the delay, which for the three delay durations happened on 1, 13 and 29% of the occasions, respectively. In the second experiment the number of spontaneous switches varied randomly from 1 to 4 within sessions. The delay between the final switch and the end of continuous viewing was chosen equal to the duration of one presentation during intermittent viewing. This delay was chosen for experimental efficiency, because by design a presentation duration during intermittent viewing was short enough to minimize the occurrence

of additional switches before stimulus offset. The analysis of memory decay during continuous presentation (Figures 5.2E, 5.3A and A5.2) required more data than the other analyses, and was based on additional sessions with three naive observers of the sphere.

5.4.5 *Forced perceptual switches*

In the experiments where perception was exogenously forced away from the previous winner (Figure A5.3), unambiguous stimuli were constructed as follows. For the ambiguous sphere we added binocular disparity to the dots, defining a unique rotation direction. Brief exposure to such an unambiguous rotation direction tends to cause perception of that same rotation direction during subsequent ambiguous viewing (Jiang et al., 1998). For binocular rivalry the unambiguous stimulus consisted of one of the eyes' images in isolation, which caused dominance of the opposite eye's image during subsequent ambiguous viewing (flash suppression (Wolfe, 1984)). An effective duration of unambiguous presentation was determined per subject in pilot sessions beforehand, and amounted to 0.9 s on average for the sphere, and 0.8 s on average for binocular rivalry.

Appendix

Perceptual stabilization and bias in binocular rivalry

Figure A5.1A shows the probability for binocular rivalry of perceiving the winner of the previous intermittent presentation sequence during the current intermittent presentation sequence, for up to four intervening perceptual switches during continuous presentation. The results are similar to those for the ambiguous sphere. The probability lies above chance level throughout all four plots, and the conditions where the last percept preceding the current intermittent presentation sequence was opposite to the previous winner (one and three switches) reveal an influence of this final percept. In those conditions, the probability of perceiving the previous winner was lower during the initial intermittent presentations than during the later ones (rise indicated by arrows). One qualitative difference with the ambiguous sphere data is that for binocular rivalry the influence of the final percept of the continuous viewing episode seems stronger. Part of its effect extends to the winner of an intermittent presentation sequence. As a consequence, the probability for the current winner to equal the previous winner does not decay monotonically with the number of intervening switches (as it did for the ambiguous sphere), but lies slightly higher after even switch numbers than after odd ones.

Like those of the ambiguous sphere, our observers of binocular rivalry displayed only a modest systematic perceptual bias (the fraction of intermittent presentation sequences won by the preferred percept was 0.61 on average). This seems to conflict with the finding that systematic bias can profoundly influence binocular rivalry perception during intermittent presentation (Carter and Cavanagh, 2007). We tested if our comparatively long session duration (40 min, versus 2-10 min in Carter and Cavanagh (2007)) could explain this discrepancy. Figure A5.1B shows the time course of one typical session, plotting the winners of consecutive intermittent presentation sequences against time. Both percepts have periods (runs) during which they won several consecutive intermittent presentation sequences (the circled numbers indicate the number of consecutive intermittent presentation sequences). Of note, the typical duration of such a run is several minutes. Now suppose that one effect of a systematic bias is that the

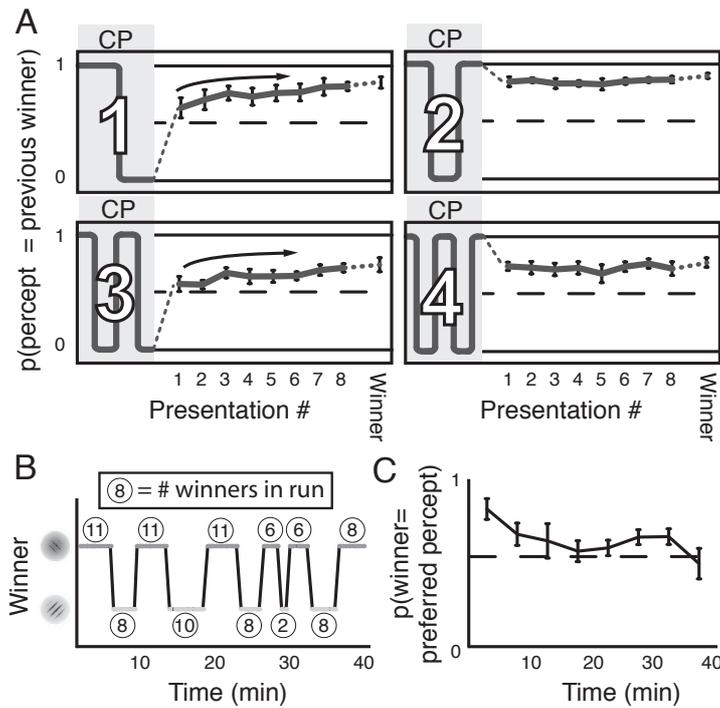


Figure A5.1: Perceptual stabilization for the binocular rivalry stimulus, in case up to four perceptual switches intervened during continuous viewing. **A.** Subject-averaged probability ($n=4$) of perceiving the winner of the previous intermittent presentation sequence during the current intermittent presentation sequence. Compare Figure 5.3C. The dominant percept during the current intermittent presentation sequence was usually identical to the winner of the previous intermittent presentation sequence. In conditions where the last percept of an intervening period of continuous viewing was opposite to the previous winner (one and three switches) the probability of perceiving the previous winner was lower during the initial presentations than during later ones (trend marked by arrows; Spearman $p < 0.01$; $\rho = 0.98$ and 0.88 for one and three switches, respectively). **B.** Winners of consecutive intermittent presentation sequences during a typical session. Runs during which the same percept won all consecutive sequences typically took several minutes. The circled numbers indicate the number of intermittent presentation sequences in such a run. **C.** Fraction of intermittent presentation sequences that were won by subjects' preferred percepts, for consecutive five-minute blocks during a forty-minute session. The fraction is higher during the first five minutes of a session than during the rest of the session (one-sided t -test; $p < 0.05$). We calculated the fraction for each subject individually before averaging, for each subject taking the percept that overall won most intermittent presentation sequences as the preferred percept.

first percept to get stabilized within a session is often the same for a given subject. Because runs of repeated dominance typically take several minutes, such a scenario would cause a profound bias in case the session duration is of the order of minutes, whereas longer sessions would reduce the influence of bias. In Figure A5.1C we investigate this idea by calculating the bias of our binocular rivalry observers for consecutive five-minute blocks within sessions. Indeed, the bias during the first block is considerably larger than the overall bias, in support of the idea that our use of long sessions may have reduced the influence of a consistent bias.

Decay of memory during continuous viewing

For one individual subject, Figure 5.2E showed the probability that the winner of the current intermittent presentation sequence equals the winner of the previous intermittent sequence, as a function of intervening continuous presentation (CP) duration. Figure A5.2 shows the same for the two other subjects who performed this experiment. The data for all three subjects are similar. For the two subjects in Figure A5.2 the probability that two consecutive winners are equal reaches chance after about 25 s and 40 s, respectively.

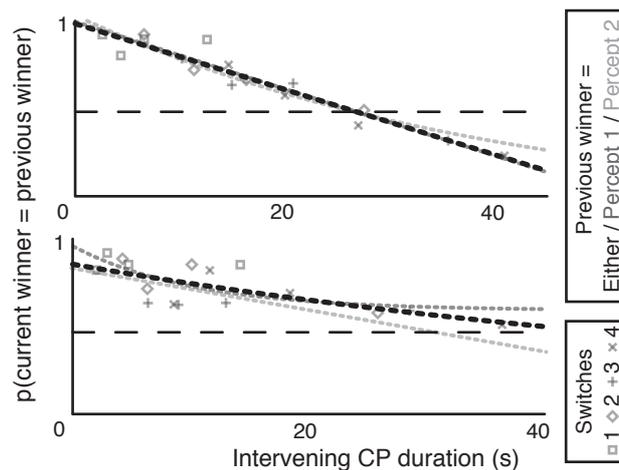


Figure A5.2: Individual subjects' probability that the winner of the current intermittent presentation sequence equals the winner of the previous intermittent sequence, as a function of intervening continuous presentation (CP) duration (cf. Figure 5.2E). The data for are similar for all three subjects. In light of the data from all three subjects together, the progression through chance in the top graph should probably be attributed to a few outliers; not to a systematic effect.

The influence of a forced perceptual switch

We performed an additional experiment where we externally forced perception away from the previous winner, instead of waiting for spontaneous switches to occur during continuous presentation. We interleaved intermittent presentation sequences, not with continuous viewing episodes, but with presentation of an unambiguous stimulus (UP; Figure A5.3A) that

imposed the percept opposite to the previous winner. For binocular rivalry we used flash suppression (Wolfe, 1984); for the ambiguous sphere we used positive priming with a disparity-disambiguated sphere (Jiang et al., 1998). Figure A5.3B depicts how this affected perception during intermittent presentation. During the initial presentation following unambiguous stimulation the probability of perceiving the previous winner was relatively low, indicating successful disruption of the winner’s dominance streak. However, during the following intermittent presentations the probability of perceiving the previous winner gradually recovered. This pattern of results is similar to that observed when spontaneous switches ended in the percept opposite to the previous winner (Figures 5.3C and A5.1A), and demonstrates that the influence of the previous winner can survive both spontaneous and forced perceptual switches.

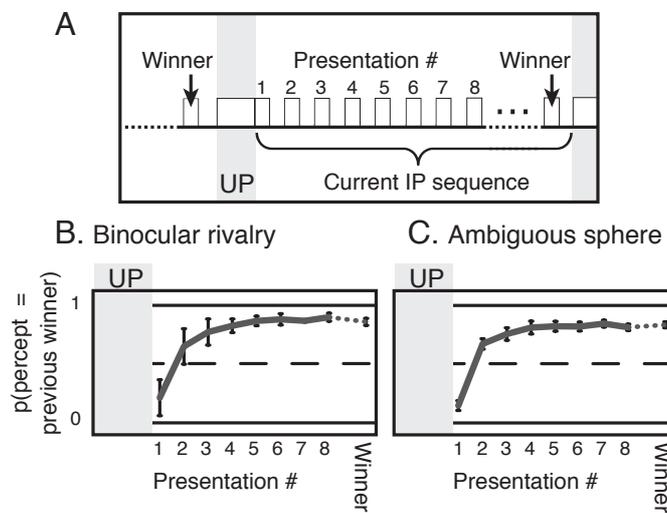


Figure A5.3: The effect of a forced perceptual switch on stabilization. A. We interleaved intermittent presentation (IP) sequences with brief presentation of a disambiguated version of the stimulus (UP). This imposed perception of the percept opposite to the previous winner. B. Results for binocular rivalry (left; $n=3$) and the ambiguous sphere (right; $n=6$). During the initial presentations following the unambiguous stimulus there was a decreased probability of perceiving the previous winner. The probability recovered during the course of an intermittent presentation sequence.

A model based on multi-timescale adaptation

Recurrence of the most recent percept after a blank has previously been accounted for using a modified version of the standard bistable oscillator model (Noest et al., 2007). In such models, the two perceptual interpretations correspond to neural populations that interact via mutual inhibition and that exhibit slow self-adaptation. In the model by Noest et al. (2007) past perception influences percept choice at stimulus reappearance via the adaptation component, which persists for some time after the model input (stimulus) has been switched off, and after the populations’ responses themselves have as a result fallen to near baseline.

Each population, indexed i and j , is modeled by two differential equations; one for its neural response H and one for its adaptation state A . We show the equations for i ; those for j are the same with indices i and j exchanged. The time derivative of response H_i is given by

$$\partial t H_i = X_i - (1 + A_i)H_i + \beta A_i - \gamma S[H_j] \quad (\text{A5.1})$$

X_i is i 's excitatory input, γ is a constant that determines the strength of cross inhibition and $S[H_j]$ is a sigmoid function of H_j (H itself can be viewed as a population-averaged membrane potential, and $S[H]$ as a population-averaged firing rate). The term $-H_i(1 + A_i)$, here called the 'shunting' term, implements a standard diminishing effect of adaptation A on activity H , causing switches during continuous viewing. The term $+\beta A_i$, here called the 'additive' term, is a key component in Noest et al. (2007) that mediates a facilitatory effect of adaptation A at stimulus onset (β is a constant). Adaptation is implemented as 'leaky integration' over activity $S[H_i]$:

$$\tau \partial t A_i = -A_i + \alpha S[H_i] \quad (\text{A5.2})$$

Here $\tau > 1$ is a time constant and α is a constant.

When both activities H_i and H_j start rising from near baseline after a blank period (during which $X_i = X_j = 0$) their race is biased by adaptations A_i and A_j . Here the 'shunting' term acts against the more adapted percept and the 'additive' term acts in its favor. Figure A5.4A illustrates how the outcome of the competition between these forces depends on the specific combination of A_i and A_j at stimulus onset. The gray and white regions indicate combinations of A_i and A_j that lead to onset dominance of percept i and j , respectively. Crucially, the model predicts a region, below the dotted line, where the 'additive' component prevails and the more adapted percept gains dominance. In earlier models, that did not incorporate the 'additive' term, this region was absent. The presence of this region allows the model to produce perceptual stabilization, where the more adapted percept keeps on regaining dominance on consecutive presentations.

With a single time scale of adaptation (equation A5.2) this model can thus explain recurrence after a blank of the most recent percept. The present findings, however, require the inclusion of adaptation on multiple timescales in the model. The most compelling reason for this is the fact that the winner of the previous intermittent presentation sequence recurs during the following intermittent presentation sequence regardless of what the final intervening percept was. This cannot be explained by the model in its present form. In the following we will explain the reason for this limitation, and we will show how it is remedied by multiple adaptation timescales.

The course of adaptations A_i and A_j during perceptual alternations in the existing model is plotted in Figure A5.4B, both as a function of time (left) and in the (A_i, A_j) -space (right) introduced in panel A. Importantly, in the right plot the course of adaptation during dominance of i (from I to II) is the same as its course during dominance of j (from II back to I), but mirrored in the main diagonal. Note that the distribution of gray in Figure A5.4A is also identical to the distribution of white mirrored in the main diagonal. These two facts combined imply the following: whenever a blank period introduced during dominance of i has characteristics that lead to recurrence of i , the model in its present form predicts that an identical blank period introduced during dominance of j will lead to recurrence of j . This is illustrated in panels C and D. Panel C shows how a blank period introduced during dominance of j can lead to recurrence of percept j . Both the delay period between the switch to j and stimulus offset (II to III), and the blank duration (III to IV) are sufficiently long for the adaptation states to end up in a white region at stimulus reappearance (IV), causing j to regain dominance. Panel D illustrates

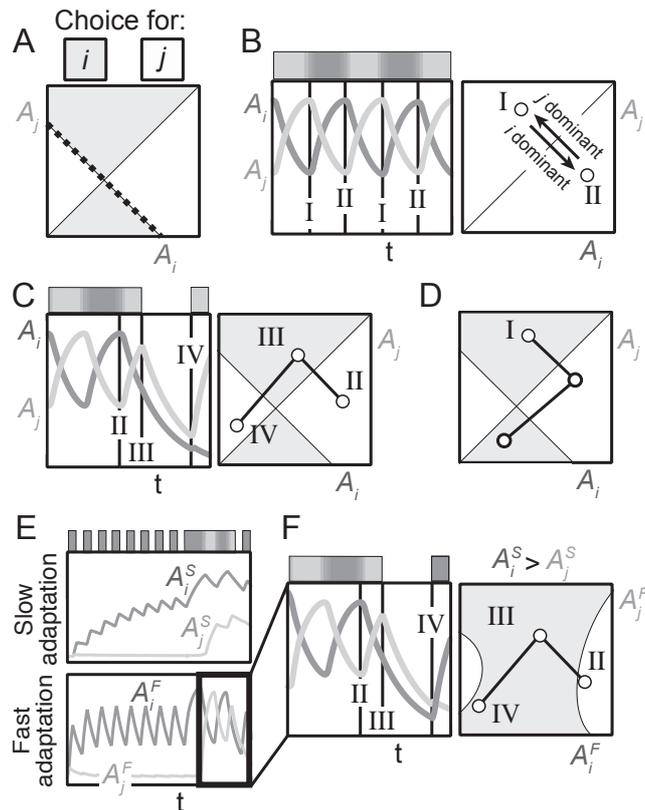


Figure A5.4: Model behavior. A. In the model by Noest et al. percept choice at stimulus onset is a function of adaptations A_i and A_j . B. The course of adaptations A_i and A_j during perceptual oscillations. C. Removing the stimulus (III) following a perceptual switch (II) to percept j can cause j to recur at stimulus reappearance (IV). D. Using identical timing, removing the stimulus after a switch to i (I) causes i to recur after the blank. E. We add a second, slower, timescale of adaptation to the model. The slower adaptations A_i^S and A_j^S remain unbalanced during continuous viewing, even after fast adaptations A_i^F and A_j^F have resumed their steady oscillation. F. Due to the remaining asymmetry in slow adaptations A_i^S and A_j^S , the dependence of percept choice on fast adaptations A_i^F and A_j^F has changed in such a way that the previous winner i regains dominance for a broader range of A_i^F , A_j^F combinations.

the effect of a blank period with identical timing, but introduced during dominance of i . The resulting course of adaptation in the (A_i, A_j) -space is the mirror image of that in panel C, and because the gray and white regions are also each other's mirror image, percept i recurs.

In order to explain our present finding that the previous winner tends to recur at stimulus reappearance, regardless of whether the stimulus was removed during dominance of i or during dominance of j , we must break the symmetry of either the adaptation trajectories, or the underlying gray-and-white landscape. The latter is what the introduction of a second adaptation time scale does.

We expand the model with a second, slower, adaptation term:

$$\partial t H_i = X_i - (1 + A_i^S + A_i^F) H_i + \beta_S A_i^S + \beta_F A_i^F - \gamma S[H_j] \quad (\text{A5.3})$$

$$\tau_F \partial t A_i^F = -A_i^F + \alpha_F S[H_i] \quad (\text{A5.4})$$

$$\tau_S \partial t A_i^S = -A_i^S + \alpha_S S[H_i] \quad (\text{A5.5})$$

These equations are identical to equations A5.1 and A5.2, except that where previously there was only one A there now are two. Subscripts and superscripts S and F indicate slow and fast adaptation, respectively. Consequently, $\tau_S > \tau_F$.

Figure A5.4E shows the course of both fast and slow adaptations during an intermittent presentation sequence and the following continuous presentation period. Slow adaptation (top graph) gradually accumulates during intermittent presentation, and stays asymmetrical throughout continuous viewing. In other words, it keeps on carrying a trace of the most recent winner, i , throughout continuous viewing, even as perception itself oscillates between i and j . Fast adaptation (bottom graph) is also asymmetrical during intermittent presentation, but A_i^F and A_j^F quickly resume their symmetrical oscillation cycles during continuous viewing. Fast adaptation during and right after the continuous presentation period is magnified in panel F. The stimulus is removed during dominance of j , and the course of fast adaptation is very similar to that in panel C. However, the outcome of the race between both percepts at stimulus reappearance (IV) is now determined by a combination of both fast and slow adaptation. Given the remaining asymmetry between slow adaptations A_i^S and A_j^S , the gray-and-white landscape as a function of A_i^F and A_j^F is no longer symmetrical. The gray region, which leads to dominance of the previous winner i , has expanded. As a consequence, i regains dominance after the blank, even though the final percept before the blank was j .

The settings during the simulations of Figure 5.4 were $\gamma=3.3$, $\beta_F=0.28$, $\beta_S=0.45$, $\alpha_F=4.0$, $\alpha_S=0.4$, $\tau_F=90$, $\tau_S=800$, $X=1.0$, all equal for i and j . For the sigmoid nonlinearity $S[H]$ we used $S[H] = \frac{H^2}{1+H^2}$ for $H > 0$, and $S[H] = 0$ otherwise. The presentation time and blank time during intermittent presentation were 50 and 90, respectively (arbitrary units). In the simulation of Figure 5.4A the delay between a switch and stimulus offset was varied between 50, 75 and 100; during the simulation of Figure 5.4B it was always 50. For our simulations we used a fourth order Runge-Kutta method with step size 0.1, applying Poisson noise to $S[H]$ on every time step. This was implemented as

$$S[H]_{new} = \frac{\text{Poisson}(S[H]_{old}C)}{C} \quad (\text{A5.6})$$

$S[H]_{old}$ and $S[H]_{new}$ are the values before and after the addition of noise, respectively. $Poisson(\mu)$ is a random value from the Poisson distribution with mean μ , and C is a constant that we set to 120. The data in Figure 5.4 were generated using independent noise on $S[H]$ in all six equations. We obtained equivalent results when applying identical noise to all three occurrences of $S[H_i]$, and to all three occurrences of $S[H_j]$.

Chapter 6

Intermittent ambiguous stimuli cause periodic perceptual alternations

This chapter in a sense continues where the previous one left off. After having investigated why perception usually does not alternate during intermittent viewing (Chapter 5), here we explore why eventually it does. I think this and the next chapter are the only ones that are really model-driven. In the previous chapters where we combined psychophysics and modeling, we invariably started off with an empirical observation, and afterward realized that we could model it. In this case the question of why perception alternates was really prompted by model analysis indicating that something interesting might be found. I enjoy the fact that by the end of this chapter we essentially have a comprehensive explanation of the entire perceptual sequence throughout a session of intermittent viewing: we have an idea of why it is that dominance repeats all the time, and we have an idea of why it alternates when finally it does. It is probably time for me to move on to a new subject.

Abstract

When viewing a stimulus that has multiple plausible real-world interpretations, perception alternates between these interpretations every few seconds. The alternations can be halted by intermittently replacing the stimulus with a blank screen. Then, the same interpretation commonly dominates on many presentations, thus stabilizing perception. Here we explore the limits of this stabilization. We show that intermittent presentation does give rise to perceptual alternations if it is continued for a sufficiently long time. Alternations during intermittent viewing typically occur only once every few minutes, and are thus much less frequent than those during continuous viewing. We show that the alternations are highly periodic, and that their frequency increases with decreasing blank duration. We also demonstrate that these alternations can be dissociated from those during continuous viewing in several ways, consistent with a distinct neural basis. We provide an interpretation that centers on a balance between priming (stabilizing) and suppressing (destabilizing) factors, which gradually shifts as traces of past perception accumulate over many minutes of stabilized perception.

6.1 Introduction

Hermann von Helmholtz argued that perception is based on a process of unconscious inference; the unconscious formation of hypotheses as to what external objects may have given rise to an observed pattern of nervous stimulation (Von Helmholtz, 1867). Few paradigms illustrate this point as convincingly as that of ambiguous perception. Here an observer is presented with a particular visual pattern that has two (or more) plausible external-world interpretations (Figure 6.1). As a result perception alternates back and forth between these alternative interpretations, as if unable to settle on one definitive hypothesis to account for the observed pattern of nervous stimulation (Leopold and Logothetis, 1999; Blake and Logothetis, 2002). Much ambiguous perception research has focused on these alternations between perceptual interpretations that occur over the course of viewing (e.g. Levelt, 1967; Fox and Herrmann, 1967; Borsellino et al., 1972).

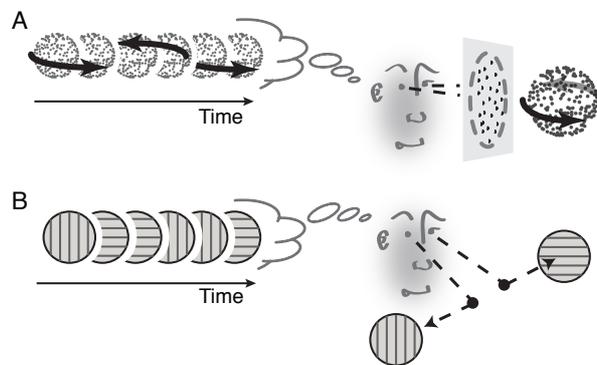


Figure 6.1: Schematic illustration of the ambiguous stimuli used in this study. *A. Structure from motion rivalry. Viewing a two-dimensional projection of a transparent structure rotating in depth, can cause perceptual alternations between two opposite rotation directions. B. Binocular rivalry. When the left and right eye view incompatible images in the same region of space, perception typically alternates between exclusive visibility of either of the two images in isolation.*

Given these frequent and unavoidable perceptual alternations during continuous presentation, it is surprising that perception of an ambiguous stimulus can be stabilized to one interpretation by the simple act of periodically removing the stimulus from view. That is, an ambiguous stimulus that is presented for brief periods, separated by seconds-long blank intervals, is generally perceived in the same configuration for many presentations in a row (Pearson and Brascamp, 2008; Orbach et al., 1963; Leopold et al., 2002). This perceptual stabilization has engendered much research in recent years, and there is now a plausible account of its origin (Noest et al., 2007; Brascamp et al., 2008). This account can essentially be broken down into two parts. First, individual presentations in this paradigm are too brief to allow perceptual alternations to occur in the manner of continuous viewing. Hence, perception generally does not change during individual presentations. Second, perception of a given interpretation leaves a facilitating trace that stimulates repeated perception of that interpretation on subsequent presentations. Hence, perception generally does not change from one presentation to the next, either.

Because of the stark contrast with continuous viewing, the first goal of research in this area

has been to understand the relative absence of alternations during intermittent viewing. Now that theory successfully explains why alternations may *not* occur in this paradigm, our attention was drawn to the few perceptual alternations that *do* remain. What are the characteristics of these alternations, and what causes them?

Here we show that perceptual alternations during intermittent viewing occur at highly regular intervals. Individual alternations can be separated by as much as ten minutes (hundreds of presentations), a duration that decreases as the blank interval that separates presentations is shortened. The presence of periodicity rules out random fluctuations as the cause of alternations during intermittent viewing, and points to a deterministic origin. The dependence of alternation frequency on blank duration argues against an autonomous clock-like oscillator, and indicates a stimulus-driven alternation mechanism.

The perceptual oscillations during intermittent viewing differ from those during continuous viewing by their timescale (tens of minutes versus seconds). The two are furthermore dissociated by a dissimilarity in the distribution of time intervals between consecutive alternations, and by a lack of correlation in individual subjects' alternation frequencies between both paradigms. Indeed, our observations are consistent with the aforementioned account of perception during intermittent viewing, which relies on mechanisms that are different from those active during continuous viewing (Noest et al., 2007; Brascamp et al., 2008). We argue that perceptual alternations during intermittent viewing occur when the facilitating trace of past perception, which initially causes stabilization, gives way to a suppressive trace, prompting opposite perception on the next presentation.

6.2 Methods

6.2.1 Stimuli and procedure

We used both binocular rivalry and structure from motion ambiguity (Figure 6.1). For binocular rivalry observers viewed square wave gratings (0.67 Michelson contrast; 3.3 periods per deg), viewed within a square aperture (side 1.4 deg). Stimuli were surrounded by a white fixation square (62 cd/m², side 1.5 deg). Background luminance was 31 cd/m², as was the mean stimulus luminance. Our structure from motion stimulus was an orthographic projection of a sphere rotating around a vertical axis (radius 1.2 deg, rotation speed 0.17 revolutions/s). It consisted of 250 white dots on a black background, presented with a red plus sign at fixation.

Sessions of intermittent presentation took 30 minutes; sessions of continuous presentation took 3 minutes. During intermittent viewing subjects were instructed to report their percept via key presses on every stimulus appearance. In case of mixed percepts they were forced to choose the most salient one. During continuous viewing subjects pressed and held the key corresponding to the dominant percept, and released all keys during mixed percepts.

Presentation durations during intermittent viewing were always 0.6 s for both stimuli. Blank durations were varied per session. For intermittent viewing of binocular rivalry, all observers started with an experiment using the default blank duration of 1.6 s. For the sphere this was 0.8 s. When we subsequently varied blank duration (Figure 6.3), we chose values around the default value for each subject individually. This choice depended on the subject's personal alternation rate at the default blank duration, relying on our preliminary observation that a reduction in blank duration could increase the alternation rate (e.g. if an observer of the sphere would perceive only two alternations per half hour at 0.8 s, we would tune the blank duration

downward, whereas if they perceived many alternations we would tune it upward). All blank durations we applied lay between 450 ms and 2.8 s. The comparison to continuous presentation of Figure 6.4C was made using the default duration of 1.6 s.

A total of twelve subjects (ten naive) participated in experiments with binocular rivalry. Seven subjects (five naive) participated in those with the ambiguous sphere.

6.2.2 Data analysis

The expected values and error margins on the power curves of Figures 6.2B and 6.3B were created on the basis of a kind of Monte Carlo analysis. For a given time series of perceptual reports we calculated the average probability of a perceptual alternation from any one presentation to the next. Taking this as a fixed probability we generated 500 artificial time series that were subsequently used to calculate the expected value and error margin. In Figure 6.2C each expected peak power value is the average of 500 simulated maxima obtained in the same way. p values were calculated from cumulative Gaussian distributions fitted to such sets of 500 simulated maxima. All power spectra were normalized by dividing by the surface area under the entire curve.

The oscillation frequencies in Figure 6.3C were normalized per subject by dividing by the frequency at the shortest blank duration. This compensated for the broad inter-subject variability in oscillation frequency that we observed. Before normalizing, all frequencies lay between 1 and 7 periods per half hour.

Similarly, for Figures 6.4A and B percept durations were normalized per subject by dividing by the mean, before pooling data for the histograms.

6.3 Results

Our first experiment was inspired by the observation that perceptual stabilization, although strikingly robust, does not extend indefinitely. After, say, a minute of stabilized perception, a period of repeated dominance of one interpretation may give way to a period of repeated dominance of the alternative interpretation. This is an interesting observation, because perceptual stabilization involves a facilitatory effect of past perceptual dominance on subsequent perception, and this facilitatory trace is known to grow stronger over presentations (Brascamp et al., 2008). That is, the more a given percept has dominated in the past, the higher the probability that it will gain dominance again on a subsequent presentation. Given this ever growing facilitation during stabilized perception, we wondered why perception would alternate at all in this paradigm.

Possible answers to this question can broadly be categorized into two groups. In one class of scenarios, alternations are caused by random fluctuations in brain activity (noise), which cause the previously unperceived interpretation to gain perceptual dominance in spite of a counteracting facilitatory trace. In that case alternations could be argued to be an inessential part of this paradigm. Alternatively, it is possible that the alternations have a deterministic origin. In that case they reflect an intrinsic aspect of the perceptual process, and their study may provide clues regarding its workings.

One way to make this first distinction is to look for periodicity in the occurrence of perceptual alternations during intermittent viewing. If alternations are caused by noise, they should

occur at random intervals and display no periodicity. If alternations are part of the deterministic behavior, in contrast, they may occur at regular intervals.

6.3.1 Alternations during intermittent viewing occur at long, regular intervals

In our first experiment subjects reported perception of an ambiguous stimulus blinking on and off intermittently. Figure 6.2 shows the main findings. Panel A shows data from one typical session, during which a binocular rivalry stimulus (Figure 6.1B) was shown for 0.6 s presentations, separated by 2 s blanks. Perception as reported by the subject is indicated on the y-axis, against time on the x-axis. Note that our subjects did not score cases of mixed dominance of both percepts, and that intermediate percept values apparent at some points in this figure are due to data smoothing (see figure caption).

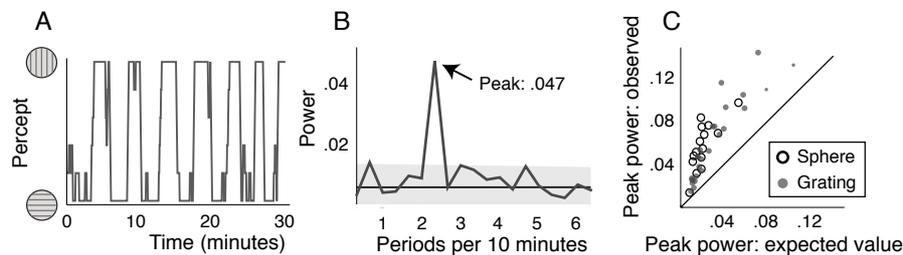


Figure 6.2: Alternations during stabilized intermittent viewing are infrequent but periodic. *A.* Perception of a binocular rivalry stimulus during a half-hour session (box car filtered with window size 5 to enhance clarity). There is a distinct impression of periodicity in the alternations. *B.* Section of the associated power spectrum. The spectrum is significantly peaked, indicating periodicity. *C.* Summary of all data. Each point corresponds to one subject and condition. In every case the power spectrum had a higher peak value (y-axis) than expected assuming random alternations (x-axis). This confirms that alternations during stabilized intermittent viewing occur periodically.

Eyeballing the time trace there is a striking impression that alternations tend to occur at regular intervals, separated by minutes-long stretches of perceptual stabilization. Such regularity can be demonstrated more formally by inspecting the power spectrum of this time series. Panel B shows part of the power spectrum of the data shown in panel A. The distinct peak indicates that alternations did tend to occur at regular intervals in this session. In this case the peak is located at a frequency of just over 2 periods per 10 minutes, or about 14 alternations per half-hour session. The black line and shaded area indicate the power that is expected in the absence of periodicity, plus or minus two standard deviations (see Methods). The peak lies significantly above the expected value ($p < .01$).

Panel C summarizes data from all sessions, using either the ambiguous sphere (open shapes; 6 subjects) or binocular rivalry (solid shapes; 10 subjects), and using off-durations ranging from 0.45 s to 2.8 s. The figure displays the observed peak value in the power spectrum as a function of the peak value that is expected in the absence of periodicity. The line $y=x$ thus indicates a lack of periodicity. All points lie above their expected value ($p < .01$ for each point apart from the two small ones on the top right), indicating that alternations tended to occur at regular intervals throughout all our sessions. The widely spaced alternations that cause this low-frequency periodicity can be distinguished from incidental brief excursions back and forth, which can also occur.

6.3.2 Oscillation frequency depends on off-duration

The periodicity observed in experiment one demonstrates that alternations during intermittent viewing are not due to random fluctuations. Instead, they must have a deterministic origin.

One may wonder whether these periodic alternations of the stabilized percept are affected by the specific choice of visual stimulation regime, or whether they might involve some autonomous internal oscillator that operates independently of visual input (e.g. Carter and Pettigrew, 2003). We therefore investigated whether the slow oscillation of the stabilized percept could be influenced by changing the rate of intermittent presentation. Each subject ran three sessions, each with a different blank duration between consecutive presentations. It is known that percept alternations during intermittent viewing become more numerous when the blank duration in between presentations is reduced (e.g. Orbach et al., 1966). Does the frequency of the minute-scale oscillations we observe here also depend on this duration?

Figure 6.3A shows perceptual reports for one observer of the ambiguous sphere. We display only data of his two extreme-most blank durations, namely 0.8 s and 2.6 s (his data for 1.6 s were intermediate). At the shorter blank duration (left) there are more perceptual alternations than at the longer blank duration (right). Moreover, irrespective of the number of incidental alternations back and forth, the frequency of the minute-scale oscillation is higher at the short blank duration than at the long one. Panel B shows this more explicitly, displaying sections of the power spectra associated with the time courses in panel A. At the short blank duration the spectrum peaks at a frequency of just over 1 period per 10 minutes (with a secondary peak at a slightly higher frequency). At the longer blank duration it peaks at a lower frequency, near half a period per 10 minutes. We will refer to the location of the peak in the power spectrum as the ‘oscillation frequency’.

Panel C summarizes how this oscillation frequency varied with blank duration for all observers. For both stimuli the oscillation frequency decreases as the blank interval becomes longer, confirming the example of panels A and B ($p < .01$ for both stimuli, one-sided t-test comparing all subject’s slopes of oscillation frequency vs. blank duration with 0).

6.3.3 Comparison to continuous presentation

The perceptual oscillations we report bear some similarity to the oscillations observed during continuous viewing. One may wonder whether the parallel between both types of oscillations goes beyond this outward similarity. In a final experiment our subjects tracked perceptual alternations during continuous presentation of a binocular rivalry stimulus. This allowed us to make a more direct comparison between the oscillation cycles as observed in both paradigms.

One hallmark characteristic of the alternation cycle during continuous presentation is the distribution of durations that separate consecutive alternations. This distribution is invariably unimodal and right-skewed (Levelt, 1967; Borsellino et al., 1972; Zhou et al., 2004; Brascamp et al., 2005). Figure 6.4A shows this distribution as obtained in our sessions of continuous viewing. Its shape is as expected on the basis of previous findings.

To see whether the oscillation cycle during intermittent viewing could be analogous to the one during continuous viewing, we investigated whether the intermittent viewing paradigm produces a similar distribution of percept durations. Figure 6.4B shows the probability histogram obtained during intermittent viewing of a binocular rivalry stimulus. Its shape is qualitatively different from that obtained during continuous presentation: instead of peaking at an intermediate duration, this distribution is monotonically decreasing. This provides a dissoci-

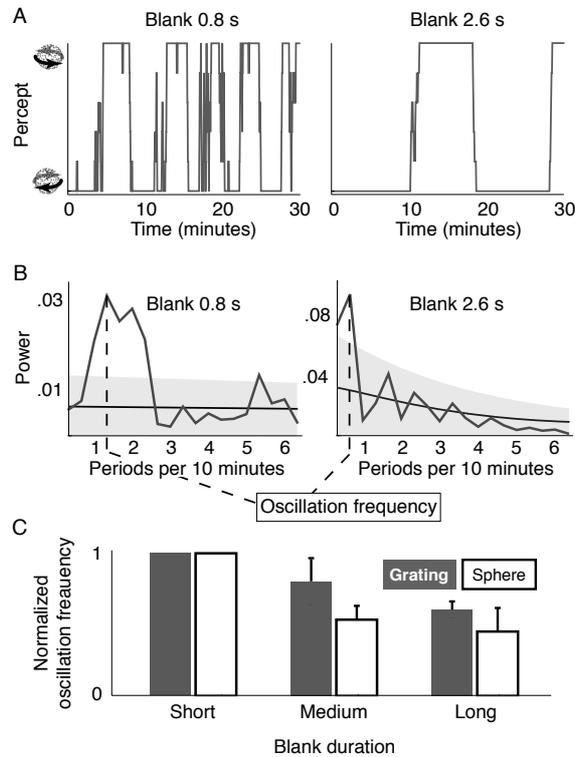


Figure 6.3: Relation between oscillation frequency and blank duration. *A.* Perception of an ambiguous sphere during a half-hour session (box car filtered with window size 5 to enhance clarity). Perception appears to oscillate faster in the condition with a shorter blank duration (left). *B.* Sections of the power spectra corresponding to panel *A.* The power peaks at a higher frequency in the condition with a shorter blank duration (left). *C.* Data for all subjects. Increasing the blank duration (*x*-axis) caused perception to oscillate at a lower frequency (*y*-axis), for both rivalry ($n=6$) and the sphere ($n=4$). Error bars are standard errors of the mean.

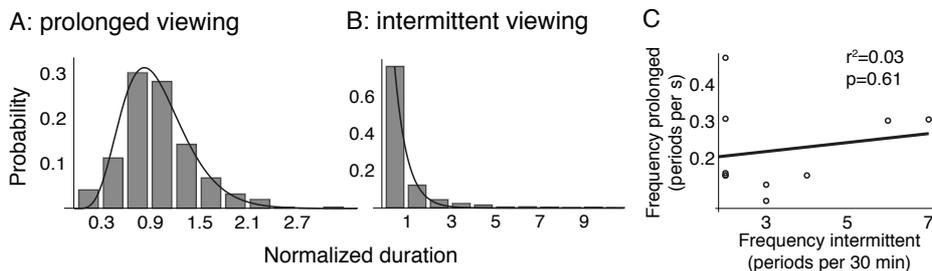


Figure 6.4: Comparison between the alternation cycles during intermittent presentation and during continuous presentation, respectively. *A* and *B.* The distribution of percept durations during continuous presentation is unimodal and right-skewed, with an exponential-like tail. During intermittent viewing this distribution is monotonically decreasing, with a tail that is thicker than exponential. *C.* Inter-subject variation in percept durations during intermittent viewing is not obviously correlated with inter-subject variation during continuous viewing.

ation between the perceptual oscillations observed during intermittent viewing and the ones associated with continuous viewing.

On top of this qualitative difference we mark one further, quantitative difference between the distributions of Figures 6.4A and 6.4B. Distributions obtained during continuous presentation have a rightward tail that is approximately exponential. Accordingly, these distributions can be reasonably described using a gamma distribution, which has an exponential tail (Pestman, 1998). This is illustrated by a fitted gamma distribution in Figure 6.4A. In contrast, the rightward tail of the distribution obtained during intermittent viewing has a falloff that is much slower than exponential. This is illustrated by the exponential distribution drawn in Figure 6.4B. This curve was fitted to the two left-most bars. If the data were approximately exponentially distributed, the curve should also be in reasonable agreement with the rest of the bars. That this is not the case can be appreciated by considering that all bars beyond a value of 5 on the x-axis still reach to a height of about 0.01, even though around an x-axis value of 5 the fitted curve falls below 0.000001.

As a second comparison between intermittent and continuous presentation, we examined inter-subject variability in the frequency of perceptual alternations. There is a considerable spread between subjects in the average perceptual alternation frequency, which provides a handle to compare various forms of ambiguous perception (Carter and Pettigrew, 2003; Pressnitzer and Hupé, 2006). If the relative alternation frequency of a subject in one paradigm tends to correlate with their alternation frequency in another paradigm, this is an indication that both may draw on similar neural mechanisms. Panel C plots alternation frequencies of nine individual subjects during continuous viewing against their alternation frequencies during intermittent viewing. We do not observe a significant correlation between both variables. This provides a further indication that perceptual alternations during intermittent viewing are fundamentally different from those observed during continuous viewing.

6.3.4 *Alternations between facilitation and suppression*

What causes perceptual alternations during intermittent viewing? A plausible account is provided by an existing model that explains many aspects of perception during intermittent viewing (Noest et al., 2007; Brascamp et al., 2008; Klink et al., 2008). According to this model, perceptual stabilization comes about when perception of an ambiguous stimulus leaves a facilitating trace that causes the same interpretation to regain dominance at the onset of a subsequent presentation.

It is important at this point to discuss a phenomenon that is quite different from perceptual stabilization, and that is also observed during intermittent presentation of an ambiguous stimulus. This phenomenon is observed in slightly different presentation regimes than those that prompt stabilization. In these regimes perception on every presentation of an ambiguous stimulus tends to be opposite to that on the previous presentation, rather than the same (Orbach et al., 1966; Noest et al., 2007; Klink et al., 2008). This thus constitutes an instance of perceptual destabilization.

The difference between presentation regimes that give rise to perceptual stabilization and those that result in destabilization lies in the duration of the blank interval between presentations. In general, blank durations above about a second produce robust stabilization, whereas those below, say, half a second give rise to destabilization (At extremely short blank durations, below around 50 ms, behavior changes again, as the presentation regime becomes effectively

continuous).

The model parsimoniously explains the existence of both stabilization and destabilization. It posits that the trace left by prior perception is in fact twofold in nature. Rather than merely facilitating repeated dominance of the same interpretation, the trace has both a facilitating component and a suppressive component. In principle, a trace of past perception can therefore cause both stabilization and destabilization. According to the model, which of the two effects wins out is determined by the *magnitude* of the trace. This idea is interesting because it changes the categorical difference between facilitation and suppression that is observed at the output level to a difference in degree on the neural level, explaining many experimental findings (see also Discussion). If a trace of past perception is weak it tends to facilitate, whereas if it is strong it tends to suppress. As described by the model, the reason for this is that the two opposing effects scale differently with the magnitude of the trace, and the facilitatory effect can not keep pace with the suppressive one at higher levels (Noest et al. (2007); see also Appendix).

Figures 6.5A and B give a simplified illustration of perceptual stabilization and perceptual destabilization in the context of this model. A full account is provided in the Appendix. The top panels show neural activity associated with perception of an ambiguous sphere. The dark curve denotes activity associated with perception of a leftward rotating sphere, and the light gray curve denotes activity associated with perception of a rightward rotating sphere. The bottom panels show the accumulation and decay of traces of past perception. During perception of the leftward rotating sphere the dark trace builds up, whereas the light one builds up during perception of the rightward rotating sphere. During the intervening blank periods both traces decay. To schematically illustrate the difference between high and low levels of these traces, the dashed line separates levels that cause facilitation from higher levels that cause suppression.

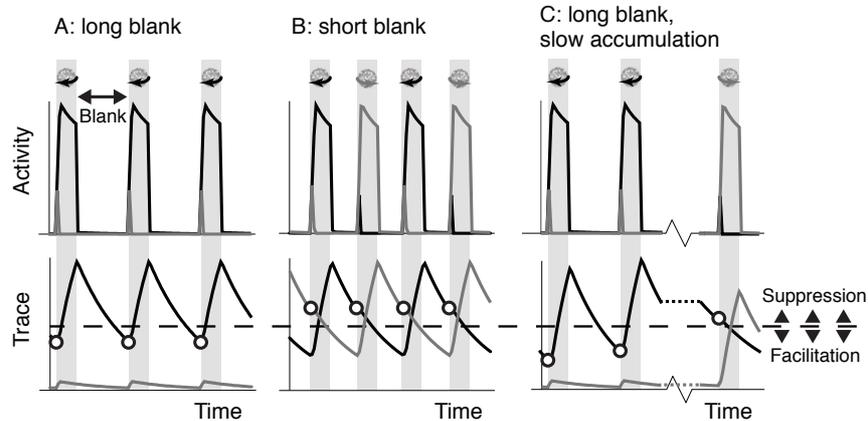


Figure 6.5: Schematic illustration of our model account of perception during intermittent viewing. Top panels: activity of the two percepts' neural representations. Bottom panels: resulting traces that determine perception during intermittent viewing. The central idea is that weak traces strengthen a representation and cause repeated dominance, whereas strong traces weaken a representation and cause suppression. A. Long blank durations allow traces to fall to a facilitatory level (below dashed line) between consecutive presentations, resulting in stabilization. B. Brief blank durations do not allow sufficient decay between presentations, leading to perceptual alternations. C. A trace that is composed of both fast and slow components can steadily grow from one presentation to the next. Then, blank durations that initially allow perceptual stabilization can, after many presentations, lead to an alternation when slow components lift the trace out of its facilitatory range.

On the first presentation in panel A the leftward turning sphere dominates. Accordingly, its trace accumulates. What is important for deciding which percept will be experienced throughout the following presentation, is the level of the trace at the next stimulus onset (marked by a open ball). In this example, the blank duration is long enough to allow the trace to fall to a facilitatory level before the next stimulus onset. Thus, this blank duration causes perceptual stabilization.

Panel B shows the situation when the blank interval between presentations is not long enough to allow the trace to fall to a facilitatory level before the next stimulus is shown. Traces of the most recent percept consistently lie in their suppressive range at the onset of the next presentation (the open ball lies above the dashed line). As a result, perception consistently alternates from one presentation to the next. This is an example of perceptual destabilization.

This reasoning can be generalized to account for our findings. It has been shown that traces of past perception do not just accumulate and decay on the time scale of individual presentations, such as in Figures 6.5A and B, but also on slower timescales (Brascamp et al., 2008). The essential difference is that slower traces do not rise and fall back to the same level over one presentation cycle. Instead, they can take many presentations of repeated perception to accumulate. What decides perception at the onset of a new stimulus is the sum of various traces that have built up on different timescales; some during the past few seconds, others during the past minutes.

Figure 6.5C illustrates schematically how this helps explain the findings in our present study. Like panel A, this panel shows the situation during perceptual stabilization. In contrast to panel A, which considered only fast traces of past perception, panel C takes into account slow traces as well. Accordingly, the curves in the bottom panel represent the sum of both fast and slow traces. Therefore, the dark curve, associated with the stabilized percept, slowly rises with every consecutive presentation, on top of rising and falling with every presentation cycle. It can now be readily seen that a perceptual alternation is expected on the first stimulus onset at which the total accumulated trace has risen out of the the facilitatory range, into the suppressive range. In 6.5C this occurs on the last depicted presentation.

Our experiments show two key characteristics of perceptual alternations during intermittent viewing. First, they occur at regular intervals and, second, the duration of these intervals decreases as the blank duration between presentations is reduced. The model account illustrated in Figure 6.5 is consistent with both these findings. According to this account, the duration of the intervals that separate two consecutive perceptual alternations is determined by the number of presentations required for the accumulating trace to rise out of the facilitatory range. This explains why these intervals are approximately constant at a given setting. The observation that these intervals become shorter as the blank duration is reduced is also consistent with the present account, which predicts them to become shorter for two reasons. First, at shorter blank durations the stimulus is present during a larger fraction of the time. Consequently, the slow traces will take less time to accumulate. A second factor is that a short blank duration allows less time for the fast trace to decay before the next stimulus onset. As a result, a smaller contribution of the slowly accumulating trace is required before the combined trace has a suppressive effect. Both factors cause more frequent alternations at lower blank durations. In the Appendix we provide simulation results that confirm that the model can account for our experimental findings.

6.4 Discussion

Perception of an appearing ambiguous stimulus is influenced by traces of past perception. What is the nature of these traces? The computational model we employ assumes that they take the form of sensitivity changes in neurons that code competing interpretations of an ambiguous stimulus. When a perceptual interpretation is dominant the sensitivity of the neurons that code that interpretation is altered. The choice between both percepts at the reappearance of a stimulus is determined by competition between the neurons coding the alternative interpretations. The outcome of this competition is heavily dependent on the relative sensitivities of the competing neurons, and therefore on past perception. This idea has proven very useful in explaining various findings regarding perception of intermittently presented ambiguous stimuli. Previously, variations of the model we employ have accounted for perceptual stabilization and perceptual destabilization (Noest et al., 2007; Klink et al., 2008), the observed influence of both short-term and long-term perceptual history on subsequent perception (Brascamp et al., 2008) and the interaction between perceptual history and volitional control (Klink et al., 2008). The explanatory power of this account indicates that it may accurately capture the essential interactions underlying perception of intermittently presented ambiguous stimuli.

An essential characteristic of this account is that it supposes that during perceptual memory, information on past perception is retained in the form of neural *sensitivity*, rather than as *activity*. In that sense the memory trace is thus more akin to a latent capacity to generate a percept upon renewed stimulation, than to an active maintenance of the percept during stimulus absence. A similar distinction between activity-based storage and sensitivity-based storage has been noted in the context of working memory, where it was proposed that synaptic facilitation could mediate working memory in the absence of delay period spiking activity (Mongillo et al., 2008).

It was previously shown that perception of intermittently presented ambiguous stimuli depends on perceptual traces that accumulate on various timescales; some on the order of seconds and others on the order of a minute (Brascamp et al., 2008). Our present results indicate that the slowest of these traces of past perception can accumulate over many minutes, in some cases even more than ten minutes. This is required to account for the observation that consecutive alternations are separated by up to fifteen minutes in some of our sessions. Is it plausible that neurons that code an interpretation of an ambiguous stimulus undergo such slow changes in sensitivity? In general, stimulation of sensory neurons causes both very fast and very slow changes in the neurons' response characteristics (Ohzawa et al., 1985; Bonds, 1991; Müller et al., 1999; Ulanovsky et al., 2004; Albrecht et al., 1984). Also, changes in psychophysically measured detection thresholds, which presumably reflect changes in sensory neurons' response characteristics, are still in progress after more than ten minutes of stimulation (Rose and Lowe, 1982). A similar observation, of particular importance in the present context, is that the time course of the perceptual alternation cycle systematically changes during a session of continuous presentation of an ambiguous stimulus, pointing to progressive changes in the underlying neural structures (Lehky, 1995; Suzuki and Grabowecky, 2002; Hupé and Rubin, 2003; Van Ee, 2005; Suzuki and Grabowecky, 2007). Some of these changes in the alternation cycle are still ongoing even after thirty minutes of viewing (Hollins and Hudnell, 1980). In sum, it is plausible that neurons involved in processing ambiguous stimuli undergo sensitivity changes on many timescales when they are activated, including timescales slow enough to account for our findings.

Figure 6.4C shows no correlation between a subject's alternation frequency during intermit-

tent viewing and their alternation frequency during continuous viewing. We were concerned, however, that a potential correlation might be masked by the two individual points in the top left of the graph. These points correspond to two subjects who combine a relatively high alternation frequency during continuous viewing with a low alternation frequency during intermittent viewing. We wondered whether in these two subjects the low number of alternations during intermittent viewing might be caused by a severe bias toward one of the two percepts. Such a bias could prevent alternations during intermittent viewing, while leaving continuous viewing relatively unaffected (Carter and Cavanagh, 2007). To control for this possibility we calculated each subject's bias as the fraction of all presentations on which the more predominant percept dominated. For these two subjects the fractions were 0.6 and 0.62, respectively, compared to 0.59 averaged over all subjects. We conclude that bias has had little influence on these data points, supporting the conclusion that percept durations are not correlated between the two paradigms.

The overall lack of evidence in Figure 6.4 for a direct association between the perceptual oscillation cycle during intermittent viewing and the one observed during continuous viewing, is in agreement with current models, which ascribe alternations in both paradigms to distinct neural events (Pearson and Brascamp, 2008). It is also consistent with other empirical findings that dissociate both paradigms (Carter and Cavanagh, 2007; Brascamp et al., 2008). The model we employ accurately predicts perception during intermittent viewing as well as perception during continuous viewing within the same neural network. However, an analysis of its dynamics shows that the neural events that prompt perceptual alternations are different in both paradigms (Noest et al., 2007), in spite of the shared neural substrate. In other words, the same model network is brought into two distinct modes by the two stimulation regimes. In short, during continuous presentation perception depends on mutual inhibition between neurons that code alternative interpretations, and gradual weakening (adaptation) of the currently dominant neurons. The model produces an alternation when neurons that code a dominant interpretation have been weakened to a point where the alternative interpretation (with the aid of noise) is able to escape from suppression and becomes dominant in turn. In contrast, perceptual alternations during intermittent presentation depend on a percept selection process triggered by the onset of a new stimulus. At stimulus onset the stimulus has been absent for a while, and neither of the interpretations is dominant or suppressed. Instead, both interpretations start to compete for dominance simultaneously, during a race which is biased by traces of past perception. When this competition is decided in favor of the interpretation that was suppressed on the previous presentation, this constitutes a perceptual alternation. Our present findings further strengthen the view that alternations during intermittent viewing are not directly related to those during continuous viewing.

A key property of the model we use to interpret our data is that it predicts traces of past perception to have a facilitatory effect on percept selection at moderate levels, but a suppressive effect at higher levels (Figure 6.5). This idea is consistent with many psychophysical observations. In the context of intermittently viewed ambiguous stimuli, it explains both our present findings and the transition from perceptual stabilization to perceptual destabilization at shorter blank intervals (Figure 6.5). It also accounts for several findings regarding the effect of *unambiguous* stimuli on subsequent perception of an ambiguous stimulus. When a disambiguated version of an ambiguous stimulus is viewed prior to an ambiguous stimulus, this can either cause the second stimulus to be perceived in the same way as the first one, or it can prompt dominance of the alternative interpretation. The distinction between the two effects lies in what

may be called the energy of the unambiguous stimulus: brief or low contrast prior stimuli facilitate repeated dominance, whereas long or high contrast prior stimuli suppress the previewed interpretation (Long et al., 1992; Kanai and Verstraten, 2005; Brascamp et al., 2007; Pearson et al., 2008). This is consistent with the idea of a transition from facilitation to suppression as a trace of past perception grows stronger. Indeed, the same model employed here can also account for the effects of unambiguous prior stimuli in certain conditions (Brascamp et al., 2007). It is known that traces left by unambiguous and ambiguous stimuli are not identical (Pearson and Clifford, 2005; Kanai and Verstraten, 2005). In general, unambiguous prior stimuli have a stronger tendency to cause suppression, whereas ambiguous prior stimuli are more commonly associated with facilitation. However, both types of stimuli obey the same principle regarding the relation between the strength of a trace and the direction of its effect: weak traces cause facilitation but strong traces cause suppression.

Our results indicate that a facilitatory trace of past perception may prompt repeated dominance of the same percept for minutes, only to then change into a suppressive trace that causes a perceptual alternation. One may wonder what functional mechanism our experiments probe. On a functional level, it is not obvious why the visual system would employ two opposing effects of past perception, and why one effect would overtake the other after a certain viewing period. Here we would like to speculate that facilitation and suppression correspond to two distinct requirements that the visual system has to meet. Suppression in our paradigm is the result of diminished neural sensitivity following prior activation. Such diminished sensitivity is ubiquitous in neural systems (Ohzawa et al., 1985; Bonds, 1991; Albrecht et al., 1984). Functionally, this so-called ‘gain control’ is believed to help tune neurons’ sensitivities to the intensity of prevailing input, thereby increasing the range of input intensities across which neurons can code effectively (Ohzawa et al., 1985; Abbott et al., 1997). Perceptual facilitation, on the other hand, corresponds to enhanced neural sensitivity in response to prior stimulation. We suggest that this enhanced sensitivity has the functional purpose of *increasing the system’s readiness to respond to newly appearing stimuli that were of behavioral relevance in the past*. Although this idea must remain speculative at this point, it is consistent with several characteristics of perceptual facilitation.

First, in order to promote the ability to act on newly appearing input, the effect of enhanced sensitivity should specifically be associated with stimulus onsets. This is indeed the case. Past perceptual dominance can facilitate repeated dominance at stimulus reappearance, but during continuous viewing it invariably has the opposite effect of shortening dominance episodes (Blake et al., 1990, 2003; Fang and He, 2004). Indeed, during a period of continuous viewing which immediately follows a period of intermittent viewing, the percept that was stabilized during intermittent viewing enjoys briefer dominance episodes than the alternative interpretation (JB, unpublished observations). In that situation the same percept can thus be said to simultaneously be ‘preferred’ during intermittent viewing and ‘non-preferred’ during continuous viewing. The same observation applies to the effect of unambiguous prior stimuli. Perception of an appearing ambiguous stimulus can be biased in favor of an interpretation that was previewed using a disambiguated version of that stimulus. In contrast, the dominance durations of the previewed interpretation during continuous viewing are invariably shorter than those of the alternative interpretation (Blake et al., 1980; Fang and He, 2004; Long and Toppino, 1994; Long et al., 1992). Thus, facilitation seems to be restricted to stimulus onsets, whereas suppression is prominently observed during continuous stimulation as well. This is consistent with the idea that, in terms of its function in natural vision, facilitation is specifically

related to the processing of appearing stimuli, rather than sustained stimuli. Incidentally, this distinction between appearing and continuous stimuli is adequately captured in the model we employ, where the facilitatory effect essentially constitutes a (subthreshold) offset in activity, whereas the suppressive effect takes the form of a gain reduction (Noest et al. (2007), see also Appendix). The offset can provide a decisive head-start at the onset of stimulation, but when responses have reached a steady-state level it is swamped by gain reduction.

The second part of our conjecture is that facilitation enhances processing of stimuli specifically if they were of behavioral relevance in the past. From a functional perspective, it seems reasonable that processing should not indiscriminately be enhanced for any stimulus that was presented before, but specifically for stimuli that were relevant. One fairly automatized way in which the visual system may make a rough distinction between previously relevant stimuli and previously irrelevant stimuli is by observing the outcome of past visual selection processes, because stimuli that previously gained preferential processing had some asset that dissociated them from competing stimuli. This reasoning is similar to that employed to interpret instances of sequential priming in attentional selection (Maljkovic and Nakayama, 2000; Kristjánsson, 2006). Thus, we propose that facilitation is particularly pronounced if the prior stimulus gained some form of preferential processing, whereas suppression occurs more indiscriminately. This idea provides some context for the observation that ambiguous prior stimuli have a stronger tendency to facilitate than unambiguous prior stimuli. Unambiguous stimuli do not involve clear preferential processing of the perceived interpretation, which is consistent with the fact that they leave only a modestly facilitatory trace.

In sum, we speculate that intermittently presented ambiguous stimuli probe both a facilitatory capacity that is functionally associated with processing of relevant stimuli when they reappear, and a suppressive capacity that is related to gain control. Two ways in which the visual system may segregate these effects functionally is by means of a difference in their moment of action (facilitation acts specifically at stimulus onset; suppression also acts during steady-state processing), and a difference in their dependence on stimulus relevance (facilitation acts specifically on relevant stimuli; suppression acts more indiscriminately).

6.5 Conclusion

We demonstrate that intermittent presentation of an ambiguous stimulus can give rise to periodic alternations between both perceptual interpretations. These alternations are separated by minutes to tens of minutes, and they become more frequent when the blank interval between presentations is reduced. Our interpretation of these findings centers on the idea that a trace of past perception that initially facilitates, can turn into a suppressive one as it accumulates higher. We discuss the functional distinction between suppression and facilitation in visual processing. These results are consistent with current ideas regarding perception of intermittently presented ambiguous stimuli. Furthermore, they indicate that perception of these stimuli is dependent on traces that accumulate over more than ten minutes.

Appendix

Model account of perceptual memory

Intermittent presentation of an ambiguous stimulus is a common method to study history effects in ambiguous perception. The first indication that perceptual memory can operate both in a facilitatory domain and in a suppressive domain, came from such experiments. In these experiments the blank duration between presentations was held constant throughout a session. Then, seconds-long blank durations will lead to facilitation, resulting in repeated perception of the same interpretation (Orbach et al., 1963; Leopold et al., 2002). This is the paradigm used in this study. If blank durations consistently lie below about half a second, on the other hand, this causes suppression throughout a session, resulting in alternating perception on consecutive presentations (Orbach et al., 1966; Noest et al., 2007; Klink et al., 2008).

Noest et al. (2007) accommodated those findings using a model that forms the backbone of our present account. It consists of two neural populations, each representing one perceptual interpretation. The two interact via cross inhibition, and their activities are modulated by a slow signal that depends on past activity: the memory trace. Each population is represented by two differential equations, one describing activity H , and another one for the slow history-dependent signal, called A . The two populations are indexed as i and j , and we show only the equations for i . Those for j are the same with indices i and j exchanged. The time derivative of activity H of population i given by:

$$\partial t H_i = X_i - (1 + A_i)H_i + \beta A_i - \gamma S[H_j] \quad (\text{A6.1})$$

Here X_i is i 's excitatory input, γ is a constant that determines the strength of cross inhibition and $S[H_j]$ is a sigmoid function of H_j (H itself can be viewed as a population-averaged membrane potential, and $S[H]$ as a population-averaged firing rate). The slow trace A affects activity H in two distinct ways: when A rises it reduces the response gain of this population via the term $-(1 + A_i)H_i$, whereas via the term $+\beta A_i$ it elevates the population's resting activity. Noest et al. (2007) used standard 'leaky integration' of spiking activity $S[H]$ to model the slow memory signal:

$$\tau \partial t A_i = -A_i + \alpha S[H_i] \quad (\text{A6.2})$$

Here $\tau > 1$ is a time constant, chosen such that A rises and falls over a large fraction of its total range within a single stimulus presentation and blank interval, respectively. α is another constant that modulates the maximum level of A .

This architecture gives rise to either facilitatory memory or suppressive memory, depending on the exact combination of A_i and A_j . This is illustrated in Figure A6.1A, where the gray area indicates regions of suppression, whereas the white area corresponds to facilitation. The existence of both perceptual stabilization at long blank intervals and perceptual alternations at short blank intervals, fits well with this layout. Figure A6.1B shows the course of the system in terms of this landscape during intermittent presentation with long blank intervals. Long intervals allow the system to sink into the facilitatory domain before every stimulus onset. Thus, they give rise to facilitatory memory, and the same percept appears on every presentation (in this case percept i). If short blank intervals are applied instead, the system does not have time to fall into the domain of facilitatory memory before the stimulus reappears. Instead,

the system remains in the domain of suppressive memory, resulting in opposite percepts on consecutive presentations. This is illustrated in Figure A6.1C.

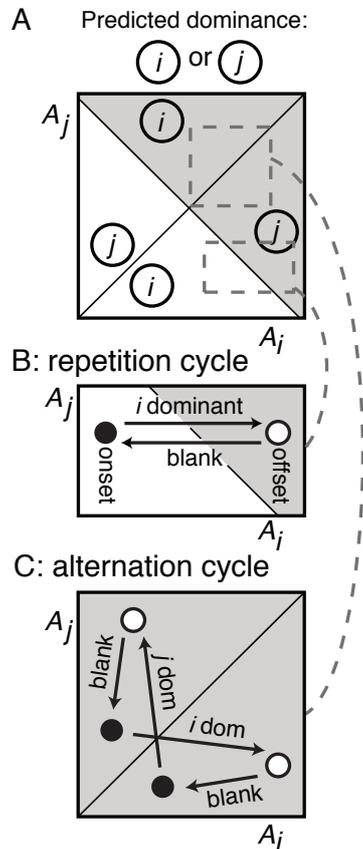


Figure A6.1: Perception at stimulus appearance depends on the values of traces A_i and A_j . A. General layout of the 'percept choice landscape'. At moderate levels of traces A_i and A_j , the population with the highest A value gains dominance (white). This may be called a facilitatory effect of A . At higher levels of traces A_i and A_j the population with the lowest A value gains dominance (gray). Thus, here A has a suppressive effect. B. Long blank durations lead to percept repetition. C. Short blank durations cause perception to alternate on consecutive presentations.

This basic model accounts for the existence of both facilitatory memory and suppressive memory. However, it does not explain our present finding that the system can switch between the two modes over the course of a session, even though the presentation paradigm remains unchanged. To account for that finding the model has to be expanded in two ways.

First, the perceptual oscillations we observe have a period length of up to twenty minutes in some conditions. Thus, a percept may dominate repeatedly for about ten minutes (barring the incidental perceptual alternation back and forth) before memory switches from the facilitatory regime to the suppressive regime. This implies that A may accumulate toward the border between the two regimes over the course of hundreds of presentations: much slower than the time constant in equation A6.2 allows. Thus, the model has to be expanded to include multiple A terms; faster ones responsible for up and down movements through the landscape during every single stimulus on-off cycle, and slower ones that cause the gradual accumulation of A

over many presentations. Independent empirical evidence has previously prompted the same conclusion that both fast and slow timescales are required (Brascamp et al., 2008). That previous work indicated slow processes with a time constant on the order of a minute; the extremely long period durations in the present work show that even slower processes must be present as well. We incorporate these requirements by replacing equation A6.2 with three copies of the same equation, each for a different timescale.

$$\tau_1 \partial t A_{i,1} = -A_{i,1} + \alpha_1 S[H_i] \quad (\text{A6.3})$$

$$\tau_2 \partial t A_{i,2} = -A_{i,2} + \alpha_2 S[H_i] \quad (\text{A6.4})$$

$$\tau_3 \partial t A_{i,3} = -A_{i,3} + \alpha_3 S[H_i] \quad (\text{A6.5})$$

Here τ_1 , τ_2 and τ_3 are three different time constants; one on the order of seconds, one on the order of minutes and one on the order of tens of minutes. In equation A6.1 the A terms should now be replaced with the sum of A_1 , A_2 and A_3 , written as ΣA :

$$\partial t H_i = X_i - (1 + \Sigma A_i) H_i + \beta \Sigma A_i - \gamma S[H_j] \quad (\text{A6.6})$$

With this extension the model can account for the accumulation of slow changes over the course of multiple minutes. This is illustrated in Figure A6.2A. Each dot corresponds to a stimulus onset. From the lightest to the darkest shade they depict the progression through time from the very first onset in this sequence to the onset at which suppressive memory takes effect.

The second extension implied by our findings is more intricate than the first. Its necessity is illustrated in Figure A6.2B. This panel is essentially the continuation of panel A. It shows how the system would continue after having reached the suppressive memory regime, if we do not modify the model any further. As soon interpretation j gains dominance over interpretation i , the system crosses to the opposite side of the landscape (arrow marked ' j dom'). During the subsequent blank interval ΣA levels drop again (arrow marked 'blank'). However, at the vast majority of parameter settings using this model, this drop does not move the system into the facilitatory memory regime. Instead, it stays in the suppressive regime, and dominance alternates back to interpretation i on the next presentation (arrow marked ' i dom'). Thus, once a slow trace has built up sufficiently to reach the suppressive regime, it tends to stick in that regime on subsequent presentations. As a result, the model without further modification predicts that a period of prolonged stabilized perception (Figure A6.2A) will be followed by an indefinite period of constantly alternating perception (Figure A6.2B). These alternations are analogous to what is observed at brief blank durations (Figure A6.1C), but in this case they occur at long blank durations due to an accumulated slow trace that lifts the system out of the facilitatory regime.

The scenario of Figure A6.2B is not what we observe in our experiments. Instead, after a period of stabilized perception of one interpretation has ended, we observe stabilized perception of the opposite interpretation. Geometrically speaking, this observation can be reconciled with our model, if we modify it such that the border between the facilitatory regime and the suppressive regime in the landscape does not run straight, but *convex*. Any type of convexity will tend to shift the model's behavior in the direction indicated by the empirical data. This is illustrated in Figures A6.2C and D. By definition a convex border runs lower near the x and

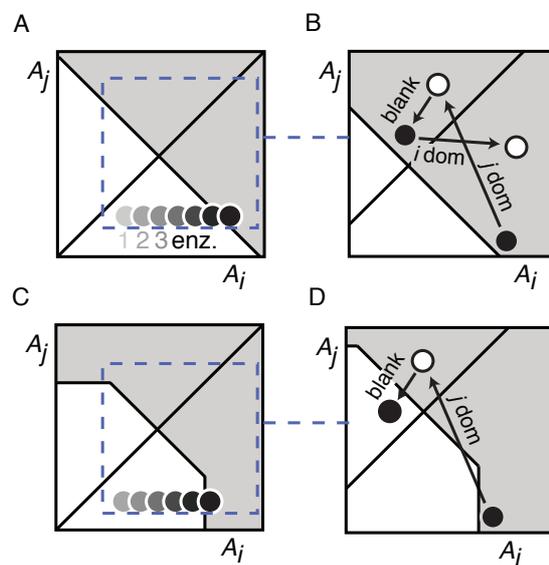


Figure A6.2: A convex border between facilitation and suppression allows slow perceptual alternations. A. Each consecutive stimulus onset (balls) during stabilization lies higher up in the landscape, due to slow accumulation of ΣA . B. In a model based on equation A6.6, prolonged stabilization tends to give way to ongoing alternations, which is not what is empirically observed. C and D. Modifying the model such that the border between facilitation (white) and suppression (gray) is convex rather than straight, allows prolonged stabilization of one percept to take turns with prolonged stabilization of the other percept.

y axis than it does near the center of the landscape. Moreover, toward the end of a stabilized sequence the system is located close to one of the axes, because prolonged suppression has allowed the ΣA level of the non-stabilized interpretation to fall very low. Thus, the system crosses the border between the facilitatory domain and the suppressive domain at a location where that border runs relatively low (panel C). Panel D shows the subsequent course of events during dominance of interpretation j and the following blank interval. Due to the convexity, these events land the system in a region where the border between facilitation and suppression runs relatively high. This allows the system to fall back into the facilitatory domain, resulting in stabilized perception of the opposite percept, instead of continued perceptual alternations (cf. panel B). In other words, if the border between the facilitatory domain and the suppressive domain is convex, a slow trace may build up sufficiently to terminate a sequence of perceptual stabilization, without automatically preventing stabilization of the opposite percept as well.

In the model this convexity can be achieved in a number of different ways, even without changing any circuitry, by replacing one of the linear terms in equation A6.6 with a nonlinear term. Thus, we can account for our experimental findings using an expanded version of the model by (Noest et al., 2007), which includes memory traces on multiple time scales, as well as a nonlinearity that causes a convex border between the facilitatory regime and the suppressive regime in the percept choice landscape. This is demonstrated in Figure A6.3, which shows simulation results with such a model.

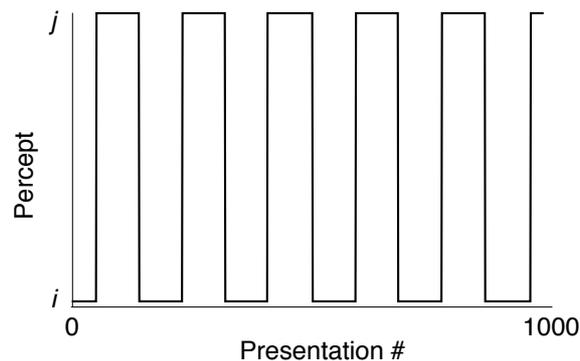


Figure A6.3: *Simulation results. With the addition of multiple timescales of accumulation and a nonlinearity, the model by Noest and colleagues can account for the occurrence of slow perceptual alternations during intermittent viewing. In this simulation an alternation period takes about 175 presentations; comparable to our experimental findings. Parameter values for this simulation were $\gamma=3.3$, $\beta=0.21$, $\beta'=0.04$, $\alpha_1=3$, $\alpha_2=0.5$, $\alpha_3=0.5$, $\tau_1=100$, $\tau_2=2000$, $\tau_3=4000$ and $X=1$. The stimulus presentation duration was 50; the blank duration was 140 (arbitrary units). The simulation was performed using a fourth order Runge-Kutta method with step size 0.1.*

What does a convex border between the facilitatory domain and the suppressive domain mean in terms of neural interactions? In fact, in our model the convexity can be achieved in a variety of ways, all involving the substitution of some linear term in equation A6.6 with a nonlinear alternative. On the one hand, this means that convexity of the border is neurally plausible, because the assumption of linearity is often not met in neural systems. On the other hand it means that this particular finding is not a very rigorous model constraint and does not force any specific conclusion in terms of the underlying interactions. In our simulations we obtained a convex border by changing the way ΣA modulates the resting level of H or, in other

words, by changing the $+\beta\Sigma A$ term:

$$\partial t H_i = X_i - (1 + \Sigma A_i) H_i + \max[0, \beta \Sigma A_i - \beta' \Sigma A_j] - \gamma S[H_i] \quad (\text{A6.7})$$

Here $\max[0, value]$ represents a threshold nonlinearity that returns $value$ if $value > 0$, and 0 otherwise. β' is a constant. Thus, this equation implements a kind of cross interaction between trace ΣA_i , and trace ΣA_j . In the original model the resting level of H_i was modulated upward by its own memory trace ΣA_i . In this expanded version it is both modulated upward by ΣA_j and simultaneously modulated downward by the trace of the competing percept, ΣA_j . Antagonistic interactions between oppositely tuned populations are common in neural systems (e.g. Snowden et al., 1991; DeAngelis et al., 1992).

It is worth commenting on the expansion from one timescale in equation A6.2 to three timescales in equations A6.3-A6.5. This expansion doubles the number of equations in the model, which appears to be a dramatic increase in the model's free dimensions and thus a reduction of its parsimony. While this model is certainly less parsimonious than the original one (which was based on less data), several factors should be noted. First, the expansion merely copies equation A6.2 to several timescales and thus changes little conceptually. Second, like many neural processes, the evolution of ΣA may well conform to a power law rather than to an exponential process (Anderson, 2001; Drew and Abbott, 2006). This casts a different light on the addition of timescales to our model, because capturing a power law function in differential equations necessitates an approximation using a series of exponential equations with different time constants.

Chapter 7

Perceptual memory: toward neural measurement

This chapter is the only one that is not written as a typical scientific paper. Instead, in this chapter I give an overview of ideas about perceptual stabilization that we have developed over the last couple of years, and their implications. I did this mostly because I feel our understanding of the phenomenon has advanced to a point where it can be used to guide the development of new questions and experiments.

Abstract

Being exposed to ambiguous visual input typically results in a highly unstable and unpredictable perceptual experience. Perception erratically alternates between possible interpretations, unable to settle on a definitive solution. However, inconclusive sensory information can sometimes be disambiguated by temporal context. When an ambiguous image reappears after temporary absence, its perception is often determined by traces left by previous perception. Recent years have seen considerable progress in the understanding of this perceptual memory. Computational models now reproduce many psychophysical observations and indicate what neural interactions may underly them. In contrast, there is little direct evidence of neural activity associated with perceptual memory. Here we analyze the behavior of a plausible model of perceptual memory, and we perform a psychophysical experiment that further tightens the link between model and empirical data. We subsequently explore the extent to which the model can guide a search for neural correlates of perceptual memory. We formulate several predictions that may help ensure an effective search, including an explicit prediction of neural spiking activity that may be observed in association with perceptual memory.

7.1 Introduction

Memory and perception are classically considered distinct domains of neural function. In recent years, however, a paradigm has gained interest in which the two are intimately linked. This paradigm shows that perception of certain, so-called ambiguous, images depends critically on a kind of memory (Pearson and Brascamp, 2008). Ambiguous images convey inconclusive information to the eyes, allowing room for various perceptual interpretations (Figure 7.1). They cause a highly unpredictable sensory experience, prompting perception to alternate erratically between alternative interpretations (Leopold and Logothetis, 1999; Blake and Logothetis, 2002). However, if an ambiguous stimulus has been seen before, memory traces left by previous encounters can make for highly predictable perception when the stimulus reappears. Then, perception of the reappearing stimulus is automatically determined by perception in the past. Apparently, perceptual memory traces can supplement sensory information, allowing the brain to choose an interpretation.

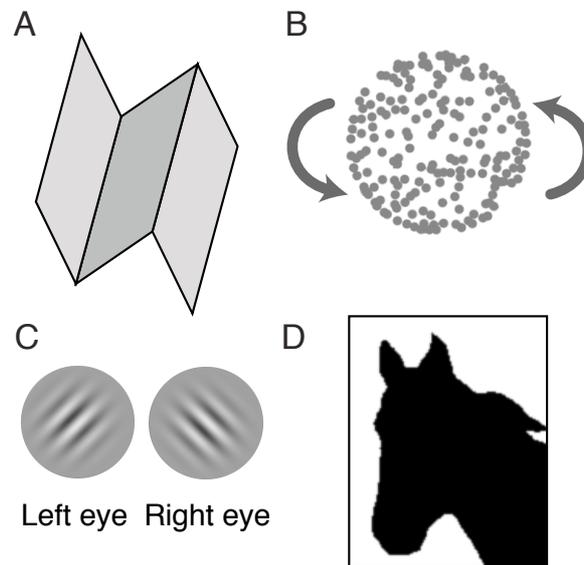


Figure 7.1: *Examples of ambiguous stimuli. A. The two light gray surfaces may face either left or right. B. One frame of a dynamic structure from motion stimulus, which consists of a two-dimensional projection of a transparent sphere rotating in depth. The direction of rotation is ambiguous. C. In binocular rivalry the two eyes view highly dissimilar images. Observers experience alternating perception of either image in isolation. D. The horse can be seen either as coming toward the observer or as moving away.*

Understanding this remarkable phenomenon can be a valuable step for research in the area of intersection between perception and memory. Curiously, when inquiring into the present level of understanding of this phenomenon, one may get highly disparate answers depending on the perspective of the speaker. On the one hand, there has been considerable progress in formalizing the neural interactions that underly the observed interplay between perception and memory. This has led to an effective description on a system dynamics level, that is supported by a broad range of psychophysical observations (Brascamp et al., 2008; Klink et al., 2008; Noest et al., 2007; Brascamp et al., in prep). In terms of system level computations, therefore, few of

the currently available empirical data remain unexplained. On the other hand, in terms of the actual neural concomitants of these computations, our present knowledge remains limited. There are only few published studies reporting neural measurements in this paradigm. One neuroimaging study implies the involvement of specialized visual areas in memory retention, while also suggesting a correlation between frontoparietal activity and the impact of memory on perception (Sterzer and Rees, 2008). A role for early visual areas is implied by work using event-related brain potentials (ERP) (Kornmeier and Bach, 2004, 2005).

A significant limitation of theoretical work based on psychophysical results, is that it uses virtual neural components to account for observations on the output level. How events within this theoretical construction, and corresponding perceptual phenomena, map onto actual activity of cells in the brain is uncertain. However, an important merit of computational work in this area is that it has allowed the construction of an explicit theory of perceptual memory. This has provided a framework to lay bare implications and pitfalls that are difficult to identify using less formal methods, and to formulate testable predictions that are indispensable for better scientific understanding. Theory and neural measurement are two complementary approaches, whose integration is required to comprehend the brain processes involved in perceptual memory.

Here we make use of the power provided by a modeling approach, to predict a complex and counterintuitive pattern of data that may be observed in psychophysical experiments of perceptual memory. This prediction hinges on anticipated changes in the state of virtual neurons, that would occur over the course of an experiment. Testing it requires an experimenter's ability to keep an accurate model of changes in an observer's brain during a session. We subsequently perform an experiment that tests this prediction, and demonstrate that it is accurate. Given the improbability of this data pattern when viewed outside of the context of computational theory, this finding provides strong evidence that current theory effectively captures the essential interactions underlying perceptual memory.

Based on further model analyses we propose specific patterns of neural activity that may accompany perceptual memory. Importantly, these analyses indicate that various apparently reasonable assumptions about the neural concomitants of perceptual memory may not be correct. Consequently, attempts at finding these neural concomitants, and of interpreting them, may be bound to fail if they are guided by common sense. Measuring the patterns that we predict requires theoretical knowledge of the dynamics of the system responsible for perceptual memory, as well as the ability to maintain a model of this system as it changes during a session. We demonstrate that our present psychophysical experiment, in combination with model analysis, provides the tools to meet these requirements, and to start designing theory-guided experiments into the neural concomitants of perceptual memory.

This work further deepens our theoretical understanding of perceptual memory. Moreover, it provides a starting point to translate the present level of advancement in the theoretical domain into an equivalent level of knowledge of the neural activity involved. This integration of theory and measurement can move the field toward a comprehensive understanding of the interaction between visual input and memory that underlies perceptual experience.

7.2 Results

7.2.1 Model description I: perception of an appearing stimulus

Our present model describes processing of a stimulus that has two perceptual interpretations. It assigns to each interpretation a separate neuronal representation. The two inhibit each other's activities, which assures a high activity of only one representation at a time, corresponding to perception of the associated interpretation. The response characteristics of the representations undergo slow adaptive changes during stimulation. This adaptation allows the outcome of competition between the two representations to vary depending on their prior activity. As a result, through adaptation past perceptual dominance can influence subsequent perception. This architecture is schematically illustrated in Figure 7.2A.

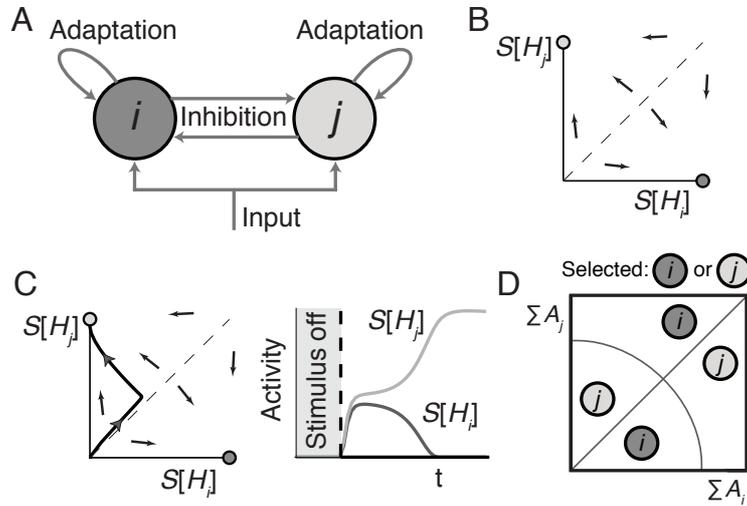


Figure 7.2: Basic characteristics of the present model. A. The model is based on an architecture consisting of two mutually inhibitory neural representations (i and j) that exhibit slow self adaptation. B. When a stimulus is present the model has two stable states. In one stable state the activity $S[H_j]$ of representation j is high, whereas activity $S[H_i]$ of representation i is low (light gray ball), corresponding to dominance of percept j . In the other stable state (dark gray ball) the opposite is true. C. At the onset of a stimulus, the system develops from a state of inactivity to either of the two stable states. D. Which of the two stable states will become occupied at stimulus onset, depends on the balance between the two representations' adaptation states, ΣA_i and ΣA_j .

The model we employ has been developed on the basis of a range of psychophysically observed characteristics of perceptual memory (Brascamp et al., 2008; Klink et al., 2008; Noest et al., 2007; Brascamp et al., in prep). Its behavior has previously been analyzed in detail (Brascamp et al., 2008; Noest et al., 2007). Here we summarize its key characteristics without proof. For a given adaptation state, the response properties of a representation are described by a single differential equation. Using indices i and j ,

$$\partial t H_i = X_i - (1 + \Sigma A_i) H_i + \beta \Sigma A_i - \gamma S[H_j] \quad (7.1)$$

gives the time derivative of activity H for representation i , and

$$\partial t H_j = X_j - (1 + \Sigma A_j) H_j + \beta \Sigma A_j - \gamma S[H_i] \quad (7.2)$$

does the same for representation j . Matching these equations to the diagram of Figure 7.2A, H is the activity of a representation i or j , X is the excitatory input provided by the stimulus, the term $-\gamma S[H]$ mediates mutual inhibition, and ΣA denotes the adaptation state of a representation. Adaptation state ΣA is present in two separate terms in these equations. This specific dual effect of the adaptation state of a representation on its response is the essence of the present model, and will be discussed below. Note that $S[H]$ in these equations denotes a sigmoidal transformation of H . That is, H represents a latent activation that is analogous to a cellular membrane potential, whereas $S[H]$ describes the resulting effective activity, analogous to a spiking response.

For a given combination of adaptation states ΣA_i and ΣA_j this system of two coupled differential equations develops to either of two stable states. In one stable state $S[H_i]$ is high and $S[H_j]$ is low; in the other the converse is true. The two balls in Figure 7.2B show these two stable states in a space that plots $S[H_i]$ against $S[H_j]$. The dashed line indicates a so-called separatrix: any system state above this line develops toward the top left stable state, whereas any system state below this line develops to the bottom right one. This distinction is also indicated by arrows in Figure 7.2B that show the direction in which the system develops at various locations. When the system reaches the stable state where $S[H_i]$ is high (bottom right), this corresponds to perceiving interpretation i , whereas the other stable state corresponds to perceiving interpretation j .

What decides how an ambiguous stimulus is perceived when it appears? In terms of our model, this question translates to: to which of the stable states does the system develop after input X is switched from zero to some positive value? In Figure 7.2B this percept selection process corresponds to a route that starts near the origin, and that ends at one of the two balls. One such route is drawn in 7.2C, left, which shows an example where interpretation j gains perceptual dominance. The same route, now plotted as a function of time, is plotted in 7.2C, right.

The route in Figure 7.2C, left, starts very close to the separatrix. In other words, in the brief period right after stimulus appearance, there is a very fine balance between developing toward perceptual dominance of interpretation i , and developing toward dominance of interpretation j . What factor decides which of the two interpretations is selected? In our model this factor is ΣA , the adaptation states of the competing interpretations. Indeed, if ΣA_i and ΣA_j are equal, the route after stimulus appearance starts exactly on the separatrix, and only noise can force a perceptual decision.

The model posits that an elevated adaptation state ΣA can either benefit a representation, or impair it. In equations 7.1 and 7.2, ΣA benefits a representation via the term $+\beta \Sigma A$, but simultaneously impairs it via $-H(1 + \Sigma A)$. An equivalent statement in words is that adaptation elevates the level of baseline activity H of a representation, while simultaneously reducing its response gain. Because of this dual, opposed, effect of adaptation, percept selection does not depend in a straightforward, monotonic way on ΣA . Instead, the facilitatory term $+\beta \Sigma A$ is particularly influential at moderate adaptation levels, but the suppressive term $-H(1 + \Sigma A)$ gains influence at higher levels. Figure 7.2D depicts an essential plot that summarizes how percept selection depends on adaptation levels ΣA_i and ΣA_j . Interpretation i is selected to be perceived when ΣA_i is moderately higher than ΣA_j (bottom region) or when ΣA_j is much higher than ΣA_i (top region). Conversely, interpretation j is selected at moderate levels of ΣA_j (left region), and at high levels of ΣA_i (right region).

7.2.2 Model description II: development of adaptation, and how it accounts for memory

So far we have considered the percept selection process for given levels of adaptation ΣA_i and ΣA_j , without addressing perceptual memory explicitly. Perceptual memory comes about when the adaptation levels change in response to stimulation, which provides the required link between past perception and subsequent percept selection. We can describe the evolution of an adaptation term A in response to stimulation as so-called leaky integration of activity $S[H]$. That is, saturating accumulation when $S[H]$ is positive, and passive decay when it is zero. This is illustrated in Figure 7.3A, which shows the evolution of A in response to a pulse in $S[H]$. In equation form, this amounts to

$$\tau \partial_t A_i = -A_i + \alpha S[H_i] \quad (7.3)$$

for adaptation A_i of representation i , and

$$\tau \partial_t A_j = -A_j + \alpha S[H_j] \quad (7.4)$$

for adaptation A_j of representation j . Here τ is a time constant that is larger than 1, assuring that adaptation changes more slowly than activity $S[H]$.

Although this is the basic form of changes in adaptation, psychophysical data are consistent with the notion that the total adaptation of a representation, denoted as ΣA , is the sum of several adaptation components that each evolve at different rates (Brascamp et al., 2008, in prep). Each component individually behaves as indicated by equations 7.3 and 7.4, but each has a different time constant τ . The resulting course of the summed adaptation ΣA in response to a series of pulses of $S[H]$ is illustrated in Figure 7.3B. Fast adaptation components are responsible for fast accumulation and decay during a pulse in $S[H]$ and the following pause, respectively. Slow adaptation components mediate an incremental accumulation of ΣA from one pulse to the next. To accommodate psychophysical observations, the fastest adaptation components must have a time constant on the order of seconds, whereas the slowest ones must have a time constant on the order of tens of minutes (Brascamp et al., in prep).

The model outlined above provides a context for understanding how a memory trace from prior perception can act to disambiguate an ambiguous stimulus. The key to gaining an intuition for the workings of the model lies in learning to translate perception of an ambiguous stimulus into the corresponding motions in the space introduced in Figure 7.2D. This space depicts the outcome of the percept selection process as a function of the adaptation states ΣA_i and ΣA_j .

Figure 7.3C, left, illustrates this idea for a basic stimulus sequence that consists of two consecutive presentations of the same stimulus, separated by a blank interval. This plot uses the space of Figure 7.2D to plot the course of adaptations ΣA_i and ΣA_j during the stimulus presentation sequence. At the moment the stimulus first appears (open circle marked '1') both adaptations are equal, and either interpretation may gain dominance. In this example interpretation i becomes dominant. Dominance of i corresponds to a high activity level of i but not j (cf. Figure 7.2C), and so ΣA_i rises during dominance, while ΣA_j stays low. This development is indicated by the gray arrow marked ' i perceived', which connects the moments of first stimulus appearance and subsequent disappearance. After the stimulus disappears ΣA_i starts decaying. This is indicated by the arrow marked 'blank'. This course of events causes the system to be located within the bottom middle region of this space at the moment the stimulus reappears

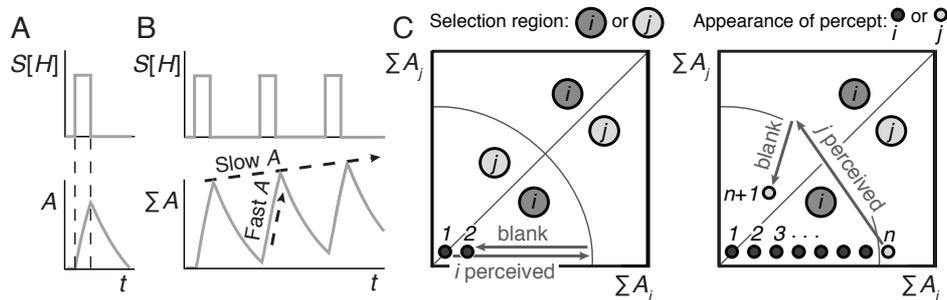


Figure 7.3: The role of adaptation during intermittent viewing of an ambiguous stimulus. A. Adaptation of a representation (bottom panel) is calculated as leaky integration of the representation's activity, $S[H]$. B. If the same percept dominates on consecutive presentations (top panel), fast adaptation components mediate a rise and fall in adaptation during every stimulus on/off cycle, while slow components cause a gradual rise from one presentation to the next. The total adaptation, ΣA is the sum of several components that evolve on separate timescales. C. Course of adaptation during dominance of interpretation i (arrow marked i perceived) and during the following blank interval (arrow marked *blank*). D. The course of adaptation during a sequence of presentations (marked '1', '2', '3' through 'n') with stabilized dominance of interpretation i . On interpretation n an alternation to j occurs because slow adaptation has lifted the system into a regime that favors perception of j at stimulus appearance.

(open circle marked '2'). Accordingly, stimulus reappearance will again give rise to perception of interpretation i . Thus, prior perception of the stimulus has determined perception on the second encounter.

Figure 7.3C, right, shows a course of events that occurs in many psychophysical studies of perceptual memory, and that is relevant for our present purpose. Such studies generally show the same ambiguous stimulus on many consecutive presentations, separated by blank intervals (Orbach et al., 1963; Leopold et al., 2002; Maier et al., 2003; Chen and He, 2004; Brascamp et al., in prep). During such prolonged sequences of intermittent presentations, observers commonly perceive the same interpretation on tens to hundreds of consecutive presentations. Such a period of stable perception of one interpretation is usually followed by a similarly long period of perceiving the alternative interpretation (Brascamp et al., in prep). Figure 7.3C, right, shows how this works out in the context of our model. Starting from the bottom left, the figure shows a sequence of consecutive stimulus appearances, marked '1', '2', '3', '...'. They are numbered in temporal order, starting from the beginning of a session. The first set of presentations in this session all give rise to dominance of interpretation i , so on every consecutive appearance ΣA_i rises slightly further (cf. Figure 7.3B). The appearance at which the alternative interpretation, j , first gains perceptual dominance, is marked with an ' n '. At this appearance, adaptation ΣA_i has risen into a region that corresponds to selection of interpretation j . During this presentation interpretation j is perceived. Accordingly, its adaptation ΣA_j rises, while the highly accumulated ΣA_i decays. This corresponds to a motion indicated by the arrow marked ' j perceived'. After the following blank period (arrow marked '*blank*') this course of events lands the system in the left region where interpretation j is selected. This appearance, marked ' $n+1$ ', marks the beginning of a prolonged sequence of appearances on which interpretation j gains dominance, as elaborated below.

7.2.3 Model prediction: psychophysics

A benefit of the above model account is that it permits explicit and testable predictions on the behavior of perceptual memory. This provides a basis for a deeper understanding of the phenomenon. Also, it allows a critical assessment of the veracity of the account, which is not always possible in the case of more informal accounts, whose implications may not be clear. The model was designed and refined on the basis of psychophysically measured characteristics of perceptual memory, and reproduces those characteristics in simulations (Brascamp et al., 2008; Klink et al., 2008; Noest et al., 2007; Brascamp et al., in prep). Here we investigate a precise and counterintuitive prediction that follows from the model in its existing form, thus putting the model to the test.

Perceptual memory may be defined as a relation between prior perception of a stimulus, and its perception on a subsequent encounter. However, illustrations such as Figure 7.3C indicate that the nature of this relation may not always be straightforward. For instance, the appearance marked ' n ' in Figure 7.3C, right, prompts dominance of interpretation j , whereas the immediately preceding appearance prompted dominance of interpretation i . This is remarkable, considering that these categorically different perceptual outcomes are the result of nearly identical sequences of prior perception. At appearance n prior perception causes a perceptual alternation to interpretation j , whereas at the preceding appearance it caused perceptual repetition of interpretation i .

Our present prediction states that whether a given appearance will prompt a perceptual alternation (relative to the previous presentation) or perceptual repetition, should depend on the duration of the blank interval that immediately precedes the appearance. Moreover, we predict that this relation between blank duration and alternation probability should change in a lawful way over the course of a session. That is, applying a particular blank duration may reduce the probability of an alternation in some stages of the session, whereas in other stages that same blank duration may promote an alternation.

This is illustrated in Figure 7.4. Figure 7.4A continues where Figure 7.3C, right, left off. The appearance marked ' $n+1$ ' is the same in both figures. In Figure 7.4A, a sequence of stable perception of interpretation i has just ended, and appearance ' $n+1$ ' marks the beginning of a sequence of j percepts. Consecutive appearances are numbered ' $n+1$ ', ' $n+2$ ', ' $n+3$ ' through ' o '. Each appearance is located slightly further to the top left, because adaptation of interpretation j keeps accumulating, while the previously accumulated adaptation of interpretation i gradually decays. Appearance ' o ' is located toward the end of this stabilized sequence of j percepts. If ΣA_j rises only slightly further from the location of appearance ' o ', the system will land in the top region where interpretation i is selected for dominance. When this happens, it will initiate a stabilized sequence of i percepts that is the mirror image of Figure 7.4A, thus perpetuating the cycle of alternating periods of opposite dominance.

To see what this means in relation to the influence of the blank duration on perceptual alternations, we zoom in on two parts of the landscape of Figure 7.4A. Figure 7.4B zooms in on the gray rectangle around appearance ' o '. Starting from appearance ' o ', dominance of interpretation j will cause further accumulation of adaptation ΣA_j . Whether this adaptation is sufficient to prompt a perceptual alternation, depends on the blank duration that follows. If the blank interval is relatively brief, the subsequent stimulus appearance is located in a region where interpretation i is selected, and a perceptual alternation to interpretation i ensues (Figure 7.4B, left). If, on the other hand, the blank interval is longer, ΣA_j has more time to decay (Figure 7.4B,

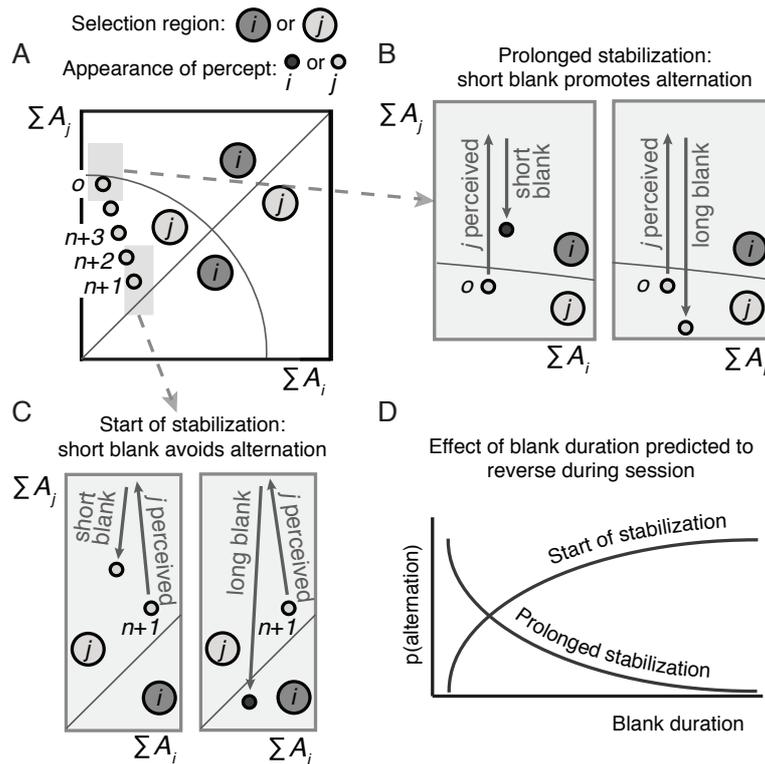


Figure 7.4: Model prediction: psychophysics. A. A period of stabilized perception of j continues from presentation $n+1$ through o , and ends when further adaptation lifts the system across the region boundary on the top left of the landscape. B. Outtake of panel A. Toward the end of the stabilization period, alternations are promoted by applying short blank durations, whereas long blank durations strengthen perceptual stability. C. Another outtake of panel A. At the beginning of the stabilization period, stability can be strengthened by applying short blank durations (left), whereas long blank durations induce alternations to i . D. Summary of the prediction illustrated in panels A-C. At the start of a period of stable perception increasing the blank duration between consecutive presentations should elevate the probability of a perceptual alternation. Toward the end of a period of stable perception increasing the blank duration should reduce the probability of an alternation.

right). As a result, the next stimulus appearance is located in a region where interpretation j is selected, resulting in percept repetition. Thus, once an interpretation has enjoyed a long period of stabilized perception, applying long blank durations should tend to prolong this stabilization even further, whereas short blank durations should prompt perceptual alternations.

Figure 7.4C zooms in on the pink rectangle in Figure 7.4A that is located around appearance ' $n+1$ '. Starting from appearance ' $n+1$ ' the system will move to the top left during stimulus presentation, as ΣA_j accumulates and ΣA_i decays. If the blank duration between presentations is sufficiently brief, the next stimulus appearance is located within a region where interpretation j is selected, causing percept repetition (Figure 7.4C, left). If, on the other hand, a longer blank interval is used, the system will sink into a region where interpretation i is selected, thus prompting a perceptual alternation to interpretation i (Figure 7.4C, right). This means that, if an interpretation has not yet enjoyed dominance on many consecutive presentations, applying short blank durations should help it gain dominance again, whereas long blank durations should cause perception to alternate to the opposite interpretation.

In sum, we predict that the influence of blank duration at the beginning of a sequence of stabilized perception should be the *opposite* of its influence toward the end of a sequence of stabilized perception.

It is worth noting that the arrows in Figures 7.4B-C have a vertical extent that is much larger than their horizontal extent. The reason for this is that this figure depicts a time period during which interpretation j repeatedly dominates over interpretation i . Accordingly, the vertical extent of the arrows is mainly due to relatively fast components of ΣA_j , which build up and decay during every presentation and subsequent blank, respectively. Interpretation i is not perceived in this time period. Therefore, its adaptation ΣA_i is largely made up of residual slower components that have accumulated during the prior stabilization period of interpretation i (which was depicted in Figure 7.3C, right). The gradual decay of these slow components is responsible for the modest horizontal extent of the arrows in Figures 7.4B-C.

Figure 7.4D summarizes the prediction that follows from this model analysis. We anticipate that the relation between blank duration and the occurrence of perceptual alternations reverses over the course of a session, in a predictable manner. After an interpretation has enjoyed stabilized dominance on many presentations, increasing the blank duration should reduce the probability of perceptual alternations. If, in contrast, an interpretation is just starting to become stabilized, increasing the blank duration should increase the probability of a perceptual alternation. We performed simulations that confirm that this analysis of model behavior is accurate (Appendix).

7.2.4 Psychophysical test

We performed a psychophysical experiment to test the above prediction. Our ambiguous stimulus consisted of a patch of left tilted gratings presented to one eye, and a patch of right tilted gratings simultaneously presented to the corresponding region of the other eye (Figure 7.1C, Figure 7.5A). This causes alternating perception of either of the images in isolation, a phenomenon called binocular rivalry (Blake and Logothetis, 2002). The temporal characteristics of binocular rivalry memory are typical for perceptual memory of ambiguous stimuli in general (Leopold et al., 2002; Brascamp et al., 2008).

Subjects viewed many consecutive presentations of the same stimulus, separated by blank intervals, and reported their percept on every presentation. The duration of each blank interval

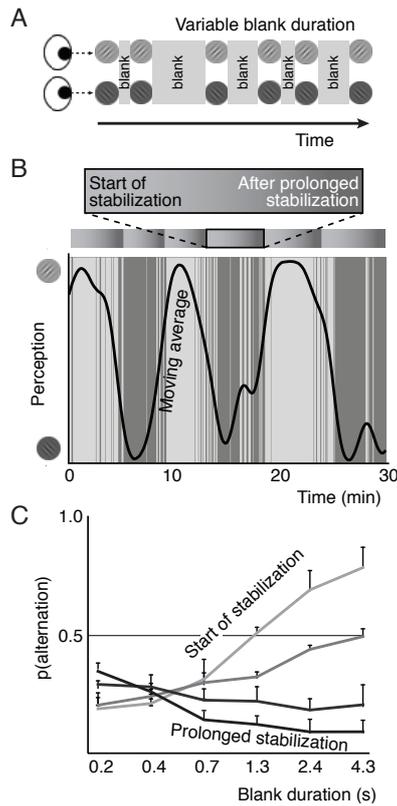


Figure 7.5: Psychophysical experiment. A. Subjects viewed a binocular rivalry stimulus presented intermittently, with randomly drawn blank durations between presentations. B. A typical session. Perception alternated erratically between both interpretations (indicated by the background shade), but there were underlying periods during which the same interpretation dominated on the vast majority of presentations (indicated by the moving average). We separated data according to the position within such a period of relative stability, ranging from 'at the start of stabilization' to 'after prolonged stabilization'. This is schematically indicated above the plot. C. Relation between blank duration (x-axis) and the probability of a perceptual alternation from a given presentation to the next (y-axis). At the beginning of stabilization (light curve) long blank durations promote alternations. After prolonged stabilization (dark curve) long blank durations promote stability. Intermediate shades correspond to intermediate stages of the stabilization sequence. These findings agree with the model prediction. Error bars are standard errors of the mean ($n=4$).

was drawn randomly from a set ranging from 215 ms to 4.2 s (Figure 7.5A).

Figure 7.5B plots perception during a typical session. The background shade of the plot indicates perception as reported by the observer. There are frequent alternations between the competing perceptual interpretations. In spite of this, there are prolonged periods during which the same percept is seen on the great majority of presentations. This is more clearly expressed by the smooth curve, which shows a moving average of reported percepts. Minutes-long periods during which the rightward tilted gratings dominate are alternated with minutes-long periods during which the leftward tilted gratings dominate. These prolonged periods of (relatively) stabilized perception are characteristic of the intermittent presentation paradigm. In terms of our model description they correspond to sequences of appearances like the ones shown in Figure 7.3C, right, and Figure 7.4A.

In order to test our prediction, we separated our data into subsets according to the location within a period of stabilized perception. These subsets ranged from ‘at the start of stabilization’ to ‘after prolonged stabilization’. This is schematically illustrated at the top of Figure 7.4B, and explained in detail in the Methods section.

For each subset separately we calculated the probability that perception would alternate to the non-stabilized percept on a given stimulus appearance, as a function of the blank duration preceding that appearance. Subject-averaged results are shown in Figure 7.4C. Like in Figure 7.4B, the four shades from light to dark correspond to the progression from the beginning of a stabilization period to the end of a stabilization period. The lightest shaded curve shows that at the beginning of a stabilization period the probability of an alternation increases with increasing off-duration. The darkest shaded curve shows the opposite effect: if a percept has already been stabilized for a long time, a long blank interval tends to keep it stabilized, whereas a brief interval will more likely yield an alternation. The intermediate shades have intermediate courses.

This experiment confirms that the relation between blank duration and alternation probability changes in the predicted manner over the course of a period of stabilized perception. It is hard to see how this detailed and counterintuitive finding could be expected outside of the context of the present model account. Therefore, we see it as strong evidence that the present model captures the essential dynamics of the system responsible for perception of appearing ambiguous stimuli.

7.2.5 Model prediction: neural activity

As illustrated by the above experiment, theoretical understanding in this area has now developed to a point where even remarkably complex patterns in psychophysical data are interpretable in terms of the dynamics of the underlying system. On the other hand, knowledge of the actual neural events responsible for these dynamics remains scarce. Can we use the predictive power of present theoretical understanding to guide the search for neural correlates of perceptual memory?

Below we will provide a prediction of specific patterns of neural firing activity that may accompany memory of ambiguous figures, as well as describing the type of presentation paradigm and analysis that may help identifying them. Before we do that, it is useful to point out a number of apparently plausible assumptions on the nature of neural correlates of perceptual memory, which may require revision in the light of present theoretical understanding.

Perceptual memory of ambiguous stimuli is most naturally described as a dependence of

perception following a blank period on perception preceding that blank period. This indicates a form of information retention across the blank period. Therefore, a search for neural correlates might reasonably be aimed at correlating activity measured during an intervening blank period with perception immediately before and/or after the blank period. This strategy has been used in an fMRI study of perceptual memory (Sterzer and Rees, 2008), and preliminary neurophysiological work also indicates that informative neural data may be obtained in this way (Maier et al., 2002). It is noteworthy, however, that current theoretical knowledge suggests two reasons why this approach, in spite of its intuitive appeal, might not be the method of choice in the search for the neural concomitants of perceptual memory.

First, psychophysical characteristics of perceptual memory are consistent with the idea that memory is embodied in the adaptation states of neurons involved in percept competition (Figure 7.2). This idea of memory as an altered adaptation state has implications for the interpretability of measurements taken during a blank interval. Importantly, as we presently understand it, the perceptual consequence of an altered adaptation state depends critically on the specific way in which adaptation interacts with the competition process at stimulus appearance (Figure 7.2). How this perceptual consequence would be associated with effects that adaptation has during the preceding blank period, is not obvious. In terms of neural spiking activity during stimulus absence, prolonged adaptation can cause a reduction of baseline firing rate (e.g. Albrecht et al., 1984). It is therefore conceivable that neurons whose activity is associated with perceiving a given interpretation, have a lower baseline firing rate following prolonged perceptual stabilization of that interpretation. In terms of predictive power of upcoming perception, however, activity measured during a preceding blank period may not be as informative as might intuitively be suspected.

Second, whether neural activity is measured during a blank interval or otherwise, there appears to be a serious limitation to approaches that attempt to correlate measured activity with immediately preceding or subsequent perception. If an experiment involves repeated presentation of a stimulus, the adaptation levels that arguably mediate perceptual memory (Figure 7.2) are not determined solely by immediately preceding perception, but reflect a minutes-long perceptual history (Figures 7.3C and 7.4A). This implies that immediately preceding perception may not correlate strongly with potential activity patterns associated with perceptual memory. Similarly, a particular outcome of the competition process does not correspond uniquely to a specific combination of adaptation levels. Instead, there is a broad range of adaptation levels that all yield the same perceptual outcome (Figure 7.2D). Therefore, perception immediately following the moment a neural measurement was taken, may not be very informative regarding the state of the neural system either. These points are well illustrated by Figure 7.3C, right. Here, the state of the neural system at the moment of the stimulus appearance marked 'n' is almost identical to the state at the moment of the preceding appearance, yet the perceptual outcome on both appearances is opposite¹. Conversely, the system state traverses a broad range of values from appearances '1' up to 'n', yet the perceptual outcome is identical in all cases. In sum, we suggest that it is of limited value to use perception on any small number of presentations as the independent variable in analyses aimed at finding the neural concomitants of perceptual memory.

¹In the idealized space of Figure 7.3C two system states that lead to opposite perceptual outcomes can lie arbitrarily close together. In practice noise will ensure that the boundaries between regions of opposite perception are more fuzzy, which puts a limit on the possible degree of similarity between two neural states that reliably produce opposite perception.

Analysis of the present model suggests an alternative approach to finding neural concomitants of perceptual memory. This approach hinges on two ideas. First, rather than measuring activity during blank intervals, we suggest recording responses to the *appearance* of an ambiguous stimulus. Second, instead of basing data analysis on a correlation with immediately preceding or immediately following perception, we propose a method to effectively incorporate the entire perceptual history during a session.

This approach is based on model analysis of the initial upstroke in activity in response to the onset of input. The initial part of the curve in Figure 7.2C, right, is an example of such an upstroke. As illustrated in Figure 7.6A, the shape of the upstroke systematically changes as a function of the adaptation state. Figure 7.6A again shows the familiar landscape that is a function of the adaptation states of the two competing neural representations. This time, the insets in the landscape show onset responses that are predicted in various regions of the landscape. Activity $S[H]$ during the first tens of milliseconds after stimulus onset is plotted in a dark shade for the representation of the darker, leftward tilted grating, and in a light shade for the representation of the lighter, rightward tilted grating.

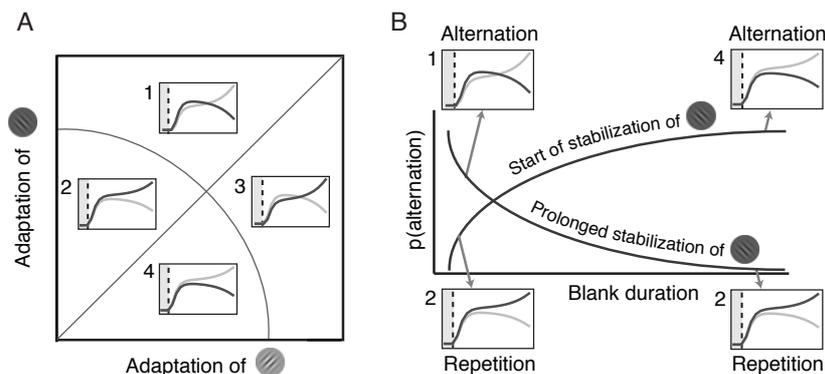


Figure 7.6: Model prediction: neural activity. We suggest that the signature of perceptual memory lies in the neural response to the onset of an ambiguous stimulus, measured as a function of slowly accumulated adaptation variables. A. The shape of the onset response to an ambiguous stimulus depends systematically on the long-term system state. Insets show model activity of the two representations during the first tens of milliseconds after stimulus appearance. A qualitative difference between the onsets at low adaptation states (insets 2 and 4) and those at high adaptation states (insets 1 and 3) is that the latter show a crossover between the responses of the two representations during their upstrokes. B. Perceptual alternations and repetitions in our psychophysical experiment of Figure 7.5 can be identified with specific locations within the landscape of panel A. This panel shows the resulting mapping between our psychophysical experiment and predicted onset responses.

The quantitative shape of the onsets changes throughout the landscape. Figure 7.6A illustrates only the most conspicuous, qualitative, prediction, which concerns the top right region of this landscape (insets marked '1' and '3'). Here, the trace associated with the interpretation that will become dominant on this presentation, initially runs *below* the trace associated with the alternative interpretation. Consequently, in this part of the landscape we predict a crossover of both curves during the first tens of milliseconds after stimulus appearance.

By itself, this theoretical prediction does not directly suggest an approach to measuring neural concomitants of perceptual memory. One might record onset responses to appearing ambiguous stimuli, but to put our prediction to the test, one also needs a method to estimate

the adaptation states of the putative competing representations at the moment the recording was taken. In other words, one has to be able to separate measured responses according to position within the space of Figure 7.6A.

Fortunately, the psychophysical experiment described in Figures 7.4 and 7.5 provides a method to perform such a separation. As shown in Figure 7.4, our theoretical prediction in that experiment was based on anticipated behavior in various regions in the landscape of Figure 7.6A. The fact that our prediction matched the observed perceptual outcome, suggests that our estimation of the evolving position within this landscape throughout experimental sessions was reasonably accurate.

Figure 7.6B illustrates how the psychophysical experiment of Figures 7.4 and 7.5 relates to the present prediction on neural activity. Like Figure 7.4D, it schematically shows how the probability of a perceptual alternation varies with blank duration at various moments throughout a sequence of stabilized perception. Importantly, perceptual alternations and perceptual repetitions at various stages of stabilization are associated with well defined locations within the adaptation landscape (Figure 7.4). If they were not, it is unclear how our prediction of psychophysical outcome, which was based on anticipated behavior at those locations, could match the empirical observations. For instance, if the leftward tilted gratings have just become perceptually stabilized, an alternation to the rightward tilted gratings following a long blank interval typically corresponds to a position in the bottom section of the landscape near inset '4' (Figure 7.4C). More generally, the predicted patterns of onset responses can be associated with our psychophysical data in the manner indicated in Figure 7.6B. In other words, the same analysis that allowed us to adequately separate our data into groups like 'at the start of stabilization' and 'after prolonged stabilization', can be used to sort onset responses measured during perceptual stabilization. The methods described in the present work may thus allow our prediction on a neural concomitant of perceptual memory to be tested.

Note that the inset marked '3' in Figure 7.6A is not featured in Figure 7.6B. This is because it is associated with prolonged perceptual stabilization of the rightward tilted gratings, a situation that is not depicted in Figure 7.6B.

7.3 Discussion

We have suggested what type of data pattern one might look for in search of neural concomitants of perceptual memory. Another question is what brain areas to target for such neural measurements. Our prediction centers on activity levels of the putative neural representations of competing interpretations of an ambiguous stimulus. Although for binocular rivalry these representations have in the past been identified with monocular neurons in the primary visual cortex (Blake, 1989), more recent observations do not support this notion (Leopold and Logothetis, 1996). Rather, if interactions such as formalized in our model indeed underly the behavior that is observed in psychophysical experiments, they appear to be a property of distributed neural networks; not of any local neural population (Blake and Logothetis, 2002; Tong et al., 2006).

One key requirement for a cell to potentially show the memory-related activity we predict, is for its activity to be associated specifically with perception of one interpretation of an ambiguous stimulus, and not the other. For various kinds of ambiguous stimuli such cells have been identified, and their anatomical locations vary with the particular stimulus at hand. For instance, for binocular rivalry such cells are present in cortical areas V1, V2, V4, V5/MT (mid-

dle temporal area), and IT (inferior temporal area) of macaque monkeys, in increasing order of abundance (Leopold and Logothetis, 1996; Sheinberg and Logothetis, 1997; Logothetis, 1998). In the case of another commonly studied type of rivalry called structure from motion rivalry (Figure 7.1B), V5/MT of macaque monkeys harbors such cells (Bradley et al., 1998), and MST (medial superior temporal area) is also a candidate area (Saito et al., 1986).

Regarding the neural prediction of Figure 7.6, one might wish to gain an intuition as to why we predict a crossover in onset responses in certain regions of the adaptation landscape (insets 1 and 3). For this one has to recall the two effects of adaptation ΣA on response $S[H]$ that are implemented in equations 1 and 2. First, adaptation allows a representation to respond more readily to the onset of input. As a result, the onset response of a more adapted population starts earlier. Second, adaptation causes a reduced response gain of a representation. Accordingly, the slope of the onset response of a more adapted population is shallower. In the bottom regions of the adaptation landscape the reduced response latency allows the more adapted representation to win the competition at stimulus onset (insets '2' and '4'). In the top regions, however, the more adapted representation loses the competition. This is not because the more adapted representation does not have a reduced response latency in these regions, but because the reduced response latency is outweighed by the reduced response gain. As a result, in this area of the landscape the onset response of the more adapted representation starts quickly, but later it is overtaken by the response of the other representation (insets '1' and '3').

A process that is conceptually analogous to perceptual selection of an interpretation of an appearing ambiguous stimulus, is attentional selection of an item among distracting items. Attentional selection has been proposed to depend on a winner-take-all competition between elements in a scene (Fecteau and Munoz, 2003). Similarities between the memory phenomena observed in ambiguous figure perception and memory in attentional selection, have previously prompted the suggestion that there may be parallels between their neural bases (Brascamp et al., 2008). In attentional selection research, computational modeling and neurophysiological measurements have recently started to converge toward an integrated framework (Fecteau and Munoz, 2003). Could knowledge from attentional selection research aid our understanding of perceptual memory in the present paradigm? An intriguing, although by no means conclusive, observation in this context is that neural traces measured in macaque FEF (frontal eye fields) in association with attentional selection are quite similar to the traces we predict in association with perceptual selection (Schall and Thompson, 1999). That is, when comparing cells with receptive fields centered on a search target and cells centered on a distractor, the onset of a search array typically triggers a simultaneous rise in activity of both cell types, until after a brief plateau the activity of cells centered on the target rises further while the activity of distractor cells returns to baseline (cf. Figure 7.6). The shape of these traces, moreover, is influenced by factors such as trial history (Bichot and Schall, 2002). These observations raise the possibility that similar competitive interactions, though not necessarily with an overlapping anatomical basis, underly both attentional and perceptual selection.

Previous work has shown that the effect of a blank interval inserted between two presentations of an ambiguous stimulus depends on the duration of the blank interval. As blank durations are increased from hundreds of milliseconds to multiple seconds, the probability of an alternation at the moment of stimulus reappearance falls from above chance level to below chance level (Orbach et al., 1966; Noest et al., 2007; Klink et al., 2008). In other words, brief blank intervals prompt alternations, whereas long blank intervals cause repetitions. How can this be reconciled with the data in Figure 7.5, which show that the relation between blank duration and

alternation probability depends on long-term perceptual history? The key difference between previous experiments and the present one is that here we varied the blank duration within sessions, whereas previous work used blocked conditions. During a session of consistently long blank intervals, adaptation is sufficiently low to prompt percept repetition on the vast majority of presentations. Only incidentally does adaptation accumulate high enough to cause an alternation to the opposite percept (Figure 7.3C, right; see also Brascamp et al. (in prep)). During sessions of consistently short blank intervals, on the other hand, adaptation has little time to decay during stimulus absence. As a result, adaptation is relatively high throughout a session, resulting in many perceptual alternations. The main difference between our experiment and previous ones, therefore, is that we could dissociate the effect of a single blank duration from that of accumulating adaptation during previous presentations, whereas previous experiments, using a blocked design, measured both effects simultaneously.

7.4 Conclusion

There is presently a relative paucity of data on neural activity associated with perceptual memory. In contrast, the phenomenon seems fairly well understood from a theoretical perspective. Here we summarize the workings of a computational model that successfully accounts for known perceptual memory data, and we report an experiment that tightens the link between model components and psychophysically observable variables. We then explore the extent to which present theoretical understanding can help guide experiments into the neural origin of perceptual memory, and come to several recommendations. First, we suggest that finding neural concomitants of perceptual memory requires measuring in the period immediately after the appearance of an ambiguous stimulus, rather than during memory retention. Second, we propose that data analysis should be based on minute-scale percept history, rather than on isolated perceptual events. Finally, we provide an explicit prediction of neural activity associated with perceptual memory, and propose an experimental approach that could bring such activity to light.

7.5 Methods

7.5.1 Stimuli and procedure

Four subjects (three naive) took part in our psychophysical experiment. They viewed square wave gratings (0.67 Michelson contrast; 3.3 periods per deg), within a square aperture (side 1.4 deg). Stimuli were surrounded by a white fixation square (62 cd/m^2 , side 1.5 deg). Background luminance was 31 cd/m^2 , as was the mean stimulus luminance. Stimuli were viewed through a mirror stereoscope at a visual distance of 47 cm. The intermittent presentation cycle consisted of 0.6 s presentations, alternated with blank intervals of random duration. Each blank duration within a session was drawn from a set of six durations ranging from 0.22 s to 4.2 s, equally spaced on a log axis. In order to adequately study the effect of blank duration we wished to sample this entire range. However, presenting many very short or very long durations would interfere with the evolution of the slow alternation cycle (Brascamp et al., in prep). We therefore adjusted the relative probabilities of each of the six durations being drawn, so that intermediate durations occurred more often than short or long ones. The two central durations each had a

probability of 23%, the outermost pair each had a probability of 10%, and the remaining two each had a probability of 17%. Subjects were instructed to report perception via key presses on each presentation during 30 minute sessions.

7.5.2 Data analysis

Our model predicts that the relation between blank duration and alternation probability is affected by the accumulated amount of slow adaptation 7.4. Also, adaptation of a given neural representation is supposed to accumulate when the corresponding percept is dominant. As a measure of a representation's slow adaptation state we essentially calculated a low-pass filtered version of the corresponding dominance periods during a session. That is, a representation's slow adaptation state A_S was calculated by numerically integrating

$$\tau \partial_t A_S = -A_S + K \quad (7.5)$$

over the time series of perceptual reports. On every integration step the value of K was set to 1 if the corresponding percept was dominant. It was set to 0 both if the corresponding percept was suppressed and during blank intervals. The time constant τ was set to 50 s. Figure 7.7 shows an example of a perceptual time course during a session (background shade), and the corresponding values of slow adaptation as calculated using equation 7.5 (curves).

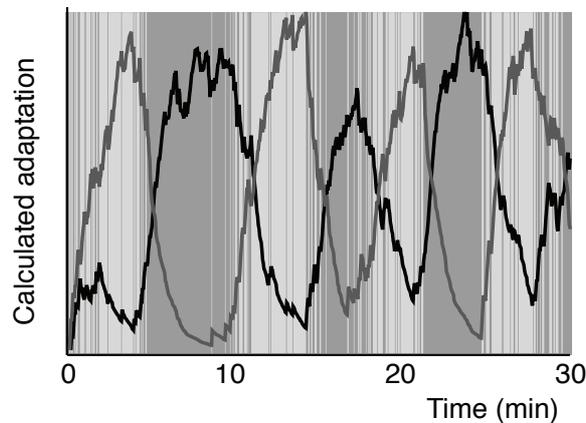


Figure 7.7: Data analysis. The background shade indicates perception during a typical experimental session. The shaded curves indicate slow adaptation of the representation of the lighter, rightward tilted grating (light shade) and of the darker, leftward tilted grating (dark shade), calculated on the basis of equation 7.5. The difference between these two curves was used as a regressor to categorize our experimental data into groups corresponding to various stages within a stabilization sequence.

For every time point within a session we calculated the difference between the two calculated slow adaptation values. The first 75 presentations of each session were considered a startup period in terms of slow adaptation, and were removed from further analysis. We then separated all remaining stimulus presentations within a session into four groups according to the four 25% percentiles of the associated difference values. This produced the four categories in Figure 7.5C, ranging from ‘at the start of stabilization’ to ‘after prolonged stabilization’. The

category labeled ‘at the start of stabilization’ is associated with an adaptation value of the dominant representation that is much lower than that of the non-dominant representation; the one labeled ‘after prolonged stabilization’ corresponds to relatively high adaptation value of the dominant representation.

Appendix

We wished to ensure that our model analysis of the relation between blank duration and alternation probability (Figure 7.4) agreed with actual model behavior. We therefore performed a simulation using the same model as used in Brascamp et al. (in prep).

Simulation

For activity H_i of representation i the model has

$$\partial t H_i = X_i - (1 + \Sigma A_i) H_i + \max[0, \beta \Sigma A_i - \beta' \Sigma A_j] - \gamma S[H_i] \quad (\text{A7.1})$$

This model incarnation is slightly more elaborate than the one discussed in the main text, which is necessary to account for long term perceptual alternations during intermittent viewing (Brascamp et al., in prep). Also, whereas in the main text we did not specify the exact number of timescales comprising adaptation ΣA , for this simulation we fixed the number at three. Thus, adaptation ΣA_i of representation i becomes

$$\tau_1 \partial t A_{i,1} = -A_{i,1} + \alpha_1 S[H_i] \quad (\text{A7.2})$$

$$\tau_2 \partial t A_{i,2} = -A_{i,2} + \alpha_2 S[H_i] \quad (\text{A7.3})$$

$$\tau_3 \partial t A_{i,3} = -A_{i,3} + \alpha_3 S[H_i] \quad (\text{A7.4})$$

where

$$\Sigma A_i = A_{i,1} + A_{i,2} + A_{i,3} \quad (\text{A7.5})$$

For activity and adaptation of interpretation j the equations are the same as equations A7.1 through A7.5, but with subscripts i and j exchanged.

Parameter settings for our simulation were: $\gamma=3.3$, $\beta=0.21$, $\beta'=0.04$, $\alpha_1=3$, $\alpha_2=0.5$, $\alpha_3=0.5$, $\tau_1=100$, $\tau_2=2000$, $\tau_3=4000$ and $X=1$. Stimulus presentation duration was 50 (arbitrary units). Blank durations were drawn randomly from six values spaced equally along a log axis between 87 and 432 (arbitrary units). As in our psychophysical experiment, the two central durations were drawn most often (each constituted 25% of all blank intervals), the longest and shortest duration were drawn least often (8% each), and the intermediate durations fell in between at 17%. The simulation was performed using a fourth order Runge-Kutta method with step size 0.1. On every time step we applied Poisson noise on $S[H]$. This was implemented as

$$S[H]_{new} = \frac{\text{Poisson}(S[H]_{old}C)}{C} \quad (\text{A7.6})$$

$S[H]_{old}$ and $S[H]_{new}$ are the values before and after the addition of noise, respectively. $Poisson(\mu)$ is a random value from the Poisson distribution with mean μ , and C is a constant that we set to 2000.

Analysis

After performing the simulation we calculated a measure of slow adaptation based on simulated perception, analogous to our procedure in analyzing the empirical data. That is, we low-pass filtered our simulated perceptual data according to equation 7.5, using a time constant τ of 5000 (arbitrary units). Subsequently, we separated all stimulus appearances according to the associated values of calculated slow adaptation, creating four categories based on the four 25% percentiles. The simulation results are depicted in Figure A7.1. They show that the model behaves as indicated by the analysis of Figure 7.4.

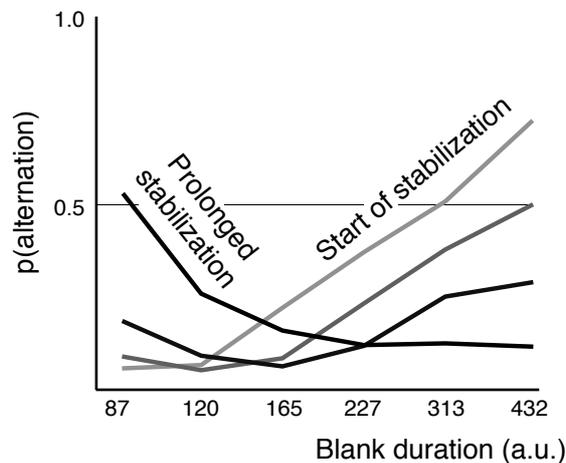


Figure A7.1: Simulation results. Our model produces the same relation between blank duration, long-term stabilization state and alternation probability as observed in our psychophysical experiment (cf. Figure 7.5). Durations are in arbitrary units (a.u.).

Summary and discussion

Summary per chapter

Chapter 1. Distributions of alternation rates in various forms of bistable perception

While watching an ambiguous stimulus, consecutive perceptual alternations are separated by periods whose durations follow a unimodal, rightward skewed distribution. We investigated the similarity of this distribution to a number of analytical distributions. A noteworthy outcome of this comparison is that the gamma distribution, which is commonly used to describe ambiguity data, did not fit our data better than several other distributions did. Also, a detailed analysis demonstrated closely resembling distributions associated with very different kinds of ambiguous stimuli. This suggests that, although these stimuli undoubtedly activate at least partially distinct neural populations, the mechanisms underlying perceptual alternations are similar in all cases.

Chapter 2. The time course of binocular rivalry implies a fundamental role of noise

This chapter fits in a tradition of studying the dependence of the time course of alternation cycle on stimulus parameters. We analyzed the temporal characteristics of the binocular rivalry alternation cycle as a function of the contrasts of the two competing images. In terms of phenomenology, the chapter provides a fairly comprehensive description of the relation between contrast and both dominance durations and transition durations (which both had partially been described before), and of the relation between contrast and the occurrence of 'return transitions'. In terms of the mechanisms behind perceptual alternations, this chapter is part of a sequence of recent papers that demonstrate the importance of random fluctuations in brain activity in triggering perceptual alternations (Kim et al., 2005; Moreno Bote et al., 2007).

Chapter 3. Sensory memory for ambiguous vision

This chapter summarizes the literature that describes how perception of an appearing ambiguous stimulus depends on prior perception of that same stimulus (perceptual memory). One main conclusion is that the psychophysical characteristics of this kind of perceptual memory are consistent with a neural substrate within sensory areas. Also, the chapter points out that empirical evidence is consistent with computational models that attribute memory to a persistent facilitatory trace that biases competition between neural populations that represent competing interpretations.

Chapter 4. Flash suppression and flash facilitation in binocular rivalry

We investigated how perception of an appearing binocular rivalry stimulus depends on prior exposure to one of the two competing images in isolation. The main finding in this chapter is that prior exposure to one of the images can allow that image to gain initial dominance in rivalry (facilitation), while it can also have the opposite effect of causing dominance of the competing image (suppression). Moreover, there is a continuous progression from facilitation to suppression as the ‘energy’ of the prior stimulus is increased, either by raising its contrast or by prolonging its duration. These findings appear related to observations regarding the effects of ambiguous, rather than unambiguous, prior stimuli. We conclude by proposing that the distinction between traces left by ambiguous stimuli and traces left by unambiguous stimuli is gradual rather than categorical. This supports the idea that ambiguous stimuli do not receive exceptional treatment by the brain, but rather draw on the same mechanisms as unambiguous stimuli.

Chapter 5. Multi-timescale perceptual history resolves visual ambiguity

Here we investigated how perception of an appearing ambiguous stimulus depends on prior perception of that same stimulus. The interesting finding here is that perception is determined both by transient traces left by recent perceptual events, and by more persistent traces left by remote perceptual events. More specifically, when relatively persistent traces are temporarily masked by more transient traces, they are not erased. Instead, the more persistent effects resurface as soon as the transient traces fade. This existence of parallel traces on separate timescales indicates that perceptual memory reflects not one, but multiple independent neural processes. One aspect that distinguishes this chapter from most previous work on this subject is that it treats perceptual memory essentially as a type of sensory aftereffect, rather than as persistence of a percept. In other words: the stance taken here is that memory involves persistence of altered sensitivity, rather than persistence of spiking activity.

Chapter 6. Intermittent ambiguous stimuli cause periodic perceptual alternations

This chapter demonstrates a slow and regular perceptual alternation cycle during intermittent viewing of an ambiguous stimulus. This cycle consists of individual periods of up to ten minutes during which the same percept dominates on virtually every presentation. The frequency of alternations increases as the blank duration between presentations is reduced, underscoring the stimulus driven nature of this cycle. The alternation cycle observed here differs from the one observed during continuous presentation in several respects, indicating the two rely on distinct neural events. The same computational model that was used to accommodate the findings of Chapter 5, also accommodates those of this chapter. We now have a plausible account of the system dynamics responsible for all perceptual events during a session of intermittent viewing.

Chapter 7. Perceptual memory: toward neural measurement

There is presently an imbalance between the degree to which computational analysis allows us to understand the dynamics of perceptual memory (relatively well), and what we actually know about neural activity associated with this phenomenon (relatively little). In this chapter

we summarize the workings of the computational model of perceptual memory that was used in Chapters 4, 5 and 6, and we describe a psychophysical experiment that tests a prediction by the model. We then use knowledge from both the model and this psychophysical experiment to from predictions that can guide the search for neural concomitants of perceptual memory. We identify several pitfalls that may interfere with this search, and describe how to circumvent them using the best of present knowledge from psychophysics and computational modeling.

General discussion

Chapters 1 and 2 of this thesis stand in relative isolation, and they may not require a discussion in a broader context than provided in the chapters themselves and in the summaries above. The remainder of the thesis, however, deals with essentially a single subject from a number of different angles. This is the question of how temporal context affects the processing of ambiguous stimuli. Most main implications and conclusions have been mentioned at various points throughout the text, but it may be useful to provide a comprehensive discussion here. By focusing on the findings and implications of this thesis itself, this discussion will have only limited overlap with Chapter 3, which reviewed the overall literature on this subject, and at a point in time when some of the work in this thesis had not yet been published.

The influence of history on the perception of an appearing ambiguous image

PHENOMENOLOGY When an ambiguous image appears, its perception often depends on previous perception of that same or a similar image. This can be either a facilitatory relation, when a previously encountered interpretation has a greater chance of being perceived, or a suppressive relation, when a previously encountered interpretation is more likely to become suppressed. The facilitatory effect has been termed ‘perceptual memory’. By extension, in Chapter 3 we chose to use this term for any history effect observed in ambiguous perception. Results in this thesis indicate that conditions that prompt facilitation and those that cause suppression can be very similar. If a prior stimulus is itself also ambiguous, the direction of its effect depends on stimulus timing (Chapter 7), as well as long-term history (Chapters 6 and 7). If a prior stimulus is unambiguous, the direction of its effect depends on stimulus timing and parameters such as its contrast (Chapter 4). Whether the effect of an ambiguous prior stimulus also depends on factors such as contrast, is presently unknown.

The trace of a prior stimulus gradually develops during presentation of the prior stimulus. Initially, this causes a steady growth in the strength of its facilitatory effect. As the trace evolves further, however, its effect changes from facilitatory to suppressive. This is true, both when the prior stimulus itself is ambiguous (Chapter 5), and when it is unambiguous (Chapter 4).

Ambiguous stimuli leave a trace whose evolution displays both fairly fast (seconds) and much slower (minutes) components (Chapters 5, 6 and 7). These components are mutually independent, so that interference with the trace on a short timescale does not affect the more persistent components. This indicates that history effects in ambiguous vision may reflect multiple distinct neural mechanisms, which jointly affect perception. Whether traces of unambiguous stimuli also display this type of multi-timescale behavior has not been tested, but it seems likely.

MECHANISM: WHAT CONSTITUTES A PERCEPTUAL MEMORY? Prior stimulation arguably exerts its effect on subsequent perception via sensitivity changes in neurons that respond both to the prior stimulus and to the subsequent ambiguous stimulus. During presentation of the prior stimulus the responding neurons get adapted. If the prior stimulus differentially activates neurons associated with either interpretation of the subsequent ambiguous image, this can cause an imbalance in their sensitivities, and thus affect perception. Storage in the form of sensitivity can be contrasted with storage in the form of (spiking) activity. A similar distinction has recently been pointed out with regard to working memory (Mongillo et al., 2008). Several observations within and outside of this thesis support the idea that perceptual traces in the present paradigm are composed of imbalances in sensitivity.

- The effect of a prior stimulus gradually accumulates during its presentation, and gradually decays afterward. Of particular interest are the facts that (i) after a sequence of alternating perception, the net effect of the resulting trace depends on the overall dominance ratio between the two percepts (Chapter 5); (ii) traces accumulate on several timescales simultaneously (Chapters 5 and 6); (iii) a sufficiently accumulated trace of dominance of a given percept still exerts its effect after multiple dominance periods of the alternative percept have intervened (Chapter 5) and (iv) the magnitude of a trace associated with a given percept can be fairly well predicted by ‘leaky integration’ over dominance periods of that percept (Chapter 7). These characteristics are more consistent with an adaptation-like mechanism of ongoing sensitivity shifts in response to stimulation than with, for instance, a mechanism centered on persistent activity after a stimulus has been removed.
- As partly discussed in Chapter 3, prior stimuli which are not expected to cause differential adaptation in neurons associated with the two interpretations of a particular ambiguous stimulus, usually do not affect subsequent perception of that stimulus. For instance, perceptual competition between opposite directions of rotation in depth is little affected by prior presentation of a static stimulus or of motion in the frontoparallel plane (Maier et al., 2003). This type of finding is consistent with the idea that differential adaptation of the neurons involved in competition is required for a prior stimulus to exert an effect.
- If, as proposed in Chapter 3, these traces originate in sensory areas, persistent activity is not the most obvious option for their neural basis. Persistent activity after stimulus removal is observed in sensory neurons (e.g. Pasternak and Greenlee, 2005), but it is not typical. Sensitivity changes in response to stimulation, on the other hand, are ubiquitous.
- Existing computational models that can reproduce empirical data, use sensitivity changes as the basis of history effects (Chapters 4 through 7). Although it is possible for a model based on incorrect principles to behave similar to its biological counterpart, parallels in behavior are an indication of parallels in internal workings.

If temporal context does interact with perception of ambiguous stimuli in the way we propose, the question what constitutes ‘a perceptual memory’ can essentially be answered in either of two ways. One way of looking at it is that the imbalance in adaptation state is the memory. This is a concise answer, and accurate in the sense that it is this imbalance that persists over time. An alternative answer is that a perceptual memory consists in the interaction between adaptation state and the perceptual competition process triggered by an ambiguous stimulus. One benefit of this latter definition is that it is more naturally consistent with the finding that very similar

adaptation states can have opposite perceptual outcomes, which is somewhat unsatisfactory if we define the adaptation state itself as the memory (Chapters 6 and 7). Moreover, by explicitly pointing to the ‘readout’ by means of an ambiguous stimulus, the latter definition makes it clearer that otherwise analogous forms of memory may have a somewhat different outward appearance if they do not involve an ambiguous test stimulus. This opens up ways to finding related phenomena that do not involve perceptual ambiguity (see below).

RELEVANCE OUTSIDE OF THE LABORATORY Making perceptual sense out of retinal stimulation is an active process. The observed effects of prior stimulation and perception may play a role in streamlining this process. In the broadest terms, the sensitivity changes that underly history effects in our paradigm may allow the visual system to exploit statistical regularities in input sequences, and thereby increase the efficiency with which it can organize new input when it appears.

In Chapters 5 and 6 we suggest a more specific implementation of this idea. We propose that the observed facilitatory effects may reflect an elevated readiness for the visual system to select items that were relevant in the recent past, as these are likely to be relevant again in the near future. We propose that for an automatic mechanism – that is, one that does not involve complex inference – to be helpful in this regard, it may be based on the simple principle that relevant items are those that the visual system selected previously. In a sense the visual system would then rely on more complex selection processes that it employed on previous encounters, and use this type of priming to avoid having to invoke those processes repeatedly. Similar reasoning has been followed in the interpretation of certain forms of attention priming (Maljkovic and Nakayama, 2000; Kristjánsson, 2006). The idea is that the visual system achieves this in a passive and automatized manner by means of sensitivity changes in networks that are functionally located beyond some selection bottleneck, so that only those items get ‘primed’ that made it through the selection initially. In our paradigm, repeated dominance of a given interpretation would thus be caused by sensitivity changes in networks that get activated specifically during perceptual dominance of that interpretation.

One interpretation of this idea is that the history effects we probe have a function in perceptual organization of visual input, which inherently contains some degree of ambiguity. Once the visual system has settled on a particular interpretation of sensory information, these effects could then facilitate subsequent organization when similar input appears. In a broader context, we may draw parallels with other forms of visual selection besides the selection of a perceptual interpretation. In Chapters 5 and 6 we discuss selection by means of attention or eye movements. Those types of selection have been proposed to rely on competitive processes remarkably similar to the ones we envision (Desimone and Duncan, 1995). Also, they display history effects that are comparable to those observed in the present paradigm (Maljkovic and Nakayama, 2000; Kristjánsson, 2006). It is presently not clear to what extent attentional selection and perceptual selection have an overlapping neural basis (e.g. Koch and Tsuchiya, 2007; Mitchell et al., 2004; Leopold and Logothetis, 1999; Kanai et al., 2006). Regardless, the conceptual similarity between these various functions raises the possibility that the brain may employ analogous mechanisms for each of them even if the anatomical substrate is different.

One factor that masks the parallels between history effects in attentional selection and those in our paradigm, is that those in attentional selection are manifested as changes in response time (Maljkovic and Nakayama, 1994, 2000; McPeck et al., 1999), whereas we measure perceptual dominance. Based on present understanding of memory in our paradigm, one may argue

that this does not reflect a difference between the nature of memory in both paradigms, but a difference in the extent to which an observer's intention can overrule a biasing memory trace. Specifically, in the case of perceptual ambiguity, a memory trace can often determine selection for dominance because it is the main biasing factor. In contrast, in a typical attentional selection paradigm subjects are instructed to attend to a particular item, and are well able to do so. A history-driven bias, therefore, is usually not able to overrule this top-down force and thereby determine selection, but it may accelerate or decelerate the selection process depending on its direction relative to the top-down bias.

The above conceptualization centers on a possible functional role of the facilitatory effect of past dominance. It does not address the suppressive effect. The observed suppression is more pronounced during prolonged viewing than at stimulus appearance. Also, it is stronger following exposure to unambiguous stimuli than following exposure to ambiguous stimuli. The functional interpretation presented above, in contrast, centers on selection processes engaged by appearing stimuli, and on history effects that are due to prior selection (of a perceptual interpretation). For these reasons, we argue in Chapter 6 that the primary functional role of suppression might be unrelated to the kind of selection processes discussed above. Suppressive effects are more commonly observed in perception than facilitatory ones. Although they have been studied extensively, there is presently no final answer as to their functional role (e.g. Schwartz et al., 2007). One of the more widespread notions is that they have to do with 'gain control', and enhance the system's ability to code efficiently over a broad range of inputs (Ohzawa et al., 1985; Abbott et al., 1997).

INDICATIONS FOR FUTURE RESEARCH There are many ideas and hypotheses about the role of temporal context in ambiguous vision that have yet to be explored. Here I summarize some open questions that I think are particularly interesting.

A direct way to investigate potential relations between history effects in ambiguous perception and other memory phenomena, is to experimentally test whether a trace built up in one paradigm has a measurable effect in the other. For instance, it is worth investigating whether traces left by prior perceptual dominance would affect subsequent detection or discrimination thresholds. This kind of transfer between paradigms would help get a clearer view of the neural origin of the perceptual memory phenomenon.

Arguably of specific interest is the question whether any transfer can be observed between perceptual selection paradigms and attentional selection paradigms. Would prior attentional selection facilitate subsequent perceptual selection, and vice versa? In a sense, one half of this question has already been answered by studies where attention was drawn to an image, and subsequently that image engaged in binocular rivalry with an image in the other eye (Mitchell et al., 2004; Chong and Blake, 2005; Kamphuisen et al., 2007). Those studies did demonstrate a priming effect of prior attentional selection on subsequent dominance, and it would be interesting to follow up on those results. The converse experiment, where one would test if prior dominance allows an item to gain attentional selection more swiftly, has yet to be performed.

If it is true that the memory traces we observe allow the visual system to exploit statistical regularities in input sequences, then the time course of buildup and decay should in some way match the temporal statistics of input. For instance, the fact that traces accumulate over minutes (Chapters 5 through 7), indicates that our normal visual environment contains relevant correlations on those slow timescales. It would be interesting to see how flexible the 'temporal weighting kernel' is, with which the visual system weighs past events. That is, we might

be able to perform experiments akin to perceptual learning studies, in which we determine the time course of buildup and decay of perceptual memory, both before and after a 'learning' experiment. In this learning experiment subjects would perform a task on visual input that contains relevant information on a well defined range of timescales. This might alter the visual system's weighting of information on those timescales. An effect on the subsequent time course of perceptual memory could help us understand how the visual system learns to make effective use of temporal context in its processing of input.

In Chapter 7 we stressed the need for measurements into the neural basis of the present phenomena. The psychophysical and theoretical basis laid in this thesis, as well as in previous work, should help guide an effective search for this neural basis.

Samenvatting

Hoofdstuk 1. Distributies van alternatiefrequenties in verschillende soorten bistabiele perceptie

Tijdens het observeren van een ambigue stimulus volgt de duur van afzonderlijke dominantieperiodes een unimodale distributie waarvan de rechter staart dikker is dan de linker. Wij vergeleken de vorm van deze distributie met die van verschillende theoretische distributies. Een opmerkelijke uitkomst is dat de zogenaamde gammadistributie, die algemeen gebruikt wordt om dit type data te beschrijven, niet beter op onze data paste dan een aantal andere theoretische distributies. Verder vonden we dat sterk verschillende ambigue stimuli allemaal distributies opleverden die tot in detail met elkaar overeen kwamen. Dit wijst erop dat bij al deze stimuli, hoewel ze ongetwijfeld (gedeeltelijk) verschillende hersengebieden aanspreken, toch een vergelijkbaar mechanisme verantwoordelijk is voor de perceptuele alternaties.

Hoofdstuk 2. Het tijdsverloop van binoculaire rivaliteit wijst op een cruciale rol van interne ruis

Dit hoofdstuk past in een traditie van onderzoek naar de invloed van stimuluseigenschappen op het tijdsverloop van de cyclus van perceptuele alternaties. Wij onderzochten hoe dit tijdsverloop bij binoculaire rivaliteit afhangt van het contrast van de twee rivaliserende afbeeldingen. Het hoofdstuk geeft een vrij complete beschrijving van de relatie tussen contrast en zowel dominantieduren als overgangsduren (beide al gedeeltelijk eerder beschreven), en ook van de relatie tussen contrast en het optreden van zogenaamde 'afgebroken overgangen'. Wat betreft de implicaties in termen van het neurale mechanisme achter binoculaire rivaliteit, past dit hoofdstuk goed in een recente reeks artikelen die laten zien dat willekeurige fluctuaties in hersenactiviteit een rol spelen in het veroorzaken van perceptuele alternaties (Kim et al., 2005; Moreno Bote et al., 2007).

Hoofdstuk 3. Sensorisch geheugen in ambigue perceptie

Dit hoofdstuk geeft een samenvatting van de literatuur die handelt over de vraag hoe perceptie van een verschijnende ambigue stimulus afhangt van voorafgaande blootstelling aan diezelfde stimulus (perceptueel geheugen). Een belangrijke conclusie is dat dit soort geheugen eigenschappen heeft die wijzen op een oorsprong binnen het sensorisch deel van het brein. Ook wordt er gewezen op de overeenstemming tussen empirische gegevens en wiskundige modellen die het geheugen toeschrijven aan een beklijvende verandering in de toestand van de neurale populaties die de alternatieve stimulusinterpretaties representeren.

Hoofdstuk 4. 'Flash suppression' en 'flash facilitation' in binoculaire rivaliteit

Hier onderzochten we hoe perceptie van een verschijnende binoculaire-rivaliteitsstimulus afhangt van eerdere blootstelling aan een van de twee afbeeldingen die samen de rivaliteitsstimulus vormen. De belangrijkste vinding is dat eerdere blootstelling aan een van de afbeeldingen soms zorgt dat deze afbeelding perceptueel dominant wordt tijdens rivaliteit (facilitatie), maar soms juist dat zij perceptueel wordt onderdrukt (onderdrukking). Bovendien vinden we dat facilitatie naadloos overgaat in onderdrukking als de 'energie' van de vooraf getoonde stimulus wordt verhoogd, hetzij door het contrast te verhogen, hetzij door de presentatieduur te verlengen. Het lijkt erop dat deze vindingen verband houden met andere vindingen met betrekking tot effecten van het vooraf tonen van een ambigue (in plaats van niet-ambigue) stimulus. Wij opperen dat sporen nagelaten door ambigue stimuli en sporen van niet-ambigue stimuli niet wezenlijk van elkaar verschillen. Dit suggereert dat ambigue stimuli geen aparte behandeling krijgen in het brein, maar dezelfde mechanismen aanspreken als niet-ambigue stimuli.

Hoofdstuk 5. Perceptuele geschiedenis op verschillende tijdschalen bepaalt de uitkomst van visuele competitie

Hier onderzochten we hoe de perceptie van een verschijnende ambigue stimulus afhangt van eerdere blootstelling aan diezelfde stimulus. De interessantste vinding in dit hoofdstuk is dat perceptie afhangt van zowel kortdurende sporen, nagelaten door recente visuele stimulatie, als beklivende sporen van stimulatie verder in het verleden. Het is zelfs zo dat meer beklivende sporen, wanneer ze tijdelijk hun werking verliezen doordat ze gemaskeerd worden door meer kortdurende sporen, niet worden uitgewist. Integendeel: hun werking is weer meetbaar zodra de kortdurende sporen zijn uitgedoofd. Dit wijst op het bestaan van verschillende onafhankelijke perceptuele sporen op verschillende tijdschalen, wat suggereert dat er een aantal verschillende neurale processen ten grondslag liggen aan perceptueel geheugen. Een verschil tussen dit hoofdstuk en eerder werk over dit onderwerp, is dat we perceptueel geheugen in wezen behandelen als een sensorisch naeffect, in plaats van als een beklivende perceptuele toestand. Met andere woorden: perceptueel geheugen komt voort uit blijvend veranderde gevoeligheid van neuronen; niet uit beklivende vuuractiviteit van neuronen.

Hoofdstuk 6. Onderbroken ambigue stimuli veroorzaken periodieke perceptuele alternaties

In dit hoofdstuk laten we zien dat onderbroken ambigue stimuli zorgen voor een trage en regelmatige cyclus van perceptuele alternaties. Afzonderlijke alternaties worden gescheiden door periodes van soms meer dan tien minuten perceptuele stabiliteit. Alternaties worden talrijker als de tijdsduur tussen afzonderlijke presentaties wordt verkort, wat erop wijst dat dit proces niet autonoom is maar gedreven wordt door de visuele input zelf. De cyclus van perceptuele alternaties die we hier beschrijven verschilt van de alternatiecyclus tijdens continue presentatie van een ambigue stimulus, wat wijst op verschillen in de neurale basis. Het wiskundige model dat we gebruikten om de vindingen van hoofdstuk 5 te duiden, verklaart ook wat we in dit hoofdstuk zien. Er is nu dus een plausibele beschrijving van de dynamica van de hersenprocessen die verantwoordelijk zijn voor perceptie van onderbroken ambigue stimuli.

Hoofdstuk 7. Perceptueel geheugen: op weg naar neurale metingen

Er bestaat momenteel een imbalans tussen hoe goed modellen de dynamica achter perceptueel geheugen lijken te beschrijven (redelijk goed), en hoeveel we weten over de eigenlijke neurale activiteit die verantwoordelijk is voor dit fenomeen (vrij weinig). In dit hoofdstuk vatten we het model samen dat we gebruikt hebben om onze waarnemingen van hoofdstuk 4, 5 and 6 te begrijpen, en we beschrijven een experiment dat gericht is op het testen van een voorspelling van dit model. Daarna gebruiken we het inzicht dat zowel model als experimenten ons verschaffen om richting te geven aan onderzoek naar de neurale basis van perceptueel geheugen. We wijzen een aantal valkuilen aan die dat onderzoek zouden kunnen belemmeren, en beschrijven hoe die te ontwijken zijn.

Bibliography

- Abbott, L. F., J. A. Varela, K. Sen and S. B. Nelson (1997) Synaptic depression and cortical gain control. *Science*, 275(5297): 220–4
- Alais, D. and R. Blake (1999) Grouping visual features during binocular rivalry. *Vision Res*, 39(26): 4341–53
- Alais, D. and R. Blake, eds. (2005) *Binocular rivalry*. Cambridge, Massachusetts; London, England: The MIT Press
- Albrecht, D. G., S. B. Farrar and D. B. Hamilton (1984) Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. *J Physiol (Lond)*, 347: 713–39
- Anderson, R. (2001) The power law as an emergent property. *Memory & cognition*, 29(7): 1061–1068
- Andrews, T. J. and D. Purves (1997) Similarities in normal and binocularly rivalrous viewing. *Proc Natl Acad Sci USA*, 94(18): 9905–8
- Attneave, F. (1968) Triangles as ambiguous figures. *The American journal of psychology*, 81(3): 447–53
- Bichot, N. P. and J. D. Schall (2002) Priming in macaque frontal cortex during popout visual search: feature-based facilitation and location-based inhibition of return. *J Neurosci*, 22(11): 4675–85
- Billock, V. A. and B. H. Tsou (2003) *Encyclopedia of Optical Engineering*, chapter Vision, nonlinear dynamic approaches. New York: Marcel Dekker
- Blake, R. (1989) A neural theory of binocular rivalry. *Psychological review*, 96(1): 145–67
- Blake, R. (2001) A primer on binocular rivalry, including current controversies. *Brain & Mind*, 2(1): 5–38
- Blake, R. and R. Fox (1974) Adaptation to invisible gratings and the site of binocular rivalry suppression. *Nature*, 249(456): 488–90
- Blake, R. and N. K. Logothetis (2002) Visual competition. *Nat Rev Neurosci*, 3(1): 13–21
- Blake, R., R. P. O'Shea and T. J. Mueller (1992) Spatial zones of binocular rivalry in central and peripheral vision. *Vis Neurosci*, 8(5): 469–78

-
- Blake, R., K. V. Sobel and L. A. Gilroy (2003) Visual motion retards alternations between conflicting perceptual interpretations. *Neuron*, 39(5): 869–78
- Blake, R., D. Tadin, K. V. Sobel, T. A. Raissian and S. C. Chong (2006) Strength of early visual adaptation depends on visual awareness. *Proc Natl Acad Sci USA*, 103(12): 4783–8
- Blake, R., D. Westendorf and R. Fox (1990) Temporal perturbations of binocular rivalry. *Perception & psychophysics*, 48(6): 593–602
- Blake, R., D. H. Westendorf and R. Overton (1980) What is suppressed during binocular rivalry? *Perception*, 9(2): 223–31
- Blake, R., L. Zimba and D. Williams (1985) Visual motion, binocular correspondence and binocular rivalry. *Biological cybernetics*, 52(6): 391–7
- Bonds, A. B. (1991) Temporal dynamics of contrast gain in single cells of the cat striate cortex. *Vis Neurosci*, 6(3): 239–55
- Borsellino, A., A. de Marco, A. Allazetta, S. Rinesi and B. Bartolini (1972) Reversal time distribution in the perception of visual ambiguous stimuli. *Kybernetik*, 10(3): 139–44
- Bossink, C. J., P. F. Stalmeier and C. M. D. Weert (1993) A test of levelt’s second proposition for binocular rivalry. *Vision Res*, 33(10): 1413–1419
- Bradley, D. C., G. C. Chang and R. A. Andersen (1998) Encoding of three-dimensional structure-from-motion by primate area mt neurons. *Nature*, 392(6677): 714–7
- Brascamp, J. W., T. H. J. Knapen, R. Kanai, A. J. Noest, R. van Ee and A. V. van den Berg (2008) Multi-timescale perceptual history resolves visual ambiguity. *PLoS ONE*, 3(1): e1497
- Brascamp, J. W., T. H. J. Knapen, R. Kanai, R. van Ee and A. V. van den Berg (2007) Flash suppression and flash facilitation in binocular rivalry. *J Vis*, 7(12): 1–12
- Brascamp, J. W., J. Pearson, R. Blake and A. V. van den Berg (in prep) Intermittent ambiguous stimuli cause periodic perceptual alternations
- Brascamp, J. W., R. van Ee, W. R. Pestman and A. V. van den Berg (2005) Distributions of alternation rates in various forms of bistable perception. *J Vis*, 5(4): 287–98
- Breitmeyer, B. (1984) *Visual masking*. Oxford, New York: Clarendon Press, Oxford University Press
- Burke, D., D. Alais and P. Wenderoth (1999) Determinants of fusion of dichoptically presented orthogonal gratings. *Perception*, 28(1): 73–88
- Campana, G., A. Cowey and V. Walsh (2002) Priming of motion direction and area v5/mt: a test of perceptual memory. *Cereb Cortex*, 12(6): 663–9
- Carandini, M., D. J. Heeger and J. A. Movshon (1997) Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci*, 17(21): 8621–44
- Carpenter, R. H. and M. L. Williams (1995) Neural computation of log likelihood in control of saccadic eye movements. *Nature*, 377(6544): 59–62

-
- Carpenter, R. H. S. (1981) *Eye movements: cognition and visual perception*, chapter Oculomotor procrastination. Hillsdale, NJ: Erlbaum
- Carter, O. L. and P. Cavanagh (2007) Onset rivalry: brief presentation isolates an early independent phase of perceptual competition. *PLoS ONE*, 2(4): e343
- Carter, O. L. and J. D. Pettigrew (2003) A common oscillator for perceptual rivalries? *Perception*, 32(3): 295–305
- Chen, X. and S. He (2004) Local factors determine the stabilization of monocular ambiguous and binocular rivalry stimuli. *Curr Biol*, 14(11): 1013–7
- Chong, S. C. and R. Blake (2005) Exogenous attention and endogenous attention influence initial dominance in binocular rivalry. *Vision Res*, 46(11): 1794–803
- Cogan, R. (1973) Distributions of durations of perception in the binocular rivalry of contours. *The Journal of general psychology*, 89(2d Half): 297–304
- Corbetta, M. and G. L. Shulman (2002) Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci*, 3(3): 201–15
- Crick, F. and C. Koch (1998) Consciousness and neuroscience. *Cerebral Cortex*, 8(2): 97
- De Marco, A., P. Penengo and A. Trabucco (1977) Stochastic models and fluctuation in reversal time of ambiguous figures. *Perception*, 6(6): 645–656
- DeAngelis, G. C., J. G. Robson, I. Ohzawa and R. D. Freeman (1992) Organization of suppression in receptive fields of neurons in cat visual cortex. *J Neurophysiol*, 68(1): 144–63
- Desimone, R. and J. Duncan (1995) Neural mechanisms of selective visual attention. *Annu Rev Neurosci*, 18: 193–222
- Dorris, M. C., M. Paré and D. P. Munoz (2000) Immediate neural plasticity shapes motor performance. *J Neurosci*, 20(1): RC52
- Drew, P. J. and L. F. Abbott (2006) Models and properties of power-law adaptation in neural systems. *J Neurophysiol*, 96(2): 826–33
- Ernst, M. O., M. S. Banks and H. H. Bühlhoff (1999) Touch can change visual slant perception. *Nat Neurosci*, 3(1): 69–73
- Ernst, M. O. and H. H. Bühlhoff (2004) Merging the senses into a robust percept. *Trends Cogn Sci*, 8(4): 162–9
- Fang, F. and S. He (2004) Stabilized structure from motion without disparity induces disparity adaptation. *Curr Biol*, 14(3): 247–51
- Fecteau, J. H. and D. P. Munoz (2003) Exploring the consequences of the previous trial. *Nat Rev Neurosci*, 4(6): 435–43
- Fox, R. and J. Herrmann (1967) Stochastic properties of binocular rivalry alternations. *Perception & psychophysics*, 2(9): 432–436

-
- Fox, R. and F. Rasche (1969) Binocular rivalry and reciprocal inhibition. *Perception & psychophysics*, 5(4): 215–217
- Freeman, A. W. (2005) Multistage model for binocular rivalry. *J Neurophysiol*, 94(6): 4412–20
- Georgeson, M. and J. M. Georgeson (1987) Facilitation and masking of briefly presented gratings: time-course and contrast dependence. *Vision Res*, 27(3): 369–79
- Gilboa, G., R. Chen and N. Brenner (2005) History-dependent multiple-time-scale dynamics in a single-neuron model. *J Neurosci*, 25(28): 6479–89
- Grossmann, J. K. and A. C. Dobbins (2005) Competition in bistable vision is attribute-specific. *Vision Res*, 46(3): 285–92
- Hollins, M. (1980) The effect of contrast on the completeness of binocular rivalry suppression. *Perception & psychophysics*, 27(6): 550–6
- Hollins, M. and K. Hudnell (1980) Adaptation of the binocular rivalry mechanism. *Invest Ophthalmol Vis Sci*, 19(9): 1117–20
- Horton, J. C. and W. F. Hoyt (1991) The representation of the visual field in human striate cortex. a revision of the classic holmes map. *Arch Ophthalmol*, 109(6): 816–24
- Huber, D. E. and R. C. O'Reilly (2003) Persistence and accommodation in short-term priming and other perceptual paradigms: temporal segregation through synaptic depression. *Cognitive science*, 27(3): 403–430
- Huk, A. C. and M. N. Shadlen (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci*, 25(45): 10420–36
- Hupé, J.-M. and N. Rubin (2003) The dynamics of bi-stable alternation in ambiguous motion displays: a fresh look at plaids. *Vision Res*, 43(5): 531–48
- Jiang, Y., A. J. Pantle and L. S. Mark (1998) Visual inertia of rotating 3-d objects. *Perception & psychophysics*, 60(2): 275–86
- Kalarickal, G. J. and J. A. Marshall (2000) Neural model of temporal and stochastic properties of binocular rivalry. *Neurocomputing*, 32-33: 843–853
- Kamphuisen, A. P., R. J. A. V. Wezel and R. van Ee (2007) Inter-ocular transfer of stimulus cueing in dominance selection at the onset of binocular rivalry. *Vision Res*, 47(9): 1142–4
- Kanai, R. and Y. Kamitani (2003) Time-locked perceptual fading induced by visual transients. *Journal of cognitive neuroscience*, 15(5): 664–72
- Kanai, R., T. H. J. Knapen, R. van Ee and F. A. J. Verstraten (2007) Disruption of implicit perceptual memory by intervening neutral stimuli. *Vision Res*, 47(20): 2675–83
- Kanai, R., N. Tsuchiya and F. A. J. Verstraten (2006) The scope and limits of top-down attention in unconscious visual processing. *Curr Biol*, 16(23): 2332–6

-
- Kanai, R. and F. A. J. Verstraten (2005) Perceptual manifestations of fast neural plasticity: motion priming, rapid motion aftereffect and perceptual sensitization. *Vision Res*, 45(25-26): 3109–16
- Kanai, R. and F. A. J. Verstraten (2006) Attentional modulation of perceptual stabilization. *Proc Biol Sci*, 273(1591): 1217–22
- Keyesers, C. and D. Perrett (2002) Visual masking and rsvp reveal neural competition. *Trends Cogn Sci*, 6(3): 120–125
- Kim, Y.-J., M. Grabowecky and S. Suzuki (2005) Stochastic resonance in binocular rivalry. *Vision Res*, 46(3): 392–406
- Klink, P. C., R. van Ee, M. M. Nijs, G. J. Brouwer and A. J. Noest (2008) Early interactions between neuronal adaptation and voluntary control determine perceptual choices in bistable vision. *J Vis*, 8(5): 1–18
- Knapen, T., R. van Ee and R. Blake (2007) Stimulus motion propels traveling waves in binocular rivalry. *PLoS ONE*, 2: e739
- Koch, C. and N. Tsuchiya (2007) Attention and consciousness: two distinct brain processes. *Trends Cogn Sci (Regul Ed)*, 11(1): 16–22
- Kornmeier, J. and M. Bach (2004) Early neural activity in necker-cube reversal: evidence for low-level processing of a gestalt phenomenon. *Psychophysiology*, 41(1): 1–8
- Kornmeier, J. and M. Bach (2005) The necker cube—an ambiguous figure disambiguated in early visual processing. *Vision Res*, 45(8): 955–60
- Kovács, I., T. V. Papathomas, M. Yang and A. Fehér (1996) When the brain changes its mind: interocular grouping during binocular rivalry. *Proc Natl Acad Sci USA*, 93(26): 15508–11
- Kreiman, G. (2001) *On the neuronal activity in the human brain during visual recognition, imagery and binocular rivalry*. Ph.D. thesis, California Institute of Technology
- Kreiman, G., I. Fried and C. Koch (2002) Single-neuron correlates of subjective vision in the human medial temporal lobe. *Proc Natl Acad Sci USA*, 99(12): 8378–83
- Kreiman, G., I. Fried and C. Koch (2005) *Binocular rivalry*, chapter Single-neuron correlates of subjective vision in the human medial temporal lobe. Cambridge, Massachusetts; London, England: The MIT Press
- Kristjánsson, A. (2006) Rapid learning in attention shifts: A review. *Vis Cogn*, 13(3): 324–362
- Laing, C. R. and C. C. Chow (2002) A spiking neuron model for binocular rivalry. *Journal of computational neuroscience*, 12(1): 39–53
- Lankheet, M. J. M. (2006) Unraveling adaptation and mutual inhibition in perceptual rivalry. *J Vis*, 6(4): 304–10
- Lee, S.-H., R. Blake and D. J. Heeger (2004) Traveling waves of activity in primary visual cortex during binocular rivalry. *Nat Neurosci*, 8(1): 22–3

-
- Lehky, S. R. (1988) An astable multivibrator model of binocular rivalry. *Perception*, 17: 215–228
- Lehky, S. R. (1995) Binocular rivalry is not chaotic. *Proc R Soc Lond, B, Biol Sci*, 259(1354): 71–76
- Leopold, D. A. and N. K. Logothetis (1996) Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature*, 379(6565): 549–553
- Leopold, D. A. and N. K. Logothetis (1999) Multistable phenomena: changing views in perception. *Trends Cogn Sci*, 3(7): 254–264
- Leopold, D. A., M. Wilke, A. Maier and N. K. Logothetis (2002) Stable perception of visually ambiguous patterns. *Nat Neurosci*, 5(6): 605–9
- Levelt, W. J. (1966) The alternation process in binocular rivalry. *British Journal of Psychology*, 67(3): 225–238
- Levelt, W. J. (1967) Note on the distribution of dominance times in binocular rivalry. *British Journal of Psychology*, 58(1): 143–5
- Liu, L., C. W. Tyler and C. M. Schor (1992) Failure of rivalry at low contrast: evidence of a suprathreshold binocular summation process. *Vision Res*, 32(8): 1471–9
- Logothetis, N. K. (1998) Single units and conscious vision. *Philos Trans R Soc Lond, B, Biol Sci*, 353(1377): 1801–18
- Logothetis, N. K., D. A. Leopold and D. L. Sheinberg (1996) What is rivalling during binocular rivalry? *Nature*, 380(6575): 621–4
- Logothetis, N. K. and J. D. Schall (1989) Neuronal correlates of subjective visual perception. *Science*, 245(4919): 761–3
- Long, G. and A. Olszweski (1999) To reverse or not to reverse: when is an ambiguous figure not ambiguous? *Am J Psych*, 112: 41–71
- Long, G. M. and C. J. Moran (2007) How to keep a reversible figure from reversing: teasing out top-down and bottom-up processes. *Perception*, 36(3): 431–45
- Long, G. M. and T. C. Toppino (1994) A new twist on the rotating-trapezoid illusion: evidence for neural-adaptation effects. *Perception*, 23(6): 619–34
- Long, G. M., T. C. Toppino and G. W. Mondin (1992) Prime time: fatigue and set effects in the perception of reversible figures. *Perception & psychophysics*, 52(6): 609–16
- Lumer, E. D., K. J. Friston and G. Rees (1998) Neural correlates of perceptual rivalry in the human brain. *Science*, 280(5371): 1930–4
- Lumer, E. D. and G. Rees (1999) Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proc Natl Acad Sci USA*, 96(4): 1669–73
- Magnussen, S. and M. W. Greenlee (1999) The psychophysics of perceptual memory. *Psychological research*, 62(2-3): 81–92

-
- Maier, A., N. K. Logothetis and D. A. Leopold (2002) Neural activity during stable perception of ambiguous displays in monkey visual cortex [abstract]. *Soc Neurosci Abstr*
- Maier, A., M. Wilke, N. K. Logothetis and D. A. Leopold (2003) Perception of temporally interleaved ambiguous patterns. *Curr Biol*, 13(13): 1076–85
- Maljkovic, V. and K. Nakayama (1994) Priming of pop-out: I. role of features. *Memory & cognition*, 22(6): 657–72
- Maljkovic, V. and K. Nakayama (2000) Priming of pop-out iii: A short-term implicit memory system beneficial for rapid target selection. *Vis Cogn*, 7: 571–595
- Maloney, L. T., M. F. D. Martello, C. Sahm and L. Spillmann (2005) Past trials influence perception of ambiguous motion quartets through pattern completion. *Proc Natl Acad Sci USA*, 102(8): 3164–9
- Martínez-Trujillo, J. and S. Treue (2002) Attentional modulation strength in cortical area mt depends on stimulus contrast. *Neuron*, 35(2): 365–70
- McLachlan, G. J. and D. Peel (2000) *Finite mixture models*. New York, Chichester, Weinheim, Brisbane, Singapore, Toronto: John Wiley and Sons
- McPeck, R. M., V. Maljkovic and K. Nakayama (1999) Saccades require focal attention and are facilitated by a short-term memory system. *Vision Res*, 39(8): 1555–66
- Meng, M. and F. Tong (2004) Can attention selectively bias bistable perception? differences between binocular rivalry and ambiguous figures. *J Vis*, 4(7): 539–51
- Mitchell, J. F., G. R. Stoner and J. H. Reynolds (2004) Object-based attention determines dominance in binocular rivalry. *Nature*, 429(6990): 410–3
- Mongillo, G., O. Barak and M. Tsodyks (2008) Synaptic theory of working memory. *Science*, 319(5869): 1543
- Moreno Bote, R., J. Rinzel and N. Rubin (2007) Noise-induced alternations in an attractor network model of perceptual bi-stability. *J Neurophysiol*, 98(3): 1125–1139
- Mueller, T. J. (1990) A physiological model of binocular rivalry. *Vis Neurosci*, 4(1): 63–73
- Mueller, T. J. and R. Blake (1989) A fresh look at the temporal dynamics of binocular rivalry. *Biological cybernetics*, 61(3): 223–232
- Müller, J. R., A. B. Metha, J. Krauskopf and P. Lennie (1999) Rapid adaptation in visual cortex to the structure of images. *Science*, 285(5432): 1405–8
- Murata, T., N. Matsui, S. Miyauchi, Y. Kakita and T. Yanagida (2003) Discrete stochastic process underlying perceptual rivalry. *Neuroreport*, 14(10): 1347–52
- Nawrot, M. and R. Blake (1989) Neural integration of information specifying structure from stereopsis and motion. *Science*, 244(4905): 716–718

-
- Necker, L. A. (1832) Observations on some remarkable optical phenomena seen in switzerland; and on an optical phenomenon which occurs on viewing a figure of a crystal or geometrical solid. *Lond Edinb Phil Mag J Sci*, 1(5): 329–337
- Noest, A. J., R. van Ee, M. M. Nijs and R. J. A. V. Wezel (2007) Percept-choice sequences driven by interrupted ambiguous stimuli: a low-level neural model. *Journal of vision*, 7(8): 10
- Ohzawa, I., G. Sclar and R. D. Freeman (1985) Contrast gain control in the cat's visual system. *J Neurophysiol*, 54(3): 651–67
- Ono, H. and R. Barbeito (1985) Utrocular discrimination is not sufficient for utrocular identification. *Vision Res*, 25(2): 289–99
- Orbach, J., D. Ehrlich and H. Heath (1963) Reversibility of the necker cube. i. an examination of the concept of "satiation of orientation". *Percept Mot Skills*, 17: 439–458
- Orbach, J., E. Zucker and R. Olson (1966) Reversibility of the necker cube vii. reversal rate as a function of figure-on and figure-off durations. *Percept Mot Skills*, 22: 615–618
- O'Shea, R. P. (2004) Psychophysics: catching the old codger's eye. *Curr Biol*, 14(12): R478–9
- O'Shea, R. P., A. J. Sims and D. G. Govan (1997) The effect of spatial frequency and field size on the spread of exclusive visibility in binocular rivalry. *Vision Res*, 37(2): 175–83
- Paffen, C. L. E., D. Tadin, S. F. te Pas, R. Blake and F. A. J. Verstraten (2005) Adaptive center-surround interactions in human vision revealed during binocular rivalry. *Vision Res*, 46(5): 599–604
- Pasternak, T. and M. W. Greenlee (2005) Working memory in primate sensory systems. *Nat Rev Neurosci*, 6(2): 97–107
- Pastukhov, A. and J. Braun (2007) Temporal characteristics of priming effects on the perception of ambiguous patterns [abstract]. *J Vis*, 7(9): 367
- Pearson, J. and J. W. Brascamp (2008) Sensory memory for ambiguous vision. *Trends Cogn Sci*, 12(9): 334–341
- Pearson, J., C. Clifford and F. Tong (2008) The functional impact of mental imagery on conscious perception. *Curr Biol*, 18(13): 982–1017
- Pearson, J. and C. W. G. Clifford (2004) Determinants of visual awareness following interruptions during rivalry. *J Vis*, 4(3): 196–202
- Pearson, J. and C. W. G. Clifford (2005) Mechanisms selectively engaged in rivalry: normal vision habituates, rivalrous vision primes. *Vision Res*, 45(6): 707–14
- Pearson, J., D. Tadin and R. Blake (2007) The effects of transcranial magnetic stimulation on visual rivalry. *J Vis*, 7(7): 2.1–11
- Pearson, K. (1894) Contributions to the mathematical theory of evolution. *Philos Trans R Soc Lond, A*, 185: 71–110

-
- Pelli, D. G. (1990) *Vision: Coding and efficiency*, chapter The quantum efficiency of vision. Cambridge: Cambridge University Press
- Pestman, W. R. (1998) *Mathematical statistics, an introduction*. Berlin, New York: De Gruyter
- Pinkus, A. and A. J. Pantle (1997) Probing visual motion signals with a priming paradigm. *Vision Res*, 37(5): 541–52
- Pressnitzer, D. and J.-M. Hupé (2006) Temporal dynamics of auditory and visual bistability reveal common principles of perceptual organization. *Curr Biol*, 16(13): 1351–7
- Raghavan (1989) *Sources of visual noise*. Ph.D. thesis, Syracuse University
- Reddi, B. A. and R. H. S. Carpenter (2000) The influence of urgency on decision time. *Nat Neurosci*, 3(8): 827–30
- Reynolds, J. H., T. Pasternak and R. Desimone (2000) Attention increases sensitivity of v4 neurons. *Neuron*, 26(3): 703–14
- Rose, D. and I. Lowe (1982) Dynamics of adaptation to contrast. *Perception*, 11(5): 505–28
- Saito, H., M. Yukie, K. Tanaka, K. Hikosaka, Y. Fukada and E. Iwai (1986) Integration of direction signals of image motion in the superior temporal sulcus of the macaque monkey. *Journal of Neuroscience*, 6(1): 145
- Schall, J. D. and K. G. Thompson (1999) Neural selection and control of visually guided eye movements. *Annu Rev Neurosci*, 22: 241–59
- Schwartz, O., A. Hsu and P. Dayan (2007) Space and time in visual context. *Nat Rev Neurosci*, 8(7): 522–535
- Sclar, G., J. H. R. Maunsell and P. Lennie (1990) Coding of image contrast in central visual pathways of the macaque monkey. *Vision Res*, 30(1): 1–10
- Sengpiel, F., C. Blakemore and R. Harrad (1995) Interocular suppression in the primary visual cortex: a possible neural basis of binocular rivalry. *Vision Res*, 35(2): 179–95
- Sheinberg, D. L. and N. K. Logothetis (1997) The role of temporal cortical areas in perceptual organization. *Proc Natl Acad Sci USA*, 94(7): 3408–13
- Shiraishi, S. (1977) A test on levelt's model on binocular rivalry. *Japanese Psychological Research*, 19: 129–135
- Silver, M. A. and N. K. Logothetis (2004) Grouping and segmentation in binocular rivalry. *Vision Res*, 44(14): 1675–92
- Smith, P. L. and R. Ratcliff (2004) Psychology and neurobiology of simple decisions. *Trends Neurosci*, 27(3): 161–8
- Snowden, R. J., S. Treue, R. G. Erickson and R. A. Andersen (1991) The response of area mt and v1 neurons to transparent motion. *J Neurosci*, 11(9): 2768–85

-
- Soetens, E., A. Melis and W. Notebaert (2004) Sequence learning and sequential effects. *Psychological research*, 69(1-2): 124–37
- Sterzer, P. and G. Rees (2008) A neural basis for percept stabilization in binocular rivalry. *The Journal of Cognitive Neuroscience*, 20(3): 389
- Sterzer, P., M. O. Russ, C. Preibisch and A. Kleinschmidt (2002) Neural correlates of spontaneous direction reversals in ambiguous apparent visual motion. *Neuroimage*, 15(4): 908–916
- Stollenwerk, L. and M. Bode (2003) Lateral neural model of binocular rivalry. *Neural computation*, 15(12): 2863–82
- Suzuki, S. and M. Grabowecky (2002) Evidence for perceptual "trapping" and adaptation in multistable binocular rivalry. *Neuron*, 36(1): 143–57
- Suzuki, S. and M. Grabowecky (2007) Long-term speeding in perceptual switches mediated by attention-dependent plasticity in cortical visual processing. *Neuron*, 56(4): 741–753
- Tanaka, Y. and D. Sagi (1998) A perceptual memory for low-contrast visual signals. *Proc Natl Acad Sci USA*, 95(21): 12729–33
- Tong, F. (2001) Competing theories of binocular rivalry: a possible resolution. *Brain & Mind*, 2(1): 55–83
- Tong, F. and S. A. Engel (2001) Interocular rivalry revealed in the human cortical blind-spot representation. *Nature*, 411(6834): 195–9
- Tong, F., M. Meng and R. Blake (2006) Neural bases of binocular rivalry. *Trends Cogn Sci*, 10(11): 502–11
- Tong, F., K. Nakayama, J. T. Vaughan and N. Kanwisher (1998) Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron*, 21(4): 753–9
- Tsuchiya, N. and C. Koch (2005) Continuous flash suppression reduces negative afterimages. *Nat Neurosci*, 8(8): 1096–101
- Ulanovsky, N., L. Las, D. Farkas and I. Nelken (2004) Multiple time scales of adaptation in auditory cortex neurons. *J Neurosci*, 24(46): 10440–53
- Van Dam, L. C. J. and R. van Ee (2005) The role of (micro)saccades and blinks in perceptual bi-stability from slant rivalry. *Vision Res*, 45(18): 2417–2435
- Van den Berg, A. V. and E. M. van Loon (2005) An invariant for timing of saccades during visual search. *Vision Res*, 45(12): 1543–55
- Van Ee, R. (2005) Dynamics of perceptual bi-stability for stereoscopic slant rivalry and a comparison with grating, house-face, and necker cube rivalry. *Vision Res*, 45(1): 29–40
- Van Ee, R., L. C. J. van Dam and G. J. Brouwer (2004) Voluntary control and the dynamics of perceptual bi-stability. *Vision Res*, 45(1): 41–55
- Van Ee, R., L. C. J. van Dam and C. J. Erkelens (2002) Bi-stability in perceived slant when binocular disparity and monocular perspective specify different slants. *JOV*, 2(9): 597–607

-
- Van Loon, E. M., I. T. C. Hooge and A. V. van den Berg (2002) The timing of sequences of saccades in visual search. *Proc Biol Sci*, 269(1500): 1571–9
- Von Helmholtz, H. (1867) *Handbuch der physiologischen Optik*, volume 9 of *Allgemeine Encyclopaedie der Physik*. Leipzig: Voss
- Wade, N. J. (1975) Binocular rivalry between single lines viewed as real images and afterimages. *Perception & psychophysics*, 17(6): 571–577
- Walker, P. (1975) Stochastic properties of binocular rivalry alternations. *Perception & psychophysics*, 18: 467–473
- Watson, A. B. and D. G. Pelli (1983) Quest: a bayesian adaptive psychometric method. *Perception & psychophysics*, 33(2): 113–20
- Wheatstone, C. (1838) Contributions to the physiology of vision. i: On some remarkable, and hitherto unobserved, phenomena of binocular vision. *Philos Trans R Soc*, 1: 371–394
- Wilke, M., N. K. Logothetis and D. A. Leopold (2003) Generalized flash suppression of salient visual targets. *Neuron*, 39(6): 1043–52
- Wilson, H. (2007) Minimal physiological conditions for binocular rivalry and rivalry memory. *Vision Res*, 47(21): 2741–2750
- Wilson, H. R. (2003) Computational evidence for a rivalry hierarchy in vision. *Proc Natl Acad Sci USA*, 100(24): 14499–503
- Wilson, H. R. (2005) *Binocular rivalry*, chapter Rivalry and perceptual oscillations: a dynamical synthesis. Cambridge, Massachusetts; London, England: The MIT Press
- Wilson, H. R., R. Blake and S.-H. Lee (2001) Dynamics of travelling waves in visual perception. *Nature*, 412(6850): 907–10
- Wilton, R. N. (1985) The recency effect in the perception of ambiguous figures. *Perception*, 14(1): 53–61
- Wolfe, J. M. (1983) Influence of spatial frequency, luminance, and duration on binocular rivalry and abnormal fusion of briefly presented dichoptic stimuli. *Perception*, 12(4): 447–56
- Wolfe, J. M. (1984) Reversing ocular dominance and suppression in a single flash. *Vision Res*, 24(5): 471–8
- Yang, Y., D. Rose and R. Blake (1992) On the variety of percepts associated with dichoptic viewing of dissimilar monocular stimuli. *Perception*, 21(1): 47–62
- Zhou, Y. H., J. B. Gao, K. D. White, I. Merk and K. Yao (2004) Perceptual dominance time distributions in multistable visual perception. *Biological cybernetics*, 90(4): 256–63

Acknowledgements

I would like to thank several people who coauthored one or more chapters of this thesis.

Bert, one of the first things that comes to mind when I think of all the ways in which you contributed to my development over the years, is your ability to approach scientific subjects with an open mind. When I was disappointed about some experiment that did not work the way we expected, you often surprised me by pointing out other aspects of the data that looked interesting. You taught me not to get stuck focusing on prior expectations, but to see whatever result as a potentially relevant finding. I also greatly appreciated your loose style of supervising: providing help and feedback when needed, but allowing me to set my own course when appropriate.

Raymond, although I was not a member of your group you have been something of an unofficial 'second supervisor'. Of particular value to me were your improvements on my writing style, and also your exceptionally international outlook, which stood at the basis of my foreign adventures with Tomas.

André, without your 2007 paper on perceptual stabilization and your continuing input after that, this thesis would have been nothing like it is now. Not only does your 2007 theory form the basis of our thinking on the subject, your enthusiastic and ever incisive feedback provided essential clues and reassurance that allowed us to develop our ideas further.

Wiebe, every PhD student encounters moments when their project seems stuck and none of their efforts seem to work out as planned. I think my first such moment was when I encountered technical problems working on what is now the first chapter of this thesis. When I eventually called you for advice, you did not just solve my problem, but were instantly enthusiastic about the project, wrote software that sped things up tremendously, and remained closely involved in the rest of the project. It is a shame we never got around to any further joint projects, but working together on this one was a great pleasure.

Tomas, we briefly shared an office when we had both just started as PhD students, and since then we pretty much developed simultaneously. We designed research, performed experiments, worked at foreign institutes and wrote scientific papers together. I would like to thank you for the good times we had in Utrecht and abroad, for innumerable conversations that floated in and out of science, and for letting me use your guest room every time I missed my last train.

Ryota, Chapters 4 and 5 of this thesis have their basis in Tomas and my visit to you in Pasadena in 2006. One thing I recall is your instant and insightful feedback to results we produced. Also, I think we benefited a lot from the fact that you were a few steps ahead of us in terms of knowing your way around 'the world of science'. I would like to thank you for many interesting, instructive and enjoyable discussions, both during our period in Pasadena and since then.

Joel and Randolph, thank you for making me feel right at home during my visits to Nashville. Randolph, of the many things I valued in our conversations, I will mention only your ability to combine a detailed understanding with a big-picture perspective. It is a skill that I hope to master myself someday. Joel, we spent much time in front of a small laptop screen, meticulously reworking paragraphs in our manuscript. It was a process that I enjoyed a lot, and I think we can be proud of the end result.

Publications

Journal papers

- J.W. Brascamp, J. Pearson, R. Blake and A.V. van den Berg (in prep) Intermittent ambiguous stimuli cause periodic perceptual alternations
- J. Pearson and J.W. Brascamp (2008) Sensory memory for ambiguous vision. *Trends in Cognitive Sciences*, 12(9): 334-341
- J.W. Brascamp, T.H.J. Knapen, R. Kanai, A.J. Noest, R. van Ee and A.V. van den Berg (2008) Multi-timescale perceptual history resolves visual ambiguity. *PLoS ONE*, 3(1): e1497
- T.H.J. Knapen, R. Kanai, J.W. Brascamp, J.J.A. van Boxtel and R. van Ee (2007) Distance in feature space determines exclusivity in visual rivalry. *Vision Research*, 47(26): 3269-3275
- J.W. Brascamp, T.H.J. Knapen, R. Kanai, R. van Ee and A.V. van den Berg (2007) Flash suppression and flash facilitation in binocular rivalry. *Journal of Vision*, 7(12): 12, 1-12
- J.W. Brascamp, R. van Ee, A.J. Noest, R.H.A.H. Jacobs and A.V. van den Berg (2006) The time course of binocular rivalry reveals a fundamental role of noise. *Journal of Vision*, 6(11): 1244-1256
- R. van Ee, A.J. Noest, J.W. Brascamp and A.V. van den Berg (2006) Attentional control over either of the two competing percepts of ambiguous stimuli revealed by a two-parameter analysis: Means don't make the difference. *Vision Research*, 46(19): 3129-3141
- J.W. Brascamp, R. van Ee, W. Pestman and A.V. van den Berg (2005) Distributions of alternation rates in various forms of bistable perception. *Journal of Vision*, 5(4): 287-298
- J. Grasman, J.W. Brascamp, J.L. van Leeuwen and B. van Putten (2003) The multifractal structure of arterial trees. *Journal of Theoretical Biology*, 220(1): 75-82

Conference abstracts

- J.W. Brascamp, T.H.J. Knapen and J.J.A. van Boxtel (2008) Invisible images can leave stronger afterimages than visible images. ECVP 08, Utrecht
- M.C. de Jong, J.W. Brascamp, C. Kemner and R. van Ee (2008) Neural correlates of memory in ambiguous perception measured with EEG. ECVP 08, Utrecht

-
- J.W. Brascamp, J. Pearson, R. Blake and A.V. van den Berg (2008) Slow changes in neural state mediate percept switches in intermittent binocular rivalry. VSS 08, Naples
 - T.H.J. Knapen, J. Pearson, J.W. Brascamp, R. van Ee and R. Blake (2008) The role of frontal areas in alternations during perceptual bistability. VSS 08, Naples
 - J.W. Brascamp, T.H.J. Knapen, A.J. Noest, R. van Ee and A.V. van den Berg (2007) History dependence on multiple timescales in ambiguous figure perception. SFN 07, San Diego
 - J.W. Brascamp, T.H.J. Knapen, N.V. Stolk, R. van Ee and A.V. van den Berg (2007) Temporal characteristics of history dependence in ambiguous figure perception. ECVF 07, Arezzo
 - J.W. Brascamp, T.H.J. Knapen, R. Kanai, R. van Ee and A.V. van den Berg (2007) Perceptual memory of ambiguous figures survives spontaneous perceptual alternations. VSS 07, Sarasota
 - L.C.J. van Dam, R. Mulder, A.J. Noest, J.W. Brascamp, A.V. van den Berg and R. van Ee (2007) Sequential dependency for percept durations in binocular rivalry. VSS 07, Sarasota
 - R. van Ee, J.W. Brascamp, G.J. Brouwer, L.C.J. van Dam (2006) Endogenous control over either of the two percepts that compete for visual awareness in perceptual rivalry. ASSC 06, Oxford
 - J.W. Brascamp, R. van Ee, A.J. Noest and A.V. van den Berg (2006) Temporal dynamics of binocular rivalry imply an important role for noise. ECVF 06, St. Petersburg
 - R. van Ee, J.W. Brascamp, A.J. Noest and A.V. van den Berg (2006) Attentional control over either of the two competing percepts in the Necker cube. ECVF 06, St. Petersburg
 - J.W. Brascamp, A.J. Noest, R. van Ee and A.V. van den Berg (2006) Transition phases show the importance of noise in binocular rivalry. VSS 06, Sarasota
 - J.W. Brascamp, A.J. Noest and A.V. van den Berg (2005) The third percept in bistable perception. VSS 05, Sarasota
 - J.W. Brascamp, A.V. van den Berg and R. van Ee (2004) Shared neural circuitry for switching between perceptual states and ocular motor states? VSS 04, Sarasota
 - J.W. Brascamp, S. Kranenbarg, J.L. van Leeuwen and M. Muller (2000) Morphological restrictions in fish embryos due to limited nutrient supply and metabolic waste product deposition. SEB 00, Exeter

Curriculum vitae

Jan Willem Brascamp was born in Bennekom, the Netherlands on December 24, 1977. He went to highschool at the Revis Lyceum in Doorn, where he graduated in 1996. He went on to study biology (BSc and MSc) at Wageningen University, where he specialized in zoology and theoretical biology. In 2003 he started as a PhD student at Utrecht University's Functional Neurobiology group, supervised by Prof Bert van den Berg. As part of his PhD program he made research visits to Prof Shinsuke Shimojo's research group at the California Institute of Technology in Pasadena, CA, and to Prof Randolph Blake's group at Vanderbilt University in Nashville, TN. In October 2008 he will start as a post doctoral associate at the lab of Dr Raymond van Ee at Utrecht University, and in January 2009 he will move to Prof Randolph Blake's group at Vanderbilt University.