THE ROLE OF PREVIOUS EXPERIENCE IN CONSCIOUS PERCEPTION

Dissertation
zur Erlangung des Doktorgrades
der Naturwissenschaften

vorgelegt beim Fachbereich Psychologie und Sportwissenschaften
der Johann Wolfgang Goethe – Universität
in Frankfurt am Main

von
Caspar Martin Schwiedrzik
aus Köln

Frankfurt am Main 2011
D30

vom Fachbereich Psychologie und Sportwissenschaften der Johann Wolfgang Goethe – Universität als Dissertation angenommen.

Dekan:
Prof. Dr. Rolf van Dick

Gutachter:
Prof. Dr. Christian Fiebach, Johann Wolfgang Goethe-Universität Frankfurt a.M.
Prof. Dr. Dr. hc. mult. Wolf Singer, Max Planck Institut für Hirnforschung, Frankfurt a.M.
Prof. Dr. Jochen Kaiser, Johann Wolfgang Goethe-Universität Frankfurt a.M.

Datum der Disputation: 14.12.2011

to my beloved wife

ACKNOWLEDGEMENTS

My deepest thanks go to

… Prof. Singer for teaching me that science is not about showing something, but about finding something out. That was the best advice given and it came at the right time.

... Prof. Fiebach for being my co-supervisor at the Goethe University.

... Prof. Kaiser for his willingness to examine my thesis as a third reviewer.

...Prof. van Dick for assuring a rapid procedure, and Mrs. Siebert for her help in implementing it.

… Axel Kohler for his openness and his generous and never ending support.

… my friends at the Max Planck Institute for Brain Research and the Brain Imaging Center, in particular to Erhan Genç for being a great office mate, to Sara van Leeuwen for getting a smile on my face not only during Salad Club, to Jaan Aru for his fierce opposition to my interpretations, to Felipe Aedo-Jury for sharing the passion for soccer, to Amra Hodzic for her amazing kindness, to Arjen Alink for being the only other male in the Salad Club, and to Sergio Neuenschwander, Kerstin Schmidt, Bruss Lima, Gordon Pipa, Sylvia van Stijn, Anna Meyer, Anna-Maria Leugner, Ovidiu Jurjut, Viola Priesemann, Andreea Lazar, Sandra Anti, Chris Kell, Florian Beißner, and Tonio Heidegger for the great time.

… my past and present colleagues in the Department of Neurophysiology, the Neuroimaging group, the Brain Imaging Center, and the newly arrived ESI for creating a stimulating environment throughout the years.

… my coauthors Eugenio Rodriguez, Christian Ruff, and Notger Müller for the good collaboration.

… my Diploma students Frauke Leitner and Pascal Jacob for their enthusiasm and hard work.

… Paula Morgenstern, Steffen Angermair, Teresa Übelacker, and Georg-Friedrich Paasch for their help with data acquisition.

… Ralf Deichmann and Ulrike Nöth for their help with the fMRI sequences.

… Michael Schmid, Talis Bachmann, Bruce Bridgeman, and Michael Herzog for their advice along the way.

… Yusuf for his kind and honest words, and for protecting us in the night and on the weekends.

… my parents Susanne and Wolfgang for setting the standards, providing the basis, love, and support that brought me here.

… my brothers Jakob and Ludwig for setting the standards even higher.

… Lucia, whom I love and admire: You disentangle my twisted thoughts, you uncage my curiosity, you hold me when I fail, you keep me going, you are always there. Without you none of this would have been possible.

The research presented in this thesis was carried out at the Max Planck Institute for Brain Research, Department of Neurophysiology. It was funded by the Max Planck Society. All published material has been reproduced with permission from the respective publisher (Chapter 2: ARVO, Chapter 3: Proceedings of the National Academy of Sciences of the United States of America, Chapter 4: Journal of Neuroscience).

CONTENTS

CHAPTER 1

CHAPTER 2

CHAPTER 3

CHAPTER 4

CHAPTER 5

CHAPTER 6

A

ABSTRACT

Which factors determine whether a stimulus is consciously perceived or unconsciously processed? Here, I investigate how previous experience on two different time scales – long term experience over the course of several days, and short term experience based on the previous trial – impact conscious perception. Regarding long term experience, I investigate how perceptual learning does not only change the capacity to process stimuli, but also the capacity to consciously perceive them. To this end, subjects are trained extensively to discriminate between masked stimuli, and concurrently rate their subjective experience. Both the ability to discriminate the stimuli as well as subjective awareness of the stimuli increase as a function of training. However, these two effects are not simple byproducts of each other. On the contrary, they display different time courses, with above chance discrimination performance emerging before subjective experience; importantly, the two learning effects also rely on different circuits in the brain: Moving the stimuli outside the trained receptive field size abolishes the learning effects on discrimination ability, but preserves the learning effects on subjective awareness. This indicates that the receptive fields serving subjective experience are larger than the ones serving objective performance, and that the channels through which they receive their information are arranged in parallel. Regarding short term experience, I investigate how memory based predictions arising from information acquired on the trial before affect visibility and the neural correlates of consciousness. To this end, I vary stimulus evidence as well as predictability and acquire electroencephalographic data. A comparison of the neural processes distinguishing consciously perceived from unperceived trials with and without predictions reveals that predictions speed up processing, thus shifting the neural correlates forward in time. Thus, the neural correlates of consciousness display a previously unappreciated flexibility in time and do not arise invariably late as had been predicted by some theorists. Admittedly, however, previous experience does not always stabilize perception. Instead, previous experience can have the reverse effect: Seeing the opposite of what was there, as in so-called repulsive aftereffects. Here, I investigate what determines the direction of previous experience using multistable stimuli. In a functional magnetic resonance imaging experiment, I find that a widespread network of frontal, parietal, and ventral occipital brain areas is involved in perceptual stabilization, whereas the reverse effect is only evident in extrastriate cortex. This areal separation possibly endows the brain with the flexibility to switch between exploiting already available information and emphasizing the new. Taken together, my data show that conscious perception and its neuronal correlates display a remarkable degree of flexibility and plasticity, which should be taken into account in future theories of consciousness.

1

GENERAL INTRODUCTION

When we open our eyes in the morning after a good night of sleep, we immediately have the impression of a rich, colorful and detailed world. It feels obvious to us that we perceive everything that surrounds us. But is this actually the case? Scientific evidence suggests that our feeling of seeing everything is actually an illusion. When put to the test, we for example blatantly fail to perceive changes in color and luminance of objects if they are in motion (Suchow & Alvarez, 2011). Similarly, phenomena such as change blindness or inattentional blindness illustrate how poor we actually perform in consciously perceiving even massive changes in our environment (Kim & Blake, 2005). Thus, it seems that we see many things, but many (if not most) things we do not. What happens to the stimuli that we do not perceive? Here, evidence is accumulating that they are not lost at peripheral stages of processing, but make it all the way even up to a semantic analysis (Kouider & Dehaene, 2007). Thus, the mere processing of stimuli does not make the difference between consciously perceiving and not consciously perceiving. What then distinguishes processing from experiencing? This question has puzzled scientists and philosophers since centuries, and lately, the focus has turned to the brain as the most likely place where this distinction arises. In particular, according to present-day belief, there must be a crucial difference between the neuronal processes leading to conscious perception, and those that do not lead to conscious perception (Crick & Koch, 1990). The current research program is to state a working definition of consciousness, to operationalize it in behavioral terms, and then to measure brain activity that differentiates perceiving a stimulus from not perceiving the identical stimulus (Searle, 1998). The hope is that by understanding how brain activity differentiates conscious from unconscious processing, we also understand how consciousness arises from the activity of billions of neurons. This strategy has led to a rapidly growing number of empirical findings (reviewed in Dehaene & Changeux, 2011; Tononi & Koch, 2008) and theoretical proposals (Bachmann, 2007; Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006; Edelman, Gally, & Baars, 2011; Lamme, 2006; Melloni & Singer, 2010; Singer, 2009).

A particularly fruitful line of investigation in the “quest for consciousness” (Koch, 2004) has been to identify the conditions that allow for conscious perception. It is now common sense that we perceive strong stimuli unless we do not pay attention to them, and that we can perceive even weak stimuli if we do pay attention to them. Accordingly, conscious experience is proposed to depend on two factors: the intensity of sensory stimulation and top-down attention (Dehaene, et al., 2006). Are these the only ways to bring stimuli into awareness?

Conscious perception cannot solely depend on sensory evidence. For one, the information entering our brain through the peripheral sensory organs is in itself ambiguous: For example, it is not possible to decide whether an activation pattern on the retina of one eye represents an ellipse or a circle viewed with a certain slant if only the activation pattern itself is considered, because the retina does not represent depth information. On the other hand, stimuli are also often degraded or camouflaged (by themselves), e.g., if an animal tries to hide from its predators. Nevertheless, perception seems effortless and remarkably quick to us. How is this achieved? At the moment, the scientific community has not reached consensus on this issue: While some propose that perception is driven by a rapid feed-forward architecture (Serre, Oliva, & Poggio, 2007), others emphasize the importance of top-down signals and recurrent processing (Gilbert & Sigman, 2007; Hochstein & Ahissar, 2002; Lamme & Roelfsema, 2000). An emergent pattern, however, is an emphasis on previous experience (Engel, Fries, & Singer, 2001; Kveraga, Ghuman, & Bar, 2007). Everyday life and laboratory experiments strongly suggest that perception is greatly facilitated when one has previous experience with the stimulus to be perceived or generally knows what one is looking for. For example, identifying an object in a real world scene is massively impaired if it does not appear in its usual place, e.g., a fire hydrant on a rooftop rather than on the street (Biederman, 1972). From a constructivist point of view, the currently available sensory evidence is in fact always integrated with other, preexisting top-down information (Gregory, 1997). Interestingly, however, this latter aspect is neglected by current theories of consciousness.

Understanding the role of previous experience in conscious perception is the central topic of this thesis. In particular, I investigate how long-term experience (in the form of perceptual learning) and short-term experience (from the previous trial) affect whether we see, how we see, and what the underlying neuronal mechanisms are in the human brain. To unravel how consciousness arises, it is imperative to probe the brain according to this natural way of functioning. In its default mode, the brain integrates currently available information with prior knowledge. However, in the laboratory, we usually create highly artificial conditions with massively degraded visual input with which the subject has no prior experience. We then vary solely the amount of sensory evidence to manipulate visibility. By doing so, we restrict the brain to process sensory input in a bottom-up fashion, and by measuring only once, we don't allow the brain to adapt to this situation through learning. Although the brain is able to solve this challenge and to generate conscious perception without contextual information, a complete understanding of how consciousness comes about requires that the brain is allowed to run free, i.e., to learn and to use prior knowledge. This is because it is not the same to externally provide the brain with more input as it is to study how the brain itself generates appropriate conditions to decipher the visual world. Given that conscious perception can arise under two radically different scenarios, it is also conceivable that the neural correlates of consciousness will differ in these two settings.

SOURCES OF PRIOR KNOWLEDGE

To enable the integration of previous experience with currently available information, the brain is equipped with a variety of plasticity mechanisms, ranging from short-term adaptation to long-term structural changes in grey and white matter. It is now well accepted that the brain can undergo major changes in function and morphology not only during development, but also in the adult state. Accordingly, it has been shown that sensitivity to basic stimulus features (Ahissar & Hochstein, 1998; Fahle, 2009; E. J. Gibson, 1963; Goldstone, 1998), but also various higher cognitive functions such as visuo-spatial attention (C. S. Green & Bavelier, 2003) and working memory (Olesen, Westerberg, & Klingberg, 2004) can be trained in adult subjects. Such learning allows the brain to acquire information on a long time scale, and this affects perception. For example, computer experts that have effortfully learned to categorize personal computers on a specific rather than a general level are less prone to change blindness if the change occurs on a computer, while remaining insensitive to changes of other objects (Archambault, O'Donnell, & Schyns, 1999). But even brief encounters can have long lasting effects: For example, when we are confronted with fragmented black and white images of an object in which the figure cannot be clearly identified (such as the famous picture of the Dalmatian dog by Ronald C. James), we at first fail in perceiving the object. However, once the coherent picture is presented and after explicit recognition has taken place, the object can be clearly seen; importantly, in future encounters, the object will always continue to be perceived, although the currently available information always stays the same. These long-term learning effects are balanced by a variety of short-term effects that endow the brain with the flexibility to react to an ever changing environment. Often, the most informative source of information is what has happened just a moment ago. This is for example evident in the phenomenon of priming. Here, a briefly presented, semantically or otherwise related prime stimulus greatly improves both the speed and accuracy of subsequent target discrimination. Similarly, just holding information in working memory (Soto, Wriglesworth, Bahrami-Balani, & Humphreys, 2010) or being pre-cued to its location (Carrasco, Ling, & Read, 2004) enhances the visibility of stimuli. Finally, on the time scale of phylogenesis, our brain is adapted to integrate incoming information with perceptual principles that seem to arise from the efficient exploitation of the statistics of our environment (Geisler & Diehl, 2002). For example, already newborn babies group stimulus elements according to the Gestalt law of proximity (Farroni, Valenza, Simion, & Umiltà, 2000). This principle, namely that elements that are close to each other belong together, reflects the fact that in visual images parallel contours which are closer to each other are indeed more likely to belong to the same physical object (Brunswik & Kamiya, 1953; also see Geisler, Perry, Super, & Gallogly, 2001; Sigman, Cecchi, Gilbert, & Magnasco, 2001).

Considering that many cognitive functions can be improved through practice and/or previous experience, one can derive the prediction that conscious perception should also be malleable. This contrasts with the sometimes explicit, often tacit view that conscious perception is a static, all-or-none phenomenon (reviewed in Corso, 1963; Eriksen, 1960). In this thesis I will investigate how the integration of information on a moment-by-moment basis and how learning over the course of several days alters and determines conscious perception. These two time scales of previous experience are naturally linked: Short-term experience can be used to form templates which are immediately available and effective to aid conscious perception. If such templates are repeatedly requested and confirmed, they can be permanently laid down in the brain through learning. Furthermore, short-term experience must be integrated with knowledge already acquired through long-term learning. Thus, the creation, testing and updating of knowledge forms a continuous loop in which the most valuable predictions are stabilized in brain structure and function.

DEFINITIONS AND MEASUREMENTS

Before describing the experiments, I will define several terms and outline how I measure conscious perception: I understand conscious perception as the representation of internal or external changes that are accessible to introspection or direct report. Accordingly, I will regard an event as consciously perceived whenever a subject can report the presence or absence of a stimulus or its identity. With regard to measurement, I take the stance that task performance alone cannot be considered an exhaustive measure of awareness. Rather, subjective measures need to be taken seriously (Dehaene & Naccache, 2001; Lau, 2008a; Merikle, Smilek, & Eastwood, 2001; Wiens, 2007). This conviction is based on conceptual and empirical arguments. If one is interested in what can be considered the key aspect of conscious experience, namely the subjective impression of the subject, it is crucial to acquire subjective measures of experience, because only the subject itself has access to its own sensations. Also, it has been demonstrated in carefully designed experiments that subjective and objective measures can dissociate, not only in patients with blindsight (reviewed in Cowey, 2004; Stoerig, 2006), but also in healthy subjects (e.g., Lau & Passingham, 2006). Thus, although we need to be aware of the potential pitfalls of subjective measures (e.g., their sensitivity to response bias), they constitute our prime indicator of whether a stimulus was consciously perceived or not. In contrast, task performance is most parsimoniously explained by processing, but not necessarily experiencing stimuli.

CONSCIOUS PERCEPTION AS A RESULT OF LONG-TERM EXPERIENCE

In Chapters 2 and 3, I will present studies on how perceptual learning changes not only our capacity to process (and thus to discriminate) stimuli, but how it also changes our subjective experience of these stimuli. Perceptual learning is broadly defined as "any relatively permanent and consistent change in the perception of a stimulus array following practice or experience with this array" (E. J. Gibson, 1963). Although perceptual learning is a well-studied phenomenon, the question of how it changes our subjective experience has not been addressed. Do we actually "see more" after training? Research in this field has mostly concentrated on how subjects learn to distinguish between two or more stimulus alternatives that are repeatedly presented over the course of many thousand trials. For example, after sufficient practice, participants' sensitivity to discriminate the alignment of Vernier stimuli significantly increases, even into the hyperacuity range (Poggio, Fahle, & Edelman, 1992). Psychophysical (Karni & Sagi, 1991; Schoups, Vogels, & Orban, 1995) as well as brain imaging work in humans (Schiltz et al., 1999; S. Schwartz, Maquet, & Frith, 2002; Yotsumoto et al., 2009; Yotsumoto, Watanabe, & Sasaki, 2008) and electrophysiological studies in monkeys (Raiguel, Vogels, Mysore, & Orban, 2006; Schoups, Vogels, Qian, & Orban, 2001; Yang & Maunsell, 2004; Zohary, Celebrini, Britten, & Newsome, 1994) point to early visual areas as the most likely stage at which such perceptual learning occurs. However, although such increases in sensitivity have been dubbed "learning to see" (Crist, Li, &

Gilbert, 2001; Sagi & Tanne, 1994), this is somewhat a misnomer, as it is evident that objective performance in a task and awareness of the stimuli on which the task is performed can dissociate: For example, subjects can learn to discriminate the orientation (Seitz, Kim, & Watanabe, 2009) or motion direction (Watanabe, Nanez, & Sasaki, 2001) of stimuli without awareness of the learning material neither during nor after the training. Hence, if a subject learns to discriminate stimuli, it does not immediately follow from this that the stimuli were also experienced.

Here, I explicitly test the relationship of processing and experiencing a stimulus during long-term perceptual learning. In particular, I test whether it is necessary and/or sufficient to see a stimulus in order to discriminate it, whether learning to discriminate and learning to see exhibit the same or different learning curves, and whether they depend on the same circuits in the brain.

Conscious perception and short-term experience

In Chapter 4, I investigate how knowledge acquired on a moment-by-moment basis affects conscious perception. In particular, I study how such knowledge can be used to generate predictions about upcoming events and how this affects the neural correlates of consciousness. Despite the complexity of our visual environment and the fact that visual input tends to be occluded, ambiguous, unclear, or presented from an infinite number of different points of view, we usually have the impression to see well. This might also be because our visual environment contains many sources of prediction: objects stay in the visual field for a relatively long time; they move predictably; even newly appearing objects are mostly predictable due to their context. Predictions then serve the purpose of aiding perception of degraded input. Although theories that highlight the role of predictions in brain functioning are gaining acceptance (Engel, et al., 2001; Friston, 2010; Mumford, 1992; Rao & Ballard, 1999), current theories of consciousness do not explicitly model the influence of such prior knowledge on experience. Here, I study the effect of short-term experience and ensuing predictions by taking advantage of a phenomenon called ‘hysteresis’. ‘Hysteresis’ is a term taken from physics, which denotes the propensity of a dynamical system to remain in a state once it has acquired that state. A classic example in vision, already noted by Helmholtz (1867, p. 728f.), comes from stereopsis: If the two images in the stereoscope are initially placed such that they match and can easily be fused, and subsequently pulled apart, depth can still be perceived even if the monocular images are further apart than the interpupillary distance (Fender & Julesz, 1967). Thus, perceiving the fused images stabilizes the percept against later disruptions of sensory information. To specifically investigate the effect of previous experience on conscious perception and its neuronal correlates, I create hysteresis by gradually revealing a visual stimulus from noise until it is fully visible and subsequently decrease stimulus visibility. Under such conditions, once the stimulus has been recognized, the knowledge of stimulus identity can be used to make predictions about the upcoming stimulus which then interact with bottom-up information and aids perception. Perception is thus stabilized against the reduction in available information.

This paradigm constitutes a straight forward test for models that view conscious perception as the result of evidence accumulation, similar to what has been proposed for decision making (Smith & Ratcliff, 2004). In such models, the decision process constitutes a race between stimulus alternatives, in which evidence is gathered from noisy input until a response criterion has been reached. Providing the system with prior information could in principle lower the decision threshold (Smith & Ratcliff, 2004), accelerate the rate of evidence accumulation (Brown & Heathcote, 2005), or give one of the possible interpretations of the visual input a head start (Carpenter & Williams, 1995). In the context of conscious perception, however, the effects of prior information are not considered. Rather, it is proposed that evidence for a coherent conscious percept is always accumulated from various preconsciously operating modules until it has reached a decision bound (Dehaene, 2008). This decision is then broadcasted throughout the brain, which occurs at the end of evidence accumulation, and thus late. Such late correlates of consciousness have often been observed empirically, occurring around 300 ms post stimulus (Del Cul, Baillet, & Dehaene, 2007;

Gaillard et al., 2009). In fact, some even propose the occurrence of late activity as the sole marker of conscious experience (Lamme, 2006). However, when previous knowledge is taken into account, a straightforward prediction is that the processes leading to conscious perception should speed up accordingly, potentially shifting the neural correlate of consciousness in time. This directly leads to the question whether the neural correlate of consciousness is attached to one particular brain process, or whether this link can vary depending on how conscious experience comes about. Here, I investigate these questions using electroencephalography, because this method offers an exquisite temporal resolution to investigate the sequence of neuronal events leading to a conscious percept.

PREVIOUS EXPERIENCE CAN MAKE CONSCIOUS PERCEPTION GO TWO DIFFERENT WAYS

Admittedly, however, previous experience does not always stabilize perception. Instead, previous experience can have the reverse effect: Seeing the opposite of what was there. In particular, prolonged viewing of a stimulus often leads to so-called repulsive aftereffects. For example, exposure to one direction of motion causes subsequently viewed stimuli to appear moving in the opposite direction, an aftereffect also known as the “waterfall illusion” (Purkinje, 1820). Similar effects can be observed when viewing oriented lines (J. J. Gibson & Radner, 1937), color (Webster & Mollon, 1991), alphanumeric characters (Whitaker & McGraw, 2000), and even faces (Leopold, O’Toole, Vetter, & Blanz, 2001; Webster & MacLin, 1999). Such aftereffects are usually attributed to neuronal adaptation (e.g., Anstis, Verstraten, & Mather, 1998). It is important to note that adaptation is no longer conceived as simple neuronal fatigue, but instead as a computational strategy that serves to optimally code and transmit information (Barlow, 1990; Barlow & Földiák, 1989; Clifford, Wenderoth, & Spehar, 2000; Wainwright, 1999). Furthermore, adaptation is not confined to the single neuron, but constitutes a complex network phenomenon (Felsen et al., 2002; Gutnisky & Dragoi, 2008; Hansen & Dragoi, 2011).

Given that previous experience can apparently make conscious perception go two different ways, stabilizing what we saw before (hysteresis) or seeing the opposite of what was there (adaptation induced aftereffects), I investigate what determines the direction of the effect of previous experience in Chapter 5. To this end, I use multistable stimuli because they allow isolating the effect of conscious experience itself from the effects of mere stimulus processing in hysteresis and adaptation. When multistable stimuli are presented intermittently, perception on a given trial depends on the one hand on which interpretation has been perceived in the trial before, and on the other hand on the amount of stimulus evidence that was provided in the trial before. The former dependency allows studying hysteresis, while the latter allows studying adaptation. I track these two effects in behavior, and, using functional magnetic resonance imaging, in the brain, where I ask the question whether hysteresis and adaptation arise from the same cortical networks. This enables me to address a recent debate on whether hysteresis and adaptation can be explained by a single mechanism (Gepshtein & Kubovy, 2005) or by two mechanisms (Brascamp, Pearson, Blake, & van den Berg, 2009; Noest, van Ee, Nijs, & van Wezel, 2007; Wilson, 2007), and whether what enters consciousness is what is predicted or what was unexpected. More generally, this pertains to the still unresolved question of how the brain maintains the balance between exploiting redundancies and staying sensitive to new information.

A THEORETICAL EMBEDDING

Can the long- and short-term effects of previous experience be summarized under a single theoretical framework? Here, I propose that this is possible using a Bayesian model of perception (for recent reviews, see Geisler, 2008; Vilares & Körding, 2011; Yuille & Kersten, 2006). In this framework, perception is the result of an inference, in which the available sensory information is compared to a prediction derived from previous experience. Importantly, a central assumption is that evidence is not deterministically, but probabilistically related to perception: Given that input is noisy and ambiguous it is the task of our perceptual apparatus to estimate the likelihood by which a certain feature or object was present in a scene. Thus, a neural representation would not be of a feature or object, but an estimate of the probability that this feature or object was present. This calculation takes both the currently available information and all prior information into account, which is why the Bayesian framework is so valuable in understanding the role of previous experience in conscious perception. In formal terms, perception depends on the *probability distribution* of the available evidence on the one hand, and on a *prior* (the expectation that the world will be in a particular state) on the other hand. What we perceive is the result of the application of Bayes' rule (1763) to the likelihood function and the prior, resulting in the so-called *posterior distribution*. Empirical support for this framework stems from studies showing that human perception indeed approximates the behavior of Bayesian Ideal Observers (Fiser, Berkes, Orban, & Lengyel, 2010).

A recent proposal for a neural implementation of the Bayesian framework posits that the brain actively tries to optimize its priors in order to reduce the error that arises when priors and incoming evidence do not match (Friston, 2010; also see Lee & Mumford, 2003; Rao & Ballard, 1999). In particular, it is proposed that within a hierarchical system, higher order brain areas generate predictions (corresponding to the priors), which are tested against the incoming evidence in lower areas. The mismatch (or "prediction error") is signaled back to the higher areas, which then revise their predictions. This naturally entails an online learning process through which priors are continuously updated and refined. Thus, using previous experience on the long- and short-term scale is a crucial ingredient of perception. As will be outlined in detail in Chapter 5, the opposing effects of hysteresis and adaptation can be understood as reflecting changes in the prior and in the likelihood function, respectively. Evidence that the brain indeed performs hierarchical Bayesian inference is accumulating mostly in human brain imaging studies (Alink, Schwiedrzik, Kohler, Singer, & Muckli, 2010; Garrido, Kilner, Stephan, & Friston, 2009; S. O. Murray, Kersten, Olshausen, Schrater, & Woods, 2002; Summerfield, Trittschuh, Monti, Mesulam, & Egner, 2008), but recently also in electrophysiological recordings in animals (Berkes, Orban, Lengyel, & Fiser, 2011).

SUMMARY

In summary, the human brain is an adaptative and active system which modifies the way it processes input through learning by exploiting previous experience and ensuing predictions. This fact is clearly underappreciated by current theories about consciousness and its neuronal correlates. In my thesis, I aim to fill this ostensible gap. In particular, I investigate how previous experience on a long time scale, based on extensive perceptual learning, and how previous experience on a short time scale, i.e., predictions based on information from a previous trial, affect conscious perception at threshold. I ask whether consciousness is malleable by practice, like other cognitive functions, or whether the threshold for conscious perception is fixed. If consciousness is something that we can learn, we might also be able to recover it in case it is lost, for example in patients with brain lesions. I also examine what distinguishes processing from experiencing, and whether they rely on the same brain networks. This directly pertains to the question whether we can localize the neural correlates of consciousness to a particular place in the brain. On a short time scale, I study how previous experience might speed up conscious perception. Here, the question is whether the neural correlate of consciousness can be regarded as an event fixed in time, or whether its occurrence is flexible. Finally, I address why previous experience can make us sometimes more likely and sometimes less likely to see something (again). This serves to tackle the problem of how we see something expected while staying receptive to the unexpected.

2

SENSITIVITY AND PERCEPTUAL AWARENESS INCREASE WITH PRACTICE IN METACONTRAST MASKING

Can practice effects on unconscious stimuli lead to awareness? Can we "learn to see"? Recent evidence suggests that blindsight patients trained for an extensive period of time can learn to discriminate and consciously perceive stimuli that they were previously unaware of. So far, it is unknown whether these effects generalize to normal observers. Here we investigated practice effects in metacontrast masking. Subjects were trained for five consecutive days on the stimulus onset asynchrony (SOA) that resulted in chance performance. Our results show a linear increase in sensitivity (d') but no change in bias (c) for the trained SOA. This practice effect on sensitivity spreads to all tested SOAs. Additionally, we show that subjects rate their perceptual awareness of the target stimuli differently before and after training, exhibiting not only an increase in sensitivity, but also in the subjective awareness of the percept. Thus, subjects can indeed "learn to see".

Schwiedrzik CM, Singer W, Melloni L (2009). Journal of Vision, 9(10):18.1-18.

It is well accepted that the brain is a highly plastic organ that can undergo major changes in function and morphology not only during development, but also in an adult state. This plasticity has been intensively investigated in studies of perceptual learning and it has been shown that the sensitivity to stimulus features can be drastically improved with practice (Ahissar & Hochstein, 1998; Goldstone, 1998). Furthermore, not only basic perceptual skills are subject to practice-dependent changes, but also higher cognitive functions such as visuo-spatial attention (C. S. Green & Bavelier, 2003) or working memory (Olesen, et al., 2004). However, the potential plasticity of another characteristic of the human brain, namely its capacity to produce perceptual awareness (PA), has not been thoroughly investigated. Specifically, it is so far unknown whether practice can render a previously invisible stimulus visible, whether we can "learn to see". In a clinical context, this is an important question, as patients with acquired impairments of conscious perception, such as blindsight patients, might be trained to become conscious again. The very essence of blindsight is a dissociation between chance performance for simple yes-no responses (a subjective measure of awareness) and above chance performance in forced-choice procedures (an objective measure of awareness). Practice can lead to improved performance in forced-choice procedures in human blindsight patients (Bridgeman & Staggs, 1982; Chokron et al., 2008; Henriksson, Raninen, Näsänen, Hyvärinen, & Vanni, 2007; Raninen, Vanni, Hyvärinen, & Näsänen, 2007; Stoerig, 2006; Zihl, 1980; Zihl & Werth, 1984) as well as in monkeys with bilateral ablation of primary visual cortex (V1) (Dineen & Keating, 1981; Humphrey, 1974). Although very high levels of accuracy can be reached in forced-choice tasks, the dissociation between objective and subjective measures often remains unchanged (Sahraie et al., 1997). However, without a concomitant increase in subjective awareness the psychological strain of acquired blindness is not alleviated, and blindsight capabilities are not used in everyday life. A recent systematic study in a sample of 12 cortically blind human subjects who carried out daily discrimination training over a period of three months found not only an increase in sensitivity deep in the blind visual field, but also an increase in reported PA of the stimuli (Sahraie et al., 2006; also see Zihl & von Cramon, 1985). These results are interesting because they indicate that cortical blindness resulting from brain damage is at least partially reversible. Furthermore, the result that - apart from sensitivity - acknowledged awareness of the stimuli increases with practice points to the possibility that not only basic visual functions, but also awareness itself is trainable. If that should indeed be the case, it would be an indication that the "threshold of visual awareness" is not fixed, even in the case of permanent cortical damage. However, it is so far unknown whether such training effects generalize to normal observers.

A common practice in the study of unconscious perception is to use masking in order to present stimuli at or below the threshold of visual awareness. In metacontrast masking (Alpern, 1953), a trailing, non-overlapping mask is used to render a stimulus that precedes the mask invisible. Depending on the task and the stimuli used, a U-shaped function of performance is usually obtained, with a minimum at a positive, non-zero stimulus onset asynchrony (SOA) between stimulus and the subsequent mask (Breitmeyer & Ogmen, 2006, p.43ff). Metacontrast masking has been used to study unconscious perception in normal subjects, for example to induce conditions of "relative" blindsight (Lau & Passingham, 2006), subliminal priming (Neumann & Klotz, 1994; Vorberg, Mattler, Heinecke, Schmidt, & Schwarzbach, 2004), and to investigate the relationship between attention and awareness (Kentridge, Nijboer, & Heywood, 2008). Part of the attraction of metacontrast masking stems from the huge database accumulated over a century of behavioural research (for a recent review, see Breitmeyer & Ogmen, 2006).

Here, we used metacontrast masking to study the effects of training on sensitivity and awareness in normal subjects. We hypothesized that if the information of the target stimulus is available to the visual system although the stimulus is not consciously perceived, and if the "threshold of visual awareness" is not fixed, training should render previously invisible stimuli visible (also see Kanwisher, 2001). To test this hypothesis, we first measured sensitivity as a function of stimulus onset asynchrony (SOA) in a two alternative forced-choice (2AFC) form discrimination task in individual subjects. We then chose the SOA that yielded zero sensitivity and

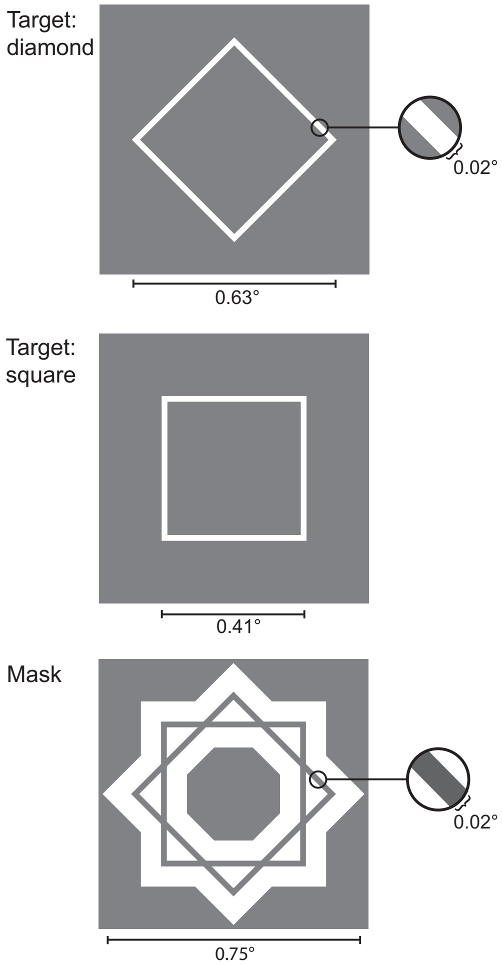

Figure 2-1. Stimuli. A square (0.41°×0.41°) and a diamond (0.63°×0.63°) were used as target stimuli. The outlines of the targets were 0.02° wide. The mask stimulus was a star-shaped figure made up of two squares and two diamonds directly neighboring the positions of the target stimuli on their inside and outside borders (gap width 0.02°).

continued to train our subjects on this SOA for five consecutive days. To test the upper limits of the effects of training, we also trained an individual subject over a period of 24 days. After training, we assessed the subjects' performance over a wide range of SOAs and at a non-trained transfer position. Additionally, we measured the PA of the stimuli that subjects reported before and after training. Using Signal Detection Theory (SDT) we find that subjects not only improve linearly in sensitivity (while their bias is largely constant), but that their reported PA of the stimuli also increases after training. These improvements are still detectable after several months. We thus hypothesize that awareness itself is indeed trainable.

METHODS

Participants

Eight subjects (five male, mean age 24, range 22-27) participated in the main experiment lasting five sessions, and one additional subject (female, age 23) was trained over a period of 24 sessions. All had normal or corrected-to-normal vision and no history of neurological and/or psychiatric disease. Handedness was assessed using the Edinburgh Inventory (Oldfield, 1971). The mean laterality quotient was H=74.49 (range 58.3-91.6), indicating that all nine subjects were right-handed. All subjects gave written informed consent.

Stimuli

Stimuli were displayed on a CRT monitor (Sony GDM 5502, resolution 1024×768) at a refresh rate of 100 Hz. Subjects viewed the screen from 1 m distance. The visible screen size was 21.05°×16.17°. Stimulus presentation and response recording were controlled by Presentation software (version 10.3, www.neurobs.com), running under Microsoft Windows 98. All stimuli were presented on a dark gray background (6.26 cd/m^2). A white fixation cross (53 cd/m^2, 0.56°×0.56°) was displayed at the center of the screen.

Stimuli were presented at 3.2° either above or below fixation (position was counterbalanced across subjects). A square (0.41°×0.41°) and a diamond (0.63°×0.63°) were used as target stimuli. The outlines of the targets were 0.02° wide and had a luminance of 34.42 cd/m^2. The mask stimulus was a star-shaped figure made up of two squares and two diamonds directly neighboring the positions of the target stimuli on their inside and outside borders. The mask outlines had a width of 0.065° on the outside and a width between 0.043° and 0.086° on the inside (with regard to the stimuli). The mask had the same luminance as the target stimuli. Enlarged versions of targets and mask can be seen in Figure 2-1.

A combination of para- (forward) and metacontrast (backward) masking has been shown to be especially effective with regard to masking (Macknik & Livingstone, 1998). We thus used such a combination to mask our target stimuli in order to create conditions where we would obtain a *d'*=0 and to avoid potential ceiling effects. Each trial started with a fixation period of variable length (between 1000 and 1500 ms). Subsequently, the paracontrast mask was presented for 50 ms. After a fixed inter-stimulus interval (ISI) of 30 ms, the target was presented for 10 ms. The metacontrast mask was presented for 50 ms at stimulus onset asynchronies (SOA) ranging from 20 ms to 160 ms (20, 40, 60, 80, 100, 120, 140, 160 ms) for the threshold estimation or at an individually determined SOA during the training sessions. Responses were recorded only after mask offset.

Procedure

The main task was to discriminate the two target stimuli (square from diamond) in a forced-choice procedure: Subjects were instructed to quickly press the button "S" on a keyboard whenever they saw the diamond and to press the button "L" whenever they saw a square. Accuracy and not speed was emphasized. Subjects were also informed that both stimulus alternatives would appear in a randomized order but with equal probabilities and that there were no catch trials. In addition to the discrimination task, subjects rated their PA of the stimuli on a three-point-scale (see below).

Subjects were asked to maintain fixation on the center of the screen throughout the experimental sessions. Instructions were given in both verbal and written form. The experiments were conducted in a darkened room. Constant head position was assured by the use of a chinrest with forehead support.

The experiment took place over a period of five consecutive days. On the first day, we assessed for each subject the SOA that would yield the maximal masking in a typical threshold experiment. This SOA was then used for the training. The first training session was conducted directly after the assessment of the critical SOA on the first day. On days two to four, only training sessions were conducted. On the fifth and last day, the last training session was administered.

After this last training session, we again assessed the masking threshold, followed by the transfer task. Each part of the experiment will be described in detail below.

Threshold estimation

Before threshold estimation, subjects carried out practice trials of "slow motion" versions of the upcoming experimental trials in order to familiarize them with the stimuli and with the task. For the threshold estimations, we presented target stimuli at SOAs ranging from 20 ms to 160 ms. Each target was presented 40 times at each SOA, yielding 80 trials per SOA and a total of 640 trials. After every 160 trials, we introduced a break of variable length. The occurrence of SOAs was randomized but counterbalanced over blocks. The sequence of target stimuli was fully randomized and no feedback was given.

In addition to target discrimination, subject rated their PA of the respective stimulus on a three-point-scale after every trial. To this end, we presented a screen with the question "How visible was the stimulus? Invisible / unclear / clearly visible" after the subject's forced-choice response. Subjects were instructed to respond with key presses for "invisible" if they did not see the stimulus, for "unclear" if they saw something but could not identify it, and for "clearly visible" if they unmistakably saw the stimulus. These questions were not aimed at the subjects' confidence in their reports, but at their phenomenal impression of the stimulus. Thus, our scale differs from previously used scales that either assessed only confidence (Persaud, McLeod, & Cowey, 2007; Wilimzig, Tsuchiya, Fahle, Einhäuser, & Koch, 2008), or a compound of confidence and PA (Wessinger, Fendrich, & Gazzaniga, 1999). This is an important distinction, because it has been shown that confidence and awareness are not equivalent measures and recover differently after repeated exposure in blindsight patients (Sahraie, Weiskrantz, & Barbur, 1998). Our operationalization of PA was discussed with the subjects before the first threshold estimation. It was further pointed out to the subjects that there were neither right nor wrong answers to the question and that they should rate only their perception of the target stimulus and not the whole stimulus sequence or the mask stimulus. The threshold estimation took approximately 1 h. Thresholds and PA were estimated before the first and after the last training session.

Training sessions

From the threshold estimation, we extracted the SOA with the lowest performance individually for each subject. This SOA was then used in the training sessions. Data from the PA rating was not used to choose the training SOA, and no rating of PA was required during training. Subjects engaged in the forced-choice discrimination task for 600 trials per training session (a total of 3000 trials). After every 100 trials, a break was introduced. Subjects received feedback (correct / incorrect) after every response as well as after every block of 100 trials (percentage correct). Correct / incorrect feedback was displayed at fixation for 1000 ms in green letters (font Arial, font size 20 pt, 37.68 cd/m^2) or red letters (9.05 cd/m^2), respectively. A training session lasted approximately 30 min.

To ensure stable high motivation across subjects, we developed a pay-off scheme based on monetary reward. Subjects would receive €5 for each of the training sessions. If they improved in their performance by at least 10%, they could earn an additional €2 per day. However, if they showed no gain or even a decrease in performance as compared to the previous day, they would lose €2. Performance levels (in percentage correct) were written down after each training session and subjects were reminded of their previous performance as well as the pay-off scheme before each training session. Also, subjects were reinstructed about the discrimination task.

Transfer task

On the last day of the experiment, we assessed whether the exercise on the trained stimulus position would transfer to another position. To this end, we switched the stimulus position to its mirror position above or below the fixation cross. For instance, if a subject had been trained on the position above the fixation cross, the new stimulus position would now be below the fixation cross. Subjects had no prior knowledge that the stimulus position would be switched. We employed the same SOAs and the same procedure as during the other training sessions (see "Training sessions"). Subjects could earn an additional €2 if they reached at least the performance of the last session on the originally trained position.

Retest

During the retest session, which took place between five and ten months after the last training session, we again assessed objective and subjective thresholds, followed by a test of sensitivity at the transfer position (procedure as for the initial transfer task). Subjects received 15€ for their participation in the retest.

Analysis

In order to calculate *d'* and *c* values, square trials were considered signal trials and diamond trials were considered noise trials (Wickens, 2002, p.114). For the threshold data, this yielded 40 signal and 40 noise trials per SOA. We calculated *d'* and *c* for each SOA independently, thus the values indicate discriminability of signal and noise trials at a given SOA and not discriminability between SOAs. For the training data, 300 signal and 300 noise trials were available per session. To correct for extreme false alarm or hit rate proportions, we used the loglinear correction. In this approach, 0.5 is added to the number of hits and false alarms, and 1 is added to the number of signal trials and the number of noise trials. The loglinear correction reduces bias in the calculation of *d'* when corrections have to be applied in order to avoid infinite *z* score values (Hautus, 1995).

2AFC tasks are usually considered to be bias-free. However, this assumption does not always hold true (Macmillan & Creelman, 2005, p.170ff). Furthermore, it has been shown that perceptual learning can lead to changes in sensitivity as well as bias (Wenger, Copeland, Bittner, & Thomas, 2008; Wenger & Rasche, 2006; also see Seitz, Nanez, Holloway, Koyama,

& Watanabe, 2005). Thus, we included a measure of bias into our analysis. For the calculation of bias, three possible measures can be calculated: the criterion location *c* (the distance between the criterion and the point at which the signal over the noise distribution cross), the relative bias *c'* (the criterion location relative to the corresponding *d'*), and the likelihood ratio β (the ratio of the heights of the signal and the noise distribution). We chose to calculate *c* for our purposes, because it is the only measure of bias that is orthogonal to but associated with sensitivity (*d'*) and fulfills the monotonicity condition at and below chance levels (Macmillan & Creelman, 1990). Furthermore, the range of *c* is similar to that of *d'*; this eases the comparison of the change in the two measures.

Data from the PA rating could not be analyzed in a SDT framework, because we did not in-

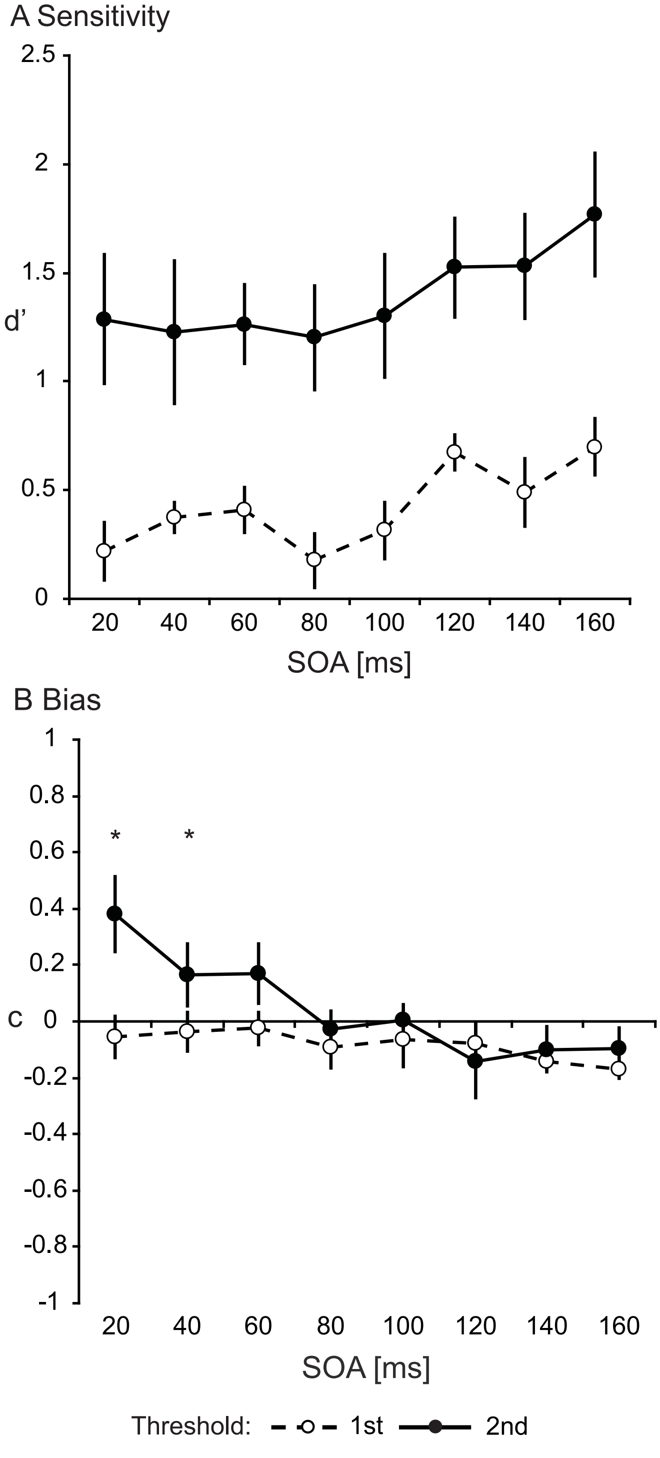

Figure 2-2. Average sensitivity (*d'*) and bias (*c*) from threshold 1 and threshold 2. Errors bars represent the standard error of the mean.

clude catch trials in our paradigm. Thus, given that the PA rating aimed at subjective impression of both the square and the diamond target stimuli, we did not have noise trials in this measure. Further, the data could not be analyzed as a type 2 task as the decision axis was not the decision in the 2AFC task but the stimulus itself (Galvin, Podd, Drga, & Whitmore, 2003). We instead calculated mean PA rating values at each SOA for correct responses, errors, and pooled correct responses and errors (see below).

In all analyses of variance (ANOVA) with more than one degree of freedom we used the Greenhouse-Geisser correction (McCall & Appelbaum, 1973). We report adjusted degrees of freedom and adjusted p values.

RESULTS

Discrimination performance

First threshold

Metacontrast masking yields either functions that show a linear relationship to SOA (type A masking) with maximal masking at a SOA of 0 ms or U-shaped functions of SOA (type B masking) with a maximal effect at a time point greater than 0 ms. Usually, these functions are obtained from the hit rates. In order to assess whether our combination of para- and metacontrast masking yielded comparable functions in our paradigm with combined forward and backward masking, we first analyzed d' and c values in repeated measures ANOVAs with the factor SOA (20, 40, 60, 80, 100, 120, 140, 160 ms). For d', we indeed found a significant effect of SOA ($F(3.394, 23.757)=5.620$, $p=0.0036$, $\eta^2=0.445$), whereas for c, the effect of SOA did not reach significance ($F(4.003, 28.018)=1.178$, $p=0.342$, $\eta^2=0.144$). This indicates that the experimental manipulation of the SOA led to changes in sensitivity, but not to changes in bias. As can be seen from Figure 2-2, both d' and c vary with SOA in a complex way. To ease the interpretation of the results, we aligned the SOAs to the SOA at which the individual subject was subsequently trained. We then centered our analysis on the trained SOA per subject and further included one neighboring SOAs above and below, respectively. Five subjects were trained at a SOA of 80 ms, two at a SOA of 100 ms, and one at a SOA of 60 ms. As can be seen from Figure 2-3, the reordering yielded a clearly U-shaped function of SOA for d', but not for c.

An additional ANOVA of this reduced sample of SOA (untrained below, trained, untrained above) confirmed the significant effect of SOA for d' ($F(1.513, 10.591)=11.186$, $p=0.004$, $\eta^2=0.615$) and its non-significant effect for c ($F(1.504,10.526)=0.378$, $p=0.636$, $\eta^2=0.051$). Additionally, a within subjects contrast showed a significant quadratic trend of SOA for d' ($F(1, 7)=50.178$, $p=0.0002$, $\eta^2=0.878$). Thus, the obtained functions of sensitivity are comparable to U-shaped type B metacontrast masking functions.

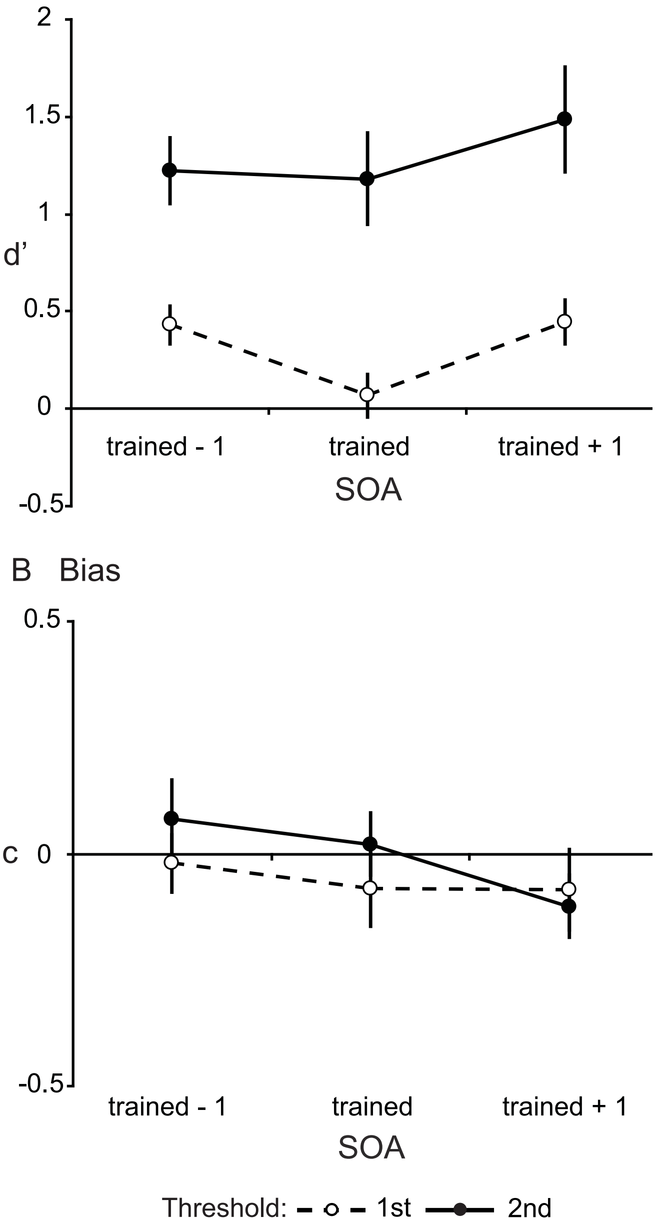

Figure 2-3. Sensitivity (d') and bias (c) from threshold 1 and threshold 2, individually aligned to the trained SOA. Errors bars represent the standard error of the mean.

To make sure that subjects were trained on a SOA that was at the objective threshold of conscious perception, we ran a two-sided *t*-test of their sensitivity values at that SOA against 0. We indeed found that the d' of the trained SOA was not significantly different from 0 (mean difference 0.067, $T(7)=0.620$, $p=0.555$), indicating that subjects were objectively unaware of the stimuli. By contrast, the sensitivity at the two neighboring untrained SOAs was significantly higher than 0 (below: mean difference 0.4323, $T(7)=4.448$, $p=0.003$; above: mean difference 0.4472, $T(7)=3.907$, $p=0.0058$). Thus, at these SOAs the subjects were objectively aware of the stimuli.

A two-sided t-test of the corresponding values of c against 0 was also not significant at any of the three SOAs (all $p>=0.3842$). This indicates that subjects were not biased to respond with either response alternative at these SOAs, as would be expected from a 2AFC task with equal probabilities for the two response alternatives.

Training sessions

To assess whether subjects improved in their performance throughout the training sessions, we entered *d'* and *c* values of the five training sessions into a repeated measures ANOVA. For *d'*, we found a significant effect of session ($F(1.824, 12.767)=15.578$, $p=0.0004$, $\eta^2=0.690$); a within subjects contrast yielded a significant linear trend ($F(1, 7)=21.737$, $p=0.0023$, $\eta^2=0.756$) and a quadratic trend that approached significance ($F(1, 7)=5.214$, $p=0.0564$, $\eta^2=0.427$). Figure 2-4 shows that subjects improved linearly in their sensitivity in the course of the training sessions, with a trend towards reaching a plateau for the two final sessions. Sensitivity increased on average by a *d'* of 1.16 (*SD* 0.74) from the first to the fifth session.

For *c*, we did not find a significant effect of session ($F(3.283, 22.983)=0.852$, $p=0.488$, $\eta^2=0.109$). In fact, none of the *c* values obtained for the individual sessions differed significantly from 0 (two-sided *t*-tests, all $p>=0.2655$, uncorrected). Thus, while subjects improved linearly in sensitivity, their bias to respond with one or the other response alternative did not change during the training (also see Figure 2-4).

The average *d'* for the first training session was 0.51 (*SD* 0.32) and significantly different from 0 ($T(7)=4.470$, $p=0.0029$). To address the question of when during the training session the subjects crossed the objective threshold, we split up the first session into six blocks

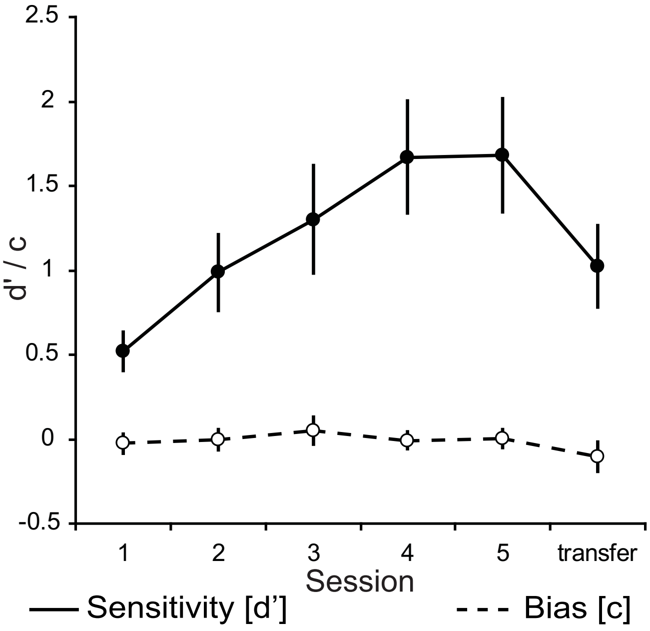

Figure 2-4. Sensitivity (*d'*) and bias (*c*) over the training sessions and for the transfer task. Errors bars represent the standard error of the mean.

of 100 trials (50 signal trials and 50 noise trials) each. After applying the loglinear correction, we tested *d'* for each block against 0 using two-sided paired *t*-tests. Only the first block was not significantly different from 0 (mean difference 0.3614, $T(7)=2.219$, $p=0.0708$). However, when two subjects were excluded that had a $d'>1$ already in the first block, we found that for the remaining subjects, the first as well as the second block were not significantly different from 0 (first block: mean difference 0.1278, $T(5)=1.272$, $p=0.2594$; second block: mean difference 0.3009, $T(5)=1.740$, $p=0.1424$). Thus, the subjects crossed the objective threshold early during the first session.

Perceptual learning is often characterized by an initial rapid learning phase that is later followed by slow, sustained learning. To characterize the learning during the first session, we analyzed the average *d'* of all subjects in each of the six blocks by means of a linear regression. We found that on average, *d'* increased linearly not only over all sessions, but also within the first session, ranging from $d'=0.36$ (*SD* 0.48) in the first block to $d'=0.59$ (*SD* 0.38) in the sixth block. Block number explained a significant proportion of variance in *d'*, $R^2=0.6601$, $F(1,4)=7.7677$,

$p<0.05$. When we excluded the two subjects that had a $d'>1$ in the first block, similar results were obtained ($R^2=0.7906$, $F(1,4)=15.0987$, $p<0.02$). Here, the average d' ranged from $d'=0.13$ (SD 0.25) in the first block to $d'=0.49$ (SD 0.40) in the sixth block. Linear regressions for the remaining sessions 2-5 were all non-significant (all $p>0.083$). Thus, learning was characterized by a steep linear increase in sensitivity during the first session, followed by a more variable, gradual increase over the remaining sessions.

Second threshold

After the subjects had significantly improved in their sensitivity to discriminate the two stimuli during the training sessions, we now tested how this improvement affected the objective threshold as a function of SOA. We hypothesized that the training effect could either be specific to the trained SOA, affect all SOAs, or affect only a subsample of the SOAs. Thus, we analyzed the performance in the second threshold estimation for both the full range of SOAs as well as the SOAs aligned to the trained SOA.

For the full range of SOAs, we ran a repeated measures ANOVA with factors session (first, second) and SOA (20, 40, 60, 80, 100, 120, 140, 160 ms), again for both the d' and the c values. For d', we found a significant main effect of session ($F(1, 7)=27.311$, $p=0.0012$, $\eta^2=0.796$), and a significant main effect of SOA ($F(3.121, 21.849)=6.467$, $p=0.0024$, $\eta^2=0.480$), but no significant interaction ($F(3.115, 21.805)=0.359$, $p=0.7903$, $\eta^2=0.049$). Thus, training led to a significant improvement in the forced-choice discrimination task that affected all SOAs, and not only the trained SOA (see Figure 2-2).

The same analysis for c yielded a significant main effect of session ($F(1, 7)=8.084$, $p=0.0249$, $\eta^2=0.536$), a significant main effect of SOA ($F(3.535, 24.746)=8.090$, $p=0.0003$, $\eta^2=0.536$), as well as a significant interaction of session and SOA ($F(4.004, 28.027)=3.204$, $p=0.0275$, $\eta^2=0.314$). Post-hoc two-sided t-tests (uncorrected) revealed that the bias changed significantly with training only for the two shortest SOAs (20 ms: mean difference -0.4375, *T*(7)=-3.628, *p*=0.0084; 40 ms: mean difference -0.2030, *T*(7)=-2.512, *p*=0.0403; all other *p*>= 0.0651). Hence, for the two shortest SOAs, subjects became more biased to respond with "diamond" after the training (also see Figure 2-2b).

Following our approach from the analysis of the first threshold, we again aligned the SOAs with the trained SOA and analyzed these data with repeated measures ANOVAs for both *d'* and *c*. For *d'*, we confirmed the significant main effect of session (*F*(1, 7)=30.553, *p*=0.0009, η^2=0.814) and a significant main effect of SOA (*F*(1.343, 9.402)=7.702, *p*=0.0158, η^2=0.524), but no interaction between session and SOA (*F*(1.640, 11.480)=2.301, *p*=0.1494, η^2=0.114). As can be seen from Figure 2-3a, the obtained function of SOA moves up and changes its shape, which indicates that no significant masking took place anymore after the training. For c, we found no significant effect or interaction (all *p*>=0.0907).

Post-hoc t-tests revealed that the *d'* values for the trained SOA indeed significantly differed before and after the training (mean difference -1.114, *T*(7)=-5.122, *p*=0.001, two-sided, uncorrected), and that the *d'* for the trained SOA was significantly different from 0 after the training (mean difference 1.1821, *T*(7)=5.009, *p*=0.0007, one-sided, uncorrected). The corresponding estimate of *c* did not differ between the two thresholds (mean difference -0.0964, *T*(7)=-1.214, *p*=0.264, uncorrected, two-sided), nor did it differ from 0 after training (mean difference 0.0209, *T*(7)=0.318, *p*=0.760, uncorrected, two-sided). The differences in *d'* and *c* for the trained SOA between the first and second threshold are displayed in Figure 2-5 for each subject individually.

Transfer task

In order to evaluate whether the training on one stimulus position showed transfer to another stimulus position, we compared *d'* and *c* of the transfer task with the results from the first and last training sessions in paired two-sided *t*-tests. We reasoned that if the results from the transfer task were significantly lower than the results from the last session and lower than or

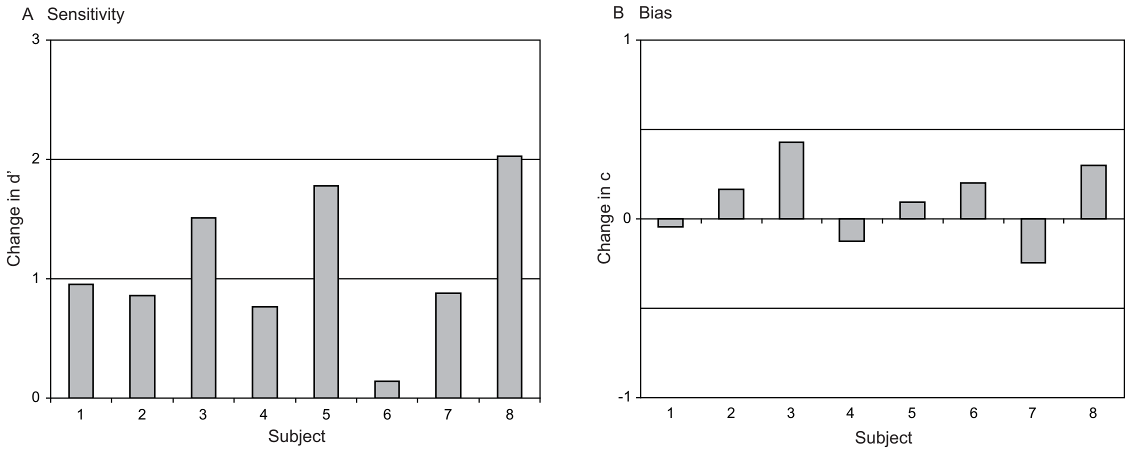

Figure 2-5. Differences in sensitivity (*d'*) and bias (*c*) for the trained SOA between threshold 1 and threshold 2 per subject.

equal to the results from the first session, this would argue for no transfer to an untrained stimulus position. However, if the results from the transfer task were equal or higher than the results from the last session, this would point towards independence of the training effects from the stimulus position. In the case of partial transfer, we would expect the results from the transfer task to lie somewhere between the results from the first and last training session.

The average *d'* reached in the transfer task was 1.02 (*SD* 0.68), while the average *c* was -0.10 (*SD* 0.24). This indicates that for *d'*, the results from the transfer task were significantly better than the results from the first training session (mean difference 0.5068, *T*(7)=-2.895, *p*=0.0232), while there was no significant difference between the transfer and the fifth training session (mean difference -0.6579, *T*(7)=2.219, *p*=0.0620). Upon further inspection of the data, we found that the individual subjects' differences between the *d'* values from the fifth session and the transfer session were negative in six out of eight cases. In the two remaining cases, the differences were positive (0.34 and 0.30). When we excluded these two outliers, we found that the remaining subjects performed significantly worse in the transfer task as compared to the last training session (mean difference -0.9854, *T*(5)=3.522, *p*=0.0169). Thus, most subjects showed partial but not full transfer of the training effects to the untrained stimulus position. Data for all subjects is displayed in Figure 2-4.

For *c*, a two-sided t-test against 0 for all subjects revealed that the bias was not significantly different from 0 in the transfer task (mean difference -0.1043, $T(7)=-1.210$, $p=0.2655$). Furthermore, the differences between the bias in the transfer task and the bias in the first and in the last session were both not significantly different from 0 either (first session: mean difference -0.0776, $T(7)=-0.854$, $p=0.4212$; fifth session: mean difference -0.1074, $T(7)=-1.209$, $p=0.2660$). This indicates that a change of the stimulus position did not lead to a significant change in bias.

Perceptual awareness rating

For the PA rating, our analyses aimed at three questions: 1. Do the PA ratings follow the same or a similar function of SOA as the objective measures that we obtained in the two alternative forced-choice task? 2. Do PA ratings increase as a result of training? 3. Do the PA ratings differ depending on accuracy? We reasoned that the mean PA rating for correct responses should be higher than the mean PA rating for errors before as well as after the training if there was a genuine increase in PA, whereas no difference in the PA rating for correct and incorrect responses would indicate biasing effects such as over- or underconfidence. These three questions where first addressed in a repeated measures ANOVA with factors session (first, second), accuracy (correct responses, errors), and SOA (20, 40, 60, 80, 100, 120, 140, 160 ms). We found a significant main effect of session ($F(1, 7)=7.869$, $p=0.0263$, $\eta^2=0.529$), a significant main effect of accuracy ($F(1, 7)=18.730$, $p=0.0034$, $\eta^2=0.728$), and an effect of SOA that only approached significance ($F(2.770, 19.391)=3.021$, $p=0.0580$, $\eta^2=0.301$). Furthermore, we found a significant interaction of session and accuracy ($F(1, 7)=21.890$, $p=0.0023$, $\eta^2=0.758$) as well as a significant interaction of accuracy and SOA ($F(3.524, 24.669)=3.509$, $p=0.0249$, $\eta^2=0.334$), but no interaction of session and SOA ($F(3.662, 25.635)=1.122$, $p=0.3652$, $\eta^2=0.138$), and no interaction of session, accuracy, and SOA ($F(3.686, 25.805)=0.729$, $p=0.5702$, $\eta^2=0.094$). These results indicate that the mean PA ratings do not differ between SOAs, but that they do

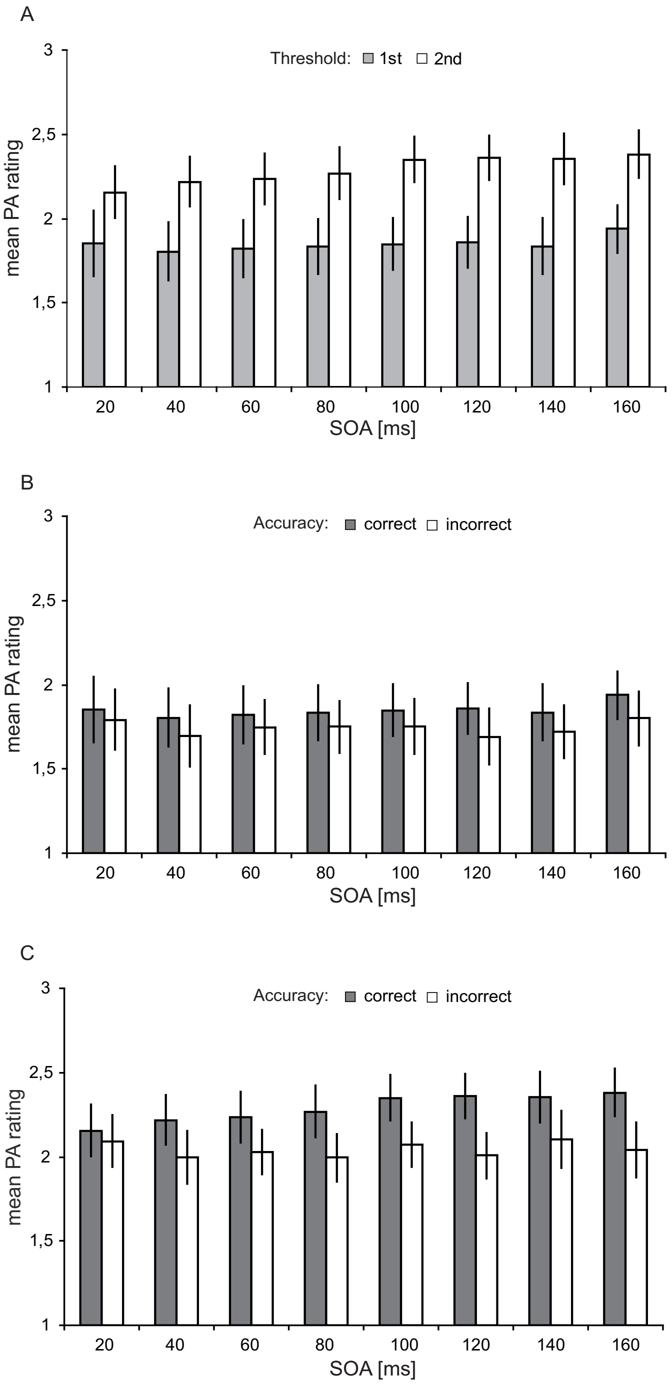

Figure 2-6. Mean perceptual awareness ratings from threshold 1 and threshold 2 by accuracy. (a) Mean perceptual awareness rating for correct responses from threshold 1 and threshold 2. (b) Mean perceptual awareness rating for correct responses and incorrect responses from threshold 1. (c) Mean perceptual awareness rating for correct responses and incorrect responses from threshold 2. Errors bars represent the standard error of the mean.

differ before and after training, and that this difference is modulated by accuracy. To further elucidate these results, we first split up our data set by session in order to investigate whether the mean PA rating was higher for correct than for incorrect responses both before as well as after the training. This was confirmed for the data from the first threshold by means of a repeated measures ANOVA with factors SOA (20, 40, 60, 80, 100, 120, 140, 160 ms) and accuracy (correct responses, errors). We found a significant effect of accuracy ($F(1, 7)=9.321$, $p=0.0185$, $\eta^2=0.571$), but no effect of SOA ($F(3.574, 25.015)=2.107$, $p=0.1158$, $\eta^2=0.231$), and no interaction ($F(3.551, 24.858)=0.771$, $p=0.5406$, $\eta^2=0.099$). Thus, the PA rating was sensitive to variations in accuracy at each SOA before the training, showing higher PA for correctly identified stimuli than for incorrectly identified stimuli (mean difference 0.1056, *SE* 0.35). The same analysis for the PA scores of the second threshold yielded identical results: We found a significant effect of accuracy ($F(1, 7)=23.6$, $p=0.0018$, $\eta^2=0.771$), but no effect of SOA ($F(3.891, 27.237)=1.858$, $p=0.1480$, $\eta^2=0.210$), and no interaction ($F(3.177, 22.236)=2.352$, $p=0.0967$, $\eta^2=0.252$). Thus, the PA scores retained their sensitivity for accuracy even after the training, with higher scores for correct responses than for errors (mean difference 0.2482, *SE* 0.51). The results from the first and second threshold are illustrated in Figure 2-6 b and c.

Based on the previous analyses, we now continued to split up the PA ratings by accuracy for further analyses. In order to answer the question whether subjects become more aware of the stimuli after the training we analyzed the mean PA rating only for correct responses in a repeated measure ANOVA with factors session (first, second) and SOA (20, 40, 60, 80, 100, 120, 140, 160 ms). We found a significant effect of session ($F(1, 7)=10.459$, $p=0.0144$, $\eta^2=0.599$) as well as a significant effect of SOA ($F(2.817, 19.719)=7.729$, $p=0.0015$, $\eta^2=0.525$), and no interaction between these two factors ($F(2.966, 20.760)=2.672$, $p=0.0745$, $\eta^2=0.276$). These results indicate that subjects became more aware of the stimuli that they also correctly identified (see Figure 2-6a). This seems to be the case not only for the trained SOA, but also for the untrained SOAs, thus paralleling the findings from the forced-choice task. To further elucidate the origin of the increase in PA for correct responses, we also calculated the proportion of "invisible", "unclear" and "clearly visible" responses for correct responses for the first and the second threshold at the trained SOA. To this end, we normalized the amount of the respective PA ratings by the amount of hits and compared these values for the two thresholds. As can be seen from Figure 2-7, the number of "unclear" correct responses did not change significantly from the first to the second threshold (mean difference 0.0557, $T(7)=0.5340$, $p=0.6099$), while the number of "invisible" responses dropped to almost the same extent (mean difference 0.2044, $T(7)=2.4506$, $p=0.0441$) as the number of "clearly seen" responses increased (mean difference 0.26, $T(7)=2.8121$, $p=0.0261$).

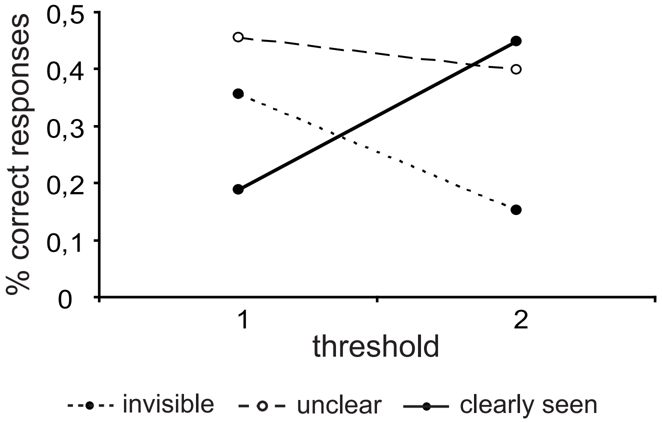

Figure 2-7. Proportion of perceptual awareness ratings for hits from threshold 1 and threshold 2.

We also conducted a repeated measure ANOVA with factors session (first, second) and SOA (20, 40, 60, 80, 100, 120, 140, 160 ms) for erroneous responses. A higher awareness for errors would be difficult to interpret, but can be taken as an indication of a more liberal use of the scale or overconfidence after the training. We only found an effect of session that approached significance ($F(1, 7)=5.361, p=0.0538, \eta^2=0.434$). Thus, the PA for incorrectly identified stimuli did not change significantly, however, an effect of overconfidence with regard to the PA rating cannot fully be ruled out, given the statistical trend in the data for errors.

Retest

Four of the eight trained subjects could be retested between five and ten months after their last training session. For the threshold performance, we compared *d'* and *c* between the three data sets over all SOAs at the single subject level. At a significance level of $p=0.05$, we found the sensitivity over all SOAs to be significantly lower in the first threshold session than in the retest in all four subjects, while the sensitivity in the second threshold session was significantly higher than or equal to the sensitivity reached in the retest in three out of four subjects. This indicates that the subjects retained some but not all of their training effects. The criterion *c* was not significantly different from 0 in three out of four subjects. For the transfer position, we found neither sensitivity nor criterion to be significantly different from the previous assessment (*d'*: mean difference -0.1607, $T(3)$=-1.681, p=0.1913; *c*: mean difference -0.0367, $T(3)$=-0.2694, p=0.8041). For the PA rating, we found that the mean PA rating for correct responses over all SOAs was significantly lower in the first threshold session than in the retest in three out of four subjects. The mean PA rating for correct responses was significantly higher in the second threshold than in the retest in three subjects as well. However, the mean PA rating for correct responses was significantly higher than the mean PA rating for errors in only two subjects, while it was significantly lower for one subject, and not significantly different for another subject. Thus, although the subjects preserved high sensitivity levels even without further practice, the differential effect of training for correct responses and errors on the PA rating was not fully retained in all subjects after a period of five to ten months.

Individual subject

We also trained an individual subject for 24 consecutive sessions (14400 trials) to obtain an approximation of the upper limit of the training effects. Figure 2-8 shows the results of the first and second threshold estimations at the upper stimulus position. The SOA yielding the lowest *d'* value in the first threshold was 120 ms (*d'*=0.06, *c*=0.03). The subject was subsequently trained on this SOA. The *d'* reached in the first session was 1.05 (*c*=0.14), and 2.48 in the last session (*c*=0.19). We used a linear regression to analyze the development of *d'* and *c* during the course of training. We found that session number explained a significant proportion of variance in *d'*, R^2=0.908, $F(1,22)$=216.960, $p<0.0001$. Session number also explained a significant proportion of variance in *c*, R^2=0.422, $F(1,22)$=16.090, p=0.001. However, session number explained only about half of the variance in *c*, while it explained most of the variance in *d'*. Also, *d'* showed substantial improvement, reaching from 1.05 to 3.04 (range 1.99). At the same time, *c* stayed close to the unbiased 0 point, reaching from -0.06 to 0.30 (range 0.37). We interpret these results as to be in agreement with the results from the bigger sample trained for five con-

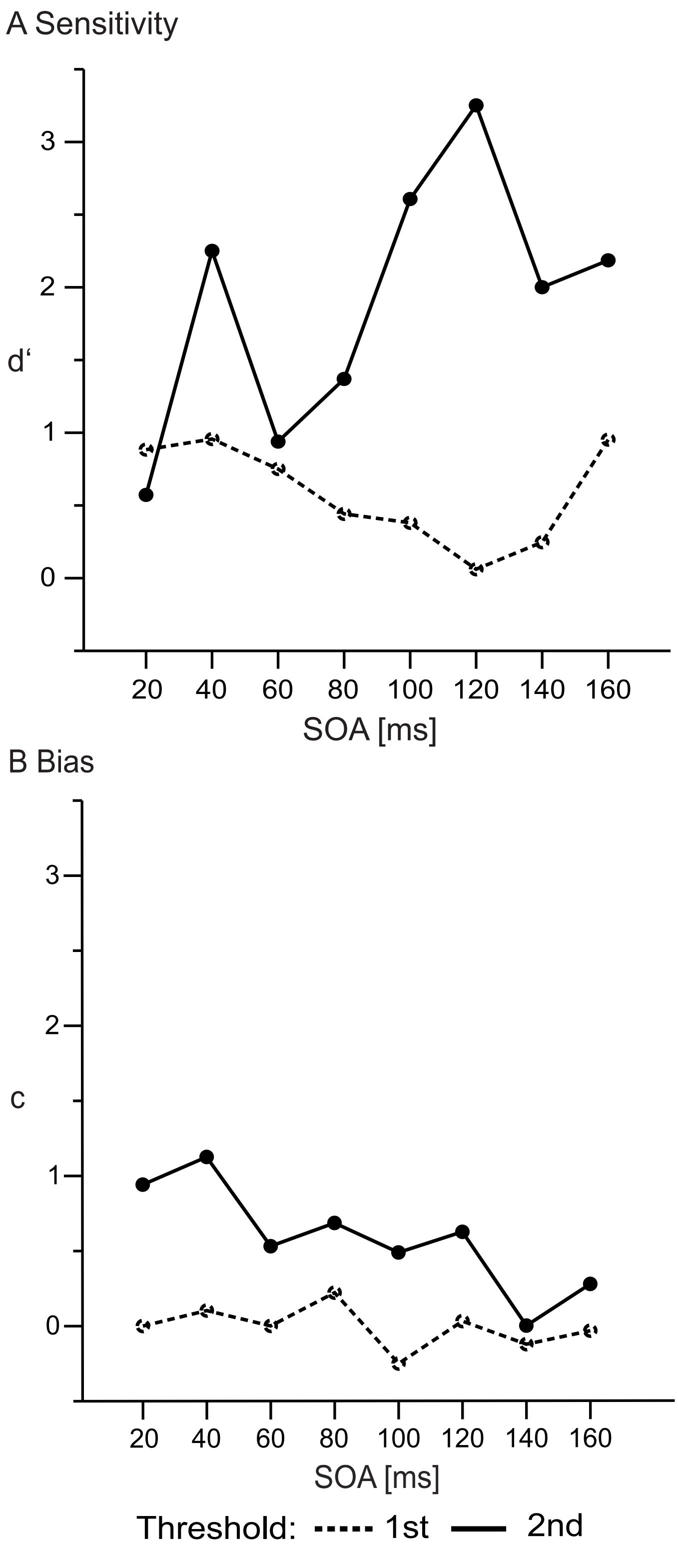

Figure 2-8. Sensitivity (*d’*) and bias (*c*) from threshold 1 and threshold 2 for an individual subject trained for 24 sessions.

secutive sessions: While there is a significant linear improvement of sensitivity, the bias shows little or no systematic change in the course of training (see Figure 2-9).

As can be seen from Figure 2-8a and in agreement with the results from the bigger sample, we found an increase in sensitivity (*d’*) at almost every SOA for the second threshold. This increase was especially pronounced for the trained SOA (first threshold *d’*=0.06, second threshold *d’*=3.25). The subject also became more biased to respond with “diamond”: Whereas the criterion at the trained SOA was 0.03 for the first threshold, it increased to 0.62 for the second threshold (see Figure 2-8b).

For the transfer, *d’* dropped to 0.62 (*c*=0.37). The results from the transfer task were lower than the results from the last session (*d’*=2.48) and lower than the results from the first session *d’*=1.05). This

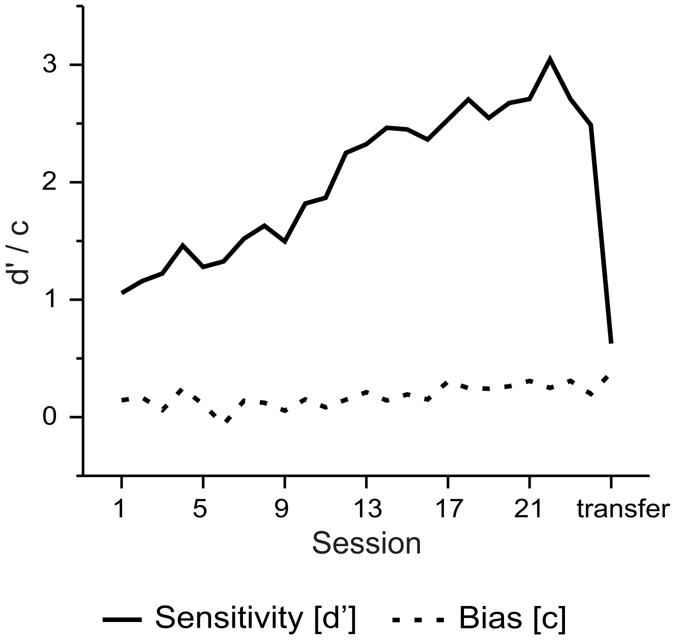

Figure 2-9. Sensitivity (*d'*) and bias (*c*) over the training sessions and for the transfer task for an individual subject trained for 24 sessions.

result points towards no transfer of the increased sensitivity to the untrained stimulus position for this subject.

To elucidate whether the individual subject became more aware of the stimuli as a result of the training, we calculated the mean PA ratings per SOA for correct responses and errors for each of the two thresholds. As in the larger sample of subjects, we first compared the PA ratings for correct responses and errors of the first threshold. This was done by means of a paired two-sided *t*-test, using each SOA as a sample. As the mean PA rating in the first threshold was higher for correct responses than for errors only in four out of eight SOAs (20, 40, 100, 160 ms), the result of this test was not significant (mean difference 0.0114, $T(7)=0.492$, $p=0.6378$). Nevertheless, we went on to compare the PA rating for correct responses and the PA rating for errors of the first and second threshold, again using paired two-sided *t*-tests using each SOA as a sample. The PA rating for correct responses was significantly higher in the second than in the first threshold (mean difference -0.2189, $T(7)=5.287$, $p=0.0011$), while the PA rating for errors did not change (mean difference -0.0930, $T(7)=2.003$, $p=0.0852$). Also, the PA rating for correct responses was significantly higher than the PA rating for errors in the second threshold (mean difference 0.1373, $T(7)=4.608$, $p=0.0025$). Thus, the individual subject trained for 24 sessions became more aware of the correctly identified stimuli, while she showed no change in PA for incorrectly identified stimuli after the training. The subject was also retested after eight months. At a significance level of $p=0.05$, the sensitivity was significantly lower in the first threshold session than in the retest, while there was no significant difference between the second threshold session and the retest. Thus, the training effect on sensitivity was retained. However, the criterion c was significantly higher than 0 over all SOAs in the retest, thus indicating that the subject was now biased to respond with “diamond”. For the PA rating, we found that the mean rating for correct responses was significantly higher during the retest than during the first threshold session, and significantly higher during the second threshold session that during the retest. The differences in the mean PA rating for errors and the mean PA rating for correct responses only approached significance ($p=0.0585$). These results indicate that a limited amount of the improvements in PA were retained even after eight months, which is in agreement with most of the other retested subjects.

DISCUSSION

The results of this experiment show that subjects’ sensitivity to discriminate between two metacontrast-masked stimuli improves rapidly with practice in a linear manner. This effect is not specific to the trained timing of the stimuli, but generalizes to other timings as well. Furthermore, some limited training effects can be observed for a mirror-symmetric position relative to the horizontal meridian. Critically, when subjects are asked to report their PA of the masked stimuli, an increase in subjective PA can be measured after the training. The training effects for both the objective as well as the subjective measure persist to a limited extent over a period of several months. We thus hypothesize that awareness is trainable, a conclusion that is in accordance with recent findings in blindsight patients (Sahraie, et al., 2006).

Improvements in objective measures

A number of studies have reported improvements in performance after training under conditions of impaired conscious perception in normal observers. For example, training effects have been shown to exist in forward (Coyne, 1981) and backward pattern masking paradigms (Braff, Saccuzzo, Ingram, McNeill, & Langford, 1980; Hertzog, Williams, & Walsh, 1976; Maehara & Goryo, 2003; Schiller, 1965; Schubö, Schlaghecken, & Meinecke, 2001; Ward & Ross, 1977; Wolford & Kim, 1992; Wolford, Marchak, & Hughes, 1988), as well as under conditions of crowding (Chung, 2007; Huckauf & Nazir, 2007). Our results are also consistent with previous reports on learning effects in metacontrast masking. Hernandez and Lefton (1977) observed improvements in sensitivity and no changes in response bias over twelve sessions in a foveal metacontrast detection task. Similarly, Hogben and Di Lollo (1984) reported improvements in a target location identification task over five sessions. However, none of these studies specifically investigated whether subjects can become objectively aware of stimuli that they were previously unaware of. Accordingly, initial thresholds were above $d'=0$ in both previous metacontrast masking experiments (although no statistics are provided regarding this issue). Additionally, in these two studies subjects were not trained on the SOA that yielded the worst performance, but concurrently on all SOAs. Thus, learning could have taken place at a SOA for which the target stimulus was consciously perceivable, and then transferred to other SOAs.

Contrary to our and the above findings, two studies investigating subliminal priming under metacontrast masking conditions observed no change in discrimination performance ($d'=0$) after training with feedback. In a study by Klotz and Neumann (1999), subjects were required to discriminate whether or not a specific target stimulus was present in any of two locations. Even though feedback was given and motivation was upheld by financial reward, subjects failed to improve significantly in their discrimination performance over up to 640 trials. In a similar study Vorberg, Mattler, Heinecke, Schmidt, and Schwarzbach (2004) found no improvements in the ability to discriminate the direction of a metacontrast masked target over more than 3000 trials with error feedback. What might account for the differences between these two studies and our results? In the study of Klotz and Neumann subjects had to attend to two potential target locations, whereas subjects were attending to just one location in our task. It has been shown that perceptual learning can occur without awareness of the target stimuli under certain conditions (Watanabe, et al., 2001). However, if attentional resources are further limited (by the attentional blink) learning does not occur anymore (Seitz, Lefebvre, Watanabe, & Jolicoeur, 2005). Thus, the combination of impaired awareness and distributed attention might have precluded improvements in discrimination performance. In the study by Vorberg and colleagues, subjects were instructed to withhold their responses for 600 ms, whereas in our study, subjects could respond as soon as the mask turned off. It has been argued that subject can be briefly aware of a stimulus, but that delaying the response to a stimulus can lead to rapid forgetting and thus to an underestimation of initial awareness (Dennett & Kinsbourne, 1992; Lachter & Durgin, 1999; Lachter, Durgin, & Washington, 2000). If improvements in discrimination performance do not go along with an improvement in memory, this difference in task might account for the difference between our study and the study by Vorberg and colleagues.

How might the improvements in sensitivity observed in our study come about? One possibility is that subjects learned to suppress the masks. Studies in pattern masking, where it is possible to change the mask while keeping the target stimulus constant, have shown that subjects can learn to suppress specific masks with training, and that this effect breaks down when the masks, but not when the target stimuli are changed (Schubö, et al., 2001; Wolford, et al., 1988). These findings are in line with recent evidence from monkey inferotemporal cortex, where suppression of the single unit response to the mask was the most consistent finding when monkeys were trained to identify masked objects (Op de Beeck, Wagemans, & Vogels, 2007). In metacontrast masking, it has been shown that the suppression of the metacontrast mask by a second mask can lead to target recovery (Breitmeyer & Ogmen, 2006, p.254ff). Thus, it is not unlikely that mask suppression is contributing to the observed improvements in target discrimination.

Another possibility is that the signal of the target stimulus was strengthened (Gold, Bennett, & Sekuler, 1999). In discrimination tasks, an ideal observer would correlate input with templates of the possible targets, and respond according to the best fit. Refinement of the templates leads to less overlap, thus increasing discriminability. Accordingly, in monkeys trained to discriminate orientations the orientation selectivity of V1 neurons increases (Schoups, et al., 2001).

Learning could also have taken place in the time domain. There is good evidence that observers can improve substantially in their ability to discriminate short time intervals between visual stimuli (Westheimer, 1999) and two-flash discrimination performance improves after practicing masked letter recognition (Wolford, et al., 1988). Similar conclusions were drawn by Ventura (1980) with respect to training effects on brightness ratings under foveal metacontrast conditions. Thus, subjects may have learned to discriminate visual events in time that they perceived as one perceptual event before the training in our paradigm as well. Alternatively, subjects may have learned to use the fixed-interval forward mask as a temporal marker for the upcoming stimulus, thus facilitating the allocation of temporal attention in a specific time window (Rolke, 2008; Rolke & Hofmann, 2007). Attention has been shown to attenuate metacontrast masking (Boyer & Ro, 2007; Tata, 2002; also see Breitmeyer & Ogmen, 2006, p.243ff for a critical discussion of attentional effects in metacontrast masking), and effects of training on attention have also been hypothesized to be the basis of perceptual learning (Vidnyanszky & Sohn, 2005). All four proposed mechanisms can in principle explain the improvement in objective performance at the trained position, as well as the transfer effects from the trained SOA to other SOAs. The location specificity of training effects in metacontrast masking has so far not been investigated. Transfer to other stimulus locations is often used to make inferences upon the cortical area where perceptual learning takes place. If transfer is abolished by small stimulus displacements, this is usually taken as evidence for training effects in a lower visual area with small receptive fields (e.g., Karni & Sagi, 1991). In our study, we found evidence for a limited transfer of improvements in sensitivity to another stimulus location. However, because of the large separation of the tested stimulus locations, the rather special position of the stimuli on the vertical meridian, and the mirror-symmetric transfer location no inferences on the likely site of learning are possible (Dill, 2002).

Whatever the specific mechanism is, the crossing of the “threshold of visual awareness” with training indicates that some attribute of the representation of the target stimuli must have changed in order to become accessible to conscious perception. It has been hypothesized that conscious representations are stronger, more stable and more distinct than unconscious representations (Cleeremans, 2008). In principle, training could have led to a strengthening, stabilization and (most importantly) increased distinctiveness of the representation of the target stimulus. In this framework, a representation becomes available to phenomenal consciousness and report once it has reached sufficient quality through learning. Alternatively, mechanisms that read out representations might have become more tuned to the elusive representations of the target stimuli, or more able to use specific features of the target stimuli to discriminate between them. Both accounts require that a representation of the masked target stimulus was actually available even though masking took place. This assumption is supported by at least three lines of evidence: First, even when the target stimulus is perceptually invisible due to metacontrast masking, forced-choice detection can still result in above chance performance (Schiller & Smith, 1966; but see Otto, Ogmen, & Herzog, 2006). Second, the orientation of metacontrast masked targets can be recovered from V1 activity using functional magnetic resonance imaging (Haynes & Rees, 2005). Third, experiments in subliminal priming under metacontrast masking conditions consistently find evidence for access to target information in the absence of conscious target perception (Klotz & Neumann, 1999; Neumann & Klotz, 1994; Vorberg, et al., 2004).

Improvements in subjective measures

Our subjects showed improved discrimination of stimuli as a result of training. When adopting an operationalization of consciousness that relies purely on objective, forced-choice measures, these data already suggest that subjects can indeed become aware of stimuli that they were not aware of before the training. However, although commonly used, such an operationalization of PA has repeatedly been criticized for not taking into account the subjective experience of the subject (e.g., Wiens, 2007). Importantly, even when both objective and subjective measures of awareness are assessed, they do not always yield the same results. For example, Lau and Passingham (2006) found that subjects rated their PA of metacontrast masked stimuli higher for longer SOAs than for shorter SOAs at identical levels of accuracy. Similarly, the blindsight patient GY is able to correctly discriminate motion directions at low and at high motion speed in his blind hemifield while reporting awareness only in the latter condition (Sahraie, et al., 1997). Thus, subjects' accuracy does not always directly reflect their subjective awareness. Accordingly, improvements in objective measures of awareness after training do not necessarily go along with improvements in subjective awareness: Studies in blindsight patients show that remarkable levels of accuracy can be reached in objective tasks in the absence of acknowledged awareness (for a review, see Cowey, 2004). Similarly, investigations of implicit sequence learning indicate that objective performance can improve with practice without conscious awareness of the learned sequences (Destrebecqz & Cleeremans, 2001). In masking of faces by faces, subjects ability to discriminate between fearful and non fearful faces increases over time as indexed by objective measures; however, subjective awareness does not change significantly (Szczepanowski & Pessoa, 2007). These studies further corroborate the notion that high performance on a forced-choice task does not necessarily warrant the conclusion that subjects are aware. Although a dissociation of objective and subjective performance was not the focus of the current study, it is thus possible that subjects improved in our objective task without concomitant changes in subjective awareness.

Depending on the task and stimuli employed, metacontrast masking has been shown to not only impair the discriminability of briefly presented target stimuli, but also to change the subjective experience of these stimuli, for example their brightness (e.g., Petry, 1978). Here, we employed PA ratings asking for stimulus clarity to test for such effects. However, we did not find U-shaped functions of these ratings with SOA. This might be due to the high initial difficulty of the task over all SOAs, a strong central tendency in the rating responses, and/or the fact that metacontrast masked targets are not invisible in the strict sense, but only degraded to such an extent that discrimination becomes impossible. Irrespective of the question whether the PA rating yields a U-shaped function, the ratings vary with accuracy before and after the training. This validates that our ratings capture an aspect of the subjective experience of the subjects, because on average, correct responses should correlate with higher subjective visibility than incorrect responses. Critically, we find an increase in subjective visibility after training which is specific for correct responses. This effect indicates that the subjective quality of the percept changes, irrespective of the fact that the stimulus might have been noticeable (but not identifiable) before the training. Whether this change result from the changes in sensitivity or whether it occurs independently of those remains an open question.

It has also been proposed that PA depends on criterion setting (Lau, 2008b). In this framework, higher order representations of signal and noise are learned from internal signal and noise distributions. The internal signal and noise distributions as well as the learned higher order distributions are separated by a criterion. Whether we report perceiving or do not report perceiving a stimulus (PA) depends on the criterion that separates the learned higher order signal distribution from the learned higher order noise distribution. The underlying internal signal and noise distributions represent the efficiency of information processing, not PA. According to the theory, the internal signal distributions can be captured in 2AFC tasks and the resulting sensitivity measure, whereas PA ratings tap into the higher order representations. Since the higher order representations need to be learned, they will not always reflect the underlying distributions of internal signal and noise properly, and can thus lead to a bias in reported PA, while leaving the sensitivity measure unaffected. In this case, perceptual learning could lead to changes in sensitivity that are accompanied by changes in PA (more efficient information processing and higher PA), changes in sensitivity that are not accompanied by changes in PA (more efficient information processing in the absence of changes in PA), and changes in PA only. In a 2AFC discrimination task with equal stimulus probabilities, optimal performance can only be reached if a stable criterion is established around the point where the two signal distributions intersect, which is indeed what we observe in our data. This criterion, however, is not the criterion that is relevant for reporting PA of the stimulus (although it may be important in learning higher order representations). It thus remains conceivable that perceptual learning leads to a change in criterion for reporting PA without a concomitant change in criterion for the 2AFC task.

Implications for the study of consciousness

An absolute threshold is an unrealistic assumption in the study of consciousness (Eriksen, 1960). However, especially in the field of subliminal perception, a threshold of conscious perception has to be established in order to demonstrate perception without awareness. This is and has been a highly controversial issue. One approach to overcome these controversies has been the implementation of so called "objective" measures of awareness and the application of SDT (Holender, 1986). As in our study, a measure of sensitivity (such as *d'*) is calculated as an aggregate over many trials, and a *t*-test of this measure against 0 is taken as an indication of whether a stimulus was consciously perceived or not. However, our study shows that even a *d'*=0 does not allow for the conclusion that stimuli are presented at a *fixed* threshold of conscious perception. If an adequate number of trials and feedback are provided, and sufficiently high motivation is assured subjects can become "aware" of stimuli that they were previously unaware of. Interestingly, blindsight studies in monkeys usually involve heavy training schedules after surgery, which might explain the rarity of cases where the animals have not been found to show blindsight abilities (e.g., Dineen & Keating, 1981). But even without explicit training, spontaneous improvements in performance have for example been reported after continuous testing over many years in human blindsight patients (e.g., Trevethan, Sahraie, & Weiskrantz, 2007).

Furthermore, being able to become aware also poses a challenge to studies that aim to contrast consciously perceived and not consciously perceived trials (Baars, 1997, p.18ff). Differences observed between perceived and unperceived conditions might not be due to differences between conscious and unconscious perception, but the result of insufficiently trained tasks. Investigating the transition from unaware to aware states might be a valuable alternative to the contrastive approach. Further studies will be needed to show how the amount of training maps onto stages of awareness, and whether awareness *itself* is a trainable function or whether it is training in other functions such as perception or attention which renders previously invisible stimuli visible. Whatever the case may be, our results show that awareness can be modified though training, and that subject can indeed “learn to see”.

3

SUBJECTIVE AND OBJECTIVE LEARNING EFFECTS DISSOCIATE IN SPACE AND IN TIME

Perceptual learning not only improves sensitivity, it also changes our subjective experience. However, the question how these two learning effects relate is largely unexplored. Here we investigate how subjects learn to see initially indiscriminable metacontrast masked shapes. We find that sensitivity and subjective awareness increase with training. However, sensitivity and subjective awareness dissociate in space: Learning effects on performance are lost when the task is performed at an untrained location in another quadrant, whereas learning effects on subjective awareness are maintained. This indicates that improvements in shape sensitivity involve visual areas up to V4, whereas changes in subjective awareness involve other brain regions. Furthermore, subjective awareness dissociates from sensitivity in time: In an early phase of perceptual learning, subjects perform above chance on trials which they rate as subjectively invisible. Later, this phenomenon disappears. Subjective awareness is thus neither necessary nor sufficient for achieving above-chance objective performance.

Schwiedrzik CM, Singer W, Melloni L (2011). Proceedings of the National Academy of Sciences of the USA, 108(11):4506-11.

Our perceptual apparatus is constantly shaped by experience. This has been shown, i. a., in experiments investigating perceptual learning, where practice on a sensory task leads to increases in perceptual sensitivity (Goldstone, 1998). Although perceptual learning is a well studied phenomenon, the question of how it changes subjective awareness has rarely been addressed. Do we actually “see more” after training? Apart from anecdotal reports pointing in this direction, a quantitative analysis of the effects of learning on subjective awareness is largely missing (Fahle, 2009). This is due to the fact that studies in perceptual learning have almost exclusively focussed on objective task performance (i.e., sensitivity in terms of Signal Detection Theory (D. M. Green & Swets, 1966)). However, studies from a different line of research, namely those investigating conscious perception, have found that objective performance in a task and awareness of the stimuli on which the task is performed can dissociate. Such is the case in blindsight (Stoerig, 2006), but also in normal observers (Lau & Passingham, 2006; Meeres & Graves, 1990; Schärli, Brugger, Regard, Mohr, & Landis, 2003; but see Azzopardi & Cowey, 1998). This shows that these two aspects of perception cannot be treated as being equivalent. We have recently shown that subjects can be trained to perform on and to see stimuli that are initially invisible to them (Schwiedrzik, Singer, & Melloni, 2009). We thus hypothesize that awareness is trainable, a conclusion that is in accordance with recent findings in blindsight patients (Sahraie, et al., 2006). However, the time course of those learning effects has not been explored, i.e., whether the improvements in sensitivity and subjective awareness are dependent on each other. In particular, changes in sensitivity could be a prerequisite for changes in subjective awareness. Alternatively, it could be necessary to subjectively see a stimulus in order for changes in sensitivity to occur. Last but not least, training could affect sensitivity and subjective awareness in parallel without any mutual dependence between these two aspects of perception. The latter question is related to another issue which is still a matter of debate, namely whether objective performance and subjective awareness depend on the same circuits in the normal brain.

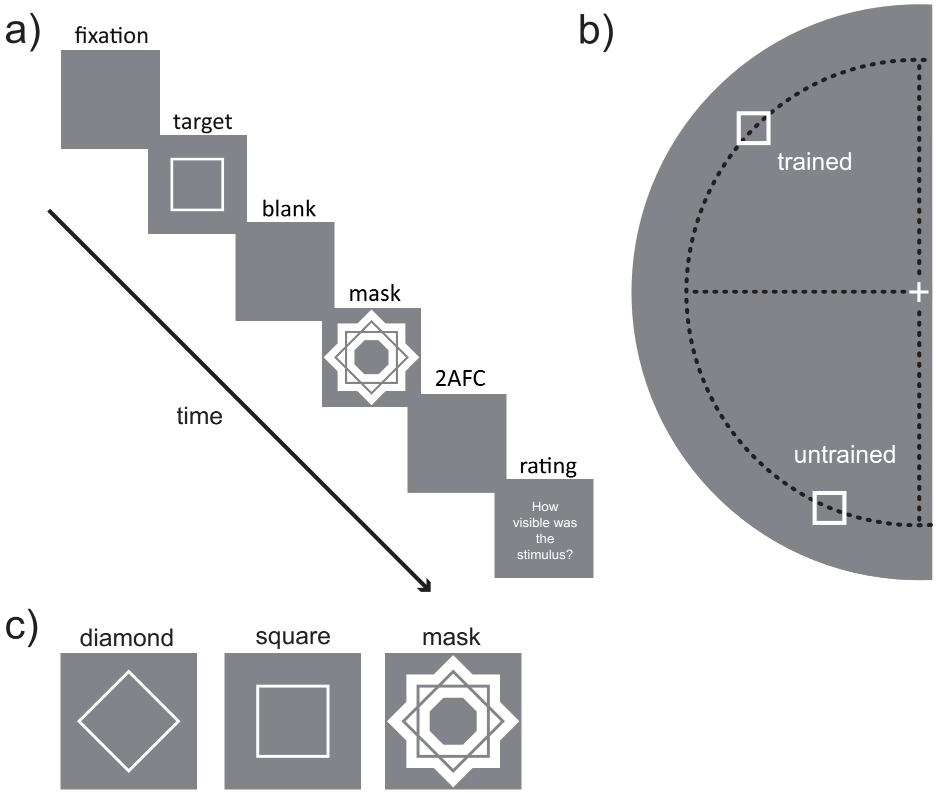

Figure 3-1. Stimuli and experimental procedures. (a) Each trial started with a fixation period (1000-1500 ms). The target was presented for 10 ms. The mask was presented for 50 ms at SOAs between 20 ms and 150 ms. Subjects then decided between square and diamond in a two alternative forced-choice (2AFC) task, followed by a rating of subjective visibility of the target stimulus. (b) The training location was in the upper left quadrant at 4° eccentricity. The isoeccentric transfer location was in the lower left quadrant, 6.6° from the trained location. (c) A diamond and a square served as target stimuli, the mask was star-shaped and did not overlap with the target contours.

Here, we address both questions by training subjects on a shape discrimination task under metacontrast masking conditions (Figure 3-1). Stimulus parameters are such that the shapes are objectively indiscriminable before training. Through continuous practice, subjects learn to discriminate a square from a diamond, thus crossing the objective threshold (d'=0). Additionally, subjects rate their subjective awareness of the stimuli on a trial-by-trial basis. This allows us to measure the time courses of learning effects on sensitivity and subjective awareness concurrently. To localize the learning effects in the brain by psychophysical techniques, we change the stimulus position after the final training session to a new stimulus location. Based on an estimate of receptive field (RF) size, this allows us to test the hypothesis that early visual areas up to V4 are the sites of objective learning effects.

We find that changes in sensitivity and changes in subjective awareness indeed dissociate: Improved sensitivity does not transfer into another quadrant, whereas subjective awareness does, indicating that subjective awareness and objective performance depend on different brain regions. Furthermore, in an early phase of perceptual learning, subjects perform above chance on trials which they rate as subjectively invisible. This effect disappears with practice. Taken together, these results support the notion that thresholds of awareness are not fixed, that subjective awareness is neither necessary nor sufficient for changes in sensitivity, and that the cortical loci of learning effects in subjective awareness and in sensitivity are not identical.

RESULTS

Subjective and objective learning effects dissociate in space: Thresholds and transfer

Thresholds as a function of stimulus onset asynchrony (SOA) were first assessed before training. For each subject, the SOA yielding zero sensitivity ($d'=0$) in this first threshold measurement was subsequently used for training (see below). In order to evaluate whether improvements in performance were confined to the trained SOA or generalized to untrained SOAs, we also evaluated the full threshold function after training. In addition, we changed the location of the stimulus to a position 6.6° away from the trained location to another quadrant at isoeccentricity after the training and the final threshold assessment to probe which brain regions are involved in the learning effects. In particular, this was done to test whether area V4, an important intermediate stage in the analysis of shape (Pasupathy, 2006), or any higher area in the ventral stream was the locus of perceptual learning in our study. Finally, to address the potential role of feedback in perceptual learning, half of the subjects received blockwise percentage correct feedback during training.

To evaluate whether subjects' sensitivity, response bias, and subjective awareness changed as a function of training, we first entered their d', c, and mean Perceptual Awareness

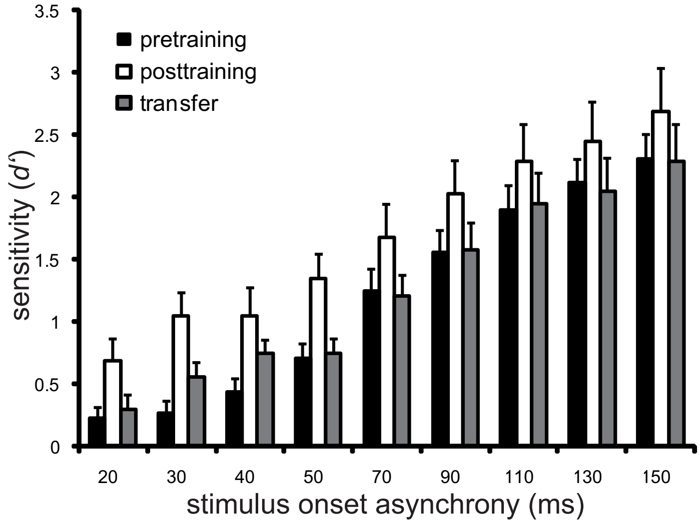

Figure 3-2. Sensitivity (*d'*) as a function of SOA before and after the training phase and at the untrained location. Sensitivity always increased linearly with SOA. Linear fits of the mean *d'* to the SOAs were highly significant at each threshold (all $R^2 > 0.9$, all $P < 0.01$). Error bars represent the *SE* of the estimated marginal mean.

Scale (PAS, Ramsøy & Overgaard, 2004) rating from the threshold assessments separately into repeated measures analyses of variance (rmANOVA) with the factors session (before training, after training, transfer) and SOA (for *d'* and *c*: 20, 30, 40, 50, 70, 90, 110, 130, 150 ms, for PAS: 20-50 ms, see below). For the PAS data we included the factor accuracy (correct, incorrect). We also tested whether the presence of blockwise feedback affected any of the measures (between subjects factor feedback). On the PAS, 1 refers to "No Experience", 2 to "Brief glimpse (a feeling that something has been shown)", 3 to "Almost clear experience (ambiguous experience of the stimulus)", and 4 to "Clear experience", whereby PAS>=2 indexes 'subjective detection' and PAS=4 indexes 'subjective discrimination' (see below).

Objective learning effects

Sensitivity (*d'*) increased linearly with SOA in all sessions (effect of SOA, $F(1.381, 27.618)=44.794$, $p<0.01$, $\eta^2=0.691$), a signature of type A metacontrast masking (Figure 3-2). Learning clearly affected sensitivity (effect of session, $F(1.603, 32.065)=9.744$, $p<0.01$, $\eta^2=0.328$): Mean *d'* was significantly higher after the training than before (mean difference 0.498, *SE* 0.141, *p*<0.01) and as compared to the transfer position (Figure 3-3a, mean difference 0.425, *SE* 0.087, *p*<0.01). Thus, training on a discrimination task results in above chance performance even if subjects initially show zero sensitivity. However, there was no significant difference between sensitivity before training and after training at the transfer position (mean difference -0.073, *SE* 0.131, *p*>0.58, uncorrected). Thus, subjects' sensitivity increased at the trained position, but this improvement did not transfer to another quadrant at isoeccentricity. This indicates that changes in sensitivity involve visual areas up to area V4 (see Discussion). Separate analyses restricted to the trained SOA and on the subsample of SOAs that we used for the analysis of the PAS data confirmed these results. Blockwise feedback did not affect performance on this task (no significant main effect of feedback or interactions with feedback, all *p*>0.41).

As for response bias (*c*), we found that different SOAs were not associated with different response biases, and importantly, response bias did not change as a result of training nor with the transfer (no significant main effects or interactions, all *p*>0.16).

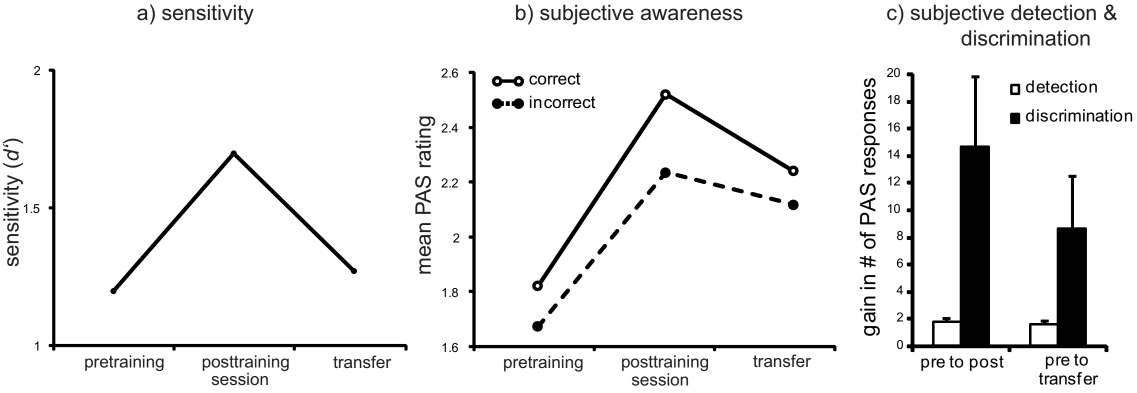

Figure 3-3. Threshold sessions. (a) Average sensitivity (*d'*) and (b) subjective awareness (mean PAS rating) for correct and incorrect trials on the four short SOAs in the threshold sessions before and after training and at the untrained location (transfer). (c) Average learning induced gain in subjective detection ([# of trials with PAS>=2 posttraining] divided by [# of trials with PAS>=2 pretraining]) and subjective discrimination ([# of trials with PAS=4 posttraining] divided by [# of trials with PAS=4 pretraining]) on the four short SOAs at the trained and transfer location. Error bars represent the standard error of the mean.

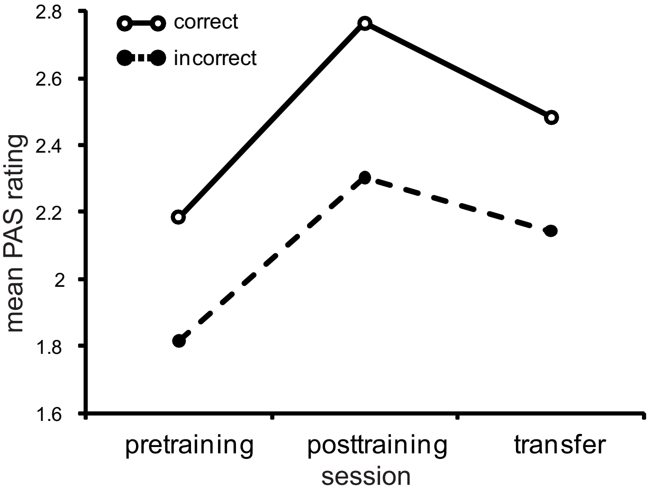

Figure 3-4. Average subjective awareness for all nine SOAs in the threshold assessments before and after the training phase and at the untrained location. Empty cells were replaced by the average PAS rating from the respective group (feedback/no feedback).

Subjective learning effects

For the PAS data, only the first four SOAs (20-50 ms) were entered into the analyses because 17 out of 22 subjects had less than 10 wrong answers in at least one of the longer SOAs (70-150 ms), which would have rendered the factor accuracy unreliable (see Figure 3-4 for all SOAs). Ratings of subjective awareness increased as a function of SOA ($F(1.412, 28.240)=10.150, p<0.01, \eta^2=0.337$) and session ($F(1.410, 28.203)=29.917, p<0.01, \eta^2=0.599$), and this increase was especially pronounced on correct trials (Figure 3-3b, session × accuracy interaction, $F(1.764, 35.275)=6.654, p<0.01, \eta^2=0.250$). Importantly, the mean PAS rating for correct responses was higher than for incorrect responses in each session (all $p<0.01$). Although ratings for both correct and incorrect responses increased after training, this increase was substantially bigger for correct responses (mean difference between PAS rating for correct and incorrect responses pretraining: 0.37, *SE* 0.043, $p<0.01$; posttraining: 0.459, *SE* 0.082, $p<0.01$). Comparable results were obtained when we plotted the rate of incorrect trials with high PAS ratings against the rate of correct trials with high PAS ratings for three levels of visibility to obtain Receiver Operating Characteristics (*ROC)* and calculated the Area under the curve (*AUC*), an index of how well PAS ratings predict accuracy (see Materials and Methods). The *AUC* was above chance (0.5) before and after the training both for the trained SOA (Figure 3-5, pretraining: mean difference 0.0656, *T*(21)=5.896, *p*<0.01; posttraining: mean difference 0.1260, *T*(21)=6.907, *p*<0.01, all one-sided) as well as for the average over the 4 short SOAs (pretraining: mean difference 0.0866, *T*(21)=7.579, *p*<0.01; posttraining: mean difference 0.1329, *T*(21)=7.441, *p*<0.01, as above). This pattern of results rules out that increases in PAS ratings were due to a bias to indiscriminately respond with higher ratings as learning progresses, since such behavior would erase differences between correct and incorrect trials. We found no effect of or interaction with feedback (all *p*>0.41).

A dissociation between sensitivity and subjective awareness was observed for the transfer task. Whereas moving the stimulus to another quadrant resulted in a significant drop of sensitivity, PAS ratings for correct responses remained above pretraining levels (Figure 3-3b). Crucially, the difference between correct and incorrect responses remained significant (mean difference transfer: 0.34, *SE* 0.059, *p*<0.01), and PAS ratings still predicted accuracy when tested against chance (Figure 3-5, *AUC* trained SOA: mean difference 0.1031, *T*(21)=4.976, *p*<0.01; *AUC* 4 short SOAs: mean difference 0.0911, *T*(21)=6.723, *p*<0.01, all one-sided), confirming that subjects continued using the PAS ratings in a meaningful way. Thus, the RFs underlying objective task performance appear to be smaller than the RFs supporting subjective awareness, indicating that training affects different brain regions.

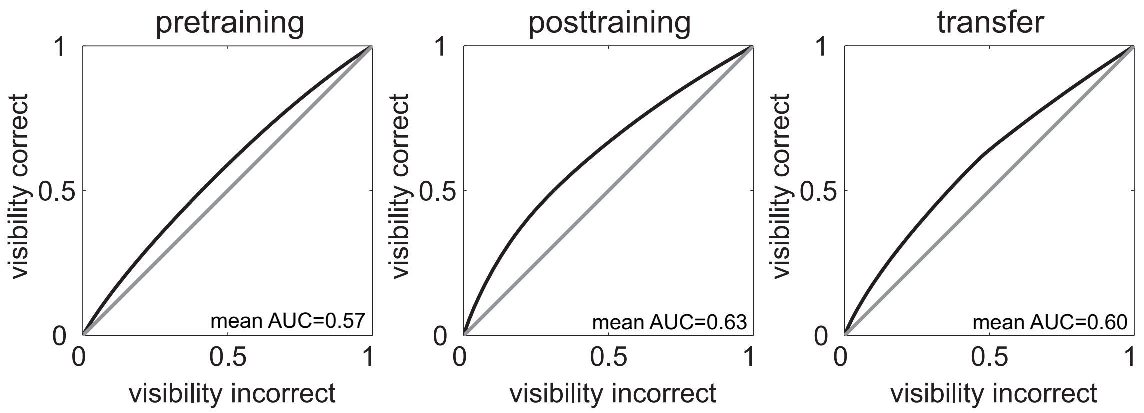

Figure 3-5. Average *ROC* for the trained SOA for pretraining (Left) and posttraining (Center) and at the transfer location (Right). The *ROC* curves were obtained by fitting a line of 1000 points to the three inflection points by means of maximum-likelihood estimation. The resulting *AUC* was significantly different from 0.5 (chance) in all three sessions (all *P* < 0.01, one-sided, Bonferroni corrected).

It is unlikely that the high PAS ratings at the untrained location are explainable by a carry-over effect from the postraining threshold measurement to the transfer location. If a carry-over effect were in place then the results at the untrained location should resemble in some way the pattern of results obtained at the source of the carry-over effect, i.e., the posttraining session. This was not the case: Neither the absolute PAS ratings for correct and incorrect trials nor the difference between them (trained SOA: mean difference 0.11, *SE* 0.04, $T(21)=2.2924$, $p<0.05$; average of the 4 short SOAs: mean difference 0.16, *SE* 0.04, $T(21)=3.4907$, $p<0.01$) were identical in the posttraining and transfer session (Figure 3-3b).

The previous analysis revealed a gradual increase in subjective awareness with practice. However, it is unclear whether the observed increments reflect heightened subjective awareness of discriminant features (i.e., a clearer impression of shape), or heightened subjective detection ability (i.e., seeing something as opposed to seeing nothing) (Dienes & Seth, 2010), since by design, the PAS encompasses both. A dissociation between subjective awareness and objective performance could then correspond to differences in task, namely detection versus discrimination. In order to compare the subjective and objective task when both rely on shape information, we divided the PAS ratings into two categories: trials where subjects reported a clear impression of the shape of the stimulus (PAS=4), and trials where subjects at least reported seeing something (PAS>=2). For brevity, we refer to these trials as 'subjectively discriminated' and 'subjectively detected', respectively. The proportion of subjectively detected stimuli increased from pretraining to posttraining (trained SOA mean difference -0.284, *SE* 0.052, $p<0.01$; 4 short SOAs: mean difference -0.252, *SE* 0.047, $p<0.01$), and was above pretraining levels at the untrained location (trained SOA mean difference -0.196, *SE* 0.057, $p<0.01$; 4 short SOAs: mean difference -0.185, *SE* 0.054, $p<0.01$). The proportion of subjectively discriminated trials also increased from pretraining to posttraining (trained SOA mean difference -0.152, *SE* 0.048, $p<0.01$; 4 short SOAs: mean difference -0.159, *SE* 0.047, $p<0.01$, all one-sided, Bonferroni corrected for the number of sessions). Most importantly, it remained above pretraining levels at the untrained transfer location (trained SOA: mean difference -0.072, *SE* 0.030, $p<0.05$; 4 short SOAs: mean difference -0.083, *SE* 0.031, $p<0.05$, as above), although not fully reaching posttraining levels (trained SOA: mean difference 0.080, *SE* 0.022, $p<0.01$; 4 short SOAs: mean difference 0.076, *SE* 0.021, $p<0.01$, as above). This shows that the dissociation between sensitivity and subjective awareness holds even when objective discrimination and subjective discrimination are directly compared. To assess whether learning effects in subjective discrimination exceeded learning effects in subjective detection, we calculated the respective gain for the trained and transfer location. Indeed, the gain for discrimination was substantially larger than for detection at both locations (Figure 3-3c, measure × session interaction, trained SOA: $F(1, 21)=5.117$, $p<0.05$, $\eta^2=0.196$; 4 short SOAs: $F(1, 21)=8.376$, $p<0.01$, $\eta^2=0.285$). Thus, subjective discrimination increased with training and transferred to the untrained location, and this learning effect exceeded the learning effect in subjective detection.

Subjective and objective learning effects dissociate in time: Training sessions

Objective learning effects

Subjects were trained for five consecutive sessions at an SOA that was initially at the objective threshold, as confirmed by a *d'* not significantly different from 0 before training neither in the feedback group (mean difference: 0.15, $T(10)=2.129$, $p>0.05$, two-sided, uncorrected) nor in the no feedback group (mean difference: 0.12, $T(10)=1.873$, $p>0.05$, two-sided, uncorrected). We have previously shown that practice on this task leads to increases in sensitivity without changes in response bias (Schwiedrzik, et al., 2009). Here, we confirmed these results in a rmANOVA with the within subjects factor training session (session 1-5) and the between subjects factor feedback. We found that sensitivity increased with session ($F(2.573, 51.454)=15.830$, $p<0.01$, $\eta^2=0.442$), whereas response bias did not change significantly ($F(2.179, 43.576)=0.737$, $p>0.49$, $\eta^2=0.036$). Sensitivity increased linearly ($F(1, 20)=29.752$,

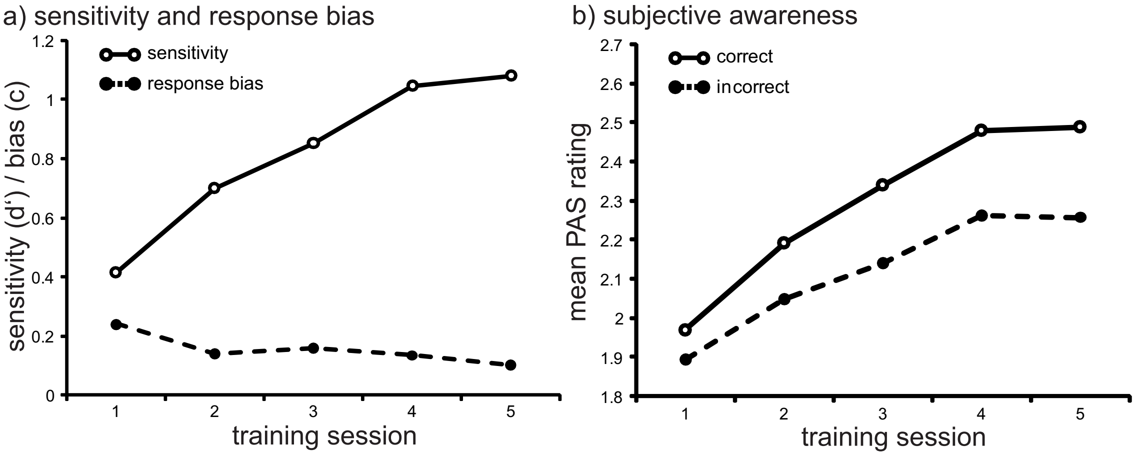

Figure 3-6. Training sessions. (a) Sensitivity (*d'*) and response bias (*c*) at the trained SOA during the training sessions. (b) Subjective awareness (mean PAS rating) for correct and incorrect trials during the training sessions.

$p<0.01$, $\eta^2=0.598$), on average by a d' of 0.66 (*SD* 0.56) from the first to the fifth session (Figure 3-6a). Blockwise feedback did not affect sensitivity nor response bias in this task (all $p>0.11$).

Subjective learning effects

To examine the time course of subjective awareness as a function of practice, we asked subjects to rate the stimulus visibility on a trial-by-trial basis. Subjective awareness ratings increased over sessions ($F(1.437, 28.742)=12.017$, $p<0.01$, $\eta^2=0.375$), with a more pronounced increase for correct responses (accuracy × session, $F(1.941, 38.828)=7.672$, $p<0.01$, $\eta^2=0.276$), but no effect of feedback (all $p>0.26$). When considering subjectively detected (PAS>=2) and subjectively discriminated (PAS=4) trials separately, we found that the proportions of both increased with session (detection: $F(1.337, 28.085)=8.218$, $p<0.01$, $\eta^2=0.281$, discrimination: ($F(1.467, 30.816)=7.384$, $p<0.01$, $\eta^2=0.260$)). Thus, practice on initially indiscriminable stimuli leads to improvements in objective performance and subjective awareness (Figure 3-6). The *AUC* per session was always above chance (Figure 3-7, all $p<0.01$, one-sided), and even increased linearly with session ($F(1, 20)=15.680$, $p<0.01$, $\eta^2=0.439$). Thus, the improvements in subjective awareness do not simply reflect a change in criterion, since a bias to indiscriminately

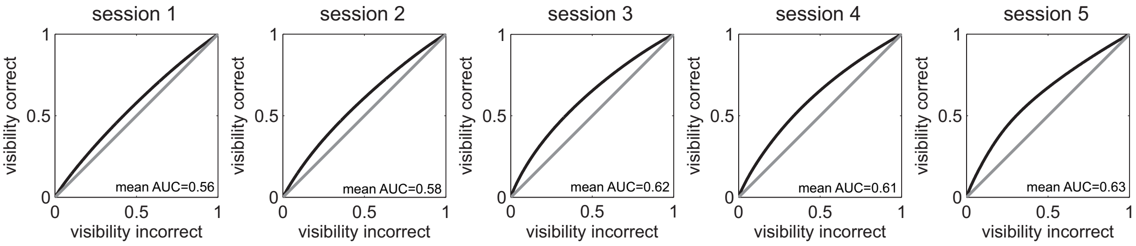

Figure 3-7. Average *ROC* per training session. The *AUC* was always above chance (all $P < 0.01$, one-sided, Bonferroni corrected) and increased linearly with session [$F(1, 20) = 15.680$, $P < 0.01$, $\eta^2 = 0.439$].

use higher ratings as learning progresses would not lead to a differential increase in the ratings for correct and incorrect responses or to an increase in the *AUC*.

We then turned to the question whether objective performance was dependent on subjective awareness, i.e. whether subjects had to consciously perceive the stimuli in order to respond correctly, and whether this changed over the course of learning. We separately analyzed the percentage of correct responses for trials rated as invisible (PAS=1) and for trials rated as clearly visible (PAS=4) per session. We found that subjects performed significantly better than chance for trials that they rated invisible (mean percentage correct: 0.57, $T(21)=2.865$, $p=0.0465$, two-sided, Bonferroni corrected for the number of sessions), but only during the first session (Figure 3-8, all other $p>0.88$, as above)[1]. This suggests that subjective awareness is not necessary for above-chance objective performance at the earliest stage of learning. For trials that were rated as clearly visible, the percentage of correct responses was always higher than chance (all $p<0.02$, as above). Since the number of PAS=1 and PAS=4 trials develops in opposite directions over sessions (Figure 3-8), the above chance performance for PAS=4 trials in session 1 rules out the possibility that the lack of above chance performance for PAS=1 trials in sessions 2-5 is due to the progressively smaller number of trials. Taken together, subjects were aware of some trials (correct trials |

---

[1] Missing values were replaced by the mean of the respective group (feedback/no feedback). This did not change the overall pattern of results. There were no missing values for PAS=1 trials in session 1.

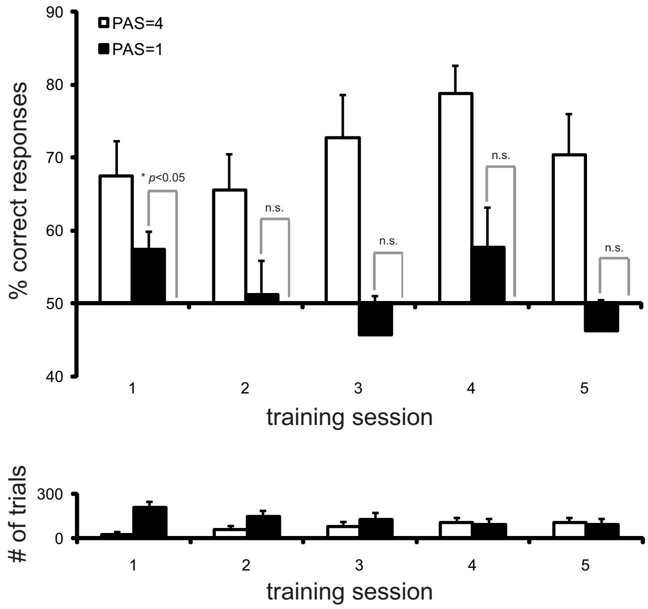

Figure 3-8. Objective performance on invisible and clearly visible trials during training. (upper panel) Percentage of correct responses on subjectively invisible (PAS=1) and subjectively clearly visible (PAS=4) trials per training session. (lower panel) Average number of invisible (PAS=1) and clearly visible (PAS=4) trials per session. Error bars represent the standard error of the mean, 'n.s.' is not significant at an α level of 0.05, Bonferroni corrected.

PAS=4) even in session 1; concomitantly, they were apparently able to utilize information that fell below their subjective threshold for the discrimination task (correct trials | PAS=1). However, although subjects continued to report that some stimuli were invisible in session 2-5, they did not continue to perform above chance in these trials. Thus, it appears that with increasing awareness, information below the subjective threshold is either not available or not used anymore.

DISCUSSION

We found that both sensitivity and subjective awareness change in perceptual learning. However, learning does not affect these two aspects of perception in the same way. This evidences that they should not be treated as being equivalent. Subjective awareness is not *sufficient* for achieving or maintaining objective performance: At an untrained location in another quadrant, sensitivity drops back to pretraining levels, whereas learning effects on subjective awareness are preserved. This differential generalization of objective performance and subjective awareness across the retinotopic map indicates that the cortical loci of learning effects in subjective awareness and sensitivity are not identical. Furthermore, subjective awareness is also not *necessary* for above-chance objective performance, as performance on trials which were rated as subjectively invisible was above chance for the discrimination task. However, this is only the case in the first training session. Thus, the immediate use of subjectively unavailable information vanishes as perceptual learning progresses. Still, changes in sensitivity can occur even when the stimuli used for training do not cross the threshold of subjective awareness. Taken together, these findings support the notion that perceptual thresholds are not fixed, and that in order to fully characterize perceptual learning, both objective and subjective measures need to be considered.

Progression of sensitivity and subjective awareness

Both sensitivity and subjective awareness increased with training. Comparing the pretraining with the posttraining thresholds, we found that the learning effects are not limited to the trained SOA, but spread in time to all other tested SOAs. The training related improvements can have diverse reasons. Given that subjects could not infer the precise moment of appearance of the target (randomized fixation period, randomized SOAs), learning effects cannot be attributed to an improved focusing of attention in time. Another possibility is that subjects learned to suppress the mask, as has been found for pattern masking (Schubö, et al., 2001; Wolford, et al., 1988). This cannot be directly tested in metacontrast masking since the mask cannot be changed for a given target. However, mask suppression by a second mask has been shown to lead to target recovery in metacontrast masking (Breitmeyer, Rudd, & Dunn, 1981). Alternatively or additionally, the representation of the target stimuli might have been strengthened through learning (Gold, et al., 1999).

The changes in objective performance were accompanied by changes in subjective awareness. The more pronounced increase in subjective awareness on correct than on incorrect trials and the increase of the *AUC* with sessions show that the higher PAS ratings were not due to an indiscriminate response bias towards higher scores as learning progresses. Notably, the learning curves for sensitivity and subjective awareness dissociate. Only in the first session, information remaining below the subjective threshold leads to above chance performance. This effect vanishes in the remaining sessions. Previous studies in blindsight patients as well as in normal observers have found similar dissociations within sessions (Lau & Passingham, 2006; Meeres & Graves, 1990; Schärli, et al., 2003; Stoerig, 2006). Thus, under certain conditions, visual stimuli allow for correct behavior without concomitant subjective experience. Learning studies have also reported dissociations between sensitivity and subjective awareness: For example, normal observers can learn to discriminate motion directions (Watanabe, et al., 2001), orientations (Seitz, et al., 2009), and emotional expressions in masked faces (Szczepanowski & Pessoa, 2007) without awareness. However, contrary to our results, practice did not improve subjective awareness in these studies.

The above chance performance on trials which subjects rated as invisible suggests that information which has entered the visual system can be utilized even if this information is not accessible to subjective awareness. A change in subjective awareness requires that the initial representation of the stimulus is transformed to make it accessible for subjective report. It has been proposed that this requires the formation of a higher order representation of the stimulus (Lau, 2008b). In this framework, higher order representations are learned from the underlying signal and noise distributions upon which the discrimination task is performed. Learning-induced changes in the underlying distributions (as captured by *d'*) and/or changes in the signal and noise distributions of the higher order representations can then lead to increased subjective awareness. It has also been hypothesized that conscious representations are stronger, more stable and more distinct than unconscious representations (Cleeremans, 2008). A recent fMRI study lends support to this proposal by showing that activity patterns in the temporal lobe elicited by consciously perceived stimuli are less variable than those evoked by unperceived stimuli (Schurger, Pereira, Treisman, & Cohen, 2010; but see Clifford, 2010). Accordingly, once a representation has reached sufficient quality and stability through learning, it becomes accessible to report. Our study expands previous findings by showing that the immediate use of subjectively unavailable information vanishes once a consciously accessible representation has been established.

Interestingly, objective performance on clearly seen trials was not at ceiling throughout the training sessions. This is consistent with previous masking studies (Boehler, Schoenfeld, Heinze, & Hopf, 2008; Del Cul, Dehaene, Reyes, Bravo, & Slachevsky, 2009; Summerfield, Jack, & Burgess, 2002). This finding can be accounted for by a model in which subjective awareness and objective performance do not depend on each other. However, other factors, e.g., motor errors when reporting the target stimulus in the 2AFC task and/or subjective awareness or illusory percepts, could also account for part of the variance. In masking, it is believed that illusory percepts can arise when top-down processes incorrectly complete the highly degraded target information (Summerfield, et al., 2002). Such illusory percepts should prevail with short

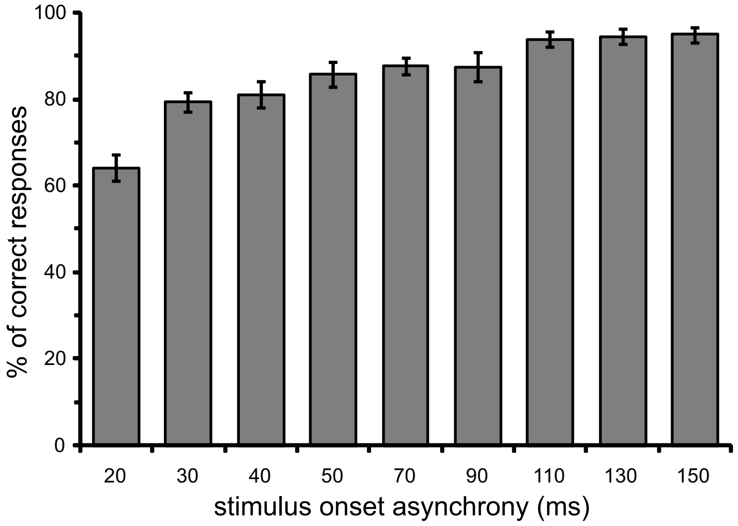

Figure 3-9. Average percentage of correct PAS=4 responses per SOA. Objective performance on clearly seen trials increases with SOA [main effect of SOA: $F(4.295, 85.891) = 21.567$, $P < 0.01$, $\eta^2 = 0.519$] and is at ceiling for the longest SOAs, which indicates that the relationship between objective performance and subjective awareness depends on the amount of available bottom-up information. Such behavior would be predicted by a model in which subjective awareness and objective performance rely on parallel channels with common input but independent source of noise: The more bottom-up information is available, the lower the influence of uncorrelated noise in the two channels. Suboptimal performance on clearly seen trials can also be attributed to motor errors or illusory percepts. Error bars represent the *SE* of the estimated marginal mean.

SOAs when bottom-up information is scarce and the influence of top-down mechanisms is strong (Figure 3-9).

Different effects of a change of stimulus location

A further dissociation between sensitivity and subjective awareness was observed for the transfer task. After the training and the final threshold assessment, we changed the location of the stimulus to a position 6.6° away from the trained location to another quadrant at isoeccentricity. At 4° eccentricity, RFs in human area V4 are smaller than 6° (Kastner et al., 2001). Furthermore, RFs in area V4 usually do not cross the horizontal meridian (Boussaoud, Desimone, & Ungerleider, 1991; Gattass, Sousa, & Gross, 1988), whereas the RFs at later stages in the ventral stream, such as those of the anterior inferotemporal cortex, are on average much larger and not constrained to an individual quadrant (Desimone & Gross, 1979). The fact that all learning effects in sensitivity were lost when the stimulus location was changed indicates that learning most likely involved area V4 and/or preceding areas. In support of this notion, electrophysiological experiments in monkeys have found suppressive effects of metacontrast masks on neuronal activity of V4 neurons (Kondo & Komatsu, 2000).

The change in stimulus location did not have the same effect on subjective awareness. In fact, subjective awareness remained close to posttraining levels at the new location. This was even the case when only considering trials on which subjects could subjectively discriminate the stimuli, showing that the dissociation between objective performance and subjective awareness was not solely due to a transfer of subjective detection ability or a difference in task. Crucially, an unspecific carry-over effect from the postraining threshold measurement cannot easily explain the results since such an effect should either preserve the absolute PAS ratings and/or the relative difference between PAS ratings for correct and incorrect responses, which was not the case. Thus, sensitivity and subjective awareness generalized differently across the visual field, implying that the respective learning effects involve different brain regions. Similar results have been obtained for visual priming: Subliminal visual primes (affecting sensitivity) are only effective when displayed in the same quadrant as the supraliminal targets, whereas supraliminal primes (affecting sensitivity and subjective awareness) are translation invariant (Bar & Biederman, 1999). Furthermore, our results indicate that subjective awareness is not only not necessary for correct responses, it is also not sufficient to support generalization of learning effects to the untrained location.

A region suggested to be particularly relevant for subjective awareness is the dorsolateral prefrontal cortex (DLPFC, Lau & Passingham, 2006). As in our study, the authors compared matched performance levels for differences in subjective awareness. This was achieved by comparing two SOAs in type B metacontrast masking. The DLPFC showed higher BOLD activity at matched performance levels for SOAs that led to more 'seen' responses than another SOA with less 'seen' responses. The RFs of visually responsive neurons in the DLPFC are larger than 6.6° at 4° eccentricity (Rainer, Asaad, & Miller, 1998; Suzuki & Azuma, 1983) which would be compatible with the generalization of subjective awareness across space. Also, learning to discriminate visual objects leads to a sparsening of population activity and a sharpening of the tuning of individual neurons in this area (Rainer & Miller, 2000). Thus, given its RF characteristics and malleability by practice, the DLPFC is a possible site of learning effects for subjective awareness.

No measurable effect of blockwise feedback

The factor feedback remained insignificant in all analyses, i.e., blockwise feedback did not influence the development of sensitivity, response bias, or subjective awareness. Previous evidence on the role of feedback in perceptual learning is inconclusive. While some studies have found that feedback is necessary (trial-by-trial; Seitz, Nanez, Holloway, Tsushima, & Watanabe, 2006) or at least beneficial (blockwise; Herzog & Fahle, 1997), others have found no effects (trial-by-trial; Fahle & Edelman, 1993; Petrov, Dosher, & Lu, 2006). Our results indicate that repeated performance on the stimuli is sufficient for learning under metacontrast masking conditions, as has also been found for learning effects in blindsight patients (Sahraie, et al., 2006). However, it remains possible that trial-by-trial feedback is more effective in driving learning in our task.

Conclusions

Our results show that perceptual thresholds, objective or subjective, are not fixed, but can be changed through practice. Awareness is trainable. However, the dissociations between subjective awareness and objective performance show that the respective improvements are not simple byproducts of each other. In particular, their neuronal substrate seems to differ. Thus, if we want to understand how perceptual learning affects not only sensory processing but also higher cognitive functions, and unravel the brain regions involved in such training effects, we will have to go beyond assessing performance in isolation.

MATERIALS AND METHODS

Participants

22 subjects (8 male, mean age 24, range 19-30) participated in the experiment after giving written informed consent. All but 3 subjects were right-handed as assessed with the Edinburgh Inventory (Oldfield, 1971); all reported normal or corrected-to-normal vision and no history of neurological and/or psychiatric disease.

Stimuli

Stimuli were displayed on a CRT monitor (HP p1230, resolution 1024 × 768, visible screen size 30° × 22.9°) at a refresh rate of 100 Hz. Subjects viewed the screen from 75 cm distance. Background luminance was 3.18 cd/m^2. A square (0.35° × 0.35°) and a diamond (0.53° × 0.53°) were used as target stimuli. The outlines of the targets were 0.02° wide and had a luminance of 25.74 cd/m^2. The mask (25.74 cd/m^2) was star-shaped (Figure 3-1c). Its inner edges were contiguous with the target stimuli from both sides without spatial overlap. In the main experiment, all stimuli were presented in the upper left quadrant at 4° eccentricity. For the transfer, stimuli were presented in the lower left quadrant, 6.6° from the trained position (center-to-center) at isoeccentricity (Figure 3-1b). A fixation cross was always present at the center of the screen. Each trial started with a fixation period of 1000 to 1500 ms. Subsequently, the target was presented for 10 ms. The mask was presented for 50 ms at stimulus onset asynchronies (SOAs) ranging from 20 ms to 150 ms (20, 30, 40, 50, 70, 90, 110, 130, 150 ms) for the threshold estimation or at an individually determined SOA during the training sessions (Figure 3-1a).

Procedure

Subjects had to discriminate whether they saw a square or a diamond by pressing one of two buttons on a keyboard (counterbalanced within groups). Additionally, they had to rate the subjective visibility of the respective stimulus on the four point Perceptual Awareness Scale (PAS, Ramsøy & Overgaard, 2004) on a trial-by-trial basis by a button press. On this scale, 1 corresponds to “No Experience”, 2 to “Brief glimpse (a feeling that something has been shown)”, 3 to “Almost clear experience (ambiguous experience of the stimulus)”, and 4 to “Clear experience of the stimulus”. Subjects were asked to maintain fixation on the center of the screen throughout the experimental sessions. The experiments were conducted in a darkened, sound attenuating chamber. Constant head position was assured by the use of a chinrest with forehead support.

The experiment took place on five consecutive days. On the first day, we determined at which SOA a given subject performed the discrimination task at chance (20 ms: 7 subjects, 30 ms: 10 subjects, 40 ms: 4 subjects, 50 ms: 1 subject). This SOA was then used for the training. The first training session was conducted directly after the threshold measurement. On day two to four, only training sessions were conducted. On the fifth day, the last training session took place. After this training session, we again assessed the masking threshold, followed by a threshold measurement at the transfer position.

For all threshold measurements, each target was presented 40 times at each SOA, yielding 80 trials per SOA and a total of 720 trials. After every 180 trials, we introduced a break of variable length. The occurrence of SOAs was randomized and counterbalanced over blocks. The sequence of target stimuli was fully randomized and no feedback was given.

After the initial threshold measurement, subjects were randomly assigned to either a feedback or a no feedback group. Subjects in the feedback group would receive blockwise percentage correct feedback during the training sessions. Subjects completed 600 trials per training session (a total of 3000 trials). After every 100 trials, a break was introduced. The number of squares and diamonds was balanced per block. Subjects were paid 15 € per hour. To assure constant motivation during the training sessions, subjects received a bonus of 2 € if they improved by 10% from the previous training session, or a fee of 2 € if they did not improve.

Analysis

Squares were considered signal trials and diamonds were considered noise trials. This yielded 40 signal and 40 noise trials per SOA for the threshold estimates and 300 signal and 300 noise trials per training sessions. For the calculation of sensitivity (*d'*) and response bias (*c*), we used the loglinear correction to correct for extreme false alarm or hit rate proportions (Hautus, 1995). For subjective awareness, we calculated the mean PAS rating for correct and for incorrect responses, respectively. The mean PAS rating is suited to assess how subjective awareness changes gradually with learning. Furthermore, we split the PAS ratings into trials on which the target stimulus was minimally detected (PAS>=2) and trials in which the subjects clearly saw the target stimulus (PAS=4) and thus, task-relevant information was subjectively available. For brevity, we refer to these splits as 'subjective detection' and 'subjective discrimination', respectively. Splitting the data this way allows us to investigate whether increases in subjective awareness are due to improved subjective detection and/or improved subjective discrimination. Furthermore, it allows for a more stringent comparison of the objective discrimination task with the subjective ratings, since objective discrimination can be directly compared to subjective discrimination.

To directly evaluate the size of the learning effects in subjective detection and subjective discrimination, we calculated the respective gain at the trained and untrained location, which takes into account the pretraining levels of subjective awareness: ([# of trials with PAS>=2 for posttraining or transfer] divided by [# of trials with PAS>=2 pretraining]) for subjective detection, and ([# of trials with PAS=4 for posttraining or transfer] divided by [# of trials with PAS=4 pretraining]) for subjective discrimination.

We also plotted the rate of incorrect trials against the rate of correct trials at three levels of visibility to obtain *ROC* curves. This allows us to determine how well PAS ratings predict accuracy. Here, the 'hit rate' refers to the percentage of correct trials with a high PAS rating, and the 'false alarm rate' refers to the percentage of incorrect trials with a high PAS rating. This is similar to the procedure known as Type II *ROC* analysis (Galvin, et al., 2003), where the relationship between confidence in one's response and the accuracy of that response is investigated by plotting the rate of correct responses with a high confidence rating (hits) against the rate of incorrect responses with a high confidence rating (false alarms). By varying the criterion at which a PAS rating is considered high (PAS=1 vs. PAS=2, 3, 4; PAS=1, 2 vs. PAS=3, 4; PAS=1, 2, 3 vs. PAS=4), we obtain three inflection points, to which we fitted *ROC* curves using a proper binormal model (Metz & Pan, 1999; Pesce & Metz, 2007) in ROCkit (Kurt Rossmann Laboratories for Radiologic Image Research, http://xray.bsd.uchicago.edu/krl/). This model assumes likelihood ratio as the decision variable. From the *ROC* curves, we calculated the Area under the curve (*AUC*), which indexes how well correct responses can be distinguished from incorrect responses based on PAS ratings. Cells for which the maximum likelihood estimation in ROCkit did not converge were replaced by the average *AUC* of the respective group for further analyses.

In all rmANOVA with more than one degree of freedom we used the Greenhouse-

Geisser correction and report adjusted degrees of freedom and p values. All p values for t-tests are Bonferroni corrected for multiple comparisons, unless otherwise stated.

4

EXPECTATIONS CHANGE THE SIGNATURES AND TIMING OF ELECTROPHYSIOLOGICAL CORRELATES OF PERCEPTUAL AWARENESS

Previous experience allows the brain to predict what comes next. How these expectations affect conscious experience is poorly understood. In particular, it is unknown whether and when expectations interact with sensory evidence in granting access to conscious perception, and how this is reflected electrophysiologically. Here we parametrically manipulate sensory evidence and expectations while measuring event-related potentials (ERPs) in human subjects to assess the time course of evoked responses which correlate with subjective visibility, the properties of the stimuli, and/or perceptual expectations. We found that expectations lower the threshold of conscious perception and reduce the latency of neuronal signatures differentiating seen and unseen stimuli. Without expectations, this differentiation occurs around 300 ms and with expectations around 200 ms post stimulus in occipito-parietal sensors. The amplitude of this differentiating response component (P2) decreases as visibility increases, regardless of whether this increase is due to enhanced sensory evidence and/or the gradual build up of perceptual expectations. Importantly, at matched performance levels, responses to seen and unseen stimuli differed regardless of the physical stimulus properties. These findings indicate that the latency of the neuronal correlates of access to consciousness depend on whether access is driven by stimulus saliency or by a combination of expectations and sensory evidence.

Melloni L, Schwiedrzik CM, Müller N, Rodriguez E, Singer W (2011). Journal of Neuroscience, 31(4):1386-96.

Conscious perception is not solely determined by stimulus saliency. Strong stimuli can remain unnoticed if attention is deployed elsewhere as shown in the attentional blink or change blindness paradigms, and weak sensory stimuli can be readily perceived if they are attended to (Carrasco, et al., 2004). Consequently, perceptual awareness is proposed to depend on two factors (Dehaene, et al., 2006): the intensity of sensory stimulation and top-down attention, which enhances sensory processing. However, evidence suggests that attention may not be the only top-down factor that determines perception. Everyday experience indicates that recognition is greatly facilitated if one knows what to expect. In laboratory settings, when subjects are confronted with fragmented black and white images of an object, they may fail in perceiving the object. However, once the object has been identified it pops out and will henceforth be recognized immediately (the Eureka effect, Ahissar & Hochstein, 2004; Dolan et al., 1997). Similarly, providing subjects with the name or category of the object in anticipation of the stimulus lowers the threshold for its visibility (Eger, Henson, Driver, & Dolan, 2007; Esterman & Yantis, 2010). In patients with visual extinction due to parietal lesions, information held in working memory can reduce the effects of extinction on awareness (Soto & Humphreys, 2006). These observations suggest that perceptual thresholds depend not only on attention but also on a priori knowledge. Such top-down effects agree with theories that consider perception as a Bayesian process in which sensory information is matched with priors, i.e., with prior knowledge and expectations (Friston, 2010; Gregory, 1997; Kveraga, et al., 2007). An increasingly popular implementation of this Bayesian integration is predictive coding, in which the brain constructively predicts upcoming sensory input rather than passively registering it (Friston, 2005; Rao & Ballard, 1999): Top-down predictions "explain away" lower level representations through recurrent interactions such that the mismatch between expected and observed evidence (prediction error) is minimized. Accordingly, weaker neural responses to anticipated than to unexpected stimuli have been reported (Alink, et al., 2010; Garrido, et al., 2009).

Here, we test whether the threshold of perceptual awareness is modulated by previous experience. In particular, we examined how sensory evidence and top-down expectations, respectively, influence the threshold of awareness, and whether the two factors modulate brain activity differently. To this end, we measured electroencephalographic (EEG) activity in a visual paradigm in which we generated perceptual hysteresis by parametrically increasing and then decreasing the sensory evidence for an initially hidden stimulus. Under these conditions, an expectation is built up once the subject has perceived the stimulus and this in turn increases the visibility of subsequent lower contrast stimuli. This allows for contrasting brain states with and without expectations and with perceived and non-perceived stimuli for identical stimulation conditions. The results confirm that expectations lower the threshold for perceptual awareness and provide new evidence on the electrophysiological signatures associated with conscious perception. Importantly, the amount of prior knowledge and ensuing expectations determine whether the electrophysiological signatures of awareness occur early or late after stimulus presentation.

Materials and Methods

Participants

Sixteen volunteers participated in the main EEG study (mean age 26± 2 years; 10 female; 14 right-handed). Of those, nine participants also took part in a control EEG experiment (5 female, 8 right-handed). Ten additional participants took part in a behavioural control experiment (mean age 26± 5 years; 5 female; 9 right-handed). All participants had normal or corrected-to-normal vision and had no history of neurological or psychiatric disorders. Participants gave written informed consent before the experiment. The study protocol was conducted in accordance with the declaration of Helsinki and approved by the local ethics committee.

Material

Stimuli were presented on a TFT monitor (resolution 800×600) at a refresh rate of 60 Hz and located approximately 1 meter from the subjects. The visible screen size subtended

16.197°×12.615° in the horizontal and vertical plane. Presentation software (version 10.3) was used for stimulus presentation and response collection.

Stimuli were either letters, numbers or symbols embedded into a background field of random noise (Figure 4-1a). We parametrically manipulated stimulus visibility by modifying the dot density of the target shape, while keeping the dot density of the background constant. This results in a parametric modulation of the signal to noise ratio, based on which we created six different visibility levels. For the lowest visibility level target dot density was similar to that of the surrounding background, creating the perception of just a field of random dots. Visibility was increased by linearly decreasing the dot density of the target. Differences in dot density between target and background provide a grouping cue for segmentation allowing to perceive the illusory contours of the target. To avoid sensory adaptation the location of the dots and of the target were randomized on each trial (i.e., the presentation of each stimulus). All stimuli were created using Matlab (The MathWorks, Inc., Natick, MA, USA).

Stimuli were displayed at the center of the screen surrounded by a grey background and comprised 13.134°×10.626°. Forty different stimuli were used (25 letters, 9 numbers and 6 symbols). In addition, target stimuli were presented in three different sizes, with an absolute vertical extent of the target of 75 pixels, 100 pixels, and 125 pixels, corresponding to a height of approximately 1.637°, 2.182°, 2.726° visual angle, respectively.

Procedure

In all experiments, the subjects' task was to rate the visibility of the target stimuli on the four-point Perceptual Awareness Scale (PAS, Overgaard, Rote, Mouridsen, & Ramsøy, 2006): "1" corresponds to "No experience of the stimulus, "2" corresponds to "Brief glimpse of the stimulus but could not recognize what it was", "3" corresponds to an "Almost clear impression of the stimulus" and "4" corresponds to a "Clear Impression of the stimulus". These questions were aimed at the subjects' fine-grained phenomenal impression of the stimulus. Thus, our scale dif-

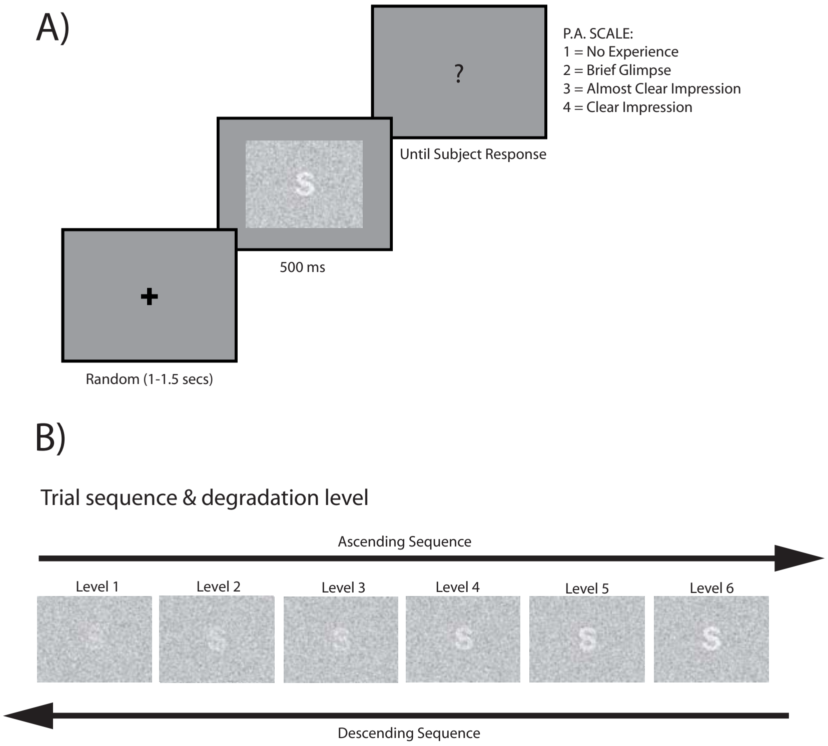

Figure 4-1. Experimental design. (a) The task was to rate the visibility of target stimuli embedded in a field of random noise. After a fixation cross of variable length was presented, the target stimuli appeared for 500 ms, followed by a question mark indicating subject to give their visibility ratings. To assess visibility we used the Perceptual Awareness Scale (Overgaard, et al., 2006). (b) Bottom-up information was manipulated by degrading the stimulus visibility. Six degradation levels were used. Degradation was achieved by decreasing the dot density of the target stimuli while keeping dot density of the background constant. The successive increases and decreases in stimulus degradation are referred to as "sequence". For the main experiment, the same stimulus was used during a given sequence, whereas in the control EEG experiment, the stimulus was varied every trial.

fers from previously used scales that assessed only confidence (e.g., Wilimzig, et al., 2008). We provided subjects with verbal as well as written instructions. In order to assure similar criteria between participants, we extensively discussed the operationalization of PA with the subjects before the experiment started. In particular, we emphasized the distinction between point 2 and point 3 on the scale, the criterion separating no recognition from explicit recognition of the target. Furthermore, we pointed out that they should rate only their perception of the target, and that there were neither right nor wrong answers. Responses were given by button presses on a keyboard, with a stress on accuracy rather than speed. In order to verify that subjects used the PA scale properly, some of the stimuli were clearly visible which should lead to high visibility ratings, whereas highly degraded stimuli that could not be perceived should be reported as invisible. This allowed us to assess the reliability of the subjects' judgments.

A trial started with a fixation cross presented on a grey background for a random duration of 1000-1500 ms. Then, the random dot field containing the target (letter, number or symbol) was presented for a duration of 500 ms. After that, a response screen containing a question mark indicated to the subjects that they should rate the visibility of the target (Figure 4-1a).

To study the differential effects of sensory evidence and expectancy, in a first experiment, we presented the targets in a sequence of 11 trials (for simplicity we will refer to this as 'sequence'). The same target was used throughout the sequence. The visibility of the target progressively increased, from trial 1 to 6 and then decreased again until trial 11 (Figure 4-1b). Thus, once the target became visible on the ascending part of the sequence, subjects could generate target specific expectancies, the effect of which could be tested in the descending part (hysteresis effect; Kleinschmidt, Büchel, Hutton, Friston, & Frackowiak, 2002). The generation of target specific expectancies can also be regarded as a form of priming (Grill-Spector, Henson, & Martin, 2006). We presented 120 sequences that were divided into six blocks (20 sequences each). After each block a break of variable length was introduced.

To distinguish the specific contribution of the expectancy from the sequential change in sensory evidence, we ran a control experiment in which we used the same sequences but changed the identity of the target in each trial, preventing the generation of target specific expectancies. If changes in visibility are related to expectancy and prediction about stimulus identity, and not to a mere sequence effect, no signs of hysteresis should be observed in this case. Stimuli were identical to those used in the previous study and were repeated for an equal number of times throughout the experiment. In total, 120 sequences were presented, divided into six blocks. For both studies we collected EEG data while subjects performed the task.

To further evaluate whether target specific expectations lead to changes in the threshold of perceptual awareness or changes in response bias, we ran a behavioral experiment in which we presented stimuli, only in descending order of visibility, in three conditions: (i) random sequence, (ii) predictable sequence, and (iii) interrupted sequence. In the random sequence, target identity was randomly changed for each successive degradation level, precluding any form of target specific expectations. In the predictable sequence, the same target was presented at all degradation levels. In the interrupted sequence, one unpredicted target per sequence was introduced with a probability of 20% at one of the degradation levels. This creates a scenario of high certainty regarding the future target (80% probability) which is unfulfilled in some trials. In other words, predicted and interrupted sequences only differ in that one unexpected target was introduced per sequence. The six degradation levels were presented in descending visibility order only (from degradation level L6 to L1) promoting the development of target specific expectations. Random and predictable sequence differed from the second position in the sequence onwards, whereas predictable and interrupted sequences could differ from the third position in the sequence. This was done to assure that subjects develop target expectations and thereby allowing us to evaluate the associated visibility to an unexpected target. In total, we presented 180 sequences: 20 random, 80 expected and 80 interrupted sequences (20 unexpected trials per degradation level), divided into four blocks. The three types of sequences were randomized throughout the experiment. The same stimuli were used for all sequences, and they were the same as those used in the two experiments described above.

Our rationale for these control experiments was that if subjects have a tendency to respond with higher visibility ratings (response bias) simply because they think they know what will be seen next, we would expect higher ratings for unexpected stimuli in the interrupted sequence. Conversely, if expectations truly augment visibility, the ratings of these unexpected stimuli should be the same as those presented under the random sequence.

Data Acquisition

All experiments was conducted in an electrically shielded, sound-attenuated, and dimly lit cabin. A TFT monitor was placed outside the cabin behind an electrically shielded window.

All devices inside the cabin were battery operated to avoid interference of the line frequency. For study 1 and 2, EEG was recorded with a QuickAmp amplifier (Brain Products GmbH, Munich, Germany) using a custom 128-channel elastic cap (Easy Cap GmbH, Herrsching-Breitbrunn, Germany), with average reference and a forehead ground electrode. The electrodes were equally spaced over the whole head and symmetrically placed between both hemispheres. The inferior row of electrodes was placed down to a line from the low cheeks back to several centimeters below the inion, thus providing sufficient coverage of the lower parts of the head. Two additional electrodes were placed on the infra-orbital ridges of the left and right eye, respectively, and two further electrodes were placed on the neck below the edge of the electrode cap. Data were sampled at 1000 Hz and analogue filtered between 0.01 and 100 Hz during recording. Electrode impedances were kept below 10 kΩ. Data were digitally saved on a computer outside the cabin for later offline analysis.

Data analysis

Behavioral data and curve fitting

Although we applied a 4-point visibility scale, for analysis the 4-point scale was re-coded into a single response category (visibility), collapsing responses 1+2 and responses 3+4 as "invisible" and "visible", respectively. This dichotomy was based on explicit recognition of the target in the stimulus. Then, the percentage of the rating "visible" was calculated for each stimulus in the sequence. For the main and control EEG experiment we analyzed our data using an unconstrained repeated measure ANOVA and a constrained parametric model. In the unconstrained ANOVA we treated degradation level and order as two separate factors with six and two levels (ascending and descending sequence) respectively. In the constrained model we created a new dependent variable (inflexion point) by fitting, per subject, a sigmoid function of degradation level to the behavioral data to compute an inflexion point. We then tried to explain differences in the inflexion point data in terms of a repeated measure ANOVA with one factor (ascending vs. descending sequence). The rating scores for visibility were fitted with a generalized sigmoid function defined as:

$$f(x) = \frac{1}{1+e^{-a(x-\varphi)}} + b \quad ,$$

where *a* determines the slope (steepness) of the sigmoid, *b* the offset, and *j* the horizontal shift of the function (the threshold). The sigmoid function was fitted to best represent the data points by using a gradient descent method to minimize approximation error. Before the fitting, values were normalized to the interval [0..1]. For our purposes, the most important parameter of the sigmoid was *j,* because it determines the threshold, i.e., the point on the horizontal axis at which the function reaches half of its maximum. If predictions influence visibility this point should shift along the horizontal axis for visibility functions computed from ascending and descending sequences, respectively.

For the behavioral control experiment, we ran an unconstrained repeated measure ANOVA with factors sequence type (random, predictable, interrupted) and degradation level (4 levels, from L4 to L1). Only four degradation levels were included since only in those interruptions of expectations could occur (as required to promote target specific expectations).

In all analyses of variance with more than one degree of freedom we used the Greenhouse-Geisser correction. We report adjusted degrees of freedom and adjusted *p* values.

In order to investigate whether the increase in visibility observed in the main EEG experiment is the consequence of a top-down perceptual expectation (the generation of target specific expectancies) as opposed to mere passive repetition of the stimuli, we run a control analysis taking advantage of the fact that the number of repetitions for each stimulus was the same in the main EEG and control EEG experiment. For the control EEG experiment, we tested whether visibility increased systematically with the number of repetitions. Each stimulus was recoded depending on whether it had been previously presented (old) or not (novel). Old stimuli were further subdivided depending on the number of previous repetitions (from 1 up to 10). Thus, we could explicitly test whether the number of repetitions had an effect on recognition in a repeated measures ANOVA with factors repetition, order and degradation level. We further compared zero repetitions (new) with the maximal amount of repetitions (10, old) and when repetitions were regrouped into few (1-3), intermediate (4-6) and many repetitions (7-9).

Analysis of event-related potentials

For the analysis of event-related potentials (ERPs) of the main and control experiment, the continuous EEG signal was band-pass filtered (0.3-20 Hz) with a phase shift-free Butterworth filter (time constant [s] = 0.530526, slope [db/oct] = 24). For each degradation level, the filtered signal was then segmented in series of 1200 ms long epochs. Each epoch started 200 ms before the onset of the stimuli. Trials containing voltages exceeding ± 100 µV, or gradients exceeding ± 50 µV were rejected. The automatic artefact rejection was supplemented by visual inspection of every trial to reject cases with electrode drifts, eye-movements or electromyographic activity. Artifact free trials were averaged per degradation level in synchrony with the onset of the stimuli, and baseline corrected over a 200 ms window. Since the different degradation levels led to different ratios of seen vs. unseen trials we did not consider behavioural performance in this analysis. We solely compared the electrophysiological responses between different degradation levels, assuming that differences reflect both the degradation level and differences in visibility. Linear response modulations are expected for components reflecting physical properties of the stimuli, whereas non-linear responses should reflect the behaviorally assessed visibility (as in Figure 4-2a). Additionally, we directly compared seen and unseen trials (see below). Mean amplitudes were analyzed for P1 (110-130 ms), N1 (130-150ms), P2 (180-270ms) and P3 (300-500 ms) components. We focused on those ERP components because previous studies have related them either to visual awareness (for a review see Koivisto & Revonsuo, 2010) or to predictions (Garrido, et al., 2009). To increase statistical power, mean amplitude was computed over groups of electrodes representative of the topography of each scalp component. Regions of interest (ROIs) were selected based on the compound waveform of all conditions to minimize selection bias favouring a particular condition. For P1 we selected a group of occipito-parietal electrodes (P5, P7, PO7, PO5, P6, P8, PO6, PO8), for N1 the occipito-central electrodes (O1, OZ, O2, O9, Iz, O10), for P2 the occipito-parietal electrodes (P5, P7, PO7, PO5, PO9, O1, O9, Iz, Oz, O2, O10, P6, P8, PO6, PO8, PO10), and for P3 the parieto-central electrodes (CP3, CP1, P1, P3, Pz, CPz, PO1, PO2, P2, P4, CP2, CP4). All EEG processing was done using Brain Vision Analyzer 1 (Brain Products GmbH, Munich, Germany). The Matlab toolbox EEGLAB was used for visualization and topographic plots (Delorme & Makeig, 2004).

We also directly studied amplitude modulations in relation to seen and unseen stimuli for the P1, N1, P2 and P3 components. For this analysis, only data from the main experiment were included. We further restricted our analysis to Level 4 stimuli in the ascending and to Level 3 stimuli in the descending part of the stimulation sequence because the ratio between seen and unseen stimuli was similar for these levels (~50%). Preprocessing steps were the same as those described above for the whole sequence. Trials in which subjects identified the stimuli (visibility rating 3 and 4) were categorized as seen, and trials in which subjects could not identify the stimuli (visibility rating 1 and 2) were classified as unseen. Pairwise comparisons confirmed that after artifact rejection the ratio of seen vs. unseen trials were similar for Level 4 ($T(15)=0.816$, $p=0.427$) and Level 3 ($T(15)=1.020$, $p=0.324$).

As for the behavioral data, we analyzed the electrophysiological data - mean amplitude per ERP component - with an unconstrained ANOVA treating degradation level and order as two separate factors. Thus, both for the behavioral and electrophysiological data, we looked for the main effect of degradation, order and their interaction. For some components (P1, P2 and P3), region of interest was included as a factor. In all analyses of variance with more than one degree of freedom we again used the Greenhouse-Geisser correction (Picton et al., 2000). We report adjusted degrees of freedom and adjusted p values.

Acknowledging that the preselection of electrodes and components might have precluded observing additional components related to awareness (e.g., with a different time course than the ones already investigated) and to further evaluate the time points at which changes in electric field strength distinguish between seen and unseen trials at threshold, we calculated global field power (GFP, Lehmann & Skrandies, 1980) for level 4 ascending sequence (AS) and level 3 descending sequence (DS). GFP allows for investigation of amplitude differences between conditions without any preselection of electrodes, and thus to contrast conditions across the whole scalp. GFP is equivalent to the spatial standard deviation of the scalp electric field and is calculated as the square root of the mean of the squared value recorded at each electrode (vs. the average reference). To statistically test for differences in GFP between experimental conditions we performed a series of paired nonparametric statistical analyses based on randomization tests (Manly, 1991). The empirical probability distribution of the differences in GFP between seen and unseen conditions was created separately for L4(AS) and L3(DS) by randomly shuffling the conditions within subjects in the original data, and recomputing the GFP for these shuffled conditions. This procedure was repeated 65536 times (2^n, n being the number of subjects; in our case 16) so as to be able to estimate the probability ($p<0.01$) that the difference between the two distributions is observed by chance (M. M. Murray, Brunet, & Michel, 2008). We applied an additional criterion of temporal stability for ten consecutive time points, corresponding to 10 ms at 1000 Hz sampling rate (for a similar approach see M. M. Murray, et al., 2008; Pourtois, Delplanque, Michel, & Vuilleumier, 2008).

RESULTS

Behavioral results

Figure 4-2a shows the percentage of responses signaling a stimulus as seen as a function of the degradation level of the stimuli and their place in the presentation sequence, when the target is predictable. A significant effect of degradation level was observed confirming that our manipulation of visibility was successful (degradation level: $F(1.702, 25.535) = 407.375$, $p<0.0001$): The percentage of "seen stimuli" increases significantly as degradation level decreas-

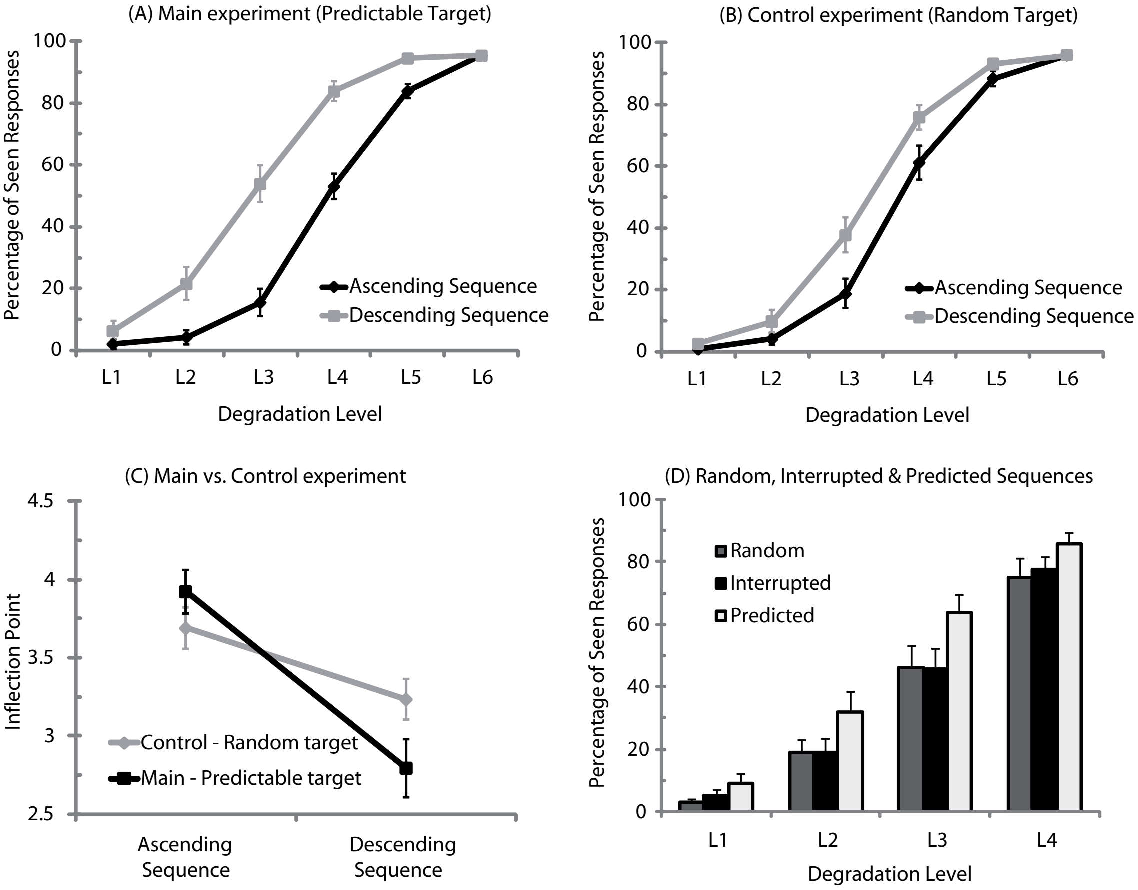

Figure 4-2. Behavioural results. Percentage of seen stimuli per degradation level for ascending sequences and descending sequences for the main experiment with predictable sequences (a) and the control EEG experiment with random sequences (b). Black lines refer to ascending sequences, gray lines to descending sequences. (c) Subjective threshold data (inflection point) for the main (predictable target) and control (random target) EEG experiment. (d) Percentage of seen stimuli per degradation level for random, interrupted and predicted sequences. In the random sequences, target stimuli varied at each degradation level leading to uncertainty regarding stimulus identity. In the predictable sequences, the same target was used throughout the sequence, thus generating high certainty about stimulus identity. In the interrupted sequences subjects had high certainty regarding the identity of the upcoming target (80% probability) that was in 20% of the cases unexpectedly unfulfilled. Visibility ratings in the interrupted sequences correspond to ratings given to the unexpected stimuli.

es. For all subjects, the curve relating visibility to degradation levels was well fitted by a sigmoid and exhibits a sharp transition between degradation level 3 and 5 from judgments “mainly invisible” to “mainly visible”. Confirming the non-linearity in the visibility function we found that the rate of increase in subjective visibility was higher between degradation level 3 (L3) and degradation level 5 (L5) - around the threshold of visibility in the ascending sequence-, than for the increase between L2-L3 and L5-L6 ($T(15) = -10.227$, $p<0.0001$) (for a similar approach, see Del Cul, et al., 2007). A clear hysteresis effect was observed between the ascending and the descending part of the stimulus series ($F(1,15)=178.819$, $p<0.0001$): subjects reported to detect a stimulus more often in the descending (59%) than in the ascending sequence (42%). The increase in visibility due to the order effect (hysteresis) was not the same for all degradation levels (degradation level*order: $F(2.064,30.956)=41.900$, $p<0.0001$), and this effect is explained by a significant quadratic function ($F(1,15)=126,667$, $p<0.0001$). Pairwise comparisons contrasting the same physical stimulus (degradation level) during the ascending and descending sequence revealed that all stimuli but L1 – that was clearly not distinguishable - showed hysteresis. In order to quantify this effect we fitted sigmoid curves, per subjects, to the ascending and descending series and set the subjective threshold at the inflection point of the sigmoid. The threshold for the ascending sequence was at degradation level 3.96 and for the descending sequence at 2.76. Thus, the gain in visibility caused by prior knowledge of the stimulus corresponds to 1.2 degradation steps. Accordingly, the threshold of visibility was reached at degradation level 4 in the ascending (53.02% seen stimuli) and at degradation level 3 in the descending part (54.01% seen stimuli).

To evaluate whether the hysteresis effect results from a memory–based prediction of stimulus identity or from the sequential presentation of the stimuli we used our control experiment in which the target stimuli differed in each trial, both for the ascending and for the descending sequence. We ran an ANOVA with factors predictable, degradation level and order. As expected, in the ascending part of the curve, main and control experiment showed comparable percentages of seen stimuli (Figure 4-2b). However, during the descending sequence the percentage of seen responses was significantly higher in the main (59.36%) than in the control experiment (52.42%) (interaction predictable and order: $F(1,8)=28.849$, $p=0.001$; interaction predictable*order*degradation level: $F(2.069,16.551)=10.273$, $p=0.001$). Similar results were obtained when comparing the inflection point data for the control and main experiment (Figure 4-2c). Visibility thresholds were similar in the ascending part, but differed significantly in the descending part ($T(8)=3.363$, $p=0.01$). In summary, the psychophysical functions exhibit clear signs of perceptual hysteresis, suggesting that the buildup of an internal representation of the upcoming stimulus enhances visibility.

Given that the degree of uncertainty regarding the stimulus identity differs between predictable and random sequences, it can be argued that the introduction of target specific expectations do not lead to real changes in the threshold of perceptual awareness but simply promote subjects to change their response criterion across the predictable sequences. To directly evaluate this possibility we turned to our behavioral experiment in which stimuli were presented, only in descending visibility order, in three conditions: (i) random sequence, (ii) predictable sequence, and (iii) interrupted sequence. The critical manipulation was the degree of certainty regarding the identity of the upcoming stimuli. Target stimuli either randomly varied per degradation level (random sequences), creating high uncertainty about the stimulus identity or were kept the same throughout the sequence (predictable sequence), creating high certainty about the stimulus identity. Crucially, for the interrupted sequence subjects had high certainty regarding the identity of the upcoming target (80% probability) that was unexpectedly unfulfilled. As can be seen in Figure 4-2d, visibility ratings for the unexpected stimuli in the interrupted sequences were the same as for stimuli in the random sequences at each degradation level; both showed lower visibility ratings than for the predicted stimuli in the predictable sequence (sequence type: $F(1.635,14.712)=19.358$, $p<0.001$; degradation level: $F(2.122,19.100)=163.902$, $p<0.001$; sequence type*degradation level: $F(3.321,29.891)=1.98$, $p=0.133$). This shows that the increases in visibility ratings reflect a true change in the threshold of visibility due to target specific expectations and not a simple response bias, since such a bias would carry over to unexpected targets during the interrupted sequences.

An additional analysis revealed that the mere number of repetitions of a stimulus cannot explain the improved identification levels observed in the main experiment (see Figure 4-3).

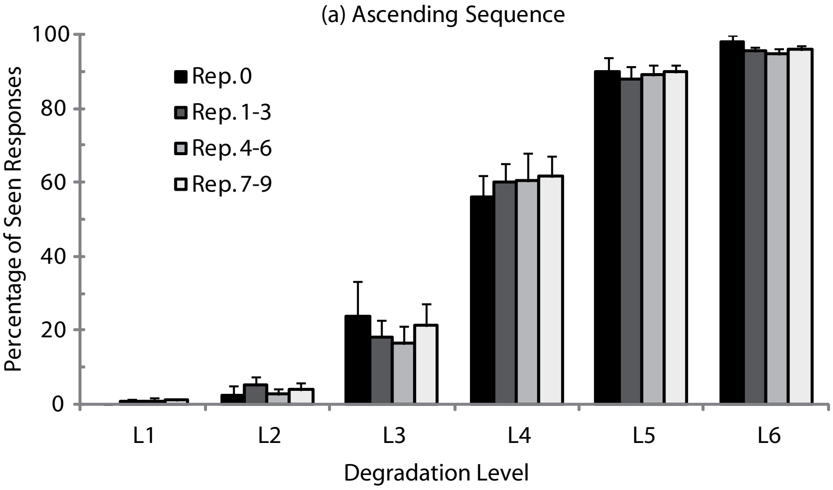
(a) Ascending Sequence

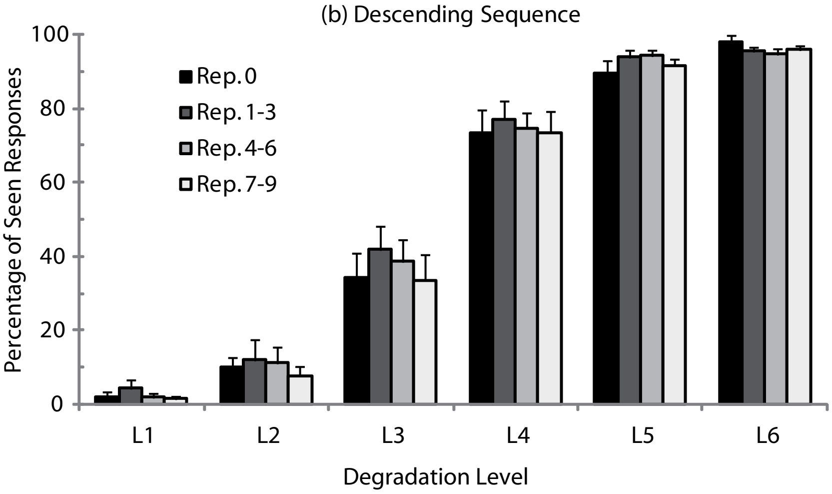
(b) Descending Sequence

Figure 4-3. Percentage of seen responses as a function of degradation level and repetition for the ascending and descending sequence. No effect of or interaction with repetition were found in a repeated measures ANOVA with factors repetition, order and degradation level [repetition $F(3.798,30.381)=0.600$, $p=0.657$; repetition x order $F(3.373,26.984)=1.187$, $p=0.152$; repetition x degradation level $F(6.338, 50.704)=0.663$, $p=0.688$; repetition x order x degradation level $F(5.965,47.716)=0.908$, $p=0.497$]. Further analyses confirmed this results when we compared zero repetitions (new) with the maximal amount of repetitions (10, old) (all $p>0.392$), or when repetitions were regrouped into few (1-3), intermediate (4-6) and many repetitions (7-9) (all $p>0.1$). The above figure exemplifies that new stimuli lead to similar recognition levels as stimuli repeated between 7 and 9 times, both in the ascending and the descending sequence. In sum, we find no behavioral effect of mere passive repetition in the control experiment. Thus, since the number of repetitions was the same in both experiments (control and main), we conclude that mere repetition cannot explain the increased visibility observed in the main experiment. Of note is that mere repetition of stimuli can lead to facilitatory effects (e.g., identification) when intervening stimuli are present, even when prime and test are presented several days apart (van Turennout, Ellmore, & Martin, 2000). This casts doubt upon the possible concern that facilitatory effects in the control experiment might not have been observed because of the intervening items. Table 4-1 summarizes the mean visibility ratings and the mean number of trials per repetition.

Repetition	Visibility Mean	Visibility Std.Error	Count Mean	Count Std.Error
Rep0	.48	.03	10.98	.06
Rep1	.50	.03	10.87	.07
Rep2	.49	.03	10.93	.06
Rep3	.49	.03	10.83	.04
Rep4	.48	.02	11.00	.07
Rep5	.49	.03	10.94	.07
Rep6	.48	.03	10.83	.07
Rep7	.49	.03	10.89	.09
Rep9	.49	.03	10.86	.07
Rep9	.47	.02	10.97	.05
Rep10	.48	.02	10.89	.03

Table 4-1. Mean visibility and mean amount of trials per repetition (the amount of trials did not differ significantly between conditions, all $p>0.4$).

Event-related potentials

We focused on two questions: i) which ERP components correlate with the sharp transition in visibility and ii) how ERP components are affected by the change in visibility caused by the buildup of prior expectations. To determine brain activity (ERP components) that directly relates to *subjective visibility* we used the following criteria: a) ERP components should invariably display differential amplitudes for seen and unseen stimuli irrespective of whether visibility is solely caused by stimulus saliency or by a combination of saliency and expectancy, and (b) exhibit a similar response profile in terms of their amplitude as the behaviorally assessed visibility. In particular, they should exhibit non-linear amplitude modulations with a sharp transition at the threshold of visibility. In contrast, linear response modulations were expected for components reflecting *physical properties* of the stimuli.

Inspection of the ERPs revealed four major components: P1, N1, P2 and P3. For each of these components we examined whether they exhibited amplitude modulations related to the physical properties of the stimuli (degradation level) or subjective visibility, the order of the stimuli (hysteresis), or the interaction between stimulus visibility and order.

The P1 component showed no modulation neither for degradation nor order and also no interaction between these factors (all $p>0.2$, Figure 4-4a). There was only a trend towards reduced amplitude for the stimuli with lowest visibility. This confirms previous results showing that P1 is not related to perceptual awareness (Del Cul, et al., 2007; Sergent, Baillet, & Dehaene, 2005). The N1 component, when averaged across the ascending and descending sequence, respectively, had a reduced amplitude for the descending as compared to the ascending part (effect of order $F(1, 15)=6.122$, $p=0.026$, Figure 4-4b). N1 changes further showed an interaction between degradation level and order (degradation level x order: $F(3.335, 50.030)=2.822$, $p=0.043$), suggesting that the modulation of this component differs for the ascending and descending part. A clear effect of degradation existed for the ascending part ($F(3.647, 54.702)=3.386$, $p=0.018$). The N1 amplitude decreased linearly with increasing visibility ($F(1, 15)=9.639$, $p=0.007$) as would be predicted by a decrease in prediction error under predictive coding (Garrido, et al., 2009). In contrast, no effect of degradation was observed for the descending part ($p>0.9$). Thus, in the ascending part the N1 modulation mainly reflects differences in sensory stimulation, while in the descending part it seems to reflect expectations. To test whether expectations differentially modulate responses to the same stimuli we performed pairwise comparisons between responses to the same physical stimuli (degradation level) in the ascending and descending parts. This analysis revealed amplitude modulations only for those stimuli that were barely visible but not for the highly visible stimuli [(L1 ($p=0.058$), L2 ($p=0.010$), L3 ($p=0.037$), L4 ($p=0.316$), L5 ($p=0.116$)]. This suggests that top-down effects selectively affect the processing of stimuli with poor visibility (bottom-up information). In fact, the generation of a memory trace seems to change responses to sensory stimulation in such a way that once expectancy is established, stimuli evoke similar N1 amplitudes irrespective of degradation level. In summary, the amplitude of the N1 component is modulated by expectations indicating that prior knowledge affects responses already 130 ms after stimulus presentation. Importantly, however, the modulation of the N1 component does not reflect

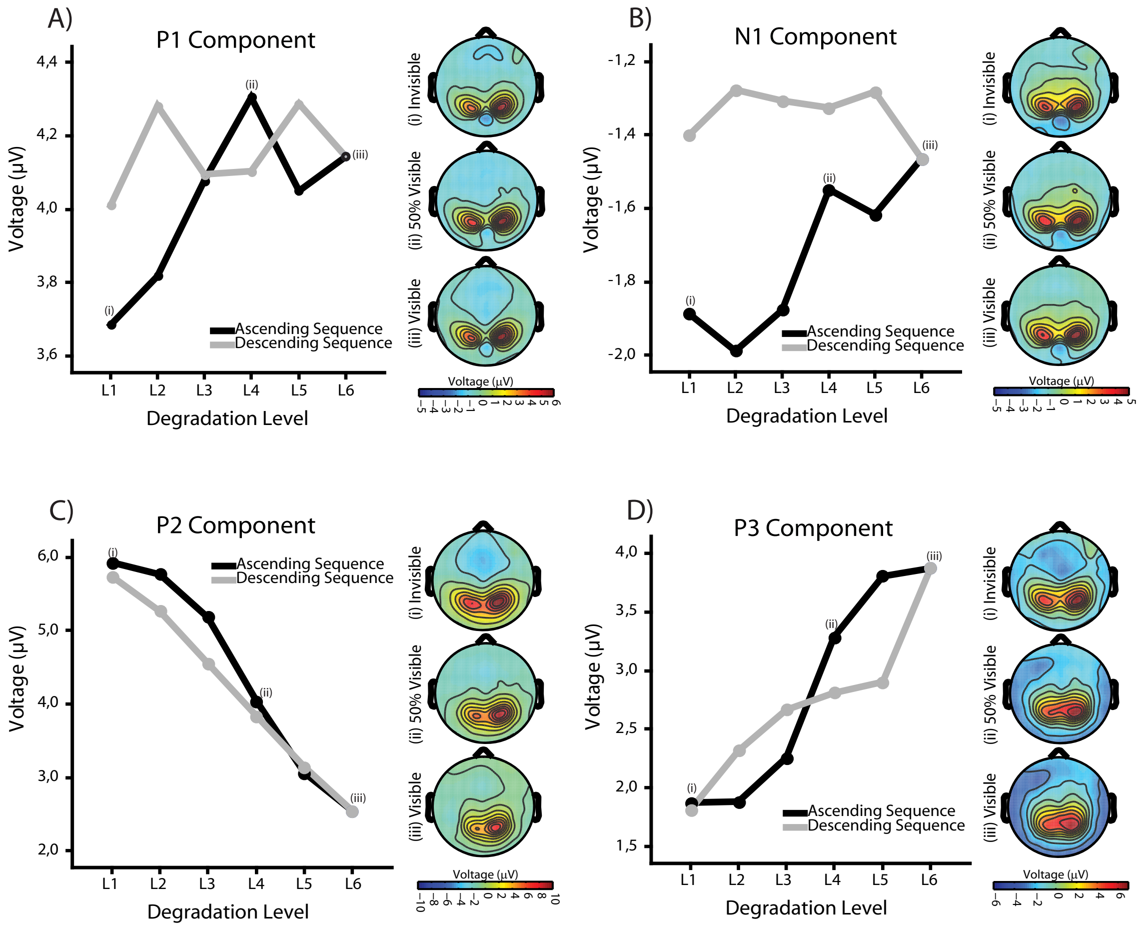

Figure 4-4. Electrophysiological Results (Main Experiment). Amplitude modulations as a function of degradation level and order (ascending sequence/descending sequence) for P1 (a) N1 (b) P2 (c) and P3 (d). Scalp topographies are shown for ascending sequence at degradation level 1 (invisible), degradation level 4 (threshold of visibility) and degradation level 6 (highest visibility).

subjective visibility. The latter follows a sigmoidal function while the N1 component follows the linear change of stimulus degradation, at least in the ascending sequence.

The other components, P2 and P3, showed clear but different amplitude modulations related to degradation level and subjective visibility. For the P2 component a within subject ANOVA with factor degradation level, order and electrode location (region of interest, ROI) (3 levels: left, right and central sensors) revealed that P2 amplitude decreases significantly with decreasing degradation (degradation level: $F(1.530,22.954)=52.667$, $p<0.001$), and this relation can be well described by a cubic function (*F*(1,15)=15.019, *p*=0.001) (see Figure 4-4c and 4-5), thus resembling the non-linear behavior of the psychometric curve. This suggests that the amplitude of the P2 component is inversely related to the visibility of the stimuli. P2 amplitude was also significantly smaller for the descending than for the ascending sequence (order: *F*(1,15)=12.641, *p*=0.003), suggesting that perceptual expectations also attenuate this ERP component. However, the attenuation of the P2 component by prior knowledge depends on the degradation level (degradation level*order: *F*(3.449,51.738)=6.228, *p*=0.001). Pairwise comparisons between responses to the same physical stimuli (degradation level) in the ascending and descending sequence revealed that only responses to barely visible stimuli showed lower amplitudes whereas responses to clearly visible stimuli showed similar amplitudes (L1 (*p*=0.169), L2 (*p*=0.003), L3 (*p*=0.001), L4 (*p*=0.116), L5 (*p*=0.479), Figure 4-4c). This indicates that top-down effects on P2 amplitudes are stronger when bottom-up saliency of the stimuli is low. This suggests that P2 attenuation is mainly related to visibility rather than degradation or expectancy per se. There was also a lateralization effect. P2 amplitudes were higher for the right than for central and left sensors (*F*(1.708,25.615)=10.251, *p*=0.001, Figure 4-5), which is in agreement with previous findings (Freunberger, Klimesch, Doppelmayr, & Holler, 2007). Scalp topographies of P2 modulation are shown in Figure 4-6. Both the lateralization effect and the decreases in amplitude with increasing visibility are clearly visible. It is important to note that contrary to the modulations observed in N1 amplitude where an effect of degradation level was only observed for the ascending sequence, P2 showed amplitude modulations as a function of visibility in both the ascending and descending sequences. Thus, P2 is the first ERP component whose amplitude modulation corresponds well to the behaviorally assessed visibility levels, following the nonlinear function of visibility in both the ascending and the descending sequences.

The P3 component differed from the P2 component in that it showed no main effect of order. Its amplitude increased significantly with decreasing degradation (*F*(2.387,35.798)=35.593, *p*<0.001). In addition, a significant interaction between degradation level and order was found

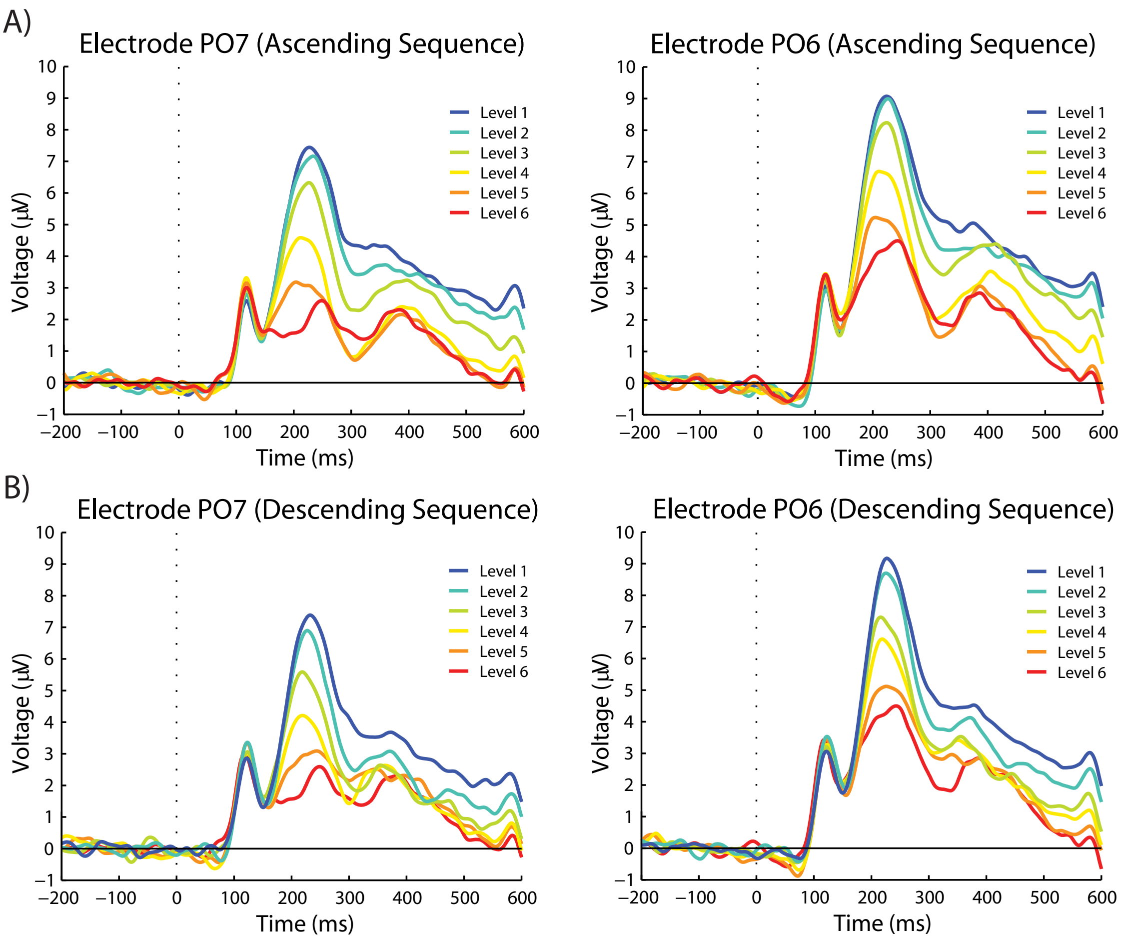

Figure 4-5. Waveform P2 component (Main Experiment). Left (PO7) and right (PO6) parieto-occipital electrodes where P2 modulations can be clearly distinguished. The upper panel shows amplitude modulations per degradation level during the ascending sequence. The lower panel shows amplitude modulations during the descending sequence.

($F(2.588,38.826)=7.179$, $p=0.001$). In the ascending sequence, stimuli which were judged as visible (4 to 6) elicited a significantly higher ($T(15)=-7.072$, $p<0.001$) amplitude than invisible stimuli (1 to 3), with a sharp transition at the visibility threshold between degradation 3 and 4 (Figure 4-4d). In the ascending sequences, no differences in amplitude were found across stimuli judged as visible (all $p>0.08$) or invisible (all $p>0.4$), respectively. This non-linearity replicates previous findings relating P3 amplitude to subjective visibility (Del Cul, et al., 2007). In contrast, when the analysis was carried out on the descending sequence, the separation between visible and invisible stimuli disappeared. Instead, all stimuli but stimulus 1 elicited a P3 response

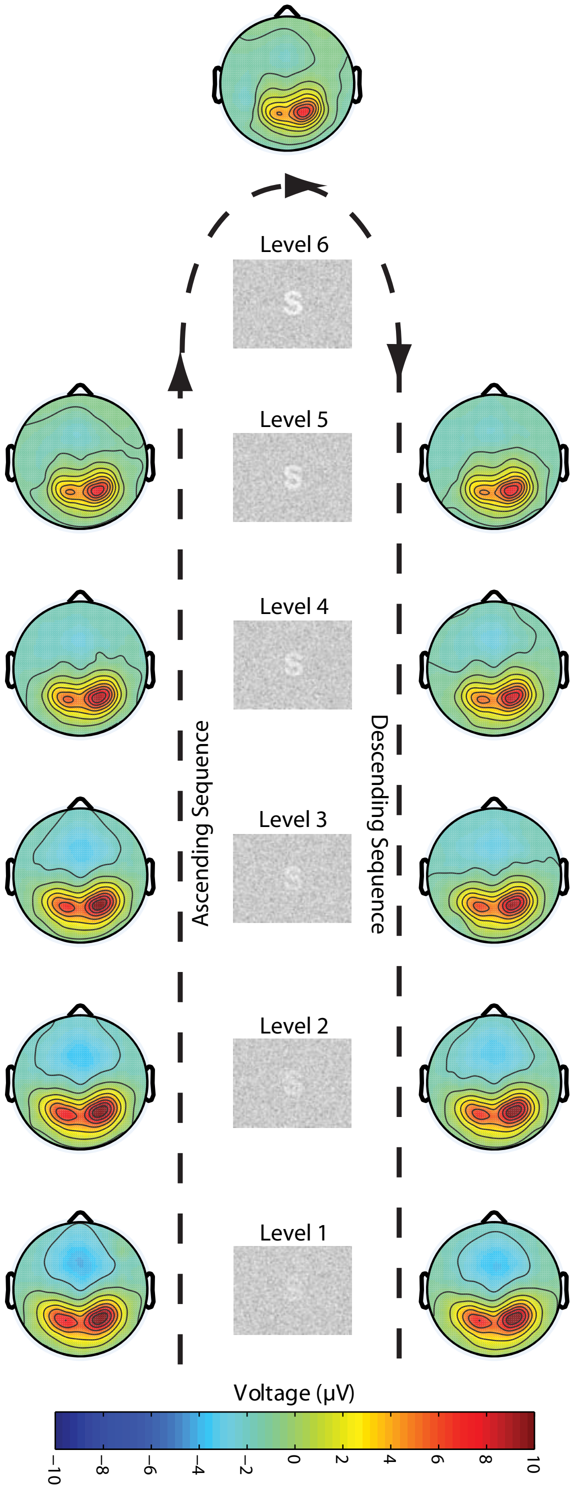

Figure 4-6. P2 Scalp Topographies (Main Experiment). Scalp topographies of P2 amplitude per degradation level and ascending and descending sequence. Clear amplitudes decrements can be observed as a function of visibility.

of similar amplitude, despite marked differences in visibility and physical properties (for L1 against all others, all *p*>0.03 ; for comparisons of L2 to L5, all *p*>0.2; Figure 4-4d). As for the P2 component, there was a lateralization effect. Responses in the right sensors were more pronounced, especially for stimuli with lower visibility.

To further confirm these results we analyzed the different ERP components in the condition where the target stimuli differed for every step in the ascending and descending sequences such that perceptual expectations could not be built up.

As displayed in Figure 4-7, none of the components showed variations in amplitude when comparing responses to stimuli at equal degradation levels in the ascending and descending sequence (all *p*>0.2). In contrast, a clear effect of degradation level was observed for all components (P1: *F*(2.678, 21.423)=4.209, *p*=0.020; N1: *F*(2.925, 23.402)=4.227, *p*=0.016; P2: *F*(1.611, 12.886)=34.809, *p*<0.001; P3: *F*(2.618, 20.940)=17.836, *p*<0.001). With decreasing degradation P1 amplitude increased (*F*(1, 8)=11.459, *p*=0.010), while the amplitude of the N1 and P2 decreased (*F*(1, 8)=11.510, *p*=0.009, *F*(1, 8)=8.804, *p*=0.018, respectively). For the P1 and N1 component the functions relating degradation with amplitude were linear, for the P2 component this function was quadratic. The quadratic relation is mainly explained by a saturation effect at degradation 1 (Figure 4-7c). If this degradation level is removed, the modulation of the P2 component becomes also linear both for the ascending and descending sequence. The modulation of the P3 component followed a cubic function (*F*(1, 8)=6.386, *p*=0.035), with a sharp transition between seen and unseen stimuli, thus correlating well with the psychometric function of visibility. Furthermore, two clear clusters of amplitudes emerged dissociating seen from unseen stimuli: P3 amplitudes were similar for the groups of stimuli with low (1-3) (all *p*>0.05) and high visibility (4-6) (all *p*>0.05) respectively, whereas clear amplitude differences existed between the groups. This was true for both the increasing and

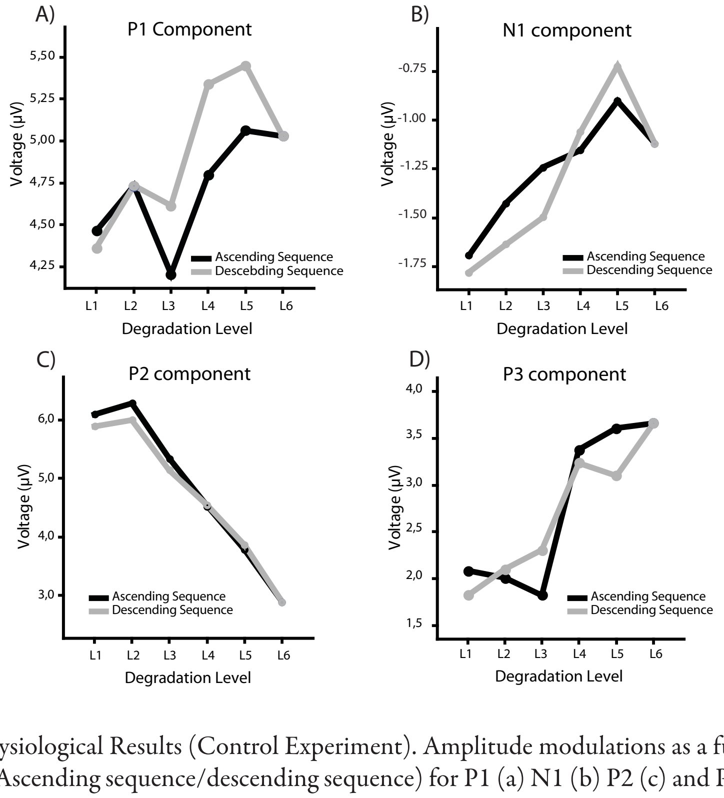

Figure 4-7. Electrophysiological Results (Control Experiment). Amplitude modulations as a function of degradation level and order (Ascending sequence/descending sequence) for P1 (a) N1 (b) P2 (c) and P3 (d).

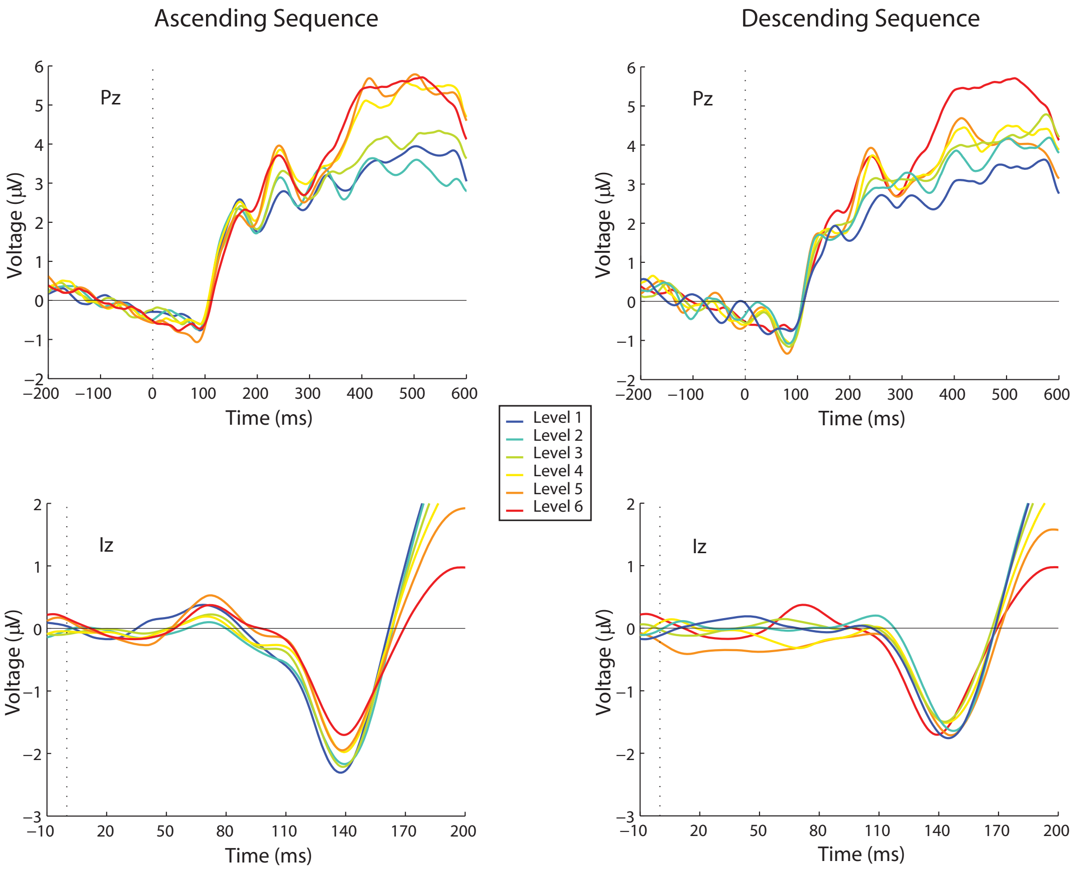

Figure 4-8. Waveforms of N1 and P3 components (Main Experiment). Iz and Pz electrodes showing the N1 and P3 modulations, respectively. Left panels show amplitude modulations per degradation level during the ascending sequence, right panels show amplitude modulations during the descending sequence. Note that the time scale for electrode Iz is different than for electrode Pz. This was done to allow for a better appreciation of the N1 component.

for the decreasing sequence (increasing: $T(8)=-5.481$, $p=0.001$; decreasing: $T(8)=-5.193$, $p=0.001$). Waveform traces can be seen in Figure 4-8 for the main and in Figure 4-9 for the control experiment.

Effects of visibility at threshold

Finally, we directly tested the factor visibility (seen vs. unseen responses) for the data obtained from the main experiment to determine which ERP component varied with subjective visibility when the visual stimulus remains constant. To this end, we used only stimuli at the threshold of subjective visibility: stimuli with degradation level 4 for the ascending sequence, and degra-

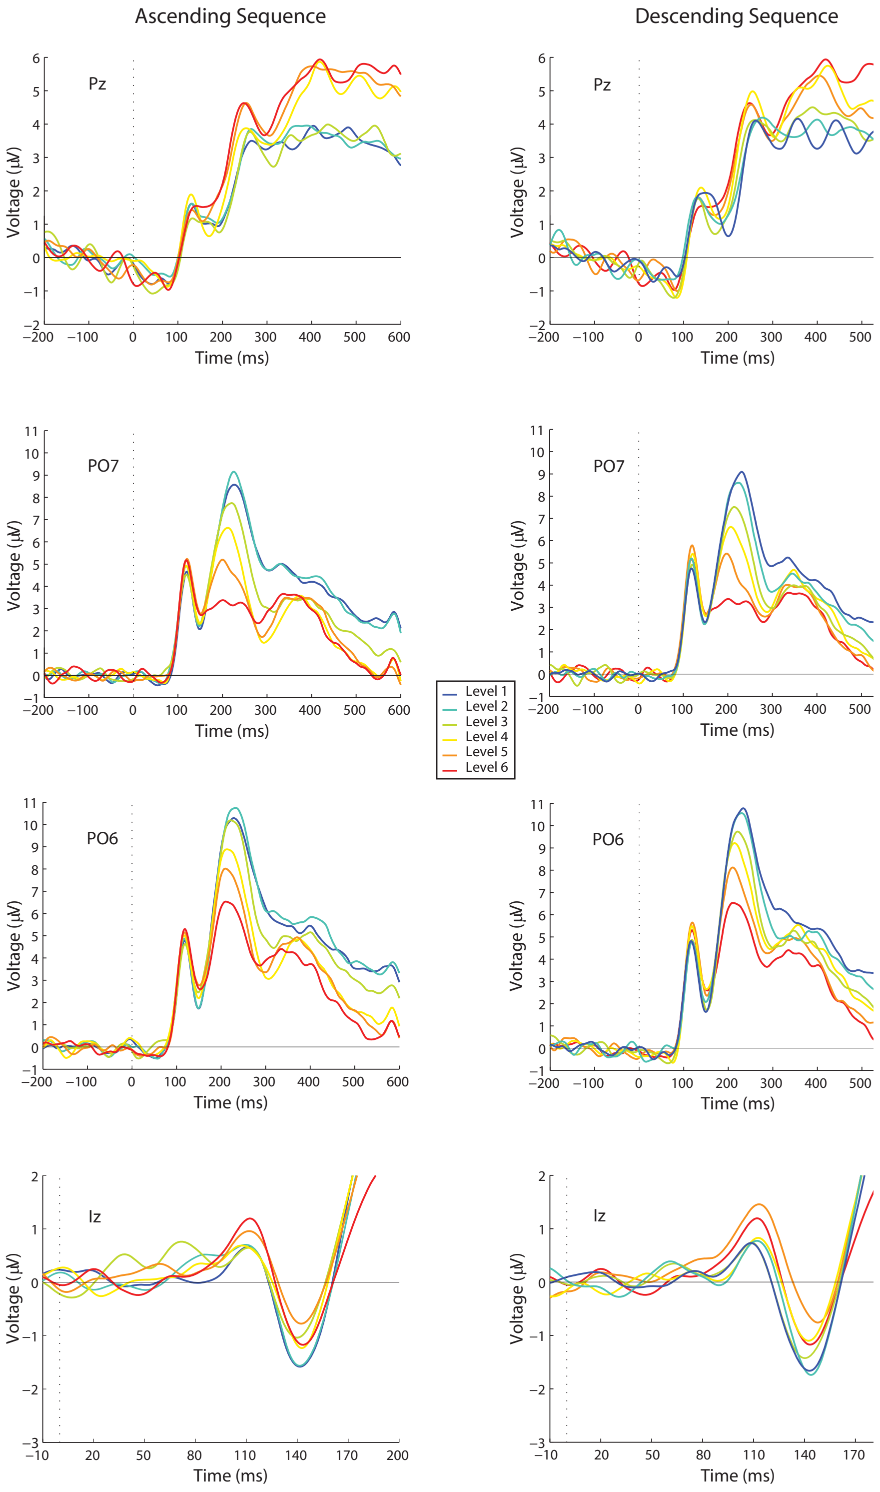

Figure 4-9. Waveforms of P1, N1, P2, and P3 components (Control Experiment). Electrodes Pz, PO7, PO6 and Iz, where P3, P2, N1 and P1 components and their amplitude modulations as a function of degradation level are clearly distinguishable. Left panels show amplitude modulations per degradation level during the ascending sequence, right panels show amplitude modulations during the descending sequence.

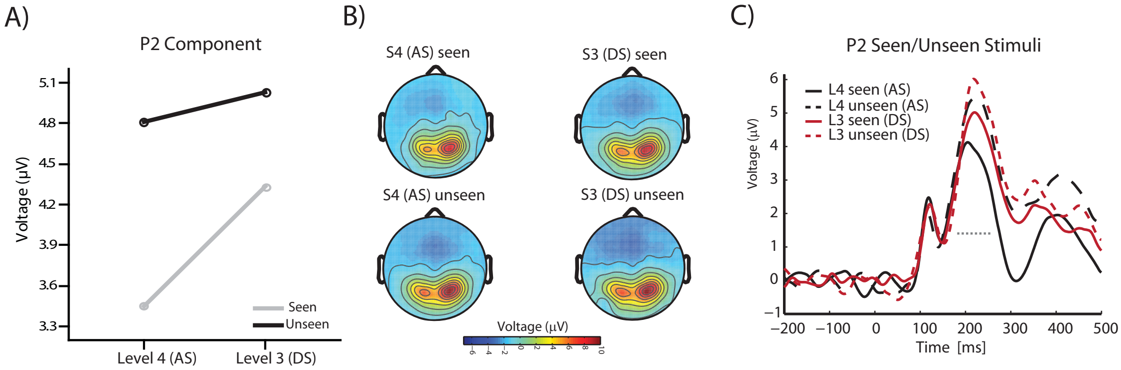

Figure 4-10. Effect of visibility (seen versus unseen trials) at threshold for the P2 component. (a) Lower amplitudes were recorded for seen as compared to unseen trials for both stimuli at threshold. That is, for degradation level 4 during the ascending sequence (AS) and degradation level 3 during the descending sequence (DS). Scalp topographies for seen and unseen trials (b) and their corresponding waveforms (c). Period of significant difference is marked with a dashed line.

dation level 3 for the descending sequence. These stimuli produced similar numbers of seen and unseen trials even though the physical properties of the stimuli differed. We reasoned that ERP components that differentiate between seen and unseen trials at *both* degradation levels are likely to be related to visibility and not to expectations or differences in sensory stimulation. The only ERP component that differentiates between seen and unseen trials at both degradation levels was P2 (see Figure 4-10). Significant amplitude reductions for visible stimuli as compared to invisible stimuli were found during the ascending sequence for degradation level 4 (seen: 3.450 µV; unseen: 4.810 µV; $T(15)=-4.711$, $p=0.00028$) and during the descending sequence for level 3 (seen: 4.331 µV; unseen: 5.030 µV; $T(15)=-3.091$, $p=0.007$). None of the other components (P1, N1 or P3) distinguished visible from invisible trials at both degradation levels (interaction visibility x degradation level: P1 $F(1, 15)=7.748$, $p=0.014$; P3 $F(1, 15)=20.429$, $p=0.0004$). The N1 component showed no modulations in relation to visibility during the ascending and descending sequence (N1 all $p>0.4$). The P1 and P3 component showed higher amplitudes for seen as compared to unseen stimuli only for the ascending sequence (P1 $T(15)=3.721$, $p=0.002$; P3 $T(15)=6.601$, $p>0.001$). Additional time-resolved analyses support the conclusion that the difference between seen and unseen trials is captured by the amplitude of the investigated ERP components (see Figure 4-11).

DISCUSSION

Our results indicate that expectations alter the threshold of visibility. Subjects identified more stimuli in the descending than in the ascending sequences. Since stimuli in both sequences were identical, changes in visibility must be attributed to predictions about the upcoming stimulus. This was reflected in a shift of the psychometric function. Control experiments ruled out that the effects were due to response bias. Predictions have been shown to aid perception (Biederman, 1972; Snodgrass & Feenan, 1990). For example, prior knowledge about the direction and velocity of moving targets enhances their detectability (Schwiedrzik, Alink, Kohler, Singer, & Muckli, 2007; Sekuler & Ball, 1977). Similarly, previous exposure aids the recognition of incomplete or ambiguous figures (Dolan, et al., 1997; Kleinschmidt, et al., 2002). Predictions can also have detrimental effects on perception when they are wrong (Bruner & Potter, 1964) or not updated (as in change blindness, Rensink, 2000).

Signatures of visibility and expectations

By parametrically manipulating stimulus evidence and expectations, we distinguished brain activity related to visibility from activity reflecting sensory properties and/or perceptual expectations of the upcoming stimuli. Our results indicate that even the earliest ERPs change in the presence of expectations. However, not all reflect visibility. Such components should a) differentiate between seen and unseen stimuli irrespective of whether visibility was solely due to stimulus saliency or a combination of saliency and expectancy, and b) show a non-linear dependence of their amplitude on degradation as the psychophysical function. In contrast, linear response modulations reflect physical stimulus properties.

The control EEG experiment where no predictions about stimulus identity could be established revealed that P1 and N1 were the only ERPs that followed a linear profile: P1 amplitude increased with decreasing degradation, while N1 amplitude decreased. In the presence of expectations, P1 amplitude was not modulated by physical stimulus properties. N1, however,

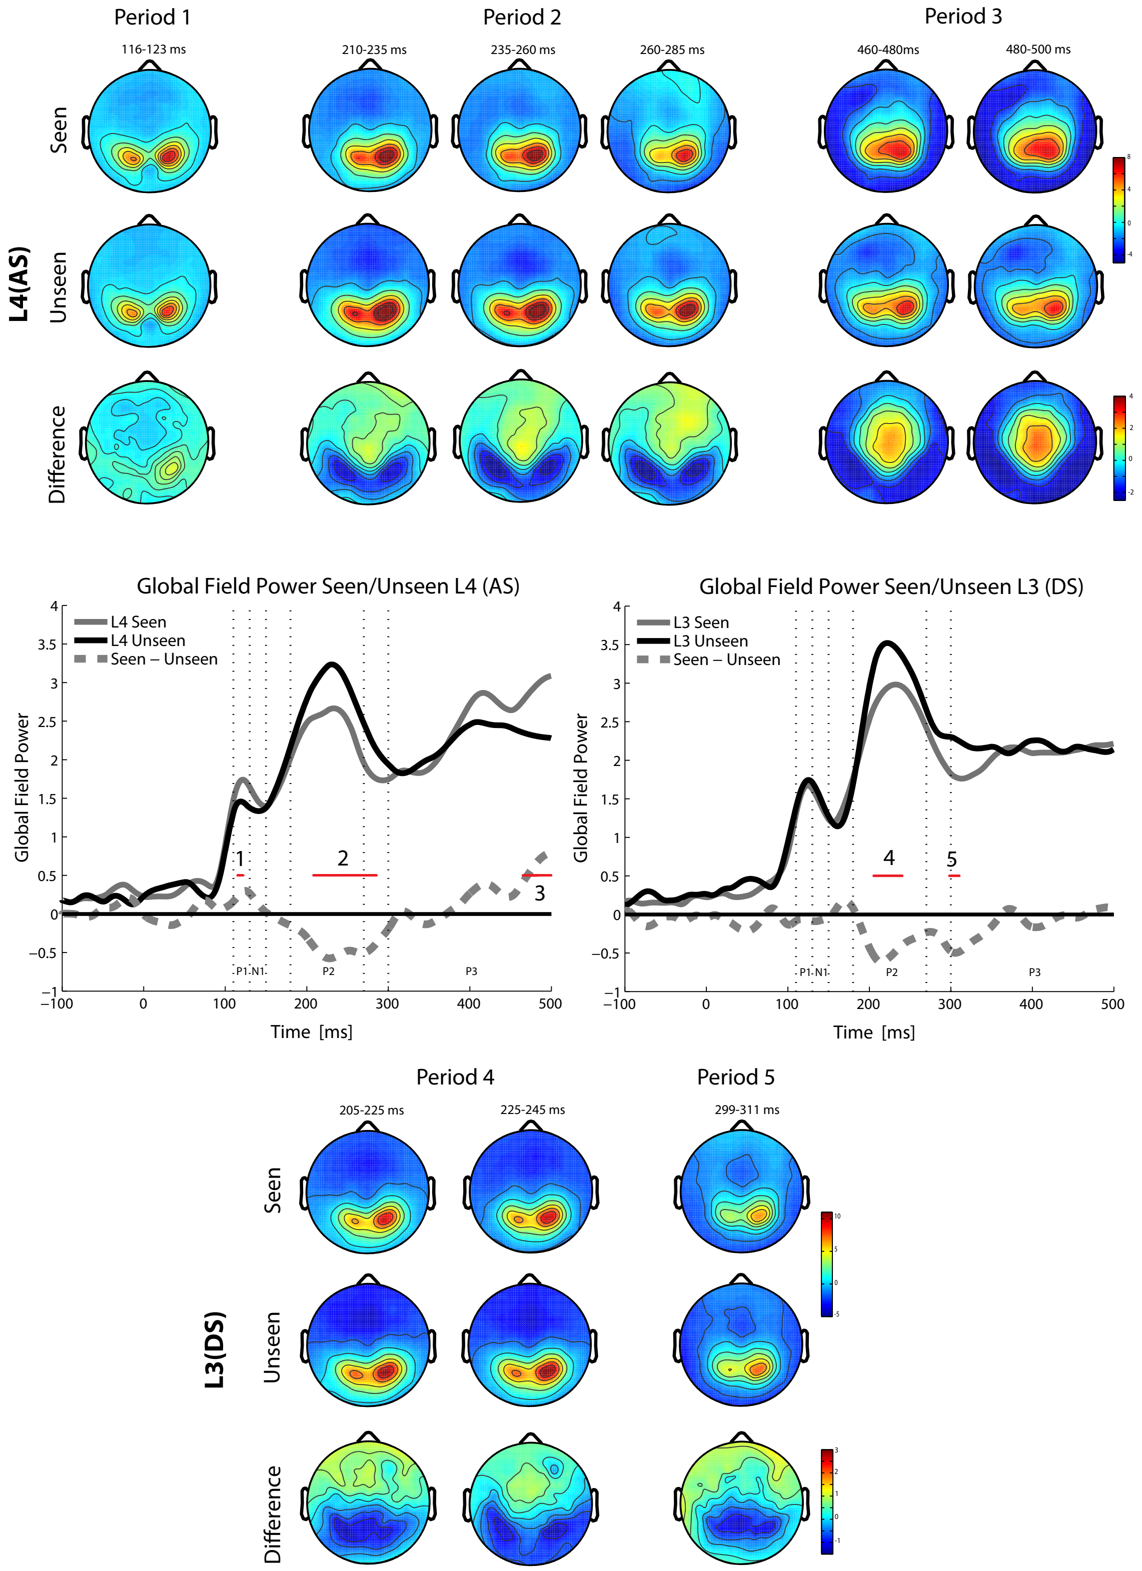
Period 1
116-123 ms
Period 2
210-235 ms
235-260 ms
260-285 ms
Period 3
460-480ms
480-500 ms
L4(AS)
Seen
Unseen
Difference
8
4
0
-4
4
2
0
-2
Global Field Power Seen/Unseen L4 (AS)
L4 Seen
L4 Unseen
Seen – Unseen
Global Field Power
Time [ms]
1
2
3
P1
N1
P2
P3
Global Field Power Seen/Unseen L3 (DS)
L3 Seen
L3 Unseen
Seen – Unseen
4
5
4
3.5
3
2.5
2
1.5
1
0.5
0
–0.5
–1
–100
0
100
200
300
400
500
Period 4
205-225 ms
225-245 ms
Period 5
299-311 ms
L3(DS)
10
5
0
-5
3
2
1
0
-1

showed linear amplitude decrements as a function of degradation, but only during the ascending sequence. During the descending sequence, N1 amplitude was low and not further modulated. This suggests that N1 is reduced once a hypothesis is built and subsequently remains low, irrespective of visibility. This is consistent with predictive coding which anticipates decrements in prediction error as a function of predictability (i.e., top-down predictions that explain away bottom-up inputs via recurrent interactions). Simulations of predictive coding in the auditory domain show that unpredicted stimuli evoke large N1 components, reflecting prediction error when expecting the predicted standard (Garrido, et al., 2009). The neuronal processes associated with the visual N1 represent, then, the first stage at which top-down predictions are compared to bottom-up information (Hopf, Vogel, Woodman, Heinze, & Luck, 2002). Studies investigating the effect of predicting words by previous text (Dambacher, Rolfs, Gollner, Kliegl, & Jacobs, 2009) and target location by previous experience in visual search (Chaumon, Drouet, & Tallon-Baudry, 2008) similarly found effects of predictions around 100 ms post stimulus. Perceptual expectations seem to increase visibility if they are verified by bottom-up information, shortly after input is available.

These effects cannot be explained by attention: Spatial and temporal attention were kept constant in our paradigm. It could be argued that attentional load declines as contrast

◂◂ Figure 4-11. Global field power (GFP) and scalp topographies for seen, unseen and difference between seen and unseen conditions at threshold - level 4 ascending sequence (AS) and level 3 descending sequence (DS). Periods of significant differences in GFP between conditions are marked in red and labelled period 1-5. Horizontal dashed lines refer to the periods in which ERP-components (P1, N1, P2 and P3) were investigated in a previous analysis. For L4 (AS), GFP differed significantly between seen and unseen conditions in three periods: 116-123 ms (period 1), 209-286 ms (period 2), and 466-500 ms (period 3). For L3 (DS), GFP differed significantly between seen and unseen conditions during two periods: 206-241 ms (period 4) and 299-311 ms (period 5). All five periods temporally coincide with the ERP components P1, P2 and P3, which were investigated in the previous analyses. GFP did not reveal an additional component related to awareness: evidence for such a component would be obtained if differences in GFP were present in time windows not corresponding to the originally investigated ERP components or if a difference was present continuously throughout several components without inverting sign. Scalp topographies at the moment of significant differences between seen and unseen trials are presented in the upper and lower panel. They show that seen and unseen trials do not differ in field topography, and accordingly the scalp topography of the difference wave is maximal where the maximal power of the ERP components is found. Altogether, this suggests that the difference between seen and unseen trials is captured by the amplitude of the investigated ERP components.

increases, because the number of distracting dots decreases as the target stimulus evolves. N1 amplitude has been shown to be higher for high as compared to low load (Rorden, Guerrini, Swainson, Lazzeri, & Baylis, 2008). In the main experiment, load decreases during the ascending and increases during the descending sequence. However, N1 amplitude remains low throughout the descending sequence. Thus, attentional load cannot fully explain the observed N1 modulations. Furthermore, a hallmark of attention is the increase in neural activity elicited by attended relative to unattended stimuli (Kastner & Ungerleider, 2000; Luck & Hillyard, 1995; Mangun, 1995). However, the main effect of expectations on N1 was a reduction in amplitude, similar to what is observed for the auditory mismatch negativity (Garrido, et al., 2009). Amplitude attenuation for predicted stimuli have also been reported in fMRI experiments (Alink, et al., 2010; Summerfield, et al., 2008). Thus, attention and expectations boost saliency, but their signatures are reversed (Summerfield & Egner, 2009).

Components following N1 also showed an effect of expectations. However, these effects were related to visibility. At threshold, P2 exhibited lower amplitudes for seen as compared to unseen stimuli independently of physical stimulus properties. Also, P2 amplitude decreased as visibility increased, with a sharp transition at threshold. Furthermore, during the descending sequence, when expectations are firmly established, P2 amplitude was lower than during the ascending sequence, when responses depended mainly on stimulus properties. Thus, the neuronal processes associated with P2 are related to perceptual awareness and as the latter, modulated by prior knowledge.

The inverse relationship between visibility and P2 amplitude is proposed to reflect reentrant mechanisms by which perceptual expectations are compared with bottom-up information (Freunberger, et al., 2007; Kotsoni, Csibra, Mareschal, & Johnson, 2007). Higher areas send predictions to lower areas which detect divergence between prediction and evidence, generating an error signal that increases with mismatch (Friston, 2010; Mumford, 1992). If a prediction fails and a stimulus is not seen, the error signal is highest. Di Lollo et al. (2000)

and Lamme (2006) propose that feedback-feedforward interactions are the core of conscious perception. Preconscious, top-down hypotheses about incoming stimuli are sent to lower areas where they are iteratively compared with the input (Di Lollo, et al., 2000). Consequently, conscious perception is the result of a hypothesis test that iterates until information is consistent across higher and lower areas.

The proposal that P2 indexes reentrant activity is substantiated by source analysis indicating that P2 reflects reactivation of early visual areas including V1 (Di Russo, Aprile, Spitoni, & Spinelli, 2008), latency data from intracranial ERPs (Olson, Chun, & Allison, 2001), and increased flow of mutual information from extrastriate to striate areas at around 220 ms (Hinrichs, Noesselt, & Heinze, 2008). For object substitution masking thought to target reentrant processing, Kotsoni et al. (2007) found that seen trials elicited lower P2 amplitudes than unseen trials, a pattern also reflected by BOLD activity in V1 (Weidner, Shah, & Fink, 2006).

P3 amplitude also correlated with visibility. P3 amplitudes were higher for seen than for unseen trials. However, this was only so during the ascending sequence when perception relied mostly on sensory stimulation. In the descending phase when expectations had been established, P3 amplitude did not differentiate seen from unseen trials. Moreover, the nonlinear increase of P3 amplitude as a function of visibility was only present during the ascending sequence. Thus, P3 fulfills all criteria for a signature of visibility, but only in the absence of expectations. Previous studies have found correlations between P3 amplitude and visibility (Del Cul, et al., 2007; Sergent, et al., 2005). In a masking paradigm, Del Cul et al. (2007) found that P3 was the only component whose amplitude distinguished visible and invisible stimuli and followed visibility sigmoidally. Had we not manipulated expectations, we would have drawn similar conclusions. However, controlling expectancy we found that already P2 was related to visibility. Considering P3 latency, it is likely that it reflects post-perceptual processes *related* to visibility (e.g., update of memory) and not visibility as such, at least in our experiment. However, in our control experiment, when no predictions were generated, only P3 showed a sharp transition at threshold, clearly distinguishing between seen and unseen stimuli. Thus, in the absence of expectations, P3 reflects visibility more closely than P2.

Signatures of visibility: early or late?

Previous studies have suggested either early (~100 ms) (Pins & Ffytche, 2003; Roeber et al., 2008) or late (~300 ms) (Del Cul, et al., 2007; Sergent, et al., 2005) activity correlated with visibility. Our results however suggest that the signatures of visibility are not bound to processes with a strict latency, but depend on the presence of expectations (Banquet & Grossberg, 1987). Latencies are short (indexed by P2, ~200 ms) when expectations are present, and long (indexed by P3, ~300 ms) when expectation are absent. Similar changes in latency have been reported in fMRI experiments as a function of previous exposure (James, Humphrey, Gati, Menon, & Goodale, 2000; Kleinschmidt, et al., 2002).

An open question is why expectations shorten the latency of signatures of visibility. One possibility is that previous experience speeds up the accumulation of evidence necessary for visibility (James & Gauthier, 2006). Predictions could also take the form of a decisional bias, i.e., a change in the weight of a particular stimulus (higher prior evidence) over others. Thus, the starting point for the accumulation of evidence lies closer to the decision bound, accelerating the collection of information. Biasing signals correlating with detectability have been observed in lateral occipital complex before stimulus onset (Stokes, Thompson, Nobre, & Duncan, 2009).

In light of our results it appears necessary to reinvestigate the neuronal correlates of consciousness, taking into account how cognitive functions (attention, expectations, memory, etc.) influence the timing (and potentially other features) of processes required for access to consciousness. A thorough understanding of the neuronal correlates of consciousness might require a departure from the strategy of merely comparing seen with unseen conditions, and instead necessitates a proper characterization of the interactions among all cognitive processes that ultimately lead to conscious experience (Melloni & Singer, 2010).

5

UNTANGLING PERCEPTUAL MEMORY: HYSTERESIS AND ADAPTATION MAP INTO SEPARATE CORTICAL NETWORKS

The fact that previous experience plays a pivotal role in perception is well documented. Critically, previous experience can affect perception in two opposing ways: It can either render us more likely to perceive the same stimulus again (hysteresis), or more likely to see something else (as in various aftereffects, thought to result from adaptation). However, it is currently unknown what determines the direction of such experience-dependent perceptual effects and whether they result from the same neuronal mechanism. To address these questions we combine functional magnetic resonance imaging with a paradigm with which we can dissociate hysteresis from adaptation in behavior. We provide clear evidence that adaptation and hysteresis map into two distinct cortical networks: Hysteresis engages a widespread network including higher order parietal and frontal regions. In contrast, adaptation is confined to early sensory, extrastriate areas only. Thus, although affecting our perception concurrently, hysteresis and adaptation are functionally dissociated in the human brain: Higher order areas stabilize perception, while early sensory areas display adaptation. This separation may explain how the brain maintains the balance between exploiting redundancies and staying sensitive to new information.

Schwiedrzik CM, Ruff CC, Leitner F, Singer W, Melloni L (in preparation).

Our perception depends on the currently available sensory information, but also strongly on previous experience (Fecteau & Munoz, 2003; O. Schwartz, Hsu, & Dayan, 2007). This is evident, for instance, in the waterfall illusion, where prolonged exposure to one direction of motion causes subsequently viewed stimuli to appear moving in the opposite direction (Purkinje, 1820). Such repulsive aftereffects are commonly attributed to neuronal adaptation (e.g., Anstis, et al., 1998). Interestingly, however, previous experience can also have the opposite effect: After brief exposure to a moving stimulus, the subsequent ambiguous probe appears to move in the same direction (Kanai & Verstraten, 2005; also see Long, Toppino, & Mondin, 1992), a phenomenon known as "hysteresis" or "priming". Whereas adaptation has been proposed to prepare the brain for the uptake of new information (Barlow, 1990), hysteresis might help to stabilize percepts in face of constantly changing low-level stimulus features (Kleinschmidt, et al., 2002). What then determines the direction of the effect of previous experience?

Here, we address this question in the context of multistable stimuli. For long, multistable perception has been considered to be memoryless (R. Blake, Fox, & McIntyre, 1971; Borsellino, De Marco, Allazetta, Rinesi, & Bartolini, 1972; R. Fox & Herrmann, 1967). However, it is now clear that when multistable stimuli are presented intermittently, perception on any given trial will depend on hysteresis and adaptation (for a review, see Pearson & Brascamp, 2008). Despite much effort devoted to the study of such effects of previous experience in multistable stimuli, it is still poorly understood whether and how hysteresis and adaptation jointly and concurrently determine perception and what the underlying neuronal mechanisms are. One type of models postulate that both effects can be explained by a single mechanism (Gepshtein & Kubovy, 2005), most prominently the amount of neuronal adaptation (R. Blake, Sobel, & Gilroy, 2003; Chen & He, 2004; Orbach, Ehrlich, & Heath, 1963), whereas other models propose two distinct mechanisms that separately account for either hysteresis or adaptation (Brascamp, et al., 2009; Noest, et al., 2007; Wilson, 2007). Furthermore, the neuronal stage at which hysteresis and adaptation occur is under debate: While some authors have ascribed hysteresis and adaptation to early sensory processing (Brascamp, et al., 2009; Noest, et al., 2007; Wilson, 2007), other accounts suggest that both effects arise from interactions between early and higher-order processing stages (Gigante, Mattia, Braun, & Del Giudice, 2009; Kanai & Verstraten, 2005, 2006; Maier, Wilke, Logothetis, & Leopold, 2003; Pastukhov & Braun, 2008; Sterzer & Rees, 2008). Progress on these questions has been hampered by the fact that none of the previous brain imaging studies (Kleinschmidt, et al., 2002; Raemaekers, van der Schaaf, van Ee, & van Wezel, 2009; Sterzer & Rees, 2008) have experimentally separated hysteresis from adaptation, hence studying mixtures of both effects that may be difficult to compare and to interpret. Here we resolve these limitations by measuring brain activity with functional magnetic resonance imaging (fMRI) while participants complete a paradigm that enables us to investigate hysteresis and adaptation concurrently but to dissociate them experimentally. With this paradigm, we are thus able to effectively isolate the brain areas showing effects of hysteresis from those showing adaptation. We find that adaptation and hysteresis cannot be accounted for by the same mechanism, because they involve different cortical networks: Hysteresis is reflected by activity not confined to visual areas in which the relevant feature (orientation) is explicitly coded, but involves higher-order areas in parietal and frontal cortices. In contrast, adaptation is confined to activity in early visual, extrastriate areas only. We propose that these results are fundamentally at odds with models that define both hysteresis and adaptation as expressions of the same underlying neuronal process. Instead, our data seem more consistent with a Bayesian model of perceptual memory in which hysteresis is the result of a change of the prior coded in higher-level cortices, whereas adaptation changes the available stimulus evidence computed in sensory cortex. Keeping these two effects of previous experience separated in this way may allow the brain to maintain the balance between exploiting redundancies and staying sensitive to new information.

METHODS

Participants

Thirty healthy human subjects (mean age 25.1, range 21-38, 11 male) participated in the psychophysical part of this study. Twenty nine were right handed, as assessed with the Edinburgh Inventory (Oldfield, 1971). All subjects had normal or corrected-to-normal vision, reported no history of neurological or psychiatric disease, and gave written informed consent before participation. Twenty of these subjects (mean age 24.6, range 21-29, 7 male, all right handed) also participated in the fMRI study. Two further subjects participated in the fMRI study but their data were excluded due to technical artifact. Subjects received monetary compensation for their participation.

Procedure and stimuli

For the psychophysical experiments, stimuli were displayed on a CRT monitor (HP p1230, resolution 1024×768, visible screen size 30°×22.9° at 75 cm distance, 150 Hz refresh rate). Inside the scanner, a video goggle system was employed (Resonance Technology MR Vision 2000, resolution 800×600, visible screen size 30°×22.5° at 1.2 m virtual distance, 60 Hz refresh rate). Stimuli were generated using Matlab (R2007a, The MathWorks). Stimulus presentation and response collection were controlled by Presentation software (v13.1, Neurobehavioral Systems). All stimuli were presented on a gray background (7.62 cd/m^2); a red fixation circle was continuously present in the center of the screen (6.72 cd/m^2).

We adapted a paradigm originally introduced by Gepshtein and Kubovy (2005), in which we used sequences of stimuli with different degrees of multistability: a rectangular and a hexagonal dot lattice. Dot lattices (11.5° diameter) were made up from ordered spatial arrangements of white Gaussian blobs ('dots', 0.25° diameter), presented at the center of the screen. Dots were aligned along several orientations, depending on the dot lattice geometry. Dots are perceptually grouped along the shortest inter-dot distance according to the Gestalt law of proximity (Kubovy, Holcombe, & Wagemans, 1998), hence giving the impression of oriented lines. Rectangular dot lattices are bistable, i.e., observers are equally likely to observe 0° or 90° orientations if the dots are evenly spaced. Hexagonal dot lattices are tristable; when dots are uniformly spaced, observers are equally likely to see 0°, 60°, or 120° orientations (Kubovy & Wagemans, 1995).

The likelihood to perceive the 0° [p(0°)] or 90° [p(90°)] orientation in rectangular dot lattices can be manipulated by changing the aspect ratio (AR), i.e., the inter-dot-distance along the 0° and 90° orientation (Figure 5-1a). The shorter the distance between the dots in one of the orientations, the more likely is the perception of the organization along that orientation. Hence, to examine carry-over effects between rectangular and hexagonal dot lattices, we systematically manipulated the AR of the rectangular stimulus while keeping the tristable stimulus identical and maximally instable (Kubovy & Wagemans, 1995). We used seven ARs: 1.3^{-1}, 1.2^{-1}, 1.1^{-1}, 1, 1.1, 1.2, and 1.3. At AR=1, p(0°)=p(90°) with an inter-dot-distance of 1°. Thus, the stimulus is fully bistable. At AR<1, p(0°)>p(90°). At AR>1, p(0°)<p(90°). The product of the inter-dot-distances for 0° and 90° was kept invariant (~1), so as to ensure roughly equal dot density for all ARs. The total area of any dot lattice was smoothed with a Gaussian (SD 0.15) to give the impression that the lattice lines continue behind a virtual aperture (Nikolaev, Gepshtein, Gong, & van Leeuwen, 2010). This avoids interactions of the oriented lines with the edges of the stimuli. The exact position of the dot lattice was pseudo-randomly jittered within the aperture (0°-1.15°) to prevent dots of subsequent displays to occupy systematically related portions of space.

On each trial, we presented two dot lattices, the first of which was rectangular (800 ms) and the second hexagonal (300 ms) (Figure 5-1b). Both lattices were oriented along the same orientation (henceforth referred to as 0°), but this orientation was randomly varied from trial to trial, covering 90° in 1° steps. After each presentation of a dot lattice, subjects chose the orientation they had perceived from a screen displaying four alternative orientations (0°, 90° and the two diagonal orientations for the rectangular dot lattices; 0°, 60°, 120° and the unlikely

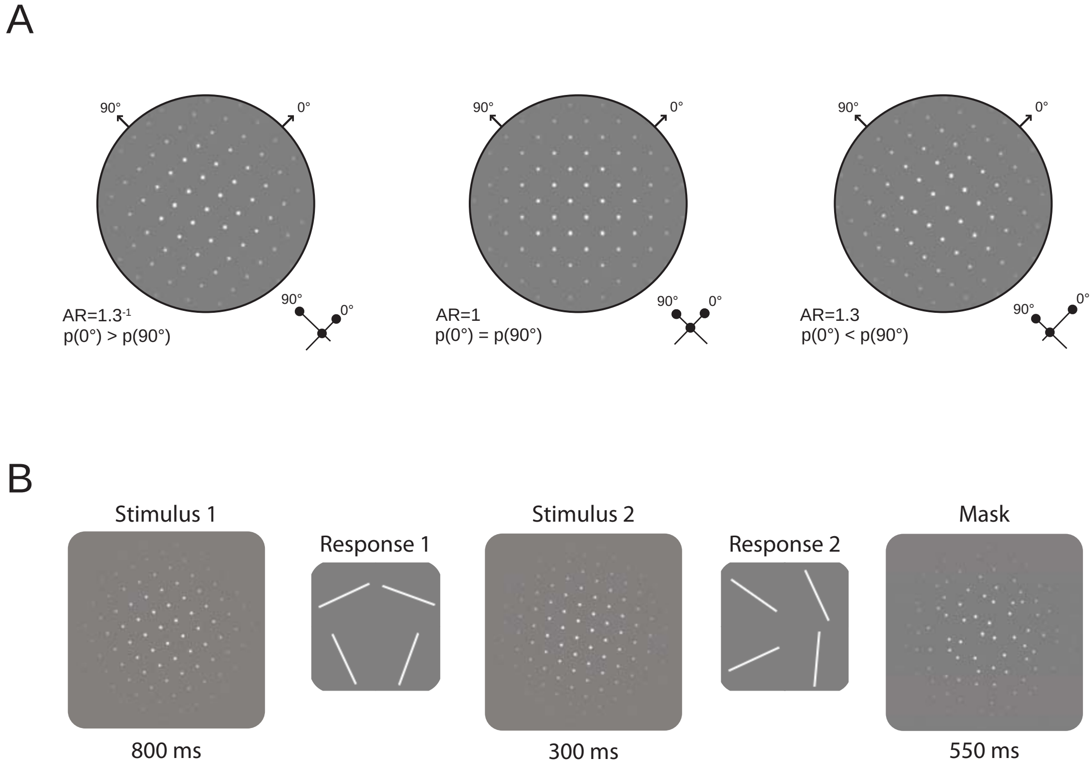

Figure 5-1. Stimuli and experimental procedure. (a) Examples of rectangular dot lattices with different aspect ratios (AR), i.e., the inter-dot distance along the 0° and 90° orientations. Small arrows indicate the two dominant percepts, 0° and 90°. The likelihood of perceiving a particular orientation depends on the shortest inter-dot distance. At AR=1.3-1 (left panel), the inter-dot distance along the 0° orientation is shorter (see inset). Thus perception favors 0° over 90° and subjects tend to perceive the dot array as tilted to the right (p(0°)>p(90°)). In contrast, at AR=1.3 (right panel) perception favors 90° over 0°, and the array is likely to be perceived as tilted to the left (p(0°)<p(90°)). At AR=1 (middle panel), 0° and 90° orientations are equally likely, thus, the stimulus is bistable (p(0°)=p(90°)). (b) On a given trial, we first presented a rectangular dot lattice with varying AR followed by a response screen on which subjects chose the perceived orientation. Then, we presented a hexagonal dot lattice (with fixed AR), also followed by a response screen. Last, we presented a dynamic random dot mask to avoid carry-over effects to the next trial. All stimuli were presented centrally and subjects were instructed to fixate a red fixation dot at the center of the screen. Rectangular dot lattices are bistable, i.e., observers are equally likely to observe 0° or 90° orientations if the dots are evenly spaced. Hexagonal dot lattices are tristable; when dots are uniformly spaced, observers are equally likely to see 0°, 60°, or 120° orientations. Hysteresis can be investigate by assessing how the likelihood to perceive 0° in the second stimulus depends on whether 0° were already perceived in the first stimulus (hysteresis). Adaptation can be investigated by assessing how the likelihood to perceive 0° in the second stimulus depends on AR, i.e., the amount of stimulus evidence for 0° in the first stimulus.

90° orientation for the hexagonal dot lattices). The position at which the alternatives appeared (and thus the response buttons) was fully randomized within and between trials. Subjects were instructed to be accurate, to fixate on a central fixation dot (as ascertained by eye tracking, see below), and to report the first perceived orientation in case the percept switched during the presentation of an individual dot lattice. Note that subjects were never required to judge whether they had seen the same orientation in the two stimuli. After the response to the hexagonal dot lattice, a dynamic random dot mask (updated at 25 Hz) was presented for 550 ms to avoid afterimages and between-trial carry-over effects.

Subjects completed nine blocks of 70 trials. Conditions were pseudo-randomized such that each AR occurred equally often during each block. We monitored eye position using a binocular infrared eye tracker in all psychophysical experiments (SR Research Eyelink 1000, 1000 Hz sampling rate). The eye tracker was recalibrated using a standard nine point calibration procedure at the beginning of each block. The experiments were conducted in a darkened, sound-attenuating chamber. Constant head position was assured by the use of a chinrest with forehead support. Before the experiments started, subjects received 30 practice trials.

Procedure and stimuli - functional magnetic resonance imaging

Stimuli and procedure during fMRI were essentially identical to the psychophysical experiment, with the following exceptions: We reduced the number of ARs to three (AR=1, 1.1, 1.2). Thus, rectangular dot lattices would either be bistable or biased towards 90°. We also introduced trials in which the second stimulus was omitted. In such a "partial trial design", these trials enable us to separate the hemodynamic response to the second stimulus from the convolved responses to both stimuli (Ollinger, Corbetta, & Shulman, 2001; Ollinger, Shulman, & Corbetta, 2001; see below).

Subjects completed ten blocks of 72 trials in two sessions. Each block consisted of 30 trials with two stimuli, 30 trials with one stimulus, and 10 "null" baseline trials (during which only the fixation cross was visible on the display). The trial types were presented in pseudo-random order, avoiding long sequences of identical trial types (this was achieved by shuffling conditions in sequences of 14 stimuli, 5 times per block). Furthermore, the first trial was randomly chosen from all available conditions and later discarded from all analyses to eliminate T1 saturation artifacts in the fMRI data. The last trial in each block was always a baseline trial in order to not curtail the hemodynamic response for the last event in a block. Subjects were instructed to respond accurately within 1.7 s. This was practiced inside the scanner before the experiment started.

Parameters - functional magnetic resonance imaging

MRI data was acquired on a 3T scanner (Siemens Allegra), using a 4 channel head coil. For the main experiment, we acquired 216 volumes of 29 slices per run of functional ($T2^*$-weighted) echoplanar imaging (EPI) data (TR=2.5 s, TE=30 ms, voxel size 3×3×3 mm, gap thickness 0.3 mm). Anatomical images were acquired using a T1-weighted magnetization prepared rapid gradient echo (MPRAGE) sequence (160 slices, TR=2.25 s, TE=4.38 ms, voxel size 1×1×1 mm). All sequences covered the whole brain.

Data analyses - psychophysics

Behavioral data was analyzed with a logistic regression using Generalized Estimating Equations (GEE, Liang & Zeger, 1986) in SPSS (v17 and v18, SPSS Inc.). GEE is an extension of the generalized linear model developed for the analysis of repeated-measures designs (for an introduction to GEE see Hanley et al. (2003)). Following Gepshtein and Kubovy (2005), we restricted our analyses to the response alternatives with equal likelihood at AR=1, i.e., cases were subjects responded either 0° or 90° to the rectangular dot lattices, and 0°, 60° or 120° to the hexagonal dot lattices. However, note that analyses using all response alternatives yielded comparable results. Data was sorted by subject, block, and trial number. To account for the correlations between successive trials we used a working correlation matrix with a 1^{st} order autoregressive relationship.

Data analyses - eye tracking

To ensure that none of the behavioral effects were ascribable to eye movements, we acquired eye tracking data during the psychophysical experiments (n=26). Temporal windows for analyses of fixation stability were defined from 200 ms before stimulus onset until stimulus offset (800 ms for stimulus 1, 300 ms for stimulus 2). A trial was excluded from the analyses if a blink occurred within 100 ms before stimulus onset or after stimulus offset. We then calculated the percentage of eye position samples falling into a square window around the fixation dot for two levels of precision (window size 1.5°×1.5° or 2°×2°). Repeated measures analyses of variance (rmANOVA) were run separately at each level of precision for stimulus 1 and stimulus 2, with factors eye (left, right) and hysteresis (0°-0°, 90°-0°) to control for the effect of fixation stability on the hysteresis effect, and with factors eye (left, right) and aspect ratio (1.3^{-1}, 1.2^{-1}, 1.1^{-1}, 1, 1.1, 1.2, 1.3) to control for the effect of fixation stability on the adaptation effect.

Data analyses - functional magnetic resonance imaging

MRI data was analyzed in Brain Voyager QX (v2.1, Brain Innovation), SPSS, and in Matlab using the Brain Voyager Toolbox and custom code. Prior to preprocessing, the first three volumes of each functional run were excluded to prevent T1 saturation effects. The remaining data was slice scan time corrected with cubic spline interpolation, 3D motion corrected using trilinear/sinc interpolation, and temporally high-pass filtered at 0.01 Hz. Functional and anatomical data were brought into ACPC space using cubic spline interpolation and then transformed into standard Talairach space (Talairach & Tournoux, 1988) using trilinear interpolation. For the whole brain analyses, data was spatially smoothed with a Gaussian kernel (8 mm FWHM). To create inflated surface reconstructions, the grey-white matter boundary in the structural scans was segmented, reconstructed, smoothed and inflated (Kriegeskorte & Goebel, 2001) separately for each hemisphere. For the main experiment, blood oxygenation level dependent (BOLD) responses were estimated using a deconvolution model in a random effects (RFX) general linear model (GLM). For each condition, we defined appropriately placed series of eight finite impulse response (FIR) predictors (one per volume) to model the 20 seconds BOLD response following the onset of each trial. Contrasts were run over the 3rd and 4th predictor (corresponding to 5 and 7.5 s post stimulus), thus covering the peak of the BOLD response while accounting for variability in the peak time in different subjects and brain regions (Handwerker, Ollinger, & D'Esposito, 2004). To reduce the number of voxel-by-voxel comparisons, we restricted the analyses by use of a cortex mask based on the individual grey-white matter boundary. This reduced the number of voxels to 47405. For follow-up region of interest (ROI) analyses, we extracted the mean deconvolved time courses of the clusters identified in the whole brain analysis per subject from the unsmoothed data.

To identify cortical regions involved in perceptual hysteresis, we first defined a GLM with seven independent predictors (each modeled with the FIR set comprising 8 post stimulus periods): (1) Hysteresis (two stimuli trials where subjects responded 0° to both stimuli); (2) no hysteresis (two stimuli trials where subjects responded 90° to the first stimulus and 0° to the second); (3) 0° (one stimulus trials where subjects responded 0°); (4) 90° (one stimulus trials where subjects responded 90°); (5) a predictor for two stimuli trial where subjects chose one of the remaining orientations on the first, second stimulus or for both stimuli; (6) a predictor for one stimulus trials where subjects chose a diagonal; and (7) a predictor for trials where subjects failed to respond. We then created beta maps per subject for two contrasts: [hysteresis vs. no hysteresis] and [0° vs. 90°]. In a second level analysis, we contrasted [hysteresis vs. no hysteresis] > [0° vs. 90°] to reveal brain areas with significant changes related to hysteresis. This map was thresholded at a voxel level of $p<0.01$, $T(19)=2.860$, and cluster size thresholded at $p<0.05$ (5000 iterations) to correct for multiple comparisons. The resulting cluster size threshold was 263 mm^3/297 voxels. Since perceptual memory can only build up with the presentation of the first stimulus, solely areas showing a significant BOLD response to the first stimulus (post-hoc *t*-test on the mean of 0° and 90° one stimulus trials, $p<0.05$, uncorrected) were considered to be directly involved in perceptual hysteresis.

To identify regions expressing adaptation, we first defined a GLM with seven independent predictors (each modeled with the FIR set comprising 8 post stimulus periods): one predictor per AR for one and two stimuli trials, respectively, and one for trials where subjects failed to respond. We then created beta maps per subject for the contrasts [AR=1 vs. baseline], [AR=1.1 vs. baseline] and [AR=1.2 vs. baseline] for trials with two stimuli, and ran an F-test to compare between the three conditions. The resulting map was thresholded at a voxel level of $p<0.05$, $F(2,38)=3.244$, and cluster size thresholded at $p<0.05$ (5000 iterations) to correct for multiple comparisons. The resulting cluster size threshold was 749 mm^3/783 voxels. For each of the clusters, we extracted the deconvolved time courses per subject. Areas showing a significant positive BOLD response to the first stimulus (post-hoc *t*-test on the mean of 0° and 90° one stimulus trials, $p<0.05$, uncorrected) and a linear trend (as assessed by linear trend analysis, Howell, 2002, p. 408f) of the peak BOLD response over the three ARs for the second stimulus were considered to be involved in perceptual adaptation. Note that a trend analysis requires that the factor for which the trend is investigated is significant in the preceding omnibus *F*-test, but that the presence or absence of a linear trend is statistically independent of this preceding test. Thus, ROI analyses were statistically independent of the contrasts performed for the whole brain.

Eye tracking was not available inside the scanner, however, to assess whether eye movements could account for the observed effects, we used a method introduced by Beauchamp (2003). This method exploits the fact that major signal changes during eye movements can be observed in the eyeballs themselves. Thus, differential eye movements in different conditions can be assessed by comparing signal changes in the eyeballs. We first identified for each functional run voxels with time course jumps >=4 SD. The resulting maps were averaged per subject. Based on the average maps, we then defined ROIs for each eyeball in each subject (mean number of voxels left eye: 108.2, *SE* 3.93; mean number of voxels right eye: 111.86, *SE* 2.79), and joined these ROIs into one ROI per subject. If ROIs could not be identified this way, we defined a sphere (radius 3 voxels, size 123 voxels) around the group-average peak voxel of the respective eyeball. This was done for six eyeballs. ROIs could be defined for 19 subjects. These ROIs were used as control regions for all other contrasts.

RESULTS

Behavioral results

For the behavioral data, we first analyzed the responses to the rectangular dot lattices (R1) using a logistic regression with aspect ratio (AR) as the predictor. AR reflects the inter-dot distance, such that grouping by proximity is more likely to occur along the orientation with the shortest inter-dot distance. Thus, grouping along the 0° orientation is expected for AR<1 and along the 90° orientation for AR>1 (Gepshtein & Kubovy, 2005; Kubovy, et al., 1998; Nikolaev, et al., 2010). Confirming this prediction, our data (shown in Figure 5-2a) reveal that the probability of responding 0° decreased as AR increased (Wald's $\chi^2(1)=111.880$, $p<0.01$, Table 5-1). We also confirmed bistability of the rectangular dot lattice at AR=1, as the probability to respond 0° or 90° was close to chance[2].

For the responses to the second stimulus (R2) – the hexagonal dot lattice – we modeled R2 as a function of the AR of the first stimulus and the response given to the first stimulus (R1). The interaction term was not included, as a previous model had revealed no significant interaction between AR and R1 (Wald's $\chi^2(1)=0.251$, $p=0.617$). Both AR and R1 predicted R2 (AR: Wald's $\chi^2(1)=58.797$, R2: Wald's $\chi^2(1)=146.186$, both $p<0.01$, Table 5-1)[3]. Figure 5-2b shows the carry-over effect from R1 to R2: The black line in the figure shows trials in which subjects perceived 0° in the first dot lattice and continued perceiving 0° in the second dot lattice, thus exhibiting perceptual hysteresis. The gray line shows trials in which subjects first perceived 90° and then switched to perceiving 0° in the second stimulus. Importantly, the two lines run in parallel, showing that perceptual hysteresis does not interact with or depend on the AR of the first stimulus. The vertical separation between the two lines reflects the size of the perceptual hysteresis effect, indicating that

---

[2] Regression models per subject showed that AR was a significant predictor at $p<0.01$ and that its slope was negative in all 30 subjects. Exact binomial tests per subject restricted to AR=1 showed that at $p<0.05$, 12 subjects favored 0° over 90°, 14 subjects chose 0° and 90° with equal likelihood, and 4 subjects favored 90° over 0°. This imbalance is reflected in the slight offset towards responding 0° at AR=1 in the group data.

[3] Regression models per subject showed that both predictors were significant at $p<0.01$ in all subjects, and both predictors were positive in 28 out of 30 subjects (2 subjects had a negative slope for AR).

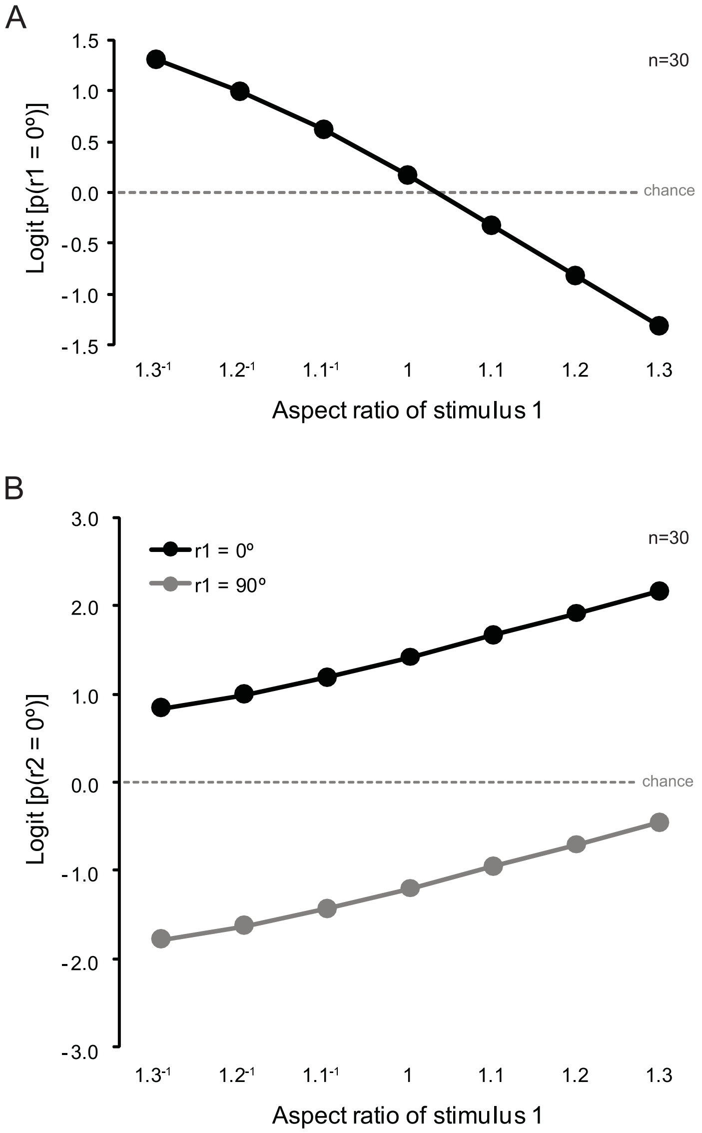

Figure 5-2. Behavioral results. (a) Responses to the first stimulus (logit). The likelihood to perceive 0° decreased as a function of aspect ratio. At AR=1.3^{-1}, subjects were more likely to perceive 0°, while at AR=1.3, subjects were more likely to perceive 90°. (b) Responses to the second stimulus (logit) as a function of the aspect ratio of the first stimulus. The likelihood to perceive 0° increases with AR, in contrast to the responses to the first stimulus. This inversion is the signature of adaptation. The black line shows hysteresis trials (subjects perceived 0° in both stimuli), the grey line shows no hysteresis trials (subjects perceived 90° in the first stimulus and 0° in the second stimulus). The vertical separation between the two lines indicates the size of the hysteresis effect. Note that both lines run parallel indicating that hysteresis and adaptation are independent.

subjects were likely to perceive the same orientation in the first and in the second lattice. The effect of adaptation becomes apparent when comparing the slope in Figure 5-2a with Figure 5-2b: While the likelihood to respond 0° decreases with AR for the first stimulus, it increases with AR for the

Response	Predictor	β (log odds)	SE β	Wald's χ^2	df	p	e^{β} (odds ratio)
R1	constant	5.091	0.4966	105.087	1	<0.01	162.524
	aspect ratio	-4.924	0.4655	111.880	1	<0.01	0.007
R2	constant	-4.399	0.4228	108.232	1	<0.01	0.012
	aspect ratio	2.509	0.3271	58.797	1	<0.01	12.287
	R1	2.623	0.2169	146.186	1	<0.01	13.776
R1 (fMRI)	constant	7.293	0.9470	59.308	1	<0.01	1469.697
	aspect ratio	-6.789	0.8444	64.583	1	<0.01	0.001
R2 (fMRI)	constant	-6.923	.9151	57.230	1	<0.01	0.001
	aspect ratio	4.843	0.6738	51.665	1	<0.01	126.837
	R1	3.048	0.4237	51.774	1	<0.01	21.080
R1 (one stimulus)	constant	6.550	0.7739	71.630	1	<0.01	698.995
	aspect ratio	-6.146	0.6946	78.286	1	<0.01	0.002

Table 5-1. Logistic regression analysis of the behavioral data of the psychophysical and fMRI experiments.

second stimulus. Thus, the more likely subjects were to perceive 0° or 90° in the first dot lattice, the less likely they were to perceive it in the second lattice. Such negative correlations are characteristic of adaptation effects. Interestingly, adaptation was even present for orientations that were not perceived. For instance, in the black line in Figure 5-2b, adaptation to 90° is evident by the positive slope (i.e., the inverse slope of Figure 5-2a). However, on these trials, subjects twice reported perceiving 0°. This indicates that both 90° and 0° orientations are always processed, regardless of whether they are perceived or not (R. Blake & Fox, 1974; Gepshtein & Kubovy, 2005; Hock, Schöner, & Hochstein, 1996). None of these results could be explained by differences in fixation stability (no significant main effects or interactions at $p<0.05$, Greenhouse-Geisser corrected). Importantly, all results could be reproduced in the scanner (Table 5-1). Here, we also found that AR significantly predicted whether subjects would respond 0° or 90° on trials with only one stimulus (Wald's $\chi^2(1)=78.286$, $p<0.01$, Table 5-1).

Imaging results

Having established that hysteresis and adaptation are independently expressed in the behavioral response patterns, we now turned to the brain imaging data to investigate the cortical areas underlying each effect. To this end, we employed a logic that closely followed the analyses of the behavioral data. To identify regions involved in hysteresis, we contrasted trials in which subjects expressed hysteresis (subjects reported twice the same percept, i.e., 0°) with trials in which they did not express hysteresis (subjects reported a perceptual switch, i.e., 90° followed by 0°). Given the hemodynamic lag, BOLD responses to the first stimulus are convolved with the responses to the second stimulus on trials with two stimuli. Thus, in order to discard the difference between perceiving the 0° or 90° orientation in the first stimulus, we also contrasted these two percepts on trials with only one stimulus and subtracted this difference from the difference between hysteresis and no hysteresis trials. The result of the contrast [hysteresis vs. no hysteresis] > [0° vs. 90°] was thus matched for the percept and the response to the second stimulus and was not confounded by potential differences in perceiving or reporting 0° or 90° in the first stimulus. Figure 5-3 shows that hysteresis is expressed in a distributed network of brain areas spanning ventral visual areas (including bilateral fusiform gyrus), superior parietal (bilateral intraparietal sulcus [IPS]), and frontal cortices (right anterior insula, right pre-supplementary motor area [pre-SMA], and right dorsomedial prefrontal cortex [dmPFC]). The time course of fMRI signal change in these regions is displayed in Figure 5-3, and illustrates graphically that in all regions the absolute response on hysteresis trials was smaller than on no hysteresis trials, resembling repetition suppression (Grill-Spector, et al., 2006). It is interesting to note that expected stimuli are often found to elicit smaller BOLD responses than unexpected stimuli (e.g., Alink, et al., 2010), and have thus been proposed to reflect an effect of top-down perceptual expectations (Friston, 2005; Summerfield & Egner, 2009).

While the previous results suggest a network of brain areas involved in perceptual hysteresis, in the next analysis we attempted to tie the contribution of those regions more directly to percept maintenance. To that end, we took advantage of the interindividual variability in the size of the behavioral hysteresis effect (log odds) and correlated this with the difference in activation between hysteresis and no hysteresis trials for each brain region. Regions showing a cor-

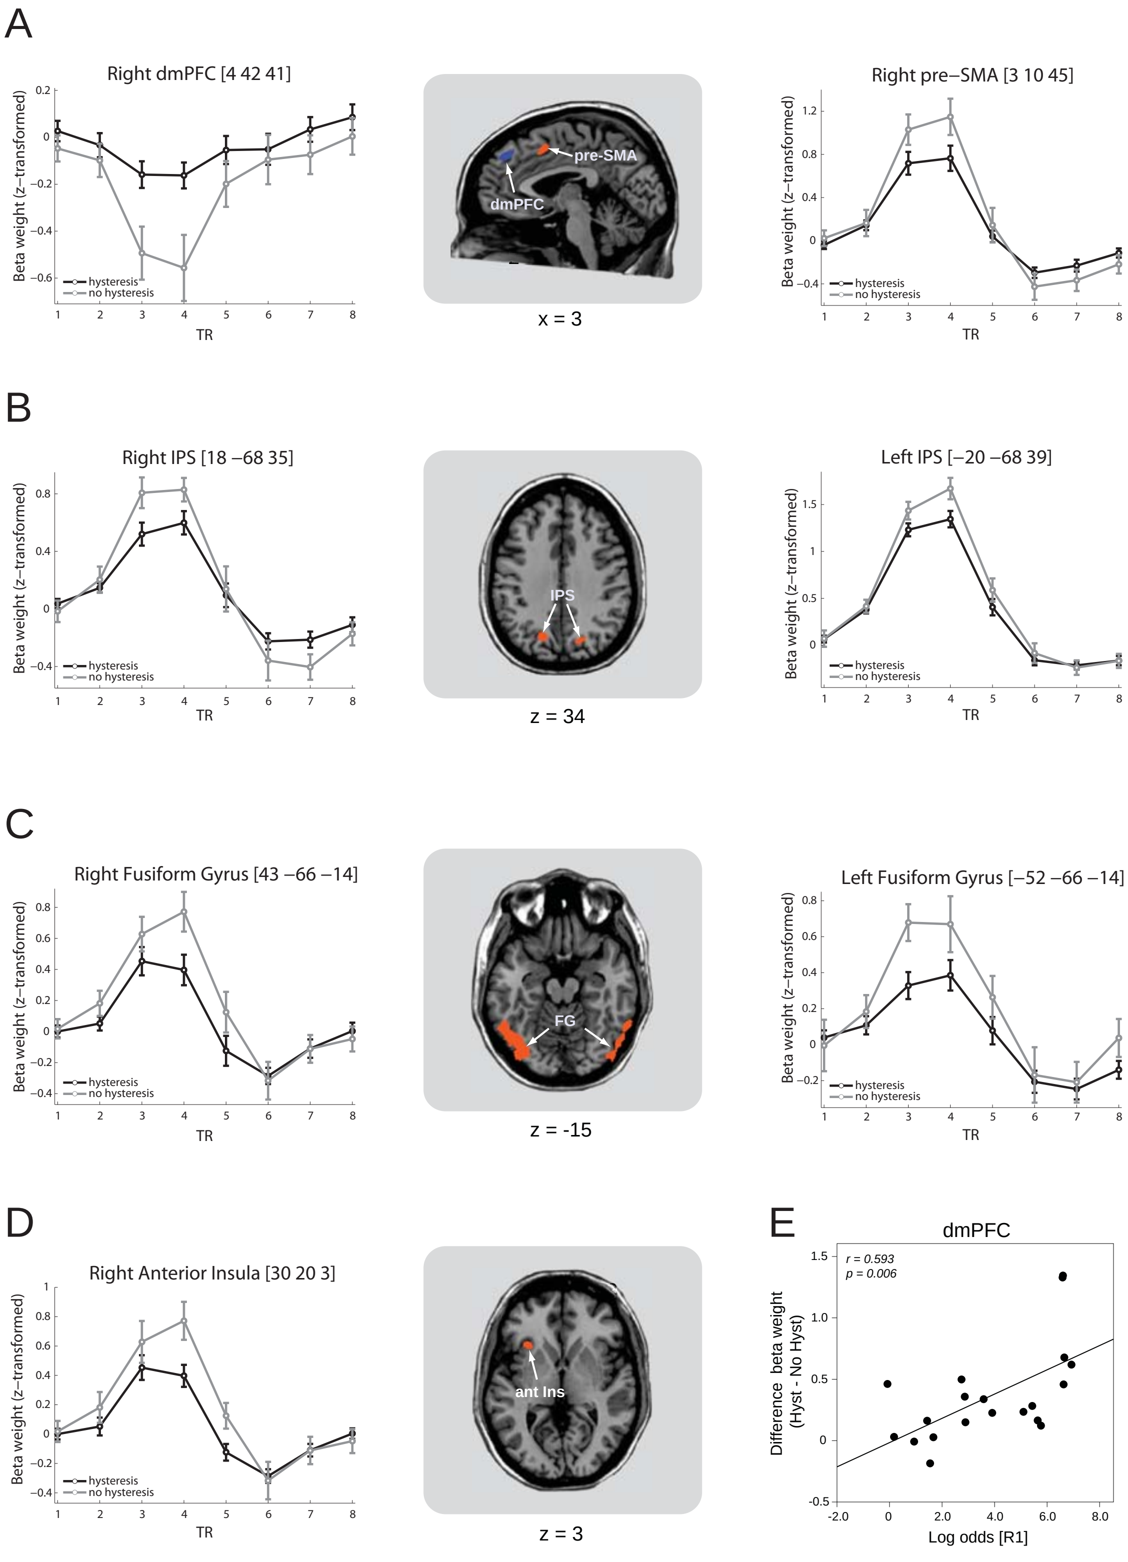
A
Right dmPFC [4 42 41]
Beta weight (z−transformed)
hysteresis
no hysteresis
TR
pre-SMA
dmPFC
x = 3
Right pre−SMA [3 10 45]
B
Right IPS [18 −68 35]
IPS
z = 34
Left IPS [−20 −68 39]
C
Right Fusiform Gyrus [43 −66 −14]
FG
z = -15
Left Fusiform Gyrus [−52 −66 −14]
D
Right Anterior Insula [30 20 3]
ant Ins
z = 3
E
dmPFC
r = 0.593
p = 0.006
Difference beta weight (Hyst - No Hyst)
Log odds [R1]

relation with behavior are likely to carry out central computations that determine the likelihood of percept maintenance. From all the areas identified in the hysteresis contrast that also showed stimulus-evoked responses, only the right dorsomedial prefrontal cortex positively correlated with the propensity to maintain perception across successive stimuli (r=0.59, p=0.0058, Figure 5-3e; all other p>0.1).

To identify regions expressing adaptation we searched for areas that showed a linear effect of AR, as observed in the behavioral data. To that end, we sorted all trials with two stimuli according to the aspect ratio of the first stimulus. We then identified voxels which showed a significant difference between ARs. Subsequently, we tested which of these areas showed a significant positive BOLD response to the first stimulus (on trials with only one stimulus) and a significant linear increase of the peak BOLD response with AR (on trials with two stimuli). This procedure was motivated by the following rationale: The three ARs used were parametrically biased towards perceiving 90° in the first stimulus. Thus, assuming that neurons are adapted by the presentation of the first stimulus and that the response to the second stimulus should reflect their state of adaptation, adaptation to 0° should be lowest for AR=1.2 and highest for AR=1. This entails that the BOLD response in any region displaying this pattern of adaptation should be highest for AR=1.2 and lowest for AR=1. The only region exhibiting this profile was found in the left occipital cortex (Figure 5-4, $F(1,19)=12.672$, p=0.002, $\eta^2=0.400$), slightly dorsal from the occipital pole, whose location [-17 -100 5] is in good agreement with the localization of human V2/V3 by Shipp et al. (1995) and Wohlschläger et al. (2005). This result was

◂◂ Figure 5-3. fMRI results hysteresis: Results of the contrast [hysteresis vs. no hysteresis] > [0° vs. 90°]. (a) Sagittal view of the right dmPFC and the right pre-SMA, and their deconvolved time courses. Note that although the dmPFC shows a negative time course, the absolute BOLD response is higher for no hysteresis than for hysteresis trials. (b) Transversal view of bilateral IPS and respective deconvolved time courses. (c) Transversal view of the bilateral fusiform gyrus and respective deconvolved time courses. (d) Transversal view of the right anterior insula and deconvolved time course. (e) Correlation between the size of the individual hysteresis effect (log odds of R1) and the difference between the peak BOLD amplitude for hysteresis and no hysteresis trials in the right dmPFC. Error bars represent the standard error of the mean, coordinates are given in standard Talairach space. Results are overlaid on the Montreal Neurological Institute standard brain, in radiological convention (left is right). Time axes are in TRs (1 TR = 2.5 s).

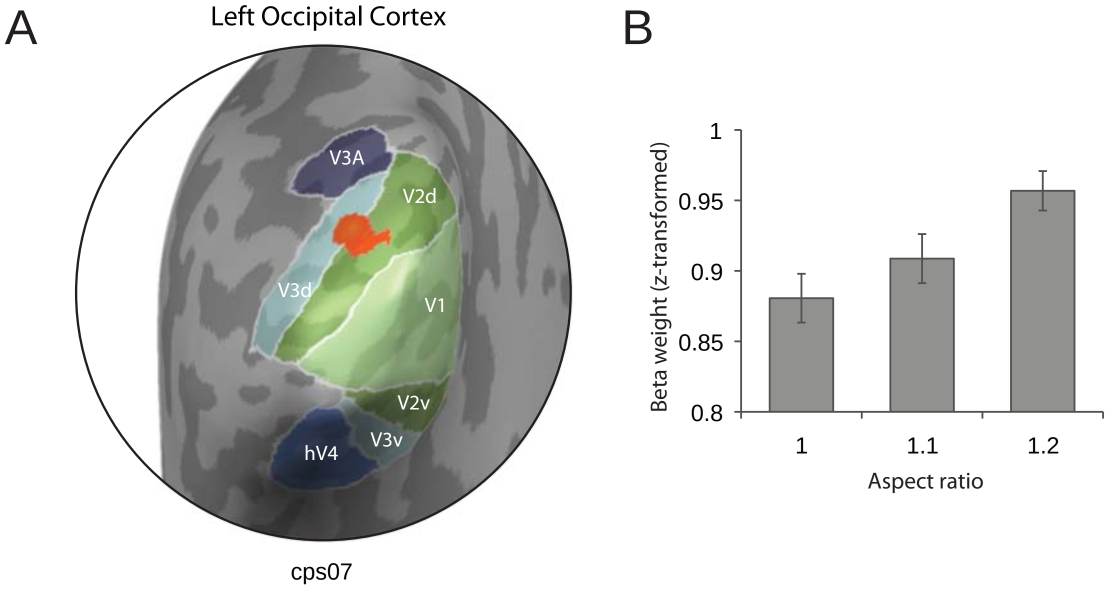

Figure 5-4. fMRI results adaptation. (a) Group activity in the dorsal occipital cortex, overlaid on the retinotopic map of the left inflated hemisphere of a representative subject (cps07) for whom retinotopic mapping was available. The cluster showing a linear effect of aspect ratio encompasses dorsal V2 and V3. Light and dark gray depict gyral and sulcal surfaces, respectively. White lines indicate the borders of early visual areas shown in different colors. (b) Peak BOLD amplitude per aspect ratio on trials with two stimuli. Error bars represent the standard error of the mean, corrected for between-subject variability (Cousineau, 2005; Morey, 2008).

also obtained when the ROI analyses were based on the same data as for the hysteresis contrast, i.e., only trials in which subjects responded 0° or 90° to the first and 0° to the second stimulus ($F(1,19)=10.587$, $p=0.004$, $\eta^2=0.358$). It is interesting to note that this location coincides with the representation of the lower right quadrant of the visual field, which exhibits the highest switch rate in binocular rivalry (Chen & He, 2003).

We finally investigated whether areas exhibiting hysteresis and adaptation would also express the respective other effect. To this end, we tested whether the areas expressing adaptation showed a difference between hysteresis and no hysteresis trials, and whether the areas expressing hysteresis showed a linear effect of AR. To assure that our conclusions were based on the same set of data, we used only the trials in which subjects responded 0° or 90° to the first and 0° to the second stimulus for this analysis. None of the effects were significant, mirroring the behavioral results in which hysteresis and adaptation also did not interact (no significant linear trend in hysteresis areas, all $p>0.2$; no significant difference between hysteresis and no hysteresis in V2/V3, $p>0.3$).

None of the effects could be accounted for by eye movements. When we ran the same contrasts as for the main analyses on ROIs defined for the eyeballs (see Methods), we did not observe any significant effect at $p<0.05$. This is in line with the eye tracking data from the psychophysical experiment.

Discussion

Taken together, our results reveal that hysteresis and adaptation are functionally dissociated in the human brain. Whereas a widespread network of frontal, parietal, and ventral occipital brain areas is involved in perceptual hysteresis, BOLD responses following the behavioral profile of the adaptation effect are only evident in extrastriate cortex. Previous psychophysical work and computational models have explained hysteresis and adaptation as resulting from the same mechanism, either the amount of neuronal adaptation alone (R. Blake, et al., 2003; Chen & He, 2004; Orbach, et al., 1963), or a "persistent bias" (Gepshtein & Kubovy, 2005). The functional division observed in our data directly refutes these interpretations. More recent models which propose separate mechanisms for hysteresis and adaptation seem to account better for our results. However, several of these models co-localize the mechanisms for hysteresis and adaptation to the same early sensory area (Brascamp, et al., 2009; Noest, et al., 2007; Wilson, 2007), an idea that also seems at odds with our data. A parsimonious explanation of the coexistence of hysteresis and adaptation can be given if both phenomena are understood in a Bayesian framework. In this framework, perception is the result of an inference, in which the available sensory information is compared to a prediction derived from previous experience (Yuille & Kersten, 2006). In our case, such a prediction can be derived from the first stimulus and is subsequently tested against the evidence provided in the second stimulus. Several models propose that this process indeed involves an interaction between higher-level brain areas generating predictions and earlier brain areas testing these predictions (Friston, 2005; Mumford, 1992; Rao & Ballard, 1999). As will be outlined below, hysteresis can be understood as the result of a prediction, whereas adaptation changes the available sensory evidence (Stocker & Simoncelli, 2006). Thus, a Bayesian model may explain the different perceptual effects of adaptation and hysteresis, as well as their separation into different brain areas.

A Bayesian account of hysteresis and adaptation

Both the multistability (Hohwy, Roepstorff, & Friston, 2008; Sundareswara & Schrater, 2008) and the divergent effects of previous experience (Stocker & Simoncelli, 2006) can be accounted for by the Bayesian framework. Importantly, the Bayesian framework assumes that perception depends on the probability distribution of the available evidence on the one hand, and on a prior (the expectation that the world will be in a particular state) on the other hand. The probability distribution of sensory evidence (the likelihood function) in multistable stimuli is multimodal, reflecting the different possible interpretations. What we perceive is the result of a multiplication of the likelihood function with the prior, i.e., the posterior distribution. This posterior distribution is again multimodal, but given the interaction with the prior, one of the peaks will be maximal, and this is the perceptual interpretation that gains dominance. Importantly, previous evidence for one interpretation (e.g., a previous percept) changes the prior distribution towards this interpretation, thus inducing hysteresis. Adaptation, however, does not affect the prior, but the likelihood function (i.e., the sensory evidence) itself: It narrows the conditional probability density around the adapted interpretation (Stocker & Simoncelli, 2006). Such an asymmetric change of the shape of the likelihood function produces the repulsive aftereffect typical for adaptation. Thus, a Bayesian framework of multistable perception is compatible with a dissociation between adaptation and hysteresis in neural space: One network computes the prior, whereas another, sensory network changes its function through adaptation, thus changing the likelihood function.

The cortical network expressing hysteresis

Understood in a Bayesian framework as outlined above, our contrast [hysteresis vs. no hysteresis] should reveal brain areas involved in the generation and testing of a prior against the incoming evidence. Following recent models of how the brain implements Bayesian inference, we expect that this occurs in an interaction between lower sensory and higher, prediction generating areas (Friston, 2005; Mumford, 1992; Rao & Ballard, 1999). The network expressing differential activity for hysteresis versus no hysteresis trials comprised the bilateral fusiform gyrus, bilateral intraparietal sulcus (IPS), right pre-SMA, right anterior insula, as well as the right dorsomedial prefrontal cortex (the mesial aspect of Brodmann Area 8). In our paradigm, the sensory feature for which a prior was generated and which was perceptually stabilized was the orientation of the dot lattices. Several of the areas we have identified are indeed processing orientation, in particular the fusiform gyrus (Orban, Dupont, Vogels, Bormans, & Mortelmans, 1997) and the IPS (Faillenot, Decety, & Jeannerod, 1999; Faillenot, Sunaert, Van Hecke, & Orban, 2001; Shikata et al., 2001). Furthermore, a recent study in macaque monkeys found that neurons in the IPS participate in grouping dots into oriented lines (Yokoi & Komatsu, 2009), which is another basic operation required to perceive orientation in our stimuli. These areas are thus prime candidates for the testing of predictions against incoming evidence. Regarding the higher order areas, it has been proposed that perceptual hysteresis might be closely linked to working memory (Maier, et al., 2003; Sterzer & Rees, 2008), which would not be predicted by a Bayesian account. The pre-SMA (Petit, Courtney, Ungerleider, & Haxby, 1998) and also the IPS (Todd & Marois, 2004) have been shown to be active in working memory tasks, however, the remaining areas do not show strong overlap with regions that are specifically involved in working memory for orientations, in particular because of the relative lack of lateral frontal activity in our study (Cornette, Dupont, & Orban, 2002; Cornette, Dupont, Salmon, & Orban, 2001). This might be explained by the fact that hysteresis is a form of implicit memory, whereas working memory is explicit. Importantly, Bayesian inference in perception is implicit. The IPS

(Corbetta & Shulman, 2002), pre-SMA (M. D. Fox, Corbetta, Snyder, Vincent, & Raichle, 2006) and the right anterior insula (Eckert et al., 2009; Sridharan, Levitin, & Menon, 2008) are also implicated in the control of top-down attention. Importantly, Kanai and Verstraten (2006) have shown in a perceptual memory paradigm employing ambiguous motion that if attention is distracted during the presentation of the first stimulus or even in the blank period before the presentation of the second stimulus, perceptual stabilization fails. Thus, top-down attention seems to be essential to instantiate perceptual memory. In a Bayesian framework, spatial attention can be understood as a prediction of where a particular stimulus (or percept) will appear (Chikkerur, Serre, Tan, & Poggio, 2010). A test or update of the prior would then evoke transient activity in these areas, as has been found for perceptual switches in investigating binocular rivalry (Lumer, Friston, & Rees, 1998) and apparent motion (Muckli et al., 2002; Müller et al., 2005; Sterzer, Russ, Preibisch, & Kleinschmidt, 2002). Last, the only region in the network expressing perceptual hysteresis showing a correlation with the individual size of the hysteresis effect was the right dorsomedial prefrontal cortex, indicating that it plays a central role in the processes underlying perceptual hysteresis. Crucially, this area has been shown to be involved in generating predictions under uncertainty (Volz, Schubotz, & von Cramon, 2003). In a Bayesian framework, this area would serve the crucial function of generating and updating predictions about the upcoming stimulus, which are then tested in the remaining areas.

A local network expressing adaptation

In contrast to hysteresis, adaptation was only evident in early extrastriate cortex, likely dorsal areas V2/V3. Areas V2/V3 contain orientation tuned neurons (Felleman & Van Essen, 1987; Levitt, Kiper, & Movshon, 1994). More importantly, however, these neurons also respond to dots which are perceptually bound into oriented lines (Peterhans, Heider, & Baumann, 2005) and more generally to illusory contours (Peterhans & von der Heydt, 1991; von der Heydt, Peterhans, & Baumgartner, 1984). Such a binding operation is required in order for dot lattices to appear as having an orientation. The involvement of an early sensory area fits well with the Bayesian framework, which predicts that adaptation does not change the prior, but the shape of the distribution of the available evidence (Stocker & Simoncelli, 2006), which is thought to be computed by early sensory areas. Interestingly, the fact that adaptation was evident at an earlier level of the visual hierarchy than hysteresis is in accordance with a previous finding (Gepshtein & Kubovy, 2005): Adaptation effects in multistable dot lattices were found to be orientation specific, i.e., the effect of AR on the second lattice vanishes when the orientation of the main axis differs by more than 15°. In contrast, hysteresis persists up to offsets of 30°, again speaking for a dissociation of the underlying mechanisms and for an involvement of regions with less precise orientation tuning in hysteresis. The absence of a hysteresis effect in areas V2/V3 indicates that the priors are not tested against the evidence in these early areas, only at a later stage. In fact, Figure 5-2b shows that hysteresis can overrule the effects of adaptation. For example, it is possible to stabilize a 0° percept even at $AR=1.3^{-1}$, at which strong adaptation to 0° should lead to a repulsive after effect. Thus, it is conceivable that the output of areas V2/V3 serves as input for subsequent processing stages, e.g., in the fusiform gyrus or IPS, where the hypothesis test is carried out. In fact, if hysteresis serves to stabilize percepts against constantly changing low-level stimulus features (Kleinschmidt, et al., 2002), the larger receptive fields of higher visual areas are advantageous as compared to small receptive fields in earlier visual areas because they allow for a certain degree of invariance. Psychophysical evidence suggests larger receptive fields involved in hysteresis by showing that hysteresis is retinotopically organized, but extends beyond the exact location where the first stimulus was presented (Knapen, Brascamp, Adams, & Graf, 2009). Interestingly, our psychophysical data indicates that the adaptation effect persists even if the adapting orientation is not consciously perceived, whereas hysteresis is by definition a conscious perceptual process. This dichotomy might be explained by the fact that hysteresis involves distributed processing, whereas adaptation is a purely local phenomenon: Several current theories of conscious perception predict that the spatial scale at which neuronal processing occurs determines whether content enters awareness or not (Dehaene, et al., 2006; Edelman, et al., 2011; Melloni & Singer, 2010).

Conclusions

Hysteresis and adaptation in multistable stimuli illustrate how we generally use previous experience to guide perception: The brain needs to cope with constantly changing and often unreliable or ambiguous input, and exploiting the regularities of the world is an efficient strategy to achieve this goal. Reoccurring patterns are informative, because they allow predicting what will happen next. This is reflected by the hysteresis effect. On the other hand, reoccurring patterns are redundant and could be discarded from further processing, thus emphasizing new information. This is evident in the repulsive aftereffects resulting from adaptation. Apparently, the brain is optimized for both strategies. Our study shows that they are implemented in different neural circuits. This separation possibly endows the brain with the flexibility to switch between two modes, one that emphasizes the new, and one that exploits already available information. Such flexibility is crucial in a non-stationary environment where one should always be prepared to predict the unexpected.

6

GENERAL DISCUSSION

In the following, I will discuss the major findings of this thesis, and their implication for the study of consciousness. In particular, I will present a dual route model that can explain the commonalities and dissociation in subjective and objective learning effects presented in Chapters 2 and 3, as well as dissociations between conscious perception and unconscious processing reported in the literature. Furthermore, I will discuss how a Bayesian account of brain processing implicated by the data presented in Chapters 4 and 5 can account for the difference between conscious and unconscious processing, and how this relates to current theories of consciousness.

A DUAL ROUTE MODEL FOR LEARNING EFFECTS ON PROCESSING AND EXPERIENCING

The data presented in Chapters 2 and 3 provide evidence for a plastic threshold of conscious perception. In particular, not only the processing of initially indiscriminable stimuli improves with practice, but training also leads to a change in subjective experience of the masked stimuli. These practice effects persist up to several months. Importantly, I found that improvements in subjective experience and objective performance are not simple byproducts of each other: They differ in their learning time courses, but strikingly also in the spatial extend of the training effects. The differential transfer across the retinotopic map allows for the conclusion that the receptive field structure underlying both learning effects is different and thus, that the brain areas in which the learning effects take place must also differ. Contrary to what might be intuitively expected, subjective experience was neither necessary nor sufficient for objective performance; also, objective performance was not sufficient to explain subjective experience. Together, the data speak for the intriguing possibility that the learning effects on objective performance and subjective experience occur in parallel channels with independent sources of noise.

It is often assumed that conscious perception arises at the end of a serially progressing sequence of processing steps (Dehaene & Naccache, 2001; Lau, 2008b). Since subjective awareness of a stimulus and objective performance on a task are usually highly correlated, such

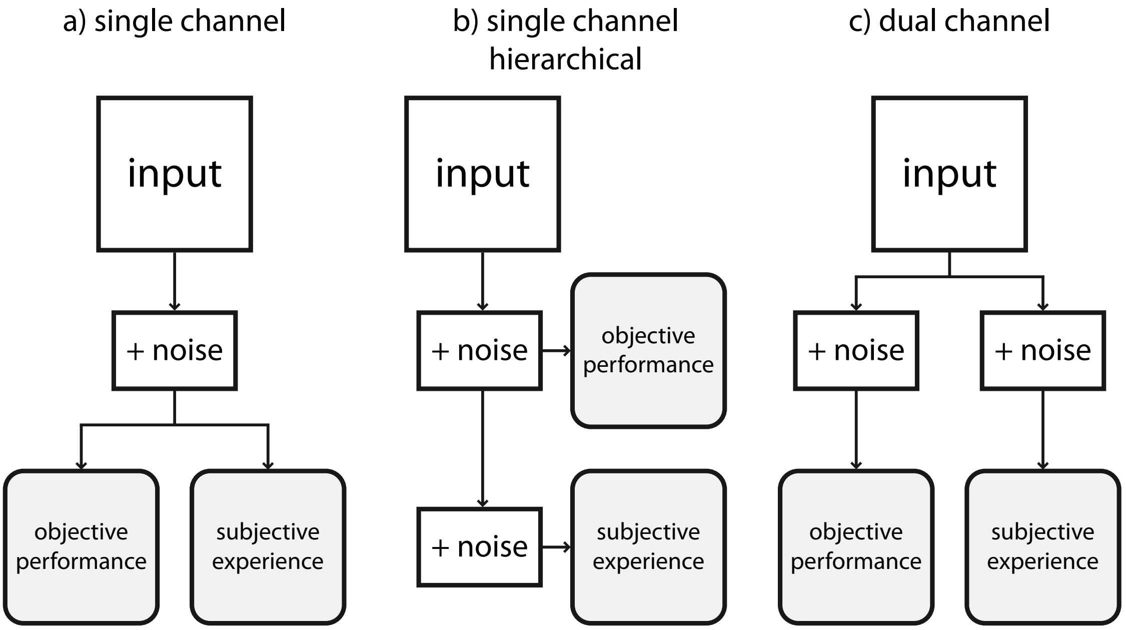

Figure 6-1. Illustrations of three different models of the relationship between objective performance and subjective experience. (a) Single channel model. The output of a single channel determines both subjective experience and objective performance. This model predicts a perfect correlation between subjective experience and objective performance. (b) Hierarchical single channel model. The input is first processed by unconsciously operating modules. The output of these modules determines objective performance. In a second step, it is further processed to give rise to subjective experience. This additional processing adds more noise, thus, subjective experience will always be noisier than objective performance. (c) Dual channel model. Unconscious computations and computations giving rise to conscious experience proceed in parallel. They are based on the same input but have independent source of noise. The common input accounts for the correlation between subjective experience and objective performance, while the independent sources of noise account for the dissociations between them.

proposals fare well with a single route model (Figure 6-1a) in which input is first processed by unconsciously operating modules and then fed into downstream processors which ultimately give rise to conscious experience and objective performance. However, models of this kind often fail to explain dissociations between objective performance and subjective awareness. This can be achieved by adding a second source of noise that renders conscious perception always noisier than the unconscious processing of the input (Lau, 2011), thus accounting for phenomena like blindsight (Figure 6-1b). However, my data speak for a parallel model, in which unconscious computations and computations giving rise to conscious experience proceed in parallel, based on the same input but with independent sources of noise (Figure 6-1c). Physiologically, this could be thought of as a mapping of sensory input to relatively small receptive fields coding for shapes (square or diamond) in one area (likely V4), and to larger receptive fields coding for the subjective quality of the percept in another area. Both types of receptive fields receive information from one sensory input stage. The learning-induced change happens not at the sensory input stage, but affects the mapping from the sensory input to the receptive fields coding for shape or the subjective quality of the percept (for example through channel weighting, Dosher & Lu, 1999). If the stimulus is moved to another position, the mapping to the smaller receptive fields breaks down, whereas the mapping to the larger receptive field is preserved.

My results are in line with a growing number of empirical findings pointing to a dissociation of subjective experience and objective performance in the healthy human brain (Bar & Biederman, 1999; Lau & Passingham, 2006; Rounis, Maniscalco, Rothwell, Passingham, & Lau, 2010; Sahraie, et al., 1997). Neuroimaging data points to the dorsolateral prefrontal cortex (DLPFC) as an important node in the neural processes underlying subjective experience (Lau & Passingham, 2006; Sahraie, et al., 1997). The DLPFC exhibits a receptive field structure which is compatible with our results (Rainer, et al., 1998; Suzuki & Azuma, 1983) and shows plastic changes in response to perceptual learning tasks (Rainer & Miller, 2000). Importantly, electrophysiological experiments in monkeys show that the pattern of changes observed in V4 as a result of learning is quite different from that observed in the DLPFC even under identical stimulation conditions (Rainer, Lee, & Logothetis, 2004). For example, while neurons in the DLPFC respond to objects independently of their degradation level after training, neurons in V4 respond differentially to degraded and undegraded learned objects. Also, the average firing rate in the DLPFC is reduced after training, while the firing rate of V4 neurons does not change. This indicates that the prefrontral cortex does not simply inherit the physiological learning effects from V4. The existence of multiple parallel connections from striate and extrastriate cortex to the prefrontal cortex (Barbas, 1988; Barbas & Mesulam, 1985; Young, 1992) lends further credibility to a dual route model as I propose.

Alternatively to the DLPFC, it has been proposed that regions in the anterior inferotemporal cortex underlie subjective experience (Bar & Biederman, 1999; Bar et al., 2001). For example, Bar and colleagues (2001) found in an fMRI experiment that only the fusiform gyrus follows the behavioral profile of subjective recognition of visually presented, degraded objects, while for instance the BOLD response of area V4 was indistinguishable for recognized and unrecognized objects. Plasticity in monkey inferotemporal cortex (reviewed in Hoffman & Logothetis, 2009) and its human homologue (Kourtzi, Betts, Sarkheil, & Welchman, 2005) is well established, and the receptive fields in the anterior inferotemporal cortex would also be able to support the spatial transfer of 6.6° in our experiment (Desimone & Gross, 1979). Tracing studies in monkeys have shown that not all information enters the inferotemporal cortex via area V4, but that there is direct input from areas V2 and V3 to areas TEO (Distler, Boussaoud, Desimone, & Ungerleider, 1993) and perhaps even TE (Baizer, Ungerleider, & Desimone, 1991). Thus, a dual route model could also be established along ventral visual areas.

The differential transfer of learning effects across the retinotopic map is also compatible with theories that pose that mere processing occurs in isolated modules, while experiencing requires global availability (Dehaene & Naccache, 2001) or long-distance synchrony between brain areas (Melloni & Singer, 2010). Here, the learning effects on processing would be the result of local changes in V4 (as above), however, the transfer of learning effects on subjective experience would be the result of the global distribution of information that is thought to be the signature of conscious access (Dehaene & Naccache, 2001). Modeling work suggests that the parallel processing of information for objective performance and subjective experience can be envisaged as a race between a conscious route and an unconscious route of processing (Del Cul, et al., 2009). If the route for subjective experience wins and information is eventually globally available, it is not longer bound to a specific receptive field size. Still, this model requires that information can become globally available even without sufficient processing in V4, and thus constitutes a dual route model as well.

Taken together, a dual route model offers a parsimonious explanation for the results of Chapter 2 and 3. Furthermore, the data substantiate recent modeling work which had only suggested that dual route models of conscious perception and unconscious processing are indeed feasible, but did not provide empirical evidence (Del Cul, et al., 2009; Pasquali, Timmermans, & Cleeremans, 2010). Last, the dual route model proposed here can also account for other results reported in the literature. For example, blindsight in normal observers (Lau & Passingham, 2006; Meeres & Graves, 1990) can arise in this model if there is more noise in the channel leading to the area(s) involved in subjective experience than in the channels responsible for merely processing the stimulus to subserve objective performance. Also, apparently illusory percepts (Boehler, et al., 2008; Del Cul, et al., 2009; Summerfield, et al., 2002), i.e., high subjective visibility ratings on objectively incorrect trials, can be explained by higher noise in the channel for objective performance and lower noise in the channel for conscious perception. Neuroimaging studies are now on the way to map out the two channels suggested by psychophysics and to identify their temporal relationship.

Conscious vs. unconscious processing and Bayesian inference

The data presented in Chapters 4 and 5 fit well with current theories of perception that frame the brain as a Bayesian inference or Helmholtz machine (Dayan, Hinton, Neal, & Zemel, 1995; Di Lollo, et al., 2000; Friston, 2010; Lee & Mumford, 2003; Rao & Ballard, 1999). In particular, these theories highlight that perception is the result of (or at least in accordance with) an unconscious inference, which utilizes both the available evidence and priors (built in and flexibly acquired) to arrive at a probable percept. However, in the context of conscious perception, these theories do not specify how such Bayesian computations are related to conscious or unconscious processing. Rather, the concept of hypothesis testing is claimed to be a general mechanism of brain function (e.g., Friston, 2010), i.e., it equally applies to conscious and unconscious processing. Given that we know that only some processing leads to conscious experience, is it possible to account for the difference between conscious and unconscious processing with a Bayesian model of the brain? Here, I will consider three possibilities which map into three theoretical views of conscious perception currently discussed in the scientific community.

1. Higher order thought

So called "higher order thought" theories of consciousness postulate that to be conscious of something (a stimulus), the representation of the stimulus itself needs to be the target of yet another, higher order representation (Lycan, 2009). Thus, one is conscious *by virtue* of a higher order (meta-)representation, and not the representation of the stimulus itself. There is fierce philosophical debate regarding the cogency of this argument (Block, 2011; Rosenthal, 2002), however, several neuroscientific and cognitive theories of consciousness have strongly embraced this view (Cleeremans, 2008; Dienes, 2004; Lau, 2008b; Rolls, 2007; Singer, 2000). Does higher order thought theorizing allow for a distinction between conscious and unconscious processing within a Bayesian framework of brain function? Since higher order thought theories postulate two representations, one of the stimulus, and a meta-representation of the stimulus, Bayesian inference could underlie the computations leading to both of them. The brain would then iteratively test and update hypotheses about the sensory evidence, and in parallel or in a second step test the prediction that one has a meta-representation of the stimulus. Thus, Bayesian processing would be a general currency of conscious and unconscious processing, but it would be the *type* of representation that differentiates them. A similar proposal has been put forward by Lau (2008b). Here, higher order representations are learned from internal, unconscious signal and noise distributions of the stimulus. Decisions on both levels (signal and noise distributions as well as the higher order distributions) are made according to Bayesian decision theory relative to a respective criterion. Since the higher order representations need to be learned, they can in principle misrepresent the underlying signal and noise distributions, which is why the proposal can account for dissociations between subjective experience and objective performance.

2. Global Workspace models

One of the leading neuroscientific theories of consciousness posits that consciousness arises when a Global Neuronal Workspace has been ignited (Dehaene & Changeux, 2011; Dehaene & Naccache, 2001). In particular, unconscious processing is thought to occur in parallel in local modules, and if the information processed by these modules is strong enough or amplified through directed attention, it is dispatched to a network of neurons with long axons, preferentially located in the frontal and parietal cortices, that make the information "globally available" (Baars, 1997). Several other models make less strong predictions about the involvement of specific brain areas, but in principle agree that conscious access depends on global availability of information (Lamme, 2006; Melloni & Singer, 2010). In these frameworks, Bayesian inference could be a mechanism by which information is processed unconsciously in local modules. Bayesian inference could also operate at the global level, estimating a coherent solution from the distributed available evidence. Conscious experience would then refer to this global solution. Such a model would also account well for the observation that conscious perception is usually phenomenally unified, and not of disparate features. A similar solution has been proposed by Dehaene (2008) using an accumulator model. Here, it is postulated that sensory evidence is accumulated in parallel in local modules, which transmit their 'decisions' to the Global Neuronal Workspace, which in turn accumulates these decisions and then dispatches a global decision to the rest of the brain. Awareness is then of this global decision. However, it has to be noted that optimal Bayesian decision making takes *all* available information into account: If an object is composed of multiple features, independent processors can in principle find Bayesian solutions for the complete object even if only using the information available to them (e.g., one solution based only on orientation, one solution based only on color, etc.). However, if the processors are informed of each other, the solution they find together is usually better than the sum of the individual solutions (Adelson & Pentland, 1996). Human subjects have indeed been found to combine cues in a Bayesian optimal fashion (e.g., Ernst & Banks, 2002; Jacobs, 1999). Avoiding local maxima this way can be computationally formalized as "particle filtering" (Lee & Mumford, 2003), and could in principle occur vertically throughout a processing hierarchy as well as laterally between modules on the same hierarchical level.

3. Partial awareness

A relatively recent proposal is that in fact, we have conscious access not only to the global, unified solution, but also independently to all the individual parts making up that solution (Kouider, de Gardelle, Sackur, & Dupoux, 2010). This so-called "partial awareness hypothesis" posits that representations at each level of the processing hierarchy can be accessed independently; this allows on the one hand access to global meaning without detail (higher level in the hierarchy) and on the other hand access to low level features, which in turn accounts for the subjective impression of phenomenal richness. Awareness is usually complete, which amounts to having access to all levels, but can also be partial, i.e., access is possible to some, but not to other levels of the hierarchy. At the other extreme, subjects can also be completely unaware, i.e., they do not have access to any level of the processing hierarchy. Here, Bayesian inference could again operate at each level of the hierarchy, and one would be only aware of the outcome of the Bayesian operation that one has currently access to. However, the partial awareness hypothesis currently does not specify whether it is possible to be aware of several levels of the hierarchy in parallel, or whether conscious access is always limited to one level at a time. Thus, it actually fails to fully account for consciousness.

Taken together, Bayesian models of perception can in theory be expanded to account for a distinction between conscious and unconscious processing. Which of the above solutions (or any other) is the correct solution is still an open question for research. In a way, the number of possible solutions to expanding the Bayesian framework to incorporate conscious and unconscious processing reflects the multitude of proposals that currently exist to explain conscious perception in the first place. It also has to be taken into account that although Bayesian models are currently very popular and can account for a very wide range of data, there is also evidence that can be interpreted to speak against Bayesian optimal perception. Anderson, O'Vari and Barth (2011) recently reported that when moving occluders are used to generate illusory contours, subjects perceive multiple surfaces, although one surface would be sufficient to "explain" the sensory input. Since the additional surface seems superfluous, it can be argued that perception is not Bayesian optimal in this case (but see Fleming, 2011). Thus, further research is also needed on this side to establish whether the Bayesian approach is indeed ideal to understand perception and its neural underpinnings (also see Colombo & Seriès, 2011).

CONCLUSION

Despite the fact that previous experience is a key factor in several important models of brain function, its role in conscious perception has so far been severely neglected. This thesis was aimed at filling this ostensible gap. Studying the role of previous experience in conscious perception revealed several important characteristics of conscious perception itself. First, my data show that the threshold of conscious perception is not fixed, since I found that conscious perception is indeed malleable both by long-term and short-term experience: We can learn to see initially indiscriminable stimuli, and we can use our prior information to continue seeing stimuli even if they are massively degraded. This questions a wide spread belief and poses a serious challenge to research that needs to establish fixed thresholds in order to investigate unconscious perception. Second, my data show that the neural correlates of consciousness do not invariantly occur late but can arise with flexible latencies, depending on how conscious experience comes about. This finding raises the possibility that yet other factors, like attention (Noguchi, Tanabe, Sadato, Hoshiyama, & Kakigi, 2007), exert similar effects, moving the neural correlates of consciousness forward or backward in time. Third, conscious perception can be affected by previous experience in at least two ways, rendering us more likely to perceive the same stimulus again (hysteresis), or more likely to see something else (adaptation). The role of previous experience is thus not only to stabilize conscious perception by exploiting experience-based predictions, but also to retain sensitivity for new information. Overall, conscious perception exhibits a remarkable degree of plasticity that is not but should be taken into account by current theories of consciousness.

In addition, my data also speak against theories that equate performance with experience. In accord with studies on brain damaged patients, my data evidence that conscious experience and the mere processing of stimuli are distinct. An important new proposal that can be derived from my studies is that the neural processes underlying these two functions rely on parallel channels. Although the brain is widely conceived as processing input in a highly parallel fashion, this notion had not been fully taken on by theories distinguishing conscious from unconscious processing. Furthermore, the distinction has to be taken into account not only when theorizing about consciousness, but also when measuring it.

Beyond that, the paradigms used in this thesis may offer an alternative way of studying the neuronal correlates of consciousness. Rather than contrasting seen with unseen conditions by manipulating stimulus evidence or relying on the spontaneous fluctuations of brain activity, one can train the brain to go from unconscious processing to conscious perception without changing anything but the brain itself. This way, we might understand how the brain *itself* generates appropriate conditions to decipher the visual world and to generate a conscious percept. This is in the end the question that really has to be solved: How does the brain bring consciousness onto itself?

Z

ZUSAMMENFASSUNG

Welche Faktoren bestimmen, ob ein Reiz bewußt wahrgenommen oder unbewußt verarbeitet wird? In der vorliegenden Arbeit wird untersucht, wie sich bestehendes Wissen auf die bewußte Wahrnehmung und auf die ihr zugrundeliegenden neuronalen Prozesse auswirkt. Ferner wird untersucht, ob bewußte Wahrnehmung erlernbar ist. Die Ergebnisse zeigen, daß bewußte Wahrnehmung nicht nur davon abhängt, wieviel Information ein Reiz liefert. Vielmehr ist sie das Ergebnis eines plastischen und integrativen Prozesses, im Zuge dessen aktuell eingehende Information mit zuvor erworbenem Wissen interagiert.

Die Frage, wie aus der Aktivität von Neuronen bewußte Wahrnehmung entsteht, ist eines der großen Rätsel der Neurowissenschaften (Koch, 2004). Wir haben den Eindruck, daß wir alles, was sich im uns umgebenden Raum befindet, klar wahrnehmen, doch dieser Eindruck täuscht. Man weiß inzwischen, daß ein Großteil der Information, die durch unsere Sinnesorgane aufgenommen wird, verarbeitet wird, ohne daß wir uns ihrer bewußt werden (Kouider & Dehaene, 2007). Was den Unterschied zwischen bewußter und unbewußter Wahrnehmung ausmacht, konnte bislang nicht abschließend geklärt werden. Fest steht, daß starke Reize das Bewußtsein normalerweise stets erreichen. Schwächere Reize werden hingegen häufig nur dann wahrgenommen, wenn ihnen aktiv Aufmerksamkeit geschenkt wird (Dehaene, et al., 2006).

In den hier vorgestellten Studien wird der Frage nachgegangen, ob es neben den Faktoren Intensität und Aufmerksamkeit noch andere Faktoren gibt, die bestimmen, ob ein Reiz bewußt wahrgenommen oder unbewußt verarbeitet wird. Insbesondere wird untersucht, wie sich eine vorherige Erfahrung auf die subjektive Wahrnehmung auswirkt und welche neuronalen Prozesse und Schaltkreise dem zugrunde liegen.

In den Kapiteln 2 und 3 wird untersucht, ob Wahrnehmung durch langfristiges, systematisches Üben beeinflußt werden kann. Es ist inzwischen etabliert, daß die Verarbeitung von Reizen in der Hirnrinde auch noch im Erwachsenenalter hochgradig plastisch ist (Ahissar & Hochstein, 1998; Fahle, 2009; E. J. Gibson, 1963; Goldstone, 1998). So können Versuchspersonen trainiert werden, sogar extrem kleine Unterschiede zwischen Stimuli zu unterscheiden (zum

Beispiel Poggio, et al., 1992). Allerdings ist bisher nicht untersucht worden, ob sich ein solches Training auch auf die subjektive Wahrnehmung der Reize auswirkt oder ob die Verbesserungen in der Unterscheidbarkeit der Reize allein auf unbewußte Verarbeitungsprozesse zurückführbar sind. Diese Frage ist auch klinisch relevant: So können manche Schlaganfallpatienten, die in Folge einer Schädigung der Sehrinde in einem Teil Ihres Gesichtsfeldes erblindet sind, Reize unterscheiden, die in den erblindeten Teil des Gesichtsfelds fallen. Allerdings geben diese Patienten an, daß sie die Reize nicht bewußt wahrnehmen. Diese Dissoziation zwischen Wahrnehmung und Verarbeitung wird als „blindsight" bezeichnet (Cowey, 2004; Stoerig, 2006). Rehabilitationsversuche haben bislang ergeben, daß durch gezieltes Üben die Fähigkeit der Patienten verbessert werden kann, die in das erblindete Gesichtsfeld fallenden Reize zu verarbeiten, und dies in seltenen Fällen auch mit zunehmender subjektiver Wahrnehmung der Reize einhergeht (Sahraie, et al., 2006; Zihl & von Cramon, 1985). Diese Ergebnisse deuten darauf hin, daß bewußte Wahrnehmung ähnlich wie andere kognitive Fähigkeiten durch gezieltes Training veränderbar ist. In gesunden Versuchspersonen kann dieser Frage nachgegangen werden, wenn an sich gut erkennbare Reize durch schnell aufeinanderfolgende Darbietung in ihrer Sichtbarkeit eingeschränkt werden („Maskierung"). So können Bedingungen hergestellt werden, in denen die Versuchsperson die dargebotenen Zielreize nicht auseinanderhalten kann. Gleichzeitig geben die Versuchspersonen an, der Reiz sei für sie unsichtbar. Wie in den Kapiteln 2 und 3 gezeigt wird, steigt im Zuge eines mehrtägigen Trainings die Fähigkeit, die Zielreize zu unterscheiden. Die Frage ist nun, ob das Training auch einen Zuwachs der subjektiven Wahrnehmung bedingt. Tatsächlich geben die Versuchspersonen als Folge des Trainings an, die Zielreize klarer wahrzunehmen. Dies zeigt, daß bewußte Wahrnehmung durch gezieltes Training verbessert werden kann, sogar wenn der Stimulus anfänglich objektiv unsichtbar ist. Hier zeigt sich, wie flexibel die neuronalen Prozesse sind, die der bewußten Wahrnehmung zugrunde liegen. Allerdings ist weder die bewußte Wahrnehmung die triviale Folge besserer Verarbeitung, noch die bessere Verarbeitung die Folge gesteigerter bewußter Wahrnehmung. Vielmehr zeigen die in Kapitel 3 vorgestellten Experimente, daß bewußte Wahrnehmung und die (unbewußte) Verarbeitung der Reize parallel vor sich gehen und daß die Hirnareale, die beiden Funktionen zugrunde liegen, nicht identisch sind. Diese Ergebnisse eröffnen nicht nur ein neues Forschungsgebiet zur Plastizität bewußter Wahrnehmung und zum Zusammenhang von bewußter Wahrnehmung und der Verarbeitung von Stimuli, sondern liefern auch wichtige Informationen für Rehabilitationsprogramme, deren Ziel darin besteht, die Plastizität des Gehirns auszunutzen, um beispielsweise Wahrnehmungsdefizite nach Hirnläsionen zu mildern.

Tritt die Plastizität der bewußten Wahrnehmung nur als Folge langfristigen Trainings zutage? In Kapitel 4 wird untersucht, wie sich Informationen, die kurz vor dem aktuell dargebotenen Stimulus erworben wurden, auf dessen bewußte Wahrnehmung auswirken. Wahrnehmung ist in den seltensten Fällen allein durch die Information, die ein Reiz trägt, bestimmt. Vielmehr wird die gerade eingehende Information mit bereits bestehenden Informationen integriert. So ist es zum Beispiel einfacher, ein Tier in seiner natürlichen Umgebung trotz Tarnung zu erkennen, wenn man weiß, welches Tier man sucht. Hier wird die bereits vorhandene Information über das Tier genutzt, um eine Vorhersage zu machen und so die Suche zu beschleunigen. In Kapitel 4 gehe ich der Frage nach, wie Vorhersagbarkeit mit der Stärke eines Reizes interagiert und wie dies zur bewußten Wahrnehmung des Reizes führt. Die Vorhersage beruht dabei auf Informationen, die erst kurz zuvor erworben wurden. Im hier genutzten Paradigma werden zunächst visuelle Reize mittels Rauschens unkenntlich gemacht. Durch schrittweise Reduktion des Rauschanteils wird der Reiz langsam sichtbarer, bis er schließlich klar erkennbar ist. Nach jedem Schritt gibt die Versuchsperson an, ob sie den Reiz wahrgenommen hat. Anschließend wird der Rauschanteil erneut schrittweise erhöht. Hat die Versuchsperson den Stimulus einmal erkannt, so erlaubt dies ihr, den Stimulus in der Folge auch trotz höheren Rauschanteils bewußt wahrzunehmen. Dieser die Wahrnehmung begünstigende Effekt wird als „Hysterese“ bezeichnet. Mittels Elektroenzephalographie kann man nun die neuronalen Prozesse vergleichen, die wahrgenommene von nicht wahrgenommenen Reizen unterscheiden. Auf diese Weise läßt sich der Effekt erfassen, den die zwischenzeitlich klare Erkennbarkeit auf die neuronalen Prozesse hat, die der bewußten Wahrnehmung zugrunde liegen: In der Phase, in der die Versuchsperson den Reiz noch nicht klar wahrgenommen hat, die bewußte Wahrnehmung also nur durch die im Reiz selbst vorhandene Information bestimmt wird, finden sich die ersten Unterschiede in den neuronalen Prozessen zwischen gesehenen und nicht gesehenen Durchgängen etwa 300 ms nach der Darbietung des Stimulus. Hingegen findet sich diese Differenzierung schon bei 200 ms, wenn die Versuchsperson den Reiz vorher schon einmal klar erkannt hat und infolgedessen voraussagen kann, wie der Stimulus aussieht. Dies zeigt, daß vorherige Erfahrung einen starken Einfluß darauf hat, ob wir einen Stimulus bewußt wahrnehmen. Zudem deutet die zeitliche Verschiebung der neuronalen Korrelate der bewußten Wahrnehmung von 300 ms ohne Vorhersage zu 200 ms mit Vorhersage darauf hin, daß die Prozesse, die der bewußten Wahrnehmung zugrunde liegen, keine feste zeitliche Beziehung zum Reiz haben. Sie hängen vielmehr davon ab, wie schnell aktuelle und bereits bestehende Informationen konvergieren. Hier zeigt sich erneut, daß die Prozesse, die zur bewußten Wahrnehmung führen, wesentlich flexibler sind als bisher angenommen. Die Ergebnisse stellen somit Theorien, die vorhersagen, daß die neuronalen Korrelate der bewußten Wahrnehmung unweigerlich spät in der Verarbeitungshierarchie auftreten (Dehaene & Naccache, 2001; Lamme, 2006), direkt in Frage.

Vorherige Erfahrung führt allerdings nicht immer dazu, daß die bewußte Wahrnehmung eines Stimulus erleichtert oder stabilisiert wird. Vielmehr kann vorherige Erfahrung auch das Gegenteil bewirken. Betrachtet man zum Beispiel für längere Zeit einen Stimulus, der sich in eine Richtung bewegt, und lenkt dann seinen Blick auf einen statischen Stimulus, so scheint sich letzterer in die gegensätzliche Richtung des zuerst betrachteten Reizes zu bewegen (Purkinje, 1820). Solche repulsiven Effekte werden allgemein auf neuronale Adaptation zurückgeführt (zum Beispiel Anstis, et al., 1998). In Kapitel 5 gehe ich der Frage nach, ob die stabilisierende Wirkung vorheriger Erfahrung („Hysterese“) und die destabilisierende Wirkung vorheriger Erfahrung („Adaptation“) auf denselben Mechanismus zurückführbar sind. Zu diesem Zweck verwende ich sogenannte „multistabile Stimuli". Diese Stimuli sind so konstruiert, daß sie mehrere perzeptuelle Interpretationen zulassen. Der Stimulus allein bestimmt dann nicht, welche Interpretation bewußt wahrgenommen wird. Vielmehr ist die Wahrnehmung hier stark davon abhängig, welche Interpretation kurz zuvor wahrgenommen wurde (Pearson & Brascamp, 2008). Mittels funktioneller Magnetresonanztomographie untersuche ich, ob die Verhaltenseffekte Hysterese und Adaptation auf Aktivität in denselben Hirnarealen zurückgehen. Die Ergebnisse zeigen, daß Hysterese mit Aktivität in einem weitverzweigten Netzwerk korreliert, das visuelle, aber auch parietale und frontale Areale einschließt, wobei der mediale präfrontale Cortex eine entscheidende Rolle zu spielen scheint. Im Gegensatz dazu ist der Adaptationseffekt nur mit lokaler Aktivität in frühen visuellen Arealen korreliert (V2/V3). Diese räumliche Aufteilung liegt vermutlich der Fähigkeit des Gehirns, eine Balance zwischen Vorhersagbarkeit und Sensitivität für Neues aufrechtzuerhalten, zugrunde.

Zusammenfassend zeigen die hier vorgestellten Studien, daß die Prozesse, die der bewußten Wahrnehmung zugrunde liegen, flexibel und sogar plastisch sind. Zum einen kann langfristiges Training dazu führen, daß Reize, die unter normalen Umständen nicht bewußt wahrgenommen werden, Zugang zu unserem subjektiven Empfinden finden (Kapitel 2 und 3). Zum anderen können Erfahrungen, die kurz vor dem aktuellen Reiz gemacht wurden, die Wahrnehmung und ihre neuronalen Korrelate verändern (Kapitel 3 und 4). Diese Befunde stellen verschiedene Annahmen, die häufig bezüglich der bewußten Wahrnehmung gemacht werden, in Frage: Zum einen zeigt sich, daß die Schwelle der bewußten Wahrnehmung nicht fest, sondern flexibel ist. Zudem sind auch die neuronalen Korrelate der bewußten Wahrnehmung keinesfalls zeitlich stabil, sondern können durch vorherige Erfahrung beschleunigt werden. Dies eröffnet die Möglichkeit, daß auch andere kognitive Prozesse wie zum Beispiel Aufmerksamkeit (Noguchi, et al., 2007), die neuronalen Korrelate der bewußten Wahrnehmung zeitlich beschleunigen oder auch verlangsamen könnten. Schließlich zeigen meine Daten auch, daß Theorien, in denen die reine Verarbeitung von Reizen mit ihrer bewußten Wahrnehmung gleichgesetzt wird, revidiert werden müssen. Die Ergebnisse stehen somit im Einklang mit Daten von Schlaganfallpatienten, bei denen bewußte Wahrnehmung und unbewußte Verarbeitung ebenfalls auseinanderfallen. Ein von mir vorgeschlagenes Modell, in dem die Prozesse der unbewußten Verarbeitung und die Prozesse, die zur bewußten Wahrnehmung führen, parallel ablaufen, kann diese Dissoziationen erklären. Zu guter Letzt stellen die hier genutzten Paradigmen eine neue Möglichkeit, die neuronalen Korrelate der bewußten Wahrnehmung zu untersuchen, dar. Anstatt Bedingungen zu vergleichen, in denen die bewußte Wahrnehmung von der Stärke des Reizes oder von spontanen Fluktuationen der Hirnaktivität abhängt, kann man das Gehirn trainieren, die Grenze zwischen unbewußter Verarbeitung und bewußter Wahrnehmung zu überschreiten. Auf diese Art besteht die Chance zu verstehen, wie das Gehirn selbst Bewußtsein erzeugt. Dies ist letztendlich die entscheidende Frage.

Dieser Text beruht zum Teil auf folgender Publikation: Melloni L, Schwiedrzik CM (2011). Bewußte Wahrnehmung als dynamischer und plastischer Prozeß. Max Planck Gesellschaft Jahrbuch 2011, http://www.mpg.de/166008/jahrbuecher.

R

REFERENCES

Adelson, E. H., & Pentland, A. P. (1996). The perception of shading and reflectance. In D. Knill & W. Richards (Eds.), *Perception as Bayesian inference* (pp. 409-423). New York, NY: Cambridge University Press.

Ahissar, M., & Hochstein, S. (1998). Perceptual learning. In V. Walsh & J. Kulikowski (Eds.), *Perceptual constancy. Why things look as they do* (pp. 455-498). Cambridge, MA: Cambridge University Press.

Ahissar, M., & Hochstein, S. (2004). The reverse hierarchy theory of visual perceptual learning. *Trends Cogn Sci, 8*(10), 457-464.

Alink, A., Schwiedrzik, C. M., Kohler, A., Singer, W., & Muckli, L. (2010). Stimulus predictability reduces responses in primary visual cortex. *J Neurosci, 30*(8), 2960-2966.

Alpern, M. (1953). Metacontrast. *J Opt Soc Am, 43*(8), 648-657.

Anderson, B. L., O'Vari, J., & Barth, H. (2011). Non-Bayesian contour synthesis. *Curr Biol, 21*(6), 492-496.

Anstis, S., Verstraten, F. A., & Mather, G. (1998). The motion aftereffect. *Trends Cogn Sci, 2*(3), 111-117.

Archambault, A., O'Donnell, C., & Schyns, P. G. (1999). Blind to object change: when learning the same object at different levels of categorization modifies its perception. *Psychol Sci, 10*(3), 249-255.

Azzopardi, P., & Cowey, A. (1998). Blindsight and visual awareness. *Conscious Cogn, 7*(3), 292-311.

Baars, B. J. (1997). *In the theater of consciousness: The workspace of the mind*. Oxford: Oxford University Press.

Bachmann, T. (2007). Binding binding: departure points for a different version of the perceptual retouch theory. *Adv Cogn Psychol, 3*(1-2), 41-55.

Baizer, J. S., Ungerleider, L. G., & Desimone, R. (1991). Organization of visual inputs to the inferior temporal and posterior parietal cortex in macaques. *J Neurosci, 11*(1), 168-190.

Banquet, J. P., & Grossberg, S. (1987). Probing cognitive processes through the structure of event-related potentials during learning: an experimental and theoretical analysis. *Appl Opt, 26*(23), 4931-4946.

Bar, M., & Biederman, I. (1999). Localizing the cortical region mediating visual awareness of object identity. *Proc Natl Acad Sci U S A, 96*(4), 1790-1793.

Bar, M., Tootell, R. B., Schacter, D. L., Greve, D. N., Fischl, B., Mendola, J. D., et al. (2001). Cortical mechanisms specific to explicit visual object recognition. *Neuron, 29*(2), 529-535.

Barbas, H. (1988). Anatomic organization of basoventral and mediodorsal visual recipient prefrontal regions in the rhesus monkey. *J Comp Neurol, 276*(3), 313-342.

Barbas, H., & Mesulam, M. M. (1985). Cortical afferent input to the principalis region of the rhesus monkey. *Neuroscience, 15*(3), 619-637.

Barlow, H. B. (1990). A theory about the functional role and synaptic mechanism of visual aftereffects. In C. Blakemore (Ed.), *Vision: coding and efficiency* (pp. 363-375). Cambridge, UK: Cambridge University Press.

Barlow, H. B., & Földiák, P. (1989). Adaptation and decorrelation in the cortex. In R. Durbin, C. Miall & G. Mitchison (Eds.), *The computing neuron* (pp. 54-72). Wokingham, UK: Addison-Wesley.

Bayes, T. (1763). An essay towards solving a problem in the doctrine of chances. *Philos Trans R Soc, 53*, 269-271.

Beauchamp, M. S. (2003). Detection of eye movements from fMRI data. *Magn Reson Med, 49*(2), 376-380.

Berkes, P., Orban, G., Lengyel, M., & Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science, 331*(6013), 83-87.

Biederman, I. (1972). Perceiving real-world scenes. *Science, 177*(43), 77-80.

Blake, R., & Fox, R. (1974). Adaptation to invisible gratings and the site of binocular rivalry suppression. *Nature, 249*(456), 488-490.

Blake, R., Fox, R., & McIntyre, C. (1971). Stochastic properties of stabilized-image binocular rivalry alternations. *J Exp Psychol, 88*(3), 327-332.

Blake, R., Sobel, K. V., & Gilroy, L. A. (2003). Visual motion retards alternations between conflicting perceptual interpretations. *Neuron, 39*(5), 869-878.

Block, N. (2011). The higher order approach to consciousness is defunct. *Analysis, Epub ahead of print*.

Boehler, C. N., Schoenfeld, M. A., Heinze, H. J., & Hopf, J. M. (2008). Rapid recurrent processing gates awareness in primary visual cortex. *Proc Natl Acad Sci U S A, 105*(25), 8742-8747.

Borsellino, A., De Marco, A., Allazetta, A., Rinesi, S., & Bartolini, B. (1972). Reversal time distribution in the perception of visual ambiguous stimuli. *Kybernetik, 10*(3), 139-144.

Boussaoud, D., Desimone, R., & Ungerleider, L. G. (1991). Visual topography of area TEO in the macaque. *J Comp Neurol, 306*(4), 554-575.

Boyer, J., & Ro, T. (2007). Attention attenuates metacontrast masking. *Cognition, 104*(1), 135-149.

Braff, D., Saccuzzo, D., Ingram, R., McNeill, B., & Langford, R. (1980). Effect of practice on visual backward masking. *Percept Mot Skills, 50*(2), 511-517.

Brascamp, J. W., Pearson, J., Blake, R., & van den Berg, A. V. (2009). Intermittent ambiguous stimuli: implicit memory causes periodic perceptual alternations. *J Vis, 9*(3), 3 1-23.

Breitmeyer, B., & Ogmen, H. (2006). *Visual masking: Time slices through conscious and unconscious vision* (2nd ed.). Oxford: Oxford University Press.

Breitmeyer, B., Rudd, M., & Dunn, K. (1981). Metacontrast investigations of sustained-transient channel inhibitory interactions. *J Exp Psychol Hum Percept Perform, 7*(4), 770-779.

Bridgeman, B., & Staggs, D. (1982). Plasticity in human blindsight. *Vision Res, 22*(9), 1199-1203.

Brown, S., & Heathcote, A. (2005). Practice increases the efficiency of evidence accumulation in perceptual choice. *J Exp Psychol Hum Percept Perform, 31*(2), 289-298.

Bruner, J. S., & Potter, M. C. (1964). Interference in visual recognition. *Science, 144*, 424-425.

Brunswik, E., & Kamiya, J. (1953). Ecological cue-validity of proximity and of other Gestalt factors. *Am J Psychol, 66*(1), 20-32.

Carpenter, R. H., & Williams, M. L. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature, 377*(6544), 59-62.

Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nat Neurosci, 7*(3), 308-313.

Chaumon, M., Drouet, V., & Tallon-Baudry, C. (2008). Unconscious associative memory affects visual processing before 100 ms. *J Vis, 8*(3), 10 11-10.

Chen, X., & He, S. (2003). Temporal characteristics of binocular rivalry: visual field asymmetries. *Vision Res, 43*(21), 2207-2212.

Chen, X., & He, S. (2004). Local factors determine the stabilization of monocular ambiguous and binocular rivalry stimuli. *Curr Biol, 14*(11), 1013-1017.

Chikkerur, S., Serre, T., Tan, C., & Poggio, T. (2010). What and where: a Bayesian inference theory of attention. *Vision Res, 50*(22), 2233-2247.

Chokron, S., Perez, C., Obadia, M., Gaudry, I., Laloum, L., & Gout, O. (2008). From blindsight to sight: Cognitive rehabilitation of visual field defects. *Restor Neurol Neurosci, 26*(4-5), 305-320.

Chung, S. T. (2007). Learning to identify crowded letters: Does it improve reading speed? *Vision Res, 47*(25), 3150-3159.

Cleeremans, A. (2008). Consciousness: The radical plasticity thesis. *Prog Brain Res, 168*, 19-33.

Clifford, C. W. (2010). Consciousness: Reading the neural signature. *Curr Biol, 20*(2), R61-62.

Clifford, C. W., Wenderoth, P., & Spehar, B. (2000). A functional angle on some after-effects in cortical vision. *Proc Biol Sci, 267*(1454), 1705-1710.

Colombo, M., & Seriès, P. (2011). Bayes in the brain. On Bayesian modelling in neuroscience. *Br J Philos Sci, Epub ahead of print*.

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci, 3*(3), 201-215.

Cornette, L., Dupont, P., & Orban, G. A. (2002). The neural substrate of orientation short-term memory and resistance to distractor items. *Eur J Neurosci, 15*(1), 165-175.

Cornette, L., Dupont, P., Salmon, E., & Orban, G. A. (2001). The neural substrate of orientation working memory. *J Cogn Neurosci, 13*(6), 813-828.

Corso, J. F. (1963). A theoretico-historical review of the threshold concept. *Psychol Bull, 60*, 356-370.

Cousineau, D. (2005). Confidence intervals in within-subjects designs: a simpler solution to Loftus and Masson's method. *Tutor Quant Methods Psychol, 1*(1), 42-45.

Cowey, A. (2004). The 30th Sir Frederick Bartlett lecture. Fact, artefact, and myth about blindsight. *Q J Exp Psychol A, 57*(4), 577-609.

Coyne, A. C. (1981). Age differences and practice in forward visual masking. *J Gerontol, 36*(6), 730-732.

Crick, F., & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Semin Neurosci, 2*, 263-275.

Crist, R. E., Li, W., & Gilbert, C. D. (2001). Learning to see: Experience and attention in primary visual cortex. *Nat Neurosci, 4*(5), 519-525.

Dambacher, M., Rolfs, M., Gollner, K., Kliegl, R., & Jacobs, A. M. (2009). Event-related potentials reveal rapid verification of predicted visual input. *PLoS One, 4*(3), e5047.

Dayan, P., Hinton, G. E., Neal, R. M., & Zemel, R. S. (1995). The Helmholtz machine. *Neural Comput, 7*(5), 889-904.

Dehaene, S. (2008). Conscious and nonconscious processes: distinct forms of evidence accumulation? In C. Engel & W. Singer (Eds.), *Better than conscious? Decision making, the human mind, and implications for institutions* (pp. 21-49). Cambridge, MA: MIT Press.

Dehaene, S., & Changeux, J. P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron, 70*(2), 200-227.

Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends Cogn Sci, 10*(5), 204-211.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition, 79*(1-2), 1-37.

Del Cul, A., Baillet, S., & Dehaene, S. (2007). Brain dynamics underlying the nonlinear threshold for access to consciousness. *PLoS Biol, 5*(10), e260.

Del Cul, A., Dehaene, S., Reyes, P., Bravo, E., & Slachevsky, A. (2009). Causal role of prefrontal cortex in the threshold for access to consciousness. *Brain, 132*(Pt 9), 2531-2540.

Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J Neurosci Methods, 134*(1), 9-21.

Dennett, D. C., & Kinsbourne, M. (1992). Time and the observer: The where and when of consciousness in the brain. *Behav Brain Sci, 15*(2), 183-247.

Desimone, R., & Gross, C. G. (1979). Visual areas in the temporal cortex of the macaque. *Brain Res, 178*(2-3), 363-380.

Destrebecqz, A., & Cleeremans, A. (2001). Can sequence learning be implicit? New evidence with the process dissociation procedure. *Psychon Bull Rev, 8*(2), 343-350.

Di Lollo, V., Enns, J. T., & Rensink, R. A. (2000). Competition for consciousness among visual events: the psychophysics of reentrant visual processes. *J Exp Psychol Gen, 129*(4), 481-507.

Di Russo, F., Aprile, T., Spitoni, G., & Spinelli, D. (2008). Impaired visual processing of contralesional stimuli in neglect patients: a visual-evoked potential study. *Brain, 131*(Pt 3), 842-854.

Dienes, Z. (2004). Assumptions of subjective measures of unconscious mental states: higher order thoughts and bias. *J Conscious Stud, 11*(9), 25-45.

Dienes, Z., & Seth, A. K. (2010). Measuring any conscious content versus measuring the relevant conscious content: Comment on Sandberg et al. *Conscious Cogn, 19*(4), 1079-1080.

Dill, M. (2002). Specificity versus invariance of perceptual learning: The example of position. In M. Fahle & T. Poggio (Eds.), *Perceptual learning* (pp. 219-231). Cambridge, MA: MIT Press.

Dineen, J., & Keating, E. G. (1981). The primate visual system after bilateral removal of striate cortex. Survival of complex pattern vision. *Exp Brain Res, 41*(3-4), 338-345.

Distler, C., Boussaoud, D., Desimone, R., & Ungerleider, L. G. (1993). Cortical connections of inferior temporal area TEO in macaque monkeys. *J Comp Neurol, 334*(1), 125-150.

Dolan, R. J., Fink, G. R., Rolls, E., Booth, M., Holmes, A., Frackowiak, R. S., et al. (1997). How the brain learns to see objects and faces in an impoverished context. *Nature, 389*(6651), 596-599.

Dosher, B. A., & Lu, Z. L. (1999). Mechanisms of perceptual learning. *Vision Res, 39*(19), 3197-3221.

Eckert, M. A., Menon, V., Walczak, A., Ahlstrom, J., Denslow, S., Horwitz, A., et al. (2009). At the heart of the ventral attention system: the right anterior insula. *Hum Brain Mapp, 30*(8), 2530-2541.

Edelman, G. M., Gally, J. A., & Baars, B. J. (2011). Biology of consciousness. *Front Psychology, 2:4.*

Eger, E., Henson, R. N., Driver, J., & Dolan, R. J. (2007). Mechanisms of top-down facilitation in perception of visual objects studied by FMRI. *Cereb Cortex, 17*(9), 2123-2133.

Engel, A. K., Fries, P., & Singer, W. (2001). Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci, 2*(10), 704-716.

Eriksen, C. W. (1960). Discrimination and learning without awareness: A methodological survey and evaluation. *Psychol Rev, 67*, 279-300.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415*(6870), 429-433.

Esterman, M., & Yantis, S. (2010). Perceptual expectation evokes category-selective cortical activity. *Cereb Cortex, 20*(5), 1245-1253.

Fahle, M. (2009). Perceptual learning. In T. Bayne, A. Cleeremans & P. Wilken (Eds.), *The Oxford companion to consciousness* (pp. 407-409). Oxford: Oxford University Press.

Fahle, M., & Edelman, S. (1993). Long-term learning in vernier acuity: Effects of stimulus orientation, range and of feedback. *Vision Res, 33*(3), 397-412.

Faillenot, I., Decety, J., & Jeannerod, M. (1999). Human brain activity related to the perception of spatial features of objects. *Neuroimage, 10*(2), 114-124.

Faillenot, I., Sunaert, S., Van Hecke, P., & Orban, G. A. (2001). Orientation discrimination of objects and gratings compared: an fMRI study. *Eur J Neurosci, 13*(3), 585-596.

Farroni, T., Valenza, E., Simion, F., & Umiltà, C. (2000). Configural processing at birth: evidence for perceptual organisation. *Perception, 29*(3), 355-372.

Fecteau, J. H., & Munoz, D. P. (2003). Exploring the consequences of the previous trial. *Nat Rev Neurosci, 4*(6), 435-443.

Felleman, D. J., & Van Essen, D. C. (1987). Receptive field properties of neurons in area V3 of macaque monkey extrastriate cortex. *J Neurophysiol, 57*(4), 889-920.

Felsen, G., Shen, Y. S., Yao, H., Spor, G., Li, C., & Dan, Y. (2002). Dynamic modification of cortical orientation tuning mediated by recurrent connections. *Neuron, 36*(5), 945-954.

Fender, D., & Julesz, B. (1967). Extension of Panum's fusional area in binocularly stabilized vision. *J Opt Soc Am, 57*(6), 819-830.

Fiser, J., Berkes, P., Orban, G., & Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn Sci, 14*(3), 119-130.

Fleming, R. W. (2011). Visual perception: bizarre contours go against the odds. *Curr Biol, 21*(7), R259-261.

Fox, M. D., Corbetta, M., Snyder, A. Z., Vincent, J. L., & Raichle, M. E. (2006). Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proc Natl Acad Sci U S A, 103*(26), 10046-10051.

Fox, R., & Herrmann, J. (1967). Stochastic properties of binocular rivalry alternations. *Percept Psychophys, 2*(9), 432-436.

Freunberger, R., Klimesch, W., Doppelmayr, M., & Holler, Y. (2007). Visual P2 component is related to theta phase-locking. *Neurosci Lett, 426*(3), 181-186.

Friston, K. (2005). A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci, 360*(1456), 815-836.

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nat Rev Neurosci, 11*(2), 127-138.

Gaillard, R., Dehaene, S., Adam, C., Clemenceau, S., Hasboun, D., Baulac, M., et al. (2009). Converging intracranial markers of conscious access. *PLoS Biol, 7*(3), e61.

Galvin, S. J., Podd, J. V., Drga, V., & Whitmore, J. (2003). Type 2 tasks in the theory of signal detectability: Discrimination between correct and incorrect decisions. *Psychon Bull Rev, 10*(4), 843-876.

Garrido, M. I., Kilner, J. M., Stephan, K. E., & Friston, K. J. (2009). The mismatch negativity: a review of underlying mechanisms. *Clin Neurophysiol, 120*(3), 453-463.

Gattass, R., Sousa, A. P., & Gross, C. G. (1988). Visuotopic organization and extent of V3 and V4 of the macaque. *J Neurosci, 8*(6), 1831-1845.

Geisler, W. S. (2008). Visual perception and the statistical properties of natural scenes. *Annu Rev Psychol, 59*, 167-192.

Geisler, W. S., & Diehl, R. L. (2002). Bayesian natural selection and the evolution of perceptual systems. *Philos Trans R Soc Lond B Biol Sci, 357*(1420), 419-448.

Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res, 41*(6), 711-724.

Gepshtein, S., & Kubovy, M. (2005). Stability and change in perception: spatial organization in temporal context. *Exp Brain Res, 160*(4), 487-495.

Gibson, E. J. (1963). Perceptual learning. *Annu Rev Psychol, 14*, 29-56.

Gibson, J. J., & Radner, M. (1937). Adaptation, after-effect, and contrast in the perception of tilted lines. I. Quantitative studies. *J Exp Psychol, 20*(5), 453-467.

Gigante, G., Mattia, M., Braun, J., & Del Giudice, P. (2009). Bistable perception modeled as competing stochastic integrations at two levels. *PLoS Comput Biol, 5*(7), e1000430.

Gilbert, C. D., & Sigman, M. (2007). Brain states: top-down influences in sensory processing. *Neuron, 54*(5), 677-696.

Gold, J., Bennett, P. J., & Sekuler, A. B. (1999). Signal but not noise changes with perceptual learning. *Nature, 402*(6758), 176-178.

Goldstone, R. L. (1998). Perceptual learning. *Annu Rev Psychol, 49*, 585-612.

Green, C. S., & Bavelier, D. (2003). Action video game modifies visual selective attention. *Nature, 423*(6939), 534-537.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York, NY: John Wiley and Sons, Inc.

Gregory, R. L. (1997). Knowledge in perception and illusion. *Philos Trans R Soc Lond B Biol Sci, 352*(1358), 1121-1127.

Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn Sci, 10*(1), 14-23.

Gutnisky, D. A., & Dragoi, V. (2008). Adaptive coding of visual information in neural populations. *Nature, 452*(7184), 220-224.

Handwerker, D. A., Ollinger, J. M., & D'Esposito, M. (2004). Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *Neuroimage, 21*(4), 1639-1651.

Hanley, J. A., Negassa, A., Edwardes, M. D., & Forrester, J. E. (2003). Statistical analysis of correlated data using generalized estimating equations: an orientation. *Am J Epidemiol, 157*(4), 364-375.

Hansen, B. J., & Dragoi, V. (2011). Adaptation-induced synchronization in laminar cortical circuits. *Proc Natl Acad Sci U S A, 108*(26), 10720-10725.

Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values of d'. *Behav Res Meth Instrum Comput, 27*(1), 46-51.

Haynes, J. D., & Rees, G. (2005). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci, 8*(5), 686-691.

Helmholtz, H. (1867). *Handbuch der physiologischen Optik*. Leipzig: Leopold Voss.

Henriksson, L., Raninen, A., Näsänen, R., Hyvärinen, L., & Vanni, S. (2007). Training-induced cortical representation of a hemianopic hemifield. *J Neurol Neurosurg Psychiatry, 78*(1), 74-81.

Hernandez, L. L., & Lefton, L. A. (1977). Metacontrast as measured under a signal detection model. *Perception, 6*(6), 696-702.

Hertzog, C. K., Williams, M. V., & Walsh, D. A. (1976). The effect of practice on age differences in central perceptual processing. *J Gerontol, 31*(4), 428-433.

Herzog, M. H., & Fahle, M. (1997). The role of feedback in learning a vernier discrimination task. *Vision Res, 37*(15), 2133-2141.

Hinrichs, H., Noesselt, T., & Heinze, H. J. (2008). Directed information flow: a model free measure to analyze causal interactions in event related EEG-MEG-experiments. *Hum Brain Mapp, 29*(2), 193-206.

Hochstein, S., & Ahissar, M. (2002). View from the top: hierarchies and reverse hierarchies in the visual system. *Neuron, 36*(5), 791-804.

Hock, H. S., Schöner, G., & Hochstein, S. (1996). Perceptual stability and the selective adaptation of perceived and unperceived motion directions. *Vision Res, 36*(20), 3311-3323.

Hoffman, K. L., & Logothetis, N. K. (2009). Cortical mechanisms of sensory learning and object recognition. *Philos Trans R Soc Lond B Biol Sci, 364*(1515), 321-329.

Hogben, J. H., & Di Lollo, V. (1984). Practice reduces suppression in metacontrast and in apparent motion. *Percept Psychophys, 35*(5), 441-445.

Hohwy, J., Roepstorff, A., & Friston, K. (2008). Predictive coding explains binocular rivalry: an epistemological review. *Cognition, 108*(3), 687-701.

Holender, D. (1986). Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal. *Behav Brain Sci, 9*, 1-66.

Hopf, J. M., Vogel, E., Woodman, G., Heinze, H. J., & Luck, S. J. (2002). Localizing visual discrimination processes in time and space. *J Neurophysiol, 88*(4), 2088-2095.

Howell, D. C. (2002). *Statistical methods for psychology* (5th ed.). Pacific Grove, CA: Duxbury/Thomson Learning.

Huckauf, A., & Nazir, T. A. (2007). How odgcrnwi becomes crowding: Stimulus-specific learning reduces crowding. *J Vis, 7*(2), 18 11-12.

Humphrey, N. K. (1974). Vision in a monkey without striate cortex: a case study. *Perception, 3*(3), 241-255.

Jacobs, R. A. (1999). Optimal integration of texture and motion cues to depth. *Vision Res, 39*(21), 3621-3629.

James, T. W., & Gauthier, I. (2006). Repetition-induced changes in BOLD response reflect accumulation of neural activity. *Hum Brain Mapp, 27*(1), 37-46.

James, T. W., Humphrey, G. K., Gati, J. S., Menon, R. S., & Goodale, M. A. (2000). The effects of visual object priming on brain activation before and after recognition. *Curr Biol, 10*(17), 1017-1024.

Kanai, R., & Verstraten, F. A. (2005). Perceptual manifestations of fast neural plasticity: motion priming, rapid motion aftereffect and perceptual sensitization. *Vision Res, 45*(25-26), 3109-3116.

Kanai, R., & Verstraten, F. A. (2006). Attentional modulation of perceptual stabilization. *Proc Biol Sci, 273*(1591), 1217-1222.

Kanwisher, N. (2001). Neural events and perceptual awareness. *Cognition, 79*(1-2), 89-113.

Karni, A., & Sagi, D. (1991). Where practice makes perfect in texture discrimination: Evidence for primary visual cortex plasticity. *Proc Natl Acad Sci U S A, 88*(11), 4966-4970.

Kastner, S., De Weerd, P., Pinsk, M. A., Elizondo, M. I., Desimone, R., & Ungerleider, L. G. (2001). Modulation of sensory suppression: Implications for receptive field sizes in the human visual cortex. *J Neurophysiol, 86*(3), 1398-1411.

Kastner, S., & Ungerleider, L. G. (2000). Mechanisms of visual attention in the human cortex. *Annu Rev Neurosci, 23*, 315-341.

Kentridge, R. W., Nijboer, T. C., & Heywood, C. A. (2008). Attended but unseen: visual attention is not sufficient for visual awareness. *Neuropsychologia, 46*(3), 864-869.

Kim, C. Y., & Blake, R. (2005). Psychophysical magic: rendering the visible 'invisible'. *Trends Cogn Sci, 9*(8), 381-388.

Kleinschmidt, A., Büchel, C., Hutton, C., Friston, K. J., & Frackowiak, R. S. (2002). The neural structures expressing perceptual hysteresis in visual letter recognition. *Neuron, 34*(4), 659-666.

Klotz, W., & Neumann, O. (1999). Motor activation without conscious discrimination in metacontrast masking. *J Exp Psychol Hum Percept Perform, 25*(4), 976-992.

Knapen, T., Brascamp, J., Adams, W. J., & Graf, E. W. (2009). The spatial scale of perceptual memory in ambiguous figure perception. *J Vis, 9*(13), 16 11-12.

Koch, C. (2004). *The quest for consciousness: a neurobiological approach*. Denver, CO: Roberts and Co.

Koivisto, M., & Revonsuo, A. (2010). Event-related brain potential correlates of visual awareness. *Neurosci Biobehav Rev, 34*(6), 922-934.

Kondo, H., & Komatsu, H. (2000). Suppression on neuronal responses by a metacontrast masking stimulus in monkey V4. *Neurosci Res, 36*(1), 27-33.

Kotsoni, E., Csibra, G., Mareschal, D., & Johnson, M. H. (2007). Electrophysiological correlates of common-onset visual masking. *Neuropsychologia, 45*(10), 2285-2293.

Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends Cogn Sci, 14*(7), 301-307.

Kouider, S., & Dehaene, S. (2007). Levels of processing during non-conscious perception: A critical review of visual masking. *Philos Trans R Soc Lond B Biol Sci, 362*(1481), 857-875.

Kourtzi, Z., Betts, L. R., Sarkheil, P., & Welchman, A. E. (2005). Distributed neural plasticity for shape learning in the human visual cortex. *PLoS Biol, 3*(7), e204.

Kriegeskorte, N., & Goebel, R. (2001). An efficient algorithm for topologically correct segmentation of the cortical sheet in anatomical MR volumes. *Neuroimage, 14*(2), 329-346.

Kubovy, M., Holcombe, A. O., & Wagemans, J. (1998). On the lawfulness of grouping by proximity. *Cogn Psychol, 35*(1), 71-98.

Kubovy, M., & Wagemans, J. (1995). Grouping by proximity and multistability in dot lattices: a quantitative Gestalt theory. *Psychol Sci, 6*(4), 225-234.

Kveraga, K., Ghuman, A. S., & Bar, M. (2007). Top-down predictions in the cognitive brain. *Brain Cogn, 65*(2), 145-168.

Lachter, J., & Durgin, F. (1999). Metacontrast masking functions: A question of speed? *J Exp Psychol Hum Percept Perform, 25*(4), 936-947.

Lachter, J., Durgin, F., & Washington, T. (2000). Disappearing percepts: Evidence for retention failure in metacontrast masking. *Vis Cognit, 7*(1/2/3), 269-279.

Lamme, V. A. (2006). Towards a true neural stance on consciousness. *Trends Cogn Sci, 10*(11), 494-501.

Lamme, V. A., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci, 23*(11), 571-579.

Lau, H. C. (2008a). Are we studying consciousness yet? In L. Weiskrantz & M. Davies (Eds.), *Frontiers of consciousness: The Chichele lectures* (pp. 245-258). Oxford: Oxford University Press.

Lau, H. C. (2008b). A higher order Bayesian decision theory of consciousness. *Prog Brain Res, 168*, 35-48.

Lau, H. C. (2011). Theoretical motivations for investigating the neural correlates of consciousness. *WIREs Cogn Sci, 2*(1), 1-7.

Lau, H. C., & Passingham, R. E. (2006). Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proc Natl Acad Sci U S A, 103*(49), 18763-18768.

Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *J Opt Soc Am A Opt Image Sci Vis, 20*(7), 1434-1448.

Lehmann, D., & Skrandies, W. (1980). Reference-free identification of components of checkerboard-evoked multichannel potential fields. *Electroencephalogr Clin Neurophysiol, 48*(6), 609-621.

Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nat Neurosci, 4*(1), 89-94.

Levitt, J. B., Kiper, D. C., & Movshon, J. A. (1994). Receptive fields and functional architecture of macaque V2. *J Neurophysiol, 71*(6), 2517-2542.

Liang, K. Y., & Zeger, S. L. (1986). Longitudinal data analysis using generalized linear models. *Biometrika, 73*(1), 13-22.

Long, G. M., Toppino, T. C., & Mondin, G. W. (1992). Prime time: fatigue and set effects in the perception of reversible figures. *Percept Psychophys, 52*(6), 609-616.

Luck, S. J., & Hillyard, S. A. (1995). The role of attention in feature detection and conjunction discrimination: an electrophysiological analysis. *Int J Neurosci, 80*(1-4), 281-297.

Lumer, E. D., Friston, K. J., & Rees, G. (1998). Neural correlates of perceptual rivalry in the human brain. *Science, 280*(5371), 1930-1934.

Lycan, W. G. (2009). Higher-order representation theories of consciousness. In T. Bayne, A. Cleeremans & P. Wilken (Eds.), *The Oxford companion to consciousness* (pp. 346-350). Oxford, UK: Oxford University Press.

Macknik, S. L., & Livingstone, M. S. (1998). Neuronal correlates of visibility and invisibility in the primate visual system. *Nat Neurosci, 1*(2), 144-149.

Macmillan, N. A., & Creelman, C. D. (1990). Response bias: Characteristics of detection theory, threshold theory, and "nonparametric" indexes. *Psychol Bull, 107*(3), 401-413.

Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide* (2nd ed.). Mahwah, NJ: Lawrence Erlbaum Associates.

Maehara, G., & Goryo, K. (2003). Perceptual learning in visual backward pattern masking. *Percept Mot Skills, 97*(3 Pt 2), 1137-1149.

Maier, A., Wilke, M., Logothetis, N. K., & Leopold, D. A. (2003). Perception of temporally interleaved ambiguous patterns. *Curr Biol, 13*(13), 1076-1085.

Mangun, G. R. (1995). Neural mechanisms of visual selective attention. *Psychophysiology, 32*(1), 4-18.

Manly, B. F. (1991). *Randomization and Monte Carlo methods in biology*. London, UK: Chapman & Hall.

McCall, R. B., & Appelbaum, M. I. (1973). Bias in the analysis of repeated-measures designs: Some alternative approaches. *Child Dev, 44*(3), 401-415.

Meeres, S. L., & Graves, R. E. (1990). Localization of unseen visual stimuli by humans with normal vision. *Neuropsychologia, 28*(12), 1231-1237.

Melloni, L., & Singer, W. (2010). Distinct characteristics of conscious experience are met by large scale neuronal synchronization. In E. Perry, D. Collerton, F. E. N. LeBeau & H. Ashton (Eds.), *New horizons in the neuroscience of consciousness* (pp. 17-28). Amsterdam, NL: John Benjamins.

Merikle, P. M., Smilek, D., & Eastwood, J. D. (2001). Perception without awareness: Perspectives from cognitive psychology. *Cognition, 79*(1-2), 115-134.

Metz, C. E., & Pan, X. (1999). "Proper" binormal ROC curves: Theory and maximum-likelihood estimation. *J Math Psychol, 43*(1), 1-33.

Morey, R. D. (2008). Confidence intervals from normalized data: a correction to Cousineau (2005). *Tutor Quant Methods Psychol, 4*(2), 61-64.

Muckli, L., Kriegeskorte, N., Lanfermann, H., Zanella, F. E., Singer, W., & Goebel, R. (2002). Apparent motion: event-related functional magnetic resonance imaging of perceptual switches and states. *J Neurosci, 22*(9), RC219.

Müller, T. J., Federspiel, A., Horn, H., Lovblad, K., Lehmann, C., Dierks, T., et al. (2005). The neurophysiological time pattern of illusionary visual perceptual transitions: a simultaneous EEG and fMRI study. *Int J Psychophysiol, 55*(3), 299-312.

Mumford, D. (1992). On the computational architecture of the neocortex. II. The role of cortico-cortical loops. *Biol Cybern, 66*(3), 241-251.

Murray, M. M., Brunet, D., & Michel, C. M. (2008). Topographic ERP analyses: a step-by-step tutorial review. *Brain Topogr, 20*(4), 249-264.

Murray, S. O., Kersten, D., Olshausen, B. A., Schrater, P., & Woods, D. L. (2002). Shape perception reduces activity in human primary visual cortex. *Proc Natl Acad Sci U S A, 99*(23), 15164-15169.

Neumann, O., & Klotz, W. (1994). Motor responses to non-reportable, masked stimuli: Where is the limit of direct parameter specification? In C. Umiltà & M. Moskovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing* (pp. 123-150). Cambridge, MA: MIT Press.

Nikolaev, A. R., Gepshtein, S., Gong, P., & van Leeuwen, C. (2010). Duration of coherence intervals in electrical brain activity in perceptual organization. *Cereb Cortex, 20*(2), 365-382.

Noest, A. J., van Ee, R., Nijs, M. M., & van Wezel, R. J. (2007). Percept-choice sequences driven by interrupted ambiguous stimuli: a low-level neural model. *J Vis, 7*(8), 10.

Noguchi, Y., Tanabe, H. C., Sadato, N., Hoshiyama, M., & Kakigi, R. (2007). Voluntary attention changes the speed of perceptual neural processing. *Eur J Neurosci, 25*(10), 3163-3172.

Oldfield, R. C. (1971). The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia, 9*(1), 97-113.

Olesen, P. J., Westerberg, H., & Klingberg, T. (2004). Increased prefrontal and parietal activity after training of working memory. *Nat Neurosci, 7*(1), 75-79.

Ollinger, J. M., Corbetta, M., & Shulman, G. L. (2001). Separating processes within a trial in event-related functional MRI. II. Analysis. *Neuroimage, 13*(1), 218-229.

Ollinger, J. M., Shulman, G. L., & Corbetta, M. (2001). Separating processes within a trial in event-related functional MRI. I. The method. *Neuroimage, 13*(1), 210-217.

Olson, I. R., Chun, M. M., & Allison, T. (2001). Contextual guidance of attention: human intracranial event-related potential evidence for feedback modulation in anatomically early temporally late stages of visual processing. *Brain, 124*(Pt 7), 1417-1425.

Op de Beeck, H. P., Wagemans, J., & Vogels, R. (2007). Effects of perceptual learning in visual backward masking on the responses of macaque inferior temporal neurons. *Neuroscience, 145*(2), 775-789.

Orbach, J., Ehrlich, D., & Heath, H. A. (1963). Reversibility of the Necker Cube. I. An examination of the concept of "satiation of orientation". *Percept Mot Skills, 17*, 439-458.

Orban, G. A., Dupont, P., Vogels, R., Bormans, G., & Mortelmans, L. (1997). Human brain activity related to orientation discrimination tasks. *Eur J Neurosci, 9*(2), 246-259.

Otto, T. U., Ogmen, H., & Herzog, M. H. (2006). The flight path of the phoenix--the visible trace of invisible elements in human vision. *J Vis, 6*(10), 7 1079-1086.

Overgaard, M., Rote, J., Mouridsen, K., & Ramsøy, T. Z. (2006). Is conscious perception gradual or dichotomous? A comparison of report methodologies during a visual task. *Conscious Cogn, 15*(4), 700-708.

Pasquali, A., Timmermans, B., & Cleeremans, A. (2010). Know thyself: metacognitive networks and measures of consciousness. *Cognition, 117*(2), 182-190.

Pastukhov, A., & Braun, J. (2008). A short-term memory of multi-stable perception. *J Vis, 8*(13), 7 1-14.

Pasupathy, A. (2006). Neural basis of shape representation in the primate brain. *Prog Brain Res, 154*, 293-313.

Pearson, J., & Brascamp, J. (2008). Sensory memory for ambiguous vision. *Trends Cogn Sci, 12*(9), 334-341.

Persaud, N., McLeod, P., & Cowey, A. (2007). Post-decision wagering objectively measures awareness. *Nat Neurosci, 10*(2), 257-261.

Pesce, L. L., & Metz, C. E. (2007). Reliable and computationally efficient maximum-likelihood estimation of "proper" binormal ROC curves. *Acad Radiol, 14*(7), 814-829.

Peterhans, E., Heider, B., & Baumann, R. (2005). Neurons in monkey visual cortex detect lines defined by coherent motion of dots. *Eur J Neurosci, 21*(4), 1091-1100.

Peterhans, E., & von der Heydt, R. (1991). Subjective contours - bridging the gap between psychophysics and physiology. *Trends Neurosci, 14*(3), 112-119.

Petit, L., Courtney, S. M., Ungerleider, L. G., & Haxby, J. V. (1998). Sustained activity in the medial wall during working memory delays. *J Neurosci, 18*(22), 9429-9437.

Petrov, A. A., Dosher, B. A., & Lu, Z. L. (2006). Perceptual learning without feedback in non-stationary contexts: Data and model. *Vision Res, 46*(19), 3177-3197.

Petry, S. (1978). Perceptual changes during metacontrast. *Vision Res, 18*(10), 1337-1341.

Picton, T. W., Bentin, S., Berg, P., Donchin, E., Hillyard, S. A., Johnson, R., Jr., et al. (2000). Guidelines for using human event-related potentials to study cognition: recording standards and publication criteria. *Psychophysiology, 37*(2), 127-152.

Pins, D., & Ffytche, D. (2003). The neural correlates of conscious vision. *Cereb Cortex, 13*(5), 461-474.

Poggio, T., Fahle, M., & Edelman, S. (1992). Fast perceptual learning in visual hyperacuity. *Science, 256*(5059), 1018-1021.

Pourtois, G., Delplanque, S., Michel, C., & Vuilleumier, P. (2008). Beyond conventional event-related brain potential (ERP): exploring the time-course of visual emotion processing using topographic and principal component analyses. *Brain Topogr, 20*(4), 265-277.

Purkinje, J. E. (1820). Beiträge zur näheren Kenntnis des Schwindels aus heautognostischen Daten. *Med Jb Kaiser Königl Österr Staates Wien, 6*, 79-125.

Raemaekers, M., van der Schaaf, M. E., van Ee, R., & van Wezel, R. J. (2009). Widespread fMRI activity differences between perceptual states in visual rivalry are correlated with differences in observer biases. *Brain Res, 1252*, 161-171.

Raiguel, S., Vogels, R., Mysore, S. G., & Orban, G. A. (2006). Learning to see the difference specifically alters the most informative V4 neurons. *J Neurosci, 26*(24), 6589-6602.

Rainer, G., Asaad, W. F., & Miller, E. K. (1998). Memory fields of neurons in the primate prefrontal cortex. *Proc Natl Acad Sci U S A, 95*(25), 15008-15013.

Rainer, G., Lee, H., & Logothetis, N. K. (2004). The effect of learning on the function of monkey extrastriate visual cortex. *PLoS Biol, 2*(2), E44.

Rainer, G., & Miller, E. K. (2000). Effects of visual experience on the representation of objects in the prefrontal cortex. *Neuron, 27*(1), 179-189.

Ramsøy, T. Z., & Overgaard, M. (2004). Introspection and subliminal perception. *Phenomenology and the Cognitive Sciences, 3*, 1-23.

Raninen, A., Vanni, S., Hyvärinen, L., & Näsänen, R. (2007). Temporal sensitivity in a hemianopic visual field can be improved by long-term training using flicker stimulation. *J Neurol Neurosurg Psychiatry, 78*(1), 66-73.

Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci, 2*(1), 79-87.

Rensink, R. A. (2000). Change blindness: Implications for the nature of visual attention. In M. Jenkin & L. Harris (Eds.), *Vision and attention* (pp. 169-188). New York: Springer.

Roeber, U., Widmann, A., Trujillo-Barreto, N. J., Herrmann, C. S., O'Shea, R. P., & Schroger, E. (2008). Early correlates of visual awareness in the human brain: Time and place from event-related brain potentials. *J Vis, 8*(3), 21 21-12.

Rolke, B. (2008). Temporal preparation facilitates perceptual identification of letters. *Percept Psychophys, 70*(7), 1305-1313.

Rolke, B., & Hofmann, P. (2007). Temporal uncertainty degrades perceptual processing. *Psychon Bull Rev, 14*(3), 522-526.

Rolls, E. T. (2007). A computational neuroscience approach to consciousness. *Neural Netw, 20*(9), 962-982.

Rorden, C., Guerrini, C., Swainson, R., Lazzeri, M., & Baylis, G. C. (2008). Event related potentials reveal that increasing perceptual load leads to increased responses for target stimuli and decreased responses for irrelevant stimuli. *Front Hum Neurosci, 2*, 4.

Rosenthal, D. M. (2002). How many kinds of consciousness? *Conscious Cogn, 11*(4), 653-665.

Rounis, E., Maniscalco, B., Rothwell, J. C., Passingham, R. E., & Lau, H. C. (2010). Theta-burst transcranial magnetic stimulation to the prefrontal cortex impairs metacognitive visual awareness. *Cognit Neurosci, 1*(3), 165-175.

Sagi, D., & Tanne, D. (1994). Perceptual learning: Learning to see. *Curr Opin Neurobiol, 4*(2), 195-199.

Sahraie, A., Trevethan, C. T., MacLeod, M. J., Murray, A. D., Olson, J. A., & Weiskrantz, L. (2006). Increased sensitivity after repeated stimulation of residual spatial channels in blindsight. *Proc Natl Acad Sci U S A, 103*(40), 14971-14976.

Sahraie, A., Weiskrantz, L., & Barbur, J. L. (1998). Awareness and confidence ratings in motion perception without geniculo-striate projection. *Behav Brain Res, 96*(1-2), 71-77.

Sahraie, A., Weiskrantz, L., Barbur, J. L., Simmons, A., Williams, S. C., & Brammer, M. J. (1997). Pattern of neuronal activity associated with conscious and unconscious processing of visual signals. *Proc Natl Acad Sci U S A, 94*(17), 9406-9411.

Schärli, H., Brugger, P., Regard, M., Mohr, C., & Landis, T. (2003). Localisation of "unseen" visual stimuli: Blindsight in normal observers? *Swiss J Psychol Schweiz Z Psychol Rev Suisse Psychol, 62*(3), 159-165.

Schiller, P. H. (1965). Monoptic and dichoptic visual masking by patterns and flashes. *J Exp Psychol, 69*(2), 193-199.

Schiller, P. H., & Smith, M. C. (1966). Detection in metacontrast. *J Exp Psychol, 71*(1), 32-39.

Schiltz, C., Bodart, J. M., Dubois, S., Dejardin, S., Michel, C., Roucoux, A., et al. (1999). Neuronal mechanisms of perceptual learning: changes in human brain activity with training in orientation discrimination. *Neuroimage, 9*(1), 46-62.

Schoups, A., Vogels, R., & Orban, G. A. (1995). Human perceptual learning in identifying the oblique orientation: retinotopy, orientation specificity and monocularity. *J Physiol, 483 (Pt 3)*, 797-810.

Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature, 412*(6846), 549-553.

Schubö, A., Schlaghecken, F., & Meinecke, C. (2001). Learning to ignore the mask in texture segmentation tasks. *J Exp Psychol Hum Percept Perform, 27*(4), 919-931.

Schurger, A., Pereira, F., Treisman, A., & Cohen, J. D. (2010). Reproducibility distinguishes conscious from nonconscious neural representations. *Science, 327*(5961), 97-99.

Schwartz, O., Hsu, A., & Dayan, P. (2007). Space and time in visual context. *Nat Rev Neurosci, 8*(7), 522-535.

Schwartz, S., Maquet, P., & Frith, C. (2002). Neural correlates of perceptual learning: a functional MRI study of visual texture discrimination. *Proc Natl Acad Sci U S A, 99*(26), 17137-17142.

Schwiedrzik, C. M., Alink, A., Kohler, A., Singer, W., & Muckli, L. (2007). A spatio-temporal interaction on the apparent motion trace. *Vision Res, 47*(28), 3424-3433.

Schwiedrzik, C. M., Singer, W., & Melloni, L. (2009). Sensitivity and perceptual awareness increase with practice in metacontrast masking. *J Vis, 9*(10), 18 11-18.

Searle, J. R. (1998). How to study consciousness scientifically. *Philos Trans R Soc Lond B Biol Sci, 353*(1377), 1935-1942.

Seitz, A. R., Kim, D., & Watanabe, T. (2009). Rewards evoke learning of unconsciously processed visual stimuli in adult humans. *Neuron, 61*(5), 700-707.

Seitz, A. R., Lefebvre, C., Watanabe, T., & Jolicoeur, P. (2005). Requirement for high-level processing in subliminal learning. *Curr Biol, 15*(18), R753-755.

Seitz, A. R., Nanez, J. E., Holloway, S. R., Koyama, S., & Watanabe, T. (2005). Seeing what is not there shows the costs of perceptual learning. *Proc Natl Acad Sci U S A, 102*(25), 9080-9085.

Seitz, A. R., Nanez, J. E., Sr., Holloway, S., Tsushima, Y., & Watanabe, T. (2006). Two cases requiring external reinforcement in perceptual learning. *J Vis, 6*(9), 966-973.

Sekuler, R., & Ball, K. (1977). Mental set alters visibility of moving targets. *Science, 198*(4312), 60-62.

Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nat Neurosci, 8*(10), 1391-1400.

Serre, T., Oliva, A., & Poggio, T. (2007). A feedforward architecture accounts for rapid categorization. *Proc Natl Acad Sci U S A, 104*(15), 6424-6429.

Shikata, E., Hamzei, F., Glauche, V., Knab, R., Dettmers, C., Weiller, C., et al. (2001). Surface orientation discrimination activates caudal and anterior intraparietal sulcus in humans: an event-related fMRI study. *J Neurophysiol, 85*(3), 1309-1314.

Shipp, S., Watson, J. D., Frackowiak, R. S., & Zeki, S. (1995). Retinotopic maps in human prestriate visual cortex: the demarcation of areas V2 and V3. *Neuroimage, 2*(2), 125-132.

Sigman, M., Cecchi, G. A., Gilbert, C. D., & Magnasco, M. O. (2001). On a common circle: natural scenes and Gestalt rules. *Proc Natl Acad Sci U S A, 98*(4), 1935-1940.

Singer, W. (2000). Phenomenal awareness and consciousness from a neurobiological perspective. In T. Metzinger (Ed.), *Neural correlates of consciousness: empirical and conceptual questions* (pp. 121-137). Cambridge, MA: MIT Press.

Singer, W. (2009). Consciousness and neural synchronization. In S. Laureys & G. Tononi (Eds.), *The neurology of consciousness. Cognitive neuroscience and neuropathology* (pp. 43-52). New York, NY: Academic Press.

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends Neurosci, 27*(3), 161-168.

Snodgrass, J. G., & Feenan, K. (1990). Priming effects in picture fragment completion: support for the perceptual closure hypothesis. *J Exp Psychol Gen, 119*(3), 276-296.

Soto, D., & Humphreys, G. W. (2006). Seeing the content of the mind: enhanced awareness through working memory in patients with visual extinction. *Proc Natl Acad Sci U S A, 103*(12), 4789-4792.

Soto, D., Wriglesworth, A., Bahrami-Balani, A., & Humphreys, G. W. (2010). Working memory enhances visual perception: evidence from signal detection analysis. *J Exp Psychol Learn Mem Cogn, 36*(2), 441-456.

Sridharan, D., Levitin, D. J., & Menon, V. (2008). A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. *Proc Natl Acad Sci U S A, 105*(34), 12569-12574.

Sterzer, P., & Rees, G. (2008). A neural basis for percept stabilization in binocular rivalry. *J Cogn Neurosci, 20*(3), 389-399.

Sterzer, P., Russ, M. O., Preibisch, C., & Kleinschmidt, A. (2002). Neural correlates of spontaneous direction reversals in ambiguous apparent visual motion. *Neuroimage, 15*(4), 908-916.

Stocker, A. A., & Simoncelli, E. P. (2006). Sensory adaptation within a Bayesian framework for perception. In Y. Weiss, B. Schölkopf & J. Platt (Eds.), *Advances in neural information processing systems 18* (pp. 1291-1298). Cambridge, MA: MIT Press.

Stoerig, P. (2006). Blindsight, conscious vision, and the role of primary visual cortex. *Prog Brain Res, 155*, 217-234.

Stokes, M., Thompson, R., Nobre, A. C., & Duncan, J. (2009). Shape-specific preparatory activity mediates attention to targets in human visual cortex. *Proc Natl Acad Sci U S A, 106*(46), 19569-19574.

Suchow, J. W., & Alvarez, G. A. (2011). Motion silences awareness of visual change. *Curr Biol, 21*(2), 140-143.

Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends Cogn Sci, 13*(9), 403-409.

Summerfield, C., Jack, A. I., & Burgess, A. P. (2002). Induced gamma activity is associated with conscious awareness of pattern masked nouns. *Int J Psychophysiol, 44*(2), 93-100.

Summerfield, C., Trittschuh, E. H., Monti, J. M., Mesulam, M. M., & Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nat Neurosci, 11*(9), 1004-1006.

Sundareswara, R., & Schrater, P. R. (2008). Perceptual multistability predicted by search model for Bayesian decisions. *J Vis, 8*(5), 12 11-19.

Suzuki, H., & Azuma, M. (1983). Topographic studies on visual neurons in the dorsolateral prefrontal cortex of the monkey. *Exp Brain Res, 53*(1), 47-58.

Szczepanowski, R., & Pessoa, L. (2007). Fear perception: Can objective and subjective awareness measures be dissociated? *J Vis, 7*(4), 10 11-17.

Talairach, J., & Tournoux, P. (1988). *Co-planar stereotactic atlas of the human brain. 3-dimensional proportional systems: an approach to cerebral imaging*. New York, NY: Thieme Medical Publishers.

Tata, M. S. (2002). Attend to it now or lose it forever: Selective attention, metacontrast masking, and object substitution. *Percept Psychophys, 64*(7), 1028-1038.

Todd, J. J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature, 428*(6984), 751-754.

Tononi, G., & Koch, C. (2008). The neural correlates of consciousness: an update. *Ann N Y Acad Sci, 1124*, 239-261.

Trevethan, C. T., Sahraie, A., & Weiskrantz, L. (2007). Form discrimination in a case of blindsight. *Neuropsychologia, 45*(9), 2092-2103.

van Turennout, M., Ellmore, T., & Martin, A. (2000). Long-lasting cortical plasticity in the object naming system. *Nat Neurosci, 3*(12), 1329-1334.

Ventura, J. (1980). Foveal metacontrast: I. Criterion content and practice effects. *J Exp Psychol Hum Percept Perform, 6*(3), 473-485.

Vidnyanszky, Z., & Sohn, W. (2005). Learning to suppress task-irrelevant visual stimuli with attention. *Vision Res, 45*(6), 677-685.

Vilares, I., & Körding, K. (2011). Bayesian models: the structure of the world, uncertainty, behavior, and the brain. *Ann N Y Acad Sci, 1224*(1), 22-39.

Volz, K. G., Schubotz, R. I., & von Cramon, D. Y. (2003). Predicting events of varying probability: uncertainty investigated by fMRI. *Neuroimage, 19*(2 Pt 1), 271-280.

von der Heydt, R., Peterhans, E., & Baumgartner, G. (1984). Illusory contours and cortical neuron responses. *Science, 224*(4654), 1260-1262.

Vorberg, D., Mattler, U., Heinecke, A., Schmidt, T., & Schwarzbach, J. (2004). Invariant time course of priming with and without awareness. In C. Kaernbach, E. Schröger & H. Müller (Eds.), *Psychophysics beyond sensation. Laws and invariants of human cognition* (pp. 271-288). Mahwah, NJ: Lawrence Erlbaum Associates.

Wainwright, M. J. (1999). Visual adaptation as optimal information transmission. *Vision Res, 39*(23), 3960-3974.

Ward, T. B., & Ross, L. E. (1977). Laterality differences and practice effects under central backward masking conditions. *Memory & Cognition, 5*(2), 221-226.

Watanabe, T., Nanez, J. E., & Sasaki, Y. (2001). Perceptual learning without perception. *Nature, 413*(6858), 844-848.

Webster, M. A., & MacLin, O. H. (1999). Figural aftereffects in the perception of faces. *Psychon Bull Rev, 6*(4), 647-653.

Webster, M. A., & Mollon, J. D. (1991). Changes in colour appearance following post-receptoral adaptation. *Nature, 349*(6306), 235-238.

Weidner, R., Shah, N. J., & Fink, G. R. (2006). The neural basis of perceptual hypothesis generation and testing. *J Cogn Neurosci, 18*(2), 258-266.

Wenger, M. J., Copeland, A. M., Bittner, J. L., & Thomas, R. D. (2008). Evidence for criterion shifts in visual perceptual learning: Data and implications. *Percept Psychophys, 70*(7), 1248-1273.

Wenger, M. J., & Rasche, C. (2006). Perceptual learning in contrast detection: Presence and cost of shifts in response criteria. *Psychon Bull Rev, 13*(4), 656-661.

Wessinger, C. M., Fendrich, R., & Gazzaniga, M. S. (1999). Variability of residual vision in hemianopic subjects. *Restor Neurol Neurosci, 15*(2-3), 243-253.

Westheimer, G. (1999). Discrimination of short time intervals by the human observer. *Exp Brain Res, 129*(1), 121-126.

Whitaker, D., & McGraw, P. V. (2000). Long-term visual experience recalibrates human orientation perception. *Nat Neurosci, 3*(1), 13.

Wickens, T. D. (2002). *Elementary signal detection theory*. Oxford: Oxford University Press.

Wiens, S. (2007). Concepts of visual consciousness and their measurement. *Adv Cog Psychol, 3*(1-2), 349-359.

Wilimzig, C., Tsuchiya, N., Fahle, M., Einhäuser, W., & Koch, C. (2008). Spatial attention increases performance but not subjective confidence in a discrimination task. *J Vis, 8*(5), 7 1-10.

Wilson, H. R. (2007). Minimal physiological conditions for binocular rivalry and rivalry memory. *Vision Res, 47*(21), 2741-2750.

Wohlschläger, A. M., Specht, K., Lie, C., Mohlberg, H., Wohlschläger, A., Bente, K., et al. (2005). Linking retinotopic fMRI mapping and anatomical probability maps of human occipital areas V1 and V2. *Neuroimage, 26*(1), 73-82.

Wolford, G., & Kim, H.-Y. (1992). The role of visible persistence in backward masking. In A. F. Healy, S. M. Kosslyn & R. M. Shiffrin (Eds.), *From learning processes to cognitive processes. Essays in honor of William K. Estes* (Vol. 2, pp. 161-180). Hillsdale, NJ: Lawrence Erlbaum Associates.

Wolford, G., Marchak, F., & Hughes, H. (1988). Practice effects in backward masking. *J Exp Psychol Hum Percept Perform, 14*(1), 101-112.

Yang, T., & Maunsell, J. H. (2004). The effect of perceptual learning on neuronal responses in monkey visual area V4. *J Neurosci, 24*(7), 1617-1626.

Yokoi, I., & Komatsu, H. (2009). Relationship between neural responses and visual grouping in the monkey parietal cortex. *J Neurosci, 29*(42), 13210-13221.

Yotsumoto, Y., Sasaki, Y., Chan, P., Vasios, C. E., Bonmassar, G., Ito, N., et al. (2009). Location-specific cortical activation changes during sleep after training for perceptual learning. *Curr Biol, 19*(15), 1278-1282.

Yotsumoto, Y., Watanabe, T., & Sasaki, Y. (2008). Different dynamics of performance and brain activation in the time course of perceptual learning. *Neuron, 57*(6), 827-833.

Young, M. P. (1992). Objective analysis of the topological organization of the primate cortical visual system. *Nature, 358*(6382), 152-155.

Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? *Trends Cogn Sci, 10*(7), 301-308.

Zihl, J. (1980). "Blindsight": Improvement of visually guided eye movements by systematic practice in patients with cerebral blindness. *Neuropsychologia, 18*(1), 71-77.

Zihl, J., & von Cramon, D. (1985). Visual field recovery from scotoma in patients with postgeniculate damage. A review of 55 cases. *Brain, 108 (Pt 2)*, 335-365.

Zihl, J., & Werth, R. (1984). Contributions to the study of "blindsight"--II. The role of specific practice for saccadic localization in patients with postgeniculate visual field defects. *Neuropsychologia, 22*(1), 13-22.

Zohary, E., Celebrini, S., Britten, K. H., & Newsome, W. T. (1994). Neuronal plasticity that underlies improvement in perceptual performance. *Science, 263*(5151), 1289-1292.

E

ERKLÄRUNGEN

ERKLÄRUNG

Ich erkläre hiermit, daß ich mich bisher keiner Doktorprüfung unterzogen habe.

Frankfurt am Main, den (Caspar M. Schwiedrzik)

EIDESSTATTLICHE VERSICHERUNG

Ich erkläre hiermit an Eides Statt, daß ich die vorgelegte Dissertation über *"The Role of Previous Experience in Conscious Perception"* selbständig angefertigt und mich anderer Hilfsmittel als der in ihr angegebenen nicht bedient habe, insbesondere, daß alle Entlehnungen aus anderen Schriften mit Angabe der betreffenden Schrift gekennzeichnet sind. Ich versichere, nicht die Hilfe einer kommerziellen Promotionsvermittlung in Anspruch genommen zu haben

Frankfurt am Main, den (Caspar M. Schwiedrzik)

C

Curriculum vitae

PERSONAL DETAILS

Date of birth	January 14, 1983
Place of birth	Cologne, Germany
Marital status	married
Nationality	German

EDUCATION

2002	Abitur, Kurfürst Friedrich Gymnasium, Heidelberg, Germany (grade: 1.1)
2003 - 2008	Studies of Psychology at the University of Konstanz, Germany (specialization in Cognitive Neuroscience and Clinical Neuropsychology, elective courses in Biology)
2007	Diploma student at the Max Planck Institute for Brain Research, Frankfurt am Main, Germany, Department of Neurophysiology Thesis title: "Spatio-temporal interactions on the apparent motion trace" (supervisors Prof. Rockstroh, Dr. Muckli, grade: 1, equivalent to A)
2008	Diploma, University of Konstanz, Germany (grade: 1, equivalent to A)
since July 2008	PhD student at the Max Planck Institute for Brain Research, Frankfurt am Main, Germany, Department of Neurophysiology (supervisor Prof. Dr. Wolf Singer)

PUBLICATIONS

Schwiedrzik CM, Alink A, Kohler A, Singer W, Muckli L (2007). A spatio-temporal interaction on the apparent motion trace. Vision Research, 47(28):3424-33.

Schwiedrzik, C.M (2007). Spatio-temporal interactions on the apparent motion trace. (Diploma thesis, University of Konstanz).

Melloni L, Schwiedrzik CM, Wibral M, Rodriguez E, Singer W (2009). Response to: Yuval-Greenberg et al., "Transient induced gamma-band response in EEG as a manifestation of miniature saccades", Neuron 58, 429-441. Neuron, 62(1):8-10.

Schwiedrzik CM (2009). Retina or visual cortex? The site of phosphene induction by transcranial alternating current stimulation. Frontiers in Integrative Neuroscience, 18;3:6.

Melloni L*, Schwiedrzik CM*, Rodriguez E, Singer W (2009). (Micro)Saccades, corollary activity and cortical oscillations. Trends in Cognitive Sciences, 13(6): 239-245. (* shared first authorship)

Schwiedrzik CM, Singer W, Melloni L (2009). Sensitivity and perceptual awareness increase with practice in metacontrast masking. Journal of Vision, 9(10):18, 1-18.

Alink A, Schwiedrzik CM, Kohler A, Singer W, Muckli L (2010). Stimulus predictability reduces responses in primary visual cortex. Journal of Neuroscience, 30(8): 2960-2966.

Melloni L, Schwiedrzik CM, Müller N, Rodriguez E, Singer W (2011). Expectations change the signatures and timing of electrophysiological correlates of perceptual awareness. Journal of Neuroscience, 31(4): 1386-1396.

Schwiedrzik CM, Singer W, Melloni L (2011). Subjective and objective learning effects dissociate in space and in time. Proceedings of the National Academy of Sciences USA, 108(11): 4506-4511.

Melloni L, Schwiedrzik CM, Rodriguez E, Singer W (submitted). Conscious and unconscious perception, and their relation to short- and long-range synchrony.

Schwiedrzik CM, Ruff CC, Leitner F, Singer W, Melloni L (in preparation). Untangling perceptual memory: hysteresis and adaptation map into separate cortical networks.

TALKS

May 2010	Interdisciplinary Center for Neuroscience Frankfurt, Goethe University, Frankfurt am Main
November 2010	Laboratory for Cognitive Neuroscience and Neuroimaging (C. E. Schroeder), Nathan S. Kline Institute for Psychiatric Research, Orangeburg, NY, USA
November 2010	Krekelberg Neuroscience Laboratory, Center for Molecular and Behavioral Neuroscience, Rutgers University, Newark, NJ, USA
January 2011	Laboratory of Neural Systems (W. Freiwald), Rockefeller University, New York, USA

CONFERENCE CONTRIBUTIONS

Schwiedrzik CM, Alink A, Kohler A, Singer W, Muckli L. “A spatio-temporal interaction on the apparent motion trace”. Poster at the European Conference on Visual Perception (ECVP), Arezzo, Italy, 2007.

Alink A, Schwiedrzik CM, Kohler A, Singer W, Muckli L. “When apparent motion and real stimuli meet in primary visual cortex”. Talk at the Meeting of the Society for Neuroscience (SfN), San Diego, USA, 2007.

Schwiedrzik CM, Singer W, Melloni M. “Sensitivity and perceptual awareness increase with practice in metacontrast masking”. Poster at the European Conference on Visual Perception (ECVP), Utrecht, the Netherlands, 2008.

Schwiedrzik CM, Singer W, Melloni M. “Sensitivity and perceptual awareness increase with practice in metacontrast masking”. Poster & Talk at the Rovereto Attention Workshop 2008: Attention and Motor Control, Rovereto, Italy, 2008.

Schwiedrzik CM, Singer W, Melloni M. “Sensitivity and perceptual awareness increase with practice in metacontrast masking”. Poster at the NIN Conference: Perceptual Learning, Motor Learning, Automaticity, Amsterdam, the Netherlands, 2008.

Schwiedrzik CM, Singer W, Melloni M. “Sensitivity and perceptual awareness increase with practice in metacontrast masking”. Talk at the 13th Annual Meeting of the Association for the Scientific Study of Consciousness (ASSC), Berlin, Germany, 2009.

Aedo-Jury F, Schwiedrzik CM, Singer W, Melloni L. “BOLD signal in retinotopic areas correlates with objective perception but not with subjective report in perceptual learning”. Poster at 16th Annual Meeting of the Organization for Human Brain Mapping (HBM), Barcelona, Spain, 2010.

Schwiedrzik CM, Leitner F, Singer W, Ruff CC, Melloni L. “Dissociating perceptual hysteresis and adaptation in multistable dot lattices: An fMRI study”. Poster at the Meeting of the Society for Neuroscience (SfN), San Diego, USA, 2010.

Melloni L, Müller N, Schwiedrzik CM, Rodriguez E, Singer W. “Expectations alter the timing of access to awareness”. Poster at the Meeting of the Society for Neuroscience (SfN), San Diego, USA, 2010.

Aedo-Jury F, Schwiedrzik CM, Singer W, Melloni L. “Differential brain activation for increases in objective performance and subjective awareness in perceptual learning”. Poster at the Meeting of the Society for Neuroscience (SfN), San Diego, USA, 2010.

Schwiedrzik CM, Leitner F, Singer W, Ruff CC, Melloni L. “Perceptual hysteresis and adaptation are expressed in distinct cortical networks". Poster at the 11th International Conference on Cognitive Neuroscience (ICON), Mallorca, Spain, September 2011.

Aedo-Jury F, Schwiedrzik CM, Singer W, Melloni L. "Functional connectivity between objective performance and subjective visibility networks increases with training during perceptual learning". Poster at the 11th International Conference on Cognitive Neuroscience (ICON), Mallorca, Spain, September 2011.

AWARDS

Acknowledged as amongst the best 5% of graduates in psychology at the University of Konstanz in the winter semester 2008.

Poster prize at the Rovereto Attention Workshop 2008: Attention and Motor Control, Rovereto, Italy, 2008.

Poster prize at the NIN Conference: Perceptual Learning, Motor Learning, Automaticity, Amsterdam, the Netherlands, 2008.

Project award from the Latin American School for Education, Cognitive and Neural Sciences, San Pedro de Atacama, Chile, 2011.

GRANTS

09-2008 – 12-2009	Interdisciplinary Center for Neuroscience (IZNF) Frankfurt am Main grant "Plasticity through magnetic brain stimulation: Cellular mechanisms – modeling – modulation"
01-2010	Travel award from the "Brain Clocks & Rhythms" Summer School, Santiago, Chile
03-2011	Travel award from the "Latin American School for Education, Cognitive and Neural Sciences", San Pedro de Atacama, Chile