From the blind eye to the mind's eye:

How behavioral relevance determines access to visual awareness

Surya Gayet

The studies in this dissertation were funded by the Netherlands Organisation for Scientific Research (NWO, The Hague, grant number 404-10-306 to S. van der Stigchel and C. L. E. Paffen).

Assessment Committee:

Prof. dr. Kees van den Bos	Universiteit Utrecht, Utrecht
Prof. dr. Serge O. Dumoulin	Universiteit Utrecht, Utrecht
Prof. dr. Christian N. L. Olivers	Vrije Universiteit, Amsterdam
Prof. dr. Richard J. A. van Wezel	Donders Instituut, Nijmegen
Dr. Tomas Knapen	Vrije Universiteit, Amsterdam

ISBN/EAN: 978-90-393-6555-7

Cover art by Designbazen, designbazen.com

Lay-out by Ferdinand van Nispen, Citroenvlinder DTP&Vormgeving, *my*-thesis.nl

Printed by GVO drukkers & vormgevers B.V. | Ponsen & Looijen, Ede, The Netherlands

From the blind eye to the mind's eye:

How behavioral relevance determines access to visual awareness

Van blind oog tot geestesoog:

Hoe de relevantie van visuele informatie
tot bewuste waarneming leidt
(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht op gezag van de rector magnificus, prof.dr. G.J. van der Zwaan, ingevolge het besluit van het college voor promoties in het openbaar te verdedigen op vrijdag 24 juni 2016 des middags te 2.30 uur

door

Surya Gayet

geboren op 16 mei 1984
te Amsterdam

Promotoren: Prof. dr. A. Postma
Prof. dr. F.J. Verstraten

Copromotoren: Dr. S. van der Stigchel
Dr. C.L.E. Paffen

Table of contents

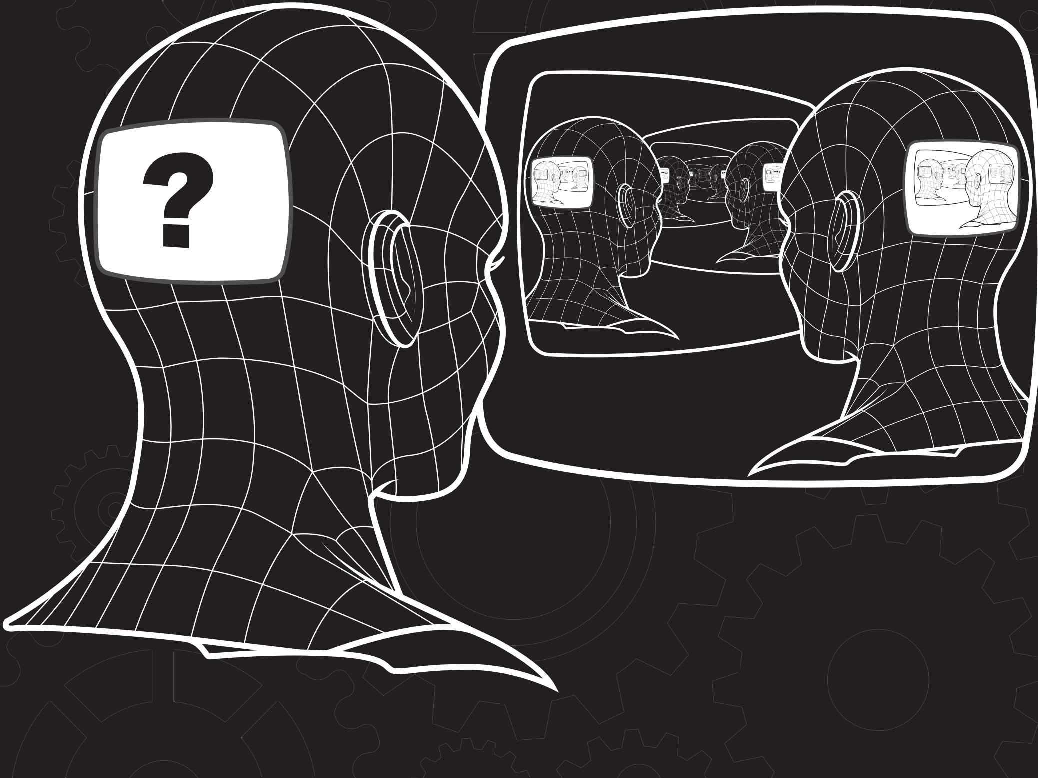

Prologue

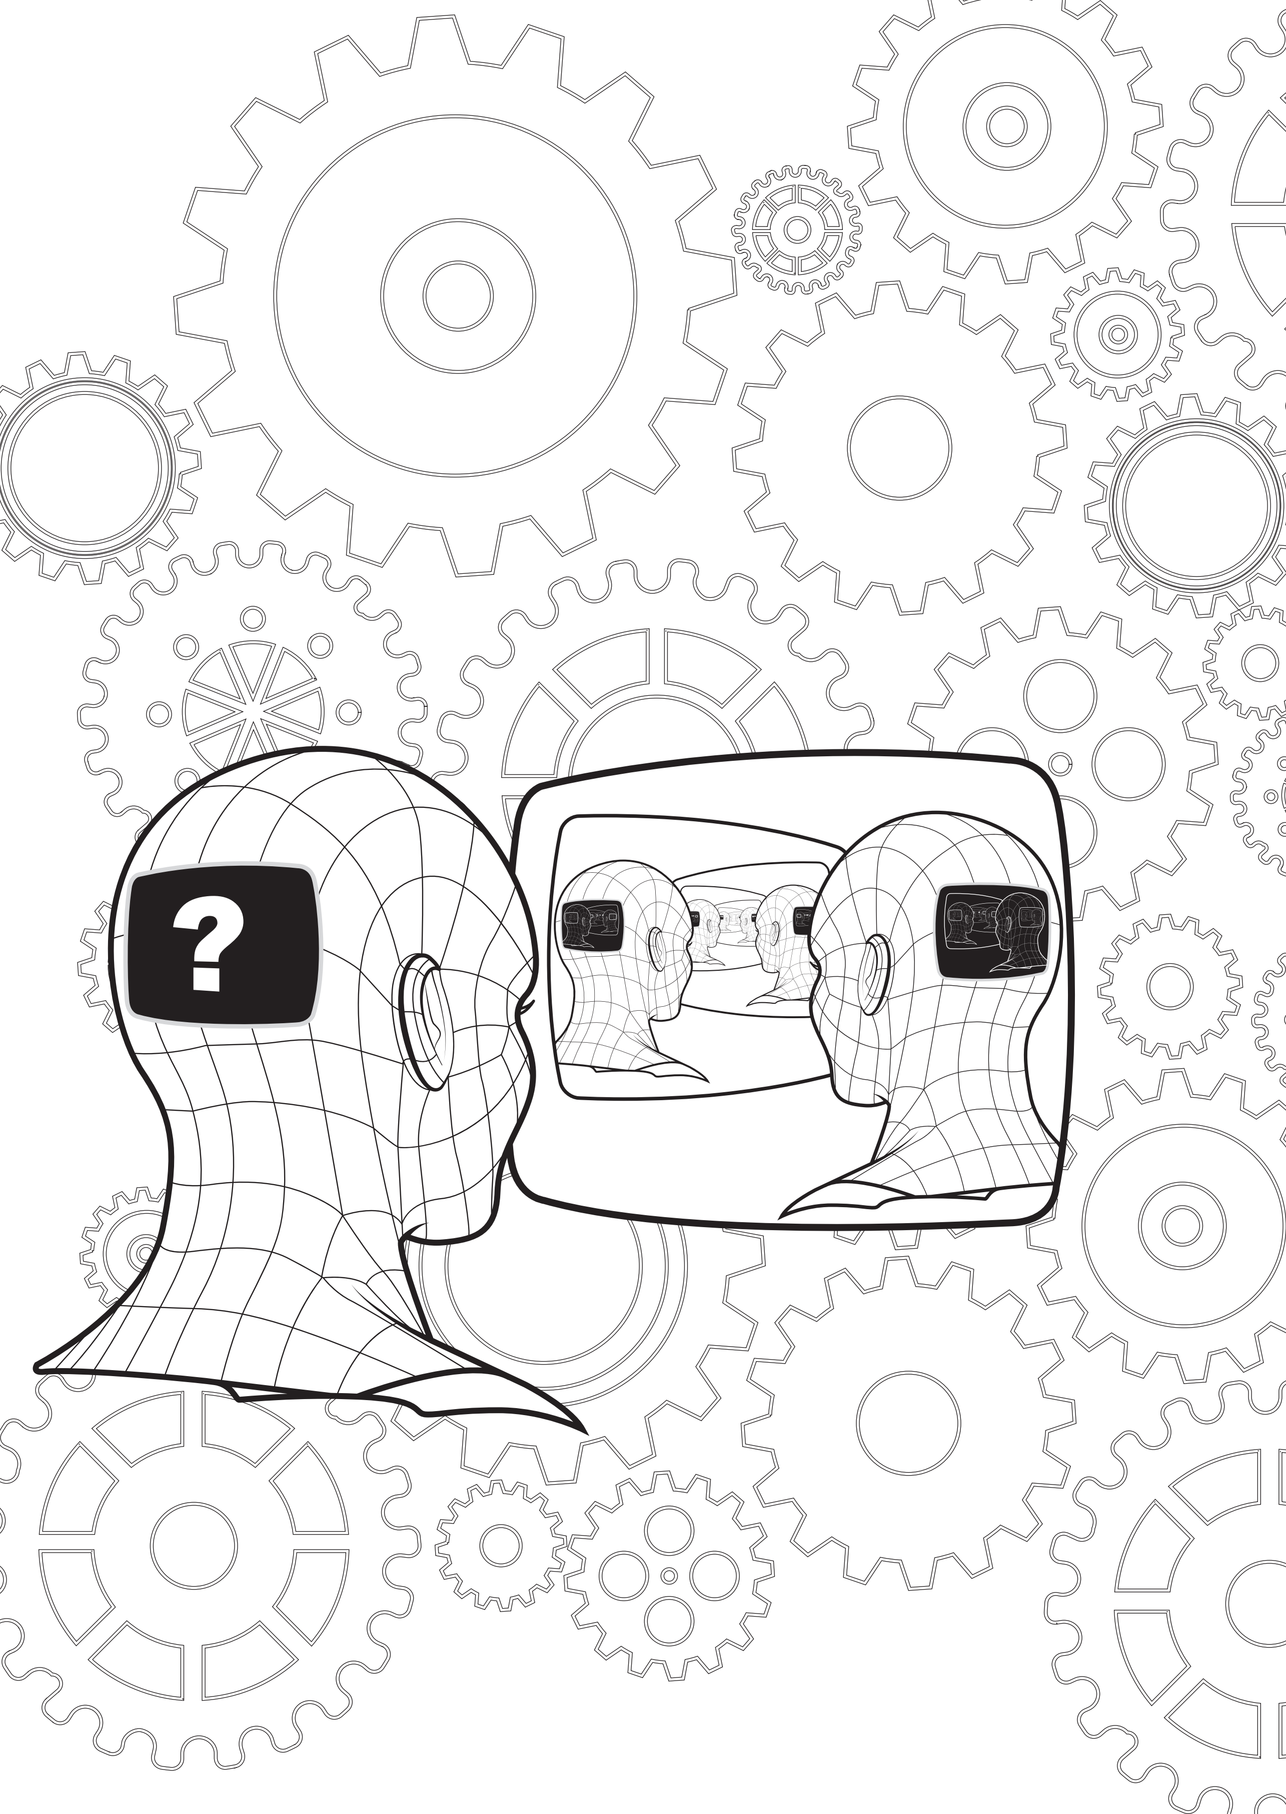

General introduction

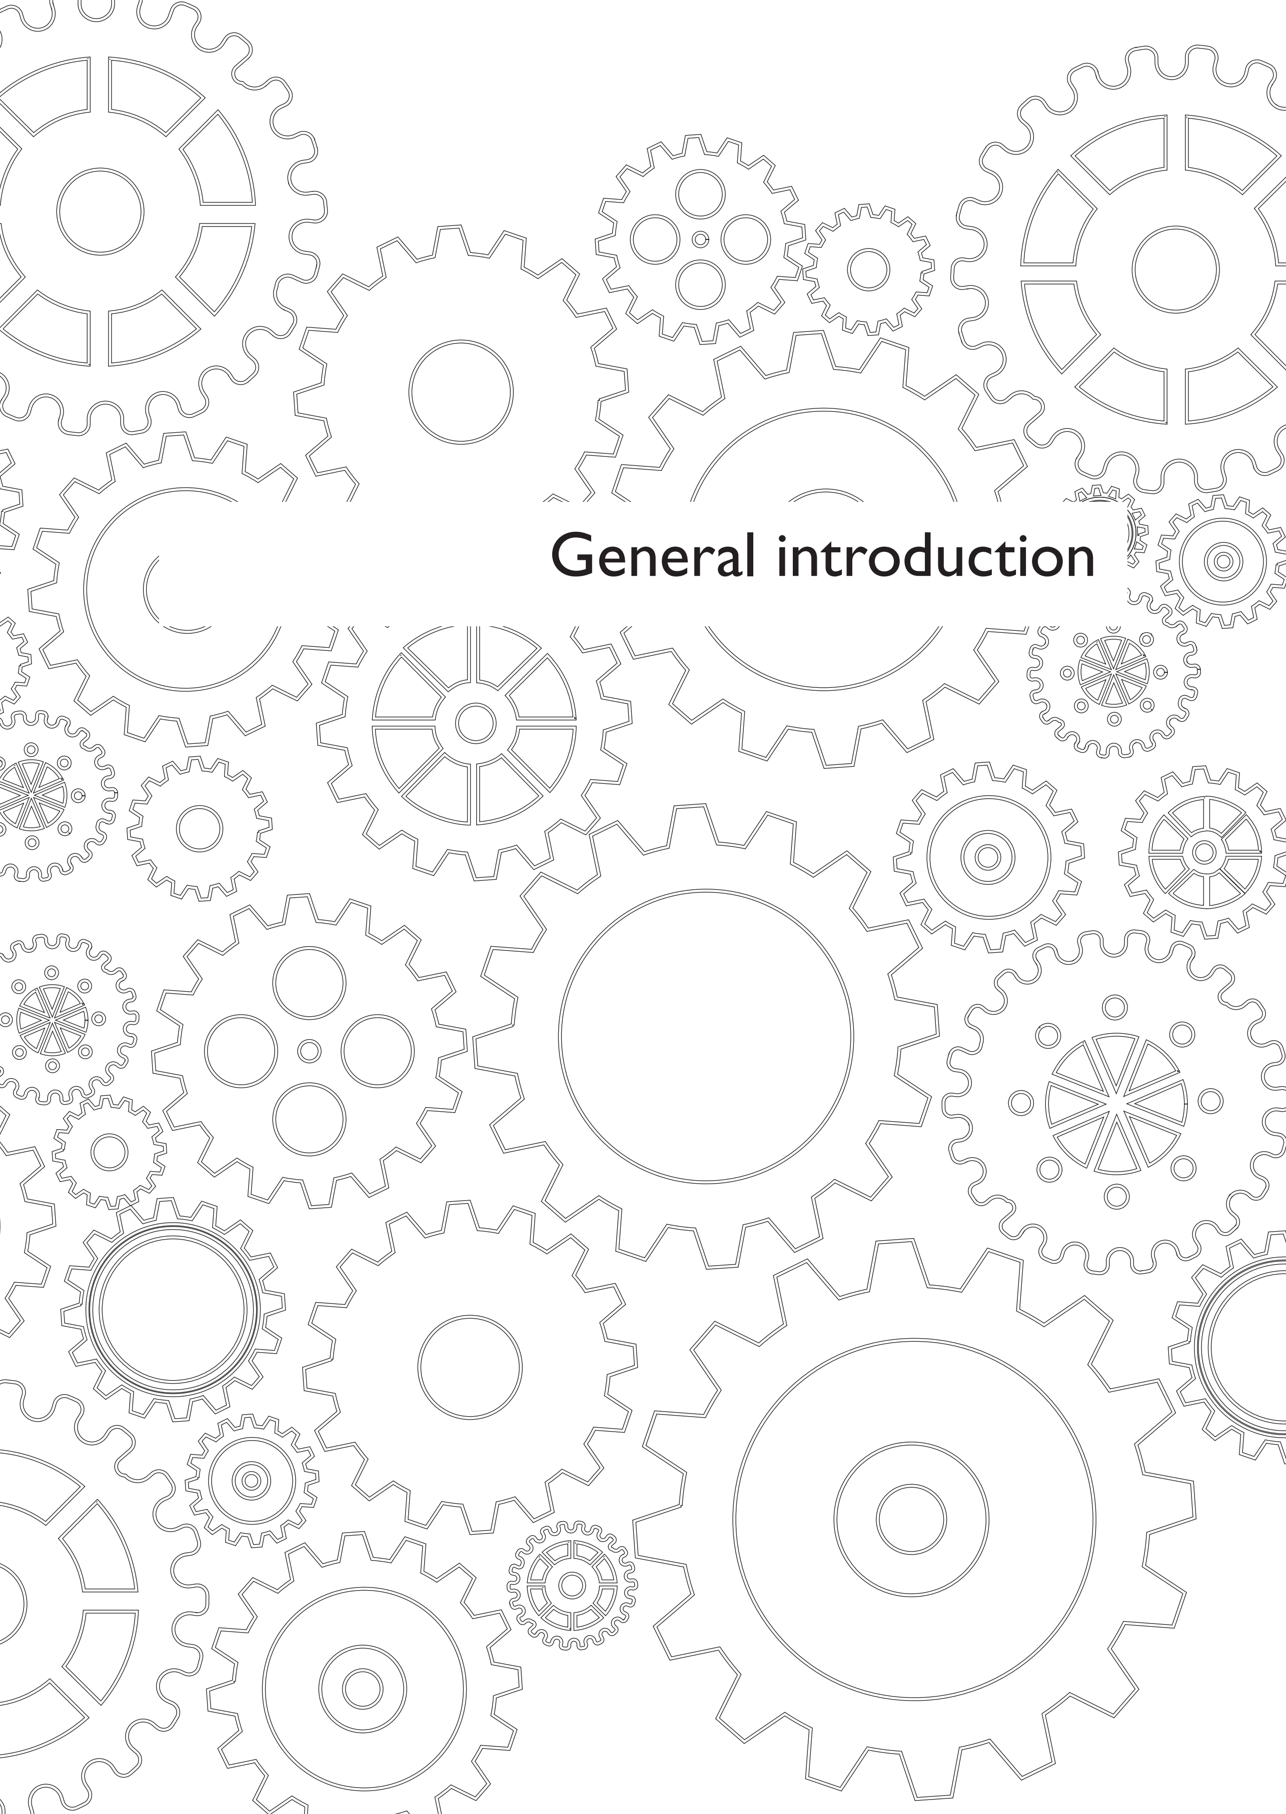

"Hey… I think I just sensed a flash of light." "So did I", a nearby voice says. "Me too", yet another voice adds, "more than once, actually. And so did most of my neighbors." A seemingly ever expanding humming of consent leaves no doubt about the sheer magnitude of this event. "What could this be?" some wonder. Eventually, they decide to send the information to the processing room, in the back, and get on with their jobs. "Those are not just flashes of light…" a voice echoes from the processing room, "those flashes follow a clear pattern. They seem to form an edge-like structure. In fact, there are many more of these edges." "And with different orientations as well" another voice intercedes. "The amount of edges and orientations is overwhelming", they jointly utter in awe. "I am not entirely sure what we are looking at" a distant voice adds "but whatever it is, it is definitely moving. If I am correct, it is moving leftwards." "The amount of moving edges is staggering" they exclaim all together, "what is this all about?" The confusion appears to reach a climax when others begin to doubt the presence of these moving edges altogether. "To be honest" this odd group of newcomers mumbles "we don't see any edges. We did notice a yellow veil though, with a little bit of black interspersed. It was coming from that direction as well. Do you think it is related to your moving edges? It most certainly is intriguing." In the excitement of the moment, all of their combined yapping and mumbling gradually increases in intensity and starts reverberating all around them. Back and forth, and louder and louder, and somehow, in an inexplicable way, the hotchpotch of individually unidentifiable exclamations amalgamates into a single coherent utterance: "Ah, finally, my train has arrived!"

"What is the origin of this strange yet familiar voice" they all wonder "which somehow deduces structure from all this chaos? Who is this mysterious entity, who undeniably relies upon our combined knowledge, yet seems unburdened by the magnitude of its contents?" "Although your questions are grammatically correct", the voice replies to their thoughts "asking who or what I am is intrinsically devoid of meaning and can therefore not be answered. You can ask me about what I do, though. And how I do it. Those are meaningful questions that are worthy of finding answers to. My job is to convey momentary experiences that summarize all your insignificant flashes of light, your boring edges, and your meaningless colors into a coherent visual world. It occurred to me that this would be more useful than, you know, counting photons all day."

At any moment in time, our brain is bombarded with sensory information. If we had to process all of this information consciously it would take us more than a lifetime to execute such simple tasks as tying our shoelaces. Consequently, only part of the visual information impinging on our retinae leads to a conscious visual experience (e.g., Edelman, & Tononi, 2000; Baars, 1997a; Dennett, 1991). This observation leads to two fundamental questions: What part of this visual information will eventually reach consciousness, and how does the visual system go about to select this information. These are the questions that I have tried to answer in the past four years, and that have been addressed empirically in the experiments that constitute my dissertation.

Before moving on to the empirical part of this dissertation, however, it is important that the reader and I have a mutual understanding on the meaning of "consciousness". This is not a trivial issue, considering that most dictionaries explain "consciousness" as "the state of being aware of something" and awareness as "the state of being conscious of something". Throughout this dissertation the terms "awareness" and "consciousness" are used interchangeably. In order to define consciousness for the present purpose, it is worthwhile mentioning that I am focusing on sensory consciousness (i.e., the experience of information gathered through the senses), and more specifically visual consciousness, as opposed to higher-level consciousness, such as self-consciousness, which refers to meta-cognitive states of the mind. In the light of this, my personal definition of consciousness is "a subjective experience that accompanies a sensory event". What I refer to as a sensory event is the occurrence of information in a sensory modality (i.e., vision, audition, etc.) that is triggered either externally, through retinal stimulation (e.g., seeing a car), or internally (e.g., imagining the color red, or dreaming of a palm tree). From the perspective of experimentation, consciousness is operationalized as the ability to convey a subjective and spontaneous report of a percept (e.g., Ramsøy, & Overgaard, 2004). This includes an observer's ability to answer such questions as: "Do you see a face or a house?", "Is some target located on the left or right side of fixation" or "did you see anything at all?". Over the years experimental researchers and philosophers have used many metaphors to describe consciousness. The famous philosopher Daniel Dennett characterized consciousness as "Fame in the brain" (Dennett, 1991), referring to the widespread brain activity associated with consciously accessible information (Dennett's view is closely related to neurobiologist Bernard Baars' Global Workspace Theory; Baars, 1998; 2005). The late biologist and Nobel laureate

Gerard Edelman coined the term "The remembered present" (Edelman, 1989), referring to the tight link between consciousness and working memory[1], which will play a prominent role in this dissertation. A similar view has been proposed by neuroscientist Victor Lamme, who related different forms of consciousness to different stages of short term memory (Lamme, 2004). The neuroscientists Guilio Tononi and Kristof Koch proposed a quantifiable measure of consciousness by equating consciousness to the amount of integrated information within a system (integrated information theory; Tononi, 2008; Tononi, & Koch, 2015).

In sum, there are almost as many theories of consciousness as there are researchers studying consciousness. Nonetheless, a number of characteristics that define consciousness, are shared by most of these theories and theorists. Firstly, consciousness is of limited capacity. That is, depending on the theory and definition, only one or a few entities of information can be accessible to consciousness at a particular moment in time. Secondly, information that is accessible to consciousness can be reported[2] behaviorally or verbally upon request (this excludes patients in a vegetative state and split brain patients; e.g., Sperry, 1984; Gazzaniga, 2005). Thirdly, it takes some time for information to reach consciousness. It is estimated that an absolute lower limit of 60-70ms is needed for visual or auditory stimuli to elicit a basic sensory experience, as constrained by neural propagation time (Efron, 1967). However, about 250 ms might be needed to consciously perceive a meaningful visual stimulus as such (e.g., recognizing that a face is a face; Koch, 2004). Fourthly,

---

[1] Visual working memory refers to the active maintenance of an image in the mind's eye, after termination of its sensory input. The human brain uses visual working memory to actively maintain relevant information online for imminent goal-directed behavior (for a review, see Baddeley, 2003).

[2] Some have argued that conscious perception overflows cognitive access (Block, 2011; Koch, & Tsuchiya, 2007; Lamme, 2003). This means that the content of our visual experience (referred to as phenomenal consciousness) comprises more visual information than we can explicitly report (referred to as access consciousness). Lamme (2003) refers to phenomenal consciousness as information that is conscious but unattended. Nothing precludes, in principle, that all visual information constituting a conscious visual experience can be attended and thereby reported. Likely, however, our means of reporting our conscious visual experience do not allow for conveying its rich phenomenal content in its entirety. This argument is well illustrated by the classic Sperling experiment (Sperling, 1960). In this experiment, observers were briefly presented with a 3 by 3 grid containing 9 letters. When they were asked to report all the letters that they had perceived, they could report about 4 or 5 letters. In another condition, a cue appeared after the grid of letters was removed from the screen, indicating which row of letters the observers should report. In this case, the observers were near perfect in reporting all three letters of any of the three rows, demonstrating that they had in fact consciously perceived all 9 letters, but were only able to report 4 or 5 when an inferior reporting strategy was used. Arguably, phenomenal consciousness therefore does not overflow access consciousness but, rather, consciousness per se overflows report. For the present purposes, I make the case that, using appropriate report paradigms, the content of consciousness can be inferred experimentally through an observer's ability to report visual information.

consciously accessible information is highly distributed throughout the brain. While neural activity elicited by non-consciously observed stimuli is restricted to sensory processing areas (e.g., visual cortex), the same stimuli engender a more widespread (Dehaene, Naccache, Cohen, Le Bihan, Mangin, Poline, & Rivière, 2001) and more synchronous (Srinivasan, Russell, Edelman, G. M., Tononi, 1999; Tononi, Srinivasan, Russell, & Edelman, 1998) pattern of brain activity when they are accessible to consciousness.

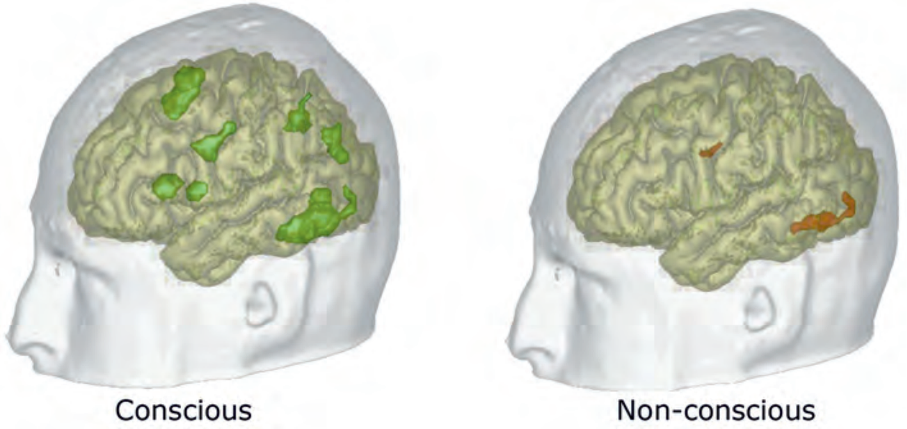

Figure 1. This image depicts the pattern of brain activity elicited by the same visual stimuli, depending on whether they gave rise to a conscious experience (left) or not (right). Typically, non-conscious visual information elicits a weaker pattern of activity that is restricted to early sensory processing areas (in this case the occipital cortex). In contrast, visual information that elicits a conscious experience elicits a widespread pattern of neural activity, which encompasses multiple cortical and subcortical brain areas. Image retrieved from Dehaene et al. (2001).

The widespread activation associated with consciously perceived information that is observed throughout multiple cortical and subcortical brain regions (see Figure 1) is an important signature of consciousness. Indeed, one can conceive how this causes limited capacity (the first characteristic), how it allows conscious information to be accessible to multiple cognitive subsystems throughout the brain (the second characteristic), and how it causes consciousness to lag behind sensory processing per se (the third characteristic). Considering the limited capacity, the number of involved brain structures, and the time it takes for visual information to reach consciousness, one may wonder why any part of our visual world has to reach consciousness at all. In light of these considerations, a number of cognitive functions have been proposed to require consciousness. Consciousness is believed to be a prerequisite for the integration of multiple visual features into a single coherent visual entity (e.g., Fahrenfort, Lamme, 2012; Treisman, 2003; Baars, 2002; Tononi, & Edelman, 1998; but, see Mudrik, Faivre, & Koch, 2014), for decision making (for reviews, see Dehaene, & Nacache, 2001; Hommel, 2007; Van Gaal, Lange, & Cohen, 2012) and planning (e.g., Crick & Koch, 2003), for performing demanding tasks, such as memorizing a telephone number (Dehaene, Kerszberg, & Changeux, 1998), and for selecting among different action plans and guiding novel behavior (Ansorge, Kunde, & Kiefer, 2014; Dehaene & Naccache, 2001; Kunde, Kiesel, & Hoffmann, 2003). These types of complex cognitive functions draw upon multiple cognitive subsystems. Appeals made on these systems are often conflicting, requiring the mediation of a distributed and flexible neural network. It has been argued that, because of this, such complex cognitive functions can only be performed with information that is accessible to consciousness (Lau, & Rosenthal, 2011; Baars, 2005; Dehaene, & Nacache, 2001). Taken together, these cognitive functions[3] that are all associated with consciousness, illustrate why information that is relevant for imminent behavior would benefit from reaching consciousness.

Conversely, one can wonder what cognitive functions are preserved in the absence of consciousness. Somnambulism, or sleepwalking, is a well-known form of prolonged absence of consciousness, in which perception and behavior occasionally remain relatively unaffected (for a review, see Plazzi, Vetrugno, Provini, & Montagna, 2005). On the extreme end of the spectrum, there are documented cases of adult somnambulists displaying such complex behavior as cooking, playing a musical instrument (e.g., Kavey, Whyte, Resor, Gidro-Frank, 1990), and even long-distance nocturnal automobile driving (e.g., Schenck, & Mahowald, 1995). On the one hand, these documented cases leave us with a sense of marvel vis-à-vis the capabilities of the non-conscious mind. Driving a car, for instance, requires fine-grained motor behavior allowing to use keys and switch gears, but additionally requires interpreting the visual world such as to recognize streets, dissociate between cars and pedestrians, avoid obstacles, and so on. On the other hand, when reading about these extreme

---

[3] As was already briefly brought up in the story at the start of the introduction, the emphasis in this dissertation lies on the cognitive functions that are associated with consciousness (what does it do, and how does it do it), rather than on unraveling the mysterious fabric that constitutes consciousness (what is it; James, 1904). Some scholars have argued that consciousness is the embodiment of these cognitive functions (Dennett, 1991). Others, however, have argued that these cognitive functions rely upon a neural state, which elicits consciousness as an epiphenomenon (leading to the so-called hard problem of consciousness; Nagel, 1974; Chalmers, 1995). By isolating consciousness from all accompanying cognitive functions and neural states, however, an entity remains that is untestable, cannot be falsified nor verified, and therefore resides outside the realm of science (Cohen, & Dennett, 2011).

cases, we are intuitively aware of some of the limitations of such non-conscious behavior. That is, we automatically assume that the somnambulist who was driving a car while sleeping, is able to do so while being awake as well. At no moment were we under the impression that the somnambulist had never seen a car before, but suddenly, in a nocturnal moment of inspiration, intuitively managed to handle the clutch, switch the gears, and learn the traffic code. Rather, we (correctly) assume that it is only well-learned, automated behavior that can be performed non-consciously.

While this conclusion seems obvious for such complex behavior as sleep driving, the limitation of non-consciously triggered behavior to well-learned and automated routines also holds for much simpler perceptual tasks (e.g., Schneider & Shiffrin, 1977). Traditional experiments demonstrate that observers attend to the location of sudden invisible stimuli in the periphery (for a review, see Mulckhuyse, & Theeuwes, 2010), and respond faster to targets that are preceded by a perceptually similar, but invisible, stimulus (e.g., Abrams & Greenwald, 2000; Damian, 2001). These behavioral effects elicited by non-conscious visual information appear to reflect automated, reflexive responses following predetermined routines. Recent findings suggest, however, that non-conscious visual processing is not restricted to neural-equivalents of the knee-tap reflex. Indeed, non-conscious visual information might trigger different routines depending on the cognitive context that is provided. For instance, non-conscious stimuli have been shown to only affect behavior to a subsequent visible target, when they happened to be of a color that was either attended (Ansorge, Horstmann, & Worschech, 2010) or happened to be useful for a concurrent task (Pan Cheng, & Luo, 2012).

Taken together, the preceding paragraphs have shown that non-conscious visual information has the potency to influence behavior. We have seen that whether or not it does so, however, might depend on the context. Part I of this dissertation, revolves around the scope and limits of how a cognitive context influences non-conscious perception (chapter 1) and the content of consciousness (chapter 2). For this purpose, we manipulated whether visual information is relevant for the goals of the observer or not. In chapter 1, we interspersed trials in which invisible arrows are presented throughout trials in which visible arrows are presented. The visible arrows could be either useful for the task at hand (i.e., 80% congruent with the location of a subsequent target) or not (i.e., 50% congruent). The results revealed that *invisible* arrows influenced response times to the targets, but only when the *visible* arrows provided an incentive to do so. Conversely, however, when we manipulated whether the *invisible* arrows were useful for the task at hand or not, this did not affect the influence of *visible* arrows on participants' response times. This demonstrates that non-conscious visual information has the potency to affect behavior, but only when a consciously accessible context provides an incentive to do so. In chapter 2, we measured the propensity of visual information to elicit a conscious experience, using a binocular rivalry paradigm (Wheatstone, 1838; Alais, & Blake, 2005; see Figure 2). The results revealed that visual information is more likely to elicit a conscious experience when it contains a feature (e.g., the color red) that was coincidentally relevant for a *concurrent* task. Whether or not this feature is relevant for the *current* task, however, did not affect the propensity of visual information to elicit a conscious experience. The converging findings from these two chapters lead us to conclude that as long as a cognitive context is provided consciously, it can codetermine the content of consciousness (chapter 2) and dictate whether or not non-conscious information influences behavior (chapter 1).

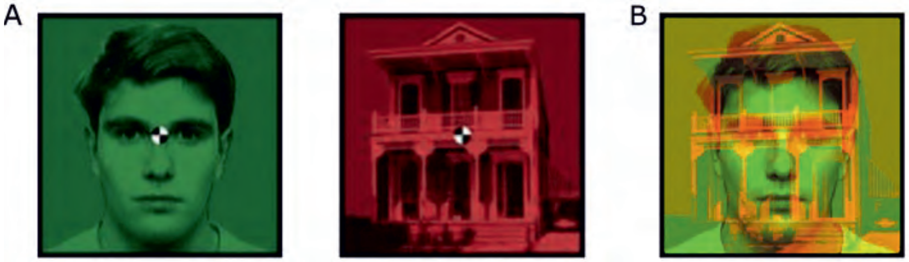

Figure 2. Example of a paradigm that allows for assessing the content of consciousness through subjective report. When each eye is presented with a different image (e.g., a house and a face), conscious perception fluctuates back and forth between the two images. This phenomenon is known as binocular rivalry (Wheatstone, 1838; Alais, & Blake, 2005). By asking observers to report what they are consciously perceiving (e.g., by pressing a key when they see the green image and another key whenever they see the red image), it allows for measuring changes in conscious perception, while the sensory stimulation itself remains constant. You can try it by looking through red/green anaglyph glasses to the image in panel B. Alternatively, you can attempt to cross-fuse the two images in panel A. For this, hold a pencil at about 10 cm from your nose, and fixate on the tip of the pencil while covertly attending the two images of panel A. You will see a third image appear, in between the red and the green image. After some practice, you will be able to get a stable view of this middle image, while maintaining fixation on the tip of your pencil. In this image you experience binocular rivalry. These images have been retrieved from Rees, Kreiman, & Koch (2002).

Considering the aforementioned cognitive functions associated with consciousness, and considering that consciousness is of limited capacity, it would serve adaptive functioning if visual information that is behaviorally relevant gains prioritized access to consciousness. Part II of this dissertation revolves around the question whether visual information that is initially not accessible to consciousness, can reach consciousness faster if it is behaviorally relevant to the observer. In chapter 3 we review an experimental paradigm that provides the means to measure (and compare) the propensity of visual stimuli to gain access to consciousness. This method, called breaking continuous flash suppression (b-CFS; Jiang, Costello, & He, 2007; Stein, Hebart, & Sterzer, 2011), consists of presenting a target stimulus to one eye, while presenting high-contrast flickering patterns to the other eye (continuous flash suppression, or CFS; Tsuchiya & Koch, 2005; a method derived from binocular rivalry). Flickering patterns in one eye cause the target stimulus in the other eye to be initially suppressed from consciousness, and the time it takes for the target to be released from suppression (e.g., such that its location or identity can be reported) provides a measure of its potency to reach access to consciousness. In the two subsequent chapters, we use two distinct methods to experimentally manipulate the behavioral relevance of basic visual stimuli. In chapter 4, we used classical fear conditioning (Mackintosh, 1983; Pavlov, 1927) to associate one of two colored rings with electric shocks, while never pairing the other annulus color with a shock. The results demonstrate that a non-conscious visual stimulus that was previously associated with electric shocks, and therefore signaled threat, reaches consciousness earlier in time than a comparable visual stimulus that was not associated with shocks. In chapter 5 we manipulated the content of visual working memory, as a means to manipulate behavioral relevance. Visual working memory refers to the active maintenance of visual information (after termination of its sensory input) for subsequent goal-directed behavior. Here, the results demonstrate that visual information that shares a feature (say, the color red) with the concurrent content of visual working memory reaches awareness faster than similar visual information that does not share this feature.

Using two radically different experimental manipulations (i.e., fear conditioning, and a visual working memory approach), we show that our visual system provides the means to propel behaviorally relevant information from the blind eye to the mind's eye. This leaves us with a final question for Part III of this dissertation: how does the visual system prioritize behaviorally relevant information over behaviorally irrelevant information in the competition for consciousness? To address this question we focused on one experimental manipulation of behavioral relevance in particular: the content of visual working memory. In chapter 6 we used computational modeling to investigate the perceptual processes leading up to the moment in time at which visual information reaches consciousness. The results reveal an initial bias towards (as opposed to a faster accumulation rate of perceptual evidence for) visual information that matches the content of visual working memory. We interpret these findings in light of a *Pre-activation* model, proposing that matching visual information taps into the same neural substrate that is already activated by the content of visual working memory, thereby reducing its threshold for reaching visual awareness. As such, we conclude that the modulation of consciousness by the content of visual working memory is capacitated by a shared neural substrate for visual representations elicited by visual information on the screen, and visual representations maintained in visual working memory. In chapter 7 we describe two functional magnetic resonance imaging (fMRI) studies that were aimed at investigating in what brain areas this interaction between perception and visual working memory takes place. A first experiment reveals that, when fully visible information matches the content of visual working memory, the pattern of brain activity elicited by this visual information is both quantitatively (i.e., BOLD response amplitude) and qualitatively (i.e., contains more multivariate information) enhanced. This modulation is observed in high-level visual areas, including regions that correspond with the lateral occipital complex and the intra-parietal sulcus. Interestingly, when visual information is removed from consciousness its neural trace is believed to be constrained to lower-level visual areas (e.g., Dehaene et al., 2001; even more so with our suppression method, Leopold & Logothetis, 1999; Yuval-Greenberg, & Heeger, 2013). Therefore, we investigated in the second experiment of chapter 7 whether visual information that remains non-conscious interacts with visual working memory as well. Strikingly, this was not the case. Rather, when visual information is completely inaccessible to consciousness, it seems to be unable to interact with the content of visual working memory.

In the discussion, at the end of this dissertation, we reconcile this latter finding from our imaging study (chapter 7, experiment 2) with the behavioral findings that we observed before (e.g., chapters 5 and 6). Simultaneously, we present a concise overview of recent studies in which the unconscious mind is proposed to be capable of cognitive functions that were hitherto reserved to the realm of consciousness, such as solving arithmetic equations (for a review, see Hassin, 2013;

but for a response, see Hesselmann & Moors, 2016). Non-conscious processing that appears to require higher level visual processing can probably be accounted for by differences in the a priori state of the observer. Based on our observations, as well as on recent findings from other labs, however, we tentatively add that the level at which visual information is suppressed from consciousness codetermines to what extent it can be processed. Arguably, at stronger levels of suppression non-conscious visual information processing is either fully abolished, or limited to the level of simple visual features. At more shallow levels of suppression, the processing of non-conscious visual information might encompass such visual characteristics as coarse stimulus configurations. Possibly, close to the threshold of consciousness, even relatively complex cognitive functions might be applied to visual information that does not give rise to a conscious percept. In the end, more complex cognitive functions seem to require consciousness, because they require the distributed activity pattern that characterizes conscious processing. Similarly, less complex cognitive functions seem to not require consciousness, as they can be performed with more focalized neural activity. Considering the graded nature of said cognitive functions, and considering the graded nature of said patterns of neural activity, we conclude that consciousness should be graded as well. That is, if we dissociate consciousness from these patterns of brain activity with which it is associated, and from these cognitive functions that it accompanies, we drag consciousness outside of the realm of empiricism.

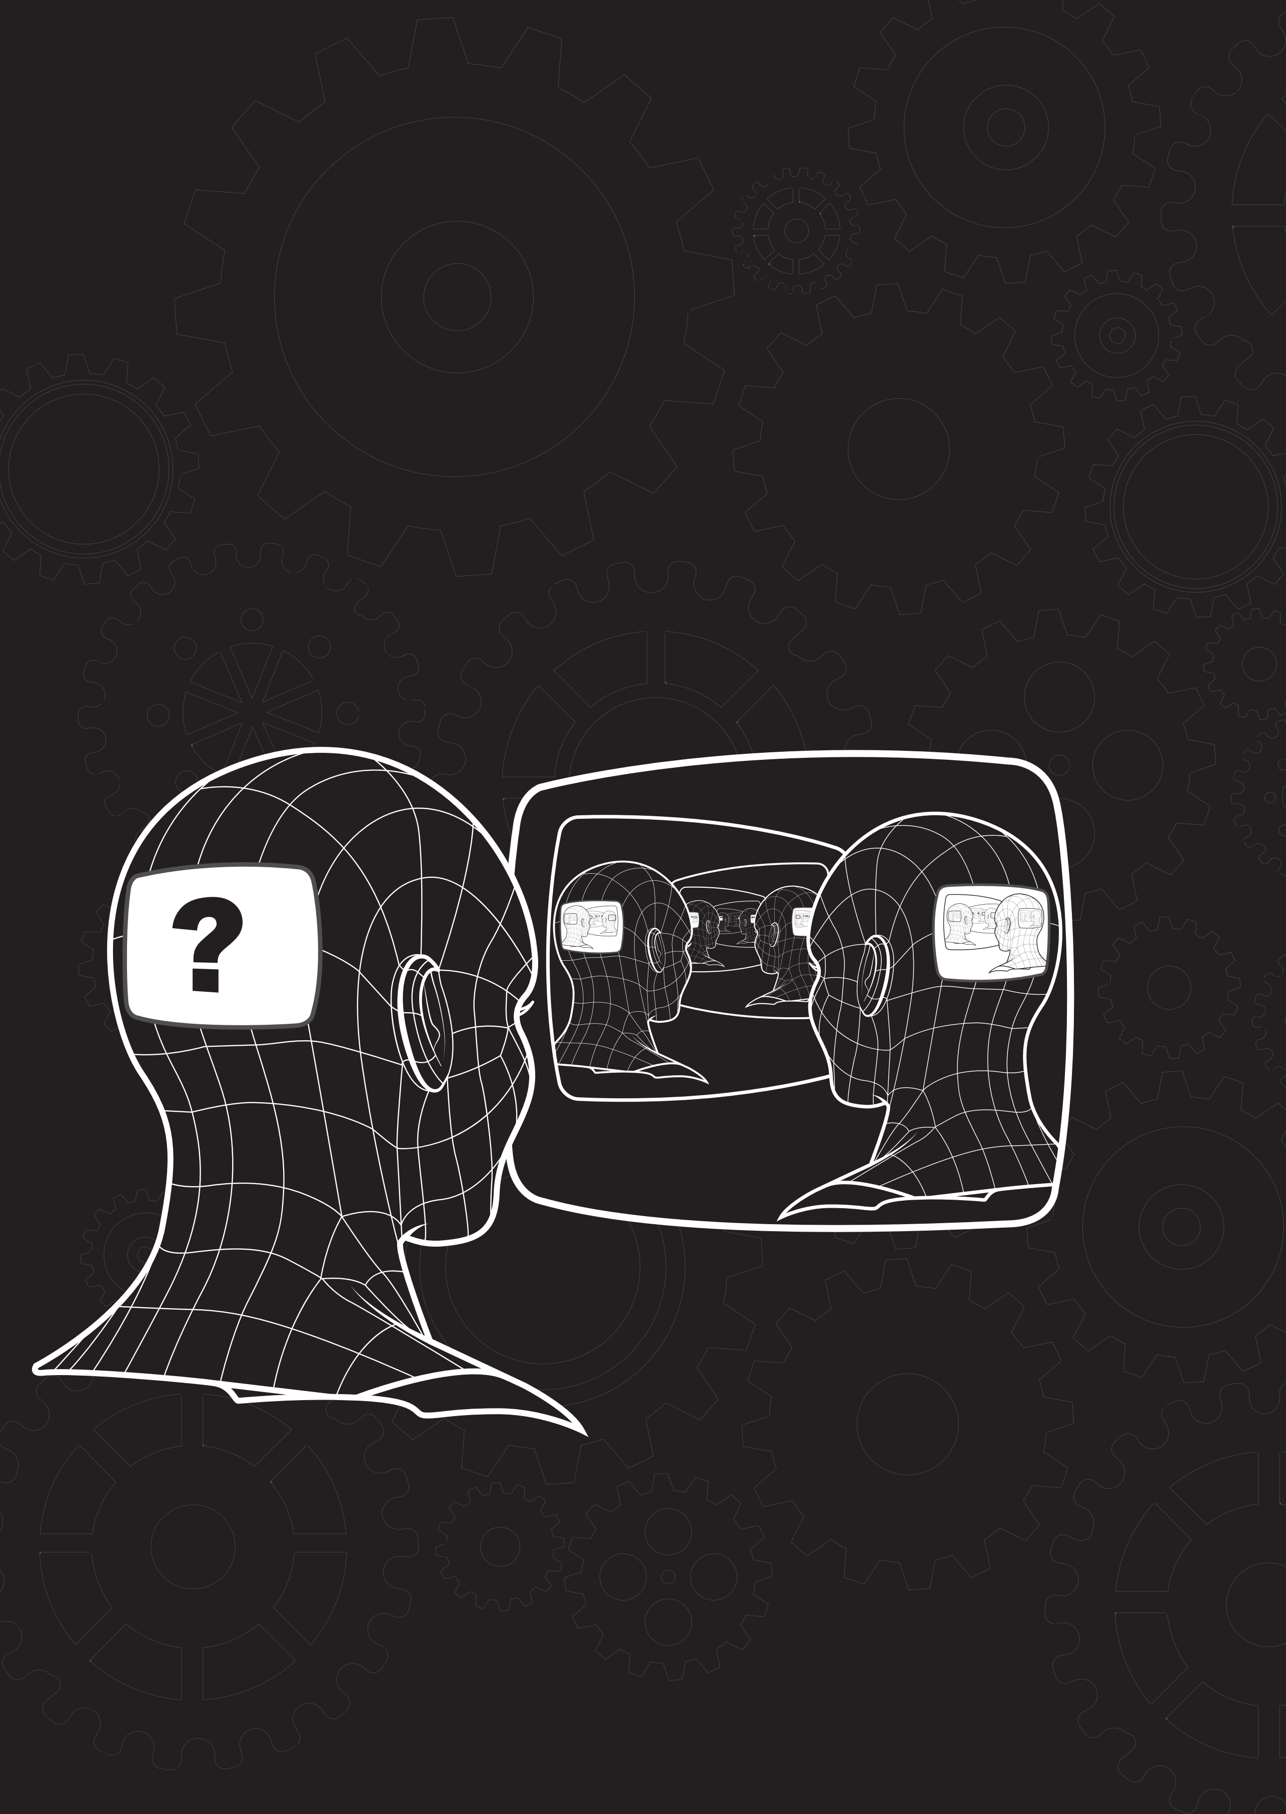

Part I

The scope and limits of non-conscious visual selection

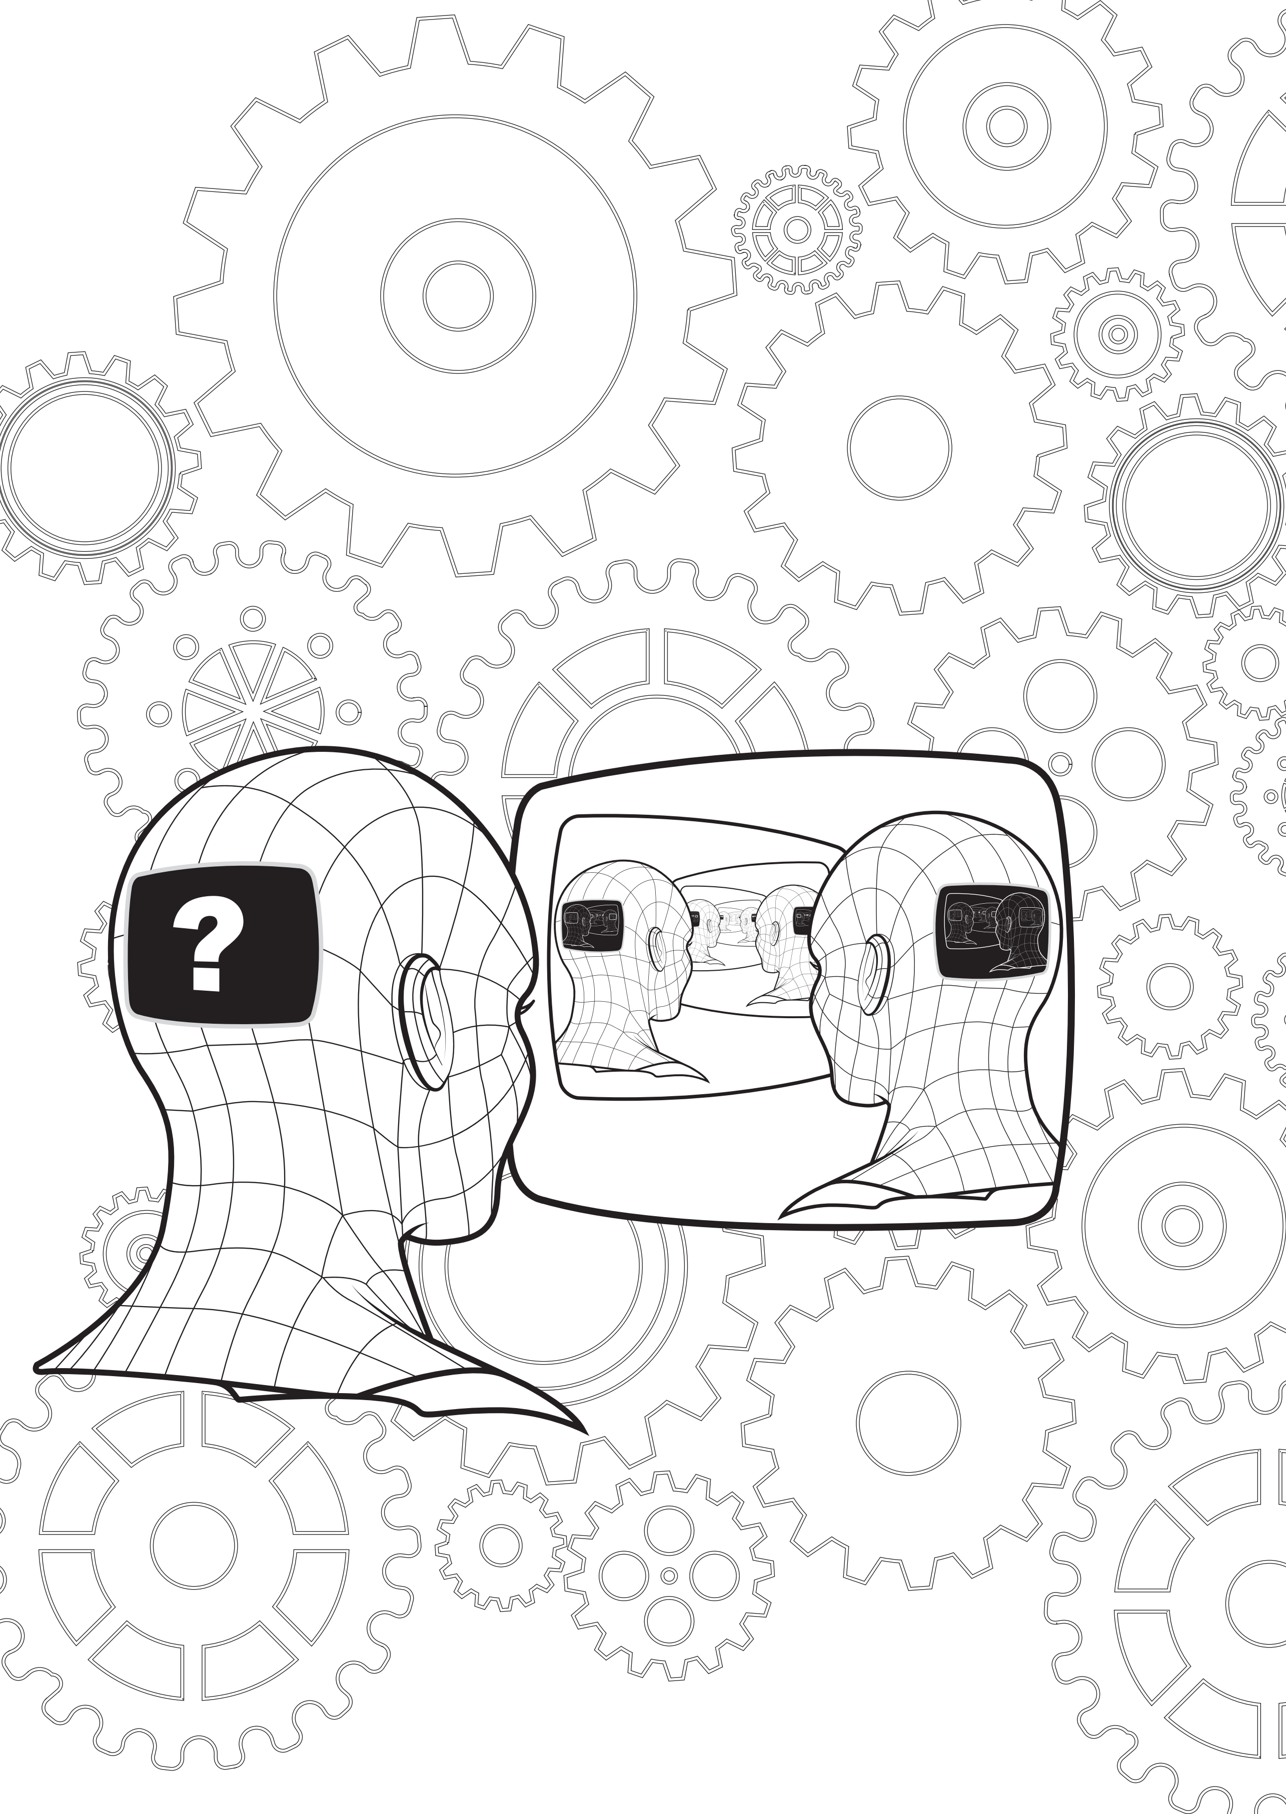

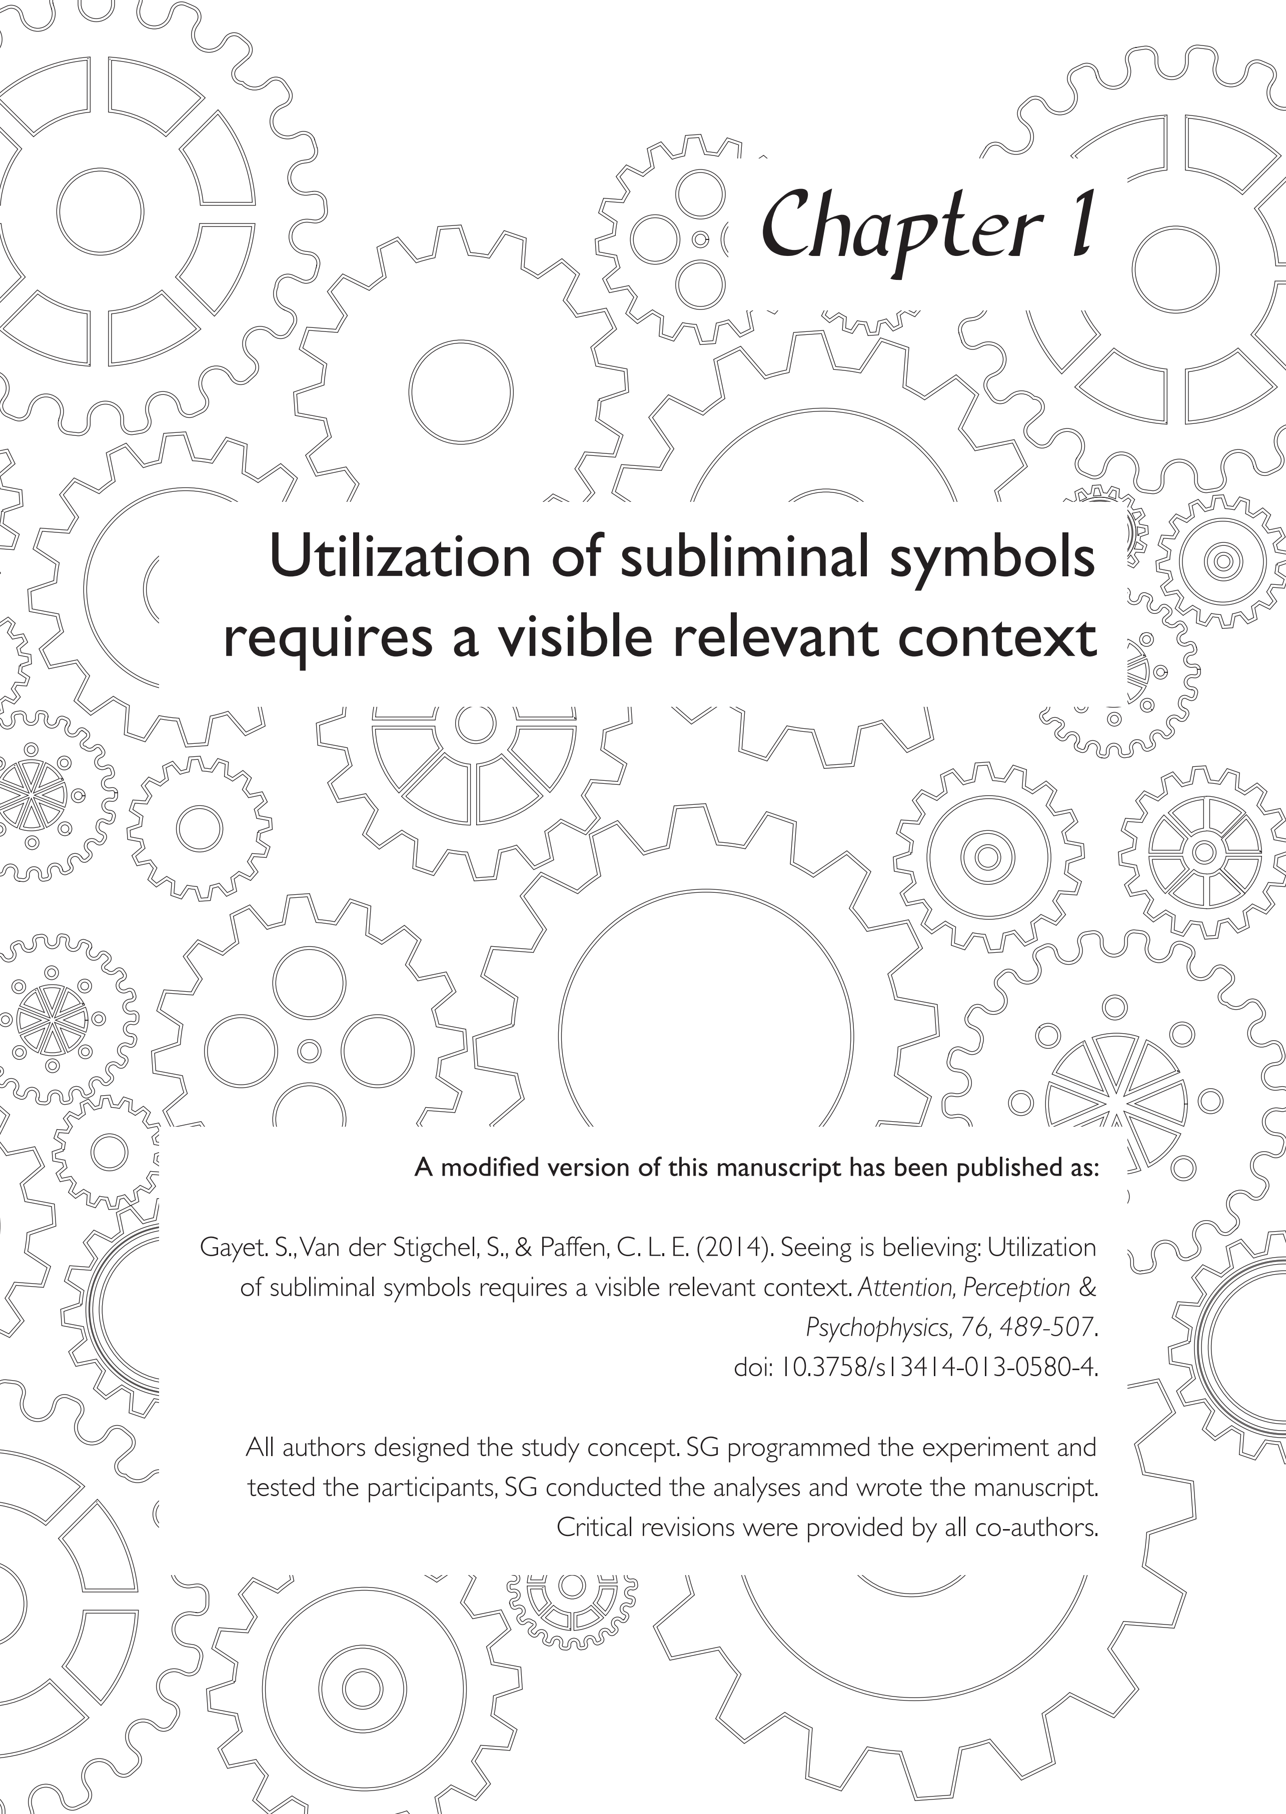

Chapter 1

Utilization of subliminal symbols requires a visible relevant context

A modified version of this manuscript has been published as:

Gayet. S., Van der Stigchel, S., & Paffen, C. L. E. (2014). Seeing is believing: Utilization of subliminal symbols requires a visible relevant context. *Attention, Perception & Psychophysics, 76, 489-507*.
doi: 10.3758/s13414-013-0580-4.

All authors designed the study concept. SG programmed the experiment and tested the participants, SG conducted the analyses and wrote the manuscript. Critical revisions were provided by all co-authors.

Abstract

Sensory input that is not available for conscious report can still affect our behavior. Recent findings suggest that such subliminal information has the potency to influence behavior in a way that is dependent on the observer's current intentions. Here we investigate whether conscious observation of stimulus relevance provides an incentive for the utilization of non-conscious stimuli. We manipulated the predictive power of directional cues to selectively affect the incentive to utilize them for a subsequent target detection task. Central arrow cues rendered invisible by interocular suppression elicited a facilitatory cueing effect, but only when intermixed with visible arrow cues that were highly predictive with respect to (i.e., 80% congruent with) the subsequent target location. When the visible cues were non-predictive (50% congruent), no subliminal cueing effect was found. An analysis of learning effects corroborates these findings; Cueing effects elicited by both visible and invisible cues increased over the course of the experiment, but only when intermixed visible cues were highly predictive. In a second experiment, we demonstrated that the intrinsic relevance of invisible cues (either 50% or 100% congruent) has no effect on the utilization of visible cues. We conclude that conscious perception is required to make statistical inferences about the relevance of symbolic cues. Once statistical information is extracted consciously, it affects subsequent non-conscious processing in a way that fits the current context. Accordingly, one of the possible functions of consciousness could be to extract general rules out of the conscious information, to provide guidelines for future behavior.

Introduction

The world around us provides us with much sensory input, most of which will fail to reach our conscious experience. Even though this subliminal information is not available to conscious report, it is now widely believed that it can nonetheless influence our behavior (e.g., Neumann, & Klotz, 1994; Eimer, & Schlaghecken, 1998; Klotz, & Neumann, 1999). It is still a matter of debate, however, how extensively this subliminal information is processed (for reviews, see: Kouider, & Dehaene, 2007; Lin, & He, 2009). The traditional view holds that subliminal stimuli can only influence behavior in an acquired, automatic manner, and are insensitive to volitional control (Posner, & Snyder, 1975; Schneider, & Shiffrin, 1977; McCormick, 1997). In this context, the physical properties of a stimulus directly influence behavior. This behavior includes phenomena such as attention shifts towards sudden unperceived onsets (for a review, see: Mulckhuyse, & Theeuwes, 2010) and stimulus-response mappings, where an invisible prime facilitates responding to a subsequent target with perceptual similarities after repeated motor responses to the visible target (Abrams, & Greenwald, 2000; Damian, 2001). Recent lines of research, however, have shown that task relevance and goal directed control settings can modulate the way subliminal stimuli affect behavior (Jaskowski, Skalska, & Verleger, 2003; Schlaghecken, & Eimer, 2004; Ansorge, & Neumann, 2005; Ansorge, & Heumann, 2006; Kiefer, & Martens, 2010). These results suggest that the way in which subliminal information influences behavior is dependent on the current cognitive state of the observer. Hence, effects of prime-target combinations might not be fixed, but rather affect behavior differently, depending on the current task demands.

The potency of task relevance to affect non-conscious processing was recently demonstrated by Ansorge, Horstmann, and Worschech (2010). In this study, four colored stimuli were meta-contrast masked by four visible non-singleton color stimuli. Participants were instructed to locate the target of a particular color and to report whether it was a diamond or a square. The masked stimulus at the target location either did or did not match the subsequent target color and was either a square or a diamond. The results showed that participants responded faster when the prime was valid (i.e., at the same position as the target) as compared with invalid (i.e., at a different position than the target). Interestingly however, this effect was restricted to cues that matched the target color. These findings suggest that the masked color singletons captured attention when they matched the observers' task set (e.g., to 'look for the red target'), but failed to capture (stimulus-driven) attention when they were task irrelevant. Indeed, the masked cues elicited an N2pc – a negative event related potential in the visual cortex contralateral to the location in space where a stimulus is attended (Luck, & Hillyard, 1994) – in task relevant, but not in task irrelevant trials. Hence, top-down incentives affected early processing of non-consciously perceived stimuli.

The *contingent involuntary orienting* hypothesis (Folk, Remmington, & Johnston, 1992) provides a plausible framework for the results described above. Folk and colleagues argue that, under conditions of spatial uncertainty, a stimulus with a feature property that is critical to the performance of the task at hand (e.g., color) will incite involuntary attentional capture. Ansorge et al. showed that this involuntary orienting to task-relevant stimuli can occur even when these stimuli are not consciously perceived. The selective utilization of subliminal stimuli as depending on predetermined conditions can also be framed within the *action trigger* hypothesis (Kunde, Kiesel, & Hoffmann, 2003; Kiesel, Kunde, & Hoffmann, 2007). While Kunde and colleagues focus on motor priming rather than attentional priming, these two frameworks have much in common in terms of primed automaticity. According to the action trigger hypothesis observers build up expectations with regard to specific (visible) stimuli, based on prior experience or task instructions. Subsequently presented matching stimuli, referred to as action triggers, directly activate the corresponding response, irrespective of their conscious identification. In contrast with the framework of Folk et al., however, these stimuli are not required to perceptually match the subsequently presented supraliminal target. Rather, stimuli are categorized on a relevant cognitive dimension, based on the current task demands. Thus, in our interpretation, whenever there is an incentive to use a stimulus, that is, when it forms an action trigger, the stimulus will trigger a response (e.g., attentional capture, eye-movement, motor response, etc.). Conscious knowledge about the context in which a behavior is performed, provides an incentive to utilize visual information. This information, whether consciously perceived or not, then affects behavior in a way that fits the current cognitive context.

In the present study our aim was to investigate whether the utilization of subliminal symbolic cues is dependent on contextual knowledge. To this end, participants performed a peripheral target detection task with either visible or suppressed central arrow cues. While the suppressed arrow cues were always 50%

congruent with, and therefore not predictive for, the subsequent target location, the visible arrow cues were either non-predictive (50% congruent) or highly predictive (80% congruent). In the latter case, the predictive value of the supraliminal (visible) arrow cues provided an incentive to utilize arrow cues, which was expected to result in faster reaction times on congruent as compared with incongruent trials (Posner, 1980). Based on the idea of action triggers, we predicted that the visible arrow cues would form a cognitive context in which an arrow predicts the location of the subsequent target. By this, the arrows become action triggers leading to faster responses on congruent trials even when invisible, that is, when they do *not* predict the subsequent target location.

Our methodology was motivated by two main considerations. First, we aimed to investigate whether prime-target congruency effects of subliminal primes extends to cases without perceptual similarities between the prime and the target. This is important since identical prime-target pairs are known to facilitate responding (Bodner, & Dypvik, 2005; Koechlin, Naccache, Block, & Dehaene, 1999), a phenomenon known as repetition priming. In the aforementioned study by Ansorge et al. (2010), special care was taken to avoid repetition priming by instructing participants to look for either of two colors. Cueing effects were still observed when the cue and the target were of different colors. However, the cue was part of the 'to-be-searched-for' task set: for instance, participants were instructed to look for the blue or red stimulus and subsequently report its shape. In the present study, participants were instructed to react to a peripheral target (requiring spatial information), and were therefore not predisposed to react to the central cue (providing symbolic information). Rather than being perceptually similar, the cue and the target had a semantic (or symbolic) relation. Whether the information contained in the cues appeared to be advantageous for subsequent behavior, depended on the statistical context.

The second and important consideration was that we aimed to show that a subliminal stimulus can affect behavior when it is, in itself, not informative for the task at hand. For this purpose, we manipulated the incentive to utilize the subliminal cues by altering the statistical context in which they were embedded, without altering the predictive value of the subliminal cues themselves. Recently, Reuss, Pohl, Kiesel, and Kunde (2012) conducted a study with intermixed unmasked and metacontrast masked arrow cues. This study revealed a facilitatory cueing effect of 7 ms when cues were predictive, but not when cues were non-predictive. However, the predictive value of both the masked cues and the unmasked cues was manipulated simultaneously. As such, it was unclear whether conscious observation of cue validity was required for non-conscious cue utilization to occur, or whether the validity in itself was enough to promote subliminal cue utilization (see the general discussion for further elaboration on this point). In the present experiment, subliminal arrow cues had no predictive value (i.e., they were 50% congruent with the subsequent target location) in either the predictive or the non-predictive condition. The only factor that varied was the incentive to utilize these cues, based on the predictive value of the intermixed supraliminal arrow cues. Using this method, we aimed to isolate the effect of cognitive context on the utilization of information presented outside of visual awareness.

To render cues invisible, we used flash suppression (Wolfe, 1984), which is derived from the phenomenon of binocular rivalry (for reviews, see: Blake, 2001; Alais, & Blake, 2005). Binocular rivalry occurs when each eye of an observer views a different image (e.g., a car for the left eye and a house for the right eye). In this situation, perception will alternate between the two images. Crucially, the image presented to one eye (the suppressed eye) is erased from perceptual awareness, while still impinging on the retina. Even though invisible, the suppressed image elicits substantial activity in dorsal areas (Fang & He, 2005) and in some cases in ventral areas (Jiang, & He, 2006) or the amygdala (Williams, Morris, McGlone, Abbott, & Mattingley, 2004; Pasley, Mayes, & Schultz, 2004). In terms of behavior, suppressed stimuli can trigger eye movements (Rothkirch, Stein, Sekutowicz, & Sterzer, 2012), direct orientation-specific grasping (Roseboom, & Arnold, 2011) and affect behavior performed on targets presented to the non-suppressed eye (e.g., Stuit, Paffen, van der Smagt, & Verstraten, 2011), despite the complete absence of conscious experience. The potency of suppressed stimuli to affect behavior makes binocular rivalry, and the different methodologies derived from this phenomenon, an ideal tool for uncovering the mechanisms of conscious and non-conscious visual processing (Koch, 2004).

This masking method has three major advantages over backward masking, a widely used method to block stimuli from visual awareness. First, the total retinal input between the supraliminal and the subliminal condition can be kept more similar; in backward masking either an extra stimulus (the mask) has to be added, or the onset asynchrony between the mask and the target has to be shortened in the subliminal condition, as compared with the supraliminal condition. Although a (contra-ocular) stimulus is needed to suppress a target with flash suppression as well, this 'suppressor' can be presented to the background (ipsi-ocular) in supraliminal trials. Hence, the presentation chronology is constant between visibility conditions. Second, the subliminal cue can be presented for up to several seconds without reaching awareness, compared to less than 50 ms with backward masking (e.g., Breitmeyer, 2007). Third, in most models, interocular suppression is accounted for by reciprocal inhibition of eye-selective channels (Tong, 2001; Tong, Meng, & Blake, 2006; Blake, 1989). As such, the feedback of information to lower visual areas remains largely uninterrupted under interocular suppression (Macknik, & Martinez-Conde, 2004), whereas it is disrupted by backward masking (Di Lollo, Enns, & Rensink, 2000; Lamme, & Roelfsema, 2000). Disrupting these feedback connections might limit processing of subliminal stimuli to simpler visuo-motor processes, for which only feedforward connections are needed (Lamme, 2001). Similarly, while CFS largely spares V1 spiking responses to suppressed stimuli (Wilke, Logothetis, & Leopold, 2006), backward masking interferes more with V1 spiking activity (Macknick, & Martinez-Conde, 2004). As such, backward masked stimuli elicit a weaker signal in early visual cortex than interocularly suppressed stimuli, thereby reducing its capacity to affect subsequent behavior (Bargh, & Morsella, 2008; Hassin, 2013; Sweeny, Grabowecky, & Suzuki, 2011).

Recently, Al-Janabi, and Finkbeiner (2012) conducted a similar experiment using eye-gaze cues and showed that backward masked gaze cues only produced a cueing effect when they appeared in the context of highly valid unmasked gaze cues. When intermixed unmasked cues were non-predictive, no cueing effect was observed for the masked cues. Although gaze cues and arrow cues both affect attention orienting when they are uninformative of the subsequent target location (Ristic, Friesen, & Kingstone, 2002), gaze cues seem to draw on the reflexive component to a greater extent than arrow cues, the latter being more strongly influenced by volitional control. For instance, the more reflexive nature of gaze cues is supported by the finding that counter-predictive gaze cues still elicit validity effects in the cued direction, whereas counter-predictive arrow cues do not (Friesen, Ristic, & Kingstone, 2004). Another argument comes from studies on inhibition of return (IOR), a phenomenon known to occur exclusively when attention is captured exogenously (Godijn, & Theeuwes, 2004; Klein, 2000; Posner, & Cohen, 1984). Under conditions of brief exposure, gaze cues can elicit IOR (Frischen, & Tipper, 2004; Frischen, Smilek, Eastwood, & Tipper, 2007). In contrast, more endogenous cues, such as arrows, do not elicit IOR (Posner, & Cohen, 1984; Rafal, Calabresi, Brennan, & Sciolto, 1989). Furthermore, uninformative central arrows shorten detection times on subsequent targets to a greater extent when they are presented as task relevant rather than task irrelevant cues (e.g., Pratt, & Hommel, 2003). Within the same experimental set-up, however, the task relevance assigned to gaze cues does not affect subsequent target detection times (Ristic, Wright, & Kingstone, 2007). Hence, in contrast to arrow cues, gaze cues are not influenced by top-down volitional control. Accordingly, neuroimaging studies revealed that attentional orienting induced by arrow cues is more related to dorsal fronto-parietal systems (Corbetta, & Shulman, 2002; Hietanen, Leppänen, Nummenmaa, & Astikainen, 2008) associated with voluntary shifts of attention, whereas gaze induced attentional orienting relies more on ventral fronto-parietal (Corbetta, & Shulman, 2002; Hietanen et al., 2008) and oculomotor (Nummenmaa, & Hietanen, 2006) systems associated with involuntary shifts of attention. These studies have therefore shown that gaze cueing is based on a strong exogenous component, which is in contrast to the effect of arrow cues, which are strongly influenced by volitional control. Whereas a wealth of studies has shown cueing effects of masked exogenous cues, the present study aimed to test whether a predominantly endogenous cue, which requires intention and interpretation, affects subsequent target detection when it is not consciously perceived.

Experiment 1

The aim of Experiment 1 was to investigate whether symbolic cues can elicit cueing effects in the absence of visual awareness, specifically when current task settings provide an incentive for cue utilization. In the supraliminal condition, in which the arrow cues were visible, it was expected that both highly predictive (e.g., Jonides, 1981; Yantis, & Jonides, 1990; Folk et al., 1992), and to a lesser extent non-predictive (Hommel, Pratt, Colzato, & Godijn, 2001; Tipples, 2002) arrow cues would provide facilitation in a target-detection task. Subliminally presented arrow cues however, were expected to facilitate target detection only if the supraliminal arrow cues were sufficiently informative. This would show that subliminal cue utilization is selective to situations in which the cognitive context provides an incentive to use them. Furthermore, we investigated the time course within which cueing effects tend to increase or decrease as a function of the (visible) informative value of the cues.

Subliminal cues were kept outside of perceptual awareness using (interocular) flash suppression, in which a high-contrast mask presented to one eye suppresses awareness of the stimulus presented to the other eye. Recent studies using various forms of interocular suppression have shown that non-consciously presented monocular (peripheral) cues can affect response times to visible target stimuli in the non-suppressed eye (e.g., Jiang, Costello, Fang, Huang, & He, 2006; Zhaoping, 2008; Self, & Roelfsema, 2010), making it an ideal method to elicit cueing effects with invisible stimuli.

The cueing method was based on that of Hommel et al. (2001) who used isosceles triangles with a 2:1 length-to-width ratio acting as central arrow cues. Leftward and rightward pointing triangles were superimposed to form an elongated ‘star of David’ on the side (see Figure 1). At cue onset, three portions of the overlapping triangles were removed, such that a single arrow cue remained. Both triangles intersected at the centroid on the center of fixation, such that all triangle stimuli (star and arrows) had equal perceptual weight on both sides of the fixation cross. This ensured that cueing effects would originate from the symbolic meaning of the arrow cue, rather than reflect an exogenous shift of attention caused by a change in the center of mass.

The star stimulus was presented to one eye, to initiate ocular dominance and enhance the strength of the subsequent flash suppressor presented to the other eye (Tsuchiya, Koch, Gilroy, & Blake, 2006). The target was presented to the eye that was made dominant by means of the flash suppression, since pop-out stimuli in a specific eye during binocular rivalry are known to enhance dominance of the ipsi-ocular percept (Ooi, & He, 1999). The complete sequence of events in a trial is depicted in Figure 1.

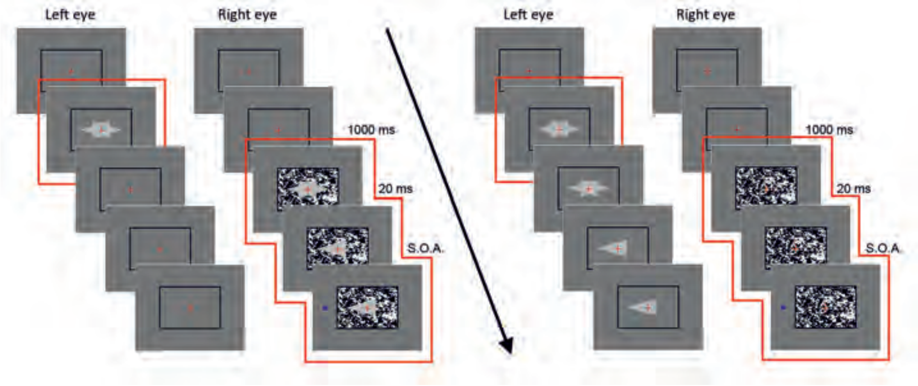

Figure 1. Sequence of events in supraliminal (left) and subliminal trials (right). The red line indicates the dominant percept. Target and cue remained on screen until a response was given.

Methods

Participants

Nineteen students from Utrecht University, age 19 to 36 (*M* = 23.4 years, *SD* = 4.36), signed informed consent before participating for course credits or payment. All participants were right-handed, had normal or corrected to normal vision and were tested for stereoscopic vision (TNO test for stereoscopic vision, 12th edition; Laméris Ootech b.v., 1972).

Apparatus and stimuli

All experiments were conducted in a dimly lit room, using an Apple dual 2-GHz PowerPC G5, fitted with a linearized 22" LaCie Electron blue IV CRT monitor (1024 x 768; 100 Hz) and an Apple keyboard, which was used for response registration. Stimulus presentation and response collection were managed using the Psychophysics Toolbox 3 (Brainard, 1997; Pelli, 1997) in MATLAB R2010a (The Mathworks, Natick, MA). The Participant's head was supported by a chinrest, on which a dichoptic mirror set-up was mounted. This resulted in an effective viewing distance of 57 cm.

All stimuli were presented on a uniform gray background with a luminance of 28.6 cd/m^2. In order to facilitate binocular fusion of the two complementary images, a black frame (91% Weber contrast) that subtended a visual angle of 5.8° × 5.8° was presented to each eye during the entire experiment. Flash suppression masks were created by (1) filtering pink (1/f) noise by a rotationally symmetric Gaussian low pass filter (sigma = 0.3) and by (2) making the resulting grayscale image binary with maximum contrast (91% Michelson). A new mask was generated on every trial. The gray arrow cue (luminance of 32.7 cd/m^2 and length of 2.3° visual angle) had a Weber contrast of 14% with the background and of 29% with the mask. The blue target stimulus (8.27 cd/m^2, $x = 0.168$, $y = 0.094$) had a radius of 0.2° of visual angle, a Weber contrast of 71%, and an eccentricity of 3.8° of visual angle with respect to the red fixation cross (10.4 cd/m^2, $x = 0.590$, $y = 0.357$).

Design and procedure

All participants took part in two experimental conditions (with non-predictive and with predictive visible cues) on different days and a control experiment at the end of the second day. The order of conditions was counterbalanced across participants. All experimental factors (visibility of the cue, predictive value of the cue, stimulus onset asynchrony) and balancing conditions (suppressed eye, target location) were fully counterbalanced within participants. The number of trials per condition is given in Table 1.

Table 1. *Overview of trials per condition for each stimulus onset asynchrony.*

Condition	Non-predictive				Predictive			
Cue visibility	Visible		Invisible		Visible		Invisible	
	240		80		240		80	
Congruency	Con.	Incon.	Con.	Incon.	Con.	Incon.	Con.	Incon.
	120	120	40	40	192	48	40	40

The experiments consisted of a target detection task, for which participants were instructed to respond as fast and accurately as possible to the location of a blue dot, by means of the left and right arrow keys. This target appeared either to the left or to the right of fixation and was preceded by a central arrow cue. Participants were informed about the predictive value of visible arrow cues, but were naïve with respect to the existence of subliminal arrow cues. The visible arrow cues (75% of the trials) were either non-predictive (50% congruent) or highly predictive (80% congruent), whereas the suppressed arrow cues (25% of the trials) were always non-predictive. The main purpose of the visible trials was to form a statistical incentive to utilize arrow cues, by creating an experimental context in which arrow cues appeared either predictive of the subsequent target location or not. Therefore, the trials that were used for the manipulation of statistical information outnumbered the (non-predictive) subliminal trials.

On both sessions participants first took part in a 24 trial practice block before starting the six experimental blocks of 160 trials each. After the second session participants performed an additional 120 trials in a control experiment that verified whether masked cues were indeed invisible. Stimulus presentation and counterbalancing was identical to the experimental task, but only subliminal trials were performed. Participants were explained that they would see a mask, and that after a while an arrow would appear. They were instructed to detect as fast and accurately as possible the direction in which the arrow pointed, and report it by means of the left and right arrow keys. A trial was aborted whenever participants failed to respond within six seconds.

Temporal differences in subliminal and supraliminal processing of emotional stimuli (e.g., Lidell, Williams, Rathjen, Shevrin, & Gordon, 2004; Williams et al., 2004) and between endogenous and exogenous cueing (e.g., Egeth, & Yantis, 1997; Mulckhuyse, & Theeuwes, 2010) suggest that subliminal stimuli might be processed in a different time course than supraliminal stimuli. Therefore different onset asynchronies (SOA) between cue and target stimulus were implemented: 100 ms, 500 ms and 900 ms.

Trial selection and inclusion criteria

Determining the effects of visual processing outside of awareness requires dissociating subliminal and supraliminal stimuli. This should be done with special care when, like in the present experiment, the subliminal stimuli are expected to influence behavior in the same direction as the supraliminal stimuli (i.e., a quantitative rather than qualitative dissociation). Therefore a number of precautions have been taken to make sure a genuine subliminal effect was measured.

First, since suppression durations in flash suppression are relatively short-lived, it is important to verify that the flash suppressed cues remained unseen until the target response is initiated. Thus, for each participant, suppression durations of subliminally presented cues should be assessed, and compared to the participant's RTs on subliminal trials. For this purpose, a particularly stringent control task (discussed in section 2.1.3. and further elaborated in the general discussion) was implemented in which participants were required to report the direction of the *cue* as fast as possible, rather than the location of the target. As cues preceded targets in the main experiment, longer RTs in response to suppressed cues in the control experiment compared to visible targets in the main experiment would indicate that the subliminal cues in the main experiment were not consciously perceived. Subliminal trials with the slowest reaction times in the experiment pose a bigger threat of being 'infected' with consciously perceived arrow cues. Therefore, for each participant, RTs in the main experiment (target detection) slower than the 5th percentile RT in the control task (i.e. the 5% fastest RTs in cue detection) were removed (the rightmost gray area in Figure 2). The onset of a monocular stimulus, in this case the target, can affect the duration of interocular suppression (Ooi, & He, 1999). Therefore, the threshold was calculated for each SOA condition separately.

Second, to retain comparable RT distributions between participants, all participants for which this procedure would result in discarding more than 5% of their trials (i.e., who might have perceived more than 5% of the cues in the subliminal condition) were excluded from further analyses.

Third, when analyzing the presence of cueing effects (i.e., the difference between RTs on congruent trials and RTs on incongruent trials) in the subliminal conditions we took into account the fraction of potentially unsuccessfully suppressed cues (i.e., the black area in Figure 2). This fraction was based on the RTs in the control experiment (cue detection). For instance, if all RTs in the cue detection task were slower than the slowest RT in the target detection task (subliminal condition) this gives an estimated 0% 'infected' cues. Accordingly, if the overlap was larger than 5%, this gives an estimated 5% 'infected' cues. Note that 5% is the maximum, since all experimental

RTs slower than this threshold were removed. These cues, for which suppression possibly failed, can be expected to elicit a cueing effect that is, at most, similar in magnitude to cueing effects of supraliminally presented cues. To compensate for this spurious effect, we subtracted the cueing effect of supraliminal cues, multiplied by the estimated fraction of unsuccessfully suppressed cues, from the cueing effect elicited by subliminal cues. This was done on an individual participant basis. For instance, one (purely illustrative) participant had an estimated 3% trials with unsuccessfully suppressed cues in the SOA 500 condition. Computation of the participant's median RTs revealed a cueing effect (incongruent versus congruent RT) of 70 ms in the SOA 500 supraliminal condition and of 15 ms in the SOA 500 subliminal condition. Thus, this participant's subliminal cueing effect in the SOA 500 condition was corrected to 15 – (70 * 0.03) = 12.9 ms for the analysis of cueing effects over all participants. As such, if the analyses reveal facilitatory cueing effects of cues in the subliminal condition, this effect could not be the result of 'infected' (i.e., unsuccessfully suppressed) cues.

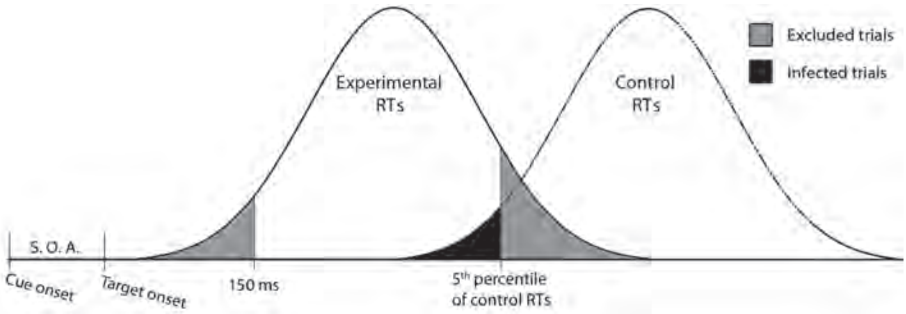

Figure 2. Schematic representation of RT distributions on subliminal trials. Experimental RTs reflect the target detection times on subliminal trials in the experimental phase, whereas Control RTs reflect the detection time of arrow directions in the control task.

Results

Analyses

The results section is comprised of an analysis of the control task, an analysis of cueing effects, an analysis of learning effects and analyses of reaction times and accuracy. All analyses are conducted using IBM SPSS Statistics 19 and start with a full factorial repeated measures ANOVA. When three-way interactions are found, subsequent ANOVAs are conducted for each SOA condition for a better understanding of the mutual relation between factors. If these new analyses reveal interactions, paired-samples t-tests are then used to determine the source of the interactions. Finally, one-sample *t*-tests are conducted to assess in which conditions RTs are affected by the factors included in our design (i.e., whether our experimental manipulations affect cueing).

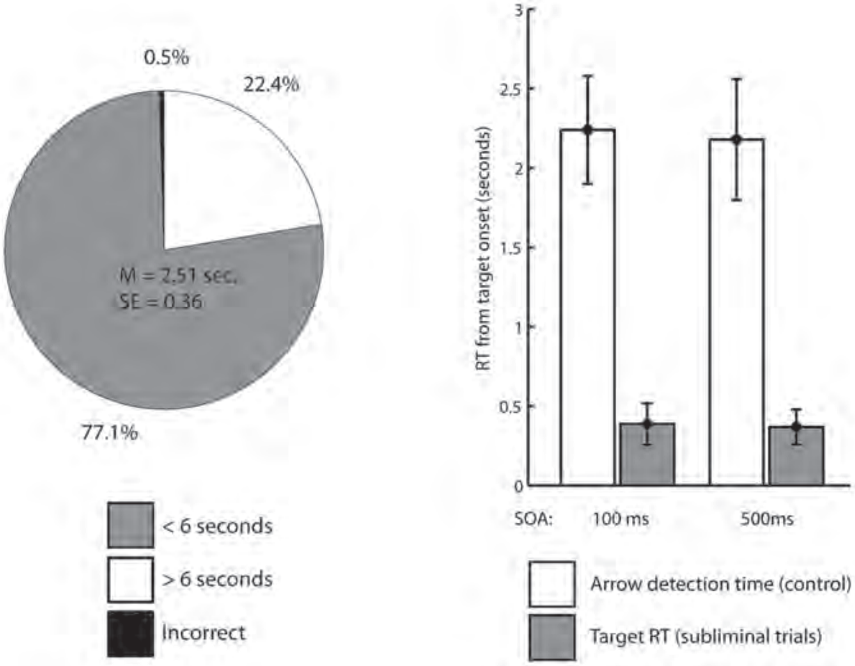

Figure 3. Results of the control task. Left: The cue remained invisible for 6 s in 22.4% of the control trials, while its direction was identified correctly within 2.51 s in 77.1%. Right: correct arrow detection times in the control task compared with target detection times on subliminal trials.

Control task

The control task was used to assess participants' detection times of the direction of subliminal arrow cues. Comparing cue detection times with reaction times in the experimental blocks gives a conservative estimate of the fraction of subliminally presented cues that were consciously perceived during the experiment. Based on the RTs in the control task, for many participants an SOA of 900 ms appeared to be too long to exclude the possibility that some suppressed arrow cues were consciously perceived. This SOA condition was therefore excluded from further analyses. Applying the exclusion criteria on the data of the SOA 100 and SOA 500 condition resulted in the exclusion of 6 out of 19 participants.

The results of the control task are depicted in Figure 3. Participants included in the analysis were unable to detect 22.4% of the arrow directions within six seconds, while 0.5% of the trials in which participants responded within six seconds yielded an incorrect response. Consequently, 77.1% of the cues were accurately detected (M = 2.5 seconds, SD = 1.3) within the time limit. Based on the participants' performance on the control task and their RTs on subliminal trials of the actual experiment (the two distributions in Figure 2), the estimate of consciously perceived arrow cues in the subliminal condition of the experiment (the black area in Figure 2) was 2.1% in the SOA 500 condition and 0.8% in the SOA 100 condition. Mean detection times in the control task did not differ between different SOAs, $t(12) = 0.48$, $p = 0.640$, $d = 0.13$, suggesting comparable suppression length of subliminal cues in both SOAs. Finally, the difference between the blocked presentation of subliminal stimuli, as used in the control experiment, and intermixed presentation (with supraliminal trials intermixed), as used in the main experiments, was found to have no impact on suppression durations (see Supplementary Materials S1).

Cueing effects

Trials in which participants reported an incorrect location of the target (0.9%), trials with RTs faster than 150 ms (0.2%) and trials with RTs slower than the participants 5th percentile RT on the control task (2.9%; varying between 0 and 5% for individual participants) were excluded from RT analyses. Subsequently, cueing effects were computed, by subtracting RTs on congruent trials from RTs on incongruent trials. All cueing effects are depicted in Figure 4. A positive cueing effect reflects facilitation in target detection elicited by the cues (i.e., faster RTs on congruent as compared with incongruent trials). Conversely, negative cueing effects would reflect inhibition. Before analyzing which conditions led to significant cueing effects, we first assess whether the difference in RT between congruent and incongruent trials is different for different SOAs. An overall factorial 2 × 2 × 2 repeated measures ANOVA was conducted, with the factors SOA (100 ms and 500 ms), cue visibility (subliminal and supraliminal) and cue relevance (predictive and non-predictive). This analysis revealed main effects of SOA, $F(1, 12) = 24.05$, $p < 0.001$, $\eta^2 = 0.67$, Relevance, $F(1, 12) = 21.52$, $p = 0.001$, $\eta^2 = .64$, and Visibility, $F(1, 12) = 58.75$, $p < 0.001$, $\eta^2 = 0.83$. These effects reflected that longer SOA's elicited a stronger cueing effect than shorter SOA's, highly predictive cues elicited a stronger cueing effect than non-predictive cues and visible cues elicited a stronger cueing effect than subliminal cues.

There was an interaction effect between SOA and Relevance, $F(1, 12) = 9.35$, $p = 0.010$, $\eta^2 = 0.44$, between SOA and Visibility, $F(1, 12) = 18.42$, $p = 0.001$, $\eta^2 = .61$, and between Relevance and Visibility, $F(1, 12) = 66.03$, $p < 0.001$, $\eta^2 = 0.85$, showing that cueing effects depended on the factors included in this design: predictive value of the visible cues, visibility of the cues and SOA. To further examine these effects subsequent repeated measure analyses were conducted for the 100 ms SOA and 500 ms SOA conditions separately, with cue relevance and cue visibility as within subject factors.

SOA 100 condition

The data from the SOA 100 condition is depicted on the left part of Figure 4. There was a main effect of cue visibility, such that supraliminal cues elicited more facilitation than subliminal cues, $F(1, 12) = 12.11$, $p = 0.005$, $\eta^2 = 0.50$. Although there was no significant main effect of cue relevance, $F(1, 12) = 3.24$, $p = 0.097$, $\eta^2 = 0.21$, an interaction was found between cue relevance and cue visibility, $F(1, 12) = 6.10$, $p = 0.029$, $\eta^2 = 0.34$, such that in the supraliminal condition predictive arrow cues ($M = 32.0$ ms, $SD = 23.4$) elicited stronger cueing effects than non-predictive arrow cues ($M = 9.6$ ms, $SD = 5.8$), $t(12) = 3.25$, $p = 0.007$, d = 0.90, whereas cueing effects onsubliminal trials did not differ between Relevance conditions, $t(12) = -0.08$, $p = 0.938$, $d = -0.02$. These findings show that even with SOAs as short as 100 ms the predictive value of the cue affects supraliminal cue utilization.

Additional one-sample t-tests revealed that supraliminal cues in the predictive, $t(12) = 4.94$, $p < 0.001$, $d = 1.37$, and non-predictive, $t(12) = 5.98$, $p < 0.001$, $d = 1.66$, condition facilitated response RTs, whereas subliminal cues in the predictive, $t(12) = 0.54$, $p = 0.60$, $d = 0.15$, and non-predictive, $t(12) = 0.85$, $p = 0.41$, $d = 0.23$, conditions did not elicit facilitation. Thus, with an SOA of 100 ms subliminal stimuli did not affect behavior, irrespective of the predictive value of the subliminal cues.

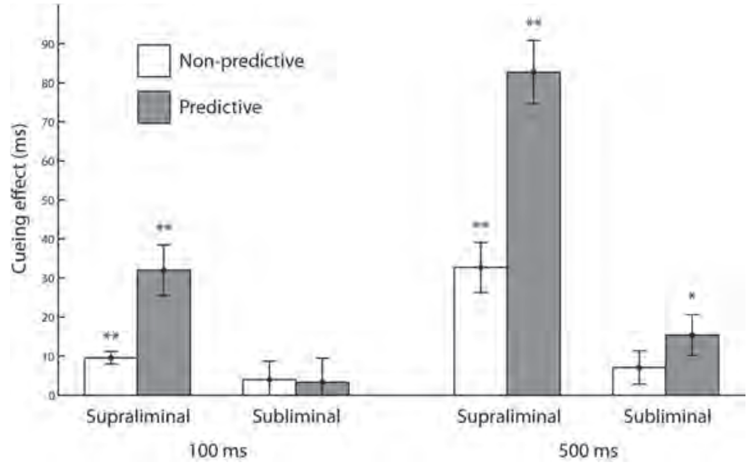

Figure 4. Cueing effects in all experimental conditions. Cueing effects were calculated by subtracting RTs on trials with congruent cues from RTs on trials with incongruent cues. Positive cueing effects reflect facilitation. Error bars depict the standard error of the mean. *$p < .05$, **$p < .005$.

SOA 500 condition

Another repeated measures ANOVA with Visibility and Relevance as within-subject factors revealed that in the 500 ms SOA condition supraliminal presentation led to stronger cueing effects than subliminal presentation, $F(1, 12) = 65.14$, $p < 0.001$, $\eta^2 = 0.84$, and predictive cues led to stronger cueing effects than non-predictive cues, $F(1, 12) = 45.13$, $p < 0.001$, $\eta^2 = 0.79$. Also, these two factors interacted, $F(1, 12) = 38.06$, $p < 0.001$, $\eta^2 = 0.76$, such that supraliminal cues elicited stronger cueing effects in the predictive as compared with the non-predictive condition, $t(12) = 7.65$, $p < 0.001$, $d = 2.12$, whereas cueing effects on subliminal cues did not differ in magnitude between relevance conditions, $t(12) = 1.97$, $p = 0.073$, $d = 0.55$.

Importantly, subsequent one sample t-tests revealed that supraliminal cues elicited facilitation when they were both predictive ($M = 82.7$ ms, $SD = 29.4$), $t(12) = 8.86$, $p < 0.001$, d = 2.81, and non-predictive ($M = 32.7$ ms, $SD = 23.1$), $t(12) = 5.04$, $p < 0.001$, d =1.42. In the subliminal condition however, the predictive cues ($M = 15.6$ ms, $SD = 18.7$) elicited facilitation, $t(12) = 2.95$, $p = 0.012$, $d = 0.82$, whereas the non-predictive cues ($M = 7.1$ ms, $SD = 15.5$) did not, $t(12) = 1.66$, $p = 0.123$, $d = 0.46$. These effects show that subliminal arrow cues only elicit facilitation when intermixed supraliminal cues reliably predict the subsequent target location. Note that all t-tests of cueing effects in and between subliminal conditions were performed using the corrected subliminal cueing effects (see section 2.1.4. for a description of this correction method). These corrections consisted of a reduction of cueing effect between 0 and 3.1 ms for individual participants (*M* = 1.9 ms, *SD* = 1.1) in the subliminal predictive condition and between 0 and 1.1 ms for individual participants (*M* = 0.5, *SD* = 0.4) in the subliminal non-predictive condition.

Learning effects

The analysis of cueing effects revealed a significant effect of subliminal facilitation in the predictive condition, whereas no such effects approached significance in the non-predictive condition. This cueing effect might have come about by task instructions that elicited a top-down incentive to either utilize or ignore arrow cues. If this were the case, cueing effects in subliminal conditions would have been present from the start and remained constant throughout the entire experiment. Alternatively, the utilization of arrow cues might have been based upon statistical evidence of the predictive value of the cue, which is accumulated throughout the experiment. The latter predicts that cueing effects in subliminal trials would increase throughout the experiment. To dissociate between these two possible origins of subliminal cueing effects, we analyzed whether cueing remained constant or increased throughout the course of the experiments. That is, we assessed whether learning occurred within an experimental session, as a function of cue relevance.

Due to the large number of counterbalanced factors, not all conditions were equally represented in each block. Rather, the whole design was repeated twice, as a result of which the first half (block 1:3) and the second half (block 4:6) were fully counterbalanced. Therefore, cueing effects were compared between the first half and the second half of both experimental sessions, by computing learning effects (i.e., cueing effect in block 1:3 subtracted from cueing effect in block 4:6). These are depicted in Figure 5.

In the 100 ms SOA condition, no learning effects were observed: for supraliminal trials neither the predictive condition, $t(12) = 0.29$, $p = 0.779$, $d = 0.08$, nor the non-predictive condition, $t(12) = -0.66$, $p = 0.523$, $d = -0.18$, yielded significant learning effects. Similarly, for subliminal trials neither the predictive, $t(12) = -0.45$, $p = 0.659$, $d = -0.13$, nor the non-predictive condition, $t(12) = -0.71$, $p = 0.493$, $d = -0.20$, yielded learning effects. In the 500 ms SOA condition, however, there was a main effect of Relevance, $F(1, 12) = 22.02$, $p = 0.001$, $\eta^2 = 0.65$, such that stronger learning effects were observed in the predictive as compared with the non-predictive condition. This learning effect did not depend on cue visibility, as there was no main effect of Visibility, *F*(1, 12) = 1.27, $p = 0.282$, $\eta^2 = 0.10$, nor was there an interaction between Visibility and Relevance on learning effect, *F*(1, 12) = 2.55, $p = 0.136$, $\eta^2 = 0.18$. Indeed, subsequent paired samples t-tests revealed that, for supraliminal trials, the learning effect was larger in the predictive compared to the non-predictive condition, $t(12) = 5.04$, $p < 0.001$, $d = 1.40$. Importantly, this was also true for subliminal trials, $t(12) = 2.74$, $p = 0.018$, $d = 0.76$. These findings show that both supraliminal and subliminal cue utilization depend on statistical relevance, as deduced from perceptual evidence (i.e., predictive visible cues).

Finally, to test for learning effects between experimental sessions, we examined whether the order in which participants took part in the non-predictive and the predictive experimental session influenced cueing effects. Session Order was entered as a between-subjects factor in a repeated measures ANOVA with the within-subject factors Visibility, Relevance and SOA. None of the interactions between Session Order and the experimental within-subject manipulations approached significance ($p > 0.6$).

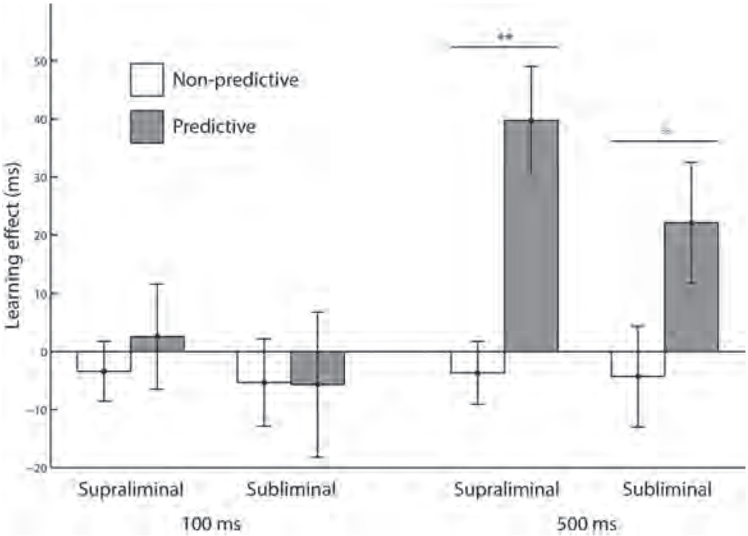

Figure 5. Learning effects in experiment 1. Learning effects were quantified as the difference in cueing effects between the first half and the second half of an experimental session. Error bars represent the standard error of the mean.

Reaction times and accuracy

The analyses of reaction times and accuracy (see Supplementary Materials S3, Table S1) are implemented to control for unexpected inconsistencies, reflecting differing response strategies between conditions. Mean reaction time on the target detection task across all experimental conditions was 389 ms (*SD* = 44.9). A repeated measures ANOVA with SOA, Relevance, Visibility and Congruency as within-subject factors revealed a main effect of SOA, $F(1, 12) = 19.83$, $p = 0.001$, $\eta^2 = 0.62$, reflecting that reaction times were faster when cues were presented 500 ms ($M = 402$ ms, $SD = 49.8$) as compared with 100 ms ($M = 377$ ms, $SD = 41.8$) before target onset. This result is likely to reflect a preparation for response execution prior to stimulus delivery (Sommer, Leuthold, & Schubert, 2001). The main effect of cue visibility was also significant, $F(1, 12) = 5.69$, $p = 0.034$, $\eta^2 = 0.32$, showing that target detection was slightly faster after supraliminal cues ($M = 387$ ms, $SD = 44.7$) as compared with subliminal cues ($M = 391$ ms, $SD = 45.4$). This latter finding probably reflects the greater cueing effects in supraliminal as compared with subliminal trials.

Mean accuracy on the target detection task over all experimental conditions was 98.8% (*SD* = 1.0%) correct responses, with a minimum of 97.3% and a maximum of 100% for individual participants. We conducted 2 × 2 × 2 repeated measure ANOVA's on the accuracy (percentage of incorrect responses on target detection) with the factors Relevance, Visibility and Congruency for each SOA condition in Experiment 1. This revealed a three-way interaction in the SOA 100 condition, $F(1, 12) = 6.65$, $p = 0.024$, $\eta^2 = 0.36$. Subsequent paired samples t-tests clarified this interaction by showing that, in the predictive condition, more errors were made after visible incongruent cues than after visible congruent cues $t(12) = 2.97$, $p = 0.012$, $d = 0.82$, whereas no such difference was observed for invisible cues $t(12) = 0.00$, $p = 1.00$, $d = 0.00$. In the non-predictive condition, however, no difference in accuracy was observed after congruent as compared with incongruent cues, in either supraliminal, $t(12) = -0.18$, $p = 0.861$, $d = -0.05$, or subliminal trials, $t(12) = 0.37$, $p = 0.721$, $d = 0.10$. In the SOA 500 condition the same three-way interaction was found, $F(1, 12) = 5.430$, $p = 0.038$, $\eta^2 = 0.31$. Again, this interaction was carried by the predictive condition, in which more errors were made after visible incongruent cues than after visible congruent cues, $t(12) = 2.65$, $p = 0.021$, $d = 0.49$, whereas no such difference was observed for invisible cues, $t(12) = -0.37$, $p = 0.721$, $d = -0.10$. In the non-predictive condition no difference in accuracy was observed after congruent as compared with incongruent cues, in either supraliminal, $t(12) =$

1.50, $p = 0.160$, $d = 0.42$, or subliminal trials, $t(12) = -1.00$, $p = 0.337$, $d = 0.28$. The finding that subliminal cues in the predictive condition did not affect accuracy, even though they did affect RTs is in line with the idea that subliminal signals are too weak to trigger a new behavior, but are strong enough to affect the time course of an initiated behavior (e.g., Van Gaal, Lamme, Fahrenfort, & Ridderinkhof, 2011). Finally, all observed effects of accuracy (reported above), reflected qualitatively similar differences of performance as the analyses of cueing effects (reported earlier); faster responses were accompanied by higher accuracy and vice versa. As such, there was no indication for a speed-accuracy trade-off (see Supplementary Materials S3, Table S1, for a complete overview of all RTs and accuracy per condition).

Discussion

The results of Experiment 1 show that visible arrow cues facilitated target detection in the cued location, irrespective of the predictive value of these arrow cues. Facilitation was observed even when an arrow appeared 100 ms before target onset (see General Discussion). Importantly, subliminal arrow cues also produced cueing effects, but only when supraliminal arrows predicted the location of the target. Also, cueing on subliminal trials was observed only when cues appeared 500 ms before target onset.

Some caution is required when interpreting these results, since it is impossible to disambiguate whether the null-result in the non-predictive condition reflects a genuine absence of cueing effects, or a lack of statistical power. This is especially relevant, since the subliminal cueing effects only marginally differed between relevance conditions. As such, a post-hoc power analysis was conducted for these two conditions, based on the effect size, the number of participants and the 0.05 (two-sided) alpha level. This revealed that the present experiment could detect the subliminal cueing effect of 15.6 ms in the predictive condition with .78 power, which is at about the power level generally considered as adequate (e.g., Cohen, 1992). In the non-predictive condition, however, the experimental power for the (non-significant) subliminal cueing effect was .33 (i.e., it had a 33% probability of being 'detected'). We conducted an additional power analysis based on the observed sample variance in the non-predictive condition (for the fixed α-level of 0.05, two-sided, and the fixed number of participants) to assess at what magnitude a cueing effect could, in theory, be reliably detected in this latter condition (for a motivation of this approach, see Thomas, 1997). This revealed that the present experiment allowed for detecting an effect of 13.2 ms with 80% certainty (i.e., with a power of 0.8), which is almost twice the magnitude of the cueing effect that was measured (7.1 ms). Thus, while the experimental power allowed for detecting a smaller effect in the non-predictive condition than in the predictive condition, a facilitatory subliminal cueing effect was only detected in the predictive condition. From this, we conclude that, in the SOA 500 condition, the occurrence of subliminal cueing effects was indeed dependent on the predictive value of the visible cues. The analysis of learning effects corroborates the findings based on cueing effects, as it reveals a qualitative rather than quantitative dissociation between cue relevance conditions. Including block as an experimental factor in the design revealed a learning pattern in the cueing effects, which was dependent on the predictive value of the (visible) cues. While subliminal cues in the predictive condition showed a learning effect that was similar to that of supraliminal cues, subliminal cues in the non-predictive condition showed no learning at all. This null-effect is unlikely to be caused by a lack of statistical power (as deductible by a post-hoc power analysis), since the mean was below zero. This finding demonstrates that the utilization of subliminal arrow cues is dependent on the gradual building up of statistical knowledge about the informative value of the supraliminal cues. Thus, the potency of symbolic cues to evoke cueing effects may vary with their relevance to the current task, as deduced from visible statistical evidence.

In this first experiment, statistical learning relied on the relevance of visible cues. Since this learning was a gradual process rather than an instant one (reflecting a direct effect caused by top-down task instructions), we next address the question whether visual awareness is necessary for statistical learning of symbol predictiveness to occur. Indeed, recent studies show that some forms of statistical learning can occur outside of visual awareness (Turk-Browne, Scholl, Chun, & Johnson, 2008) and trigger perceptual anticipation effects (Turk-Browne, Scholl, Johnson, & Chun, 2010). For this purpose we reversed our paradigm to investigate whether the predictive value of subliminal cues can influence the incentive to utilize the supraliminal cues. This would be demonstrated by a stronger cueing effect in the supraliminal condition when subliminal cues are predictive, rather than non-predictive.

Experiment 2

Methods

Participants

Eighteen participants, age 18 to 31 (*M* = 22.8 years, *SD* = 4.21), took part in Experiment 2, five of which had also participated in Experiment 1. All participants were selected for normal or corrected to normal vision, stereoscopic vision and right-handedness. Exclusion criteria were identical to Experiment 1, such that thirteen participants were included in the analyses.

Stimuli and procedure

Stimuli, test setting and presentation methods were identical to Experiment 1. However, in Experiment 2 the subliminal cues served as a manipulation of cue relevance and therefore accounted for 75% of the trials. Consequently in 25% of the trials the cue was presented supraliminally. Again, participants took part in two experimental sessions on two different days, the order of which was counterbalanced across participants. In both experiments the supraliminal cues were 50% congruent (non-predictive), whereas the subliminal cues were 50% congruent in Experiment 2a (non-predictive) and 100% congruent in Experiment 2b (predictive). Because an SOA of 900 ms appeared to be too long to guarantee the invisibility of the cues, this SOA condition was omitted from Experiment 2. The last session ended with the same control task as in Experiment 1.

If the validity of subliminally presented arrow cues elicits statistical learning, this should evoke an incentive to use them both on visible and invisible trials. Consequently, subliminal learning should be reflected in larger cueing effects on supraliminal trials in the predictive condition, as compared with the non-predictive condition. Conversely, if no such effect is found, it would suggest that statistical learning of cue validity is a process requiring visual awareness.

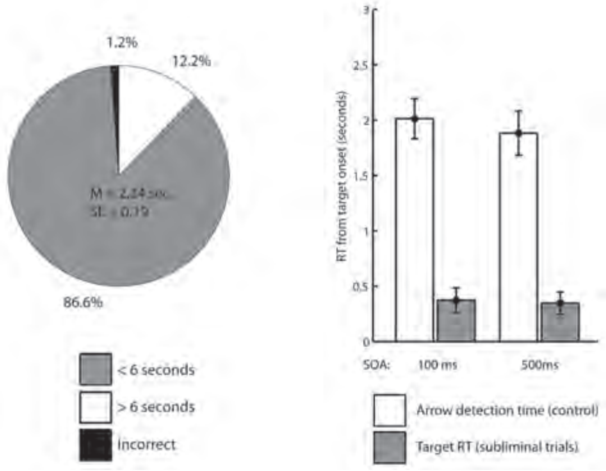

Figure 6. Results of the control task. Left: The cue remained invisible for 6 s in 12.2% of the control trials, while its direction was identified correctly within 2.24 s in 86.6%. Right: correct arrow detection times in the control task compared with target detection times on subliminal trials.

Results

Control task

The thirteen participants included in the analyses were unable to detect 12.2% of the cues within the six seconds time limit and a remaining 1.2% of control trials yielded incorrect responses. Consequently, 86.6% of the cues were accurately detected (M = 2.2 seconds, SD = 0.7) within the time limit (Figure 6). Similarly to Experiment 1, mean detection times in the control task did not differ between different SOAs $t(12) = 0.69$, $p = 0.505$, $d = 0.19$, suggesting comparable suppression times for both SOA's.

Cueing effects

Analogous to Experiment 1, trials in which participants reported an incorrect location of the target (0.7%), trials with RTs faster than 150 ms (0.1%) and trials with RTs slower than the participants 5th percentile RT on the control task (3.0%; varying between 0 and 5% for individual participants) were excluded from RT analyses. Subsequently, the RTs on congruent trials were subtracted from the RTs on incongruent trials, in all SOA × Relevance × Visibility conditions. The resulting measure reflects the cueing effect for each condition and is depicted in Figure 7. These analyses in Experiment 2 focus on the supraliminal condition, since the subliminal condition was only implemented as a manipulation of cue utilization. An analysis of the subliminal condition is provided in Supplementary Materials S2. First, an overall factorial 2 × 2 repeated measures ANOVA was conducted, with the factors SOA (100 ms and 500 ms) and Relevance (predictive and non-predictive). This analysis showed that the magnitude of cueing effects did not depend on SOA, $F(1, 12) = 3.39$, $p = 0.091$, $\eta^2 = 0.22$, or Relevance, $F(1, 12) = 0.18$, $p = 0.675$, $\eta^2 = 0.02$, nor did these factors interact, $F(1, 12) = 0.02$, $p = 0.887$, $\eta^2 = 0.00$. Moreover, the effect of Relevance did not differ between the first half and the second half of the experiment in either the short $t(12) = 1.2$, $p = 0.242$, $d = 0.34$ or the long SOA condition, $t(12) = -1.59$, $p = 0.137$, $d = -0.44$, reflecting that no learning occurred based on the predictive value of the subliminal cues.

Additional one-sample t-tests revealed that facilitatory cueing effects were observed in all supraliminal conditions; in the 100 ms SOA non-predictive condition, $t(12) = 4.63$, $p = 0.001$, $d = 1.28$, and predictive condition, $t(12) = 3.61$, $p = 0.004$, $d = 1.00$, and in the 500 ms SOA non-predictive condition, $t(12) = 3.01$, $p = 0.009$, $d = 0.86$, and predictive condition, $t(12) = 4.40$, $p = 0.001$, $d = 1.22$, such that facilitation was observed in every supraliminal condition. Subsequent paired sample t-tests revealed that these cueing effects were not affected by the predictive value of subliminal cues in either the short SOA, $t(12) = -0.25$, $p = 0.807$, $d = -0.07$, or long SOA condition, $t(12) = -0.33$, $p = 0.745$, $d = -0.09$.

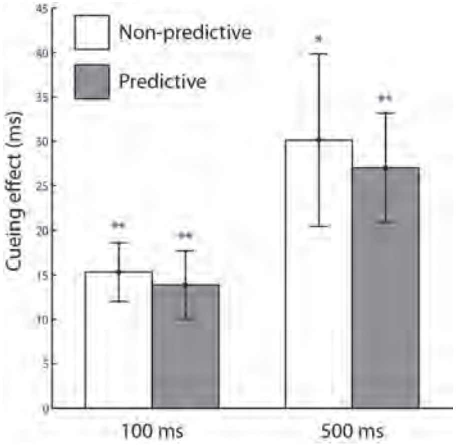

Figure 7. Cueing effects in the supraliminal condition of experiment 2. Calculated by subtracting RTs on trials with congruent cues from RTs on trials with incongruent cues. Positive cueing effects reflect facilitation. The predictive value reflects the cue validity of the subliminal arrow cues (100% or 50%). Error bars depict the standard error of the mean. * $p < .05$, ** $p < .005$.

Reaction times and accuracy

Analogous to Experiment 1, RTs and accuracy (see Supplementary Materials S3, Table S2) were analyzed, to verify that overall performance was comparable between Relevance conditions. Overall mean reaction time for Experiment 2 was 365 ms (*SD* = 40.1). A repeated measures ANOVA with the factors SOA, Relevance and Congruency showed a main effect of SOA, $F(1,12) = 27.07$, $p < 0.001$, $\eta^2 = 0.69$, showing that reaction times were shorter when cues were presented 500 ms (*M* = 349 ms, *SD* = 38.8) as compared with 100 ms (*M* = 381 ms, *SD* = 43.5) before target onset. The predictive value of subliminal cues did not affect RT on supraliminal trials, as was revealed by the absence of a main effect of Relevance, $F(1, 12) = 0.34$, $p = 0.570$, , $\eta^2 = 0.03$.

Overall response accuracy on supraliminal trials was 98.4% (*SD* = 2.1), with individual participants ranging between 93.8% and 100%. An overall repeated measures ANOVA with the factors Relevance, SOA and Congruency on the response accuracy in supraliminal trials revealed no main effect of Relevance, $F(1, 12) = 1.31$, $p = 0.273$, $\eta^2 = 0.10$, nor did the factor Relevance interact with SOA, $F(1, 12) = 0.16$, $p = 0.700$, $\eta^2 = 0.01$, or Congruency, $F(1, 12) = 1.16$, $p = 0.303$, $\eta^2 = 0.09$. As such, the predictive value of the subliminal cue affected neither the RTs, nor the accuracy on supraliminal trials (see Supplementary Materials S3, Table S2, for a complete overview of all RTs and accuracy per condition).

Discussion

Cueing effects on supraliminal trials were not affected by the predictive value of intermixed subliminal cues, which were either always congruent or had no predictive value. Hence, the predictive value of subliminal cues does not appear to elicit statistical learning, as measured by supraliminal cue utilization. In contrast, Experiment 1 showed that highly relevant visible stimuli affected both supraliminal and subliminal cueing effects and that the effects of cue relevance gradually emerged throughout the experiment. The combination of both findings shows that statistical learning of symbolic cue relevance requires visual awareness of the cues.

General discussion

The main finding of this study is that the utilization of subliminal symbolic cues is dependent on the predictive value of the cues, as deduced from visible statistical evidence. As a result, subliminal symbolic cues elicit facilitatory cueing effects on subsequent target detection, but only when intermixed supraliminal arrow cues are highly predictive of (i.e., 80% congruent with) the subsequent target location. Importantly, these effects emerged gradually throughout the experiment, as more evidence on the predictive value of arrow cues was accumulated. When the supraliminal cues were non-predictive (50% congruent) no cueing effects were found, nor did they build up over the course of the experiment. The subliminal cues themselves were always non-predictive as to the subsequent target location. Thus, when the relevance of visible arrow cues provides an incentive for cue utilization, the elicited facilitatory effect persists for arrows that are not consciously perceived. Conversely, Experiment 2 showed that the predictive value of subliminal cues (50% or 100% congruent) does not affect cueing effects of intermixed (non-predictive) visible arrow cues. We conclude that, at least in the case where supraliminal and subliminal cues are intermixed, conscious perception is required to make statistical inferences about the relevance of symbolic cues. Once this statistical information is consciously extracted, it may affect subsequent non-conscious processing in a context-dependent way.

In this study, we took special care to ensure that the observed cueing effects on subliminal trials were indeed elicited by cues that were invisible to participants. This was done by (1) estimating the fraction of cues in the subliminal condition that were unsuccessfully suppressed and (2) by correcting the cueing effects in the subliminal condition for this fraction of trials in which the cue might have been visible. Regarding the first point, it is important to note that suppressed cues were not likely to be visible based on the current data: the average reaction time in the main experiments was about 400 ms, whereas about a quarter of the cues in the control task remained unseen during the six seconds time limit. The remainder of the cues in the control task were detected after an average delay of around 2.5 seconds, more than six times as long as the mean RT in experimental trials. Moreover, we argue that the control task gave a conservative estimate of the detection times of subliminal cues. This claim is based on two arguments that rely on differences between the control task and the subliminal trials in the main experiment. First, in the experimental task (covert) attention was directed to the periphery, where the target was expected to appear. Since attention at a location away from rivalrous stimuli slows down alternations (Paffen, Alais, & Verstraten, 2006), the masked cue probably took longer to break through suppression in the experimental trials as compared with the control trials. Second, in the experimental trials attention was directed to the target, which was always presented to the eye that was made dominant with flash suppression. Directing attention to a target in a specific eye enhances the competition strength of the entire ipsi-ocular stimulus (Zhang, Jiang, & He, 2012). In our case, attending to the monocular target would have enhanced the dominance of the mask (suppressor), which was always presented to the same eye as the target. Consequently, this would have strengthened the suppression of the arrow cue (presented to the other eye) on experimental trials as compared with control trials. In sum, we believe this control task gave a low estimate of the suppression durations of subliminal cues in the actual experiment. Regarding the second point, the correction method applied in this study provides a conservative approach for detecting subliminal cueing effects. Cueing effects in the subliminal condition (Experiment 1) were corrected to account for the estimated fraction of trials in which cues were unsuccessfully suppressed. The contributions that this this fraction of trials would make to the cueing effect in the subliminal condition (if fully visible) was subtracted from the measured subliminal cueing effect. This method was thus based on the conservative assumption that unsuccessfully suppressed cues potentially elicit as much facilitation as supraliminal cues. However, it is likely that unsuccessfully suppressed cues produce less facilitation than supraliminal cues, since they are less clearly visible, as well as visible for a shorter duration than arrows in the supraliminal condition. Moreover, they might even impede subsequent target detection, since interocular conflict (mask versus arrow) is known to attract attention (Paffen, Hessels, & Van der Stigchel, 2012), in this case to the center of fixation and, as such, away from the target. The fact that a significant cueing effect is still observed using this conservative approach provides strong support of genuine non-conscious processing of suppressed arrow cues.

A possible concern that needs to be addressed relates to the finding that supraliminal arrow cues elicited cueing effects after both 100 and 500 ms SOAs, whereas subliminal cueing was only observed after an SOA of 500 ms. It could be argued that the absence of cueing effects on subliminal trials after short SOAs casts doubt on the potency of invisible symbols to elicit genuine subliminal cueing at longer SOAs. Indeed, a longer SOA provides more time for ocular dominance to switch back to the eye in which the arrow cue is presented, hence making it more likely to become visible before target onset. However, the control task demonstrated that the short and long SOA condition did not differ in arrow suppression duration, as measured from target onset. Therefore, differences in subliminal cueing effects between the two SOA conditions are unlikely to be the result of differences in suppression length. Moreover, the absence of cueing effects in the short SOA condition can be accounted for by a low-level explanation. On subliminal trials, the (suppressed) cue and the (non-suppressed) target were presented to different eyes, whereas in supraliminal trials both the cue and the target were presented to the same (non-suppressed) eye. A study by Self and Roelfsema (2010) revealed that monocular cues affect contra-ocular target detection only when the cue-target asynchrony is 150 ms or longer. The onset asynchrony of 100 ms used in the present study was probably too short to reach binocular channels and elicit intra-ocular cueing effects.

A final concern regarding the findings of Experiment 1 is that, while subliminal cueing effects were found in the predictive condition but not in the non-predictive condition, these cueing effects did not significantly differ between relevance conditions. This statistical limitation, however, does not jeopardize the conclusions of the present study. Indeed, comparable studies tend to look only at cueing effects within relevance conditions, rather than compare cueing effects between conditions (e.g., Al-Janabi, & Finkbeiner, 2012; Reuss et al., 2011). In these and similar studies, the subliminal stimulus is expected to elicit an effect that is qualitatively similar to that of its supraliminal counterpart (i.e., facilitation rather than inhibition), but is generally found to be smaller in magnitude (i.e., 5 to 15 ms; Mulckhuyse, Talsma, & Theeuwes, 2007; Reuss et al., 2011; Al-Janabi, & Finkbeiner, 2012; Palmer, & Mattler, 2013). As such, the main concern is to avoid contamination from supraliminal stimuli and to assess whether an effect is genuinely elicited by the subliminal stimuli. In the present study, this was achieved by means of the correction method described in section 2.1.4, which was more stringent in the predictive condition as compared with the non-predictive condition (i.e., corrections were computed from participants' supraliminal cueing effect, which were larger in the predictive than in the non-predictive condition). Thus, since both conditions were corrected differently, this does not allow for direct comparisons of cueing effects on subliminal trials between the two relevance conditions. In addition, it is important to point out that our general conclusion (i.e., the visible statistical context modulates the utilization of subliminal cues) does not solely rely on whether or not cueing effects are observed in the different relevance conditions. In addition to the cueing effects, the analysis of learning effects (i.e., difference in cueing effect between the first and second half of the experiment) shows a different behavioral pattern between relevance conditions. In experiment 1, a clear learning effect was observed for both supraliminal and subliminal cues in the predictive condition, whereas no learning effect was observed in the non-predictive condition. This finding reveals an additional dissociation between the predictive and non-predictive subliminal conditions. It shows that these conditions do not simply differ in the magnitude of the subliminal cueing effect, such that only one reaches the threshold of significance. Rather, it shows that two different processes are at play. In both supraliminal and subliminal trials it is the gradual accumulation of statistical evidence that either provides an increasing incentive to utilize the cues or to disregard the cues. Thus, the analysis of learning effects provides in itself compelling evidence for our general conclusion that statistical context modulates the utilization of subliminal cues.

The current findings are in line with the view that prior to non-conscious processing, conscious observation is required to construct a cognitive context which can guide behavior in a goal-directed and relevant manner. Our experimental setup was particularly well suited to demonstrate this action trigger hypothesis (Kunde et al., 2003; Kiesel et al., 2007), since it allowed to dissociate the observed relevance of a subliminal stimulus, as based on the consciously accessible context, from the actual relevance of the subliminal stimulus (which had no predictive value). The finding that subliminal arrow cues elicit cueing effects does, however, not imply that they receive full semantic analysis. In contrast with supraliminal cues, subliminal cues did not elicit cueing effects in the non-predictive condition. It appears then, that the semantic meaning of the cues was not automatically analyzed, but rather required an incentive to do so. We suggest that cues are included in a predetermined set of action triggers based on their physical properties, but only when visible cues provide a cognitive context that favors cue utilization. Thus, when the cognitive context forms an incentive to utilize subliminal information, an adequate response is triggered, irrespective of the relevance of the subliminal stimulus itself. It is important to note, however, that, while the aim of the present study was to unveil the necessary preconditions for subliminal stimuli to influence behavior (i.e., elicit cueing effects), the nature of these cueing effects remains elusive. Since the task was to respond to the location of a target rather than its identity, a cue that was informative of the target location was also informative of the required response. As such, it is impossible to unequivocally dissociate between motor priming effects (such as those described by Kunde et al., 2003) and attentional cueing (e.g., Ansorge et al., 2010; Folk et al., 1992).

Recently, Reuss et al. (2011) performed an experiment with meta-contrast masked arrow cues. Similarly to the present study, they found a facilitatory cueing effect of subliminal cues when cues were highly predictive, whereas no facilitation was observed for subliminal cues in the non-predictive condition. Their paradigm was different from ours however, in the sense that the predictive value of the subliminal cues was manipulated simultaneously with the predictive value of the supraliminal cues. Hence, this methodology did not allow for dissociating between the two possible origins of the cueing effect in the subliminal condition. Although it is debatable whether invisible arrow cues should be considered as a separate class of cues, there are two reasons to consider them as such in light of their predictive value. First, they elicit a percept that differs from that of visible cues, as a result of which different statistical relevance could be assigned to both 'cue-percepts' (compare, e.g., arrows of different colors). Second, Experiment 2 shows that participants' predictive set is not affected by the predictive value of arrow cues when they are not consciously perceived. Hence, it appears that the predictive set of visible and invisible arrow cues cannot be lumped together, but should be manipulated independently. By manipulating only the relevance of the supraliminal cues, we demonstrated that the predictive value of visible arrow cues formed an incentive to utilize arrow cues, which was then transferred to masked arrow cues. We thereby disproved the alternative explanation that the intrinsic predictive value of the subliminal cues elicited cueing effects by means of non-conscious statistical learning.

As mentioned above, the predictive value of subliminal cues did not modulate the cognitive context, such as to affect the utilization of (supraliminal) cues. This is at odds with recent studies showing that statistical learning can occur without awareness (Turk-Browne et al., 2010) and can bias spatial attention, even when completely unrelated to the current goals of the observer (Zhao, Al-Aidroos, & Turk-Browne, 2013). In contrast to these studies, however, the subliminal stimuli (either predictive or non-predictive) in the present study were intermixed with non-predictive supraliminal stimuli. Possibly, the stronger perceptual evidence provided by the latter, overruled the statistical learning of subliminal stimuli. Accordingly, the arrow cues were judged behaviorally irrelevant, based on the visible cues, irrespective of the statistical relevance of the subliminal cues.

From another perspective, the observed dissociation between supraliminal learning (Experiment 1) and the absence of subliminal learning (Experiment 2) can be framed within the *adaptation to the statistics of the environment* (ASE) model (Kinoshita, Forster, & Mozer, 2008; Kinoshita, Mozer, & Forster, 2011). This model makes predictions about speeded responses to prime-target pairs, based on a response control process that takes into account the (recent) history of events. This control process relies on response costs, which are computed from the relative cost of waiting (RT cost), and the cost of possibly making an incorrect response (accuracy cost). According to this model, swiftly reacting to primes in blocks with more incongruent prime-target pairs might result in a higher risk of making an erroneous response. Conversely, reacting to primes in blocks with more congruent prime-target pairs yields a lower risk of erroneous responses, and thus, a relative higher benefit of fast responses. In the present study, a visible statistical context provided the means to modulate cueing effects (Experiment 1), whereas an invisible statistical context did not (Experiment 2). While the history of events of Experiment 1 and Experiment 2 are similar from a statistical perspective (i.e., cues that had either a high or a low predictive value), these experiments might differ in terms of evidence accumulation. Indeed, the effects of suppressed stimuli may start to decay within hundreds of milliseconds after presentation (Kiefer, & Spitzer, 2000; Kiefer, & Brendel, 2006), such that information about these suppressed stimuli might not be accumulated over multiple trials (Humphreys, Besner, & Quinlan, 1988). Kinoshita, Mozer, and Forster (2011) showed that, whereas evidence of visible primes can accumulate over longer periods of time (i.e., a block), masked primes are mostly sensitive to recent history (i.e., trial 'n' and trial 'n – 1'). This short temporal window within which a masked prime is accessible to (non-conscious) cognition, is probably too short to infer the predictive value of prime-target pairs. In line with the ASE, the manipulation of the predictive value of subliminal cues in Experiment 2 did not provide a statistical context, based on the history of events, leading to an absence of cueing effects in this situation.

Another important finding of the present study that should be discussed in light of the ASE framework is the gradual emergence of cueing effects throughout the experimental sessions: when visible arrows were highly predictive, more facilitation was observed in the second half of the session as compared with the first half of the session, irrespective of cue visibility (i.e., supraliminal or subliminal). Crucially, these learning effects were not observed when arrow cues were non-predictive, even though facilitatory cueing effects were still observed in the supraliminal condition. The ASE predicts that, as statistical knowledge about the cues accumulates, the balance between costs (lower accuracy) versus benefits (faster RTs) will increasingly plead in favor of cue utilization in a predictive context, whereas it will remain unchanged in a non-predictive context (compared to the start of the experiment). Thus, the utilization of subliminally (and supraliminally) presented symbolic cues is affected by the gradual building up of statistical knowledge about the informative value of the (visible) cues.

At first sight, our results are at odds with those of Schall, Nawrot, Blake, and Yu (1993). In their study (as in ours), participants performed a speeded detection task, in which interocularly suppressed central arrow cues preceded peripheral targets. Participants were presented with dichoptic stimuli and were instructed to report perceptual dominance with a key press, thus initiating a trial of the actual cueing task. In contrast to our results, they found no significant cueing effects elicited by subliminal stimuli.The discrepancy between our findings and theirs can be accounted for by a number of factors. First, the longer cue-target onset asynchrony in their study (900 ms) might have resulted in the disengagement of attention from the cued location (e.g., Eimer, & Schaghecken, 2002; Barbot, & Kouider, 2011). Second, the use of binocular rivalry required participants to wait for a period of dominance and report the dominant percept with a key press to initiate a trial.The extra noise in the response times caused by this additional motor task might have reduced the power to detect cueing effects of relatively small magnitudes (~15 ms).

The existence of context-dependent subliminal processing as shown here is informative for the debate on the nature of non-conscious visual processing. On the one side is the view that non-conscious processing is as elaborate as conscious processing (e.g., Koechlin et al., 1999; Dehaene, & Naccache, 2001). On the other side is the view that non-conscious stimuli influence behavior through learned stimulus-response mapping (e.g., Abrams, & Greenwald, 2000; Damian, 2001). Both these views are unable to fully account for the present findings. In contradiction with the former view, processing of subliminal stimuli was not as extensive as processing of conscious stimuli. Experiment 2 showed that, in contrast with supraliminal cues, manipulating the relevance of subliminal cues did not affect cueing effects. Hence, conscious observation of the stimuli was required for participants to make a statistical inference about the predictive value of the cue. Stimuli presented outside of visual awareness on the other hand, were not processed sufficiently to elicit such an effect. According to the latter view, repetitively responding to a specific target builds up a stimulus-response memory trace that can subsequently be activated even when this target is not consciously accessible. As such, an acquired stimulus-response mapping between supraliminal target and response would also be triggered by the cue (i.e., a transfer from the 'to-be-responded-to' supraliminal stimulus to a similar subliminal stimulus). However, the present findings differ from stimulus-response mappings in two ways. First, the cue was a completely different stimulus, which had no perceptual similarities with the target and, as such, was not repeatedly responded to. It was only by means of the statistical context that an incentive was provided (supraliminally) to utilize the cue. Second, cues in the predictive condition elicited cueing effects in both visibility conditions, whereas a non-predictive context only elicited cueing effects for supraliminal cues, not for subliminal cues. This difference also emerged in the analysis of learning effects. Thus, the cueing effects elicited by supraliminal cues were only transferred to subliminal cues when the cognitive context provided an incentive to do so. Accordingly, acquired mappings between the peripheral target and the corresponding response cannot account for the facilitatory cueing effects as observed in Experiment 1. Rather, the present findings suggest that the processing of subliminally presented symbolic cues is the result of the gradual building up of statistical knowledge about the informative value of the (visible) cues. The inferred statistical relevance then provides an incentive for cue utilization, such that subliminal cues are only selected when they appear to be relevant for the task at hand.

Conclusion

The present paradigm allowed for disentangling the participants' incentive to utilize subliminal information from the actual relevance of the subliminal information. We showed that non-conscious processing is sufficiently elaborate for utilizing symbolic cues in a way that fits the cognitive context, even when the subliminal information is not relevant by itself. In contrast, this subliminal information is not processed extensively enough to affect behavior performed on visible stimuli by altering the cognitive context. These findings imply that the relevance of a subliminal stimulus is neither necessary nor sufficient to affect behavior. Rather, it is the conscious evaluation of visible information that provides an incentive to utilize (and predetermine the adequate response to) subsequent information, irrespective of its visibility. Accordingly, one of the possible functions of consciousness could be to extract general rules out of the perceptually available information, to provide guidelines for future behavior.

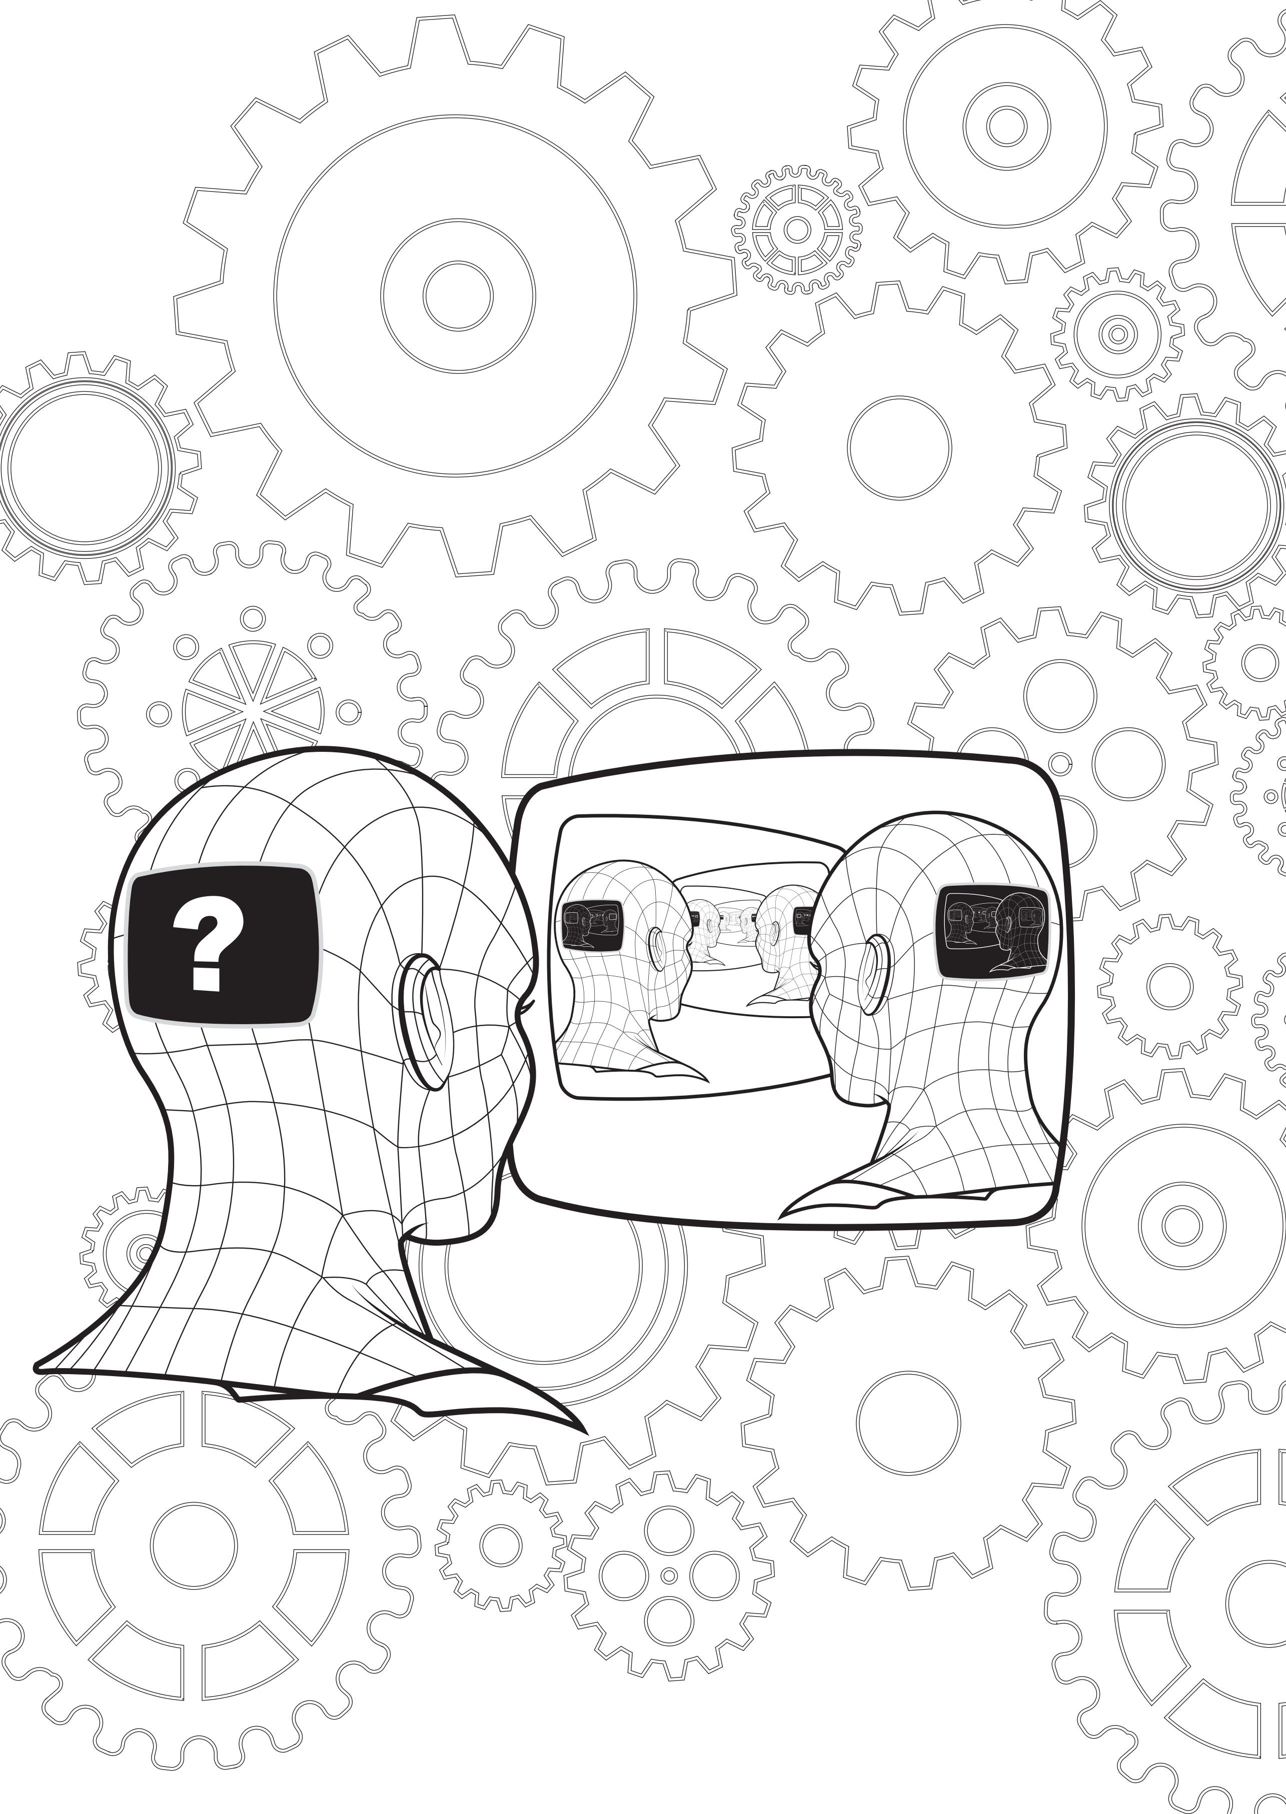

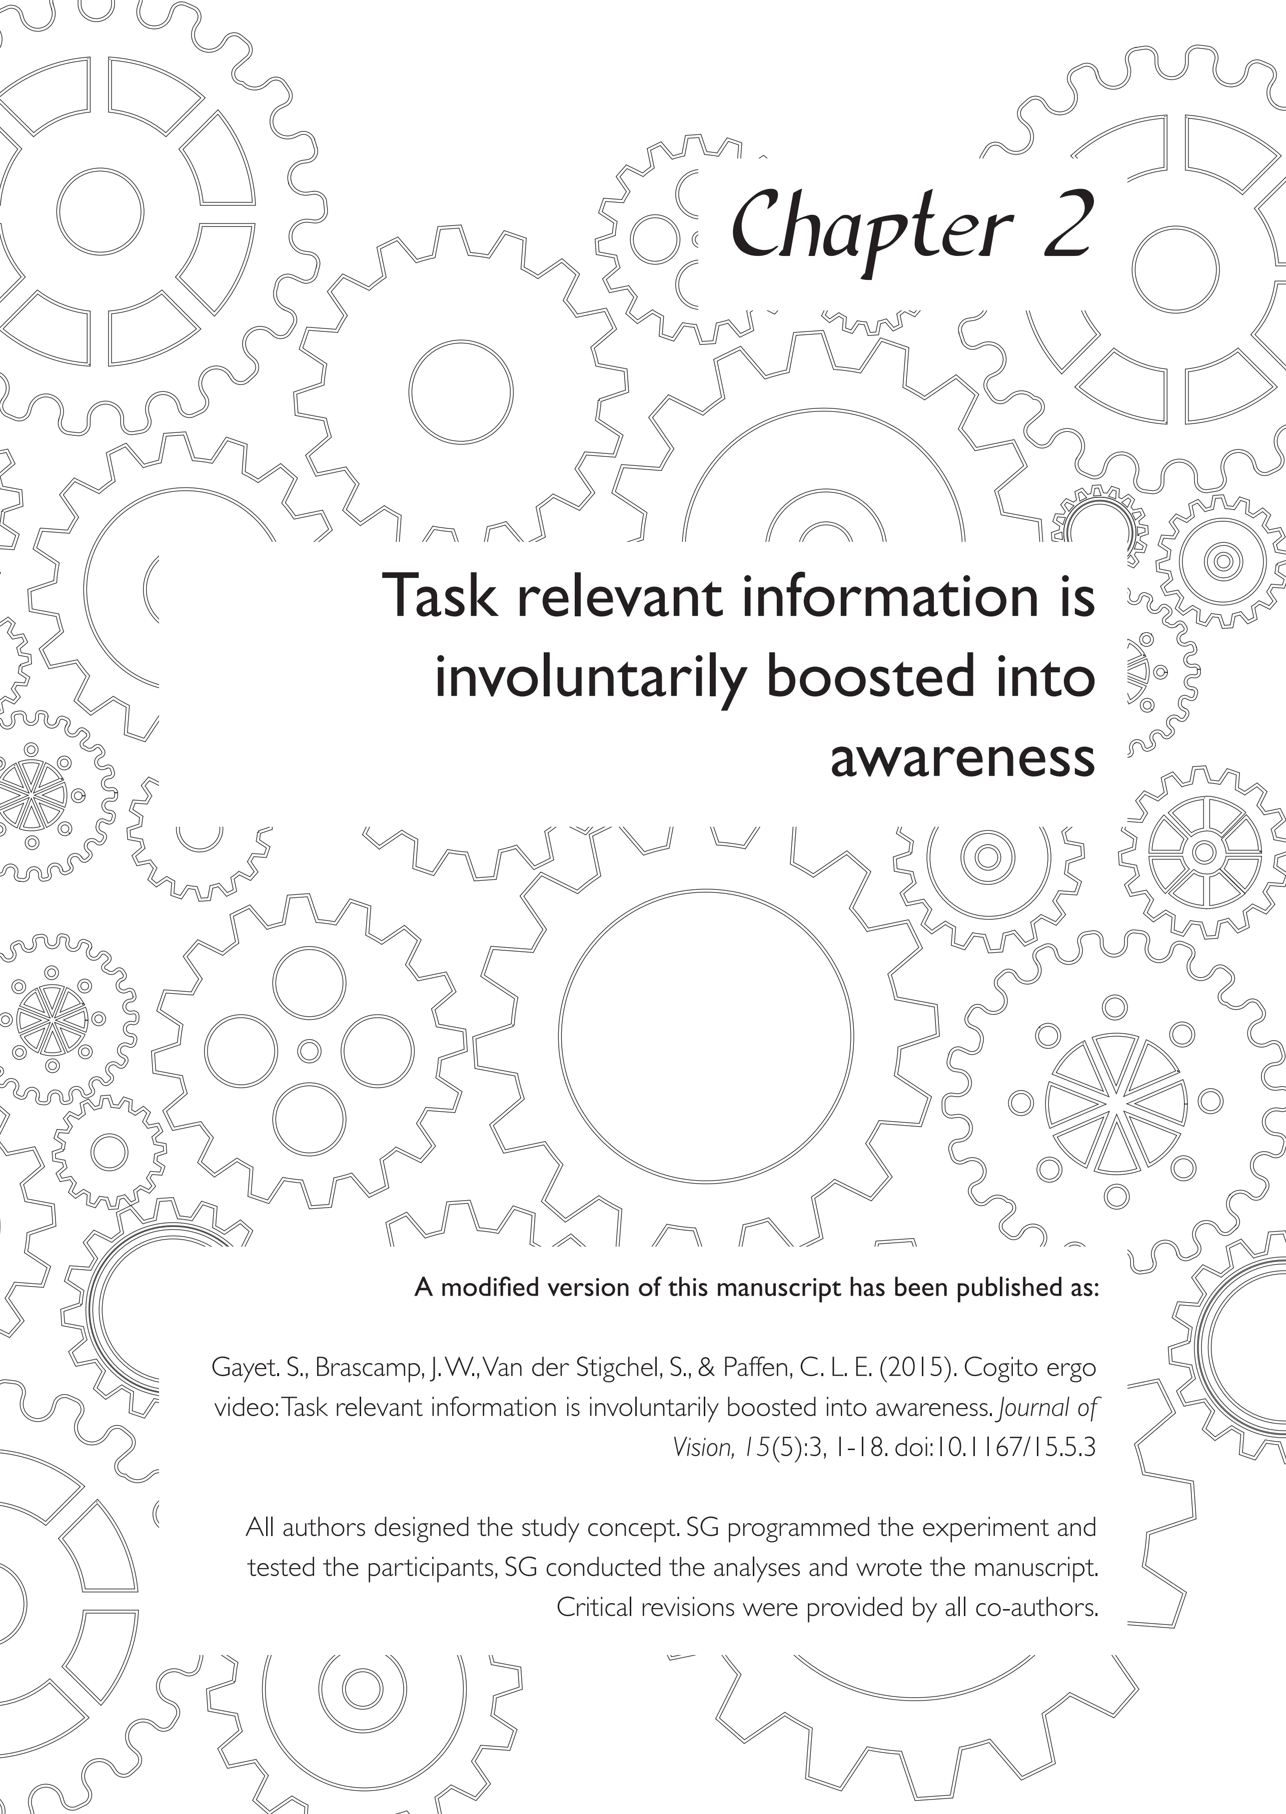

Chapter 2

Task relevant information is involuntarily boosted into awareness

A modified version of this manuscript has been published as:

Gayet. S., Brascamp, J. W., Van der Stigchel, S., & Paffen, C. L. E. (2015). Cogito ergo video: Task relevant information is involuntarily boosted into awareness. *Journal of Vision, 15*(5):3, 1-18. doi:10.1167/15.5.3

All authors designed the study concept. SG programmed the experiment and tested the participants, SG conducted the analyses and wrote the manuscript. Critical revisions were provided by all co-authors.

Abstract

Only part of the visual information that impinges upon our retinae reaches visual awareness. In a series of three experiments, we investigated how the task relevance of incoming visual information affects its access to visual awareness. On each trial, participants were instructed to memorize one out of two presented hues, drawn from different color categories (e.g., red and green), for later recall. During the retention interval, participants were presented with a differently colored grating in each eye such as to elicit binocular rivalry. A grating matched either the task relevant (memorized) color category or the task irrelevant (non-memorized) color category. We found that the rivalrous stimulus that matched the task relevant color category tended to dominate awareness over the rivalrous stimulus that matched the task irrelevant color category. This effect of task relevance persisted when participants reported the orientation of the rivalrous stimuli, even though in this case color information was completely irrelevant for the task of reporting perceptual dominance during rivalry. When participants memorized the shape of a colored stimulus, however, its color category did not affect predominance of rivalrous stimuli during retention. Taken together, these results indicate that the selection of task relevant information is under volitional control, but that visual input that matches this information is boosted into awareness irrespective of whether this is useful for the observer.

Introduction

Consciousness is believed to be a prerequisite for performing demanding tasks (Dehaene, Kerszberg, & Changeux, 1998) that require planning (Crick, & Koch, 2003), information integration (Tononi, & Edelman, 1998; Baars, 2002; Faivre, & Koch, 2014; Mudrik, Faivre, & Koch, 2014) and action selection for novel behavior (Dehaene, & Naccache, 2001; Kunde, Kiesel, & Hoffmann, 2003; Gayet, Van der Stigchel, & Paffen, 2014a; Ansorge, Kunde, & Kiefer, 2014). Due to our visual system's capacity limitations, only part of the visual information that impinges upon our retinae is selected for conscious experience (e.g., Edelman, & Tononi, 2000; Baars, 1997a; 1997b; Dennett, 1991). As a result of this processing limitation, it would seem of ecological significance to quickly separate the wheat from the chaff, such that information that is relevant for subsequent behavior predominates awareness. The aim of the present study was to investigate how such prioritization might transpire.

Our present approach centered on binocular rivalry, a phenomenon that occurs when different images are presented to each eye of an observer, causing perception to fluctuate back and forth between the two images (Wheatstone, 1838; for reviews, see Alais, & Blake, 2005; Lin, & He, 2009). During binocular rivalry, visual awareness therefore varies independently of physical stimulation, which remains constant. In the present study, we experimentally manipulated the task relevance of visual information, and measured how this affected perception during concurrent binocular rivalry. We presented observers with a dual task, comprised of an encoding phase, during which the task relevance of specific visual features (e.g., the color of a stimulus) is determined, and a perceptual selection phase, in which we measured whether one rivalrous image would predominate over the other, depending on its contingency with the predetermined task relevant feature. The separation of the encoding phase and the perceptual selection phase allowed for assessing the specific conditions under which the perceptual selection process is affected by the behavioral goals of the observer.

During binocular rivalry, and related forms of interocular competition such as continuous flash suppression (for a review, see Gayet, Van der Stigchel, & Paffen, 2014b), relevant stimuli have been shown to gain preferential access to awareness. This has been demonstrated mostly with socially relevant stimuli. For instance, faces predominate over non-faces (Stein, Hebart, & Sterzer, 2011; Stein, Peelen, & Sterzer, 2011; Bannerman, Milders, de Gelder, & Sahraie, 2008; Jiang, Costello, & He, 2007; Engel, 1956), emotional faces predominate over neutral faces (Stein, Seymour, Hebart, & Sterzer, 2014; Gray, Adams, Hedger, Newton,& Garner, 2013; Stein & Sterzer, 2012; Gray, Adams, & Garner, 2009; Yang, Zald,& Blake, 2007; Alpers & Gerdes, 2007; Pasley, Mayes, & Schultz, 2004; Coren & Russell, 1992)and faces with directed gaze predominate over faces with averted gaze (Chen & Yeh, 2012; Stein, Hebart, & Sterzer, 2011). Additionally, naturalistic images have been shown to predominate over non-naturalistic images (Baker, & Graf, 2009) and looming motion predominates over receding motion (Malek, Mendoza-Halliday, & Martinez-Trujillo, 2012; Parker, & Alais, 2007). In these studies differences in behavioral relevance of different stimuli were inevitably accompanied by differences in visual stimulus characteristics, which are known to affect interocular competition as well (e.g., Levelt, 1965; Yang, & Blake, 2011). Thus, instead of high-level cognitive effects (such as gaze direction), lower-level stimulus-characteristics (such as the specific conjunction between face curvature and pupil location) might drive the above mentioned prioritization for awareness (e.g., Chen, & Yeh, 2012). While the approach in these studies offers valuable insights in how different stimulus categories compete for visual awareness, the potential influence of differences in stimulus features between stimulus categories makes it less suitable for directly addressing the question of how relevance per se affects access to awareness. We therefore chose to strictly separate stimulus relevance from physical stimulus properties.

One way to experimentally vary the relevance of stimuli, while keeping the physical stimulation constant, is by manipulating the task instructions, such that, depending on the instruction, the same stimuli are either relevant or irrelevant for the task at hand. A number of studies have manipulated stimulus relevance while keeping the stimulus constant and measured the effect on binocular rivalry. For instance, stimuli that were previously paired with an electric shock (Alpers, Ruhleder, Walz, Muhlberger, & Pauli, 2005) or with a monetary reward (Balcetis, Dunning, & Granot, 2012) more frequently dominated perception at rivalry onset than similar stimuli that were not paired with a shock or reward. Along similar lines, a stimulus feature that was helpful for an auxiliary search task was more frequently dominant at rivalry onset than a stimulus that was not helpful for the search task (Chopin, & Mamassian, 2010).Interestingly, although these studies demonstrated that task relevant stimuli are more likely to gain initial dominance in binocular rivalry, no prolonged effects of task relevance on dominance durations during binocular rivalry were observed. This may be because manipulations of task relevance only affect the choice of the initial percept at rivalry onset, but it is also possible that a stimulus that had been task relevant prior to, rather than during, the binocular rivalry period is no longer subject to enough prioritized processing to affect perceptual selection during binocular rivalry. In sum, given our objective of assessing the influence of task relevance on access to awareness, we opted in our experiments to keep the task relevance in effect throughout the binocular rivalry period.

A particularly fruitful method to manipulate the task relevance of intrinsically neutral stimuli over a prolonged duration is the delayed match to sample task (e.g., Olivers, Meijer, & Theeuwes, 2006; Harrison, & Tong, 2009). In this task, participants are presented with two different stimuli that vary on a particular feature dimension, followed by a retro-cue that indicates which of these two stimuli should be memorized for a subsequent match to sample task. Consequently, on each trial, one stimulus is made relevant for the upcoming task, whereas the other one is not. This method has two major assets. First, it allows for controlling the period within which the stimulus is relevant for the task at hand (i.e., until the match to sample task). Second, and more importantly, the task relevant and the task irrelevant stimuli do not differ in terms of (objective) stimulus characteristics, nor do they differ in terms of initial stimulus processing depth, as it is unknown to the observer, at the time of stimulus presentation, which stimulus will be relevant for the subsequent task and which stimulus will be irrelevant. In the present set of experiments we combined the delayed match to sample task with binocular rivalry tracking to investigate the role of task relevance in the selection of information for visual awareness.

Rivalrous stimuli that match the task relevant color category are expected to predominate awareness over rivalrous stimuli that match the task irrelevant color category. This prediction stems from a number of studies, in which the delayed match to sample task was also used to manipulate task relevance of a stimulus or stimulus category. In these studies, the stimuli matching the task relevant stimulus category behave as if they were subject to an increase in signal strength, compared to the stimuli that match an irrelevant stimulus category. For instance, stimuli matching the task relevant color category capture attention (Olivers et al., 2006; van Moorselaar, Theeuwes, & Olivers, 2014) and eye movements (Hollingworth & Luck, 2009; Hollingworth, Matsukura, & Luck, 2013; Schneegans, Spencer, Schöner, Hwang, & Hollingworth, 2014; Silvis, & Van der Stigchel, 2014) in search tasks, elicit more pronounced behavioral priming effects under backwards masking (Pan, Cheng, & Luo, 2012) and appear to last longer (Pan, & Luo, 2012) than stimuli matching an irrelevant color category. Bistable perception is affected by concurrently retained information as well. The perceived direction of an ambiguously rotating structure from motion sphere was biased by the motion direction of a concurrently memorized unambiguously rotating sphere (Scocchia, Valsecchi, Gegenfurtner, & Triesch, 2013). Along similar lines, interocularly suppressed stimuli broke through continuous flash suppression faster when they matched the task relevant compared to task irrelevant color category (Gayet, Paffen, & Van der Stigchel, 2013). This study highlighted the intimate relationship between the task relevance of a stimulus category and the prioritization of matching stimuli for visual awareness. Note that the delayed match to sample task is generally referred to as a manipulation of visual working memory content. In this manuscript, we used the term 'task relevance', as it describes the experimental manipulation per se (i.e., which stimulus is relevant for the subsequent recall task) rather than 'visual working memory', which describes the anticipated behavior of the participant elicited by the experimental manipulation.

The present set of experiments was set up to investigate the way in which the current behavioral goals of the observer affect the selection of information for visual awareness. For this purpose we used a dual task paradigm, in order to separate the encoding phase, in which task relevance is determined, from the perceptual selection phase, in which the influence of the task relevant information on selection for awareness is measured. In Experiment 1 we demonstrated the effect of task relevance on perceptual dominance during binocular rivalry per se, consistent with expectations based on the literature summarized above. Rivalrous stimuli that matched the color category of the cued (i.e., task relevant) stimulus were dominant for a larger portion of the binocular rivalry period than rivalrous stimuli that matched the color category of the non-cued (i.e., task irrelevant) stimulus. Experiment 2 and Experiment 3 were set up to manipulate task relevance separately for the encoding (memorization) phase and the perceptual selection (binocular rivalry) phase. In Experiment 3, participants were instructed that either the color or the shape of a cued stimulus was relevant for the subsequent recall task. Under these circumstances, the color category of the memorized stimulus only biased perception if color was the volitionally retained feature dimension. That is, when participants had to recall the shape of a stimulus, its color category did not affect perception during concurrent binocular rivalry. This demonstrates that only the task relevant feature dimension of a stimulus was encoded such that it biased concurrent selection for visual awareness. Altering the feature dimension that is relevant for the perceptual selection task, however, revealed a different pattern of results. Once the task relevant feature (e.g., the red color category) was determined, matching rivalrous stimuli predominated awareness, both when participants reported the color (Experiment 1 and 3) and the orientation (Experiment 2) of the perceived rivalrous gratings. Importantly, in the latter case, color information per se was irrelevant and could therefore be disregarded to perform the task at hand (i.e., reporting the orientation of the stimuli). Together, these experiments showed that observers can voluntarily determine the task relevance of visual information in the encoding phase. However, once task relevance is determined, it involuntarily boosts concurrently presented matching information into awareness.

A secondary aim of the present study was to elucidate how the effect of task relevance on access to awareness of matching stimuli varied over the time course of the retention interval. This is interesting, because thus far only effects of task relevance on initial dominance have been demonstrated. The present method, however, allowed for maintaining task relevance throughout the entire binocular rivalry period. The results of Experiment 1 showed that the greater predominance of rivalrous stimuli that match the task relevant versus task irrelevant color category is stable throughout the entire retention interval. Thus, the faster access to awareness for stimuli matching the task relevant color category (as shown by Gayet et al., 2013) is not the result of an initial prioritization. Rather, the enhancement of matching information is observed as long as the task relevance is in effect.

Experiment 1

The goal of Experiment 1 was to investigate whether information that is relevant for a concurrent task will predominate during binocular rivalry. To this end participants were presented with a dual task (the full stimulus sequence of a trial is depicted in Figure 1). For the first task, participants were sequentially presented with two stimuli drawn from the color categories red or green, and a retro cue ("1" or "2") that indicated which of the two exact hues should be memorized for subsequent recall. At the end of each trial, participants were required to indicate which of three presented hues, drawn from the color category of the memorized hue, was identical to the one they had memorized at the start of that trial. During the retention interval, binocular rivalry was elicited during a ten second period, by presenting participants with a differently colored grating in each eye. The length of the retention interval was chosen as a compromise between the trial lengths in binocular rivalry studies that are typically one minute or more (e.g., Levelt, 1965) and the retention intervals of up to ten seconds used in delayed match to sample tasks (e.g., Christophel, Hebart, & Haynes, 2012; Harrison & Tong, 2009; Serences, Ester, Vogel, & Awh, 2009). Pilot studies revealed that, using the present stimuli, ten seconds would suffice for eliciting multiple dominance periods of each of the two gratings. On each trial, one of the rivalrous gratings matched the color category of the task-relevant (i.e., cued) stimulus, whereas the grating presented to the other eye matched the category of the task-irrelevant (i.e., non-cued) stimulus. During this binocular rivalry period, participants were instructed to continuously report with button presses whether they perceived a green grating, a red grating or a transition between the two. The exact hues of the rivaling stimuli were never identical to the hues used in the auxiliary memory task.

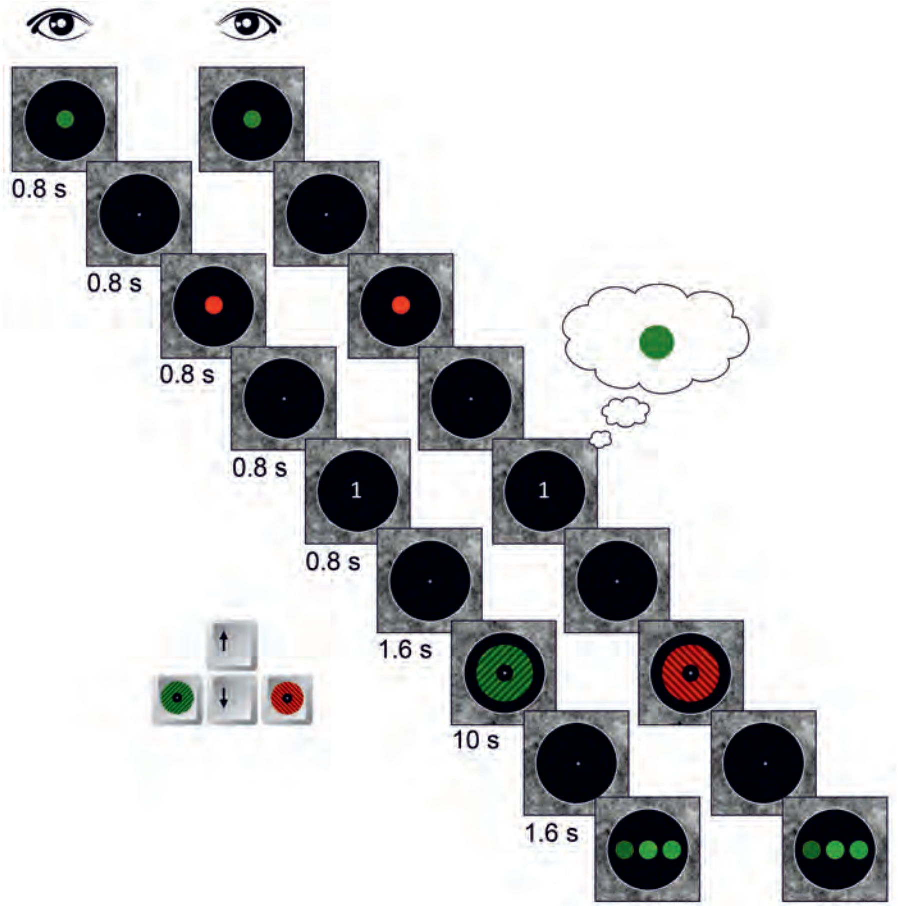

Figure 1. Stimuli and procedure of Experiment 1 and Experiment 2. Participants performed a dual task. For the memory task, participants subsequently viewed two differently colored stimuli and a retro cue ("1" or "2") that indicated which of the two stimuli should be memorized for subsequent recall. During the ten seconds retention interval, participants were presented with a different grating in each eye, such as to elicit binocular rivalry. The rivalrous stimuli could either match the color category of the cued stimulus or the color category of the non-cued stimulus. In Experiment 2, the rivalrous stimuli could also be of a third color category that was not used on that trial. Participants were instructed to continuously report the color (Experiment 1) or orientation (Experiment 2) of the rivalrous stimuli.

Methods

Participants

All experiments complied with the ethical principles set out in the Declaration of Helsinki. The participant group for this experiment consisted of 10 undergraduate students from Utrecht University, who participated for course credits or monetary reward after signing informed consent. All participants had normal or corrected-to-normal vision and were tested for color blindness with the Ishihara color blindness test plates (Ishihara, 1917), and tested for stereoscopic vision with the TNO test for stereoscopic vision (12th edition; Walraven, 1972).

Stimuli and apparatus

The experiment was conducted using an Apple dual 2-GHz PowerPC G5 equipped with a linearized 22 inch LaCie Electron Blue IV CRT monitor (1,024 × 768 pixels; 100 Hz refresh rate) and an Apple keyboard, which was used for response registration. There were no light sources in the experiment room, except for the computer monitor. Stimulus presentation and response collection were managed using the Psychophysics Toolbox 3 (Brainard, 1997; Pelli, 1997) in MATLAB (R2009b; The Mathworks, Natick, MA). A pair of displays was viewed dichoptically through a mirror stereoscope mounted on a chin rest, which kept the effective viewing distance at 57 cm. All stimuli were presented on a uniform black background (< 1 cd/m^2). To facilitate binocular fusion of the two complementary images, we surrounded a circular area of 6.7° of visual angle presented to each eye with a Brownian (i.e., $1/f^2$) noise frame that subtended 9.9 by 9.9° of visual angle and had a mean luminance of 21 cd/m^2 and a 12% RMS contrast. Also, all frames were comprised of a white fixation circle with a diameter of 0.2° of visual angle and a black fixation dot in its center (0.04° of visual angle). The retro cues consisted of the Arabic numerals "1" and "2" in white (48 cd/m^2) Arial font with a font size of 18.

The color stimuli used for the memory task were identical to those of Gayet et al. (2013). These stimuli consisted of ten non-saturated colored circles (five red and five green hues; one of each presented during each memorization phase) that were physically matched to the luminance of the saturated blue rivalrous stimulus (see below) using a PR-650 SpectraScan colorimeter-telephotometer (Photo Research, Chatsworth, CA). The resulting color patches had a mean luminance of 6.06 cd/m^2 (*SD* = 0.17) and a diameter of 1.4° of visual angle. An overview of all CIE-

Table 1. _Overview of the colors (in CIE values*) used in all experiments._

Color	X-value	Y-value	Luminance (cd/m^2)
Red 1	0.590	0.340	6.08
Red 2	0.542	0.298	5.87
Red 3	0.587	0.365	5.76
Red 4	0.510	0.279	5.84
Red 5	0.562	0.388	6.29
Green 1	0.285	0.441	6.15
Green 2	0.324	0.568	6.14
Green 3	0.262	0.479	6.22
Green 4	0.359	0.544	6.19
Green 5	0.232	0.372	6.04
Blue 1	0.181	0.145	6.09
Blue 2	0.167	0.084	5.81
Blue 3	0.170	0.145	6.22
Blue 4	0.156	0.076	5.99
Blue 5	0.182	0.187	6.28
Saturated blue†	0.152	0.071	6.02

* CIE values stands for Commission Internationale d'Eclairage values, as measured with a PR-650 SpectraScan colorimeter/telephotometer (Photo Research, Inc.).

† The saturated blue was the reference stimulus for the perceptual luminance matching (i.e., heterochromatic flicker photometry; Kaiser & Comerford, 1975; Wagner, & Boynton, 1972) with the saturated red and saturated green colors, that were used to create the square wave gratings for the binocular rivalry stimuli. Also, the luminance of the saturated blue color served as the basis for the physical luminance matching (by means of a telephotometer) of the 15 color variations for the memory task, which are described in this table.

color values is provided in Table 1. The stimuli used for the binocular rivalry task consisted of circular square wave gratings with a diameter of 4.7° of visual angle and a spatial frequency of 10 cycles per stimulus width. The colors used for the peaks of the gratings were saturated red or green. A saturated blue color(6.02 cd/m^2, $x = 0.152$, $y = 0.071$) was used to obtain perceptual isoluminance of the saturated red and saturated green colors for each subject by means of heterochromatic flicker photometry (Kaiser & Comerford, 1975; Wagner, & Boynton, 1972). The pixel intensities of the peaks of the gratings were halved to obtain the pixel intensities for the troughs of the gratings (i.e., 33% Michelson contrast). The rivalrous stimuli had a circular central recess with a diameter of 1.4° of visual angle, such as to avoid spatial overlap with the colored stimuli used for the memory task and ascertain visibility of the fixation dot.

Experimental design and procedure

The experimental conditions were fully counterbalanced and manipulated within participants. This included the order in which the color stimuli for the memory task appeared (either red followed by green or vice versa), the retro cue (memorize either the first or the second stimulus), the eyes to which the rivalrous gratings were presented (either red in the left eye and green in the right eye or vice versa) and the orientations of the rivalrous gratings (plus 45°from the vertical in the left eye and minus 45°from the vertical in the right eye or vice versa). The combinations of these conditions were repeated five times throughout the experiment, accounting for a total of 80 trials. The hues for the memory task were randomly selected in such a way that each of the five hues within a color category was presented equally often during the memorization phase of the memory task. Two factors were randomized without replacement but not counterbalanced. These were the choice of hues for the two distractor stimuli that were presented alongside the target stimulus in the recollection phase of the memory task, and the location of that target with respect to these two distractors. Finally, the stimulus-response contingencies for the binocular rivalry task were counterbalanced between participants (i.e., left arrow key for red percept and right arrow key for green percept and vice versa). The experiment was separated into 8 blocks of about 5 minutes each, and was preceded by an eight trial practice session. Participants were explicitly instructed to maintain fixation, and to avoid blinking during the ten seconds binocular rivalry period.

Data analyses

In all analyses of the binocular rivalry data, participants' reported percepts were separated into those in which the perceived rivalrous stimulus (e.g., the red grating) matched the task relevant color category (e.g., a red hue was memorized for the memory task) and those in which the rivalrous stimulus matched the task irrelevant color category (e.g., a green hue was memorized). Additionally, all transitory percepts were discarded, such that only exclusive percepts (i.e., the report of seeing either one grating or the other) were included in the analyses. Trials in which participants reported correctly and incorrectly on the memory task were both included in the analyses of the binocular rivalry data. In this method it is generally assumed (e.g., Olivers et al., 2006; van Moorselaar et al., 2014) that incorrect answers on the memory task reflect a failure to report the exact hue (i.e., choosing the wrong hue of green), rather than a failure to memorize the correct stimulus (i.e., having memorized the hue of the red stimulus instead).Because the delayed match to sample task is intrinsically difficult and, hence, brings about a high number of errors, excluding erroneous trials would drastically reduce our experimental power without increasing the diagnostic power. Compared to other binocular rivalry studies in which rivalrous stimuli are displayed for one minute or more, the ten second display time of rivalrous stimuli in the present set of experiments was relatively short. As a result of this 24.1% (*SD* = 9.2%) of the percepts were truncated. Therefore, the main analyses in the present set of experiments are the analyses of total dominance durations, which is a more robust method under these circumstances (for a similar approach, see Scocchia, Valsecchi, Gegenfurtner, & Triesch, 2014). The total dominance duration was computed as the summed duration of each rivalrous percept within a trial, separated on the basis of whether they matched the task relevant or the task irrelevant color category (i.e., the factor Relevance). In all experiments, planned comparisons were conducted between Relevance conditions. These consisted of paired-samples t-test that were declared significant if the p-value (two-sided) was smaller than the standard alpha level of 0.05 (after Bonferroni correction, in case of multiple comparisons).

Results and discussion

In the memory task, participants correctly reported the memorized color on 54.6% (*SD* = 10.3) of the trials, which is better than the 33% chance level, $t(9) = 6.55$, $p < 0.001$. The accuracy on the memory task did not differ between trials in which participants were cued to memorize the first stimulus (*M* = 55.8%, *SD* = 10.3) and trials in which they were cued to memorize the second stimulus (*M* = 53.5%, *SD* = 13.7), $t(9) = 0.56$, $p = 0.590$. This shows that the task was feasible but demanding, which is an important requirement for delayed match to sample tasks (e.g., Olivers et al., 2006).

First, a planned pair-wise comparison was conducted with the two levels of the factor Relevance, to assess whether they differed in total dominance durations (Figure 2, left). This analysis revealed that the rivalrous stimulus that matched the color category of the task-relevant stimulus was dominant for a longer portion of a trial (*M* = 3.08 seconds, *SD* = 0.54) than the stimulus that matched the color category of the task-irrelevant stimulus (*M* = 2.77 seconds, *SD* = 0.66), $t(9) = 3.61$, $p = 0.006$. Next, we wanted to make sure that this effect of Relevance is strictly brought about by the task relevance of the cued stimulus, rather than by the sequence of presentation of the colored stimuli. For this purpose, we conducted a 2 × 2 repeated measures analysis of variance with the factors Relevance (relevant versus irrelevant) and Retro Cue (i.e., "1" or "2"). In this control analysis, the absence of an interaction between the factors Relevance and Retro Cue on total dominance durations, $F(1, 9) = 0.23$, $p = 0.644$, showed that the effect of Relevance, reported above, did not depend on whether the relevant color category was that of the first or the second stimulus. This was corroborated by a main effect of Relevance, $F(1, 9) = 9.49$, $p = 0.013$, and the absence of a main effect of Retro Cue, $F(1, 9) = 0.44$, $p = 0.523$. These results demonstrate that, over an entire rivalry period, rivalrous stimuli that match a task relevant color category predominate over rivalrous stimuli that match equally accessible, but task irrelevant color category.

To further investigate this effect, the median dominance duration of individual percepts within a trial (i.e., epochs) was computed. This analysis allows to ascertain whether the effect found for total dominance durations (A) can be accounted for by a particular subset of epochs (e.g., only the first epoch), and (B) whether this effect changes throughout the retention interval. Again, this data was separated on the basis of whether the percept matched the task relevant or the task irrelevant color category (Figure 2, right). On average, participants had 4.3 (*SD*=1.4) exclusive percepts in each ten seconds rivalry period. All participants reported at least one exclusive percept on every trial. Participants had at least two exclusive percepts in 98.5% (*SD* = 2.6) of the trials, at least three exclusive percepts in 95.0% (*SD* = 6.7) of the trials, at least four percepts in 82.9% (*SD* = 15.2), at least five exclusive percepts in 63.4% (*SD* = 22.1) of the trials and at least six exclusive percepts in 36.4% (*SD* = 28.4) of the trials. Up until the fifth epoch, each condition contained at least 11 data points per participant. In the sixth epoch, however, 50% of the participants had less than 10 data points per condition, two of which had only three data points per condition. As a result of this rapid decrement in data points for increasing epoch numbers after the fifth epoch, only the first five epochs were included in the next analysis. This analysis revealed a main effect of Relevance, $F(1, 9) = 12.24$, $p = 0.007$, which means that epoch durations depended on whether percepts matched the task relevant color category (*M* =1.40 seconds, *SD* = 0.33) or the task irrelevant color category (*M* = 1.26 seconds, *SD* = 0.34). The absence of a main effect of Epoch number revealed that the different epochs did not differ in duration per se, $F(1.7, 15.3) = 1.96$, $p = 0.178$ (Greenhouse-Geisser corrected, as the assumption of sphericity was violated, $p < 0.05$). Importantly, the effect of Relevance on epoch duration did not interact with the Epoch number, $F(4, 36) = 0.29$, $p = 0.886$. This shows that the increased predominance of the rivalrous stimulus that matches the task-relevant as opposed to task irrelevant color category mentioned above is maintained throughout the retention interval.

Finally, a measure of initial dominance was obtained by computing the proportion of trials in which participants' first percept of the rivalry period matched the color category of the task relevant (i.e., cued) stimulus. This revealed no reliable effect of task relevance on initial percept, $t(9) = 0.24$, $p = 0.819$. The matching percept was reported as the first exclusive percept in 49.5% (*SD* = 6.7) of the trials. Participants had no strong bias towards reporting the grating presented to one eye more often as the initial percept than the grating presented to the other eye (48% left eye, *SD* = 5.5, range = 38.8% – 56.2%). There was however a slight bias towards reporting the green grating as the first percept more often than the red grating (30.8% red grating, *SD* = 12.5, range = 12.5% – 51.3%). This bias indicates that isoluminance achieved by means of heterochromatic flicker photometry does not ensure equal onset dominance probability in binocular rivalry.

Together, these findings demonstrate that the active retention of a color for an auxiliary task causes matching stimuli to predominate during binocular rivalry. Interestingly, the task relevance of a color category did not kick start the perception of matching rivalrous stimuli, as was demonstrated by the analysis of initial dominance. Rather, as was shown by the analysis of epoch durations, the duration of each individual epoch was lengthened when the percept matched the task relevant rather than irrelevant color category.

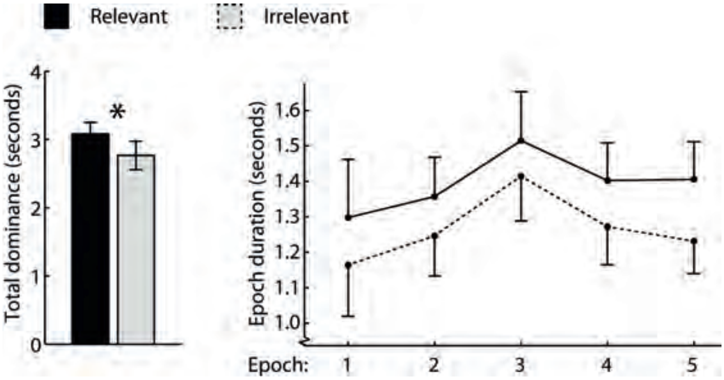

Figure 2. Results of the binocular rivalry task in Experiment 1. Left: Total dominance duration in seconds of the rivalrous stimuli that matched the task relevant (depicted in black) and the task irrelevant (depicted in light gray) color categories. Right: Median epoch duration of the rivalrous stimuli that matched the task relevant (solid line) and the task irrelevant (dotted line) color categories for the first five epochs of the binocular rivalry period. Error bars represent the standard error of the mean. $^{*}p < 0.05$.

Experiment 2

Based on the first experiment, we conclude that visual input that matches task relevant information predominates awareness over visual input that matches task irrelevant information. With the next experiment, we addressed two questions. First, we aimed to elucidate whether the effect described above is the result of an increased predominance of the rivalrous stimulus that matches the task relevant color category, or a decreased predominance of the rivalrous stimulus that matches the task irrelevant color category. For this purpose a third color condition was added, such that the rivalrous stimuli could now match the color category of the cued stimulus (task relevant), of the non-cued stimulus (task irrelevant) or of a novel color category that had not been used on that trial (task unrelated). This task unrelated color category acted as a baseline level, against which increases in dominance durations in one condition can be dissociated from decreases in dominance durations in the other condition. Second, we aimed to investigate whether the effect found in Experiment 1 would persist if participants reported the orientation rather than the color of the perceived rivalrous stimulus. This would demonstrate that the effect of a concurrently retained color category affects the perception of rivalrous stimuli based on their color, even if color is an irrelevant feature dimension for the task of reporting rivalry dominance. Additionally, requiring participants to report the orientation rather than the color of the rivalrous stimuli makes the task more robust to response biases. That is, irrespective of which grating they perceive, participants might be more prone to responding "green" when they are concurrently memorizing a (different) color from the green color category. It is less likely that memorizing a green color would bias participants towards responding to a particular orientation.

Methods

Participants

Fifteen new undergraduate students from Utrecht University participated in this experiment. All inclusion criteria from Experiment 1 applied to Experiment 2 as well. Because this experiment was comprised of three trial types rather than one (see below), we increased the number of participants to obtain comparable statistical power as in Experiment 1. Due to below chance performance on the memory task (i.e., 29% correct at a 33% chance level), one participant was excluded from further analyses and substituted by a new participant.

Stimuli and design

The stimuli and design of Experiment 2 were identical to those of Experiment 1, except for the following differences. For this experiment, five hues from a blue color category were included (see Table 1). Also, a saturated blue color for the binocular rivalry task was included. As a result of this expansion, a task unrelated condition could be implemented, in which a rivalrous stimulus matched neither the color category of the cued stimulus, nor that of the non-cued stimulus. There were now three different trial types; the two stimuli for the binocular rivalry task could match the task relevant color category and the task irrelevant color category (this condition was comparable to Experiment 1), they could match the task relevant and the task unrelated color category, or they could match the task irrelevant and task unrelated color category. All experimental conditions (i.e., trial type, color category of first stimulus, color category of second stimulus, eye of presentation, retro cue) were again fully counterbalanced within participants, adding up to a total of 72 trials (i.e., 24 per trial type). The selection of specific hues for the memory task was identical to that of Experiment 1. In contrast to Experiment 1, however, participants were now instructed to report the orientation of the perceived rivalrous gratings. Participants' key presses could now reflect exclusive perception of either a grating tilted clockwise 45° or tilted counterclockwise 45°, or reflect a transitory (i.e., non-exclusive) percept. Again, the stimulus-response contingency was counterbalanced between participants. Due to the smaller amount of trials per condition in Experiment 2 (24) compared to Experiment 1 (80), the analyses of epoch duration yielded too little data points in each Relevance condition (less than 5) from the third epoch onwards, and was therefore deemed uninformative. As in Experiment 1, the main analyses consisted of the planned pair-wise comparisons between total dominance durations of the Relevance conditions, now separated on the basis of trial type.

Results and discussion

In the memory task, participants correctly reported the memorized color on 59.1% ($SD = 6.7$) of the trials, which is better than the 33% chance level, $t(14) = 14.97$, $p < 0.001$. The accuracy on the memory task did not differ between trials in which participants were cued to memorize the first stimulus ($M = 62.0\%$, $SD = 8.2$) and trials in which they were cued to memorize the second stimulus ($M = 56.1\%$, $SD = 6.7$), $t(14) = 2.04$, $p = 0.061$.

In the binocular rivalry task, we first aimed to assess whether the effect of Relevance on total dominance duration was different for the three trial types used in this experiment. The two levels of the factor Relevance (e.g., task relevant and task irrelevant), however, differed between the three trial types (see Figure 3, left, center and right). For instance, in one trial type (e.g., left on Figure 3) the task irrelevant stimulus condition was assigned to the second level of the factor Relevance, whereas the same stimulus condition was assigned to the first level of the factor Relevance in another (i.e., right on Figure 3) trial type. The outcome of an overall analysis of variance with the three trial types would depend on the contingency that we assigned between specific stimulus conditions (i.e., task irrelevant and task unrelated) and specific levels of the factor Relevance (i.e., level 1 and level 2). To circumvent this issue, we computed the absolute difference between the total dominance durations of the two levels of each trial type (e.g., the absolute difference between the task relevant and the task irrelevant percepts). Then, we conducted a 3-level repeated measures analysis of variance with the factor Trial Type on this difference measure. This revealed a main effect of Trial Type, $F(1.357, 19.004) = 6.075$, $p = 0.016$ (Greenhouse-Geisser corrected, as the assumption of sphericity was violated, $p < 0.05$), which indicated that the effect of Relevance was different in the three trial types.

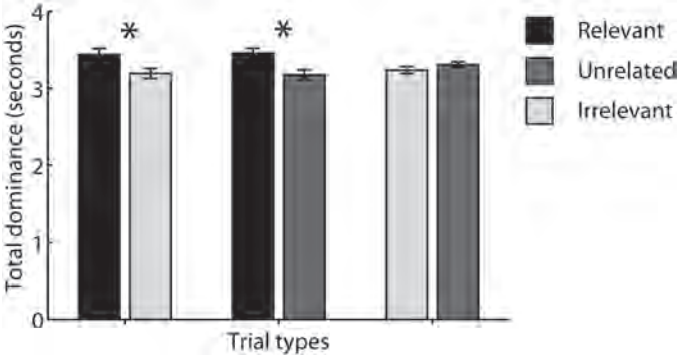

Figure 3. Results of the binocular rivalry task in Experiment 2. The height of the bars represents the total dominance duration of a percept over the ten seconds rivalry period, separated into three trial types. The black bars represent rivalrous stimuli that matched the task relevant color category; the light gray bars represent rivalrous stimuli that matched the task irrelevant color category; the light gray bars represent the rivalrous stimuli that matched the task unrelated color category. Error bars represent the standard error of the mean. * $p < 0.05$ (Bonferroni-corrected).

To investigate the specific effect of task relevance on predominance during binocular rivalry we conducted planned pair-wise comparisons between the two levels of Relevance for each of the three trial types. This revealed that rivalrous stimuli that matched the task relevant color category ($M = 3.45$ seconds, $SD = 3.06$)

were dominant for a larger portion of the binocular rivalry period than rivalrous stimuli that matched the task irrelevant color category (*M* = 3.13 seconds, *SD* = 3.17), *t*(14) = 3.19, *p* = 0.007 (significant at a Bonferroni-corrected α-level of 0.0167). This shows that the greater predominance for visual input that matches task relevant information, as found in Experiment 1, persists when a different stimulus dimension is reported (i.e., the orientation rather than the color of the rivalrous grating). Next, these analyses revealed that rivalrous stimuli that matched the task relevant color category (*M* = 3.58 seconds, *SD* = 2.20) were dominant for a larger portion of the binocular rivalry period than rivalrous stimuli that matched the task unrelated color category (*M* = 3.16 seconds, *SD* = 3.30), *t*(14) = 3.83, *p* = 0.002. In contrast, rivalrous stimuli that matched the task irrelevant color category (*M* = 3.23 seconds, *SD* = 2.16) and rivalrous stimuli that matched the task unrelated color category (*M* = 3.27 seconds, *SD* = 1.74) were dominant for an equivalent portion of the rivalry period, *t*(14) = 0.65, *p* = 0.526. Thus, the greater predominance of visual input that matches task relevant information, as found in Experiment 1, is not the result of decreased predominance of task irrelevant information but of increased predominance of task relevant information.

Finally, we aimed to investigate potential effects of task relevance on perceptual dominance at rivalry onset. We conducted a repeated measures analysis of variance on the three levels of the factor Trial Type (as for the analysis of total dominance durations above). The absence of an effect of Trial Type, *F*(2, 28) = 1.18, *p* = 0.322, showed that potential effects of task relevance on initial dominance were similar across all three trial types. Subsequent pair-wise comparisons revealed no significant difference between the rivalrous stimuli matching the relevant (*M* = 50.3%, *SD* = 12.5) and irrelevant color categories, *t*(14) = 0.09, *p* = 0.932. Rivalrous stimuli matching the relevant (*M* = 55.6%, *SD* = 10.2) color category tended to be reported more often as the initial percept than those matching the unrelated color categories, *t*(14) = 2.09, *p* = 0.055, but this tendency seems spurious when compared to the Bonferroni corrected α-level of 0.0167. Additionally, the absence of a difference between the rivalrous stimuli matching the irrelevant (*M* = 49.4%, *SD* = 14.6) and unrelated color categories, *t*(14) = 0.15, *p* = 0.885, corroborates the idea that the unrelated color category plays no special role in rivalry onset. Thus, the task relevance of a color category, as manipulated by the memory task, had no systematic influence on the perceptual dominance at rivalry onset. Participants had no strong bias towards reporting the grating presented to one eye more often as the initial percept than the grating presented to the other eye (46.0% left eye, *SD* = 14.1, range = 15.3% – 73.6%). There was however a slight bias towards reporting the green grating as the first percept more often than the red grating (28.6% red grating, *SD* = 25.0, range = 0.0% – 83.3%), as well as a bias to report the green grating more often than the blue grating (31.4% blue grating, *SD* = 20.3, range = 4.2% – 79.2%), but no preference in initial dominance was apparent between the blue grating and the red grating (49.7% red grating, *SD* = 20.1, range = 4.2% – 95.8%).

With Experiment 2, we replicated the finding that rivalrous stimuli predominate awareness when they match task relevant information. This effect was further specified by demonstrating that it was caused by an increase in predominance of the rivalrous stimuli matching task relevant information, as opposed to a decrease in predominance of the rivalrous stimuli matching task irrelevant information. This is a pattern of findings that is similar to the effect of working memory content on attentional capture (Olivers et al., 2006) and on suppression durations under continuous flash suppression (Gayet et al., 2013). Thus, boosting task relevant information, rather than suppressing task irrelevant information, might be a more general manner in which the visual system prioritizes relevant information.

More importantly, the data of Experiment 2 also revealed that the predominance of rivalrous stimuli that match task relevant color information over those that match task irrelevant color information persists when another stimulus feature dimension (i.e., orientation) is reported. Aside from eliminating potential strategic biases that participants might entrain, it also demonstrates that the effect of task relevance on perception during rivalry is not selective for the feature dimension that is used during the binocular rivalry task. Indeed, in Experiment 2, color information per se was irrelevant for the binocular rivalry task, which consisted of reporting the orientation of the rivalrous gratings. Nonetheless, the color category that was made relevant for the concurrent memory task still affected perception during binocular rivalry. As such, the task relevant information captured visual awareness automatically.

Experiment 3

Experiment 2 showed that a color category that is relevant for the memory task biases the access to awareness of colored stimuli in a concurrent perceptual report task, even when color information is irrelevant for performing that concurrent task and could therefore be disregarded completely. This raises the question whether the ability of a memorized stimulus to bias perception also generalizes to irrelevant features of that memorized stimulus, or is specific to the feature dimension that is volitionally memorized. In other words, would the color of a memorized stimulus also bias perception if the memory task entailed remembering, say, the shape of the stimulus rather than its color? To test this, we manipulated the aspect ratio (length and width) of the color stimuli used for the memory task and included two types of instructions. In one half of the experimental blocks participants were instructed to memorize the color of the cued stimulus, making 'color' a volitionally retained stimulus dimension, whereas in the other half of the experimental blocks participants were instructed to memorize the shape of these same stimuli, making 'color' an incidental feature of the cued stimulus.

Methods

Participants, stimuli and design.

Ten undergraduate students from Utrecht University participated in this experiment. All inclusion criteria from Experiment 1 applied to Experiment 3 as well. The stimuli and experimental procedure of Experiment 3 were similar to that of Experiment 1, except for the following addition (see Figure 4). The shape of the memory stimuli were drawn from two categories: either vertical ellipses or horizontal ellipses. In each category, five elliptical shapes were created that varied stochastically between 1.4° and 1.16° of visual angle (with steps of 0.08° of visual angle, i.e., two pixels), while negatively co-varying the orthogonal dimension between 0.76 and 1.0, such that the surface of the ellipses remained constant (as far as allowed for by the screen resolution). During the delayed match to sample phase of the experiment, all three stimuli varied along the critical dimension, but were identical to the memorized stimulus with respect to its irrelevant feature dimension. For instance, in the condition where participants were to memorize the shape of an ellipse of a specific hue, all three stimuli in the recollection phase were of that exact same hue, but varied along the shape dimension (and vice versa when participants had to memorize the hue of the cued ellipse).

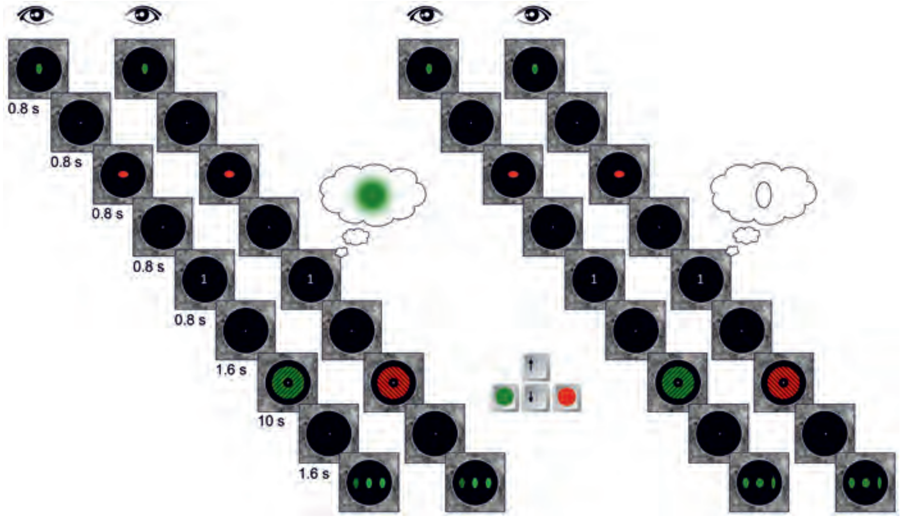

Figure 4. Stimuli and procedure of Experiment 3. Depending on the experimental block, participants were instructed to either memorize the color (left) or the shape (right) of the cued stimulus. In both conditions participants were instructed to continuously report the color of the rivalrous stimuli.

The experimental design was also identical to that of Experiment 1, except for the two following differences. First, the two shape categories of the stimuli (horizontally or vertically elongated) were fully counterbalanced within participants, along with all conditions already included in Experiment 1 (the red and green color categories of the memory task stimuli, the two retro cues and the two orientations of the rivalrous gratings). Second, in one half of the experiment participants were instructed to memorize the exact hue of the cued stimulus (as in Experiment 1), whereas in the other half of the experiment participants were instructed to memorize the exact shape of the cued stimulus. Both halves of the experimental session were preceded by eight practice trials, to get acquainted with the task and the stimulus-response contingencies. The order in which the task conditions were presented (i.e., first the color memory task and then the shape memory task or vice versa) was counterbalanced between participants. Both halves of the experimental session were divided into four blocks, totaling to 64 experimental trials. Due to the smaller amount of trials per condition in Experiment 3 (32) compared to Experiment 1 (80), the analyses of epoch duration yielded too little data points in each Relevance condition (less than 10 from the second epoch onwards and less than 5 from the third epoch onwards), and was therefore deemed uninformative. As in Experiment 1, the main analyses consisted of the planned pair-wise comparisons between total dominance durations of the Relevance conditions, now separated for both Memory Instruction conditions.

Results and discussion

In the memory task, participants correctly reported the memorized feature on 58.4% ($SD = 8.2$) of the trials, which is better than the 33% chance level, $t(9) = 8.83$, $p < 0.001$. The accuracy on the memory task did not differ between trials in which participants were instructed to memorize the color ($M = 53.8\%$, $SD = 10.6$) and trials in which they were instructed to memorize the shape of the cued stimulus ($M = 58.4\%$, $SD = 14.0$), $F(1, 9) = 0.63$, $p = 0.449$. Also, the accuracy on the memory task did not differ between trials in which participants were cued to memorize the first stimulus ($M = 57.8\%$, $SD = 14.5$) and trials in which they were cued to memorize the second stimulus ($M = 54.4\%$, $SD = 8.5$), $F(1, 9) = 0.66$, $p = 0.437$. Finally, there was no interaction between these two factors, suggesting that performance on the memory task did not vary as a function of memory instruction and retro cue, $F(1, 9) = 0.69$, $p = 0.426$.

For the analysis of the binocular rivalry data (Figure 5), we first conducted a 2 X 2 repeated measures analysis of variance on the total dominance duration with the factors Relevance (rivalrous stimulus is related to relevant versus irrelevant color category) and Memory Instruction (memorize the color versus the shape of the cued stimulus). This analysis revealed no main effect of Memory Instruction, $F(1, 9) = 0.48$, $p = 0.506$, indicating that total dominance durations of the rivalrous percepts did not depend on whether participants memorized the color or the shape of the cued stimulus. There was, however, a main effect of Relevance, $F(1, 9) = 9.81$, $p = 0.012$, which showed that the total dominance duration of a percept depended on whether it matched the color category of the cued stimulus, or the color category of the non-cued stimulus. The interaction effect between Memory Instruction and Relevance, however, did not reach significance, $F(1, 9) = 3.59$, $p = 0.091$. In the present experiment, as in Experiments 1 and 2, we aimed to investigate whether task relevant information would predominate over task irrelevant information during binocular rivalry. Our prior hypotheses for the present experiment, however, were explicitly aimed at investigating whether or not this effect would emerge in the two Memory Instruction conditions. As such, we conducted planned pair-wise comparisons between the predominance of rivalrous stimuli matching the task relevant versus the task irrelevant color category, for both Memory Instruction conditions separately (Figure 5), despite the absence of an interaction effect. These analyses allow for comparing the present results with those of Experiments 1 and 2 where similar planned pair-wise comparisons were conducted. When participants were instructed to memorize the color of the cued stimulus, the rivalrous stimulus that matched its color category (*M* = 3.64 seconds, *SD* = 0.18) was dominant for a larger portion of the ten seconds rivalry period than the rivalrous stimulus that matched the color category of the non-cued stimulus (*M* = 3.12 seconds, *SD* = 0.15), *t*(9) = 3.10, *p* = 0.0128 (significant at a Bonferroni-corrected α-level of 0.025). This replicated the findings of Experiments 1 and 2. In contrast, when participants were instructed to memorize the shape of the cued stimulus, the rivalrous stimulus that matched its color category (*M* = 3.26 seconds, *SD* = 0.18) was not dominant for a larger portion of the rivalry period than the rivalrous color that matched the color category of the non-cued stimulus (*M* = 3.31 seconds, *SD* = 0.20), *t*(9) = 0.31, *p* = 0.764. Thus, although the interaction between Relevance and Memory instruction did not reach significance, the planned comparisons suggest that the main effect of Relevance, was fully driven by the condition in which participants had to memorize the color of the cued stimulus. These planned pair-wise comparisons, similar to those used in Experiment 1 and 2, demonstrate that when the color of the cued stimulus is memorized, the matching rivalrous stimulus is dominant for about 500 ms longer than the mismatching rivalrous stimulus. However, there is no hint of a difference in dominance durations when the shape of that same stimulus is memorized (about 50 ms shorter). Together, these results show that only the intentionally stored feature dimension of a stimulus has the potency to affect subsequent perception during binocular rivalry.

Similarly to Experiments 1 and 2, we aimed to investigate potential effects of task relevance on the perceptual dominance at rivalry onset. Pair-wise comparisons revealed no differences between the rivalrous stimuli matching the relevant (*M* = 49.5%, *SD* = 4.2) and irrelevant color categories, *t*(9) = 0.86, *p* = 0.932. When participants were instructed to memorize the color of the cued stimulus there was no difference in initial dominance between the rivalrous stimulus matching the relevant (*M* = 51.2%, *SD* = 4.7) and irrelevant color categories, *t*(9) = 0.84, *p* = 0.423. Similarly, when participants were instructed to memorize the shape of the cued stimulus there was no difference in initial dominance between the rivalrous stimulus matching the relevant (*M* = 47.8%, *SD* = 6.8) and irrelevant color categories, *t*(9) = 1.02, *p* = 0.333. Moreover, these proportions did not differ between tasks, *t*(9) = 0.36, *p* = 0.730. Thus, the task relevance of a color category did not affect the initial dominance at rivalry onset. Overall, participants had no strong bias towards reporting the grating presented to one eye more often as the initial percept than the grating presented to the other eye (46.7% left eye, *SD* = 14.8, range = 32.8% – 81.3%). There was however a slight bias towards reporting the green grating as the first percept more often than the red grating, both when participants memorized the color (31.9% red grating, *SD* = 20.0, range = 0.0 – 56.3%) and the shape of the cued stimulus (34.1% red grating, *SD* = 22.0, range = 0.0 – 68.8%). This bias was comparable to that found in Experiment 1 and Experiment 2.

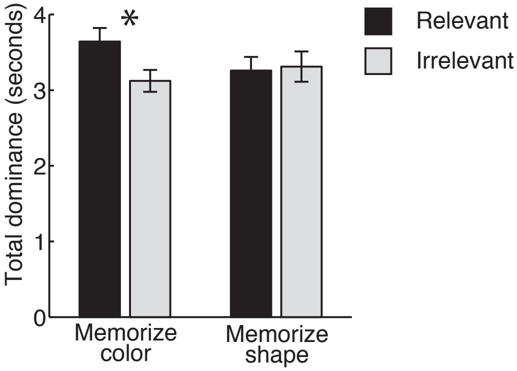

Figure 5. Results of the binocular rivalry task in Experiment 3. The height of the bars represent the total dominance duration of a percept over the ten seconds rivalry period, separated on the basis of the two task instructions (memorize the color or the shape of the cued stimulus). The black bars represent rivalrous stimuli that matched the task relevant color category; the light gray bars represent rivalrous stimuli that matched the task irrelevant color category. Errors bars represent the standard error of the mean. *$p < 0.05$ (Bonferroni-corrected).

General discussion

In the present set of experiments we investigated the influence of task relevance on the selection for visual awareness. Task relevance of stimuli was operationalized by a delayed match to sample task, in which one out of two stimuli should be memorized, while the selection for awareness was measured by means of binocular rivalry tracking during the retention interval, with stimuli that matched and mismatched the concurrently retained stimulus. All three experiments demonstrated an increased predominance for binocular rivalry stimuli that matched rather than mismatched the color category (i.e., red, green or blue) of a concurrently retained hue (e.g., a specific blue hue).

The finding that task relevant stimuli predominate awareness under conditions of binocular rivalry is in line with several other studies that used stimuli with no intrinsic relevance (Alpers et al., 2005; Balcetis et al., 2012; Chopin, & Mamassian, 2010). These studies demonstrate that stimuli that match a stimulus category in an auxiliary task are more likely to gain initial dominance in binocular rivalry. Interestingly, Alpers et al. (2005) and Chopin, and Mamassian (2010) did not detect an effect on dominance durations, and Balcetis et al. (2012) only measured initial dominance. In contrast, our study demonstrated that task relevance elicited a prolonged predominance of matching stimuli (Experiment 1, 2 and 3), that lasted throughout the rivalry period (Experiment 1). The major difference with the present set of experiments is that in their studies the stimulus features that were coupled to a relevance condition were not task-relevant during the binocular rivalry task. Possibly, a stimulus that is no longer relevant during the binocular rivalry task is no longer subject to enough prioritized processing to affect sustained perception during binocular rivalry. A compelling argument for this view comes from the self-stabilizing nature of binocular rivalry; when one eye's dominance duration is artificially lengthened (e.g., by means of task instructions) this reduces the duration of that eye's next few dominance periods (Blake, Westendorf, & Fox, 1990). As such, unless the manipulation is continuously applied, manipulations of dominance durations will tend to hamper themselves. In the study of Chopin and Mamassian, for example, a cue indicated whether on that trial participants would be required to report percept dominance, or whether they were required to perform a search task. The cue onset therefore indicated that the feature that was relevant for the search task was no longer relevant on that trial, and could thus be disregarded. In the present set of experiments, in contrast, the relevant stimulus feature had to specifically retain its relevance during the binocular rivalry task in order for participants to perform the subsequent recall task. In line with this idea, a recent study showed that memorizing a face stimulus for a subsequent recall task caused matching face stimuli to break through continuous flash suppression faster than mismatching face stimuli (Pan, Lin, Zhao, & Soto, 2014). When the recall task was performed just before the suppression task so that the face stimulus was no longer relevant, however, this effect was abolished. Therefore, it appears that the processing of stimuli under interocular competition is affected to a greater extent by a match with concurrently relevant information, compared to information that had been relevant in close temporal proximity. While these interpretations potentially explain why we did find effects of task relevance on dominance durations, it does not explain why these manipulations had no impact on perceptual selection at rivalry onset in the current study. A difference between our study and the three previously mentioned studies that did find an effect of task relevance on initial dominance (Alpers et al., 2005; Balcetis et al., 2012; Chopin, & Mamassian, 2010) is that we offered participants three rather than two options to report their percept during binocular rivalry: they could report seeing either of the two percepts as well as having an ambiguous percept. While including this third option provides more information on participants' percept during binocular rivalry, it might also hamper the detection of differences in initial dominance compared to a forced choice paradigm. Additionally, we did not provide participants with temporal constraints. Chopin and Mamassian, for instance, restarted the trial if participants had not reported a first percept within one second. Alternatively, most of the variance in initial dominance in this study was explained by factors other than whether the rivalrous stimuli matched color category of the cued stimulus or not. For instance, one participant reported the image presented to the left eye as initially dominant in 85% of the trials, whereas another participant reported the green stimulus as the initially dominant percept in 100% of the trials. In general, preferences in onset rivalry are believed to be relatively stable within participants, and very variable across participants (Carter, & Cavanagh, 2007). This is made apparent by the large between-subject differences in eye dominance preference and color preference in initial dominance in all three Experiments reported here. Such large, systematic differences leave but little room for top-down modulation by task-instruction.

The main finding of the current experiments is the asymmetry between top-down control and visual awareness. Experiment 3 revealed that the color category of the actively retained hue only affected concurrent rivalry when color was the volitionally retained feature dimension. In contrast, when color was an incidental feature, that is, when participants had to memorize the shape of a colored stimulus, dominance durations during binocular rivalry were not affected by its color category. Conversely, however, Experiment 2 revealed that prolonged dominance for stimuli matching the color category of the retained stimulus was also observed when participants reported the orientation of the rivalrous stimuli. Thus, the task relevance of a color category affected the selection of information for visual awareness, even though color information per se was irrelevant and could be disregarded to perform the task at hand. A large number of studies has shown that top-down control could affect dominance during binocular rivalry (Chong, & Blake, 2006; Chong, Tadin, & Blake, 2005; van Ee, van Dam, & Brouwer, 2005; Meng & Tong, 2004; Mitchell, Stoner, & Reynolds, 2004; Ooi, & He, 1999; Lack, 1978; for reviews, see Paffen, & Alais, 2011; Dieter, & Tadin, 2011). For instance, attending to a specific feature (e.g., Mitchell et al., 2004) or to a stimulus (e.g., Ooi, & He, 1999) presented to one eye, increases the predominance of the ipsi-ocular percept. As such, the results of Experiment 3 could be interpreted as an effect of endogenous feature-based attention on perception during binocular rivalry. In the present study, however, participants ultimately had no volitional control over which of the rivalrous stimuli should predominate. Firstly, because matching color information was propelled into awareness while participants were reporting orientation information (Experiment 2). Secondly, even in Experiment 1 and Experiment 3, where participants had to report the color of the rivalrous stimuli, the information that was boosted into awareness was totally irrelevant for the task at hand; it only matched the color category that had been determined to be task relevant beforehand. If anything, when retaining a specific hue for subsequent recall, task performance might be hindered more by prolonged perception of a similar but slightly different stimulus than by prolonged perception of a stimulus from a distinct color category. Arguably, stimuli matching the task relevant color category captured awareness, irrespective of the current behavioral goals of the observer. From this perspective, the asymmetry between top-down control and visual awareness reported here is in line with the idea that consciousness is required to determine rules that govern novel behavior (e.g., Kunde et al., 2003; Gayet et al., 2014a; Ansorge et al., 2014). Once these rules are in place, they can be implemented in non-conscious processes, in a non-flexible, automated manner. Thus, whereas participants could exert volitional control over which stimulus to retain for subsequent recall (Experiment 3), this retained information then guided the selection for awareness during binocular rivalry in a way that was insensitive to volitional control (Experiment 2).

During interocular competition, visual features of stimuli can be misbound, such that a percept comprises one feature from the stimulus presented to the left eye and another feature from the stimulus presented to the right eye. For instance, awareness of one feature dimension (e.g., flicker) can co-occur with unawareness of another feature (e.g. orientation) of the same stimulus under continuous flash suppression (Zadbood, Lee, & Blake, 2011; Yang & Blake, 2012; Mudrik, Gelbard-Sagiv, Faivre, & Koch, 2013). During binocular rivalry, a percept can, for instance, comprise of the color and motion (Andrews, & Blakemore, 1999; Carney, Shadlen, & Switkes, 1987; Creed, 1935), motion and shape (Alais, & Parker, 2006) or color and shape features presented to different eyes (Hong, & Shevell, 2006; Kang, & Shevell, 2008). Color misbinding is particularly likely if, as in the present case, differently colored rivalrous stimuli are isoluminant (Kang, & Shevell, 2008). Our Experiment 2 showed that the orientations of rivalrous stimuli that matched the task relevant color category predominated over the orientations of rivalrous stimuli that matched the task irrelevant color category. The order of magnitude of this effect (a difference of 310 ms) was comparable to that of Experiment 1 (320 ms) in which the same stimuli were used, but participants reported the color rather than the orientation of the rivalrous gratings. This allows for an alternative perspective on the present results. Experiment 3 showed that the color of the cued stimulus only affected concurrent perception when it was the volitionally retained feature dimension. In contrast, Experiment 2 showed that once the relevance of a color category was determined, it boosted rivalrous stimuli into awareness that were comprised of this color category, irrespective of whether color information was necessary to report perception. In this view, task relevancy in the encoding phase is highly selective, such that only the volitionally retained feature dimension affects concurrent perception. During perceptual selection, however, a stimulus matching the task relevant feature dimension is boosted as a whole, such that all its features (i.e., color and orientation in the present case) gain more perceptual dominance.

At first glance, our results seem at odds with a recent study in which the content of visual working memory did not affect perception during binocular rivalry (Scocchia, Valsecchi, Gegenfurtner, & Triesch, 2014). In their third experiment, stimuli were used that closely resembled the stimuli that were used in the present set of experiments.

Participants were required to memorize the spatial frequency of a grayscale sine wave grating for a delayed match to sample task. During the retention interval, differently colored orthogonal sine wave gratings were presented to each eye to elicit binocular rivalry. Participant reported the color (red or green) of the perceived grating, which either matched or mismatched the orientation of the memory stimulus. Crucially, the volitionally stored feature dimension of the memorized stimulus was its spatial frequency, whereas its orientation was an incidental feature dimension, which just happened to be part of the memorized stimulus. Our Experiment 3 revealed that only the volitionally retained feature dimension (in our case color; in their case spatial frequency) has the potency to affect concurrent perceptual selection, whereas the same feature does not affect perceptual selection if it is an incidental feature of the memorized stimulus (in our case color, while shape was memorized; in their case orientation, while spatial frequency was memorized). Our findings are in line with a functional magnetic resonance imaging study by Serences, Ester, Vogel, and Awh (2009). In this study, participants were instructed to memorize either the color or the orientation of a colored Gabor patch for delayed recall. Only the volitionally retained feature dimension could be reliably decoded from early visual areas, whereas the incidental feature dimension, that just happened to be part of the same object, could not. Consequently, the incidentally stored feature dimension which is not represented by neural activity in the visual cortex lacks the potency to interact with concurrent processing of incoming visual information. In conclusion, the null effect reported in the third experiment of Scocchia and colleagues (2014) is in line with the null effect reported in the Shape Memory condition of our Experiment 3: incidental features of a memorized stimulus do not affect perceptual selection during binocular rivalry.

In the first two experiments of the study by Scocchia and colleagues (2014) discussed above, the memory stimuli were images drawn from the stimulus categories "planes", "houses" and "faces". During the concurrent rivalry tracking task, participants reported the color (red or green) of the dominant percept, which was either an exemplar from the memorized stimulus category or from one of the other two categories. The authors found the same dominance durations for stimuli that matched and stimuli that mismatched the image category of the memorized stimulus. As such, the memory task did not affect perceptual selection during binocular rivalry. In contrast with their third experiment in which they used grating stimuli, dominance durations in the binocular rivalry task (e.g., an airplane and a face) were now measured as a function of their contingency with the volitionally retained feature dimension (e.g., an airplane). Thus, in terms of task instructions, this experiment was equivalent to our Experiment 2, in which we did find an effect of the retained stimulus category on perceptual selection. Taking together our study and that of Scocchia et al., it appears that only the active retention of lower level stimulus features (such as color) can bias concurrent perceptual selection, whereas the active retention of more complex stimuli (such as airplanes) cannot.

Conclusion

The present set of experiments demonstrates how the selection of information for visual awareness is affected by the potential behavioral relevance of that information. On the one hand, this selection process is highly specific, and thus very efficient; stimuli matching the color category of the cued stimulus predominated awareness only when color was the volitionally retained stimulus dimension. When color was an irrelevant feature dimension that just happened to be part of the cued stimulus, however, its color category did not affect concurrent selection for awareness. On the other hand, this selection process is very rigid. A color category that was previously determined as being relevant, affected concurrent perceptual selection even if it was completely uninformative for the task at hand. Taken together, these results indicate that the selection of task relevant information is under volitional control. Concurrent visual input that matches this information, however, is boosted into awareness, irrespective of whether it is useful for the current behavioral goals of the observer.

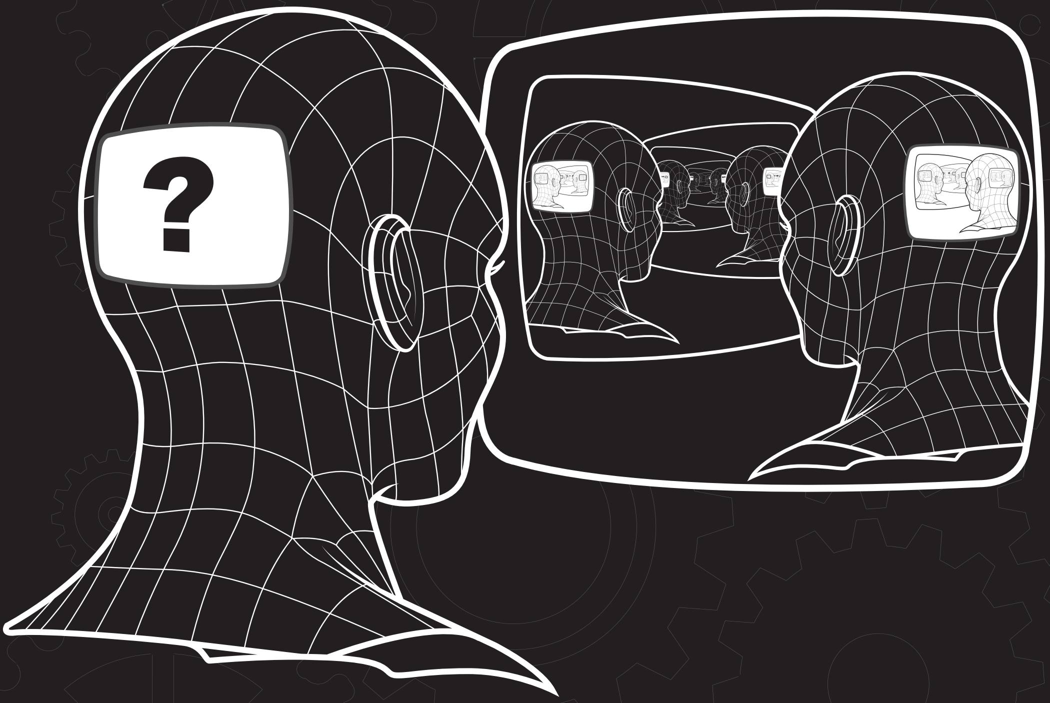

Part II

Preferential access to consciousness of behaviorally relevant visual input

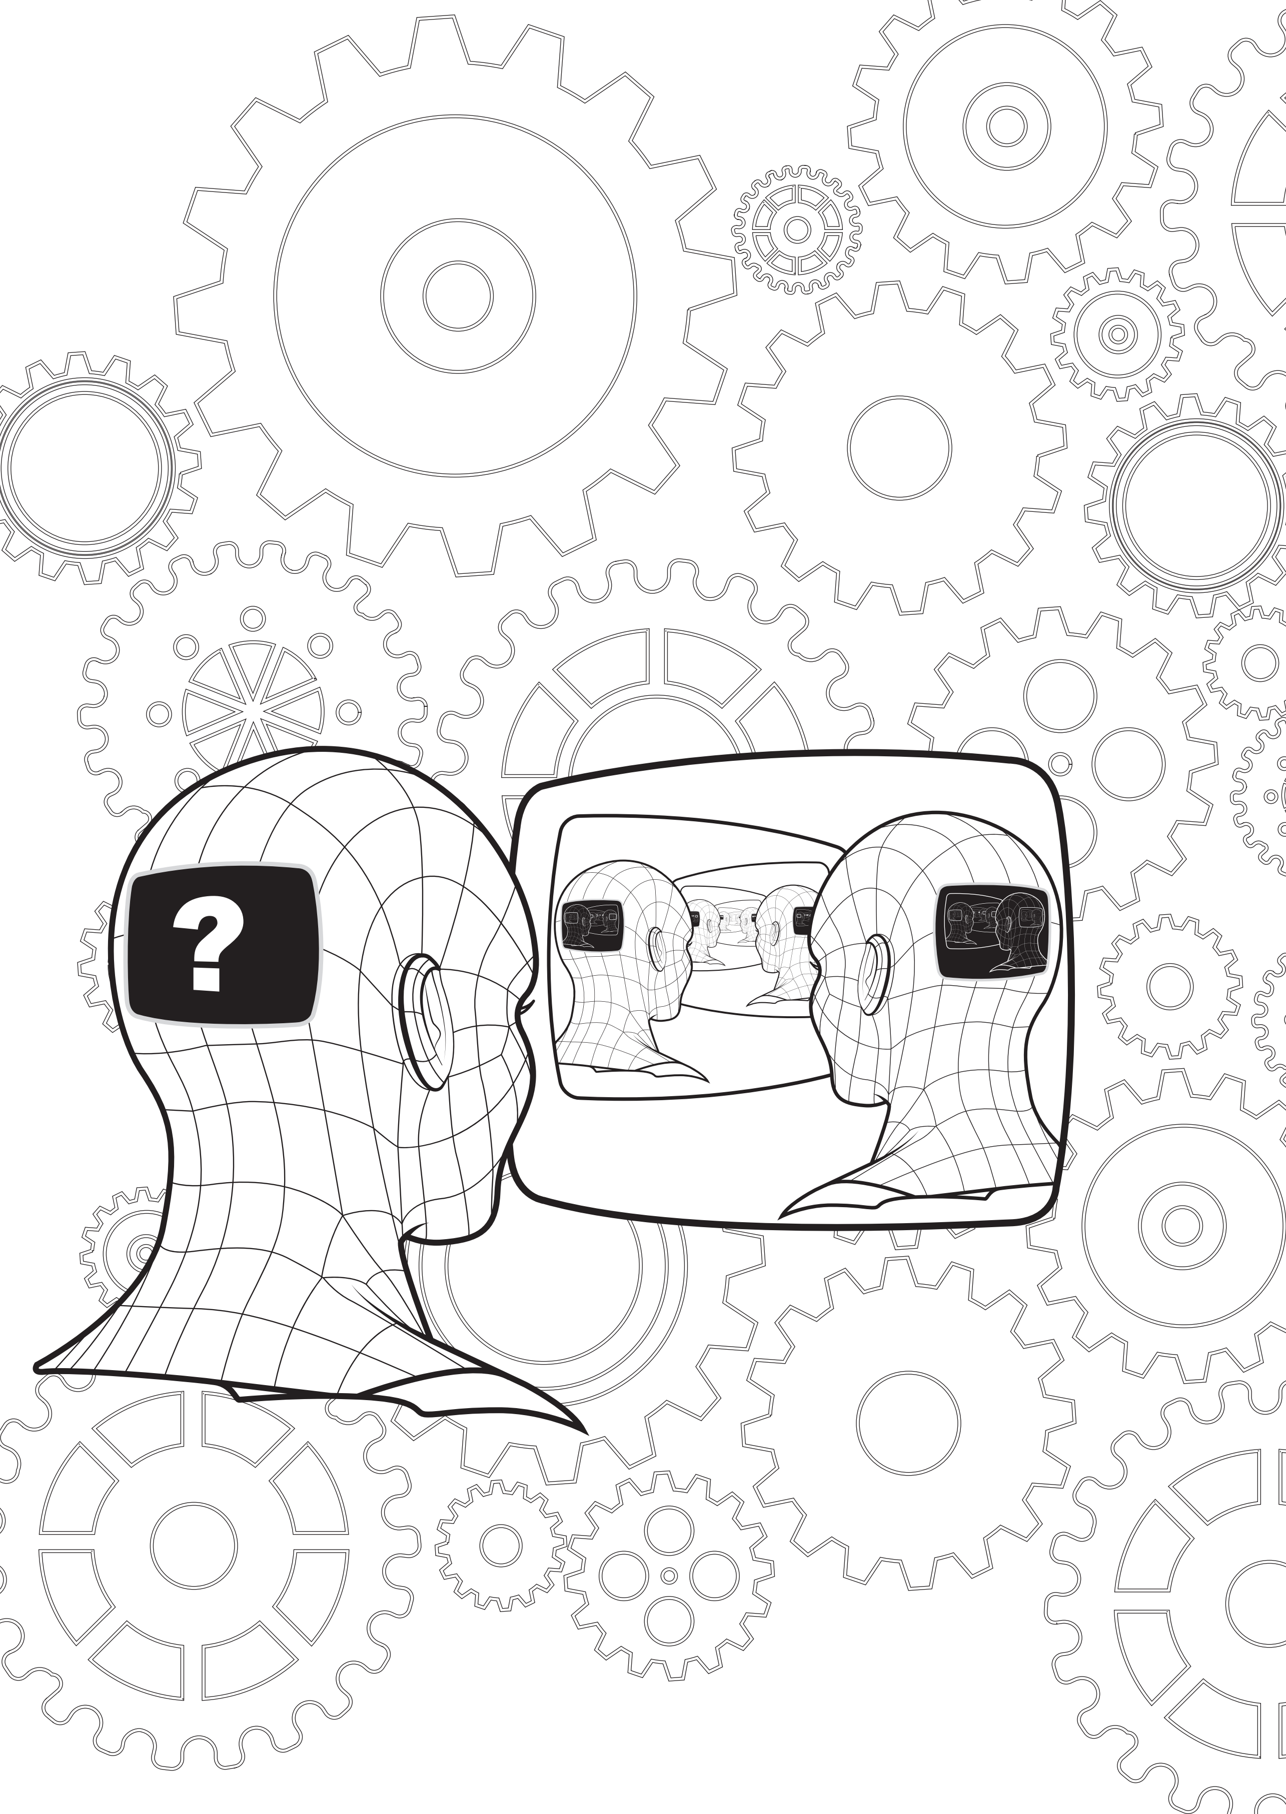

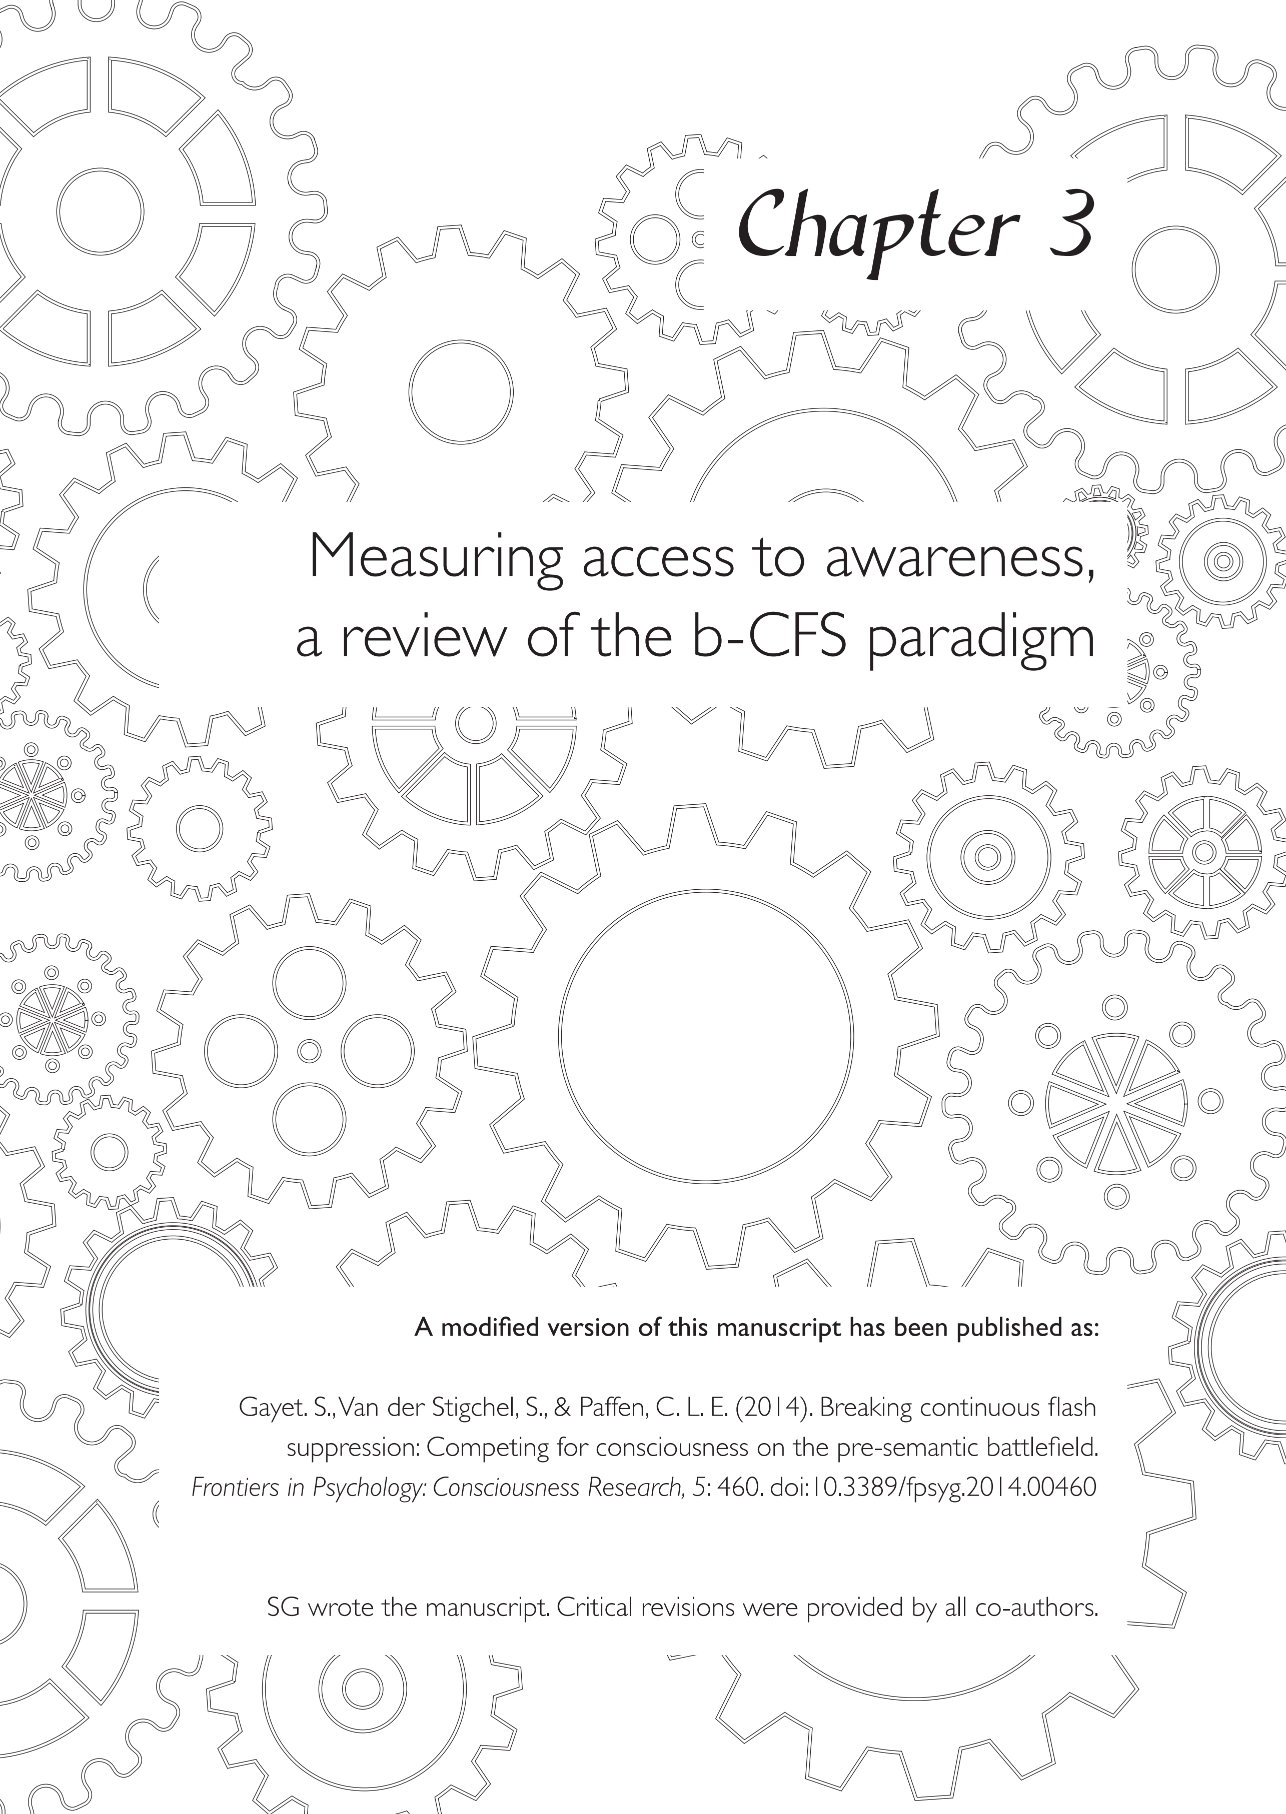

Chapter 3

Measuring access to awareness, a review of the b-CFS paradigm

A modified version of this manuscript has been published as:

Gayet. S., Van der Stigchel, S., & Paffen, C. L. E. (2014). Breaking continuous flash suppression: Competing for consciousness on the pre-semantic battlefield. *Frontiers in Psychology: Consciousness Research, 5*: 460. doi:10.3389/fpsyg.2014.00460

SG wrote the manuscript. Critical revisions were provided by all co-authors.

Abstract

Traditionally, interocular suppression is believed to disrupt high-level (i.e., semantic or conceptual) processing of the suppressed visual input. The development of a new experimental paradigm, breaking continuous flash suppression (b-CFS), has caused a resurgence of studies demonstrating high-level processing of visual information in the absence of visual awareness. In this method the time it takes for interocularly suppressed stimuli to breach the threshold of visibility, is regarded as a measure of access to awareness. The aim of the current review is twofold. First, we provide an overview of the literature using this b-CFS method, while making a distinction between two types of studies: those in which suppression durations are compared between different stimulus classes (such as upright faces versus inverted faces), and those in which suppression durations are compared for stimuli that either match or mismatch concurrently available information (such as a colored target that either matches or mismatches a color retained in working memory). Second, we aim at dissociating high-level processing from low-level (i.e., crude visual) processing of the suppressed stimuli. For this purpose, we include a thorough review of the control conditions that are used in these experiments. Additionally, we provide recommendations for proper control conditions that we deem crucial for disentangling high-level from low-level effects. Based on this review, we argue that crude visual processing suffices for explaining differences in breakthrough times reported using b-CFS. As such, we conclude that there is as yet no reason to assume that interocularly suppressed stimuli receive full semantic analysis.

Introduction

Interocular competition

When different images are presented to both eyes, observers tend to perceive only one of these images, whereas the other one does not give rise to a conscious percept (e.g., binocular rivalry, Alais, & Blake, 2005; flash suppression, Wolfe, 1984; continuous flash suppression, Tsuchiya, & Koch, 2005). Under certain conditions the suppressed image has the potency to affect behavior, but this depends on the required level of processing (for a review see, Lin, & He, 2009). For instance, the potency of low-level image properties, such as spatial frequency (Blake, Tadin, Sobel, Raissian, & Chong, 2006; Blake, & Fox, 1974), motion direction (Wade, & Wenderoth, 1978; O'Shea, & Crassini, 1981; Blake, Ahlstrom, & Alais, 1999), color (White, Petry, Riggs, & Miller, 1978) and orientation (Wade, & Wenderoth, 1978) to elicit behavioral adaptation effects is relatively unaffected by interocular suppression. Conceptual or semantic processing, however, is traditionally believed to be abolished for interocularly suppressed stimuli (e.g., Zimba, & Blake, 1983; Cave, Blake, & McNamara, 1998; Blake & Logothetis, 2002; Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006; Kang, Blake, & Woodman, 2011). In general, the extent to which neural activity reflects interocularly suppressed stimulation decreases gradually when climbing up the visual hierarchy (Blake, & Logothetis, 2002). For instance, most cells in early visual areas (80% in V1/V2 and 60% in V4/V5) respond to stimulation of either eye irrespective of the dominant percept (Logothetis, 1998). Higher processing areas such as IT, FFA and PPA, however, follow mostly (although not exclusively; Fang, & He, 2005; Jiang, & He, 2006; Sterzer, Haynes, & Rees, 2008) the dominant percept (Tong, Nakayama, Vaughan & Kanwisher, 1998). Thus, interocularly suppressed stimuli are expected to be processed at the level of features and coarse feature configurations, which we will refer to as the lower or visual processing level, but not at a semantic or conceptual level (Blake, & Logothetis, 2002), which we will refer to as higher level.

In contrast to this traditional view, studies using a novel paradigm called breaking continuous flash suppression (b-CFS; Jiang, Costello & He, 2007) seem to demonstrate that high-level processing of interocularly suppressed stimuli can occur prior to conscious experience. In the present article we aim to demonstrate that the seemingly high-level effects obtained in these b-CFS studies can be accounted for by coarse visual processing of the stimuli under continuous flash suppression (CFS). For this purpose, we provide a complete overview of all studies up to date

(30) using b-CFS. Additionally, we suggest a number of improvements to the b-CFS method that help dissociate competition at relatively high levels of processing (i.e., at a conceptual or semantic level) from competition at lower levels of processing (i.e., at a featural level, where color, orientation, etc. are processed).

Breaking continuous flash suppression

In the b-CFS paradigm, a high contrast dynamic pattern mask is presented to one eye, thereby effectively suppressing a stimulus of increasing intensity presented to the other eye. Eventually, the ocular dominance will reverse, such that the previously suppressed stimulus becomes visible. The time it takes for observers to detect the suppressed stimulus is assumed to reflect the moment in time at which the stimulus gains access to consciousness. Importantly, non-ocular factors can affect the moment at which interocularly suppressed stimuli become consciously observable (Paffen, & Alais, 2011; Blake, 2001). In light of the b-CFS paradigm, we dissociate two factors that co-determine the timing of an ocular dominance reversal. First, some stimulus classes might inherently breach the threshold of visibility faster than other stimulus classes (e.g., upright versus inverted faces; Jiang et al., 2007). Second, suppression durations can systematically differ for stimuli that either match or mismatch consciously accessible information (e.g., prime-target congruency; Costello, Jiang, Baartman, McGlennen, & He, 2009). In reviewing the b-CFS literature we propose to take into account these two distinct ways in which non-ocular factors impinge upon the selection for conscious access: manipulations of the content of the suppressed stimulus, and manipulations of the context within which the suppressed stimulus is presented. As both types of experiments have their own advantages and limitations in uncovering the nature of preconscious processes, they are discussed separately.

Effects of stimulus context

Priming

The first part of this review comprises an overview of b-CFS studies in which the detection time of identical stimuli is compared between different experimental conditions. These studies show that the same visual input can result in different suppression durations depending on the (consciously accessible) context that is provided. One widely studied way to affect the context within which information is presented is priming. This method involves presenting a stimulus prior to the b-CFS task, which is either related or unrelated to the masked target stimulus. Costello et al. (2009) showed that written words (e.g., "fire") break through suppression faster when they are preceded by a word that shares sub-word components (e.g., "tire") than when they are preceded by a word that doesn't share sub-word components. Costello and colleagues also showed that words break through suppression faster when they are preceded by a semantically related word (e.g., "burn") than when they are preceded by an unrelated word. Lupyan and Ward (2013) took this one step further by showing that this priming effect also occurs when prime and target are presented in different modalities; for instance, an image of a pumpkin broke through suppression faster after observers heard the word "pumpkin" than after hearing a word that did not match with the subsequent target. Yang and Yeh (2014) presented words under CFS, of which the onset was either accompanied by an audible white noise burst or not. Detection times were shortened by the concurrent presentation of noise bursts, but only when the audio and visual information originated from the same depth plane. Together, these priming studies reveal that visual input that matches previously perceived information breaks through suppression faster than visual input that mismatches this information. Importantly, the prime-target relation can be spatial, physical or semantic in nature, and does not require presentation in the same modality.

The content of visual working memory

Similarly to priming, the content of visual working memory is also known to affect visual processing, such that stimuli matching this content receive privileged processing compared to non-matching information (e.g., in search tasks, Olivers, Meijer, & Theeuwes, 2006). One major difference between these two methods is that visual working memory involves the active, rather than passive maintenance (i.e., rehearsal) of visual features. In experiments that manipulate the content of visual working memory, participants are instructed to retain some feature of a visual stimulus for subsequent recollection. During the retention phase, participants perform a b-CFS task in which interocularly suppressed targets either match or mismatch the information that is concurrently retained in working memory. Recently, it has been shown that target stimuli under CFS are detected faster when they match rather than mismatch a color category (Gayet, Paffen, & Van der Stigchel, 2013), an orientation (Liu, Wang, & Jiang, 2013) or a face (Pan, Lin, Zhao, & Soto, 2013) that is actively held in visual working memory. Crucially, detection times remain unaffected when the stimuli, otherwise used for the memory task, are passively viewed, as opposed to actively retained in working memory. In contrast with the priming studies discussed previously, Gayet et al. (2013) demonstrated that privileged detection of matching stimuli was only observed when the relevant stimulus dimension was retained; when participants retained the shape of a stimulus, targets that matched the color of that stimulus were not prioritized for conscious detection. Together, these working memory studies show that visual input that matches concurrently retained, task relevant information is accessible to consciousness faster than non-matching information.

Simultaneous cross modal priming

Three recent studies used a methodological approach in which the manipulation of the context was longer lasting than that of priming studies, without involving the active retention of information as in the working memory studies. In these experiments, consciously accessible, non-visual information was concurrently presented with a b-CFS task. First, Zhou, Jiang, He, and Chen (2010) demonstrated that images matching olfactory information (e.g., an image of a rose concurrently presented with the scent of a rose) break through suppression faster than images mismatching olfactory information (e.g., an image of a rose concurrently presented with the scent of butanol). Second, Alsius and Munhall (2013) showed that an interocularly suppressed talking face stimulus broke through suppression faster when an auditory sentence matched rather than mismatched the lip synchronization of the face. Finally, Salomon, Lim, Herbelin, Hesselmann, and Blanke (2013) showed an effect of proprioception on visual awareness. In their study, participants reported the orientation of an interocularly suppressed target, which was superimposed on a task-irrelevant image of a hand. This hand could either be congruent or incongruent with the participants' actual position. Targets broke through interocular suppression faster when the image of the hand matched the position of the real hand. The authors conclude that proprioception modulates the selection for conscious access of visual stimuli. Taken together, these studies show an advantage for detecting stimuli that match rather than mismatch consciously accessible information.

Visual versus conceptual analysis of suppressed stimuli

The major advantage of all b-CFS experiments described in the "stimulus context" Section is that differences in suppression durations cannot be accounted for by differences in image characteristics between conditions. This follows from the fact that in all conditions the same stimuli are used as target stimuli under CFS.

The differentiation between conditions stems purely from the relation between target stimuli and the consciously accessible context in which they are embedded. Arguably, this context biases the competition by boosting or diminishing the effective strength of the suppressed stimuli (for a similar interpretation for attention's effect on interocular suppression, see Paffen, & Alais, 2011). The authors of the papers described above generally interpret their findings in terms of pre-activation of prime related information (either semantic or physical), which biases subsequent interocular competition (e.g., Lupyan, & Ward, 2013). In this view, prime induced activity in areas further up the processing hierarchy (e.g., object selective areas) feeds back to the earlier visual cortex where the interocular competition is resolved (e.g., Tong, 2001; Blake, 1989). Note, however, that this interpretation cannot provide a satisfactory account for the semantic priming effect of Costello et al. (2009), which requires semantic analysis of the prime as well as the suppressed target. This issue is further discussed in the Discussion.

The assumption that competition occurred at the level of simple stimulus features rather than at the semantic or conceptual level was explicitly tested by Lupyan and Ward (2013) in a second experiment. Here, participants were cued with either the word "square" or "circle", before performing a b-CFS task. By using a wide range of stimulus shapes ranging on a continuum from square to circle, they found that the similarity between the target stimulus and the cued shape was negatively correlated with the detection time of the target stimulus. The authors conclude from this finding that upon hearing (or reading) a word, a visual representation of its content is automatically activated. This active representation then facilitates subsequent detection of matching visual input. As such, the effects of semantic primes on suppression durations of subsequently presented targets are visual, rather than semantic in nature. The major advantage of this interpretation is that it allows for semantic priming, in the absence of semantic analysis of the suppressed stimulus.

Further support for this idea of feature pre-activation comes from the working memory experiments described above. When observers actively retain stimulus features, such as an orientation, these features can be decoded from activity in the early visual cortex (Harrison & Tong, 2009; Serences, Ester, Vogel, & Awh, 2009). Thus, the abovementioned working memory studies allow for comparing between the situation in which prime-induced activity is retained and conditions in which prime-induced activity is discarded. The absence of an effect of the prime on suppression durations when the prime is perceived but not actively retained suggests that the prime-target congruency effects are indeed caused by pre-activation of prime induced features.

Together, the findings in this chapter show that providing a consciously accessible context prioritizes visual information that matches this context. As argued earlier, the consciously accessible context might activate a visual representation, which then interacts with the interocularly suppressed target. As such, even if the relation between the context and the suppressed target is semantic in nature, semantic analysis of the target is not *necessary* for detection times to be affected. One of the drawbacks of this type of b-CFS experiment is, however, that it does not allow for unequivocally excluding the possibility that the interocularly suppressed stimulus is processed up to a semantic level. In contrast, when comparing the potency of different stimulus classes in reaching visual awareness without providing a context, any difference in detection times between conditions (either featural or semantic in nature) reflects differences in the processing of the suppressed stimulus itself, rather than its interaction with a previously altered neural state. Studies using this approach will be discussed in the following paragraphs.

Effects of stimulus content

Visual characteristics

The second type of b-CFS experiments compares detection times between different stimulus categories. This comprises the comparison of stimulus categories that differ on the basis of relatively low-level visual properties that can be resolved in the early visual cortex, which will be discussed in this first section. For these stimulus properties, there is a tendency that more conspicuous stimuli are harder to suppress by CFS and, as such, break through suppression faster than less conspicuous stimuli. For instance, both higher contrast stimuli (Tsuchiya, & Koch, 2005) and higher spatial frequency stimuli break through CFS more readily (Tsuchiya, & Koch, 2005; Yang, & Blake, 2012). Also, certain topological properties of interocularly suppressed stimuli elicit faster detection times than others. For instance, suppressed stimuli with a hole are detected faster than open stimuli made up of the same structural elements (Meng, Cui, Zhou, Chen, & Ma, 2012). When identical stimuli follow different motion patterns, this may result in different detection thresholds as well. For instance, coherently moving dot arrays break through suppression more often than random dot arrays that are presented for the same duration (Kaunitz, Fracasso, Lingnau, & Melcher, 2013). Climbing further up the visual hierarchy, images with strong grouping cues, such as Kanisza triangles are detected faster than non-Kanisza's made up of the same constituents (Wang, Weng, & He, 2012). Together, these studies show that different stimuli yield different suppression durations, and that this effect might be linked to the saliency of the suppressed stimulus. This is in line with findings from binocular rivalry experiments, which demonstrate that the location at which a perceptual transition is initiated depends on the local saliency of the suppressed stimulus (Paffen, Naber, & Verstraten, 2008; Stuit, Verstraten, & Paffen, 2010).

Differences in suppression durations between stimulus categories can be accounted for both by properties of the suppressed stimuli per se, and by interactions between properties of the stimuli and properties of the masks (for a discussion, see Stein, Hebart, & Sterzer, 2011). We dissociate two types of interactions between the stimuli and the masks that can potentially affect suppression durations. First, increased differences between visual characteristics of the suppressed image and the CFS stimuli reduce the suppression strength. For instance, Yang and Blake (2012) showed that stimuli with oblique orientations broke through suppression faster than stimuli with cardinal orientations, when using traditional "Mondrians" as CFS stimulus (which contain only cardinal orientations). More specifically, greater similarity in spatial frequency content and orientation between the competing percepts led to stronger suppression in both b-CFS (Yang, & Blake, 2012) and binocular rivalry (Stuit, Paffen, van der Smagt, & Verstraten, 2011). Second, when the previously suppressed image (or a sub-part of it) breaks through suppression, detection is facilitated if the suppressed image and the masks are very different. As discerning a suppressed stimulus through a mask requires exceeding some threshold of certainty, stimuli with more 'proof' of being a potential target have an advantage in breaking CFS (for similar interpretations, see Kaunitz et al., 2013; and Yang, & Yeh, 2014). Such a bias could be underpinned by the phenomenon of piecemeal rivalry, which allows for perceiving local parts of the 'suppressed' stimulus (Blake, O'Shea, & Mueller, 1991; O'Shea, Sims, & Govan, 1996). Since the dominant percept is highly dynamic (i.e., the CFS masks), locally dominant stimulus parts from the non-dominant eye (in which the target is presented) are easily confused with the CFS masks, and thus disregarded. However, when piecemeal rivalry reveals stimulus parts that seem coherently related (e.g., they follow a particular pattern or movement direction), these stimulus parts may attract attention, as they are likely to be the target (e.g., collinear facilitation; Wilson, Blake, & Lee, 2001). This may affect suppression durations, since attending to a stimulus in a specific eye enhances the competition strength of the entire ipsi-ocular stimulus (Zhang, Jiang, & He, 2012; Ooi, & He, 2009).

Nonetheless, differences between aforementioned conditions do not necessitate non-conscious semantic or conceptual processing, but are based on the differentiation of stimulus properties that are generally assumed to survive interocular suppression (for an overview, see Lin, & He, 2009; Blake, & Logothetis, 2002). From the next paragraph onwards, a number of studies will be discussed in which suppression durations are affected on the basis of higher level stimulus properties (i.e., at a semantic or conceptual level). Please note that the "familiarity" and "ecological relevance" distinction, as provided below, aims at categorizing these studies based on topical similarities, rather than describing the mechanisms that drive their results.

Familiarity

Differences in detection times between stimulus categories can also arise on the basis of more high-level distinctions, such as stimulus familiarity. For instance, images of human bodies or body parts are detected faster when presented upright as compared to inverted (Stein, Sterzer, & Peelen, 2012). As the authors demonstrate that this latter effect was abolished when the images were distorted, the authors argue that the difference in detection times is accounted for by the greater familiarity of upright human bodies. Along the same lines, upright faces are detected faster than inverted faces (Jiang et al., 2007; Stein et al., 2011a; Zhou, Zhang, Liu, Yang, & Qu, 2010; Stein, Peelen, & Sterzer, 2011; Gray, Adams, Hedger, Newton, & Garner, 2013). Two of these studies (Stein et al., 2011b; Gray et al., 2013) also included a polarity inversion condition, demonstrating that detection times were fastest for normal faces (upright and normal polarity) and slowest for the most unusual face presentation condition (spatial inversion and inversed polarity), although the inversion effect was only marginally significant in the inversed polarity condition of Stein et al. (2011b). The finding that face inversion effects are dependent on (or additive to) manipulations of the contrast polarity, supports the idea that it is indeed familiarity that drives the priority for detecting upright faces. Importantly, however, Stein et al. (2011b) replicated these findings with configurations of three blobs representing two eyes and a mouth. This demonstrates that the privilege for detecting upright faces can be resolved by very crude visual processing.

Gobbini and colleagues took the manipulation of stimulus familiarity even further by showing that interocularly suppressed familiar faces are detected faster than faces of strangers (Gobbini, Gors, Halchenko, Rogers, Guntupalli, Hughes, & Cipollo, 2013). A more subtle finding comes from a study showing that faces from the own racial in-group break through suppression faster than faces from the racial out-group (Stein, 2012). That same study showed that faces of the same age group as that of the observer break through suppression faster than faces of another age group. Importantly, the differences in suppression durations between image conditions were computed relative to that of inverted faces, such that they could not be attributed to differences in low-level image properties (see the Discussion). Rather, the authors suggest that this effect is accounted for by the observer's greater visual expertise with stimuli of the own-race and own-age stimulus classes.

This facilitatory effect for detecting visual input of higher familiarity is also found for stimuli that are more recently acquired in evolutionary time, such as written language. Indeed, words in a familiar alphabet are detected faster than words in an alphabet that is unfamiliar to the observer (Jiang et al., 2007). Similarly, Chinese characters are detected faster by Chinese readers compared to the same characters that have been inversed or scrambled (Yang, & Yeh, 2011; Yang, & Yeh, 2014). Taken together, these studies show that visual input with higher stimulus familiarity is more readily detected than less familiar input. Arguably, extended experience with certain types of stimuli might facilitate subsequent detection. If so, the factor of familiarity might be the long term equivalent of the stimulus feature pre-activation as described in the "stimulus context" Section.

Ecological relevance

A number of studies demonstrate differences in detection times for stimuli that differ on the basis of ecological relevance. For instance, observers show an advantage for detecting faces turned towards the observer compared to faces turned slightly away from the observer (Gobbini, Gors, Halchenko, Hughes, & Cipollo, 2013). This difference was found to be independent of the gaze direction of the face. Similarly, faces with direct gaze break through interocular suppression faster than faces with averted gaze. This was found both for schematic faces (Chen, & Yeh, 2012) and for face photographs (Stein, Senju, Peelen, & Sterzer, 2011). This advantage for detecting faces with direct gaze could not be explained by (lower-level) effects of eye symmetry, as Stein et al. (2011c) included images of both frontal faces and laterally averted faces, such that gaze direction should be inferred by the particular combination of both face orientation and pupil position. However, the advantage in detecting stimuli with direct gaze over averted gaze persisted for inverted faces (Stein et al., 2011c; Chen, & Yeh, 2012). Gaze direction in (visible) faces is more difficultly inferred from inverted faces compared to upright faces (e.g., Vecera, & Johnson, 1995). Thus, the effect of gaze direction on detection times should be less prominent in the inverted condition than in the upright condition. The absence of this interaction between gaze direction and face inversion therefore hints toward the interpretation that crude configural differences between gaze conditions might play a causal role in eliciting these differences in detection times. For instance, Chen and Yeh (2012) propose that the specific conjunction of face curvature and pupil location is sufficient in eliciting shorter suppression durations. In line with this idea, they demonstrated in an additional experiment that the mere schematic depiction of eyes was sufficient in explaining the observed difference in detection times of full (schematic) faces.

Another ecologically potentially relevant distinction between stimulus categories is that of emotional versus non-emotional stimuli. For instance, fearful faces break through suppression faster than neutral faces (Yang, Zald, & Blake, 2007; Gray et al., 2013; Stein, Seymour, Hebart, & Sterzer, 2014) or happy faces (Tsuchiya, Moradi, Felsen, Yamazaki, & Adolphs, 2009; Yang et al., 2007; Gray et al., 2013), while happy (Yang et al., 2007) and angry faces (Gray et al., 2013) break through suppression *slower* than neutral faces. Interestingly, both types of emotional expressions break through suppression *faster* than neutral faces when schematic face images are used instead of face photographs (Stein, & Sterzer, 2012). This contradiction suggests that it is not the analysis of the emotional valence per se, but rather the visual properties of the image that affected suppression durations in these studies. In line with this lower level account, the findings of Gray et al. (2013) persisted for inverted faces and for faces with inversed polarity, while the findings of Yang et al. (2007) persisted for inverted faces and for eyes-only images. Similarly, the findings of Stein and Sterzer (2012) were fully accounted for by the relative orientation of the mouth curvature and the face contour. Finally, the findings of Stein et al. (2014) depended solely on high spatial frequency information. Since subcortical (i.e., amygdala) processing of fearful faces relies predominantly on low spatial frequency information (e.g., Ledoux, 2001), this finding suggests that pre-conscious processing of fearful faces is dependent on cortical processing. Patient SM, who has complete bilateral amygdala lesions and is unable to consciously discriminate between fearful and happy faces, showed the same advantages for detecting CFS-suppressed fearful faces over happy faces as controls did (Tsuchiya et al., 2009). As such, pre-conscious discrimination between emotional faces seems to rely more on (cortical) extraction of characteristic visual features, than on the (subcortical) analysis of the emotional valence per se. Taken together, these studies show a tendency for ecologically relevant stimuli to break through interocular suppression faster than less ecologically relevant stimuli. However, most of these effects have been shown to rely on stimulus properties, or stimulus configurations, that can be dissociated on the basis of relatively crude visual processing. In sum, semantic, conceptual or emotional analysis of interocularly suppressed stimuli is not a *necessary* condition to account for the observed differences in detection times. Rather, the extraction of purely visual information seems to sufficient to explain most of the findings discussed so far.

Climbing towards the semantic and conceptual level

As with the privilege for detecting familiar stimuli, the privilege for detecting emotional stimuli was not restricted to evolutionarily old visual input, such as faces, but was also demonstrated for words (Yang, & Yeh, 2011). Interestingly, the results of this study revealed that both (Chinese) words that *describe* a negative emotion (e.g., "anger" or "fear") and words that *induce* a negative emotion (e.g., "murder" or "abuse") were detected *later* than neutral words. Taking this idea even further, Sklar, Levy, Goldstein, Mandel, Maril, and Hassin (2012) compared suppression durations of emotionally negative expressions to suppression durations of neutral expressions. Importantly, the words that formed these expressions had no intrinsic emotional valence (e.g., "eternal" and "rest"; "eternal rest"). Nonetheless, the expressions with a negative emotional valence broke through suppression faster than neutral expressions. Interestingly, these results are at odds with that of Yang and Yeh (2011). Still, both studies demonstrate effects that require semantic processing of the words before interocular competition is resolved.

Sklar and colleagues (2012) also demonstrated that combinations of (Hebrew) words that yield incoherent expressions (e.g., "she ironed coffee") broke through suppression faster than coherent expressions (e.g., "she drank coffee"). Again, it is the semantic combination of words that determines whether an expression is coherent or incoherent, rather than the individual words themselves. This finding demonstrates that the meaning of words is indeed extracted and integrated non-consciously. Along the same lines, Mudrik, Breska, Lamy, and Deoull (2011) showed that scenes containing incongruent objects (e.g., Michael Jordan holding a watermelon) broke through suppression faster than the same scenes containing congruent objects (e.g., Michael Jordan holding a basketball). The authors stress that dissociating a coherent from an incoherent image requires the integration of an object in its semantic context; a process originally thought to require consciousness (e.g., Tononi, & Edelman, 1998). In contrast with the familiarity effects discussed earlier, the stimuli used in these last two experiments seem too complex for the differences in suppression durations to be accounted for by differences in visual experience between stimulus conditions. As such, these results imply full blown semantic analysis of interocularly suppressed stimuli.

Discussion

Assessing the level of processing

Most findings in this review can be explained by preconscious analysis of suppressed stimuli at relatively early stages of visual processing. Whether they are caused by pre-activation of primed features, by the saliency of a stimulus, or by the long time strengthening of visual representations of relevant feature configurations, these findings do not seem to require semantic or conceptual processing. In contrast, the findings discussed at the end of the "stimulus content" Section, as well as the semantic priming effect of Costello et al. (2009), seem to defy the model of early competition in interocular suppression and point to high-level analysis of the suppressed stimuli. There are, however, two reasons to plead for caution in interpreting the studies that demonstrate these high-level effects (e.g., language and scene comprehension). First, some of the results described above seem contradictory, such as the results of Sklar et al. (2012) in which negative emotional expressions yielded shorter suppression durations compared to the results of Yang and Yeh (2011) in which negative emotional words yielded longer suppression durations. In a broader sense, the overall pattern of findings of these high-level effects seems inconsistent with the pattern of findings from lower level effects. On the one hand, words and images break through suppression faster when they have a higher prevalence in the observers' visual world (i.e., when they are of higher familiarity). On the other hand, however, word combinations and complex scenes break through suppression faster when they when they are incongruent or novel, and thus are of lower familiarity. While it is conceivable that scene complexity influences the magnitude of the effect of familiarity on suppression durations, it is unexpected that scene complexity causes a reversal in the direction of the effect of familiarity on suppression durations.

Second, to demonstrate that differences in suppression durations are caused by competition at a high processing level (i.e., semantic or conceptual), it is important to implement a comparison with a condition that disrupts high-level processing, such as inversion (e.g., as used in the Sterzer lab), polarity inversion or scrambling. If the difference in suppression durations observed under normal presentation conditions is also apparent in these conditions, it is likely that the effect is caused by differences in lower level visual properties between the stimulus classes (see the Discussion). Four out of five studies that do include this type of control conditions to dissociate between competition at higher processing levels from competition at lower (visual) processing levels, demonstrated that the effect could indeed be attributed to competition at lower levels of the processing hierarchy (Chen, & Yeh, 2012; Gray et al., 2013; Stein, & Sterzer, 2011; Stein et al., 2011c). Consequently, these studies do not attribute their findings to high level processing under continuous flash suppression. Importantly, three out of four b-CFS experiments that led the authors to conclude from their data that the observed difference in suppression durations was caused by semantic or conceptual analysis of the stimuli under CFS, however, did not include such a control condition (i.e., Sklar et al., 2012; Mudrik et al., 2011; Costello et al., 2009). Thus far, the only study that convincingly demonstrates high-level competition in a b-CFS experiment, is that of Yang and Yeh (2011). In this study, the authors included an inversion condition, a scrambled condition as well as a monocular condition. This revealed that the shorter suppression durations for neutral Chinese words compared to emotional Chinese words was only apparent in the upright unscrambled dichoptic condition.

In sum, more and more studies aim at demonstrating that semantic and conceptual information might be integrated pre-consciously. However, this is hard to reconcile with studies showing that semantic priming effects are abolished under interocular suppression (e.g., Kang, Blake, & Woodman, 2009; Cave, Blake, & McNamara, 1998; Zimba, & Blake, 1983; for a review, see Lin, & He, 2009). In some studies interocular suppression is even used as a tool to disrupt semantic processing (e.g., Lupyan, & Ward, 2013). These high-level effects are also hard to reconcile with the idea that interocular competition is resolved in early visual areas such as LGN (Haynes, Deichmann, & Rees, 2005) and V1 (Polonsky, Blake, Braun, & Heeger, 2000). Although some interocularly suppressed information is known to transpire into higher visual areas (e.g., Fang & He, 2005; Jiang & He, 2006; Sterzer et al., 2008), succeeding levels in the processing hierarchy reveal less and less brain activity that reflects interocularly suppressed stimulation (Blake, & Logothethis, 2002). Moreover, CFS is known to result in greater suppression depths than more traditional methods of interocular suppression, such as flash suppression and binocular rivalry (Tsuchiya, Koch, Gilroy, & Blake, 2006). Consequently, when b-CFS is used to compare different classes of stimuli in their potency to breach the threshold of awareness, it is of utmost importance to test whether reaction times indeed reflect differences in high-level rather than low-level information in the stimuli. Additionally, irrespective of the processing level at which the competition takes place, it is crucial to assert whether reaction times indeed reflect differences in the timing at which a stimulus was available to consciousness, rather than processes arising after the stimulus became available to consciousness. These post-perceptual effects pose a threat to b-CFS experiments in which the stimulus content is manipulated as well as to experiments in which the stimulus context is manipulated. We propose that at least the following three control conditions should be included in b-CFS experiments to control for these potential pitfalls.

Control 1: Disrupting the extraction of meaning

To assess whether differences between conditions rely on high-level information (i.e., at a semantic or conceptual level), one or more conditions should be included that are known to disrupt the extraction of high-level image properties, while keeping low-level (i.e., visual) image properties relatively unaffected. This can be achieved by such manipulations as inverting the image or inverting the image polarity (e.g., Gray et al., 2013; Jiang et al. 2007; Chen, & Yeh, 2012; Stein et al., 2011; Zhou et al., 2010). These manipulations constrain the extraction of meaning from an image (Rock, 1974; Shore, & Klein, 2000), such that high-level driven effects should at least diminish under these circumstances. As such, if some image class breaks through suppression faster than another stimulus class because of high-level (i.e., semantic or conceptual) differences, the differences in detection times between these two stimulus classes should not be observed (or at least diminish) when the images are presented upside down. Conversely, if the difference in detection times between two stimulus classes does persist with inverted presentation, this suggests that there are systematic low-level visual differences between the two image classes, as these should remain unaffected by inverted presentation. In that case, the differences in low-level visual properties are the probable cause of the difference in detection times between the two stimulus classes. For this reason, rather than looking at absolute detection times for each stimulus category, it is more informative to look at the inversion effect, which is described as the difference in detection times between upright and inverted stimuli of the same stimulus category. This difference can then be divided by the detection time of inverted stimuli (as in Stein, 2012) such as to remove the between subject variability in detection times. Consequently, to assess whether a difference between image classes relies on high-level stimulus processing, it is important to demonstrate that the inversion effects (rather than the detection times per se) differ between stimulus classes.

Control 2: Stimulus reportability

Next, it is important to verify whether differences in reaction times indeed reflect differences in visual awareness. An alternative view is that differences in reaction times are driven by non-conscious processes, such that stimulus information is accessible only to the extent that it affects forced choice localization, while not being accessible to subjective report. Arguably, a stimulus fails to reach visual awareness, if it is accessible to one output system, but not to the other (Baars, 1993; Kanwisher, 2001). Thus, in order to conclude that some manipulation in a b-CFS experiment affects visual awareness, visual awareness should be measured directly. Visual awareness of a stimulus is assumed to be a prerequisite for stimulus reportability (Dehaene et al., 2006). As such, it can be operationalized as the ability to subjectively report ones percept (Weiskrantz, 1997; Dennett, 1993; Dehaene, 2001). For the present purpose, a direct way to test whether one stimulus was accessible to consciousness and the other was not, is to compare participants' ability to report the identity of two concurrent stimuli at a particular point in time. This objective measure of stimulus reportability can be implemented by presenting two stimuli of different conditions simultaneously (e.g., one at either side of fixation). After participants perform a speeded detection of the location at which (e.g., left or right of fixation) they first see a stimulus appear, they should report either the identity of the percept on the reported location, or of that on the non-reported location. If no post-detection strategic bias is involved, participants should be significantly worse at reporting the identity of the stimulus on the non-reported location as compared to that of the reported location. Conversely, if participants are equally proficient at reporting the identity of either stimulus, one may not conclude that there was a difference in conscious access between stimulus conditions.

Control 3: Post detection effects

Finally, it is important to assess whether differences in detection times indeed reflect differences in interocular suppression durations, rather than processing differences arising after conscious detection of the stimulus (e.g., a difference in response criterion). To account for these 'late' effects it is imperative to add a monocular (or binocular) control condition, in which the 'suppressed' stimulus and the CFS are presented to the same eye(s). Specifically, we advocate the use of two different monocular control conditions (as in, Gayet et al., 2013; and Costello et al., 2009). First, a monocular control condition is needed in which the presentation times are identical to that of the interocular condition, such as to keep the stimulus chronology constant (i.e., a physically similar control). The disadvantage of this condition is, however, that reaction times in this condition are much faster than in the interocular condition. Consequently, any differences in reaction times between conditions are reduced in magnitude as well, as a result of which the experimental power can be diminished (although the variance is reduced as well). Thus, it is imperative to implement a second monocular condition such, that the reaction time distributions (means and SD's) match that of the interocular condition (for further discussion on this issue, see Stein et al., 2011a). This can be achieved by (1) lengthening the ramp of the 'suppressed' stimulus, such as to mimic the longer suppression durations of trials with dichoptic presentation, and, (2) by jittering the target onset, such as to add uncertainty as to when the target will appear (i.e., a perceptually similar control). Ideally, interocular trials and monocular control trials are randomly intermixed within blocks. This has the main advantage of making the perceptual difference between dichoptic and monocular (or binocular) presentation conditions less conspicuous, due to the whimsical nature of dichoptically presented trials.

Together, these three methods provide empirical tests for (1) whether differences between stimulus conditions actually rely on high-level information, (2) whether differences in reaction times indeed reflect differences in explicit visual awareness and (3) whether reaction times were affected by processing differences emerging after conscious detection, such as changes in response criterion.

Interpreting the results of b-CFS studies

As mentioned in the introduction, the rationale underlying b-CFS experiments is that differences in suppression durations between conditions reflect different processing of stimuli prior to conscious access. An often disregarded alternative, however, is that differences between conditions may affect visual processing *during* the transitory period in which the interocularly suppressed stimuli gradually gain access to consciousness. In support of this latter idea, CFS allows for periods of partial awareness, in which some, but not all, features of a stimulus are suppressed (Yang, & Blake, 2012; Zadbood et al., 2011). Crucially, Mudrik, Gelbard-Sadiv, Faivre, and Koch (2013) demonstrated that "non-conscious" processing of faces was restricted to periods of partial awareness. This finding has two consequences for b-CFS studies: First, it indicates that detection tasks are better suited than discrimination tasks to ascertain that differences in detection times between conditions are initiated prior to a switch in ocular dominance. For instance, if the crucial manipulation involves one feature of some stimulus (e.g., color) and participants are required to report another feature of that stimulus (e.g., orientation) for the b-CFS task, it is conceivable that the color of that same stimulus was accessible to consciousness prior to its orientation. As a result, the possibility cannot be excluded that the process driving the differences in detection times between color conditions arose after the interocular competition (of that particular feature) was resolved, thereby reflecting a conscious rather than non-conscious process. Second, and more importantly, if the main goal of the experimenter is to uncover the nature of non-conscious processing, we advocate the concurrent usage of multiple suppression techniques, with an emphasis on the more traditional methods that are less susceptible to partial awareness (e.g., Mudrik et al., 2011; replicated with backward masking, Mudrik, & Koch, 2013). As suggested above, b-CFS possibly relies on processing differences during the transitory period, in which previously suppressed stimuli gain gradual access to consciousness. As such, it is hard to ascertain whether differences in stimulus processing during a transition of ocular dominance can generalize to differences in stimulus processing in the complete absence of consciousness. Despite being not as well suited as a tool to uncover non-conscious processing per se, b-CFS experiments are nonetheless very informative as a measure of access to awareness (Stein, & Sterzer, 2014). Consequently, the results of b-CFS experiments should be interpreted as such.

With the abovementioned additions to the b-CFS paradigm, we hope to provide the means to effectively dissociate situations in which competition for conscious access occurs on high-level battle grounds and thus requires conceptual or semantic processing, from situations in which the competition occurs on lower level battle grounds, such that crude visual processing of the suppressed stimuli suffices. In light of the abovementioned limitations, it should be emphasized that whether or not high-level stimulus properties exert influence on conscious access within the b-CFS paradigm does not necessarily imply that the same restrictions apply to non-conscious processing under CFS, let alone to interocular suppression in general. Thus far, however, the idea that interocularly suppressed stimuli are not analyzed up to semantic or conceptual processing levels has been mainly challenged by b-CFS experiments. The present review included 30 studies that use this experimental paradigm, of which 8 aimed to explore whether suppression durations could be affected by competition at a semantic or conceptual processing stage. Four of these studies demonstrate that these effects could be accounted for by differences in low-level visual properties, three of these studies did not include conditions that controls for differences in low-level visual properties, and as a result, only one study demonstrates high-level effects in a b-CFS task. As such, we conclude that interocular competition at a visual level is a sufficient explanation for most b-CFS studies that properly control for low-level visual differences (i.e., that include an inversion condition). As such, we should be reluctant to revise the traditional idea that semantic or conceptual analysis is abolished under interocular suppression.

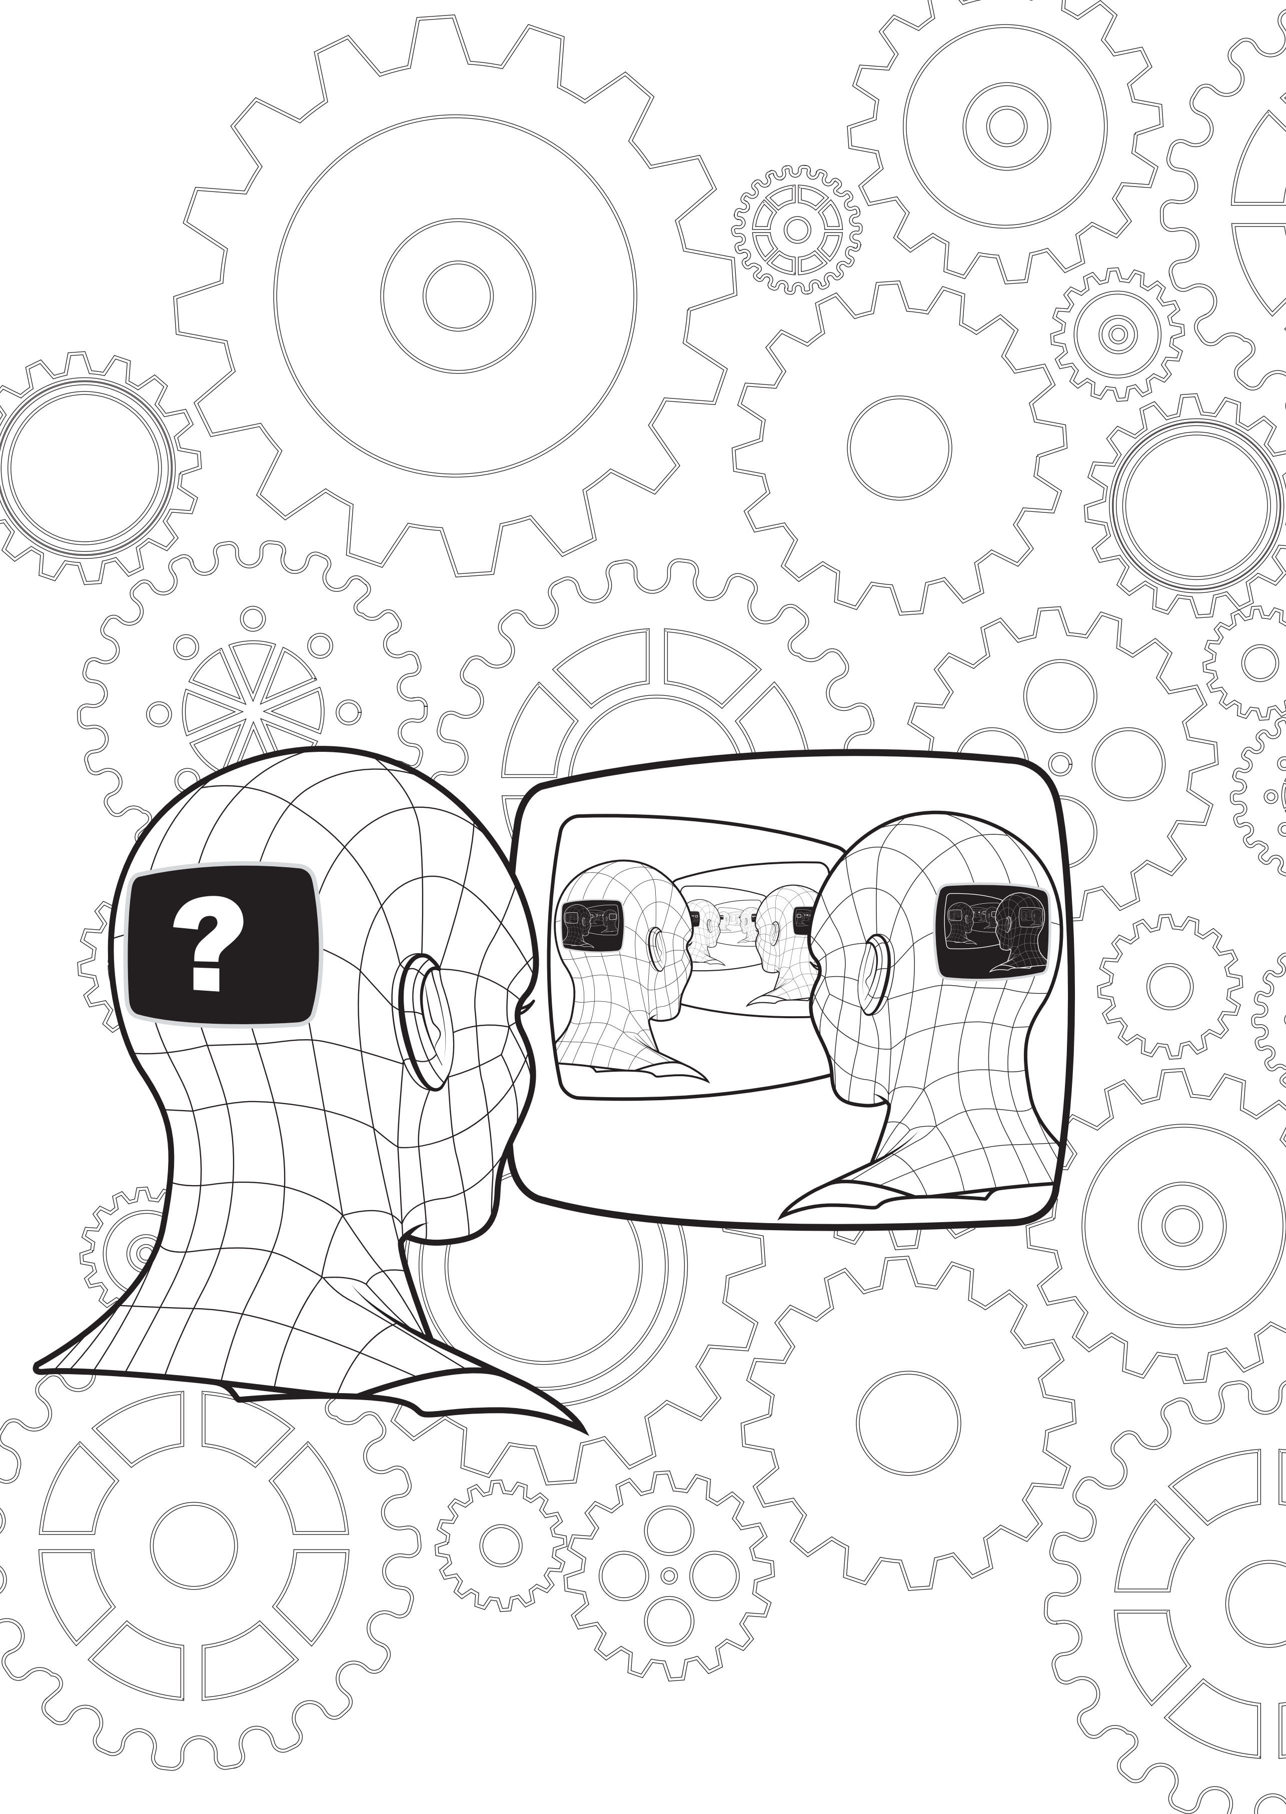

Chapter 4

Visual input signaling threat gains preferential access to awareness

A modified version of this manuscript has been published as:

Gayet, S., Paffen, C. L. E., Belopolsky, A. V., Theeuwes, J., & Van der Stigchel, S. (2016).
Fear conditioned visual input gains preferential access to awareness
in a breaking continuous flash suppression paradigm.
Cognition, 149, 77-83. doi:10.1016/j.cognition.2016.01.009

All authors designed the study concept. SG programmed the experiment and tested the participants, SG conducted the analyses and wrote the manuscript. Critical revisions were provided by all co-authors.

Abstract

Visual input that signals threat is inherently relevant for survival. Accordingly, it has been demonstrated that threatening visual input elicits faster behavioral responses than non-threatening visual input. Considering that awareness is a prerequisite for performing demanding tasks and guiding novel behavior, we hypothesized that threatening visual input would gain faster access to awareness than non-threatening visual input. In the present study, we associated one of two basic visual stimuli, that were devoid of intrinsic relevance (colored annuli), with aversive stimulation (i.e., electric shocks) following a classical fear conditioning procedure. In the subsequent test phase no more electric shocks were delivered, and a breaking continuous flash suppression task was used to measure how fast these stimuli would access awareness. The results reveal that stimuli that were previously paired with an electric shock break through suppression faster than comparable stimuli that were not paired with an electric shock.

Introduction

Detecting threatening visual input in the environment is crucial for adaptive functioning. Given that we are continuously presented with vast amounts of sensory input, any part of it that signals threat is pre-eminently relevant to the observer. Accordingly, it has been demonstrated that threatening visual stimuli, such as angry faces and spiders can be reported faster and are more difficult to ignore than non-threatening visual stimuli (e.g., Öhman, Flykt, & Esteves, 2001; for reviews, see Mather & Sutherland, 2011; Yiend, 2010). Visual awareness has been associated with such functional properties as performing demanding tasks (Dehaene, Kerszberg, & Changeux, 1998) and guiding novel behavior (Kunde, Kiesel, & Hoffmann, 2003; Gayet, Van der Stigchel, & Paffen 2014a) and is therefore valuable for selecting the appropriate set of behaviors in response to imminent threat. Accordingly, we set out to investigate whether visual stimuli that signal threat would gain faster access to awareness than stimuli that are not associated with threat.

Access to awareness was measured by means of a breaking continuous flash suppression task (b-CFS; Jiang, Costello, & He, 2007; for a review, see Gayet, Van der Stigchel, & Paffen, 2014b). In this method, a stimulus is initially interocularly suppressed by continuous flash suppression (CFS; Tsuchiya & Koch, 2005). The time it takes for this stimulus to overcome interocular suppression, so that it can be reported by an observer, provides a measure of access to awareness (Gayet et al., 2014b; Stein, Hebart, & Sterzer, 2011). The experimental manipulation of threat was obtained by associating one of two colored annuli with electric shocks (hereafter the CS+), while never pairing the other annulus color with a shock (hereafter the CS-), following a classical fear conditioning procedure (Mackintosh, 1983; Pavlov, 1927). Unlike aversive images, the usage of electric shocks has the advantage of constituting an actual threat to participants. In this context, we defined threat as a state of the world predicting an aversive event, as evidenced by prior experience. Fear conditioning allows for isolating the manipulation of threat from the visual characteristics that typically differentiate threatening from non-threatening stimuli. This is especially relevant in the present context, as differences in visual stimulus characteristics are known to affect access to awareness in a b-CFS task (Gayet et al., 2014b; Stein & Sterzer, 2012; Yang and Blake, 2012).

Recent experiments have revealed that visual input that was previously associated with an aversive event is more readily detected (Padmala & Pessoa, 2008; Phelps, Ling, & Carrasco, 2006) and attracts attention (Armony, & Dolan, 2002; Schmidt, Belopolsky & Theeuwes, 2015a) and eye movements (Mulckhuyse, Crombez, & Van der Stigchel, 2013; Schmidt, Belopolsky & Theeuwes, 2015b) to a greater extent than comparable visual input that was never associated with an aversive event. These findings lead us to hypothesize that fear conditioning increases the salience of visual input. Considering that increased stimulus contrast yields reduced suppression strength under CFS (Tsuchiya & Koch, 2005), we expected reduced suppression durations for fear-conditioned stimuli. Hence, we predicted faster response times to target gratings surrounded by a CS+ compared to a CS- annulus.

Using behavioral reports of visibility in suppression paradigms (e.g., backwards masking, bistable perception, inattentional blindness, etc.) as a measure of access to awareness brings about an important pitfall: Fear conditioning might not only affect how strongly a stimulus is suppressed, it might also affect how fast a stimulus can be responded to once it is no longer suppressed (i.e., reflecting so-called post-suppression effects). In the case of fear conditioning specifically, post-suppression effects could arise from a more liberal response bias in CS+ compared to CS- trials (requiring less perceptual evidence for deciding that a target is perceived) or increased sensitivity to CS+ compared to CS- stimuli (requiring less perceptual evidence for perceiving the target). Using interocular competition allows for dissociating between suppression (between-eyes) and post-suppression effects (within-eye). This is achieved by including a monocular control condition, in which the stimulus is presented to the eye that is already dominant. Considering that there is no interocular competition (and therefore no between-eye effect) in this control condition, it allows to selectively measure differences in response speed to the stimuli once they are already visible (i.e., post-suppression, or within-eye effects). Consequently, any difference in response times that emerges in the suppression condition but not in the monocular control condition reflects a difference in suppression durations (for discussion on this interpretation, see Gayet et al., 2014b; Stein et al., 2011). In addition to this crucial control, we designed our stimuli such as to minimize post-suppression effects. For this purpose, we created a response task (reporting the orientation of a grating) that was orthogonal to the experimental manipulation (which was tied to the color of the surrounding annulus; for a similar approach, see Salomon, Lim, Herebelin, Hesselmann, & Blanke, 2013). As a result of this, post-suppression effects that affect processing of the CS+ annulus are not expected to affect the orientation judgment task, as the grating has not been fear conditioned. Considering that we took precautions to minimize post-suppression effects, we did not expect an effect of fear conditioning in the monocular condition.

Experiment 1

Methods

Participants

Eighteen healthy participants were selected for (corrected to) normal vision, including stereoscopic vision (tested by inducing diplopia) and color perception (Ishihara, 1917). The eventual group of participants that was included in the analyses consisted of 7 males and 9 females, with an average age of 23 years (SD = 3.5). The sample size was derived from comparable b-CFS studies (e.g., Stein & Sterzer, 2012) and fear conditioning studies (Schmidt et al., 2015a, 2015b). The experimental procedure was validated by the ethical board of the VU University of Amsterdam.

Stimuli

All stimuli were presented on a black (0.2 cd/m^2) screen. Using a dichoptic mirror stereoscope, each eye was presented with a square shaped Brownian noise frame, subtending an area of 5.0 by 5.0° on the outside, that encapsulated a gray presentation area of 2.2 by 2.2° with a fixation cross (0.2 by 0.2°) in the center. The 'Mondrian' masks used to obtain CFS consisted of overlapping black, gray (41.2 cd/m^2) and white (87.0 cd/m^2) circles with diameters ranging from 0.35 to 0.70°. In each block 46 unique masks were generated. On each trial the order of the masks was shuffled, and the masks were replaced at 10 Hz. The target stimuli were comprised of saturated red (5.3 cd/m^2, *SD* = 0.8, x = 0.034, y = 0.344), green (8.6 cd/m^2, *SD* = 1.6, x = 0.152, y = 0.104) and blue (reference color at 6.4 cd/m^2, x = 0.281, y = 0.344) annuli with an outer radius of 1.2° and an inner radius of 0.7°. Flicker photometry (Kaiser & Comerford, 1975) was used to perceptually equate the luminance of the different annulus colors with the gray background (6.7 cd/m^2, *SD* = 0.8). The annuli encapsulated a sine-wave grating with a spatial frequency of 8.4 cycles/°, a Gaussian profile (*SD* = 0.35°) and a mean luminance equal to that of the gray background. The grating could have an orientation of either plus or minus 45° from the vertical midline.

The aversive stimulation used for the fear conditioning procedure consisted of 400V electric shocks with a mean amperage of 11.0 mA (*SD* = 7.5). A train of shocks with a duration of 50 microseconds each was delivered at 60 Hz, giving the sensation of a single 500 ms shock. Shocks were delivered through two electrocardiogram electrodes, connected to a Digitimer DS7A direct current stimulator, that were placed over the tibial nerve, at the medial malleolus of the right ankle.

Experimental design

The experimental design comprised the within-subject factors Conditioning (CS+ or CS- annulus), Suppression (suppression or monocular condition) as factors of interest, and the factors Eye (target presented to the left or right eye), and Orientation (grating tilted leftwards or rightwards) as factors of no interest. This resulted in 16 unique combinations of within-subject conditions, which were presented twice within the first 32 trials of the experiment, and twice within the last 32 trials of the experiment. Trial order was randomized within these two experimental halves. Eventually, the experiment was divided into four experimental blocks. The factor Color (CS+ annulus is blue or red) was a between-subject factor and was counterbalanced between participants. Taken together, each of the four Suppression x Conditioning conditions contained 16 trials. We limited the test phase of the experiment to this relatively small number of trials, as we expected that the effect of fear conditioning would extinguish after repeatedly presenting a CS+ stimulus unaccompanied by a shock.

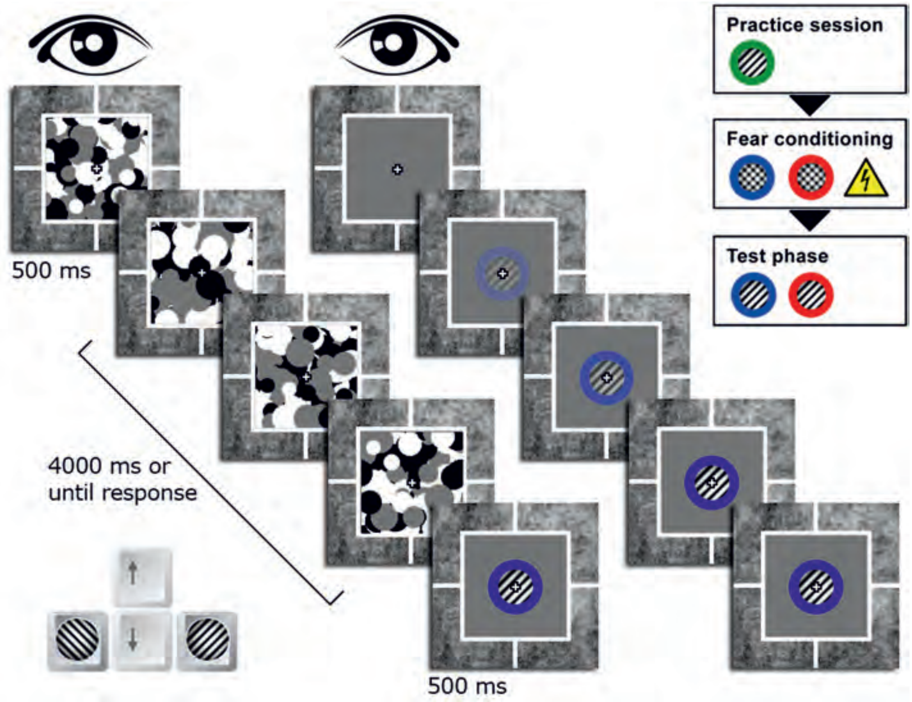

Figure 1. Schematic depiction of an experimental trial in the suppression condition. Participants were instructed to report the orientation of the grating. The same stimulus chronology was used throughout every phase of the experiment. The panel in the upper right indicates the stimuli that were used in these different phases.

Procedure

The experiment started with 32 practice trials, in which participants were instructed to report the orientation of the grating as fast and accurately as possible (see Figure 1). Each trial started with a fixation cross (500 ms), and between 300 and 600 ms after onset of the Mondrian masks the annulus and grating were presented either to the same eye as the mask (monocular condition), or to the other eye (suppression condition). These stimuli were ramped up to full intensity in either one second (suppression condition) or three seconds (monocular condition). These durations were chosen such as to elicit comparable RT distributions in the two Suppression conditions (for other studies using this approach, see Gayet et al., 2013; Stein, Hebart, & Sterzer, 2011). After a response was given, or 4000 ms had elapsed, the masks were removed from the screen, and the target was presented binocularly for 500 ms.

Next, the electrodes were attached to the ankle of the participants, and electric shocks of increasing intensity were administered following a shock workup procedure (adapted from Heitland, Groenink, Bijlsma, Oosting, & Baas, 2013). This procedure was aborted once participants rated the aversiveness of the shock as a 4 out of 5. During the subsequent 16 trials acquisition block, six out of eight trials with one annulus color (hereafter CS+) were paired with an electric shock, whereas none of the trials with the other annulus color (hereafter CS-) were paired with an electric shock. The two grating orientations (left and right tilt) were superimposed during acquisition, to avoid that participants would incidentally couple the occurrence of an electric shock to a particular orientation, rather than to a particular color. After the acquisition phase, the participants were presented with 64 experimental trials. Participants were told that no more shocks would be delivered and they were instructed to report the orientation of the gratings as fast and accurately as possible.

Data analysis

For the analysis of response times, only trials with correct responses (93.3%, *SD* = 6.1) were included. One participant was at chance level in reporting the orientation of the target grating and was therefore excluded from all analyses. Median response times were computed for each participant's Conditioning and Suppression conditions. Truncating two response time distributions at the same numerical value (i.e., at the response time deadline of 4000 ms) would cause the analysis to include more data (from the slower end of the distribution) in the condition that yielded the fastest response times, thereby compromising a fair statistical comparison between the two conditions. To circumvent this problem, we included trials in which no response was given within the 4000 ms time limit (8.3% of the trials, *SD* = 9.3) in the analysis as well, as they reflect long suppression durations. On these trials, response times were registered as "infinitely long". This approach allowed for computing median response times that include all trials, and therefore reflect the entire response time distribution, rather than the average of the response times that were within the arbitrary deadline of 4000 ms. In four participants, at least 50% of the trials in either the CS+ or the CS- condition yielded an infinitely long response (i.e., no response). As a consequence, the median response time for these participants was infinitely long as well, in at least one condition, and was therefore deemed uninformative. For these participants, we computed the median response times by including only the trials in which the target was presented to the dominant eye (i.e., the Eye condition that yielded the shortest response times). This allowed for preserving the equal prevalence of experimental conditions within these participants. One participant was excluded from further analyses as trials in neither Eye condition retained more than 50% of responses within the 4000 ms time limit. As such, the analyses in the Results section are based on 16 participants, for eight of which the shocks were associated with the red annulus color. For these participants, 7.8% (SD = 2.4) of the trials yielded either no response or an incorrect response. Additional analyses with alternative inclusion criteria are provided in supplementary materials S1 and S2.

Results

We conducted a 2x2 repeated-measures ANOVA with the within-subject factors Conditioning and Suppression, and the between-subject factor Color. This revealed a main effect of Conditioning, $F(1, 14) = 8.97$, $p = .010$, $\eta^2 = .39$, showing that, irrespective of the Suppression condition, trials in which a CS+ annulus was presented (M = 1557 ms, SD = 284) yielded faster response times than trials in which a CS- annulus was presented (M = 1718, SD = 432). The absence of a main effect of Suppression, $F(1, 14) = .20$, $p = .659$, $\eta^2 = .01$, indicated that we successfully matched the response times of trials in which the targets were interocularly suppressed (M = 1660 ms, SD = 524), and trials in which they were not (M = 1615 ms, SD = 248). This is an important requirement for making a fair comparison between effects in the suppression condition and the monocular control condition (Stein et al., 2011). The between-subject factor Color did not interact with either Suppression (p = 0.175) or Conditioning (p = 0.623). Finally, there was an interaction between Suppression and Conditioning on response times, $F(1, 14) = 13.98$, $p = .002$, $\eta^2 = .50$, but no three-way interaction with the between-subject factor Color, $F(1, 14) = .08$, $p = .780$, $\eta^2 = .01$. This shows that the interaction between Conditioning and Suppression did not depend on the specific color that was paired with electric shocks.

Subsequent paired-samples t-tests revealed that when targets were interocularly suppressed, discrimination of the oriented gratings was faster when they were surrounded by a CS+ annulus (*M* = 1498 ms, *SD* = 416) than a CS- annulus (*M* = 1822 ms, *SD* = 646), $t(15) = 3.72$, $p = .002$, Cohen's $d = 1.92$. When the grating and annulus were not interocularly suppressed (i.e., in the monocular condition), however, response times did not differ between CS+ trials (*M* = 1615 ms, *SD* = 252) and CS- trials (*M* = 1614 ms, *SD* = 271), $t(15) = .04$, $p = .968$, Cohen's $d = .02$. There was no correlation between the influence of conditioning on response times in the monocular condition and the suppression condition, $R(14) = .29$, $p = .276$. To assert that the absence of an effect of conditioning in the monocular condition reflected a null effect rather than experimental insensitivity we computed a Bayes factor (Dienes, 2012). The alternative hypothesis was modelled as a uniform distribution with a lower bound of zero and an upper bound equaling the effect of conditioning in the suppression condition. This revealed that the data was over seven times more likely to reflect a null effect than the alternative hypothesis ($B_{01} = 7.14$). Together, these findings demonstrate that stimuli that were previously paired with aversive stimulation, and therefore signal threat, are released from interocular suppression (322 ms) faster than stimuli that were not paired with a shock. The effect of conditioning on response times as a function of Suppression condition is depicted in Figure 2A (group results) and 2B (individual results).

A separate 2x2 repeated-measures ANOVA, revealed no interaction effect between the factors Suppression and Conditioning on accuracy, $F(1, 15) = .23$, $p = .638$. Consequently, the larger response time difference in the suppression condition compared to the monocular condition is not potentiated by a larger uncertainty in the suppression condition. Finally, supplementary analyses (S1 and S2) underline the robustness of the effect of fear conditioning on suppression durations.

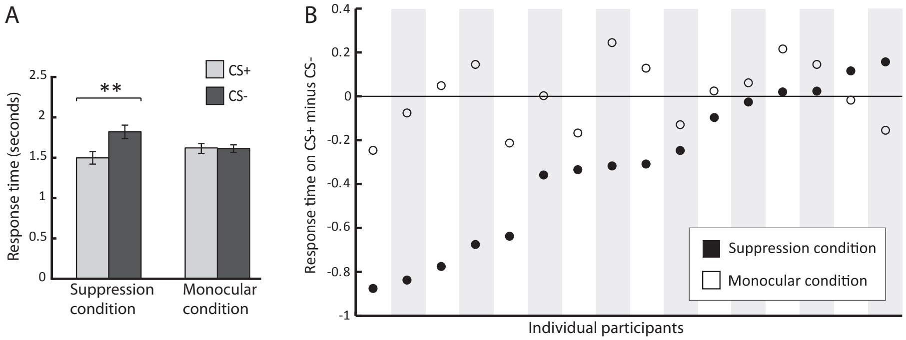

Figure 2. Group results (A) and individual results (B) for all 16 participants. In panel A, the y-axis represents the response time as a function of presentation condition (labeled on the x-axis) and fear conditioning (shades of gray). In panel B, the y-axis represents the difference in response time between CS+ trials and CS- trials for each participant (depicted along the x-axis). $^{**}p < 0.005$

Discussion

The data demonstrate that fear-conditioned stimuli are released from interocular suppression faster than equivalent stimuli that were never paired with an aversive stimulus. This speed-up was not caused by an increased response speed after the annuli were released from suppression. The correlational analysis provides further evidence that response times in the suppression condition reflect suppression durations. The absence of a correlation between the effect of fear conditioning in the suppression condition and the monocular control condition shows that participants that were faster on CS+ trials in the suppression condition were not necessarily faster on CS+ trials in the monocular control condition. This confirms that two different processes are operating in these two conditions. Additionally, the Bayesian analysis confirms that we obtained a reliable null-effect in the monocular control condition. That is, fear conditioning did not affect response times after the stimuli were released from suppression. Taken together, our experimental paradigm was successful in measuring differences in suppression durations elicited by fear conditioning, which were not contaminated by effects of fear conditioning on response times after the stimuli were released from suppression.

Experiment 2

We hypothesized that the advantage of threatening visual input to access awareness could be accounted for by an increase in salience of these stimuli. In Experiment 2, we set out to investigate the possibility that an increase in salience could indeed shorten suppression durations while leaving response times unaffected by post-suppression effects. For this purpose we manipulated the bottom-up salience of the annuli (i.e., the luminance contrast). We expected that, under conditions of interocular competition, an increase in the luminance contrast of an annulus would shorten suppression durations and thereby lead to faster response times to the target grating. Again, in the monocular condition no such difference was expected, as the eye to which annulus and grating were presented was already dominant.

Methods

In Experiment 2, the stimuli and presentation were identical to that of Experiment 1, except for one difference. Whereas the blue and green target colors and the gray background color were matched for perceptual equiluminance, as in Experiment 1, the red target color was set to its maximal luminance (27.7 Cd/m^2, x = 0.639, y = 0.344). As a result of this, the red annulus had an average Michelson luminance contrast of 61% with the gray background.

To ensure that there were no a priori differences in suppression durations of the blue and red annuli, we ran an additional control experiment (Experiment 3), which comprised 15 of the 16 participants that were included in Experiment 2. In this control experiment, the blue and red annuli were perceptually equiluminant. Here, no difference in response times was observed between trials with blue (M = 1689 ms, *SD* = 754) and red annuli (M = 1661 ms, *SD* = 702), $t(14) = 0.26$, $p = 0.798$, in the suppression condition. This shows that, at least for these participants, perceptual equiluminance matching ensures equal suppression durations in a b-CFS paradigm.

Eighteen healthy students, to which all same inclusion criteria applied as to that of the main experiment, participated in this supplemental experiment. One participant was excluded from further analyses because none of his responses were registered within the 4000 ms time limit. Another participant was excluded because he reported having repeatedly experienced diplopia throughout the experiment. The eventual group of participants consisted of 6 males and 10 females, with an average age of 22 years (SD = 3.3).

Results

The results of this experiment are presented in Figure 3. A 2x2 repeated-measures ANOVA with the within-subject factors Contrast (high or low) and Suppression (interocular suppression or monocular presentation) revealed a main effect of Contrast, $F(1, 15) = 4.65, p = .048, \eta^2 = .24$, no main effect of Suppression, $F(1, 15) = 1.39, p = .257, \eta^2 = .09$, and an interaction between Suppression and Contrast, $F(1, 15) = 7.32, p = .016, \eta^2 = .33$. Subsequent paired-samples t-tests (see Figure 2B) revealed that when targets were interocularly suppressed, discrimination of the orientated gratings was faster when they were surrounded by high contrast annuli (*M* = 1477 ms, *SD* = 605) than low contrast annuli (*M* = 1650 ms, *SD* = 712), $t(15) = 2.47, p = .026$, Cohen's *d* = 0.67. When the gratings and annuli were not interocularly suppressed, however, response times did not differ between trials in which gratings were surrounded with high contrast annuli (*M* = 1398 ms, *SD* = 266) and low contrast annuli (*M* = 1399 ms, *SD* = 258), $t(15) = .03, p = 0.979$, Cohen's *d* = .01). These findings demonstrate that increasing the luminance contrast of an annulus yields the same pattern of results as associating that annulus with a threat (main experiment).

Participants' average percentage correct over all conditions was 91.7% (*SD* = 8.6). A separate repeated measures analysis, revealed no interaction effect between the factors Suppression and Conditioning on accuracy, $F(1, 15) = 0.15, p = .708$. Again, the effect of contrast manipulation on response times mentioned above, is not accounted for by a greater uncertainty in the suppression condition compared to the monocular condition.

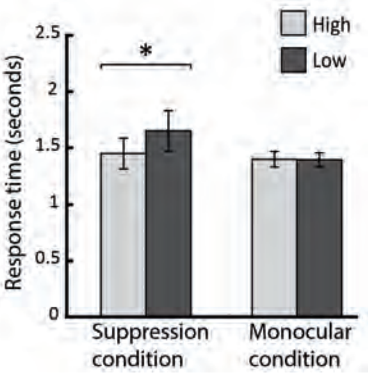

Figure 3. Group results for all 16 participants. The y-axis represents the response time as a function of presentation condition (labeled on the x-axis) and annulus contrast (shades of gray). *$p < 0.05$

Discussion

Here, we demonstrated that an increase in bottom-up salience (i.e., luminance contrast of the annuli) yielded a similar pattern of results as was obtained with fear conditioning. That is, we found shorter suppression durations for higher contrast stimuli compared to lower contrast stimuli, while the response speed to higher contrast stimuli remained unaffected. From this we conclude that (1) the combined use of an orthogonal response mapping and a monocular control condition allows for selectively measuring differences in suppression durations elicited by modulating stimulus salience and (2) increases in bottom-up stimulus salience elicit patterns of results that are similar to those engendered by fear conditioning.

General discussion

The present study reveals that a priori neutral visual stimuli gain faster access to awareness when they have been associated with electric shocks, following a classical fear conditioning procedure. We cannot ascertain whether the fear conditioning procedure successfully elicited a fear response to CS+ stimuli, as we did not include physiological response measurements. But, considering that CS+ trials and CS- trials only differed by virtue of the preceding fear conditioning procedure, it is implied that the fear conditioning procedure potentiated the faster access to awareness of CS+ compared to CS- stimuli. Based on the history of occurrence, the CS+ therefore constituted a threat, as it signaled an aversive event. Importantly, the visual stimuli were manipulated experimentally to be either threatening to the observer or not. By this, we successfully isolated the effect of threat from the visual characteristics that typically constitute the threatening stimuli. As the association of a stimulus with threat co-determined the time at which this stimulus entered awareness, it can be inferred that the perceptual system was able to differentiate threatening stimuli from non-threatening stimuli before they entered awareness. It is argued that the b-CFS method is suitable to measure access to awareness, but cannot unequivocally measure non-conscious processing (Stein & Sterzer, 2014; Gayet et al., 2014b; Stein et al., 2011). In line with our interpretation, however, earlier studies have demonstrated that threatening versus non-threatening stimulus categories (Schmack, Burk, Haynes, & Sterzer, 2015; Lipp, Kempnich, Jee, & Arnold, 2014; Jiang & He, 2006; Williams, Morris, McGlone, Abbott, & Mattingley, 2004; Morris, Öhman, & Dolan, 1998; Whalen et al., 1998; for a review, see Pessoa, 2005) as well as CS+ versus CS- stimuli (Raio, Carmel, Carrasco, & Phelps, 2012) can indeed be segregated non-consciously. The present study shows that the visual system can use this non-conscious segregation between threatening and non-threatening visual input to determine the contents of awareness.

In line with our findings, Alpers and colleagues observed increased predominance of fear conditioned stimuli in a binocular rivalry task (Alpers, Ruhleder, Walz, Mühlberger, & Pauli, 2006). While their findings are in line with the present study, they do not address the question of whether threatening information gains faster access to awareness. In binocular rivalry, it is impossible to discern whether the experimental manipulation alters dominance durations by impacting the visual processing of the perceived stimulus or that of the suppressed stimulus. As CS+

stimuli attract attention (Armony, & Dolan, 2002; Schmidt et al., 2014), it might take more time to disengage from a CS+ stimulus than from a CS- stimulus, and it might take less time to initiate report of a CS+ stimulus than a CS- stimulus, both leading to longer reported percept durations. In the present study this was not an issue, as observers initially only perceived the mask, and response speeds were therefore necessarily dependent on the suppressed stimulus rather than the dominant stimulus. Additionally, in our experiment, we were able to assert that the differences in response times were not accounted for by differences in response speed after the interocular conflict was resolved (i.e., post-suppression), by the inclusion of a monocular control condition.

In order to minimize post-suppression effects of fear conditioning on response times, we segregated the stimulus part that dictated the suppression duration (the CS+ or CS- annulus) from the stimulus part to which observers respond (the neutral grating). Chromatic annuli are more perceptually salient than low contrast sine gratings. As a result of this difference in saliency, the time point at which a switch in ocular dominance occurred was primarily driven by the time point at which the chromatic annuli broke through suppression. This is crucial, since the experimental manipulation was tied to the color of the annulus. Switches in ocular dominance likely initiated at the location of the salient stimulus (Stuit, Verstraten, & Paffen, 2010; Paffen, Naber, & Verstraten, 2008) and then spread over the initially suppressed stimulus (Kaufman, 1963) throughout the rest of the ipsi-ocular percept (Ooi, & He, 1999; Zhang, Jiang, & He, 2012). As such, faster breakthrough of the chromatic annulus resulted in faster breakthrough of the sine grating, eventually allowing observers to report its orientation. Considering that, in the present case, there was no difference in post-suppression effects (Experiment 1, monocular control condition), the difference in response times between CS+ and CS- trials reflected a difference in suppression durations.

Recent studies using a binocular rivalry task showed that monetary reward engenders an increase in dominance duration for the rewarded percept (Marx, & Einhäuser, 2015; Wilbertz, Van Slooten, & Sterzer, 2014) whereas monetary punishment engenders a decrease in dominance durations for the punished percept (Wilbertz et al., 2014). This latter finding seems at odds with the present data. In their paradigm, the act of perceiving the punished percept itself resulted in a monetary loss for the participant. In the present study, however, the CS+ stimulus signaled a potential threat. Its detection is therefore beneficial from a behavioral point of view, as detecting threat allows for selecting appropriate behavior. Taken together, our findings and those of Wilbertz and colleagues show that negative valence does not necessarily lead to faster switches in ocular dominance. Rather, visual input associated with negative valence selectively gains preferential access to awareness when it serves an adaptive function.

One possible account for the present findings is that the association of visual input with threat increases its salience, which leads to faster release from suppression. Considering that interocular competition is biased towards stimuli of higher contrast (Brascamp, van Ee, Noest, Jacobs, & van den Berg, 2006; Levelt, 1965), a stimulus that is of higher salience, although equal in contrast, is likely to be favored in interocular competition as well (e.g., Gayet, Brascamp, Van der Stigchel, & Paffen, 2015). This idea finds some substantiation in the strikingly similar results elicited by fear conditioning (Experiment 1) and by manipulation of the luminance contrast (Experiment 2) in the present study. An increase in salience could also explain the related findings that CS+ stimuli attract attention (Armony & Dolan, 2002; Schmidt et al., 2015a) and eye movements (Mulckhuyse et al., 2013; Schmidt et al., 2015b), as well as the finding that CS+ stimuli can be more readily detected than CS- stimuli (Padmala and Pessoa, 2008). Accordingly, imaging studies have reported stronger BOLD activity in the visual cortex for CS+ compared to CS- stimuli (Lim, Padmala, & Pessoa, 2009; Padmala & Pessoa, 2008).

The amygdala might play a role in modulating the cortical response to hitherto invisible threatening stimuli. A number of studies successfully detected amygdala activation of interocularly suppressed emotional stimuli (Pasley, Mayes, & Schultz, 2004; Williams et al., 2004; Jiang & He, 2006). Back-projections from the amygdala to the visual cortex, which have been directly observed in primates (Amaral, Behniea, & Kelly, 2003; Amaral & Price, 1984), could then allow for mediating the cortical response to threat stimuli (Pessoa & Adolphs, 2010). In line with this idea, amygdala activation has been shown to mediate the differential response to CS+ compared to CS- stimuli in the visual cortex (Lim et al., 2009). Framing our findings within this literature leads us to tentatively suggest that visual input signaling threat is initially detected non-consciously, as a result of which the cortical activation elicited by this visual input is enhanced. This enhanced cortical activation could potentiate the preferential access to awareness of that part of the visual world that signals threat.

Conclusion

Our results demonstrate for the first time that visual input that signals threat is not only privileged by the visual system such as to elicit stronger behavioral and neural effects, but that the very content of our consciousness is more likely to be comprised of visual information that signals threat. Considering that our experimental manipulation of threat affected the time taken to breach the threshold of awareness, we propose that threat signals were extracted from visual input *before* they reached awareness.

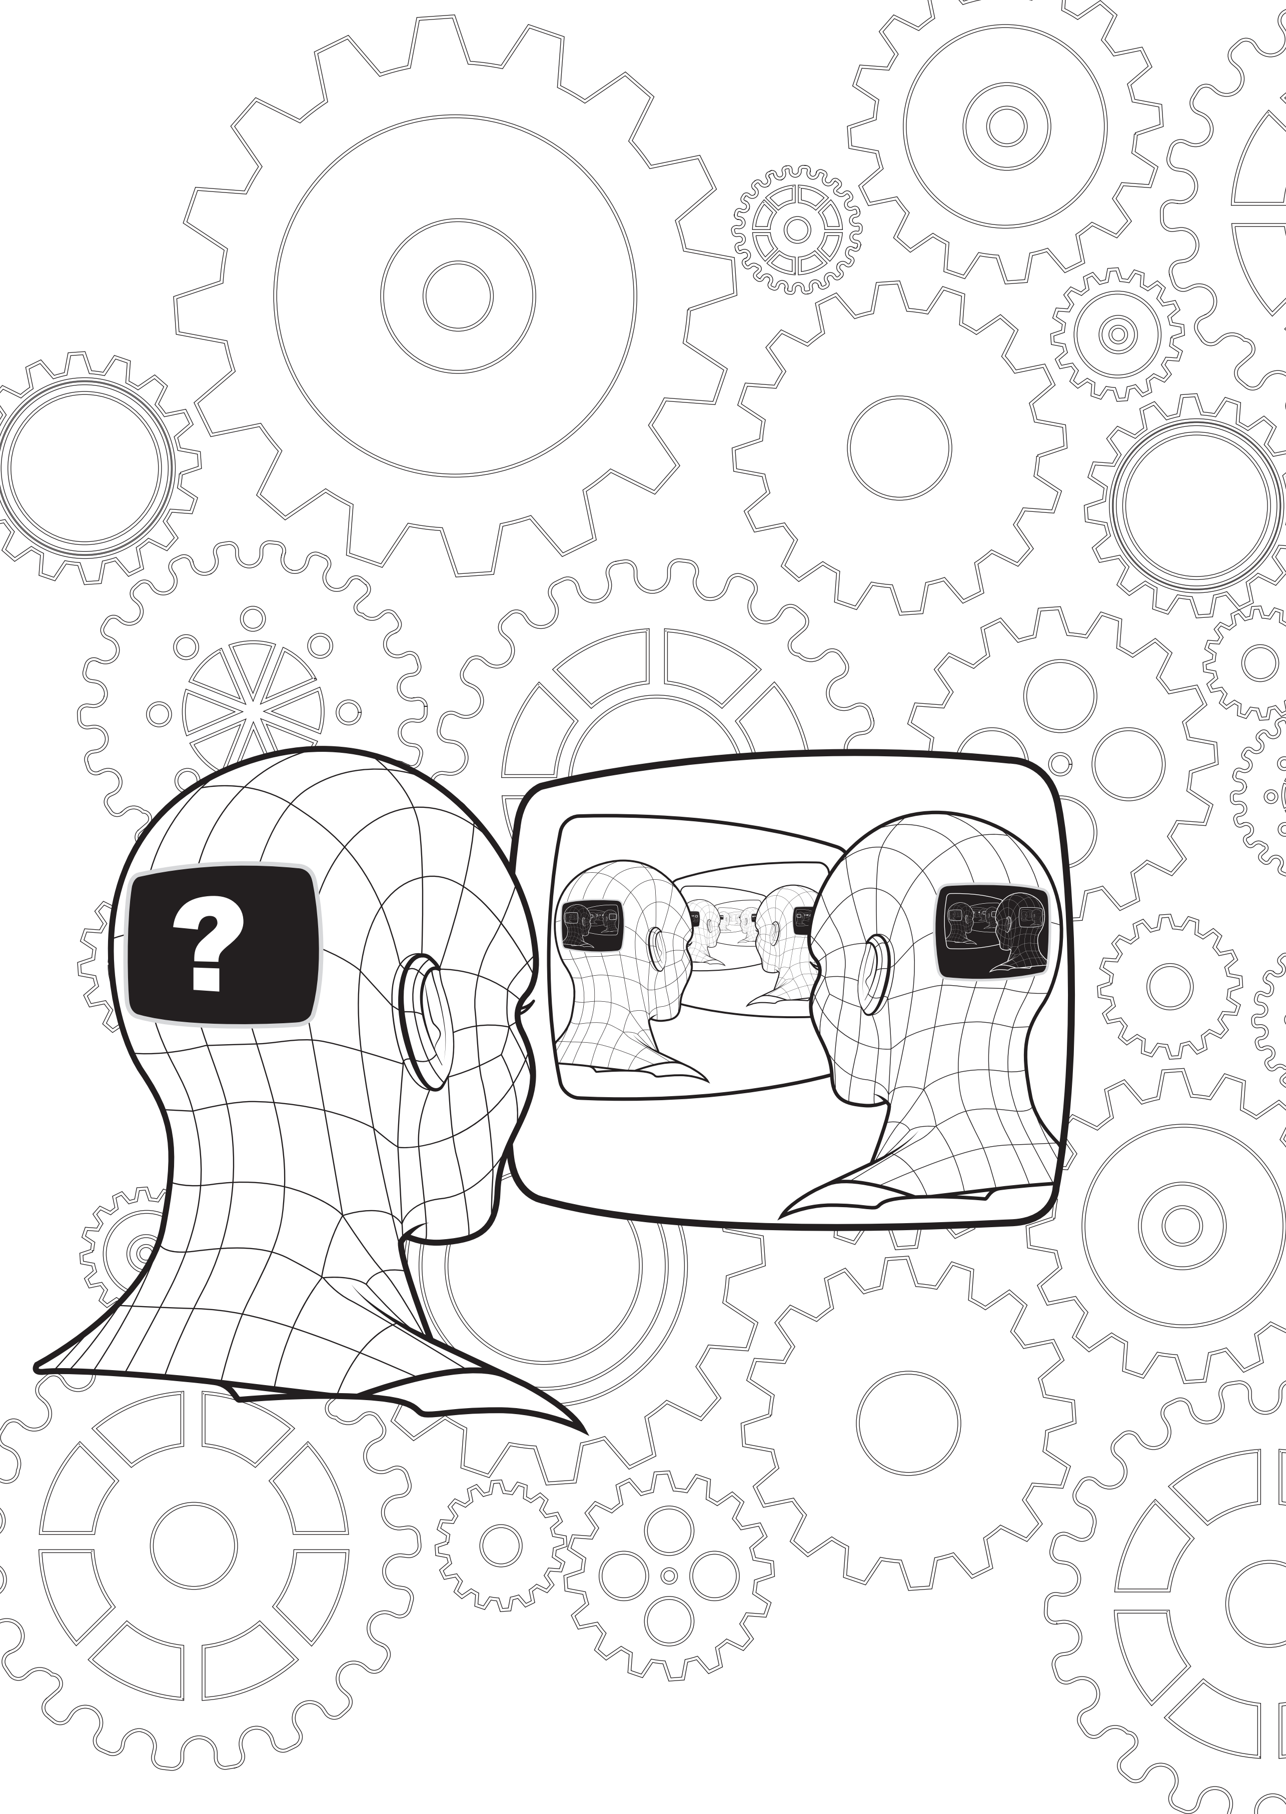

Chapter 5

The content of visual working memory modulates access to awareness

A modified version of this manuscript has been published as:

Gayet. S., Paffen, C. L. E., & Van der Stigchel, S. (2013). Information matching the content of visual working memory is prioritized for conscious access. *Psychological Science, 24*(12), 2472–2480. doi:10.1177/0956797613495882

All authors designed the study concept. SG programmed the experiment and tested the participants, SG conducted the analyses and wrote the manuscript. Critical revisions were provided by all co-authors.

Abstract

Visual working memory (VWM) is used to retain relevant information for imminent goal-directed behavior. In the experiments reported here, we found that VWM helps to prioritize relevant information that is not yet available for conscious experience. In five experiments, we demonstrated that information matching VWM content reaches visual awareness faster than does information not matching VWM content. Our findings suggest a functional link between VWM and visual awareness: The content of VWM is recruited to funnel down the vast amount of sensory input to that which is relevant for subsequent behavior and therefore requires conscious access.

Introduction

Most of the visual information that people use to interact with the physical environment does not give rise to conscious experience. However, when people are confronted with a situation that requires mental effort, such as playing a game of chess or memorizing a telephone number, they seem to be fully aware of the sensory input that is relevant for subsequent behavior (Dehaene, Kerszberg, & Changeux, 1998). This suggests that consciousness is required for handling certain situations. In accordance, a growing number of studies have associated visual awareness with different functional properties, such as decision making (e.g., Van Gaal, Lange, & Cohen, 2012) and planning (e.g., Crick, & Koch, 2003). Visual awareness, however, is of limited capacity and is capable of representing only a minute part of the available information at a particular time (e.g., Baars, 1997a, 1997b; Dennett, 1991). As a consequence of this limitation, a selection stage is needed to filter out the vast amount of incoming visual input.

The human brain uses visual working memory (VWM) to actively retain relevant information for imminent goal-directed behavior (for a review, see Baddeley, 2003). Therefore, the content of VWM would be a good candidate to help in selecting relevant information for conscious access. Two types of findings provide an indication that VWM might well play such a role. First, a stimulus presented below the threshold of awareness exerts a greater influence on behavior when it matches the content of VWM than when it does not (Pan, Cheng, & Luo, 2012). Second, VWM plays a role in selecting relevant stimuli for the attentional system by providing a bias toward stimuli matching VWM content (Hollingworth, & Luck, 2009; Hollingworth, Matsukura, & Luck, 2013; Olivers, 2009; Olivers, Meijer, & Theeuwes, 2006). Our aim in the present experiments was to examine whether the content of VWM could act on information that is not yet consciously perceived, so as to prioritize conscious access of matching, and therefore potentially relevant, information.

A recently developed method called *breaking continuous flash suppression* (b-CFS) provides the means to compare the potency of different visual stimuli to reach conscious access (e.g., Jiang, Costello, & He, 2007; Mudrik, Breska, Lamy, & Deouell, 2011; Sklar et al., 2012; Stein, Hebart, & Sterzer, 2011; Wang, Weng, & He, 2012). In this paradigm, a stimulus presented to one eye is rendered temporarily invisible by presenting a dynamic pattern to the other eye. Consequently, the time it takes for this interocularly suppressed stimulus to break into visual awareness provides a measure of prioritization for conscious access of that stimulus.

In the color-memory conditions of the present experiments, we instructed participants to retain a color in VWM and measured how this VWM content affected suppression durations of colored stimuli. Specifically, we tested whether suppressed targets that match (i.e., are congruent with) a color category held in VWM break through interocular suppression faster than targets of nonmatching (i.e., incongruent) colors.

Experiment 1

At the start of each trial in Experiment 1, a colored patch was shown, which participants were instructed to either actively memorize for a subsequent recognition task (color-memory condition) or only fixate without memorizing (passive-viewing condition). Next, a target was rendered invisible by continuous flash suppression (CFS): A target, either matching the color category of the memorized stimulus or not, gradually increased in intensity for one eye while a high-contrast dynamic pattern was presented continuously to the other eye. As soon as they detected the target, participants reported its location. In the final phase, which was omitted in the passive-viewing condition, participants reported which of three colored patches had the same specific color as the memorized stimulus.

Methods

Participants

Ten participants (age 18 to 27 years) completed Experiment 1. All participants were right-handed and had normal or corrected-to-normal vision. Before taking part in the experiment, they were tested for stereoscopic vision with the TNO test for stereoscopic vision (12th edition; Walraven, 1972). Additionally, each participant's eye dominance was determined off-line by comparing each eye's median detection time for interocularly suppressed stimuli. The presentation that led to the longest suppression durations (i.e., presenting the mask to the left eye and the target to the right eye or vice versa) was used during b-CFS.

Apparatus and stimuli

The experiment was conducted using an Apple dual 2-GHz PowerPC G5 equipped with a linearized 22' LaCie Electron Blue IV CRT monitor (1,024 × 768 pixels; 100 Hz refresh rate) and an Apple keyboard, which was used for response registration. There were no light sources in the experiment room, except for the computer monitor. Stimulus presentation and response collection were managed using the Psychophysics Toolbox 3 (Brainard, 1997; Pelli, 1997) in MATLAB (Release R2010a; The Mathworks, Natick, MA). A pair of displays was viewed dichoptically through a mirror stereoscope mounted on a chin rest, which kept the effective viewing distance at 57 cm. All stimuli were presented on a uniform gray background with a luminance of 24.5 cd/m^2. To facilitate binocular fusion of the two complementary images, we surrounded the area presented to each eye with a black frame (< 1 cd/

m^2; > 96% Weber contrast) that subtended 5.8° × 5.8° of visual angle. Instructions were presented in white letters (46.2 cd/m^2; 86% Weber contrast) in Arial font with a size of 14. All color stimuli were circles with a diameter of 1.2° of visual angle.

All 15 colors (5 red, 5 green, and 5 blue hues) for the memory task were physically matched to the luminance of the blue target stimulus using a PR-650 SpectraScan colorimeter-telephotometer (Photo Research, Chatsworth, CA). The obtained color patches had a mean luminance of 6.06 cd/m^2 (*SD* = 0.17 cd/m^2) and a 76% Weber contrast with the background.

The colored patches for the suppression task consisted of saturated red, green, or blue. The saturated-blue target stimulus (5.98 cd/m^2, $x = 0.149$, $y = 0.073$; 76% Weber contrast) was used to obtain perceptual equiluminance of the saturated-red and saturated-green target stimuli for each subject by means of heterochromatic flicker photometry (Kaiser, & Comerford, 1975; Wagner, & Boynton, 1972). In the suppression task, these patches increased in contrast, reaching full contrast at 1,500 ms. Patches were presented to the participant's nondominant eye. The dominant eye was simultaneously presented with a high-contrast dynamic pattern to evoke CFS. These pattern masks were created by (a) filtering pink (1/f) noise using a rotationally symmetric Gaussian low-pass filter ($\sigma = 3.5$) and by (b) making the resulting gray-scale image binary with maximum contrast (> 96% Michelson contrast). On every trial, 200 new masks were generated, which were presented for 10 frames each (10 Hz). Note that the target stimuli in the suppression task were never of the same identical hue as the colors used for the memory task.

Procedure

Each trial began with a fixation cross (500 ms), followed by a color stimulus (1,000 ms), which participants were told to remember for later recall (color-memory condition) or to only fixate without memorizing (passive-viewing condition). After a blank screen (2,000 ms), the suppression task was initiated. Participants were instructed to report as soon as possible whether a stimulus appeared either to the left or to the right of fixation (1.8° of eccentricity). The color of the target was irrelevant for the suppression task. This task lasted until a response was given or until 20 s elapsed, and it ended with a blank screen (500 ms). In the color-memory condition, participants were then presented with three color patches of the same color category, each with a different hue, one of which was identical to the color presented at the start of the trial. They were instructed to indicate which of these three colors had the exact same color as the memorized stimulus. The stimulus sequence of a congruent trial is depicted in Figure 1 (top left).

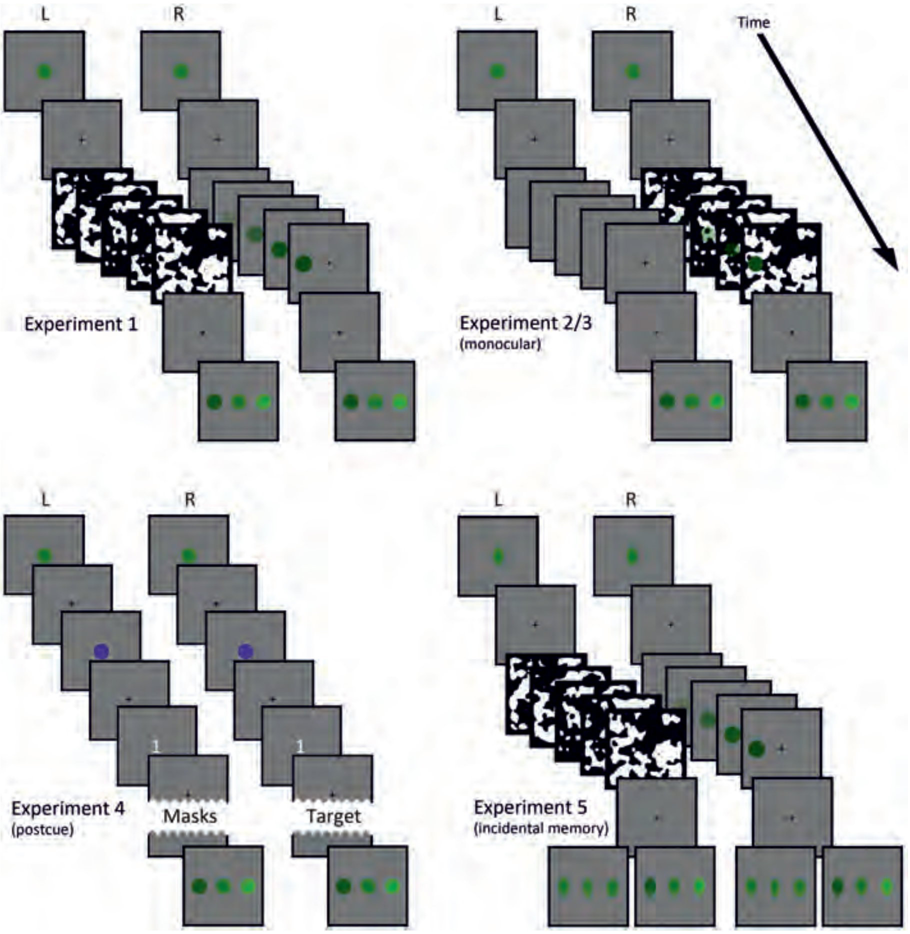

Figure 1. Stimulus sequence of a congruent trial in each of the five experiments, as presented to the left and right eye. In all experiments, observers were required to report whether a colored stimulus appeared left or right of fixation, during the presentation of the dynamic masks (i.e., a b-CFS task). Depending on the trial, this stimulus could be congruent or incongruent with a stimulus that was presented before the b-CFS task. In Experiment 1, 2, 3, and 5, participants were either instructed to memorize the color of this first stimulus for a subsequent recall task (color memory blocks) or not (passive viewing blocks). In Experiment 5, participant could also be instructed to memorize the shape of this stimulus. In Experiment 4, participants were sequentially presented with two colored stimuli and a cue (the numbers "1" or "2"). This cue instructed them to memorize either the first color or the second color.

Participants performed a total of 108 trials. Color patches for the memory task were randomly selected (but counterbalanced) from three basic color categories (red, green, or blue). Within each color category, the specific hue was randomly selected from one of five different color variations. There were 36 congruent trials, in which the target was of the same color category as the memorized stimulus, and 72 incongruent trials, in which target color and memorized color did not match; these two trial types were randomized within blocks. There were 54 color-memory trials and 54 passive-viewing trials, which were blocked. The order in which the blocks were presented was counterbalanced across participants. Both memory conditions (color memory and passive viewing) were preceded by an 8-trial practice session.

Data analysis

Only the trials in which the correct target location (left/right) was reported were included in the response-time (RT) analyses of the suppression task (less than 1% of the responses were incorrect). We determined participants' median RT for each memory condition (passive viewing and color memory) and target congruency (congruent and incongruent), and averaged them over participants. In each experiment, a 2 × 2 repeated measures analysis of variance (ANOVA) was conducted with memory condition and target congruency as factors.

Results

Results are depicted in Figure 2. The content of working memory affected suppression durations such that targets matching VWM content reached visual awareness faster than targets that did not match VWM content. This was revealed by the interaction between memory condition and target congruency, $F(1, 9) = 19.83$, $p = 0.002$. Subsequent paired-samples t tests showed that in the color-memory condition, congruent targets broke through suppression faster than did incongruent targets, $t(9) = 4.56$, $p = 0.001$, whereas this congruency effect was not observed in the passive-viewing condition ($p > 0.8$). The mean size of the congruency effect in the color-memory task was 238 ms ($SD = 165$ ms, range = 42–600 ms). Finally, accuracy on the memory task was well above the 33% chance level ($M = 68\%$ correct, $SD = 10\%$).

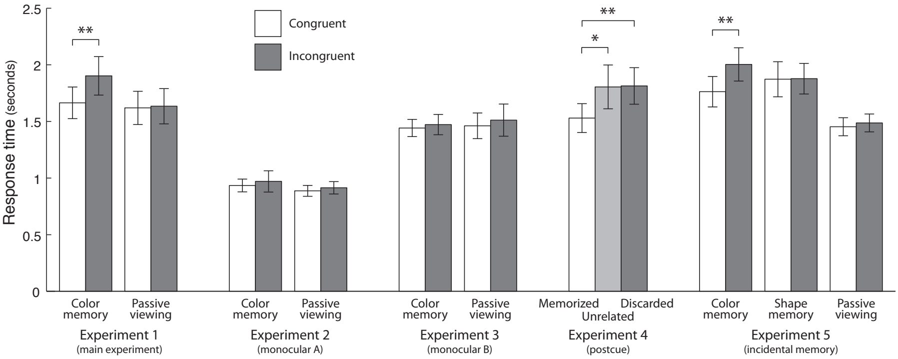

Figure 2. Mean response times as a function of experiment and condition. For Experiments 1, 2, 3, and 5, results are shown for congruent and incongruent trials in the color-memory and passive-viewing conditions. For Experiment 5, results are also shown for the shape-memory condition. For Experiment 4, results are shown for trials on which the color of the target matched the memorized color (the "memorized" condition), matched the color that participants saw but were not instructed to memorize (the "discarded" condition), and matched the color that had not been presented on that trial (the "unrelated" condition). Error bars show standard errors of the mean. *$p < 0.01$, **$p < 0.001$

The results of Experiment 1 supported our hypothesis that information matching the content of VWM is prioritized for conscious access relative to nonmatching information. The next experiments were dedicated to tackling possible confounds. Our aim was to ground the claims that we were indeed measuring conscious access (Experiments 2 and 3), manipulating VWM (Experiment 4), and dealing with prioritization (Experiment 4) and that this phenomenon is related to the relevance of the visual information participants were presented with (Experiment 5).

Experiments 2 and 3

Although differences in detection times under CFS provide a good indication that there are differences in nonconscious processing, these differences could also be accounted for by a difference in response criterion between the passive-viewing condition and the color-memory condition. For instance, a stimulus in VWM could be associated with a more liberal response tendency (i.e., a lower threshold in deciding that a stimulus has been perceived), which would result in faster RTs for matching targets than for nonmatching targets (for a discussion of this issue, see Stein et al., 2011). To ensure that the differences in RTs between congruent and incongruent trials in the color-memory condition reflect faster conscious access rather than a change in response criterion, we conducted two monocular control experiments. In these experiments, both the target and the dynamic pattern were presented to the same eye, as a result of which the target was not rendered invisible by interocular suppression. Hence, if VWM content speeded up responding to targets by other means than prioritization of conscious access, it would result in a larger congruency effect in the color-memory condition than in the passive-viewing condition in these experiments as well. Additionally, in Experiment 3, the time it took for the stimulus to increase full contrast was increased.

Methods

Experiments 2 and 3 were each completed by 10 new participants. Participants ranged in age from 18 to 31 years, were right-handed, had normal or corrected-to-normal vision, and passed the TNO test. Eye dominance was determined as in Experiment 1. In Experiment 2, all presentation times and stimuli were identical to those in Experiment 1, the only difference being that the target stimulus was presented to the same eye as the dynamic pattern (Fig. 1). Because targets were not suppressed, the distribution of detection times differed from that of Experiment 1 (for a discussion of this issue, see Stein et al., 2011). Hence, we conducted another monocular experiment (Experiment 3), in which we lengthened the ramp-up of the target, such that it would reach full contrast after 5,000 ms rather than after 1,500 ms. Additionally, we randomly varied the onset of the stimulus between 0 and 1,000 ms after CFS onset, so as to minimize anticipation effects. RTs were measured from target onset.

Results

The results of Experiments 2 and 3 show that detection times for non-suppressed color patches were not affected by the content of working memory (see Fig. 2). First, RTs on congruent trials did not differ significantly from RTs on incongruent trials either in the color-memory condition (Experiment 2: $p > 0.1$; Experiment 3: $p > 0.4$), or in the passive-viewing condition (both experiments: $p > 0.2$). Second, and more important, this congruency effect did not differ between the color-memory condition and the passive-viewing condition. This was revealed by the absence of an interaction between memory condition and target congruency in both the physically similar (Experiment 2) and the perceptually similar (Experiment 3) monocular experiments (both p's > 0.7). Mean accuracy on the memory task was 73% correct ($SD = 8\%$) in Experiment 2 and 68% correct ($SD = 9\%$) in Experiment 3. Together, these monocular experiments show that the RT differences in Experiment 1 reflected faster conscious access, rather than a difference in response criterion, in the color-memory condition than in the passive-viewing condition.

Experiment 4

In Experiment 1, participants knew they had to recall the first stimulus for later report in the color-memory condition and that they would not have to recall it in the passive-viewing condition. Hence, the cued color in the color-memory condition was not only held in working memory, but also it was potentially subject to more extensive initial processing than in the passive-viewing condition. To differentiate between these two accounts, we conducted an experiment designed to keep the initial processing of the color stimuli constant, manipulating only the content of VWM.

In addition, Experiment 1 did not allow for direct RT comparisons between the passive-viewing condition and the color-memory condition, because the latter involved a dual-task load, which slowed down target detection (e.g., Pashler, 1994). By keeping the load equal between conditions, we were able in Experiment 4 to compare RTs between memory conditions. This enabled us to distinguish facilitatory effects of congruent trials from inhibitory effects of incongruent trials.

Methods

Fifteen new participants (age 19 to 28 years) completed Experiment 4. Participants were right-handed, had normal or corrected-to-normal vision, and passed the TNO test. Eye dominance was determined as in Experiment 1. In this experiment, participants were presented with two stimuli of different colors (1,000 ms each) separated by a blank interval (1,500 ms). The second stimulus was followed by a blank screen (1,500 ms), after which a postcue appeared. This cue (the digit "1" or "2") instructed participants to actively retain either the first or the second color stimulus for later recall (Fig. 1). The suppression task and the recall phase in this experiment were identical to those in Experiment 1. Target stimuli were circular and created from the same three hues as in Experiment 1. On a given trial, the target could be from the same color category as the memorized color, from the same color category as the discarded color (i.e., the one participants saw but were not instructed to memorize), or from the unrelated color category (i.e., the one that did not appear on that trial). Because this procedure leads to two levels of congruence (congruent with memorized vs. unrelated vs. discarded stimulus) rather than one (congruent vs. incongruent), we tested more participants in this experiment to obtain comparable statistical power. Participants performed a total of 72 trials (24 per condition) over three blocks. All variables were randomized within blocks. A 3 × 2 repeated measures ANOVA was conducted with the factors condition (memorized, unrelated, discarded) and position (first stimulus is cued, second stimulus is cued).

Results

Results are depicted in Figure 2. First, the results of Experiment 4 show that VWM content alone affected suppression durations, such that targets matching VWM contents broke through interocular suppression faster than did nonmatching targets. This was revealed by a main effect of memory condition, $F(2, 28) = 10.46$, $p < 0.001$, and the absence of either a main effect of position ($p > 0.8$) or an interaction between memory condition and position ($p > 0.7$). Second, the results of this experiment confirm that information matching VWM content is prioritized rather than nonmatching information being inhibited. Subsequent paired-samples *t*-tests showed faster breakthrough times for targets matching the memorized color compared with targets matching the unrelated color, $t(14) = 2.96$, $p = .010$, as well as compared to targets matching the discarded color, $t(14) = 4.76$, $p < 0.001$. Breakthrough times for targets matching the discarded color did not differ from breakthrough times for targets matching the unrelated color ($p > 0.8$). Finally, no difference in performance was observed on the memory task between stimuli presented at Position 1 ($M = 64\%$ correct, $SD = 8\%$) and Position 2 ($M = 63\%$ correct, $SD = 10\%$), which shows that participants were equally proficient in retaining either one ($p > 0.8$).

Experiment 5

Luck and Vogel (1997) demonstrated that objects in VWM are stored as integrated wholes rather than as independent features. As such, an object in VWM contains a relevant stimulus dimension, which is volitionally memorized (as the task requires; e.g., its shape), but also contains incidentally memorized dimensions (which happen to be part of that object; e.g., its color). Experiment 5 was set up to test whether incidental memory for color would also prioritize conscious access of a target matching the memorized color category. This was achieved by having participants retain either the shape (shape-memory condition) or the color (color-memory condition) of a stimulus and assessing the effect of color congruency on suppression durations of a colored target stimulus. Accordingly, color was the relevant stimulus feature in the color-memory condition, whereas it was the irrelevant feature in the shape-memory condition.

Methods

Ten new participants (age 18 to 30 years) completed Experiment 5. Participants were right-handed, had normal or corrected-to-normal vision, and passed the TNO test. Eye dominance was determined as in Experiment 1. The stimuli and procedure of this experiment were similar to those of Experiment 1, with the addition of a condition (two blocks of 54 trials each) in which participants were instructed to retain the shape of the stimulus rather than the color. The stimuli for the memory task were either horizontal or vertical ellipses (two categories) that differed in height and width (five variations per category, ranging respectively in width from 0.6° to 1.2° and in height from 1.2° to 0.6°). Up to the recall phase of the experiment, stimulus presentation was identical in the color- and shape-memory conditions. These conditions only differed in the instruction to retain one feature or the other. In the recall phase of the shape-memory condition, three shape variations of the same shape category were presented, with all three stimuli having the same color as the memorized stimulus. Participants had to judge which shape corresponded with the one they had memorized. Conversely, in the color-memory condition, three hue variations of the same color category were presented, with all three stimuli having the same shape as the memorized stimulus (Fig. 1). As in the previous experiments, participants had to judge which color corresponded with the one they had memorized.

Results

The results of this experiment show that only the relevant feature of a stimulus stored in VWM prioritizes conscious access of a matching target stimulus (see Fig. 2). This was revealed by an interaction between memory condition (color memory, shape memory, passive viewing) and target congruency (congruent, incongruent), $F(2, 8) = 8.31$, $p = 0.011$. Subsequent paired-samples t tests showed that this interaction was caused by the congruency effect in the color-memory condition: Congruent targets broke through interocular suppression faster than did incongruent targets in the color-memory condition, $t(9) = 5.31$, $p < 0.001$, but not in the shape-memory condition ($p > 0.9$) or the passive-viewing condition ($p > 0.5$). No difference in performance was observed on the memory task between shape memory (M = 70% correct, SD = 16%) and color memory (M = 69% correct, SD = 9%), which shows that both tasks were equally demanding ($p > 0.8$).

General Discussion

The literature on VWM is mostly concerned with information that is conscious (e.g., Baars, 2003; Baddeley, 2003; but see Soto, Mäntylä, & Silvanto, 2011). Here, we showed that the content of VWM affects processing of visual information that is not yet available for conscious inspection. Crucially, we demonstrated a direct functional link between VWM and the selection of information for visual awareness: Information that matches VWM content is prioritized by the visual system so that it reaches visual awareness faster than nonmatching information does. The magnitude of this effect (about a 14% RT difference between matching and nonmatching targets) was very consistent (see Fig. 2), even when different stimuli (Experiment 5) or different stimulus sequences (Experiment 4) were used. In addition, prioritization for conscious access was selective for the relevant stimulus dimension of a memorized stimulus.

It is important to ground the claim that the b-CFS method used in the current experiments assesses conscious access. In these experiments, the suppressed target stimuli were identical in all conditions. This is different from other studies, in which differences in breakthrough times can potentially be accounted for by stimulus differences (as argued by Stein et al., 2011). Hence, if, as in the present case, suppressed stimuli are identical, an experimental manipulation that affects the time it takes for a stimulus to break through suppression reflects differential processing of that stimulus before it broke through suppression. As such, differences in suppression durations reveal differences in unconscious processing caused by the manipulation (for similar conclusions, see Gaillard et al., 2006; Stein et al., 2011). Still, detection times can also be affected by a number of cognitive or perceptual processes that arise after conscious detection of the stimuli, such as differences in response criterion. This alternative explanation was tackled in Experiments 2 and 3, which showed that VWM content did not affect detection times of non-suppressed targets.

From a neuronal perspective, the influence of VWM content on suppression durations might reflect the competition between inhibition and activation of a stimulus representation in early visual areas. Neuroimaging studies have shown that representations retained in VWM can be decoded from activity in the early visual cortex (e.g., Christophel, Hebart, & Haynes, 2012; Harrison & Tong, 2009;

Serences, Ester, Vogel, & Awh, 2009). Early visual areas have also been associated with visual awareness (e.g., Ress & Heeger, 2003; Super, Spekreijse, & Lamme, 2001; Weiskrantz, 1986; 1997) and, more specifically, with interocular suppression, which is partly accounted for by reciprocal inhibition of visual areas that retain eye-selective information (for a review, see Blake, 1989; Tong, 2001; Tong, Meng, & Blake, 2006). This spatial overlap in neuronal substrates for VWM representations and visual awareness provides the grounds for our behavioral results. In addition, this approach might explain why no color congruency effect was obtained when the shape of a colored stimulus was memorized (Experiment 5): Serences et al. (2009) showed that only the intentionally stored feature of a remembered stimulus (i.e., color but not orientation and vice versa) elicited a sustained pattern of activation in early visual areas. Thus, based on the idea of competition for visual awareness, only the intentionally stored feature has the potency to shorten suppression durations.

The present results can be framed within the global workspace model of consciousness (Dehaene et al., 1998; Dehaene & Naccache, 2001), according to which dedicated modular systems in the brain operate nonconsciously. When people face demanding tasks, however, appeals made on these systems are often conflicting, which requires the mediation of a distributed and flexible neural network. This broad neural network, then, reflects the content of consciousness (e.g., Baars, 2005; Lau & Rosenthal, 2011; Dehaene & Naccache, 2001). The implication of this model is that the current state of consciousness (reflecting strategy, intentions, etc.) affects both conscious and nonconscious processes. In accordance with this, voluntarily retained VWM content is known to affect conscious processes, such as attention allocation (Olivers et al., 2006), and nonconscious processes, such as priming (Pan et al., 2012). The present experiments provide evidence that the current conscious state (i.e., the volitional retention of information) influences visual processing in yet another way: It affects the selection of sensory information in the visual system that is not yet available for conscious report in order for it to breach the threshold of consciousness.

Conclusion

The present results reveal a functional connection between visual awareness and VWM. Whereas VWM is used to retain relevant visual information for imminent goal-directed behavior, visual awareness is needed to flexibly deal with incoming information to guide future behavior. Considering that people are continuously presented with vast amounts of sensory information, a system is needed to select and prioritize the most relevant information. The results of the present experiments suggest that VWM might well play that role in human consciousness; it funnels down the incoming sensory information to that which is relevant for imminent goal-directed behavior.

Part III

Biasing of visual input by the content of visual working memory

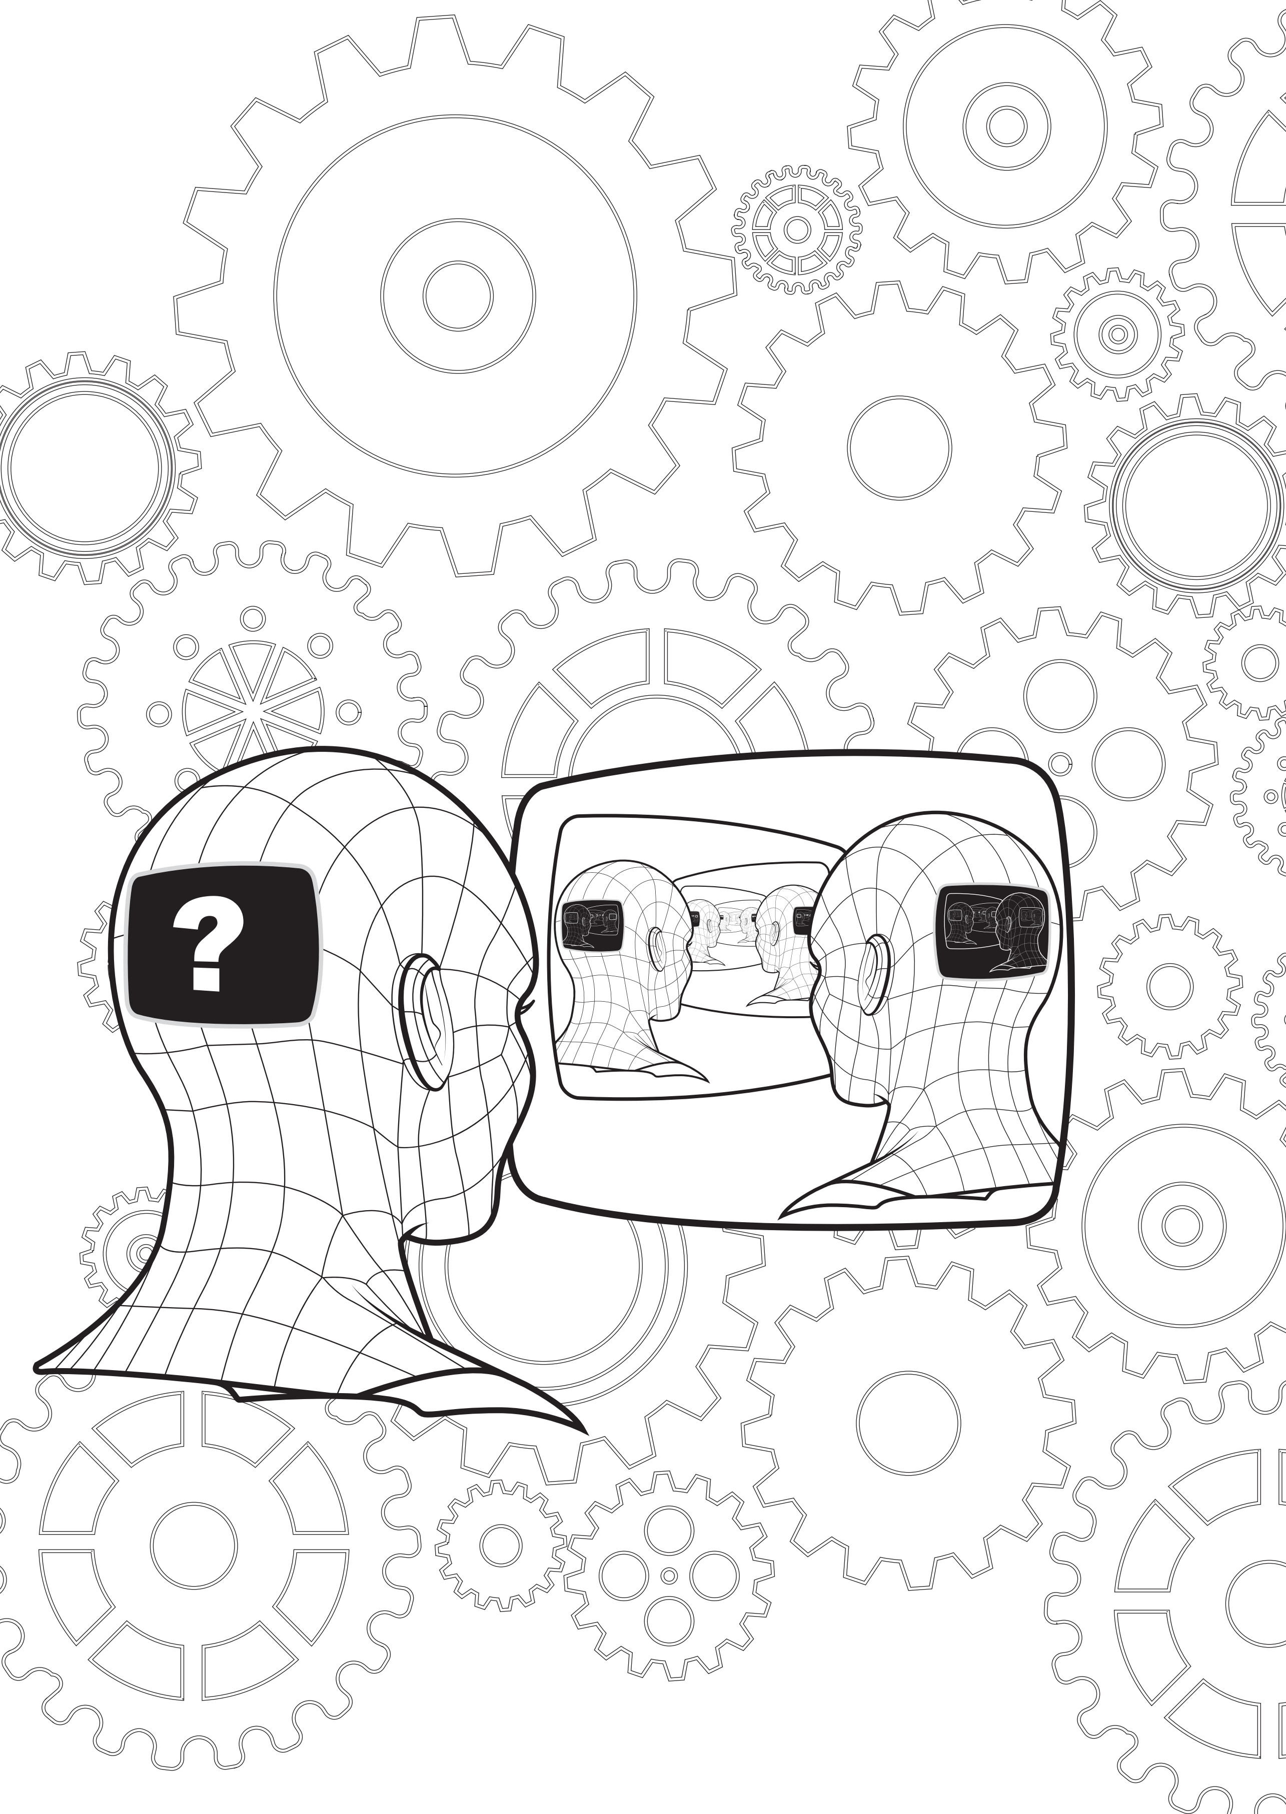

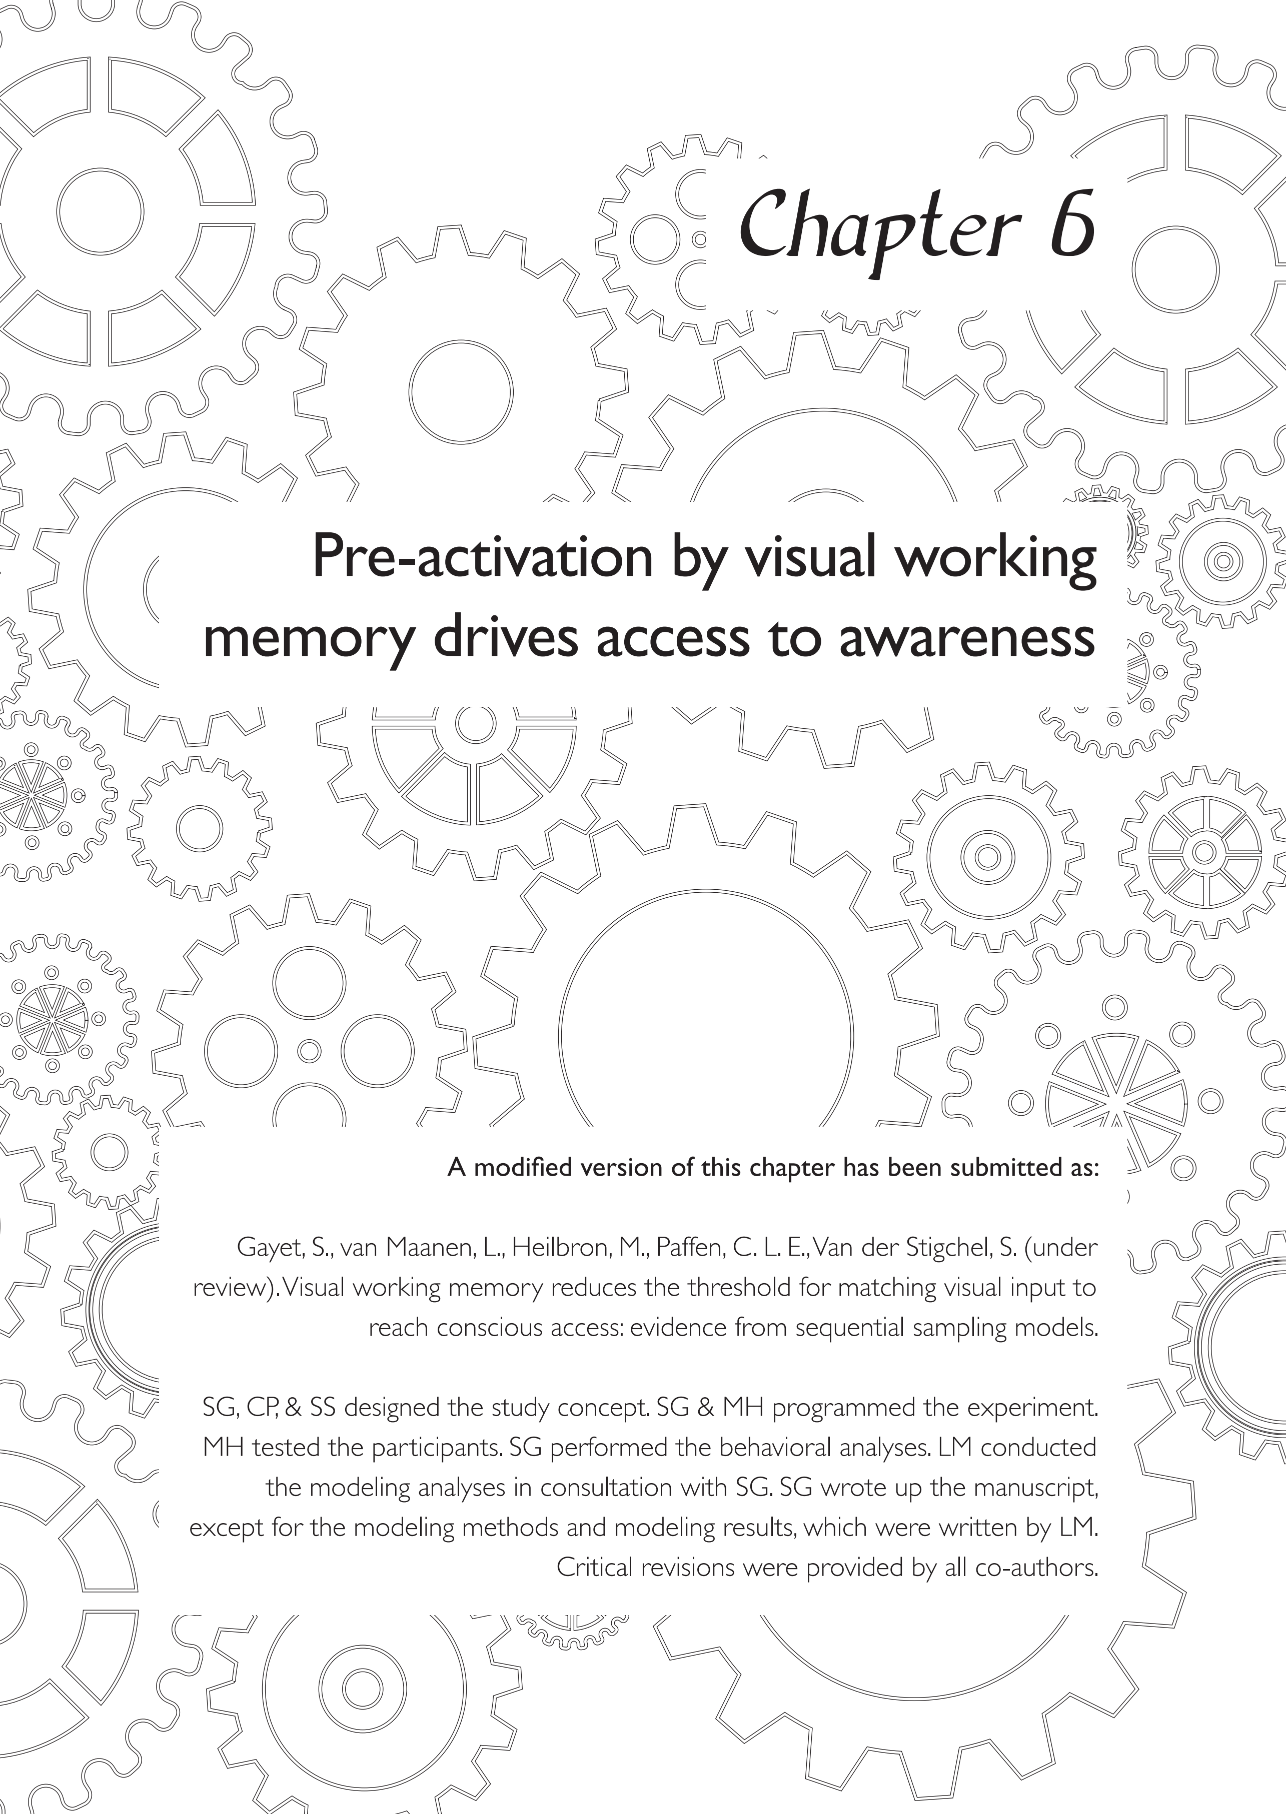

Chapter 6

Pre-activation by visual working memory drives access to awareness

A modified version of this chapter has been submitted as:

Gayet, S., van Maanen, L., Heilbron, M., Paffen, C. L. E., Van der Stigchel, S. (under review). Visual working memory reduces the threshold for matching visual input to reach conscious access: evidence from sequential sampling models.

SG, CP, & SS designed the study concept. SG & MH programmed the experiment. MH tested the participants. SG performed the behavioral analyses. LM conducted the modeling analyses in consultation with SG. SG wrote up the manuscript, except for the modeling methods and modeling results, which were written by LM. Critical revisions were provided by all co-authors.

Abstract

The content of visual working memory (VWM) affects the processing of concurrent visual input. Recently, it has been demonstrated that stimuli are released from interocular suppression faster when they match rather than mismatch a color that is memorized for subsequent recall. In order to investigate the nature of the interaction between visual representations elicited by VWM and visual representations elicited by retinal input, we modelled the perceptual processes leading up to this difference in suppression durations. We replicated the VWM modulation of suppression durations, and fitted sequential sampling models (linear ballistic accumulators) to the response time data. Model comparisons revealed that the data was best explained by a decrease in threshold for visual input that matches the content of VWM. Converging evidence was obtained by fitting similar sequential sampling models (shifted Wald model) to published datasets. Finally, to confirm that the previously observed threshold difference reflected processes occurring before rather than after the stimuli were released from suppression, we applied the same procedure to the data of an experiment in which stimuli were not interocularly suppressed. Here, we found no decrease in threshold for stimuli that match the content of VWM. We discuss our findings in light of a *Pre-activation* hypothesis, proposing that matching visual input taps into the same neural substrate that is already activated by a representation concurrently maintained in VWM, thereby reducing its threshold for reaching visual awareness.

Introduction

Working memory and consciousness are two well-studied phenomena that share a lot of similarities. For instance, both consciousness (Lamme, & Roelfsema, 2000; Sergent, Baillet, & Dehaene, 2005; Boehler, Schoenfeld, Heinze, & Hopf, 2008) and working memory (Raffone & Wolters, 2001; D'Esposito, 2007) have been proposed to rely upon recurrent cortical processing, and both consciousness (Baars, 1997c; Dennett, 1991; Edelman & Tononi, 2000) and working memory (e.g., Luck, & Vogel, 1997; Cowan, 2001) are of limited capacity. The traditional stance is that working memory only operates on information that is accessible to consciousness (e.g., Stein, Kaiser, & Hesselmann, 2016; Baars, & Franklin, 2003; Baars, 1993; but, see Soto, & Silvanto, 2014). In fact, consciousness has even been described as 'the remembered present' (Edelman, 1989), and phenomenal experience and access consciousness have been related to iconic memory and working memory respectively (Lamme, 2004). Despite the close kinship between consciousness and working memory, the question of how working memory interacts with conscious perception has not received much attention until recent years. Focusing on the visual modality, we investigated the mechanism by which the content of visual working memory (VWM) affects access to visual awareness.

The potency of visual input to reach awareness is affected by concurrent working memory maintenance in a number of distinct ways. The threshold for stimuli to reach visual awareness increases with working memory load (Maniscalco & Lau, 2015; De Loof, Verguts, Fias, & Van Opstal, 2013; Lavie, 2005) and even more so when the working memory content is visual in nature (De Loof, Poppe, Cleeremans, Gevers, & Opstal, 2015; Konstantinou, & Lavie, 2013). Recently, it has been demonstrated that the specific content of VWM also modulates access to awareness. For instance, visual input that matches rather than mismatches a color maintained in VWM is more likely to dominate in binocular rivalry (Gayet, Brascamp, Van der Stigchel, & Paffen, 2015; but, see Scocchia, Valsecchi, Gegenfurtner, & Triesch, 2014). Similarly, maintaining an unambiguous motion direction in VWM biases the perceived motion direction of an ambiguous structure from motion sphere towards the memorized motion direction (Scocchia, Valsecchi, Gegenfurtner, & Triesch, 2013).

Related findings that are of particular interest for the present purpose are those using the breaking continuous flash suppression method (b-CFS; Jiang, Costello, &

He, 2007; for a review, see Gayet, Van der Stigchel, & Paffen, 2014b). In this method, a target stimulus is initially suppressed from awareness by means of continuous flash suppression (CFS; Tsuchiya, & Koch, 2005). The time it takes for observers to report the location of the initially suppressed target provides a measure of access to awareness. Using b-CFS, it has been demonstrated that visual input that matches a color (Gayet, Paffen, & Van der Stigchel, 2013; van Moorselaar, Theeuwes, & Olivers, 2015), shape (Supplementary Materials S2 of Chapter 7), or face (Pan, Lin, Zhao, & Soto, 2014) maintained in VWM is released from continuous flash suppression faster than mismatching visual input. Considering that, in these cases, the content of VWM codetermined when visual input reached awareness, it can be deduced that the content of VWM impacts the processing of visual awareness before it reaches (full blown) awareness.

The shorter suppression durations for stimuli that match rather than mismatch the content of VWM is likely potentiated by an enhanced neural response of visual input that matches the content of VWM. Reviews on CFS (Sterzer, Stein, Ludwig, Rothkirch, & Hesselmann, 2014; Yang, Brascamp, Kang, & Blake, 2014), b-CFS (Gayet et al., 2014b) and binocular rivalry (Lin & He, 2009) reveal that interocularly suppressed information is processed only to a very limited extent. That is, only coarse processing of visual features is preserved, whereas higher level (e.g., semantic) information is severely disrupted. Considering that, when climbing up the visual processing hierarchy, a decreasing proportion of neurons represents visual input that is interocularly suppressed (Leopold & Logothetis, 1999; Yuval-Greenberg, Heeger, 2013), the content of VWM is most likely to modulate the signal elicited by the suppressed sensory input in early visual areas. The present study is aimed at dissociating between two possible ways in which representations maintained in VWM can modulate concurrent visual input in early visual processing areas. For this purpose, we replicated the finding that visual input matching the content of VWM yields faster access to awareness in a b-CFS task (Gayet et al., 2013; Pan et al., 2014; van Moorselaar et al., 2015), and compared the potency of two models in explaining the observed pattern of results.

The *Pre-activation* hypothesis proposes that representations elicited by visual input rely upon the same neural substrate as visual representation elicited by VWM. As a result of this, there early visual areas have an initial bias towards visual input that matches the content of VWM. The *Feedback* hypothesis, on the other hand, proposes that early visual processing of the visual input is modulated by higher level VWM representations. Through reciprocal connections, the suppressed visual input is continuously compared to the VWM representation, and is selectively enhanced if it matches the content of VWM. In the next sections, we elaborate on these two hypotheses, their core assumptions, and the empirical findings from which they are diluted. Next, we describe the different predictions that follow from the two hypotheses. Finally, we present model implementations of the *Pre-activation* and the *Feedback* hypotheses, allowing for comparing which of these models best describes the observed behavioral data.

According to the *Pre-activation* hypothesis, representations maintained in VWM (say, the color green) elevate the base activation level for concurrent matching visual input (a green stimulus). The *Pre-activation* hypothesis builds upon the assumption that a visual representation (the color green) draws upon the same neural population, irrespective whether its origin is retinal or mnemonic. This is in line with the recent view that content-based VWM processes reside in visual processing areas (for a discussion, see Sligte, Van Moorselaar, & Vandenbroucke, 2013) and is corroborated by recent imaging studies, in which multivariate pattern analysis allows for decoding the content of VWM from neural activity patterns in visual processing areas (e.g., Christophel, Hebart, & Haynes, 2012; Riggall, & Postle, 2012; Harrison & Tong, 2009; Serences, Ester, Vogel, & Awh, 2009). Additionally, training a classifier on the pattern of neural activity elicited by VWM content allowed for above chance classification of the stimuli that were presented on the screen, and vice versa (Harisson, & Tong, 2009). Considering a shared neural substrate for visual representations, irrespective of whether their origin is retinal or mnemonic, visual input might draw upon a neural population that was already activated by VWM maintenance. Such an additive effect of matching visual representations would result in enhanced activation of visual input that matches the concurrent content of VWM. By this, pre-activation of specific neural populations by the content of VWM should decrease the effective threshold for matching visual input to affect behavior and perception (a similar view has been proposed in a priming study by Lupyan, & Ward, 2012).

According to the *Feedback* hypothesis, the content of VWM is maintained in visual areas that are further upstream in the visual hierarchy than the visual areas where the interocular competition is resolved. Feedback-loops between these higher visual areas (containing representations of, say, the 'green color category') and the lower visual areas (processing the green stimulus) allow for enhancing visual input that complies with the categorical representation maintained in higher visual areas. This hypothesis is in line with the traditional view that the contents of VWM is coded in widespread activation over parietal and/or prefrontal regions (e.g., Pessoa, Gutierrez, Bandettini, Ungerleider, 2002; Postle, 2006). When considering the manipulation of VWM in b-CFS studies, there are at least two types of observations suggesting that the differentiation between matching and mismatching stimuli can only be construed in higher level visual areas. Firstly, in these b-CFS studies, the to-be-memorized stimulus was never presented at the same retinal location as the suppressed target stimulus presented during the delay interval (i.e., typically they are separated by 1 to 2 degrees of visual angle). From this it follows that the effect is either non-retinotopic or originates in visual areas with sufficiently large receptive fields to cover the retinotopic distance between the to-be-memorized stimulus and the suppressed target stimulus. Receptive fields in V1-4 would be unable to account for this, whereas receptive fields in object selective regions, such as LO1 might suffice, according to recent receptive field size measurements obtained in humans (Harvey & Dumoulin, 2011; note, that his method measures *population* receptive field sizes and therefore overestimates receptive field sizes). Secondly, in the b-CFS studies mentioned above, the to-be-memorized stimulus was never identical to the suppressed stimulus. Rather, matching or mismatching conditions reflected whether the suppressed stimulus was drawn from the same or from a different stimulus category than the memorized stimulus (i.e., the effect was categorical). The categorical nature of this effect leads to suggest that the effect originates in higher level visual areas (e.g., object selective areas such as LOC or non-visual areas, containing abstract coding of stimulus categories). In support of this hypothesis, it has been demonstrated that the content of VWM can sometimes be decoded from higher level visual areas as well (Christophel, Hebart, & Haynes, 2012; Riggall, & Postle, 2012). An important requirement for this hypothesis, however, is that some of the CFS suppressed visual input transpires to higher level visual areas (for a review on what information yields residual activity patterns under CFS, see Sterzer et al., 2014). The view that the modulation of suppression durations in a b-CFS paradigm relies on suppressed visual input transpiring to higher level visual areas is substantiated by the recent finding that certain priming effects under CFS only occur under conditions of partial awareness (Mudrik, Gelbard-Sagiv, Faivre, & Koch, 2013; Gelbard-Sagiv, Faivre, Mudrik, & Koch, 2016).

Both the *Pre-activation* hypothesis and the *Feedback* hypothesis predict a stronger neural response to matching visual input in (early) visual processing areas. If we consider the time course over which a stimulus is expected to be enhanced depending on its match with VWM, however, the two hypotheses provide very distinct predictions. According to the *Pre-activation* hypothesis, matching (but not mismatching) visual input should be immediately enhanced, as it draws upon a pre-activated neural population, akin to a decreased threshold. According to the *Feedback* hypothesis, the differentiation between matching and mismatching visual input should occur through gradual evidence accumulation, resulting in a positive gain modulation. Sequential sampling models (for reviews, see Forstmann, Ratcliff, & Wagenmakers, 2016; Ratcliff & Smith, 2004) provide the means to infer, from response times and response choices, the process of evidence accumulation over time that eventually leads to a perceptual decision. A core assumption in these model is that dichotomous perceptual decisions are the result of gradual evidence accumulation over time. Such models have been used before to investigate evidence accumulation in early visual processing areas (e.g., Ho, Brown, van Maanen, Forstmann, Wagenmakers, & Serences, 2012).

In the present study, we conducted a modified version of the experimental paradigm used by Gayet, et al. (2013). Similarly to the traditional paradigm, suppression durations were inferred from speeded responses to target stimuli that were initially interocularly suppressed. Participants were required to report on which side of fixation a target appeared. In the present study, there were two targets (one that was related to the memory task and one that was unrelated to the memory task), one on each side of fixation. Participants than indicated which of the two targets (left or right) broke through the suppression *first*. This critical modification allowed us to retrieve both response speed and response choice data, thereby optimizing the traditional paradigm for fitting sequential sampling models to participants' behavioral responses.

First, we replicated the finding that visual input matching the content of VWM gains preferential access to awareness in a b-CFS paradigm. Next, we implemented both the *Pre-activation* and *Feedback* hypotheses as sequential sampling models, to compare which of these quantitatively best explained the data. To ensure that our findings would generalize to the typical b-CFS paradigm, we additionally re-analyzed the data of two experiments retrieved from Gayet et al. (2013) with a sequential sampling model tailored to the type of data typically obtained in the b-CFS paradigm (the shifted Wald model, Anders, Alario, & van Maanen, in press; Heathcote, 2004). Together, our results unequivocally favor the *Pre-activation* hypothesis over the *Feedback* hypothesis. From this, we conclude that neuronal populations are pre-activated by representations maintained in VWM, as a result of which matching visual input is prioritized over mismatching visual input for gaining access to awareness.

Methods

Participants

The participant group consisted of 22 undergraduate students from Utrecht University (8 males; 21.6 years, *SD* = 3.0), who participated for course credits or monetary reward after signing informed consent. The experiment complied with all ethical guidelines set out in the Declaration of Helsinki. Participants had (corrected to) normal vision and were tested for color blindness with the Ishihara test plates (Ishihara, 1917) and for stereoscopic vision with the TNO test for stereoscopic vision (12th edition; Walraven, 1972). The number of participants was determined on the basis of prior b-CFS experiments (e.g., Gayet et al., 2015) and experiments with comparable implementations of sequential sampling models (Forstmann, et al, 2010).

Apparatus and stimuli

The experiment was conducted using an Apple dual 2-GHz PowerPC G5 equipped with a linearized 22-inch LaCie Electron Blue IV CRT monitor (1,024 × 768 pixels; 100-Hz refresh rate) and an Apple keyboard, which was used for response registration. There were no light sources in the experiment room, except for the computer monitor. Stimulus presentation and response collection were managed using the Psychophysics Toolbox 3 (Brainard, 1997; Pelli, 1997) in MATLAB (R2009b; The Mathworks, Natick, MA). A pair of displays was viewed dichoptically through a mirror stereoscope mounted on a chin rest, which kept the effective viewing distance at 57 cm. To facilitate binocular fusion of the two complementary images, we presented a circular area with a diameter of 8° to each eye, which was composed of Brownian (i.e., $1/f^2$) noise with a mean luminance of 16.08 cd/m^2. All stimuli were presented on a uniform gray background, (16.62 cd/m^2), that consisted of two circular areas with a diameter of 4.95°, superimposed on the Brownian noise. Both

Table 1. *Overview of the colors used in the main experiment.*

	CIE values*		
Color	X-value	Y-value	Luminance (cd/m^2)
Red 1	0.608	0.364	3.76
Red 2	0.632	0.345	3.50
Red 3	0.651	0.340	3.47
Red 4	0.581	0.312	3.24
Green 1	0.269	0.528	3.31
Green 2	0.287	0.600	3.53
Green 3	0.310	0.597	3.25
Green 4	0.362	0.557	3.33
Blue 1	0.166	0.137	3.45
Blue 2	0.156	0.103	3.37
Blue 3	0.159	0.079	2.87
Blue 4	0.171	0.083	2.66
Purple 1	0.228	0.123	3.47
Purple 2	0.251	0.133	3.34
Purple 3	0.293	0.151	3.40
Purple 4	0.349	0.182	3.22
Blue target**	0.147	0.068	2.64
Red target (SD)	0.635	0.343	2.61 (0.56)
Green target (SD)	0.292	0.612	4.24 (0.97)
Purple target (SD)	0.319	0.165	2.95 (0.50)

* CIE values stands for Commission Internationale d'Eclairage values, as measured from viewing distance (i.e., 57 cm) with a PR-650 SpectraScan colorimeter/telephotometer (Photo Research, Inc.).

** The saturated blue was the reference stimulus for the perceptual luminance matching (i.e., heterochromatic flicker photometry; Kaiser & Comerford, 1975; Wagner & Boynton, 1972) with the red, green and purple colors that were used for the target stimuli.

presentation areas comprised a white (32.15 cd/m^2) fixation bull's-eye composed of a white circle with a diameter of 0.2° of visual angle and a black (<0.1 cd/m^2) circle in its center with a diameter of 0.04° of visual angle. The retro cues consisted of the Arabic numerals "1" and "2" in white Arial font with a font size of 18.

The masks used for obtaining CFS were created by filtering pink (1/f) noise using a rotationally symmetric Gaussian low-pass filter ($\sigma = 3.5$) and making the resulting image binary (black and white, > 99% Michelson contrast). On every trial, 20 new masks were generated, which were presented for 100 ms each (10 Hz) in random order, with the restriction that the same mask was never presented twice in succession. During presentation, a vertical sparing with a width of 0.8° of visual angle divided the CFS presentation into two masked areas on each side of fixation. This served the purpose of increasing the independency of the two targets' suppression durations (Maruya, & Blake, 2009).

The colored stimuli used in the memory task as well as the colored target stimuli used in the b-CFS task consisted of circles with a diameter of 1.08° of visual angle. The colors used for the b-CFS task were comprised of perceptually equiluminant, saturated red, green, and blue, and purple, for which the red and blue cannons of the CRT monitor contributed equally. A saturated blue color was used to obtain perceptual equiluminance of the saturated red, green and purple colors for each subject by means of heterochromatic flicker photometry (Kaiser, & Comerford, 1975; Wagner, & Boynton, 1972). The target stimuli were presented at a fixed eccentricity of 1.62° of visual angle from fixation, at a random angular position on the left and right arcs of an imaginary circle, delimited by its main diagonals. The 20 colors used for the memory task consisted of five non-saturated hues, drawn from each of four different color categories (red, green, blue, and purple). The luminance of all hues was physically matched by means of a PR-650 SpectraScan colorimeter/telephotometer (Photo Research, Inc.). An overview of all CIE-color values of the stimuli used in this experiment is provided in Table 1.

Procedure

Participants first performed a calibration task, which was aimed at (1) allowing participants to get acquainted with the b-CFS task and (2) adjusting the intensity of the target stimuli such as to promote independent breakthrough of the two target stimuli. In case the participants experienced simultaneous breakthrough of the two targets, the experimenter reduced the opacity of the target stimuli from 100% (full opacity) to a minimum of 60% in discrete steps. In case the participant still experienced simultaneous breakthrough of the two target stimuli at this stage, the participants (N=2) could not take part in the Experiment. Next, participants performed 128 experimental trials and 10 catch trials, separated into 8 blocks. A schematic depiction of a trial is presented in Figure 1. Each trial started with two successive colored stimuli, drawn from different color categories, followed by a retrocue ("1" or "2") indicating whether the first or second color should be memorized for subsequent recall. During the retention interval, high contrast patterns started alternating at

10 Hz in one eye (now the dominant eye), temporarily suppressing visual input to the other eye (now the non-dominant eye). After a variable delay ranging between 300 and 600 ms, two target stimuli (left and right of fixation) were presented to the non-dominant eye, ramping up from zero to full intensity in 1000 ms. One of these stimuli was drawn from the same category as either the cued or the uncued color (i.e., related to the memory task), whereas the other stimulus was drawn from neither (i.e., unrelated to the memory task). Participants were instructed to report as fast as possible where (left or right) they first saw a target stimulus appear. After this response (or after 10 seconds had elapsed), two colored stimuli appeared left and right of fixation that were both drawn from the same color category as the cued color. Participants were required to indicate within 4500 ms which of these two stimuli (presented for 1500 ms) was of the exact same color as the cued color. Each trial was separated by an inter trial interval of 1600 ms.

Design and hypotheses

All factors were manipulated within-participants. The experimental design consisted of one factor of interest, Trial Type, which had two levels. In half of the trials, one of the two target stimuli always matched the category of the cued color (i.e., Memorized trial type). In the other half of the trials, one of the two target stimuli always matched the category of the uncued color (i.e., Discarded trial type). The color of the simultaneously presented second target stimulus was always drawn from the color category that was not used on that trial (i.e., the target was unrelated to the memory task). Four other factors were included as factors of no interest. This included the factors Retrocue ("1" or "2"), Eye (target presented to left or right eye), Location (unrelated stimulus presented to the left or right of fixation) and Color (the cued stimulus is drawn from the red, green, blue or purple color category). Together, these factors resulted in 64 unique combinations of conditions, of which the order was randomized. Each unique combination of conditions was presented twice, once in each half of the experiment, resulting in a total of 128 trials. Finally, on each trial, the specific hue of the cued color, the color category and the specific hue of the uncued color, and the color category of the unrelated target stimulus were determined at random, with the only constraint that the cued color, the uncued color, and the color of the unrelated target stimulus were drawn from distinct color categories.

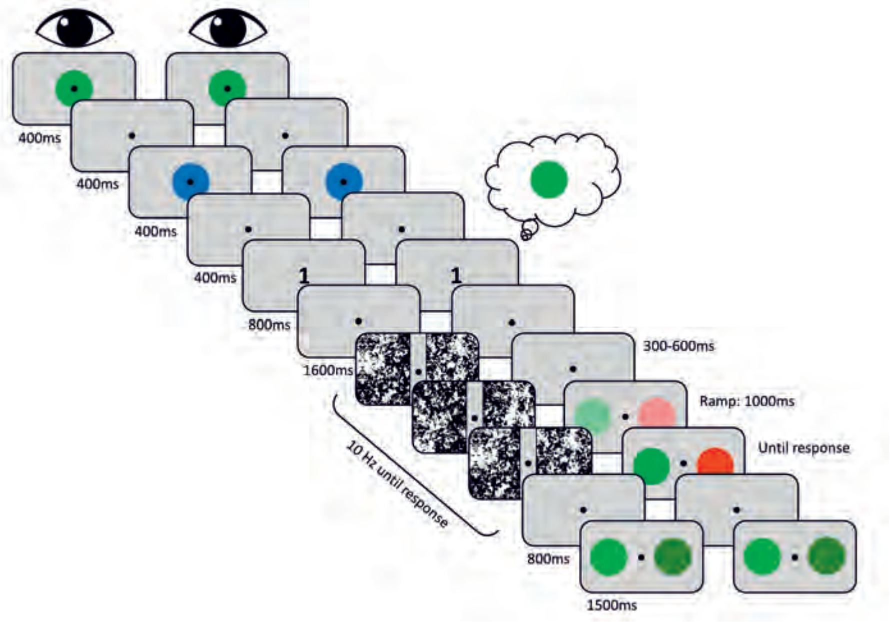

Figure 1. A schematic depiction of a trial. On each trials participants were cued to memorize one of two sequentially presented hues for subsequent recall. During the delay interval, the participants were required to report as fast as possible the location (left or right of fixation) at which a stimulus emerged through the CFS masks. One of these stimuli was always unrelated to the memory task, whereas the other stimulus was related to the memory task (i.e., either matching the cued stimulus on so-called 'Memorized trial types' or the uncued stimulus on so-called 'Discarded' trial types). Note that the specific stimulus properties used in this illustration differ from those of the actual experiment.In addition to the experimental trials described above, 10 catch trials were interspersed throughout the experimental trials (1 or 2 per block). In these catch trials, only one stimulus was presented; five times left of fixation and five times right of fixation.

Based on previous experiments (notably Gayet et al., 2013) we expected that target stimuli matching the category of the cued (i.e., memorized) color would break through suppression faster than target stimuli that mismatch this category. Assuming that this effect is engendered by the memorization rather than the presentation of the cued color, we expected that target stimuli matching the uncued color would not break through suppression faster than target stimuli drawn from an unrelated color category. Therefore we predicted that on trials with a target stimulus matching the category of the cued color (i.e., the Memorized trial type), these target stimuli would be reported (1a) more often and (1b) faster than target stimuli drawn from an unrelated color category. In line with this, we predicted that (2) these differences should not be observed in trials comprising target stimuli matching the category of the uncued color (the Discarded trial type).

Modelling analyses

We hypothesized two distinct patterns of perceptual evidence accumulation that could underlie the preferential detection of targets that match compared to targets that mismatch the concurrent content of VWM. Specifically, the *Pre-activation* hypothesis predicts a lowered effective threshold for targets that are drawn from the same color category as the cued color. The *Feedback* hypothesis, conversely, predicts a faster accumulation of perceptual evidence for targets that are drawn from the same color category as the cued color. To dissociate between these two possibilities, the *Pre-activation* and *Feedback* hypotheses were implemented as two variants of the Linear Ballistic Accumulator model (LBA; Brown & Heathcote, 2008). The LBA model assumes that evidence for each response option (e.g., the target left or right of fixation) gradually accumulates over time, until a preset response threshold is reached. The time required for the first option to reach the threshold constitutes the decision time. This is implemented in the LBA model by a set of parameters that govern the shape of the response time distributions for both response options, as well as the proportion of both response options. These parameters are the response threshold b, the mean rate of accumulation v, and a *non-decision* parameter t_0 representing components of the response time that capture other processes, such as the time required to press the response button. Additionally, variability in behavior is accounted for by allowing the accumulation rate to vary from trial to trial (represented by a normal distribution with mean v and standard deviation s), and by allowing the response threshold to vary (represented by a uniform distribution with range $[b\text{-}a, b]$, which is equivalent to varying the point at which the accumulation starts within the range [0, a]). Both the response threshold and the accumulation rate may differ across accumulators, yielding differences in response times and response choices.

The *Feedback* hypothesis is implemented as an accumulation rate difference between the Memorized and Discarded trial types. That is, all parameters are assumed to be equal for both conditions, but we included two accumulation rates ($v_{memorized}$ and $v_{discarded}$) to account for the hypothesized difference in gain. For scaling purposes, the accumulation rate for responses to the unrelated targets is set at $1\text{-}v_{memorized}$ and $1\text{-}v_{discarded}$, for each trial type respectively.

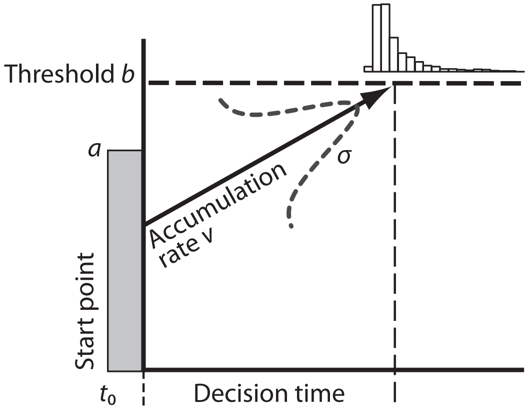

Figure 2. A schematic depiction of a single accumulator and its parameters in the LBA model. On each trial, participants provided a behavioral response that consisted of a response time and a response option (e.g., the target left or right of fixation). The LBA model assumes that evidence for each response option gradually accumulates over time, following an accumulation rate *v*, until a preset response threshold *b* is reached. The time required for the first option to reach the threshold constitutes the decision time. Here, we investigate whether the observed pattern of data is best described by either a model in which all parameters are fixed except for the threshold *b*, or a model in which all parameters are fixed except for the accumulation rate *v*.

The *Pre-activation* hypothesis is implemented by including two different response thresholds. Specifically, we estimated the overall threshold *b* across all trial types and responses, and additionally estimated the deviation of each accumulator from this threshold (Δb), in each trial type. This means that $b_{memorized}=b-\Delta b_{memorized,}$ and $b_{unrelated}=b+\Delta b_{memorized}$, with $b_{memorized}$ indicating the threshold for the target that matches the cued (i.e., memorized) color category. Similarly, in the 'Discarded' trial type, $b_{discarded}=b-\Delta b_{discarded}$, and $b_{unrelated}=b+\Delta b_{discarded}$, with $b_{discarded}$ indicating the threshold for the target that matches the color category of the uncued (i.e., discarded) stimulus. Again, for scaling purposes, the drift rate for the unrelated targets was set at *1-v*, which in this model is the same for both conditions. The present approach is similar to that of earlier studies, in which the influence of a priori information on perceptual decision making was attributed to an initial bias, on the basis of an LBA model in which only the threshold was allowed to vary (Forstmann, Brown, Dutilh, Neumann, & Wagenmakers, 2010; van Maanen, Forstmann, Keuken, Wagenmakers, & Heathcote, in press).

Because changes in the accumulation rate and threshold parameters (*v* and *b*) have different effects on the shape of the response time distributions and the choice proportions, one of these models will more closely match the observed data after the best fitting parameters are estimated, providing evidence in favor of that model. We estimated parameters by maximizing the summed log likelihood of the ventiles (i.e., the twenty 5% quantiles; Vincent, 1912; Ratcliff, 1979) of the response times, separated for both response options (i.e., target related or unrelated to the memory task), for both the 'Memorized' and 'Discarded' trial types (i.e., related target is drawn from the same color category as the cued stimulus or uncued stimulus). Maximization was done using particle swarm optimization (Clerc, 2010), to avoid local optima in fitness landscape. The LBA model implementations of the *Pre-activation* and *Feedback* hypothesis were subsequently compared by their likelihood ratio, which indicates how likely each model is relative to the other model given the data (Jeffreys, 1961). This comparison can be made directly, as both models comprise the same number of free parameters, such that no correction has to be applied for model flexibility. Finally, the model that was best at describing the observed data was compared to a control model, in which all parameters were fixed. In this case, the likelihood ratio was computed by means of a Bayesian Information Criterion (BIC) score, which allows to select the best model amongst a finite number of models (in this case three), while correcting for the number of free parameters in each model.

Results and discussion

Preliminary analyses

Three participants were excluded from further analyses for being incorrect in at least 30% of the catch trials. These participants thus repeatedly reported that a stimulus appeared first at a location where no stimulus was presented. Such responses indicate that participants either misattributed features of the CFS masks as being a target stimulus, or reported a target location prior to seeing the target. We removed all trials that yielded response times under 350 ms (1.23% of all trials) as well as those that were not responded to within the 10 seconds time window (0.04% of all trials). The remaining 19 participants were 95.3 % (*SD* = 6.8) accurate on the catch trials. Participants had an average accuracy of 73.2 % (*SD* = 6.3) on the working memory recall task, and all included participants performed above chance level, as determined by a binomial test against chance level (all *p*'s < 0.0145).

Behavioral analyses

Our first aim was to replicate the finding that visual input matching the content of VWM breaks through suppression faster than mismatching visual input. For this purpose, we computed the fraction of trials in which participants responded to the side of fixation containing the target stimulus that matched the color category of the cued stimulus (Memorized trial type). We first conducted planned one-sample t-tests to determine whether this fraction differed from chance level (i.e., 0.5, reflecting that participants responded to both target stimuli equally often). This revealed that target stimuli matching the category of the cued color were reported 7.5% more often (*M* = 34.4 out of 32 trials, *SD* = 4.1) than expected by chance, *t*(18) = 3.0629, *p* = 0.0067, *d* = 0.702 (significant at a Bonferroni corrected α-level of 0.025). In contrast, the same analysis for the Discarded trial type revealed that the number of responses to target stimuli matching the category of the uncued color was at chance level (-4.1%, *M* = 30.7 out of 32 trials, *SD* = 4.0), *t*(18) = -0.8652, *p* = 0.3983, *d* = -0.208 (not significant at a Bonferroni corrected α-level of 0.025). Additionally, a paired-samples t-test showed that these fractions differed between the Memorized trial type, in which target stimuli matched the category of the cued color, and the Discarded trial type, in which target stimuli matched the category of the uncued color, *t*(18) = 3.1397, *p* = 0.00567, *d* = 0.575. For paired comparisons, effect sizes were computed by using Cohen's d, following Morris and DeShon's (2002) equation 8, which corrects for dependence between means. The results of these analyses are depicted in Figure 3A.

Next, we aimed to assess whether the response speed to target stimuli depended on the response choice (i.e., the relation between the target stimulus and the color categories used in the memory task). For this purpose, we computed median response times for each for the two response options (target is related or unrelated to the memory task) in each of the two trial types (Memorized or Discarded trial types). This showed that, in the Memorized trial type, response speeds to target stimuli matching the category of the cued color (*M* = 1678 ms, *SD* = 645) were 170 ms faster than responses to target stimuli drawn from an unrelated color category (*M* = 1838 ms, *SD* = 841), *t*(18) = 2.4852, *p* = 0.0230, *d* = 0.816 (significant at a Bonferroni corrected α-level of 0.025). In the Discarded trial type, however, there was no difference in response speed between target stimuli matching the category of the uncued color (*M* = 1691, *SD* = 607) and target stimuli drawn from an unrelated color category (*M* = 1704, *SD* = 701), *t*(18) = 0.1267, *p* = 0.9006, *d* = 0.024. The effect of the memory task on response speeds (i.e., in the Memorized trial type) was correlated with participants' average response speed, $R = 0.72876$, $p = 0.0004$. This means that part of the variability in the response speed difference is accounted for by between-subject variability in absolute response speed (i.e., sensitivity to CFS). To remove this between-subject variability, we normalized the response times within each trial types as follows:

$$\mathrm{RT} = \frac{RT_{related} - RT_{unrelated}}{\frac{1}{2}(RT_{related} + RT_{unrelated})}$$

A similar approach has been used by Stein (2012). The resulting measure reflects the difference in response time brought about by the manipulation (i.e., the relation between the color of the target and the color category of the cued stimulus in the memory task). After normalization, one sample t-tests revealed that individual response times to target stimuli matching the color of the cued stimulus were 7.0% ($SD = 11.2$) faster than to target stimuli of an unrelated color category, $t(18) = 2.6714$, $p = 0.0156$, $d = 0.613$ (significant at a Bonferroni corrected α-level of 0.025). This reduction of response times was not apparent in the Discarded trial type ($M = 0.2\%$, $SD = 14.6$), $t(18) = 0.0614$, $p = 0.9517$, $d = 0.014$ (not significant at a Bonferroni corrected α-level of 0.025). The difference in response time between responses to the related and unrelated targets did, however, not significantly differ between trial types, as shown by a paired-samples t-test, $t(18) = 1.555$, $p = 0.1374$, $d = 0.339$. Possibly, the effect of our experimental manipulation on response times was dampened (compared to Gayet et al., 2013) due to the fact that for Memorized trials in which participants reported the unrelated target to appear first, the suppression duration was co-determined by the related target, presented to the same eye (e.g., Ooi & He, 1999).

Taken together, we replicated earlier findings (e.g., Gayet et al., 2013) by showing preferential access to awareness for visual input matching the content of VWM. The present paradigm allowed to corroborate these findings with a modified experimental paradigm, in which two stimuli directly (i.e., simultaneously) compete for access to awareness. This paradigm was specifically aimed at examining the effect of VWM on access to awareness including measures of response choice as well as response speed. In the next paragraph, these two behavioral measures will allow us to discern whether the effect of VWM on access to awareness is engendered by (1) a lowered effective threshold or (2) a faster accumulation of evidence for detecting visual input that matches rather than mismatches the content of VWM. This would provide support for either (1) the *Pre-activation* hypothesis or (2) the *Feedback* hypothesis, respectively.

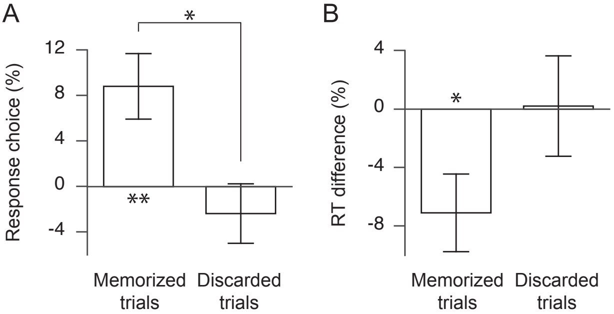

Figure 3. Behavioral results of the main experiment. On each trial, one target (i.e., response option) was unrelated to the memory task, whereas the other one was related to the memory task. The target that was related to the memory task was related to the cued stimulus on 'Memorized' trials, and to the uncued stimulus on 'Discarded' trials. Panel A depicts the percentage of trials in which participants reported seeing the target appear first that was related to the memory task. A positive value indicates that on the majority of trials the related target emerged from CFS before the unrelated target did. Panel B depicts the normalized difference in response time to targets that were related to the memory task compared to unrelated targets. A negative value reflects a faster response time to the related target. Error bars represent the standard error of the mean. *$p < 0.05$, **$p < 0.01$

Modelling analyses

We compared the implementations of the *Pre-activation* and *Feedback* hypothesis as variants of the LBA model. A model comparison supported the *Pre-activation* model over the *Feedback* model, as indicated by a likelihood ratio of 35,186. Specifically, a model with the threshold *b* as the only free parameter is 35,186 more likely to have generated the data than a model with the accumulation rate *v* as the only free parameter. Additionally, the *Pre-Activation* model was 10,459,802 times more likely to have generated the data than a control model in which no parameter was allowed to vary. Figure 4 summarizes that the *Pre-activation* model indeed captures important properties of the data. Specifically, Figure 4A shows that the *Pre-activation* model roughly captures the participants' response choices (i.e., the fraction of trials in which participants responded to the target that was related to the memory task). Importantly, the model captures the shape of the response time distributions (Figure 4B), as evidenced by the close correspondence between model predictions and data for all deciles (i.e., 10% quantiles) of the response time distribution.

The parameters of the *Pre-activation* model support the perspective that the response threshold is lower for targets that match the content of VWM (Figure 4C). One-sample t-tests against no difference show that there was a significant difference in the response threshold parameter *b* between trials in which targets were either unrelated to the memory task and trials in which target were related to the cued stimulus (i.e., in the Memorized trial type; $M_{\text{memorized}} = 82$, $SD_{\text{memorized}} = 101$), $t(18) = 3.56$, $p = 0.0022$, but not for trials in which targets were either unrelated to the memory task or were related to the uncued stimulus (i.e., in the Discarded trial type; $M_{\text{discarded}} = 9.8$, $SD_{\text{discarded}} = 109$), $t < 1$. In addition, a paired-samples t-test showed that this effect of response threshold differed between the Memorized and Discarded trial types, $t(18) = 2.51$, $p = 0.021$. This shows that the different patterns of findings in the Memorized trial types and the Discarded trial types, which only differ in the contingency between the content of VWM and the target stimuli, is accounted for by a lowered threshold for stimuli matching the content of VWM. This is the pattern of results that would be expected if the content of VWM pre-activated neural populations that process concurrently presented visual input.

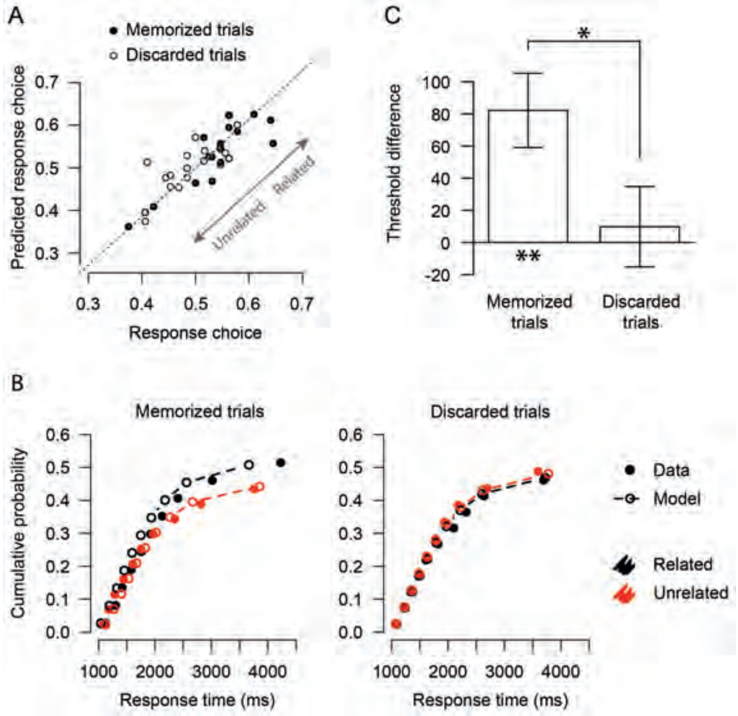

Figure 4. Model predictions and parameter estimates of the *Pre-activation* model for the main experiment. On each trial, one target (i.e., response option) was unrelated to the memory task, whereas the other one was related to the memory task. The target that was related to the memory task was related to the cued stimulus on 'Memorized' trials, and to the uncued stimulus on 'Discarded' trials. Panel A depicts the observed and predicted proportion of each response option. The grey dotted lines connect pairs of data points from the same participant. Panel B depicts the defective cumulative density of observed data and model predictions, for 10 response time quantiles (i.e., ventiles binned into deciles for illustrative purposes) averaged over participants. The two trial types are shown in different graphs. Panel C depicts the mean parameter estimates of the Pre-activation model. The bars indicate the difference in threshold parameters *b* for the related (to cued or uncued) target responses and the unrelated target responses. $*p < 0.05$, $**p < 0.01$

Generalizing to the traditional b-CFS paradigm

Introduction and methods

Using a modified version of the b-CFS paradigm, we replicated earlier findings that visual input matching the content of VWM is released from interocular suppression faster than visual input that mismatches the content of VWM. This modified paradigm allowed us to identify, that the reduced suppression durations are potentiated by a decreased threshold for stimuli matching the content of visual working memory, akin to an initial bias towards VWM matching stimuli. To ensure that the present results generalize to the typical b-CFS paradigm in which one target (i.e., response option) is provided, we re-analyzed the data of Gayet et al.'s (2013) Experiment 4. Details of the participant groups from which the data is retrieved for the supplemental modelling analyses can be found in Gayet et al. (2013). In this experiment (N = 15) only one target was presented simultaneously. The target stimulus either matched the color category of the cued stimulus (in Memorized trial types) or that of the uncued stimulus (in Discarded trial types), in 24 trials each. Trials in which the target was drawn from an unrelated color category were not included in the present analysis. This experiment does not use a binary forced choice paradigm, and therefore an LBA model is inappropriate. Instead, we fitted the response times of the correct responses (97.9%) using a one-sided accumulator model (the shifted Wald model, Anders et al., in press). This model is suitable for decomposing response time distributions into accumulation rate v, response threshold b, and non-decision time t_0 in case of only one (meaningful) response. The *Feedback* and *Pre-activation* hypotheses were again implemented by allowing separate accumulation rates and thresholds respectively, for the Memorized and Discarded trial types. Again, we compared model implementations of the *Feedback* and *Pre-activation* hypotheses using their likelihood ratio, after optimizing the parameters using particle swarm optimization.

Results

Model comparisons of the data of Experiment 4 of Gayet et al., (2013) were in line with the findings of the main experiment reported above. The likelihood ratio of the different models revealed that a model with the threshold b as the only free parameter (the *Pre-activation* model) is 19 times more likely to have generated the data than a model with the accumulation rate v as the only free parameter (the *Feedback* model). Also, the *Pre-activation* model explained the RT distribution data very well (Figure 5A), and, similarly to the main experiment, we again observed that the threshold parameter *b* is lower in Memorized than in Discarded trials (Figure 5B, $M_{\text{Memorized}} = 52$, $SD_{\text{Memorized}} = 33$, $M_{\text{Discarded}} = 61$, $SD_{\text{Discarded}} = 35$), as revealed by a paired-samples t-test, $t(14) = 3.98$, $p = 0.0014$.

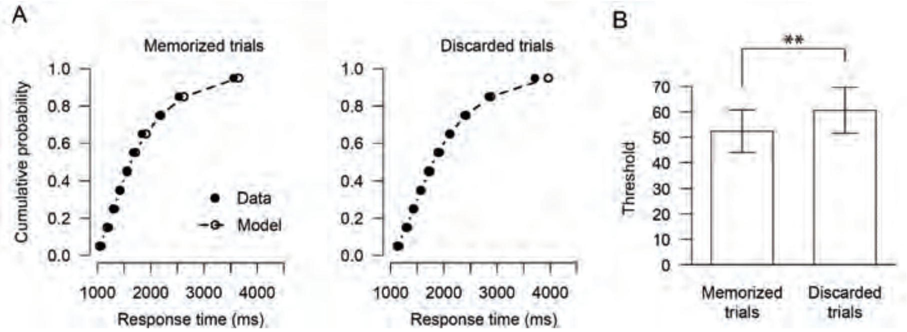

Figure 5. Model predictions and parameter estimates of the *Pre-activation* model, for a dataset retrieved from Gayet et al. (2013; Experiment 4). On each trial, one target was presented that was either related to the cued stimulus (Memorized trial type) or uncued stimulus (Discarded trial type) from the memory task. Panel A depicts the defective cumulative density functions, illustrating the probability of observing correct and incorrect responses for the data and model predictions, in 10 response time quantiles (ventiles binned into deciles for illustrative purposes), averaged over participants. The two trial types are shown in different graphs. Panel B depicts the mean parameter estimates of the Pre-activation model. The bars indicate the mean parameter estimates (and error bars) for the threshold *b* for each of the two trial types. ** $p < 0.01$

Dissociating perceptual decision threshold from response criterion

Introduction and methods

The current findings demonstrate that the faster access to awareness of target stimuli matching the content of VWM is potentiated by a lowering of the effective threshold (i.e., the threshold *b*) for matching compared to mismatching target stimuli. From this data, however, it is unclear whether the difference in threshold reflects differential processing of a matching target stimulus (1) before or (2) after the target stimulus is released from interocular suppression. Gayet et al.'s (2013) Experiment 2 and 3 demonstrate that when the target stimulus is not interocularly suppressed, response times do not differ between targets that match and mismatch the content of VWM. In these so-called monocular control experiments, the target stimuli are presented to the same eye as the CFS masks, and their opacity is gradually ramped up. Because the target stimuli are presented to the eye that already dominates perception, these experiments only capture differences in response time that are initiated after the interocular competition is resolved. As response times did not differ between trial types in these experiments, we concluded that the difference in response time that is obtained in conditions of interocular competition (e.g., the current experiment, and Gayet et al.'s Experiment 1, 4 and 5, 2013) reflects a difference in suppression duration, rather than a difference in response speed to the target stimulus once it is no longer suppressed (for more elaborate discussion on this topic, see Gayet, Paffen, Belopolsky, Theeuwes, & Van der Stigchel, 2016; Gayet et al., 2014b; Stein, Hebart, & Sterzer, 2011). Similarly, we aimed to assert that the lowered effective threshold for stimuli matching the content of VWM reflects a difference in processing before the target stimuli are visible. For this purpose, we reanalyzed the data of Gayet et al.'s (2013) Experiment 3, in which the targets were not interocularly suppressed, by fitting it using a one-sided accumulator model (the shifted Wald model, Anders et al., in press). In this experiment (N = 11) targets either matched the color category of the memorized stimulus or not (36 trials per trial type). Only correct trials (99.4%) were included in the analyses.

Results

The *Pre-activation* model, in which only the threshold *b* was allowed to vary between the two trial types, was 9 times better in describing the observed data compared to a model in which only the accumulation rate *v* was allowed to vary between the two trial types (i.e., the *Feedback* model). In this control experiment, however, the threshold *b* did not significantly differ between Memorized trials and Unrelated trials, as shown by a paired-samples t-test, $t(10) = 1.5355$, $p = 0.1557$ (see Figure 6). We advocate caution in interpreting these findings, however, as a Bayesian paired-samples t-test revealed that the null hypothesis (both threshold values are drawn from the same distribution) is only 1.3 times more likely than the alternative hypothesis (both threshold values are drawn from different distributions). This lack of evidence for either hypothesis is indicative of having insufficient power for convincingly demonstrating the absence of a difference in the threshold parameter *b* between Memorized and Unrelated trials. Nonetheless, the finding that nor the difference in response times nor the difference in the threshold parameter *b* are significantly reduced for Memorized trials in this control experiment suggests that when targets are presented to the eye that is already dominant, the effects of VWM reported in the previous experiments are, at least partly, abolished. As such, the reduction in threshold for targets matching the content of VWM reflects a processing difference that originates before the target is released from suppression.

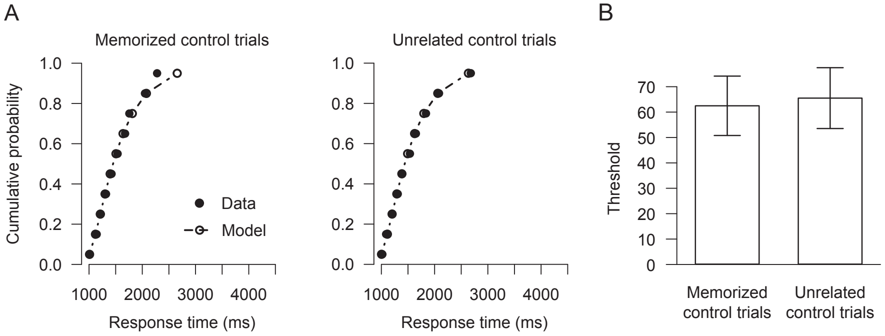

Figure 6. Model predictions and parameter estimates of the *Pre-activation* model, for a dataset retrieved from Gayet et al. (2013; Experiment 3). On each trial, one target was presented that was either drawn from the same (Memorized trial) or a different (Unrelated trial) color category than the stimulus that participants were instructed to memorize. Importantly, target stimuli in this experiment were presented to the same eye as the masks, such that no interocular suppression was elicited. Panel A depicts the defective cumulative density of observed data and model predictions, for 10 response time quantiles (i.e., ventiles binned into deciles for illustrative purposes) averaged over participants. The two trial types are shown in different graphs. Panel B depicts the mean parameter estimates of the Pre-activation model. The bars indicate the mean parameter estimates (and standard errors) for the threshold *b* for each of the two trial types.

General discussion

The content of visual working memory (VWM) affects the processing of concurrent visual input. It has been demonstrated across multiple paradigms that visual input matching the content of VWM elicits stronger behavioral responses than visual input that mismatches the content of visual working memory (e.g., Olivers, Meijer, & Theeuwes, 2006; Pan, Cheng, & Luo, 2012; Gayet, Brascamp, Van der Stigchel, & Paffen, 2015). In order to elucidate the perceptual processes leading up to this phenomenon, we conducted a modified version of the b-CFS paradigm, in which stimuli matching the content of VWM have been shown to gain accelerated access to awareness (van Moorselaar et al., 2015; Pan et al., 2014; Gayet et al., 2013). This modified paradigm allowed us to implement the predictions of two hypotheses as sequential sampling models, and compared which of these quantitatively best explained the data. The *Pre-activation* hypothesis, which predicts an initial bias towards VWM matching visual input, was implemented as a model in which the threshold was allowed to vary between matching and mismatching visual input. The *Feedback* hypothesis, which predicts a gradual enhancement of VWM matching visual input, was implemented as a model in which the slope of perceptual evidence accumulation was allowed to vary between matching and mismatching visual input. Model comparisons revealed that the data was best explained by a decrease in threshold for visual input that matches the content of VWM, thereby supporting the *Pre-activation* hypothesis. In light of this hypothesis, representations elicited by visual input rely on the same neural substrate as representations that are concurrently maintained in VWM. Consequently, the effective threshold for visual input to reach visual awareness is reduced if the visual input matches the content of VWM.

Whereas the execution of the VWM task is operated in a top-down manner, its effect on concurrent perception is not under volitional control (Gayet et al., 2015). In line with this, the Pre-activation hypothesis, provides a bottom-up account of the modulation of interocularly suppressed information by the content of VWM. That is, the modulation of visual processing areas, engendered by VWM maintenance, precedes the presentation of the stimulus. In contrast, one of the critical prerequisites for the alternative hypothesis, the *Feedback* hypothesis, is that the interocularly suppressed information should at least partly transpire beyond early visual processing areas. This goes against a broad range of findings showing that interocularly suppressed information is processed only to a limited extent. In general, the extent to which neural activity reflects interocularly suppressed stimulation decreases gradually when climbing up the visual hierarchy (Leopold, & Logothetis, 1999). Most cells in early visual areas (80% in V1/V2 and 60% in V4/V5) respond to stimulation of either eye, irrespective of which eye's input dominates perception (Logothetis, 1998). Higher processing areas, such as IT, LOC, FFA, and PPA, however, follow mostly (but not exclusively; Fang and He, 2005; Jiang and He, 2006; Sterzer, Haynes, & Rees, 2008; Sterzer, et al., 2014) the dominant percept (Tong, Nakayama, Vaughan, & Kanwisher, 1998). Similarly, behavioral studies using adaptation after-effects demonstrate that adaptation resulting from so-called low level stages of visual processing (e.g., contrast, motion direction, phase, etc.) is greatly reduced under CFS, while adaptation effects resulting from higher level stages in visual processing (e.g., complex motion, gaze direction, face gender, etc.) is abolished under CFS (for a review, see Yang, Brascamp, Kang, & Blake, 2014). Similar conclusions have been drawn for binocular rivalry (for a review, see Lin & He, 2009), which is believed to elicit weaker interocular suppression than CFS (Tsuchiya, Koch, Gilroy, & Blake, 2006). Together, this implies that CFS disrupts higher level visual processing at least to a greater extent than lower level visual processing. Conversely, any modulation of interocular suppression is more likely to impinge early on upon the processing of the suppressed visual input. The present findings provide an account for the influence of working memory on suppression duration, for which low level visual processing of interocularly suppressed information suffices.

There are, however, two possible issues with the view that pre-activation of early visual areas accounts for the preferential access to awareness of stimuli that match the content of VWM. First, the effect is categorical. That is, the target stimuli that are presented during the delay interval are never identical to those that are maintained in VWM. Rather, they are drawn from the same category (e.g., in the b-CFS paradigm, van Moorselaar et al., 2015; Gayet et al., 2013; in the binocular rivalry paradigm, Gayet et al., 2015; in attentional capture paradigms Olivers et al., 2006; van Moorselaar, Theeuwes, & Olivers, 2014). Considering the distinction between higher and lower visual areas, as described in the previous paragraph, categorical representations of stimulus features would be expected to rely on higher level visual processing areas. One explanation is that stimuli that are sufficiently similar will elicit responses in overlapping neural populations. In this view, the enhanced response to a target stimulus that matches a concurrently memorized stimulus is not caused by its inclusion in a feature category, in which case one would expect the enhancement to be all or none (i.e., the target stimulus is either drawn from the same category as the memorized stimulus or not). Rather, the enhanced response might rely on the perceptual overlap between the target stimulus and the memorized stimulus, in which case one would expect the magnitude of response enhancement to correlate with the similarity between the memorized stimulus and the target stimulus. This latter view finds support in a recent study in which participants were primed with an auditory cue (the word "square" or "circle") prior to performing a breaking continuous flash suppression task with so-called squircle stimuli (Lupyan & Ward, 2012). These are stimuli that range on a continuum that is delimited by a perfect square and a perfect circle. Shapes that matched the cue (e.g., a square after the word "square") yielded the shortest suppression durations, and suppression durations increased linearly with decreasing resemblance between the target stimulus and the cued shape. In line with this, the present findings can be explained without resorting to categorical representations of colors. Rather, the overlap in color-space between the color of the target stimulus and the color of the memorized stimulus can account for the facilitatory effect that we observe.

The second issue, is that the target stimulus was presented at (1) an unpredictable location, and (2) was never presented at the same retinal location as the to-be-memorized stimuli. In fact in some studies using this paradigm (e.g., the current study and Gayet et al., 2013) the stimuli were separated by a degree of visual angle or more, which is more than the receptive field size of foveal V1-V2 cells (e.g., Harvey & Dumoulin, 2011). Also, studies using related paradigms have shown various forms of enhanced processing of stimuli that match the content of VWM at even larger retinotopic distances between the to-be-memorized stimulus and the target stimulus (e.g., Olivers et al., 2006; van Moorselaar et al., 2014). In the particular case of color stimuli, one could argue that participants memorize the color per se, rather than the colored stimulus. That is, the color representation maintained in VWM is not necessarily spatially delimited, but rather virtually covers the entire visual field. Considering that the color of a presented stimulus can be decoded from early visual areas, including V1 (Brouwer & Heeger, 2009), maintaining a color in VWM could potentially bias interocular competition of colored stimuli in early visual areas. In the case of face stimuli (Pan et al., 2014) or shape stimuli (Chapter 7; Lupyan & Ward, 2012), however, this explanation is rather unlikely. While one could imagine retaining a color in VWM, across the visual field, this is not probable for stimuli that are defined by their spatial configuration, such as faces. For these type of stimuli, the most likely explanation is that the response enhancements of matching stimuli originate in higher visual areas with receptive field sizes that encompass the retinal distance between the target stimulus and the to-be-memorized stimulus. It is important to emphasize that while the present finding (i.e., the difference in effective threshold rather than rate of evidence accumulation) support a model in which the processing of the target stimulus and of the memorized stimulus rely on a shared neural substrate, it is not informative as to which processing areas this entails. It has been argued that the biasing of interocular competition finds its origin at different levels of the visual processing hierarchy (e.g., Blake, & Logothetis, 2002). As such, it remains unclear at what stage of the visual processing hierarchy the contents of VWM impacts the processing of concurrent matching visual input. Possibly, the contents of VWM affects the processing of visual input in later stages of the visual processing hierarchy.

We conclude from the present findings, using a VWM manipulation in a b-CFS paradigm, that visual representations elicited by VWM and visual representations engendered by retinal input draw upon the same neural substrate. This raises an important question regarding the generalizability of our findings. Here, we will make the case that the present findings are likely to generalize to other paradigms in which visual input is modulated by the content of visual working memory. Next, we will make the case that the present findings are unlikely to generalize to other experimental manipulations within the b-CFS paradigm.

The *Pre-activation* hypothesis, in contrast to the *Feedback* hypothesis, entails that the difference between the processing of visual input that either matches or mismatches the content of VWM is initiated before the visual input is presented. This follows from the finding that it was not the accumulation rate (i.e., processing efficiency) but the threshold (i.e., initial bias) that was affected by the memory task. Thus, visual input that matches the content of VWM is enhanced in a bottom-up manner, as a result of the elevated activation levels that are already in place. Considering that this processing bias precedes the b-CFS task, the present findings are not necessarily specific to b-CFS paradigms, but allow for explaining a lot of related phenomena in which facilitatory effects are observed for visual input that matches rather than mismatches the content of VWM. This includes the findings that stimuli matching the content of VWM attract attention (Soto, Heinke, Humphreys, & Blanco, 2005; Olivers, et al., 2006; van Moorselaar, Theeuwes, & Olivers, 2014) and eye movements (Hollingworth & Luck, 2009; Hollingworth, Matsukura, & Luck, 2013; Schneegans, Spencer, Schöner, Hwang, & Hollingworth, 2014; Silvis & Van der Stigchel, 2014) in search tasks, elicit more pronounced behavioral priming effects under backwards masking (Pan et al., 2012), and dominate bistable perception (Gayet, et al., 2015; Scocchia, Valsecchi, Gegenfurtner, & Triesch, 2013; but, see Scocchia, Valsecchi, Gegenfurtner, & Triesch, 2014) compared to stimuli that mismatch the contents of VWM. All these findings could potentially be accounted for by the *Pre-activation* hypothesis that we present here.

The present data revealed that VWM modulation of suppression durations in a b-CFS paradigm is driven by a reduction in effective threshold for matching visual input. This does not (necessarily) imply that other modulations of suppression durations in b-CFS paradigms are also driven by a difference in threshold between experimental conditions. For instance, the well-established finding that upright faces break through interocular suppression faster than inverted faces (e.g., Jiang, Costello, & He, 2007; Stein et al., 2011) might be driven by a higher rate of perceptual evidence accumulation (reflecting higher processing efficiency for more familiar stimuli) rather than a lower threshold (reflecting an initial bias towards any of the stimulus conditions). In contrast, differences in suppression durations between stimuli differing in low-level visual characteristics, such as luminance (Tsuchiya, & Koch, 2005) or spatial frequency (Yang, & Blake, 2012), might provide an initial difference in competitive strength, leading to a difference in threshold. To our knowledge, we present the first implementation of sequential sampling models in a b-CFS paradigm. Recent findings using the b-CFS paradigm ignited the discussion on the extent to which interocularly suppressed stimuli can be processed (for the full scope of the discussion, see Hesselmann, & Moors, 2015; Gayet et al., 2014b). We believe that sequential sampling models can be employed to gain more insights in this discussion, by modelling the behavioral data of different types of manipulations used in b-CFS experiments (including monocular control conditions). The shifted Wald model (Anders et al., in press), in particular, allows for decomposing response time distributions into an accumulation rate and a response threshold in case of only one (meaningful) response. In the present study, we provided the first application of the shifted Wald model on an existing dataset. Importantly, fitting the one-sided shifted Wald accumulator to data from the original experiment lead to a qualitatively similar patterns of results as fitting the two sided accumulators on the data obtained with the modified paradigm. As such, the shifted Wald model allows for gaining novel insights from existing b-CFS datasets, without resorting to collecting new data in a modified paradigm (tailored to the typical two sided accumulator models). Our understanding on the scopes and limits of non-conscious perception, has been clouded by the discrepancy in recent CFS findings (e.g., Gayet et al., 2014b; Hesselmann, & Moors, 2015). We make the case that the field could hugely benefit from such methods that allow for decomposing the perceptual processes that dictated what part of our visual world gains prioritized access to awareness.

Conclusion

Visual input that matches the content of VWM is released from interocular suppression earlier in time than mismatching visual input. Thus far the perceptual processes underlying this facilitatory effect were unknown. By using sequential sampling models, we were able to gain insights in the perceptual processes leading up to the moment in time at which interocularly suppressed visual input was released from suppression. Model comparisons revealed that the threshold for visual input to break through suppression is lower for visual input that matches compared to visual input that mismatches the content of VWM. This is in line with predictions of our *Pre-activation* hypothesis, which entails that the effective threshold for visual input to reach awareness is lowered when activity levels of the neural populations processing the visual input are elevated by the concurrent maintenance of a similar representation in VWM. The present finding has two implications. Firstly, it implies that visual representations that are elicited by either visual input or engendered by VWM maintenance have a shared neural substrate. Secondly, it implies that VWM impacts the processing of concurrent visual input from the onset. As such, it is unlikely that the present findings are specific to the breaking continuous flash suppression paradigm. Rather, this pre-activation account might underlie a plethora of findings demonstrating enhanced processing of visual input that matches the content of VWM.

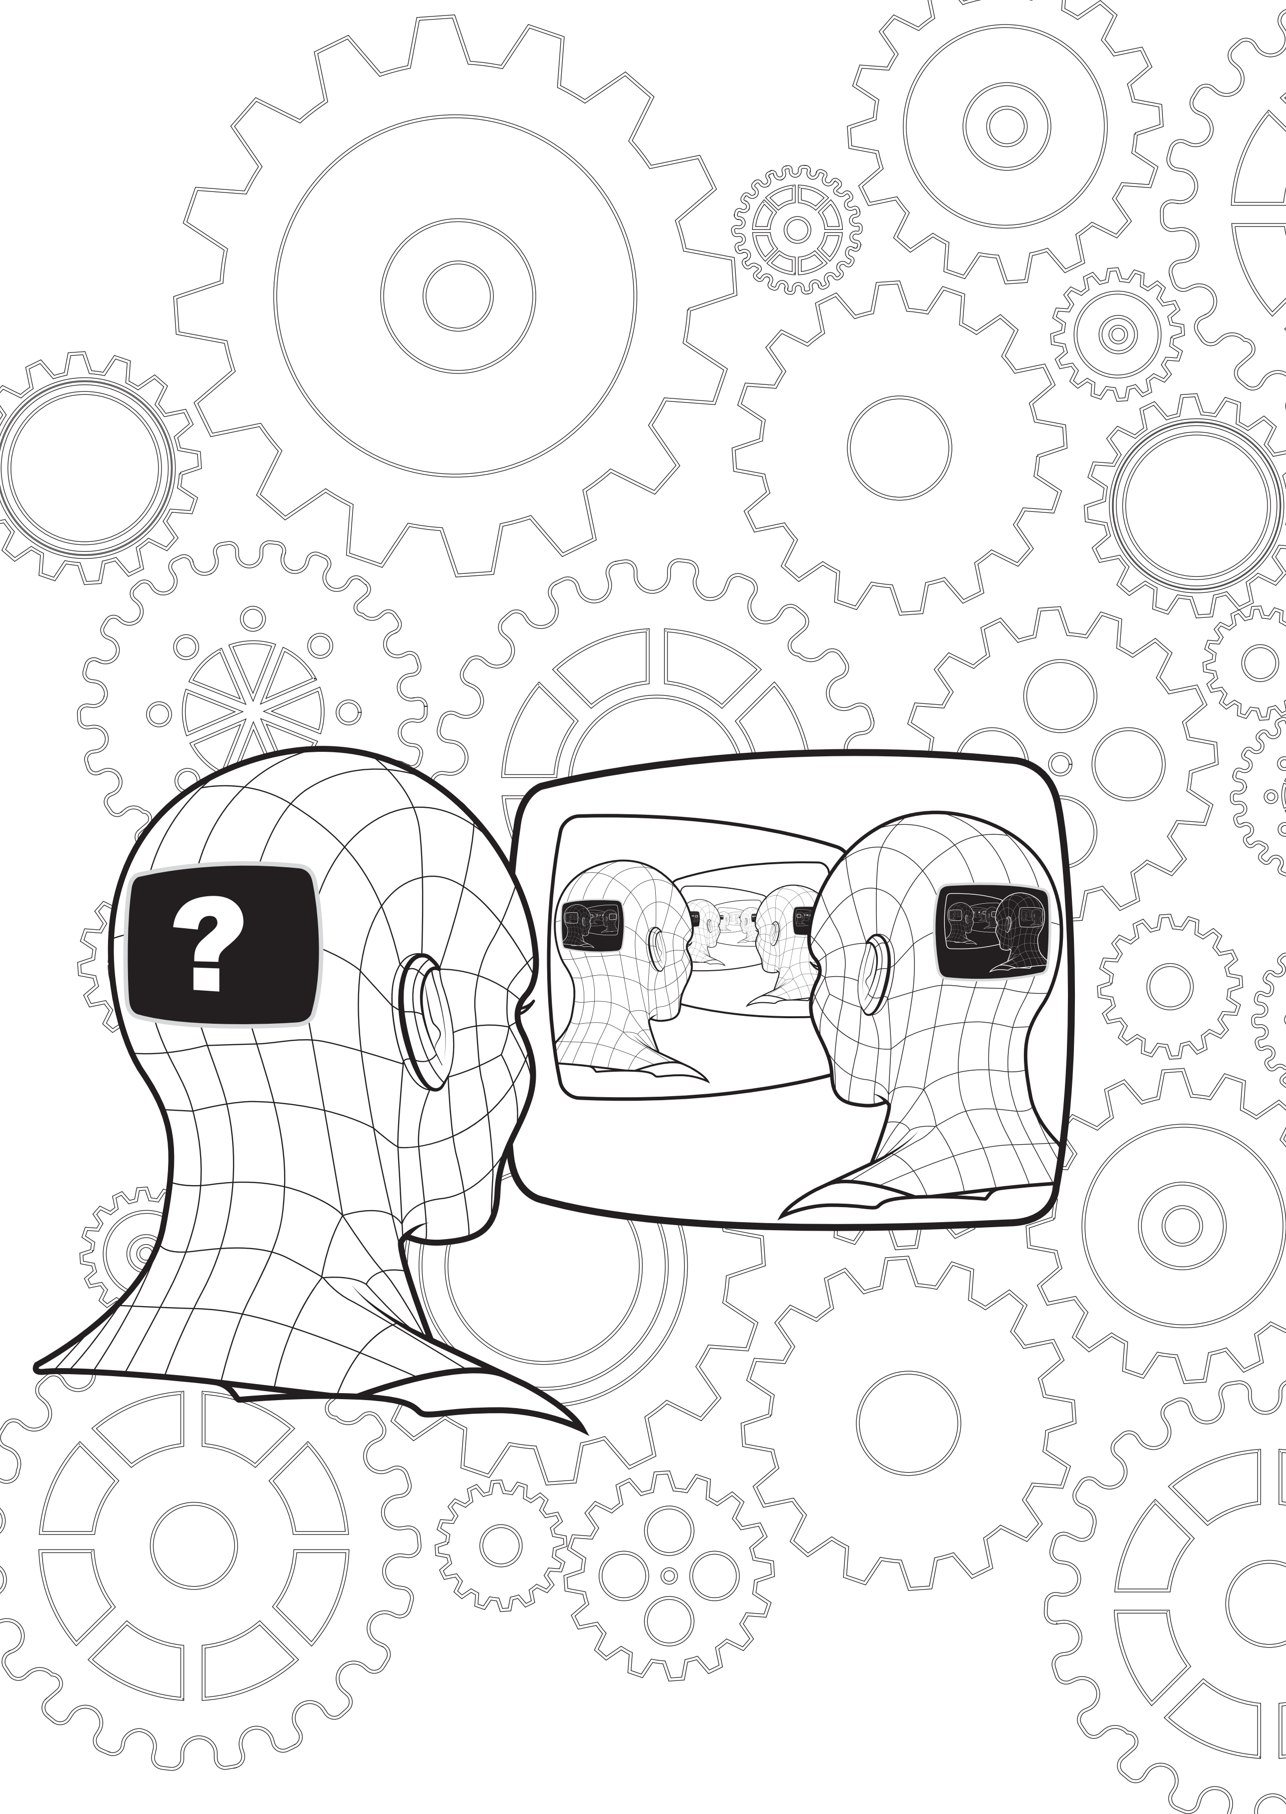

Chapter 7

Visual working memory modulates visual input in high-level visual areas

Parts of this chapter are under preparation for separate publication as:

Gayet, S., Guggenmos, M., Paffen, C. L. E., Van der Stigchel, S., Christophel, T. B., Haynes, J. D., & Sterzer, P. (in preparation). Increased BOLD response and higher decoding accuracy for visual input that matches the content of visual working memory.

Gayet, S., Guggenmos, M., Paffen, C. L. E., Van der Stigchel, S., Christophel, T. B., Haynes, J. D., & Sterzer, P. (in preparation). The content of visual working memory does not modulate the BOLD response to interocularly suppressed visual input.

Gayet, S., Paffen, C. L. E., Guggenmos, M., Sterzer, P., & Van der Stigchel, S. (in preparation). Pupillary constriction accompanies the synergetic effect of visual working memory on pre-conscious visual input.

SG, PS, CP, & SS designed the study concept. Advice was provided by MG, JH, & TC. SG programmed the experiment and tested the participants, with the assistance of MG, TC, & PS. SG conducted the analyses, in consultation with MG. TC and PS provided advice on the analyses. SG wrote the manuscript. Critical revisions were provided by all co-authors.

Abstract

Visual information that is actively maintained in visual working memory (VWM) has been shown to affect concurrent perception. A likely explanation for this, is that visual representations, elicited by either VWM or by sensory input, draw upon a shared neural substrate. Based on this interpretation, we hypothesized that visual information maintained in VWM would enhance the neural response to concurrent visual input that matches the content of VWM. To test this hypothesis, we conducted two fMRI experiments (N=15, N=20) in which participants performed a delayed match to sample task. Participants were presented with two central shape stimuli, drawn from three shape categories (rectangles, ellipses and triangles), and a retro-cue, indicating which of the two shapes should be memorized for subsequent recall. During the retention interval a shape was presented (the probe) that either matched or mismatched the memorized shape category. This allowed us to measure the neural response to the probe, dependent on its contingency with the content of VWM. The probe was either visible (Experiment 1) or rendered invisible by means of interocular suppression (Experiment 2). The results of Experiment 1 revealed that the probe elicited a stronger BOLD response (notably in superior parietal and lateral occipital cortex) when it matched rather than mismatched the content of VWM. In addition, multivariate pattern analyses revealed that classifying between shape categories of the probe yielded higher decoding accuracies when the probe matched the concurrently memorized shape. In contrast, the results of Experiment 2 revealed that, when the probe is interocularly suppressed, its neural response is not modulated by the content of VWM. Together, our results demonstrate that the neural response to visual input is enhanced when it matches the content of VWM. Considering that this effect was non-retinotopic (Experiment 1) and was not observed when visual input was interocularly suppressed (Experiment 2), we conclude that the interaction between visual input and VWM originates in high level visual areas.

Introduction

Humans navigate in a dynamic visual environment. As a result of this, it is often necessary to maintain a visual representation in the mind's eye even after the visual input has disappeared or changed. Visual working memory (VWM) is used to keep visual information available for subsequent, goal-directed behavior, after termination of its sensory input. During VWM maintenance, however, the visual system continues to receive visual input. This raises the question how the processing of visual input is affected by the concurrent content of VWM.

Behavioral experiments have demonstrated that visual input matching the content of VWM attracts attention (Soto, Heinke, Humphreys, & Blanco, 2005; Olivers, Meijer, & Theeuwes, 2006; van Moorselaar, Theeuwes, & Olivers, 2014) and eye movements (Hollingworth, & Luck, 2009; Hollingworth, Matsukura, & Luck, 2013; Schneegans, Spencer, Schöner, Hwang, & Hollingworth, 2014; Silvis & Van der Stigchel, 2014) in search tasks than visual input that mismatches the content of VWM. Along similar lines, visual input matching the content of VWM elicits more pronounced behavioral priming effects under backwards masking (Pan, Cheng, & Luo, 2012), appears to last longer (Pan, & Luo, 2012), dominates bistable perception (Gayet, Brascamp, Van der Stigchel, & Paffen, 2015; Scocchia, Valsecchi, Gegenfurtner, & Triesch, 2013; but, see Scocchia, Valsecchi, Gegenfurtner, & Triesch, 2014), and is released from interocular suppression earlier in time (Chapter 6; van Moorselaar, Theeuwes, & Olivers, 2015; Pan, Lin, Zhao, & Soto, 2014; Gayet, Paffen, & Van der Stigchel, 2013) than visual input that mismatches the content of VWM. Together, these studies point towards enhanced visual processing of visual input that matches the content of VWM.

We consider the possibility that visual representations elicited by VWM and visual representations elicited by retinal input rely upon a shared neural substrate. Such a shared neural substrate would allow for the content of visual working memory to enhance concurrent visual input, selectively when it matches, but not when it mismatches the content of visual working memory. Imaging studies using multivariate pattern analysis (MVPA) provide the means to investigate whether representations elicited by VWM and visual representations elicited by retinal input indeed rely upon the same visual processing areas. Accordingly, it has been shown that simple visual features, such as orientations, that are maintained in VWM can be decoded from early visual processing areas such as V1 (Harrison, & Tong, 2009; Serences, Ester, Vogel, & Awh, 2009), whereas more complex visual stimuli, such as objects and patterns, can be decoded from lateral occipital areas such as LOC (Xu, & Chun, 2006) or parietal areas such as IPS (Ester, Sprague, & Serences, 2015; Christophel, Hebart, & Haynes, 2012; Xu & Chun, 2006; Song & Jiang, 2006; Bettencourt & Xu, 2015). When considering the neural processing of visual representations elicited by sensory input, a similar pattern of findings is observed. For instance, visual areas that allow for decoding orientation information include V1 (Haynes, & Rees, 2006), while visual objects can be decoded from LOC (e.g., Cichy, Chen, & Haynes, 2010; Guggenmos, Thoma, Cichy, Haynes, Sterzer, & Richardson-Klavehn, 2015). Similarly, IPS is believed to play an important role in the neural processing of complex visual patterns (for a review, see Xu & Chun, 2009). Thus, visual representations rely upon the same visual processing areas, irrespective of whether they are sensory or mnemonic in origin. This view is further accentuated by a growing number of imaging studies using MVPA (Albers, Kok, Toni, Dijkerman, & de Lange, 2013; Harrison, & Tong, 2009; Rigall, & Postle, 2012) and computational modelling (Chapter 6) that directly addressed the question of whether visual representations elicited by VWM and visual representations elicited by retinal input draw upon the same neural substrate.

The present study was aimed at investigating the nature of this interaction between a visual representation elicited by visual input and a visual representation that is concurrently maintained in VWM. Based on the behavioral literature presented above, we expected to observe an enhanced neural response to visual input that matches the concurrent content of VWM. To manipulate the content of VWM, we used a delayed match-to-sample task. Participants were retro-cued to memorize one of two sequentially presented shapes (hereafter the primes), drawn from different shape categories, for subsequent recall. During the retention interval a different shape (hereafter the probe) was shown briefly, at an unpredictable moment in time and spatial location. This shape could either match the shape category of the cued prime (Memorized condition), of the uncued prime (Discarded condition), or it could be of a third shape category that was not used on that trial (Unrelated condition). The probe was task-irrelevant to the participants. The data revealed that the probe elicited a stronger BOLD response when it matched rather than mismatched the shape category that was concurrently held in VWM. In particular, this included visual processing areas corresponding to LOC and IPS. Next, we used multivariate pattern analysis (MVPA) to assess whether the visual representations elicited by our shape stimuli are enhanced (i.e., a more differentiated pattern of activity between shape categories) when the probed shape and the memorized shape are in accordance. First, decoding the shape category of the probe revealed that probes of differing shapes elicited a more differentiated pattern of activity when they matched rather than mismatched the shape category maintained in VWM. Second, decoding the shape category in VWM (during presentation of the probe) elicited a more differentiated pattern of activity when they matched rather than mismatched the shape category of the probe. Together the data demonstrate that when visual representations elicited by retinal input match concurrent visual representations elicited by VWM, the neural response to these representations is both quantitatively (i.e., stronger BOLD response) and qualitatively enhanced (i.e., more differentiated patterns of activity). This finding supports the idea that visual representations draw upon the same neural substrate, irrespective of whether their origin is retinal or mnemonic.

Behavioral studies have shown that visual input that matches the content of VWM gains faster access to awareness (e.g., Pan et al., 2014; Gayet et al., 2013). This finding suggests that the content of VWM has the potency to affect the processing of visual input before it reaches awareness Although the paradigm employed in these behavioral studies provides a measure of access to awareness, it does not allow for making claims about processing in the absence of awareness (Gayet, Van der Stigchel, & Paffen, 2014b; Stein, Hebart, & Sterzer, 2011; Stein, & Sterzer, 2014). Therefore, we conducted a second experiment that was similar to Experiment 1, but in which the probe stimulus was suppressed from awareness by continuous flash suppression (Tsuchiya, & Koch, 2005). Under these conditions, the neural response to the probe was unaffected by the concurrent content of VWM. From this we conclude that, in conditions of complete unawareness, the content of VWM does not have the potency to modulate concurrent visual input. Based on recent literature, we hypothesize in the General Discussion that VWM affects the processing of interocularly suppressed visual input when it transits from complete unawareness to full blown awareness.

Experiment 1

Methods

Participants

The group of participants in Experiment 1 consisted of 15 students (3 males, 24 years of age, *SD* = 4) from either the Humboldt University of Berlin or the Charité medical school of Berlin, participating for monetary reward. All participants had (corrected to) normal vision, and participated in at least one behavioral version of this experiment prior to taking part in the fMRI experiment. All participants gave their informed written consent prior to participating in this study, which complied with the guidelines set out in the declaration of Helsinki and was approved by the local ethics committee.

Stimuli

The shape stimuli used in all experiments (see Figure 1) consisted of plain rectangles, isosceles triangles, and ellipses of equal surface (and therefore of equal overall luminance). Within each of these three shape categories, 9 shape variants were created by varying the height-width ratio between 0.75 and 1.25 following a sigmoid function. Pilot studies revealed that this allowed for equating the perceptual dissimilarities between each successive shape within a shape category. The stimuli presented during the retention interval (i.e., the probes) were the center representatives of the three shape categories (i.e., a square, an equilateral triangle, and a circle). These stimuli were presented at a fixed eccentricity of 3 degrees of visual angle on one of 14 equally interspaced locations on the left and right arcs of an imaginary circle, delimited by its main diagonals (see Figure 2, bottom left). The other shape variations were used for the memory task, and were always presented at fixation. As such the cued (i.e., memorized) and uncued (i.e., discarded) primes were never identical to, and always at a different retinal location than the probe stimulus.

Procedure

Participants took part in 144 experimental trials during the functional scans, and 6 minutes of practice trials during the preceding structural scan. On each trial (see Figure 2), participants were subsequently presented with two primes, drawn from two different shape categories, for 400 ms each. This was followed by a 400 ms interval after which a retro-cue was presented for 800 ms. This retro-cue, the number

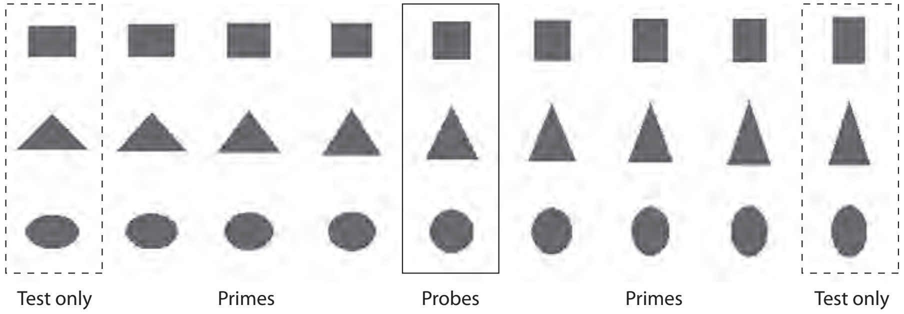

Figure 1. This is a schematic depiction of the stimulus set used in Experiment 1 and 2. The height-width ratios (but not the gray values) are exactly as used in the experiment. There were three shape categories, displayed on separate rows: rectangles, triangles and ellipses. The probes (presented during the delay interval) consisted of the center representatives of each shape category. The primes (cued or uncued) consisted of six distinct variations per shape category. All stimuli depicted in this image could be presented in the recall task of the experiment (i.e., the test phase).

"1" or "2", instructed participants to memorize either the first or the second prime for later recall. After a delay varying between 3 and 6 seconds (M = 4.5) a probe was presented for one second on one of 14 possible locations. Crucially, the probe could either match the shape category of the cued prime (hereafter referred to as the Memorized condition), it could match the shape category of the uncued prime (hereafter referred to as the Discarded condition), or it could match the shape category that was not used on that trial (hereafter referred to as the Unrelated condition). After another delay varying between 5.5 and 9.5 seconds (*M* = 7.5) two test-stimuli appeared at fixation for 1.5 seconds, one of which was identical to the cued (i.e., memorized) prime, and one of which had a slightly different height-width ratio (one step in either direction). Participants were instructed to report by means of a left-hand or right-hand button press which of these two test-stimuli was identical to the cued prime. After the participant gave a response, or after 3.5 seconds had passed, the fixation dot changed from blue to red to indicate that the trial had ended. After an inter trial interval that lasted between 1.5 and 4.5 seconds (*M* = 3) the fixation dot turned blue again, to indicate that the next trial would begin after one second.

Experimental design

The experimental design comprised the within-subject factor Congruence (Memorized, Discarded or Unrelated condition) as a main factor of interest. Factors of no interest included the shape of the probe (square, triangle or circle), the hemifield to which the probe was presented (left or right of fixation), the retro-cue (instruction to memorize first or second shape), and the correct answer in the memory task (left or right of test-stimulus is correct). Factors that were equally (and randomly) distributed over the entire experiment, but were not counterbalanced with the other factors, included the exact shape of the cued and uncued shape stimuli (one of 6 variations within each shape category), the exact shape of the incorrect answer in the memory recall task (higher or lower height-width ratio than the cued shape stimulus), and the exact angular position of the probe (one of seven positions within each hemifield, see Figure 2).

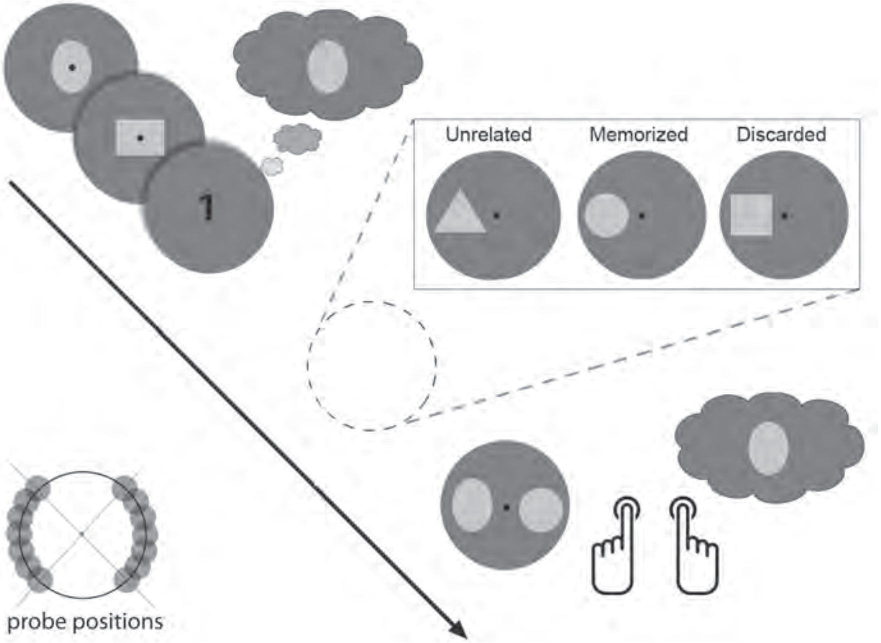

Figure 2. This is a schematic depiction of a trial in Experiment 1. Each trial started with two centrally presented shape stimuli from different shape categories (the primes), followed by a retro-cue indicating which of the two should be memorized for later recall. During the delay interval, a different shape stimulus (the probe) was presented for one second at an unpredictable timing and location (one of 14 possible locations, as seen in the bottom left of this figure). The probe was either a perfect circle, a perfect square or an equilateral triangle. As such, the probe was the center representative (see Figure 1) of either the Memorized shape category, the Discarded shape category, or the shape category that was Unrelated to the memory task. After the delay interval, two test-stimuli were presented, which were both drawn from the shape category of the cued (i.e., memorized) prime. Participants were required to report which of these was the exact shape they had been cued to memorize at the start of the trial.

Functional localizer

We conducted a separate functional localizer run after the experimental runs. The aim of the functional localizer run was primarily to locate the brain regions that are responsive to the presentation of our stimuli, relative to baseline (i.e., compared to a situation in which fixation was maintained on the background, but no additional shape stimuli were presented). A secondary aim was to functionally locate the lateral occipital complex (LOC), such as to target brain areas that are responsive to such high-level visual representations as the shape categories used in the present experiment (e.g., Grill-Spector, Kourtzi, & Kanwisher, 2001). We aimed to define LOC as those occipital voxels that respond more strongly to our shape stimuli than to scrambled shapes. Scrambled versions of the shape stimuli were obtained by randomly rearranging an 8 by 8 tiling of a square-shaped area containing the shape stimuli. The localizer was comprised of a mini-block design with presentation of intact shapes, scrambled shapes and no shapes (baseline) as blocked conditions. Each of these three mini-blocks lasted 29.4 seconds, and was separated by inter-block-intervals that varied between two and five seconds (*M* = 3.5 seconds). The sequence of three mini-blocks was repeated eight times in random order. Within a single mini-block, each of the 21 different shapes was presented once to the left hemifield and once to the right hemifield, for 450 ms followed by a fixation interval of 250 ms. Within each hemifield, each of the seven possible locations was used twice in random order. Participants were instructed to maintain fixation during the entire run and to press a button whenever they perceived the same shape twice in succession. This would occur at an unpredictable moment, three times per mini-block.

fMRI data acquisition and preprocessing

Functional MRI data were acquired on a 3 Tesla Siemens Trio (Siemens, Erlangen, Germany) equipped with a 12-channel head-coil, using a T2-weighted gradient-echo EPI sequence. The fMRI session comprised 8 experimental runs, in which 213 whole-brain volumes were acquired, and a functional localizer run, in which 242 whole-brain volumes were acquired. The fMRI runs (2000 ms repetition time, 25 ms echo time, 78° flip angle, voxel size 3 mm isotropic, 33 slices acquired in descending order, 0.75 mm inter-slice gap) were preceded by a high-resolution T1-weighted MPRAGE structural scan (192 sagittal slices, 1900 ms repetition time, 2.52 ms echo time, 9° flip angle, 256 mm field of view).

Preprocessing was performed using SPM12 (www.fil.ion.ucl.ac.uk/spm) and included slice-time correction, spatial realignment and co-registration with the structural image. For multivariate analyses, the images were smoothed with a 4 mm Gaussian Kernel. For univariate analyses, the images were smoothed with an 8 mm Gaussian Kernel, and normalized to the standard Montreal Neurological Institute (MNI) template.

Regions of interest (ROI's)
The present study was aimed at investigating whether information retained in VWM engenders enhanced visual representations of matching visual input. Considering this research question, we were particularly interested in constraining our analyses to those voxels that were responsive to the visual presentation of our stimuli, to investigate whether these voxels would be modulated by the concurrent content of VWM. For this purpose, we created our regions of interest (ROI's) based on the collection of voxels that were significant ($p_{FWE} < 0.05$) in the Stimulus > Baseline contrast of the functional localizer run. This contrast yielded 1583 voxels, which were all located in the lateral occipital cortex, the inferior and superior parietal lobules and the posterior part of the frontal lobe. Three more specific ROI's were created by intersecting this collection of stimulus responsive voxels with anatomical masks (derived from the AAL Atlas, Tzourio-Mazoyer et al., 2002), such as to separately analyze the posterior frontal, lateral occipital and parietal voxels. For the multivariate analyses, we were mostly interested in the significant voxels from the Stimulus > Baseline contrast that were situated in the lateral occipital cortex (including LOC) and the inferior and superior parietal lobules (including IPS). This resulted in an ROI comprising 1105 voxels, reverse-normalized to native subject space. We targeted these regions, as they are likely to be involved in VWM maintenance (e.g., Xu, & Chun, 2006; Christophel et al., 2012; Song, & Jiang, 2006; Bettencourt & Xu, 2015) and visual object processing (e.g., Cichy, et al., 2010; Guggenmos, et al., 2015; Kourtzi & Kanwisher, 2000), and therefore constitute the most likely locations for VWM to affect perception of our shape stimuli.

There were two main motivations for using the Stimulus > Baseline contrast from the localizer run rather than from the experimental runs. First, this allowed for determining our ROI's on the basis of a dataset that was independent from the data on which the ROI analyses were conducted. Second, considering that the stimuli were suppressed in the second fMRI experiment, using the contrast from the functional localizer run allowed for following the same procedure to create an

ROI in both fMRI experiments. Finally, as the Intact > Scrambled contrast of the functional localizer run was unsuccessful in targeting LOC this contrast was not used for fabricating ROI's.

fMRI univariate data analysis

To investigate how the neural response to the probe was affected by its contingency with the content of VWM, we first estimated a general linear model (GLM) on the smoothed (8mm) and normalized (MNI) data of individual participants. This GLM included a regressor for the probe stimulus (in either the Memorized, Discarded, or Unrelated condition) as well as non-specific regressors for the onsets of all visual stimuli (the two prime stimuli, the retrocue, and the test stimuli). These regressors were modelled as stick functions and were convolved with the canonical hemodynamic response function provided in SPM12. Additionally, six regressors for head motion were included in the GLM. The estimated beta images from the GLM were used to compute separate contrast images for the response to the probe in the Memorized, Discarded and Unrelated conditions against baseline.

First, we performed an exploratory whole brain analysis, to investigate which brain regions are modulated by the contingency between the probe and the content of VWM. For this purpose, we implemented the three contrast images (Memorized, Unrelated, Discarded) in a flexible factorial ANOVA (1 within-subject factor with 3 levels) at the group level. This allowed us to estimate whether the neural response to the probe stimulus differed between these Congruence conditions. Subsequent planned pair-wise *t*-contrasts at the group level were conducted between each pair of Congruence conditions (Memorized versus Discarded, Memorized versus Unrelated, Discarded versus Unrelated) to investigate the nature of the main effect of Congruence. Significance was determined at the cluster level (i.e., $p_{\text{uncorrected}} < 0.001$, cluster-level $p_{\text{FWE}} < 0.01$).

Second, we aimed to test our hypothesis that the response to the probe (as determined by the Stimulus > Baseline contrast from the localizer) is enhanced when it matches the content of VWM. To address this question empirically, we extracted the average parameter estimates for each of the three Congruence conditions within the four different ROI's described above (unmasked, lateral occipital cortex, inferior and superior parietal lobules and posterior frontal lobe). Repeated measure ANOVA's with the factor Congruence (and post-hoc comparisons) within each ROI allowed to assess whether those voxels that responded to the presentation of our shape stimuli in the functional localizer run additionally exhibit an increase in BOLD response when the probe matches the content of VWM. This would be reflected as a higher average parameter estimate in the Memorized condition compared to the Discarded and Unrelated conditions.

fMRI multivariate data analysis

To investigate whether visual representations retained in VWM would interact with visual representations elicited by visual input, we conducted multivariate analyses. The rationale is that if the same neural populations, and therefore the same pattern of activity, represent a specific shape (say an ellipse) irrespective of whether it stems from visual input or from VWM, the patterns of activity distinguishing two shapes (say an ellipse and a rectangle) should be more distinct if the probe and the memorized prime are in accordance than when they are in discordance. Consequently, we expected that this would be reflected as higher decoding accuracies in the Memorized condition compared to the Discarded and Unrelated conditions.

For these analyses, we first estimated a GLM for each subject, on the 4mm smoothed and non-normalized data. The GLM included shape specific (i.e., rectangle, ellipse, or triangle shape categories) regressors for the two primes, for the probe (i.e., rectangle, ellipse or triangle, in either the Memorized, Discarded or Unrelated condition), and for the two test-stimuli. The estimated beta images from the GLM were used for support vector machine (SVM) classification. SVM classification was performed with The Decoding Toolbox (Hebart, Görgen, & Haynes, 2014), using a linear SVM (libsvm). Classification was performed within the main ROI comprised of parietal and lateral occipital voxels (but see Supplementary Materials S1 for classification in the other ROI's), following a leave-one-run-out cross-validation procedure. On each iteration the classifier was trained on the beta maps of seven runs and tested on the beta maps of the remaining eighth run. Classification was done separately for the three pairs of shape categories (rectangle versus triangle, rectangle versus ellipse, and triangle versus ellipse) in each of the three Congruency conditions (i.e., Memorized, Discarded, and Unrelated). Classifier performance was then analyzed at the group level. In order to determine whether the content of VWM affected classifier performance, we used repeated measures ANOVA's (*F*-tests) with the factor Congruence as factor of interest. Next, we conducted one-sided one-sample t-tests (i.e., $\alpha = 0.1$), to test whether classifier performance was higher than the 50% chance level.

Results

Behavioral results

Participants were 62% (*SD* = 6) accurate in reporting which of two shape variations was identical to the cued (i.e., memorized) shape. A paired-samples t-test revealed that their performance did not depend on whether they were cued to memorize the first (*M* = 61%, *SD* = 8) or second (*M* = 63%, *SD* = 6) shape stimulus ($p > 0.3$). Also, within-subject *F*-tests revealed that participants' accuracy on the recall task did not significantly differ between experimental runs ($p > 0.5$, *M* ranging between 60% and 64%, SD ranging between 7 and 9), nor did it differ between the rectangle (M = 61%, SD = 9), triangle (M = 61%, SD = 7), and ellipse (M = 64%, SD = 7) Shape conditions ($p > 0.2$) or the Memorized (M = 63%, SD = 8), Discarded (M = 61%, SD = 7), and Unrelated (M = 63%, SD = 7) Congruence conditions ($p > 0.4$).

Univariate results

To assess the influence of the contents of VWM on the neural response to the probe, we first tested whether the factor Congruence significantly accounted for variation in the data (whole brain). A repeated-measures ANOVA showed a main effect of Congruence in 8 different clusters, $p_{\text{set-level}} < 0.001$. This included left and right lateral occipital cortex, left and right superior parietal lobe, right insular cortex, and three (mostly) posterior frontal regions. A complete overview of statistics and coordinates for each cluster and corresponding peak voxels is provided in Table 1. Next, we evaluated the individual pair-wise contrasts between the Congruence conditions. Following our hypothesis, we expected that the probe would elicit a stronger BOLD response when it matched, compared to when it mismatched the shape category of a concurrently memorized shape. Figure 3A reveals the clusters of voxels that elicited a stronger BOLD response in the Memorized condition compared to the Discarded (top) and Unrelated (bottom) conditions (i.e., $p < 0.001$, cluster-level $p_{\text{FWE}} < 0.01$). This included visual processing regions in the inferior division of the left (and right for the Memorized > Unrelated contrast) lateral occipital cortex extending into the inferior temporal gyrus, and in the left and right superior parietal lobe extending into the superior division of the lateral occipital cortex. In addition, this included frontal regions along the left and right pre-central gyrus extending into the pars triangularis of the right inferior frontal gyrus. All significant clusters of voxels in the different contrasts of this Experiment are described in Table 1. The reciprocal contrasts (Unrelated > Memorized, and Discarded > Memorized) revealed no significant difference in BOLD response at the cluster level. Finally, there was no differential activation between probes in the Discarded and Unrelated conditions (Discarded > Unrelated, and Unrelated > Discarded contrasts). From this, we conclude that visual input elicits a stronger neural response when it matches rather than mismatches the content of VWM.

Table 1. *Results of univariate t-contrasts between Congruence conditions at the group level*

Contrast	*p* (set-level)	Region	*p* (cluster)	*t*/*F* statistic	*k* voxels	Coordinates*		
Flexible factorial ANOVA	< 0.001	R. mid. frontal	< 0.001	4.71	416	48	30	20
		L. inf. occipital	0.001	4.68	280	-48	-62	-4
		R. post. frontal	< 0.001	4.68	435	36	2	42
		L. sup. parietal	0.001	4.35	307	-24	-58	52
		R. sup. parietal	< 0.001	4.17	502	38	-46	44
		Pre-SMA	0.008	4.12	205	6	16	54
		R. insula	0.005	4.11	227	30	22	0
		R. inf. occipital	0.008	3.68	205	52	-42	-8
Memorized > Discarded	< 0.001	L. inf. occipital	0.002	6.71	385	-48	-62	-4
		R. post. frontal	< 0.001	6.03	865	46	30	22
		L. sup. parietal	< 0.001	5.70	516	-24	-60	54
		R. sup. parietal	< 0.001	5.23	688	38	-44	44
		L. post. frontal	0.003	4.90	342	-46	0	44
Memorized > Unrelated	< 0.001	R. post. frontal	< 0.001	6.35	949	36	2	42
		L. inf. occipital	< 0.001	5.45	550	-50	-58	-10
		R. sup. parietal	0.001	5.44	969	40	-46	42
		L. sup. parietal	< 0.001	5.35	850	-24	-58	52
		R. inf. occipital	< 0.001	4.91	494	52	-50	-14
		L. post. frontal	< 0.001	4.53	929	-34	-8	52

Note. This table contains all clusters of voxels that were significantly ($p_{uncorrected} < 0.001$, cluster-level $p_{FWE} < 0.01$) modulated by the factor Congruence at the group level in Experiment 1. The upper part of the table refers to the flexible factorial ANOVA that includes all three Congruence conditions (i.e., Memorized, Unrelated and Discarded), and the other two parts refer to the pair-wise comparisons between Congruence conditions. No significant differences were found in the pair-wise comparisons that are not reported here (i.e., Discarded > Memorized, Unrelated > Memorized, Discarded Unrelated, and Unrelated > Discarded).
*Coordinates are reported in millimeters (MNI).

Next, we investigated whether the neural response to probe stimuli was modulated by the content of VWM within our functionally defined ROI's. A repeated measures ANOVA with the factors Congruence (Memorized, Unrelated, Discarded) revealed a main effect of the factor Congruence, $F(2, 13) = 12.33$, $p = 0.001$ on the average parameter estimates within all significant voxels from the Stimulus > Baseline contrast. Subsequent *t*-tests revealed that the BOLD response in these voxels was stronger for probes in the Memorized condition than in either the Discarded, $t(14) = 4.30$, $p = 0.001$, or the Unrelated condition, $t(14) = 4.56$, $p <$

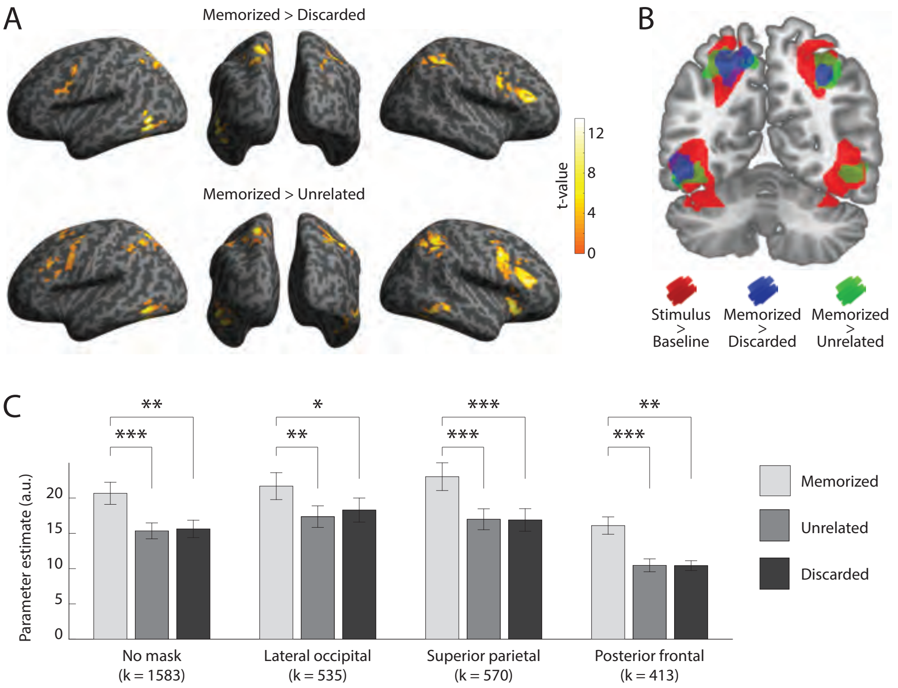

Figure 3. Results of the univariate analyses of BOLD response to the probe in Experiment 1. Panel A reveals clusters of significant voxels ($p_{\text{uncorrected}} < 0.001$, cluster-level $p_{\text{FWE}} < 0.01$) from two contrasts, projected on an inflated surface of the standard Montreal Neurological Institute (MNI) template brain. None of the other contrasts (i.e., Unrelated > Discarded, Discarded > Unrelated, Unrelated > Memorized, and Discarded > Memorized) yielded significant clusters of voxels. Panel B depicts a coronal slice encompassing visual processing areas in occipital and parietal cortex (Y-coordinate approximately 35 mm, MNI) on which the significant clusters of voxels from panel A are made binary and superimposed on the significant clusters of voxels from the Stimulus > Baseline contrast of the functional localizer run. Panel C depicts the average parameter estimates for each Congruence condition (Memorized, Unrelated and Discarded) against baseline, for different regions of interest (ROI's) at the group level. These ROI's were comprised of the significant voxels ($p < 0.05_{\text{FWE}}$) in the Stimulus > Baseline contrast of the functional localizer run (three left-most bars), separated into occipital, parietal and frontal voxels (rest of the graph). Error bars represent the standard error of the mean. *$p < 0.05$, **$p < 0.005$, ***$p < 0.0005$

0.001, but did not differ between the Discarded and the Unrelated conditions, $t(14) = 0.31$, $p = 0.764$. The same pattern of results was apparent in the three smaller ROI's. Three separate repeated measures ANOVA's revealed that the main effect of the factor Congruence was significant in the occipital ROI, $F(2, 13) = 6.86$, $p = 0.009$, in the parietal ROI, $F(2, 13) = 13.76$, $p < 0.001$, and in the frontal ROI, $F(2, 13) = 10.28$, $p = 0.002$. Subsequent *t*-tests revealed that probes elicited a stronger BOLD response in the Memorized condition than in the Discarded condition within the occipital ROI, $t(14) = 2.83$, $p = 0.013$, the parietal ROI, $t(14) = 4.95$, $p < 0.001$, and the frontal ROI, $t(14) = 3.97$, $p = 0.001$. Similarly, probes elicited a stronger BOLD response in the Memorized condition than in the Unrelated condition within the occipital ROI, $t(14) = 3.76$, $p = 0.002$, the parietal ROI, $t(14) = 4.74$, $p < 0.001$, and the frontal ROI, $t(14) = 4.68$, $p < 0.001$. Again, the BOLD response to probes in the Unrelated and Discarded condition did not differ in any of the ROI's (all p's > 0.4). Average parameter estimates for all ROIs discussed here are depicted in Figure 3C. Together, these findings indicate that those brain regions that respond to the presentation of our shape stimuli, show a larger BOLD response when these stimuli match rather than mismatch the content of VWM.

Multivariate results

To investigate whether we could decode the shape category of the probe, irrespective of its match with the shape category maintained in VWM, we averaged classifier performance over the three Congruency (Memorized, Discarded, Unrelated) conditions and Shape-pair (ellipse versus triangle, ellipse versus rectangle, triangle versus rectangle) iterations. This revealed an overall decoding accuracy at the group level of 54.1% ($SD = 5.5$), $t(14) = 2.89$, $p = 0.012$. A repeated measures ANOVA with the factors Congruence and Shape-pair revealed that classifier performance depended on Congruence, $F(2, 13) = 5.18$, $p = 0.022$, but not on Shape-pair (p > 0.5), nor on the interaction between Shape-pair and Congruence ($p > 0.1$). This shows that classifier performance depended on the contingency between the shape category of the probe and the shape category of the concurrent content of VWM (i.e., the cued prime). Subsequent t-tests revealed that classifier performance for the shape of the probe was above chance in the Memorized condition ($M = 58.0\%$, $SD = 11.0$), $t(14) = 2.79$, $p = 0.014$, but not in either the Unrelated ($M = 54.9\%$, $SD = 11.4$), $t(14) = 1.68$, $p = 0.115$, or the Discarded condition ($M = 49.4\%$, $SD = 6.2$), $t(14) = -0.36$, $p = 0.722$. Pair-wise comparisons showed that classifier performance was higher for probes in the Memorized condition than in the Discarded condition, $t(14) = 2.52$, $p = 0.024$, but did not significantly differ between either the Memorized and Unrelated conditions ($p > 0.05$) or the Unrelated and Discarded conditions ($p > 0.05$). The data further suggest that this pattern of results was mostly driven by activity in the lateral occipital cortex (see Supplementary Materials S1, Figure S1, for decoding accuracies in three different ROI's). The difference between the Memorized condition and the Discarded condition is crucial, as this contrast selectively targets the influence of VWM (this issue is further elaborated in the General Discussion). Together, these findings show that, even with identical stimulus presentation, patterns of brain activity differ more between shape categories when the visual representation elicited by VWM is in accordance with the visual representation elicited through retinal input, compared to when these visual representations are in discordance. These results are depicted in figure 4A.

In the Memorized condition, the probe and the cued (i.e., memorized) prime are of the same shape categories. In the Discarded and Unrelated conditions, however, the probe could be of either of the two other shape categories with equal probability. An alternative interpretation of the present results is, therefore, that the classifier was unable to dissociate between shape categories of the probe, but rather dissociated between shape categories of the cued (i.e., memorized) prime. This interpretation would also predict significant decoding in the Memorized condition, but chance level decoding in the Unrelated and Discarded conditions. As such, we aimed to investigate whether the increased decoding accuracy in the Memorized condition was due to (A) an enhanced neural representation of probes that match the content of VWM or (B) was the consequence of the classifier picking up the shape category in VWM rather than the shape category of the probe. For this purpose, we computed a new model, allowing the classifier to classify between shape categories maintained in VWM (i.e., the shape category of the cued prime) during the presentation of the probe. The GLM for this analysis contained the same regressors as that of the previous analysis. The only difference was the specific labeling of the shapes during the presentation of the probe (i.e., following the shape categories of the cued primes rather than that of the probes).

If, in this case, classifier performance in the Unrelated and Discarded conditions is either at chance or lower than in the Memorized condition, this shows that it is indeed the contingency between the shape category maintained in VWM and the shape category of the probe that enhanced classifier performance in the Memorized condition. In contrast, if significant classification of the shape category maintained at VWM is observed in all Congruence conditions, and does not depend on the shape category of the probe, this would indicate that the contingency between the probe and the shape category in VWM did not influence classifier performance. Rather, this would demonstrate that classification of the probe in the Memorized condition purely reflected classification of the shape category in VWM.

First, we aimed to investigate whether we could decode the shape category of the cued prime, irrespective of its match with the concurrently presented probe. For this, we averaged decoding accuracies over the three Congruency conditions. This revealed an overall decoding accuracy for the shape category of the cued prime of 54.9% (*SD* = 5.0), $t(14) = 3.81$, $p = 0.002$, which did not differ from the overall decoding accuracy for the shape category of the probe ($p > 0.5$). Next, we conducted a repeated measures ANOVA with the factors Source (decoding probe or decoding memory) and Congruence (Memorized, Discarded, and Unrelated). This analysis revealed a main effect of Congruence, $F(2, 13) = 4.15$, $p = 0.040$, but no main effect of Source ($p > 0.5$), nor did Shape and Congruence interact ($p > 0.3$). This suggests that the factor Congruence influences decoding accuracies, irrespective of whether the classifier is decoding between shape categories of the probe, or shape categories of the cued prime. To further investigate this finding we conducted one-sample t-tests to assess in which Congruence conditions the classifier performed above chance level. This revealed that the classifier was able to decode the shape category of the cued (i.e., memorized) prime during the presentation of the probe in the Memorized condition (*M* = 59.1%, *SD* = 9.0), $t(14) = 3.79$, $p = 0.002$, but not in the Unrelated (*M* = 52.7%, *SD* = 11.2; $p > 0.3$) or Discarded conditions (*M* = 53.0%, *SD* = 10.1; $p > 0.2$). Subsequent pair-wise comparisons revealed that classifier performance was marginally higher in the Memorized condition than in the Discarded condition, $t(14) = 1.92$, $p = 0.075$ (significant in one-side t-test, justified by our prior expectations on the directionality of this difference), while classifier performance did not differ between the Memorized and Unrelated conditions ($p > 0.1$), or between the Unrelated and Discarded conditions ($p > 0.9$). These findings (depicted in Figure 4B) provide evidence against the alternative explanation that classifier performance was primarily driven by the mnemonic signal in the previous analysis.

Thus, decoding between shape categories of a visual stimulus yields higher classifier performance when it matches rather than mismatches a concurrently memorized shape category. Conversely, decoding between shape categories of a memorized stimulus yields higher classifier performance when it matches a concurrently presented visual stimulus. Together, these results demonstrate that classifier performance did not exclusively rely on representations of either retinal or mnemonic origin, but rather capitalized on the increased neural response that is elicited when both are in accordance. These findings are supportive of a model in which the perceptual signal and the mnemonic signal draw upon the same underlying neural substrate, such that differing visual input elicits a more differentiated activity pattern when the visual input matches the content of VWM.

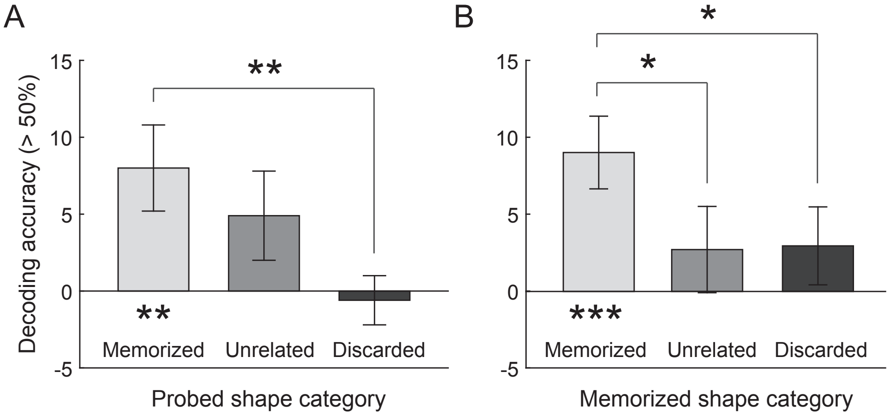

Figure 4. Panel A depicts above chance classifier performance for decoding the shape category of the probe stimulus. The shape of the probe was either of the same shape category as the cued prime (Memorized condition) or the uncued prime (Discarded condition), or it could be of the shape category that was not used on that trial (Unrelated condition). Panel B depicts above chance classifier performance for decoding the shape category of the cued prime (i.e., the shape in VWM), during the presentation of the probe. Again, the results are separated into the three Congruence conditions, describing the relation between the cued prime and the probe. In both cases classification was performed within a region of interest (ROI) that was constituted of the intersection between the Intact Stimulus > Baseline contrast from the functional localizer, and a mask containing the lateral occipital cortex and the inferior and superior parietal lobules. For decoding in other ROI's, see Figure S1 in Supplementary Materials S1. Error bars represent the standard error of the mean. *$p < 0.1$, **$p < 0.05$, ***$p < 0.005$.

Experiment 2

Introduction

The findings from Experiment 1 reveal that visual input matching the content of VWM elicits a stronger, and more informative, neural response. This observation could underlie phenomena observed in a variety of behavioral studies showing that processing of a stimulus is enhanced when it matches rather than mismatches the content of VWM. Among these is the recent finding that stimuli that are initially suppressed by continuous flash suppression (CFS; Tsuchiya, & Koch, 2005) are released from suppression earlier in time when they match compared to when they mismatch the content of VWM (Chapter 6; van Moorselaar, et al., 2015; Pan, et al., 2014; Gayet, et al., 2013). The rationale underlying this release-from-suppression paradigm (b-CFS; Jiang, Costello, & He, 2007; for a review, see Gayet, Van der Stigchel, & Paffen, 2014b) is that, considering that the contingency with VWM determines *when* a stimulus is released from suppression, it is implied that VWM affected processing of the stimulus *before* it was released from suppression (i.e., when it was still invisible). An alternative explanation is that working memory affected processing of the initially suppressed stimulus during the transitory period, that is, while the stimulus was transiting from invisibility to achieving full visibility (this view is elaborated in Gayet, et al., 2014b; Stein, Hebart, & Sterzer, 2011). In line with this idea, a recent study demonstrated that attentional manipulations of dominance durations during binocular rivalry were only effective during so-called transitory periods of unresolved ambiguity (Dieter, Melnick, & Tadin, 2015). Also, interocularly suppressed stimuli were shown to only elicit high-level behavioral priming effects under conditions of partial awareness (Gelbard-Sagiv, Faivre, Mudrik, & Koch, 2016). These findings raise the question whether the content of VWM has the potency to affect concurrent visual processing of visual input that is not (yet) available to awareness.

In order to answer this question, we used the same experimental design as in Experiment 1, while ensuring that the stimuli presented during the delay interval were suppressed from awareness, by means of continuous flash suppression. Any differential effect on BOLD response or decoding accuracy depending on the contingency between the stimulus retained in working memory and the suppressed stimulus, would indicate that the interaction between the content of VWM and a visual stimulus can occur when this stimulus is still completely suppressed from awareness.

Methods

Participants

The group of participants in Experiment 2 eventually consisted of 20 students (4 males, 25 years of age, *SD* = 4) that were gathered from either the Humboldt University of Berlin or the Charité medical school of Berlin, participating for monetary reward. Of these participants, seven participated in Experiment 1 as well. One participant was removed because of a timing error in the stimulus presentation software (all timings were off by 25%). Another participant was unable to fuse the two images into one percept, and the experiment was aborted after two runs. A third participant was removed because the participant turned out to see virtually all stimuli presented to the left visual field, but virtually none of the stimuli presented to the right visual field (50% visibility). Finally, one participant was removed after a binomial test revealed that the participant was significantly above chance in reporting the location of suppressed stimuli that were subjectively reported as invisible.

Stimuli & procedure

The procedure and stimuli of Experiment 2 were identical to those of Experiment 1, except for the following differences (see Figure 5). First, the factor Congruence now had only two levels. This allowed us to maximize our experimental power, while retaining the critical conditions (Memorized and Discarded conditions), in which stimulus presentation is identical. Second, we independently stimulated participants' left and right eye, by presenting two separate images on the screen. Participants wore prism glasses, which allowed them to fuse the two images into one coherent percept. The images were physically separated by a divider, which was positioned between the two eyes of the participants, and reached from the prism glasses all the way to the screen, on which the two images were projected.

In order to facilitate binocular combination of the two images, the two presentation areas were enclosed by a circular Brownian ($1/f^2$) noise frame, adjoined by two vertical white lines. The binocular presentation constituted of a restriction on the size of the presentation area, as a result of which the shape stimuli presented during the retention interval were now presented at an eccentricity of approximately 2 degrees of visual angle. Also, their surface area was reduced by 11% compared to that of Experiment 1.

In order to perceptually suppress the shape stimuli that were presented during the retention interval, they were presented to one eye while a high contrast dynamic pattern was presented to the other eye, thereby eliciting continuous flash suppression. These pattern masks were created by (a) filtering pink (1/f) noise using a rotationally symmetric Gaussian low-pass filter ($\sigma = 1.5$) and by (b) making the resulting gray-scale image binary with maximum contrast (as in Gayet, et al., 2013; Gayet, Van der Stigchel, & Paffen, 2014a). On every trial, 12 new masks were generated, which were successively presented at 10 Hz during the presentation of the probe. In order to strengthen suppression, CFS onset started 100 ms (one mask) prior to probe onset, and in order to minimize after-images of the probe CFS ended 100 ms (one mask) after probe offset. The eye (left or right) to which the probe was presented was added as a factor in the counterbalancing of the experimental design.

In Experiment 2, participants were required to provide two additional responses after the recall task. First, participants were required to report whether or not they perceived the probe during the CFS (forced choice "yes" or "no" with a left or right hand button press). The instructions were formulated such that a negative response would reflect a visibility rating of "1" on the Perceptual Awareness Scale (PAS; Ramsøy & Overgaard, 2004). Next, participants were requested to report (forced choice) whether the suppressed stimulus was presented to the left side or to the right side of fixation.

fMRI data acquisition and preprocessing

The fMRI data acquisition in Experiment 2 was identical to that of Experiment 1, only differing in the number of runs and volumes. The fMRI session of Experiment 2 comprised 6 experimental runs, in which 322 whole-brain volumes were acquired, and a functional localizer run, in which 271 whole-brain volumes were acquired. The preprocessing procedure of the data in Experiment 2 was identical to that of Experiment 1.

Functional localizer

Because we were unable to retrieve a clear functional localization of LOC in Experiment 1, we slightly modified the functional localizer in Experiment 2. The functional localizer now comprised mini-blocks of intact shapes, shape outlines, scrambled shape outlines and a baseline condition. The use of outlines was motivated

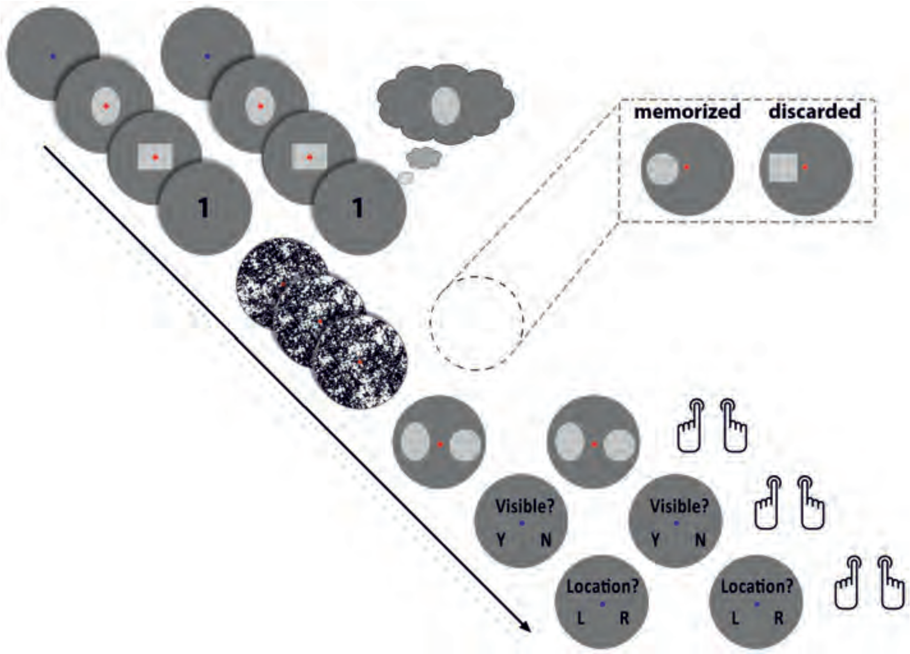

Figure 5. This is a schematic depiction of a trial in Experiment 2, as presented to the left and right eye. The probe either matched the shape category of the cued (Memorized condition) or uncued prime (Discarded condition). The probe was suppressed from awareness by presenting it to one eye, while presenting a high contrast dynamic pattern to the other eye (i.e., continuous flash suppression). After the memory recall task, which was identical to that of Experiment 1, participants were asked whether they had seen the probe (subjective assessment), and they were asked to provide a forced-choice answer as to the location of the probe (objective assessment).

by the idea that, whereas scrambled filled shapes (as used in Experiment 1) still resemble shapes, scrambled shape outlines are no longer shapes but lines. Each of these four mini-blocks lasted 12.6 seconds, and were separated by inter-block-intervals that varied between one and three seconds (*M* = 2). The sequence of four mini-blocks was repeated nine times in random order, with the only restriction that two subsequent mini-blocks were of different conditions. Within a single mini-block, shapes from each of the three shape categories were presented three times to the left hemifield and three times to the right hemifield. The exact locations and exact shape variations were drawn at random from a pool that allowed to maximally equate the prevalence of each. Participants were instructed to maintain fixation during the entire run and to press a button whenever they observed a shape of deviating brightness (about -20% luminance contrast). This would occur at an unpredictable moment, three times per mini-block.

Results and discussion

Behavioral results

In order to ensure invisibility of the stimulus, we analyzed only the trials in which the participant reported subjective invisibility of the probes (*M* = 85%, *SD* = 12). On these trials, participants were unable to report the hemifield to which the stimulus was presented with above chance level performance (*M* = 49% correct, *SD* = 5), *t*(19) = -0.862, *p* = 0.340. This allowed us to guarantee that, for these participants, probes that were reported as invisible were indeed successfully suppressed. Conversely, when participants reported subjective visibility of the probe, they were 69% accurate (*SD* = 31) in reporting the hemifield to which it was presented, which is better than chance, *t*(19) = 9.948, *p* = 0.014. Considering that for five participants this statistic was based on 4 data points or less, it is worth mentioning that over all trials with unsuccessfully suppressed probes (i.e., across participants), 80% yielded a correct response.

When probes were unsuccessfully suppressed, they matched the shape category of the cued prime (*M* = 48%, *SD* = 14) just as often as the shape category of the uncued prime, *p* > 0.5. As such, potential differences in neural response to the probes between Congruence conditions cannot be accounted for by a difference in the number of included trials in each Congruence condition. Subjective invisibility was reported in 89% (*SD* = 16) of the trials in which the probe was presented to participants' dominant eye, and in 80% (*SD* = 16) of the trials in which the probe was presented to the non-dominant eye. This shows that CFS was successful in suppressing probes in both eyes of the participants.

Finally, we analyzed the performance on the recall task, in the subjectively invisible trials only, as these trials were eventually included in the fMRI analyses. Participants were 60% (*SD* = 6) accurate in reporting which of two shape variations was identical to the cued (i.e., memorized) prime. Performance in the recall task did not depend on whether participants were cued to memorize the first (*M* = 60%, *SD* = 6) or the second (*M* = 60%, *SD* = 8) prime (*p* > 0.8), as shown by a paired-samples *t*-test. Also, repeated measures ANOVA's showed that participants' performance on the recall task did not significantly differ between experimental runs (*p* > 0.6), nor did it differ between Shape conditions (*p* > 0.9). Unlike in Experiment 1, participants performed worse on the memory task on trials in which the probe matched the shape category of the cued prime (*M* = 58%, *SD* = 7) compared to when it matched the shape category of the uncued prime (*M* = 61%, *SD* = 6), $t(19) = 2.2568$, $p = 0.019$. Thus, the invisible probe impeded performance on the recall task, when it matched compared to when it mismatched the shape category maintained in VWM. While it is expected that probes are increasingly disruptive to a concurrent memory task when they more closely resemble the to-be-memorized shape, it is unexpected that invisible probes significantly hampered recall performance, whereas visible probes did not (that is, this difference was absent in Experiment 1, and numerically in opposite direction).

Univariate results

Following the exact same procedure as in Experiment 1, we observed no difference in BOLD response between suppressed probes that matched the shape category of the cued prime (i.e., the Memorized condition) and probes that matched the shape category of the uncued prime (i.e., the Discarded condition) at the whole-brain level. In fact, no voxel survived the $p_{\text{uncorrected}} < 0.001$ threshold. This finding suggests that the content of VWM does not affect the neural response to visual input that is fully suppressed from awareness by CFS.

Next, we aimed to ensure that the absence of a differential response to probe stimuli matching the cued prime and probe stimuli matching the uncued prime was not due to experimental insensitivity but rather reflected a genuine null effect. For this purpose, we retrieved the clusters of voxels (in MNI space) that showed a significantly stronger response in the Memorized condition compared to the Discarded condition in Experiment 1, and used these as ROIs in the present analyses. Next, we extracted the average parameter estimates of the Memorized > Discarded contrast for each of the 20 participants included in Experiment 2. We then conducted Bayesian one sample *t*-tests using JASP (Love et al., 2015) to test whether the parameter estimates in these ROIs were significantly larger than 0. This would be indicative of a positive modulation of the BOLD response to probes that matched compared to probes that mismatched the content of VWM. Following Dienes (2014), a Bayes factor of 3 or more, indicating that one hypothesis is at least 3 times more likely than another hypothesis in explaining the observed data, was seen as evidence for this hypothesis. Weaker evidence (i.e., a Bayes factor smaller than 3 and larger than 1/3) is indicative of a lack of experimental power to dissociate between the two hypotheses. With this approach we first investigated whether overall, clusters of voxels that responded significantly more strongly to a probe in the Memorized condition than in the Discarded condition in Experiment 1, did so as well in Experiment 2. The Bayesian analysis revealed that the data were 6.86 times more likely to reflect the null hypothesis (the observed data is drawn from a distribution with a mean of 0), than the alternative hypothesis (the observed data is drawn from a distribution with a mean that is larger than 0). Next, we separated this ROI into three separate ROIs, corresponding to the five significant clusters of voxels that were found in the Memorized > Discarded contrast of Experiment 1 (middle row of Table 1 and top row in Figure 3A): the left lateral occipital cortex, the left and right superior parietal lobule, and the left and right posterior frontal lobe. For all ROIs, the data were between 3.90 times and 7.58 times more likely to support the null hypothesis than the alternative hypothesis (Table 2). Finally, the ROI based on the functional localizer of Experiment 2 (used for the MVPA analyses described in the next section, which was construed following the exact same procedure as in Experiment 1) was also 7.88 times more likely to reflect the null hypothesis than the alternative hypothesis. Together, these findings provide conclusive evidence that overall BOLD response to the interocularly suppressed probe was not enhanced when it matched compared to when it mismatched the content of VWM.

Table 2. Bayesian one-sample t-tests for the average parameter estimates (N = 20) in the [Memorized > Discarded] contrast of Experiment 2, for individual regions of interest (based on Experiment 1).

Region	BF_{0+}	error %	*k* voxels	Mean**	SD**
Lateral occipital left*	3.901	~ 2.152e-4	385	0.116	4.064
Superior parietal left*	6.516	~ 2.101e-4	516	-0.567	3.950
Superior parietal right*	3.998	~ 2.151e-4	688	0.078	3.607
Posterior frontal left*	5.731	~ 2.151e-4	342	-0.353	3.741
Posterior frontal right*	7.579	~ 1.841e-4	865	-0.694	3.301
All significant clusters*	6.863	~ 2.048e-4	2796	-0.528	3.203
Stimulus > baseline ROI	7.875	~ 1.719e-4	2069	-0.752	3.297

Note. Bayesian one-sample *t*-tests testing the hypothesis that the population of average parameter estimates is drawn from a distribution with a mean that is greater than 0. In this case, a Bayes factor (BF_{0+}) above 3 reflects reliable support for the null hypothesis (i.e., no difference in BOLD response to probes in the Memorized and Discarded conditions), whereas a Bayes factor smaller than 1/3 reflects reliable support for the alternative hypothesis (i.e., stronger BOLD response to probes in the Memorized compared to the Discarded condition). Bayes factors are computed in JASP (Love et al., 2015), following the method of Wagenmakers (2007). *These are clusters of voxels that showed a significantly stronger BOLD response to probes in the Memorized condition compared to the Discarded condition in Experiment 1. **Parameter estimates (in arbitrary units) for the Memorized > Discarded contrast of Experiment 2, averaged over regions of interest. SD's reflect variability between participants.

Multivariate results

Although suppressed visual input did not elicit a stronger BOLD response when it matched rather than mismatched the concurrent content of VWM, it is nonetheless possible that these different circumstances elicited differing patterns of neural activity. To investigate this possibility we conducted multivariate analyses, which followed the same procedure as those of Experiment 1. In this experiment, the classifier was unable to dissociate between shape categories of the suppressed shape stimulus per se ($M = 48.7$, $SD = 8.5$), $t(19) = -0.698$, $p > 0.1$. Also, classification performance was at chance in both the Memorized ($M = 47.9\%$, SD = 12.4), $t(19) = -0.771$, $p > 0.4$, and the Discarded condition ($M = 49.5\%$, $SD = 9.7$), $t(19) = -0.243$, $p > 0.8$, and these did not differ, $t(19) = -0.504$, $p > 0.6$. Thus, we were unable to decode the shape category of invisible stimuli, while following the procedure that did allow us to decode the shape category of visible stimuli in Experiment 1, despite having more participants and more trials per participant.

In order to ensure that the absent of a significant effect of Congruence, reported above, reflects a genuine null effect rather than experimental insensitivity, we again conducted Bayesian t-tests. First, Bayesian one-sample t-tests revealed that the overall classifier performance for decoding the shape category of the probe was 6.72 times more likely to reflect chance level classification than above chance-level classification. This was also true for the Memorized condition ($BF_{0+} = 6.97$) and the Discarded ($BF_{0+} = 5.116$) conditions separately. A Bayesian paired-samples t-test revealed that the null hypothesis (no difference in classifier performance between the Memorized and Discarded conditions) was 6.02 times more likely than the alternative hypothesis (better classifier performance in the Memorized condition than in the Discarded condition).

Next, in order to ensure that our current dataset and analysis procedure are in fact suitable at all for decoding the shape categories of such shape stimuli, we computed classifier performance for the shape categories of the test-stimuli that were presented at the end of the delay interval. Classification accuracy for these test-stimuli was 59.9% ($SD = 13.0$) in Experiment 2, which is better than chance, $t(19) = 3.46$, $p = 0.002$, and numerically better than the same analysis performed on the data of Experiment 1 ($M = 56.4$, $SD = 6.6$), $t(14) = 3.67$, $p = 0.003$. As such, the inability to decode the shape category of the suppressed stimulus in Experiment 2 is not caused by potential problems in the data or ROI selection. Rather it is caused by either the invisibility of the probe, or the presentation of the CFS.

Considering that there was no reliable decoding in the Memorized condition, however, we should remain cautious in interpreting the null effect in the multivariate results. This follows from the fact that, in this condition, the classifier not only failed to dissociate between shapes of the probe, but also failed to dissociate between shapes in VWM. This null result can be interpreted in two ways. First, the null result could reflect that the classifier is unable to dissociate between shape categories maintained in VWM unless there is a matching probe (as in Experiment 1). Second, the presentation of the CFS masks prohibited shape classification in general, as a result of which nor the shape of the probe nor the shape category in VWM can be decoded. In this latter case, the null effect does not reflect an absence of an effect of Congruence. Rather, it reflects a general decoding problem caused by the CFS. For this reason, the univariate results of Experiment 2 provide more compelling evidence that the content of VWM does not affect the processing of interocularly suppressed visual input than the multivariate results.

Behavioral experiment

Finally, we aimed to validate the stimulus set that was employed in the current study. Previously, it has been demonstrated that visual input is released from continuous flash suppression faster when it matches rather than mismatches the content of VWM in a b-CFS paradigm (Chapter 6; van Moorselaar et al., 2015; Pan et al., 2014; Gayet et al., 2013). From this, we deduced the hypothesis that the content of VWM might modulate the response to visual input while it is still interocularly suppressed. Considering that, in these studies, the stimuli were either faces or colors, we aimed to ensure that the same pattern of behavioral results can be obtained with the present stimulus set (i.e., using shape stimuli). As such, we replicated the experimental procedures of Gayet et al. (2013) with the current stimulus set (see Supplementary Materials S2, Figure S2). The results revealed that in the Memorized condition, probes were reported faster than in both the Discarded and the Unrelated conditions (Supplementary Materials S2, Figure S3), whereas no difference in response time was obtained between probes in the Discarded and Unrelated conditions. As such, the null results in the fMRI analyses of Experiment 2 cannot be accounted for by an inadequate stimulus choice.

Together, the fMRI analyses of Experiment 2 provide no evidence that information retained in VWM has the potency to affect the processing of visual input under conditions of complete interocular suppression. In fact, despite having more trials and more participants than in Experiment 1, and despite selectively analyzing those voxels that were modulated by the content of VWM in Experiment 1, Bayesian analyses confirmed that the content of VWM had no influence whatsoever on the neural response to interocularly suppressed visual input.

General Discussion

The human visual system oftentimes has to keep information online in VWM for subsequent behavior, while simultaneously continuing to process visual input. This raises the question how visual information that is maintained in VWM affects processing of concurrent visual input. In Experiment 1, we demonstrate that when visual input matches rather than mismatches the content of VWM, it elicits an enhanced neural response. Specifically, those brain regions that were responsive to the presentation of shape stimuli (relative to baseline) showed a stronger BOLD response to shape stimuli when a similar shape was concurrently maintained in VWM. Increased activity levels were observed in the inferior lateral occipital cortex (including LOC), the inferior parietal lobule (including IPS) and the posterior part of the frontal lobe (around the precentral sulcus). In addition, MVPA analyses revealed that, within lateral occipital and parietal areas, the patterns of neural activity elicited by different shape categories (i.e., triangles, rectangles and ellipses) were more distinct when the visual input was in accordance with the content of VWM. That is, classifier performance was higher when both shapes (i.e., the one maintained in memory and the one that was presented on the screen) were drawn from the same shape category, compared to when they were drawn from different shape categories. Together, the data of Experiment 1 reveal that the neural response to visual input is both quantitatively and qualitatively enhanced when it matches the content of VWM.

In Experiment 1, a region corresponding to the pars triangularis of the right inferior frontal gyrus showed a stronger BOLD response when probes matched the shape category of the cued (i.e., memorized) prime compared to the uncued (i.e., discarded) prime. This findings is interesting, as all other areas that showed an increased BOLD response when probes matched the content of VWM, corresponded to areas that showed an increased BOLD response in the Stimulus > Baseline contrast of the functional localizer as well. This means that, generally, probes matching the content of VWM elicited a pattern of neural activity that was enhanced relative to, but not different from, the pattern elicited by the same stimuli in the absence of a concurrent memory task. The pars triangularis of the right inferior frontal gyrus, however, responded specifically to visual input that matched the content of VWM, and therefore likely plays a role in mediating between representations of sensory and mnemonic origins. Whereas the left pars triangularis has been associated with memory retrieval in semantic tasks (Gabrieli, Poldrack, & Desmond, 1998), the right pars triangularis has been associated with the effortful (i.e., conscious) cognitive control of memory in general (Badre, & Wagner, 2007). More specifically, it is involved in selecting task-relevant representations among competing representations in working memory (Thompson-Schill, D'Esposito, & Kan, 1999) by inhibiting task-irrelevant representations (Badre, Poldrack, Pare-Blagoev, Insler, & Wagner, 2005). In the present case, the pars triangularis might have been involved in protecting the representation maintained in VWM from the different representation elicited by the probe presented during the delay interval. Considering that the probes were always different from the to-be-memorized shapes, it might have been more effortful to shield a working memory representation from matching probes (which resemble the memorized stimulus) than from mismatching probes (which do not so much resemble the memorized stimulus). At first sight, the fact that participants in Experiment 1 showed similar performance on the working memory recall task irrespective of the shape of the probe seems at odds with this interpretation. However, it has been shown that VWM representations in occipital cortex, but not in parietal cortex, are disrupted by visual distractors presented during the delay interval, while leaving behavioral performance on the recall task unaffected (Bettencourt, & Xu, 2015). Hence, the role of the pars triangularis might have been to shield the working memory representation in parietal areas from the visual interference elicited by the probes in occipital areas. In Experiment 2, participants performed worse on the working memory recall task when invisible probes matched the shape that was maintained in VWM. Here, no differential activation was observed in the pars triangularis between probes that matched and probes that mismatched the content of VWM. Tentatively, we hypothesize that participants were unable to shield the representation in VWM from the representation elicited by the invisible probe, as a result of which probes that matched (but slightly differed from) the content of VWM hampered performance on the recall task.

In all experiments presented here, we made use of retro-cue paradigm, in which two shape stimuli were presented, followed by a retro-cue indicating which of the two shape stimuli should be memorized for subsequent recall. This method allows for isolating the effect of VWM on the processing of visual input presented during the delay interval. When presenting only one shape for subsequent recall, differences in neural response to probes that match compared to probes that mismatch this shape, cannot unequivocally be attributed to the memorization of the shape. Instead, such differences could also be engendered by, for instance, the amount of resources allocated to the initial encoding of the to-be-memorized shape, irrespective of whether it is subsequently maintained in VWM or not. In a retro-cue paradigm, however, two different shape stimuli are presented, and initial resources that are allocated to the encoding of the two shape stimuli is identical (collapsed across trials). Hence, the only difference between the retro-cued shape stimulus and the uncued shape stimulus is that the former is actively maintained in VWM for subsequent recall. It is worth mentioning that, by jittering the delay between the to-be-memorized shape and the probe, we minimized the chance that the hemodynamic response elicited by the presentation of the probe was contaminated by the hemodynamic response elicited by the presentation of the to-be-memorized shape. This does not solve the issue of differences in initial stimulus encoding, however. Taken together, for the purpose of investigating the effect of VWM on the neural response to the probe, the crucial comparison comprises the condition in which a probe matches the shape category of the retro-cued shape stimulus (i.e., in the Memorized condition) with the condition in which a probe matches the shape category of the uncued shape stimulus (Discarded condition). These conditions are identical in terms of stimulus presentation and initial stimulus encoding, but only differ in the contingency between the prime that was memorized and the probe that was presented. This ensures that any modulation of the neural response to probes in the Memorized condition relative to the Discarded condition (as observed in both the univariate and multivariate analyses) is unequivocally engendered by the content of VWM.

From the MVPA results of Experiment 1 alone, it is not clear whether visual input that matched the to-be-memorized stimulus was enhanced, or whether visual input that matched the to-be-discarded stimulus was inhibited. This follows from the fact that decoding accuracies for visual input that was unrelated to the memory task did not differ from the two other conditions. Two arguments, however, provide support for an enhancement of the neural response to visual input that matches the to-be-memorized shape category, as opposed to an inhibition of the to-be-discarded shape category. First, the univariate analyses of Experiment 1 show that the probe elicits a stronger BOLD response when it is of the same shape category as the to-be-memorized shape, compared to when it is drawn from either an unused shape category or the shape category of to-be-discarded shape. There was no difference in BOLD response, however, to probes matching either the unused shape category or the shape category of the to-be-discarded shape. As such, visual input that matches the to-be-memorized stimulus elicits a stronger BOLD response, whereas no inhibition is observed for visual input matching the to-be-discarded stimulus. Second, the behavioral results in the present study (Supplementary Materials S2), but also behavioral results using different stimuli (Gayet et al., 2013), and different paradigms (e.g., Gayet et al., 2015; Olivers et al., 2006), consistently reveal an enhancement of the behavioral response to visual input that matches the to-be-memorized stimulus, but do not demonstrate inhibition of visual input that matches the to-be-discarded stimulus. In line with these findings, we conclude that patterns of neural activity between stimulus categories were more distinct (and therefore contained more information) when the visual representation elicited by VWM and by visual input were in accordance.

The finding that patterns of activity representing different shape categories are more distinct when visual input and VWM are in accordance, supports the idea that visual representations draw upon a common neural substrate, irrespective of whether they have a mnemonic or sensory origin (e.g., Chapter 6; Albers et al., 2013; Rigall, & Postle, 2012; Harrison, & Tong., 2009). Alternatively, it could be argued that the stronger BOLD response for visual input matching the content of VWM (univariate analyses of Experiment 1) cause a stronger signal, and therefore a stronger signal to noise ratio, which facilitates the retrieval of specific patterns of activity in the data. A stronger BOLD response was associated with better classifier performance in decoding the content of VWM in some studies (e.g., Christophel et al., 2012) but not in others (e.g., Supèr, Spekreijse, & Lamme, 2001; Harrison, & Tong, 2009). In either case, the higher decoding accuracies reflect the informative value of the activation pattern. As such, not only our classifiers, but the visual system as well could (in principle) capitalize on the contingency between visual input and VWM representations to better dissociate between stimulus categories. In sum, the present findings only provide indirect evidence for a common neural substrates for representations elicited by VWM and representations elicited by visual input. However, the present findings do demonstrate that neural representations are both quantitatively and (maybe indirectly) qualitatively enhanced by the contingency between visual input and the concurrent content of VWM.

In contrast to Experiment 1, Experiment 2 showed that the neural response to visual input is not modulated by the content of VWM when this visual input is interocularly suppressed. Despite having more trials per condition and more participants in Experiment 2, and despite directly targeting the areas that showed a significant VWM modulation in Experiment 1, Bayesian analyses confirmed that the neural response to visual input was not modulated by the concurrent content of VWM. One interpretation for the discrepancy between the findings of Experiment 1 and Experiment 2 is that visual awareness is a critical requirement for the visual system to decide whether two stimuli are from the same perceptual category or not. That is, the comparison of multiple visual features is tightly linked to the integration of multiple visual features, a cognitive function that is believed to require visual awareness (e.g., Fahrenfort, Lamme, 2012; Treisman, 2003; Baars, 2002; Tononi, & Edelman, 1998; but, see Mudrik, Faivre, & Koch, 2014). An alternative interpretation for the null result in Experiment 2 is that interocular suppression constrained the processing of the visual input to early visual areas. Indeed, most cells in early visual areas respond to stimulation of the suppressed eye as well as the dominant eye (Logothetis, 1998), whereas higher visual processing areas, such as IT, LO, FFA, and PPA follow mostly (but not exclusively; Fang and He, 2005; Jiang and He, 2006; Sterzer, Haynes, & Rees, 2008; for a review, see Sterzer, Stein, Ludwig, Rothkirch, & Hesselmann, 2014) the dominant percept (Tong, Nakayama, Vaughan, & Kanwisher, 1998). In Experiment 1, no differential activation was found in early visual areas between visual input that matched and visual input that mismatched the content of VWM. As such, it is likely that the content of VWM affected processing of visual input further upstream in the visual hierarchy, such as in LOC, an important area for the visual processing (e.g., Grill-Spector et al., 2001) and working memory maintenance of objects (e.g., Xu, & Chun, 2006), or IPS, an important area for working memory maintenance of complex stimuli (e.g., Christophel et al., 2012; Xu, & Chun, 2006; Song, & Jiang, 2006; but, see Sligte, van Moorselaar, & Vandenbroucke, 2013). In line with this, in Experiment 1, the lateral occipital component of the neural response to the probe that was modulated by the content of VWM was situated in a region that corresponds well with LOC. Similarly, the inferior/superior parietal component was situated in a region that corresponded well with IPS. Two observations support this idea that VWM affects concurrent visual processing in higher level visual areas. First, many behavioral experiments have demonstrated that the influence of VWM on visual input is not restricted by retinotopic overlap. That is, the distances between the to-be-memorized stimulus and the probe presented during the delay interval typically span multiple degrees of visual angle (e.g., Experiment 1; Gayet et al., 2013; Olivers et al., 2006; Soto et al., 2005). Considering receptive field sizes in human visual cortex (e.g., estimated by means of population receptive field mapping; Harvey, & Dumoulin, 2011) these effects are highly unlikely to originate prior to V4/V5/LO in the visual processing hierarchy. Second, in Experiment 1 as well as in the behavioral studies mentioned above, the representations in VWM were never identical to the concurrent visual input. Rather, they were congruent in the sense that they were drawn from the same stimulus category. Interestingly, LOC has been proposed to represent visual objects categorically. That is, MVPA cross-classification revealed that LOC was sensitive to differences between stimulus categories, while being relatively insensitive to such non-categorical differences as viewpoint, object size and, most importantly, stimulus exemplars within categories (Eger, Ashburner, Haynes, Dolan, & Rees, 2008; Grill-Spector, Kushnir, Edelman, Avidan, Itzchak, & Malach, 1999). Together, these observations seem to preclude involvement of early visual areas in the modulation of visual input by the content of VWM (at least with the present type of stimuli).

Thus, we conclude that the potency of representations maintained in VWM (1) arises in higher level visual areas and (2) therefore does not allow for modulating visual input that is still fully interocularly suppressed (e.g., by means of continuous flash suppression). This raises the question how the contents of VWM modulates the time it takes for initially suppressed stimuli to be released from interocular suppression (e.g., Chapter 6; Supplementary Materials S2 of the present study; van Moorselaar et al., 2015; Pan et al., 2014; Gayet et al., 2013). Crucial control conditions have demonstrated that the content of VWM indeed modulates suppression durations, rather than the time it takes for participants to respond once a probe is released from suppression (Gayet et al., 2013; for further elaboration on this control condition, see Gayet et al., 2014b; Gayet, Paffen, Belopolsky, Theeuwes, & Van der Stigchel, 2016). As such, VWM does reduce interocular suppression durations, yet does not seem to occur under conditions of full suppression. Based on this paradoxical observation, we propose that VWM has the potency to modulate the release from interocular suppression only when the suppressed visual input is just below the threshold of visual awareness. In the present case, the interocular suppression did not allow the visual input to perspire through higher level visual areas where it could interact with the content of VWM. Paradigms that measure transitions from suppression to visual awareness, however, might allow for transitory periods in which part of the interocularly suppressed information does perspire to higher level visual areas, without necessarily eliciting a clear reportable percept (Zadbood, Lee, & Blake, 2011;

Yang, & Blake, 2012). The idea that interocular suppression (and maybe access to awareness in general) is particularly sensitive to modulations around the threshold of visual awareness finds support in recent findings. For instance, it has been demonstrated that high-level perceptual priming under CFS only occurred under conditions of partial awareness (Gelbard-Sagiv et al., 2016). Partial awareness refers to the situation in which only some features can be reported, such as the color or location of the suppressed stimulus, but not others, such as face identity. Similarly, attentional modulation of binocular rivalry was observed to affect perceptual dominance only around the time of the perceptual transition between the two percepts (Dieter et al., 2015). It is unclear whether partial awareness specifically, or suppression depth more generally, determines whether top-down modulations of interocular suppression is possible (for more discussion, see Peremen & Lamy, 2014; Gayet et al., 2014b; Stein, et al., 2011). However, it must be concluded from the present data that the modulation of interocular suppression by the content of VWM, one of the most robust top-down modulations of interocular suppression, is abolished under conditions of full interocular suppression.

Conclusion

Our results demonstrate that the neural response to visual input is enhanced when it matches the content of VWM. Considering that this effect was not restricted by retinotopic overlap (Experiment 1) and was not observed when visual input was interocularly suppressed (Experiment 2), we conclude that the interaction between visual input and VWM originates in high level visual areas. The data suggest that the interaction between representations maintained in VWM and visual representations elicited by visual input originates in the lateral occipital cortex (in regions corresponding to LOC) and the inferior and superior parietal lobe (corresponding with IPS). We hypothesize that a frontal region, the pars triangularis of the right inferior frontal gyrus, is involved in shielding the working memory representation from the concurrent visual input. The present results support the hypothesis that a common neural substrate underlies the processing of visual representations, irrespective of whether their origin is retinal or mnemonic. Consequently, when the content of VWM causes an elevated activity level in specific neural populations, visual input impinging upon this same neural circuitry is both qualitatively and quantitatively enhanced, thereby affecting perception and behavior.

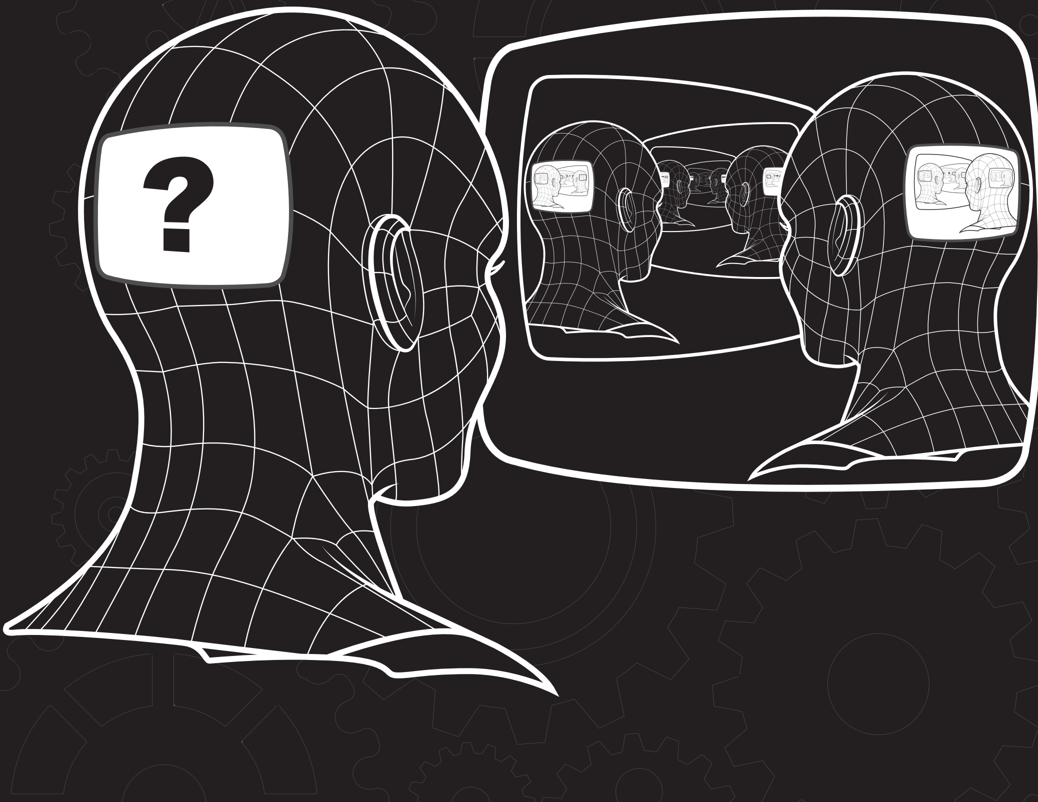

Epilogue

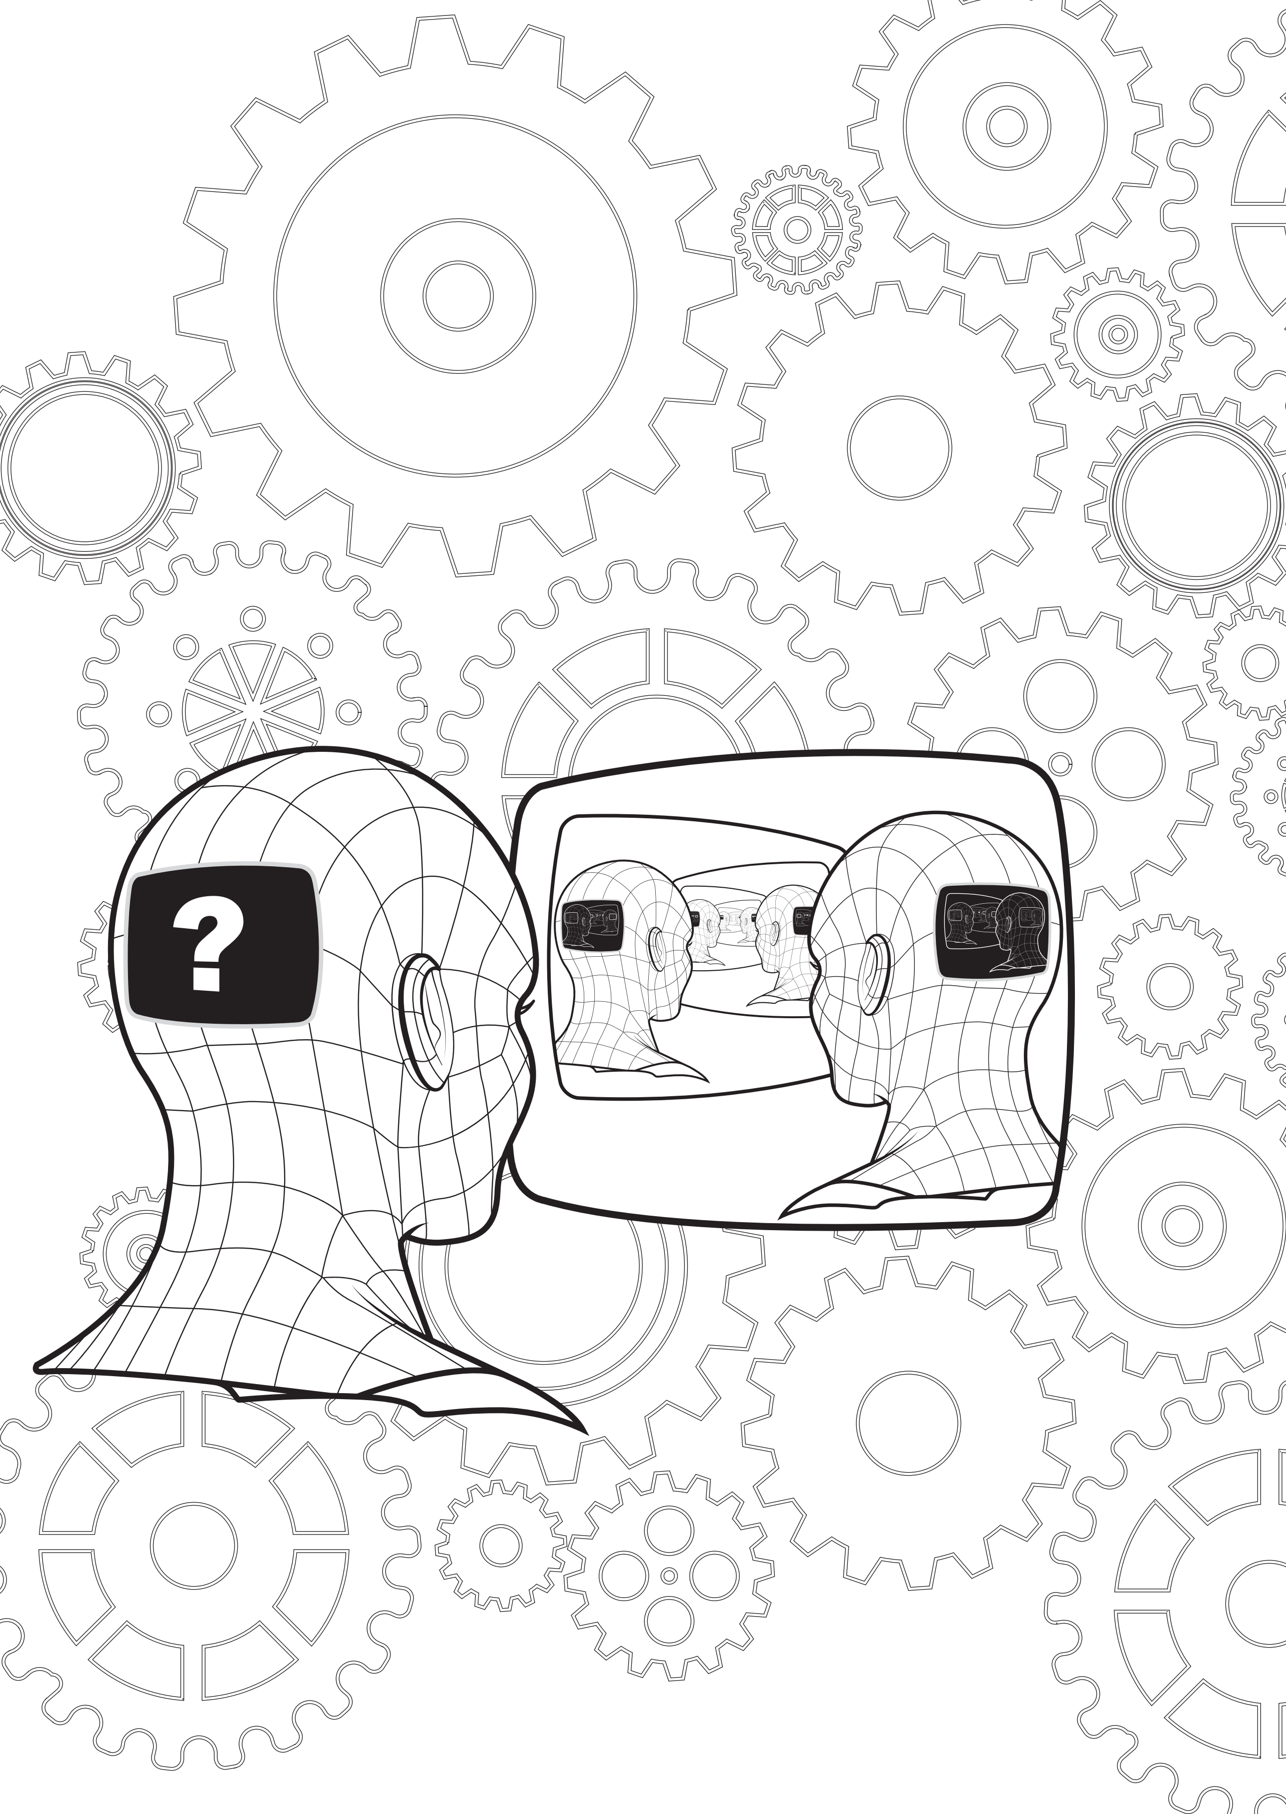

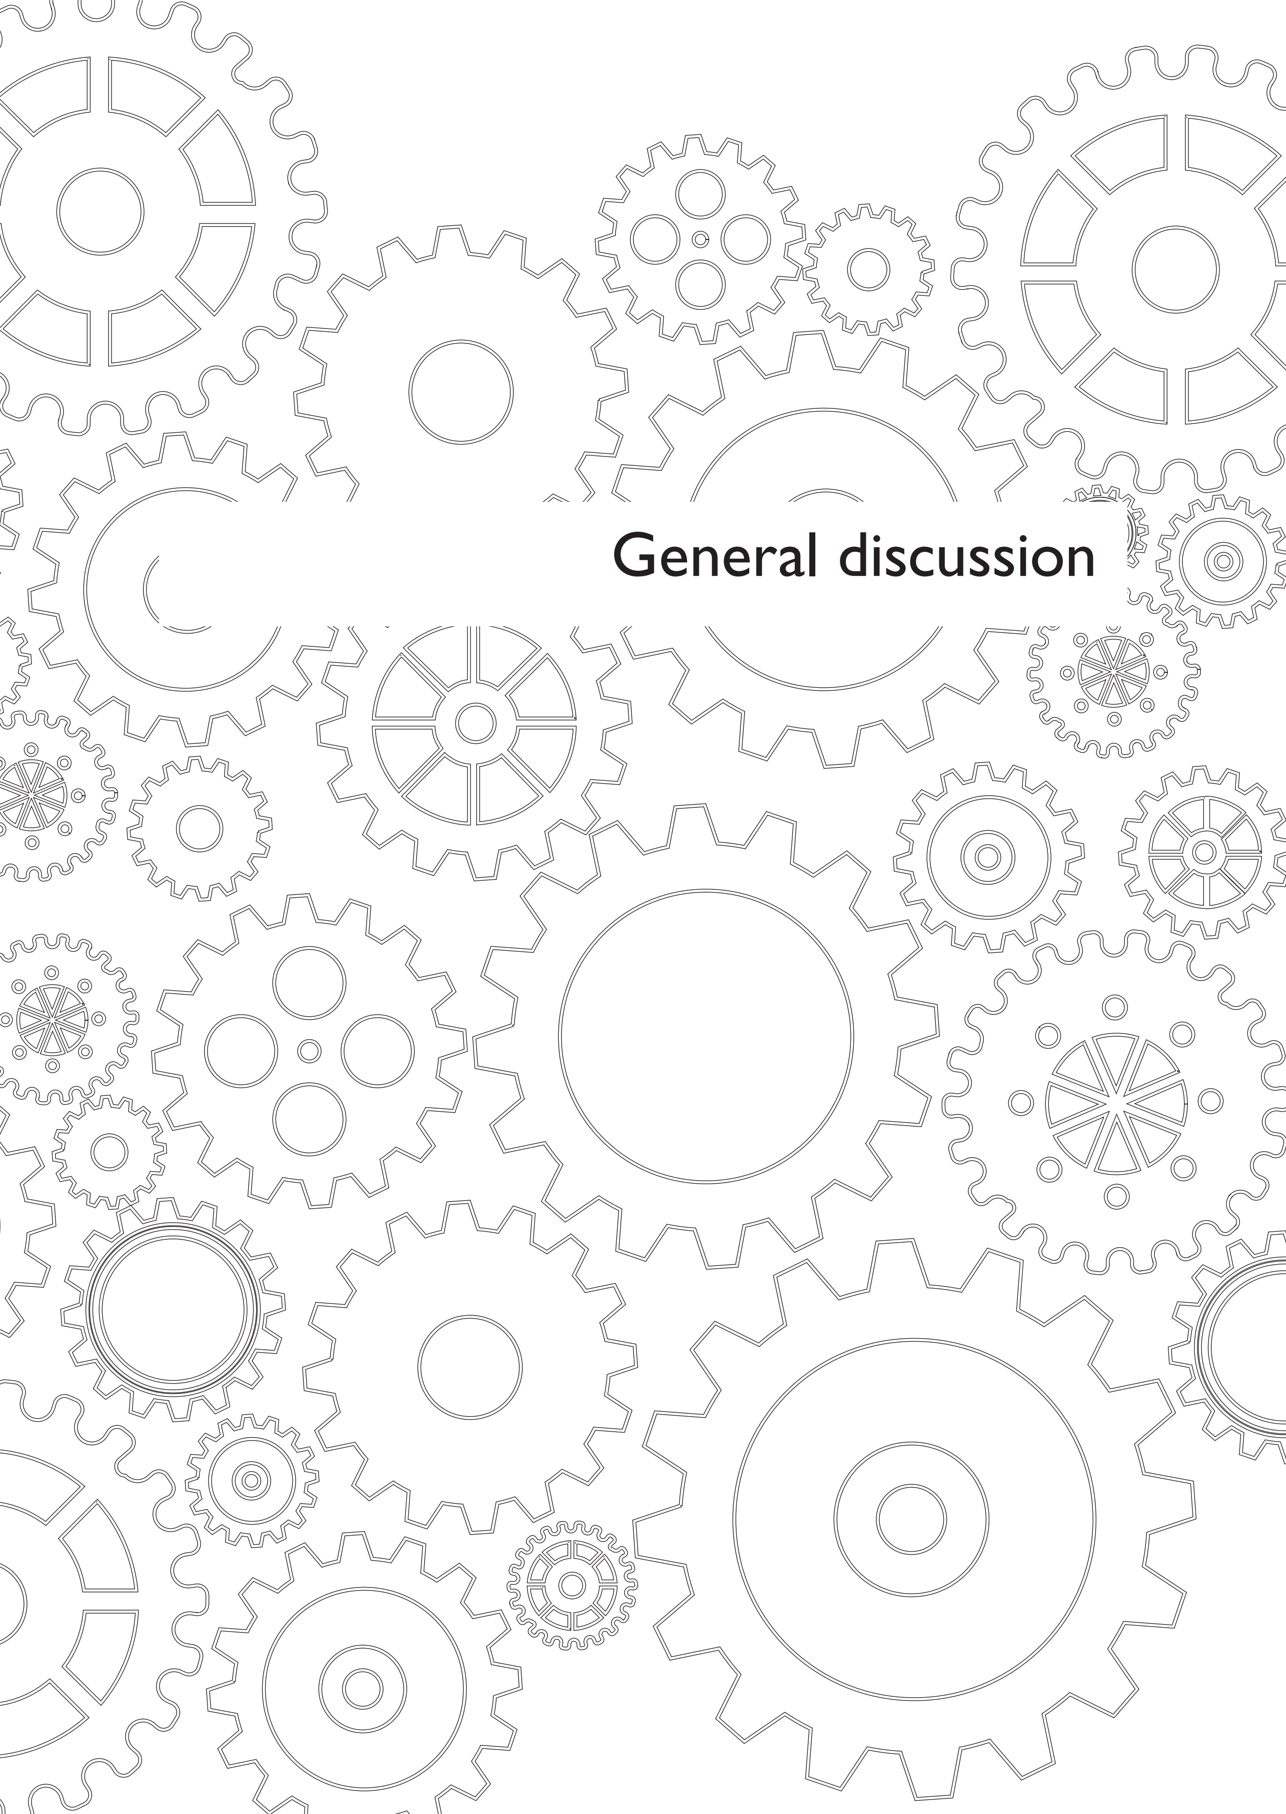

General discussion

In this final chapter, I will first discuss the theoretical insights that can be deduced from the different experimental chapters that constitute my dissertation. Subsequently, I will propose a unified mechanism that capacitates all observed phenomena, and discuss its neural implementation. Finally, I will position the main conclusions of my dissertation within the field of consciousness research.

A theoretical perspective

Considering that we are continuously presented with vast amounts of sensory input, and that only part of it will lead to a conscious experience, it is crucial for adaptive functioning that behaviorally relevant information gains prioritized access to consciousness. In line with this, we have demonstrated that the behavioral relevance of visual information (as dictated by the behavioral goals of the observer) affect the processing of visual information even before it reaches consciousness. Such behavioral goals have the potency to determine whether or not information that is inaccessible to consciousness affects behavior (chapter 1). In addition, the behavioral relevance of visual information codetermines when said visual information will be accessible to consciousness (chapters 2, 4, 5, 6, and 7). Contrary to traditional belief, non-conscious visual processing is thus not only comprised of automated perceptual reflexes. Rather, the current behavioral goals elicit an a priori state of the observer, which modulates the (perceptual, behavioral, and neural) response to non-conscious visual information. Throughout my dissertation, however, the rather rigid nature of non-conscious processing has also been demonstrated. That is, although non-conscious processing of visual information is sensitive to the behavioral goals of the observer, once the appropriate behavior is determined, non-conscious processing no longer deviates from the planned (perceptual, behavioral, and neural) response. To illustrate this paradoxical conclusion, I provide an overview below of all observations from the experimental chapters that relate to these two sides of the same coin.

In chapter 1, we demonstrated that non-conscious visual information only affected behavior when the consciously accessible context provided an incentive to do so (i.e., when it was behaviorally relevant). Specifically, invisible arrow cues affected behavior in a target detection task, but only if visible arrow cues were predictive of the subsequent target location. This demonstrates that the non-conscious information did not affect behavior reflexively. Rather, it depended on whether or not it was useful for the behavioral goals of the observer. On the flipside, however, the non-conscious visual information itself was irrelevant. That is, the non-conscious arrow cues per se were not predictive of the subsequent target location, yet they still affected behavior. From this we concluded that the behavioral goals of the observer (i.e., responding as fast as possible to a target) elicits an a priori mental state (i.e., utilize the arrow cues), which determines the response to subsequent non-conscious visual information. Once this a priori mental state is in place, it guides the response to non-conscious visual information in a reflexive manner.

This contrast, between what is determined a priori, and how this affects subsequent processing of visual information, is even more striking in Chapter 4. Here, we demonstrated that visual input that was previously associated with electric shocks reached consciousness earlier in time than visual input that was never associated with electric shocks. Thus, the visual system was able to assess which of two comparable stimuli signaled an imminent threat based on recent history, such as to codetermine the content of consciousness. Importantly, we argue that it is the behavioral relevance, and not the negative valence, that influenced access to consciousness. This is made apparent by a related study in which conscious perception of a bistable stimulus was measured, while associating one of the two percepts with a monetary loss (Wilbertz, van Slooten, & Sterzer, 2014). In contrast to our study, it would be behaviorally relevant for participants in this study *not* to consciously perceive the visual input that was associated with negative valence (i.e., a monetary punishment). Accordingly, in the study of Wilbertz et al. the percept that was associated with a monetary loss was less likely to reach consciousness. Our findings and those of Wilbert et al. taken together demonstrate that the prioritization of visual information for reaching consciousness is dependent upon the (conscious) context, and therefore seems highly flexible. Again, there is a flipside to the findings of chapter 4. Indeed, after the fear conditioning procedure, and prior to measuring access to consciousness, we informed the participants that they would receive no more electric shocks. Hence, the visual information that was previously associated with electric shocks did not signal an imminent threat, nor was it more relevant than the visual information that was never associated with electric shocks. As such, the behavioral relevance of the visual information (i.e., whether or not it was associated with shocks) capacitated an a priori state of the observer (e.g., red signals threat). This a priori state subsequently modulated conscious access in a reflexive manner, that was insensitive to the behavioral goals of the observer.

In chapters 2, 5, 6, and 7, we demonstrated that visual information gains privileged access to consciousness when it contains a feature that is relevant for a concurrent task. Specifically, observers were required to memorize a visual feature for a subsequent recall task. Visual information that was presented during the retention interval reached consciousness earlier in time (chapters 5, 6, and 7) and for longer periods (chapter 2) when it matched the concurrently memorized color. Thus, access to awareness of a specific visual feature (say, the color red) was determined by its relevance in a concurrent task. On the flipside, however, the color of the stimulus presented during the retention interval was not relevant to the participant. The participant only had to indicate the location (chapters 5, 6, and 7) or orientation (chapter 2) of this stimulus. If anything, when memorizing a specific hue of red for subsequent recall, recall performance might be hindered more by perceiving a slightly different color red than by perceiving green or blue. As such, the manipulation of relevance affected access to consciousness. However, the eventual visual input that gained faster access to consciousness because of this manipulation was not relevant at all for the behavioral goals of the observer. Again, the state of the observer was determined a priori (i.e., by memorizing a particular color) and access to awareness was modulated accordingly, in an automated manner.

In chapter 6, we used a computational modeling approach to investigate the perceptual processes leading up to the moment in time at which initially non-conscious visual information elicits a conscious experience. Sequential sampling models (for reviews, see Forstmann, Ratcliff, & Wagenmakers, 2016; Ratcliff & Smith, 2004) provide the means to infer, from response times and response choices, the process of evidence accumulation over time that eventually leads to a perceptual decision. We compared two models in their potency to explain the observed behavioral data (i.e., faster access to consciousness for visual information that matches a concurrently memorized color). One model was aimed at testing the hypothesis that visual information that matches a concurrently memorized color would have an accelerated rate of perceptual evidence accumulation. The other model was aimed at testing the hypothesis that there is an a priori bias towards visual information that matches a concurrently memorized color. Model comparisons revealed that the data was best explained by an a priori bias towards visual information that matches a concurrently memorized color. In contrast, there was no evidence for a faster accumulation rate of perceptual evidence. Thus, task-relevant colors emerged into consciousness earlier in time due to an initial bias in the state of the observer.

In sum, we consistently observe the same pattern of findings across different experimental methods (i.e., priming under traditional flash suppression, binocular rivalry, breaking continuous flash suppression, and computational modeling) and different experimental manipulations (i.e., cue usefulness, visual working memory, and classical fear conditioning): the a priori state of the observer, which depends on the behavioral goals of the observer, affects subsequent processing of visual information that is not (yet) accessible to consciousness in a reflexive manner. But how is this a priori difference in the state of the observer implemented in the visual system?

Towards a unified framework

In chapter 6 we provide a framework, which we refer to as the *Pre-activation* hypothesis. This hypothesis entails that the experimental manipulation of task-relevance elevates activity levels in neural populations that represent the relevant visual feature. As a result of these elevated activity levels, the effective threshold for reaching consciousness is lowered for visual input that matches the relevant visual feature.

Before further elaborating on the exact neural implementation entailed by this framework, I want to make the case that it is a general framework, that does not only hold for the manipulation of visual working memory that we used in chapter 6 and 7, in which we test this framework, but also holds for other manipulations of behavioral relevance. Within this dissertation, classical fear conditioning (chapter 4) and visual working memory (chapter 5, 6, 7) constitute two drastically different manipulations of behavioral relevance. Comparing these two experimental manipulations is valuable for understanding the influence of behavioral relevance per se on the propensity of visual information to access awareness. For this purpose, I provide a comparison (below) between the findings of chapter 4 (classical fear conditioning) and 5 (visual working memory), summarizing the neural and behavioral differences and similarities between the two manipulations of behavioral relevance.

From a neural perspective, we hypothesized throughout this dissertation that a stronger neural response in visual areas causes the prioritized access to consciousness of behaviorally relevant visual information in a b-CFS task. In light of the pre-activation account discussed above, we hypothesize that behavioral relevance manifests itself as a priori elevated activity level of the neural circuitry representing the task-relevant visual feature. Consequently, interocular competition in visual processing areas is biased towards visual information that matches the behaviorally relevant visual feature. From this line of reasoning, it follows that the modulation of the interocular competition itself does not differ between the different manipulations of behavioral relevance (i.e., visual working memory, or classical fear conditioning) as it is capacitated by increased activity levels in visual processing areas. It is expected, however, that the modulation of activity levels in visual processing areas is achieved differently depending on the experimental manipulation of task relevance. We will now discuss this for the fear conditioning manipulation and the visual working memory manipulation separately.

In case of the fear conditioning manipulation, we proposed in chapter 4 that the amygdala might play a role in modulating the cortical response to threatening stimuli that are not (yet) accessible to consciousness. In fact, for consciously accessible stimuli, this has been observed: the amygdala mediates the stronger neural response in visual processing areas to fear conditioned stimuli (Lim, Padmala, & Pessoa 2009). Moreover, amygdala activation has also been observed for emotional stimuli under continuous flash suppression (Jiang & He, 2006). As such, amygdala activation might also mediate the response to fear-conditioned visual information in the absence of consciousness. Currently, however, this has not been verified empirically.

In case of the visual working memory manipulation, we proposed (in chapter 5, 6, and 7) a more direct modulation of cortical activity by the behavioral relevance of the visual input. In this line of reasoning, maintaining a visual feature in visual working memory increases activity levels in the neural population representing said feature, such that concurrent visual input that draws upon the same neural population (i.e., when it matches the relevant feature) elicits a higher neural response in visual processing areas than visual input that draws upon a different neural population. This account was motivated by our finding that only the task relevant feature (e.g., only the color of a colored shape) of a to-be-memorized stimulus affected concurrent conscious access of matching visual information (chapter 2, and 5), and the finding that only the task relevant feature of a to-be-memorized stimulus is represented in visual processing areas (Serences, Ester, Vogel, & Awh, 2009). The idea that the content of visual working memory affects conscious access in this way, relies upon the assumption that visual representations elicited by visual working memory, and visual representations elicited by visual input draw upon the same underlying neural circuitry. Imaging studies using multivariate pattern analysis provide support for this assumption (Albers, Kok, Toni, Dijkerman, & de Lange, 2013; Harrison, &

Tong, 2009; Rigall, & Postle, 2012), and corroborating evidence has come from our computational modelling study in chapter 6. In line with our hypothesis, we demonstrated in chapter 7 that (consciously accessible) visual input indeed elicits a stronger neural response in visual processing areas if it matches the concurrent content of visual working memory.

Taken together, the abovementioned propositions suggest that our two manipulations of behavioral relevance were similar in how they eventually drive the competition for consciousness by modulating the responsiveness of neural populations representing the relevant visual feature. In contrast, both manipulations differ in terms of the processes leading up to this modulation. This can be further substantiated by contemplating differences and similarities between our two experimental manipulations from a behavioral perspective. Following both manipulations, targets of a behaviorally relevant color reached conscious access faster than targets of a behaviorally irrelevant color. Also, in both cases, a monocular control condition confirmed that the manipulation affected the time it took for the ocular dominance to reverse, rather than the time it took for observers to respond *after* the ocular dominance reversed. Finally, in both cases the difference in response times between behaviorally relevant and irrelevant targets correlated with observers' mean response speed. That is, participants with slower response times (i.e., longer suppression durations) showed a larger effect of behavioral relevance on conscious access. These similarities substantiate the idea that the interocular competition for consciousness itself is driven by similar processes, irrespective of the manipulation of behavioral relevance.

There are also behavioral differences between the two experimental manipulations. Most importantly, the fear-conditioning manipulation affected access to consciousness in a subsequent breaking continuous flash suppression (b-CFS) task while the manipulation was no longer applied. Additional analyses revealed that the influence of the fear conditioning procedure on the b-CFS task persisted for about 10 to 15 minutes after termination of the fear conditioning procedure. Manipulating the content of visual working memory, however, only affects access to consciousness in a b-CFS task during the retention interval. That is, the influence of the content of visual working memory on conscious access vanishes directly after the recall task (Pan, Lin, Zhao, & Soto, 2014). This difference substantiates the idea that different manipulations of behavioral relevance initially draw upon different processes. The way by which different manipulations of behavioral relevance eventually influence access to consciousness, however, might ultimately boil down to a single mechanism: pre-activating the relevant feature representation in visual processing areas.

Neural implementation

We have coined pre-activation of specific neural populations as a means for our visual system to prioritize visual input that is relevant to our behavioral goals. For a better functional understanding of how behavioral relevance impacts the selection of visual information for consciousness, it is important to investigate how this modulation is implemented at the neural level. That is, the functional level at which behavioral relevance can influence access to consciousness, depends on the extent to which visual input is processed non-consciously. For this purpose, we set out to pinpoint the exact loci along the visual processing hierarchy in which behavioral relevance can modulate (non-conscious) visual input. In this section, we will focus on visual working memory as a manipulation of behavioral relevance, and we focus on interocular suppression (more specifically, continuous flash suppression; Tsuchiya, & Koch, 2005) as a means to manipulate consciousness.

The first paper that I published during my PhD project demonstrated that visual input matching the content of visual working memory is released from continuous flash suppression earlier in time than mismatching information. Until recently, I was rather confident that this finding ultimately boiled down to neural interactions in the human primary (or at least early) visual cortex. I started to regard my participants as a source of noise that happened to surround V1. The main reasoning was that any manipulation affecting interocular competition should occur somewhere in the visual processing hierarchy where eye-selective information is still retained. Most cells in early visual areas (80% in V1/V2 and 60% in V4/V5) respond to stimulation of either eye, irrespective of which eye's input dominates perception (Logothetis, 1998). Higher processing areas, such as IT, LOC, FFA, and PPA, however, predominantly reflect the dominant percept (Tong, Nakayama, Vaughan, & Kanwisher, 1998). In addition, simple visual stimuli, such as oriented gratings, are represented in V1 both when they are elicited by retinal input (Haynes, & Rees, 2006; Harrison, & Tong, 2009) or when they are maintained in visual working memory (Harrison, & Tong, 2009; Serences et al., 2009). In light of all these considerations, early visual cortex seemed the most likely locus for visual working memory to affect visual input rendered invisible by continuous flash suppression.

Considering the abovementioned assumption, I could think of two ways in which the content of visual working memory could affect suppression durations (i.e., the time it takes for visual input to reach consciousness). The first possibility was that high-level stimulus representations in visual working memory (say, the color category red) were fed back to low level visual areas were the interocular competition was resolved (feedback framework). The second possibility was that visual working memory was stored as visual representations in the same low-level visual areas where the interocular competition was resolved (pre-activation framework). In order to differentiate between these two frameworks, we implemented them as computational models, and compared the potency of both models to explain the observed behavioral data (chapter 6). As mentioned earlier, the data was best explained by the implementation of the pre-activation framework. This framework builds upon the assumption that visual representations elicited by visual working memory and visual representations elicited by visual input draw upon the same underlying neural substrate. Although the model is not informative about *where* said neural substrate should be located in the brain, I implicitly assumed that this shared neural substrate would comprise of low level visual areas. There are two properties of the experimental paradigm, however, that are hard to reconcile with modulations of early visual areas. The first issue is that the effect of visual working memory on suppression durations is categorical. That is, the target stimulus, presented during the delay interval, is never identical to the stimulus that is maintained in visual working memory. Rather, they are drawn from the same color (chapters 5, and 6), shape (chapter 7), or face (Pan et al., 2014) *category*. This issue could be dismissed, however, by arguing that sufficiently similar stimuli could activate overlapping neural populations (some evidence for this can be found in a priming study by Lupyan, & Ward, 2013). The second issue, is that the target stimulus was never presented at the same retinal location as the to-be-memorized stimulus. In the majority of studies that employ this paradigm (e.g., chapters 5, 6, and 7), the stimuli were separated by more than a degree of visual angle, which exceeds the classical receptive field sizes of foveal V1-V2 cells (e.g., Harvey & Dumoulin, 2011). Taken together, the influence of visual working memory on suppression durations is not restricted to receptive field sizes of low level visual areas, is potentiated by categorical stimulus representations, and also applies to high-level stimuli such as faces. From this, it can be concluded that the processing areas that capacitate this effect might be further upstream in the visual processing hierarchy than I initially thought.

In chapter 7, we used an event-related fMRI approach to measure the neural response to consciously accessible visual stimuli (rectangles, ellipses and triangles), depending on their contingency with the concurrent content of visual working memory. The data revealed a stronger neural response to visual stimuli that matched the content of visual working memory, compared to stimuli that mismatched the content of visual working memory. This modulation of neural activity was observed in inferior lateral occipital areas, including brain areas corresponding to the lateral occipital complex (LOC), and in superior parietal areas, including areas corresponding to the intra-parietal sulcus (IPS). In contrast, no modulation of neural activity was observed in early visual areas. The finding that the content of visual working memory enhances the neural response to matching visual input in LOC can be related to two main properties of LOC. First, LOC has been proposed to represent visual objects categorically. That is, imaging studies have shown that LOC was sensitive to differences between stimulus categories, while being relatively insensitive to perceptual differences of exemplars within categories (Eger, Ashburner, Haynes, Dolan, & Rees, 2008; Grill-Spector, Kushnir, Edelman, Avidan, Itzchak, & Malach, 1999). Second, receptive field sizes in LOC allow for encompassing the retinal distances between the to-be-memorized stimulus and the different target locations, as typically used in the present experimental paradigm. The increased activity observed in IPS, in turn, can be related to its involvement in the working memory maintenance of complex visual stimuli (e.g., Christophel, Hebart, & Haynes, 2012; Xu, & Chun, 2006; Song, & Jiang, 2006). Nonetheless, the finding that visual working memory modulates the neural response to visual input in such high level processing areas as LOC and IPS, seems incompatible with the finding that visual working memory affects the processing of interocularly suppressed visual input. That is, the neural response to interocularly suppressed visual information is generally confined to lower level visual areas. In order to investigate these seemingly incompatible findings, we conducted a second event-related fMRI experiment in which the target stimuli, presented during the delay interval, were rendered invisible by means of continuous flash suppression. The results revealed that, when visual input is fully interocularly suppressed, its neural response is not modulated by the content of visual working memory. The neural response to non-conscious visual input in general (Dehaene, Naccache, Cohen, Le Bihan, Mangin, Poline, & Riviere, 2001), and to interocularly suppressed visual input more specifically (e.g., Tong et al., 1998; for a review see Sterzer, Stein, Ludwig, Rothkirch, & Hesselmann, 2014), is believed to be mostly restricted to early visual areas. As such, these findings corroborate the results from the first experiment in showing that the modulation of visual input by the content of visual working memory originates in higher level visual areas.

Reconciliation

We are now left with the five following pieces of a puzzle. First, visual working memory affects the neural response to concurrent visual input in high level visual processing areas. Second, visual working memory does not affect the neural response to interocularly suppressed visual input. Third, visual working memory causes reduced interocular suppression durations of matching visual input (in a b-CFS task). Fourth, control conditions confirm that visual working memory indeed modulates interocular suppression durations, rather than processing differences after the visual input is released from suppression. Fifth, the shorter suppression durations for visual input matching the content of visual working memory was driven by an a priori bias (i.e., pre-activation). Importantly, all of these findings were statistically extremely robust.

How can we reconcile the finding that visual working memory does not modulate the neural response to fully suppressed visual input, yet does modulate its access to consciousness, as observed through behavioral measures? The most likely explanation stems from the intrinsic difference between shielding visual input from consciousness (e.g., by means of continuous flash suppression) and by measuring the time it takes for visual input to reach conscious access. In the former case, the suppression of the visual input is kept relatively constant, whereas in the latter case, the visual input transits from full suppression to full visibility. This transitory period might be of particular interest for the present purpose as it is specific to interocular suppression conditions, and would therefore not emerge in the monocular no-suppression control conditions (as referred to in the fourth point of the previous paragraph). Recent observations have shown that, under conditions of prolonged continuous flash suppression, periods of partial awareness might arise, in which some stimulus properties (such as color) are available to consciousness, whereas other are not (such as orientation; Yang, & Blake, 2012; Zadbood, Lee, & Blake, 2011).

These observations indicate that suppression from consciousness by continuous flash suppression is not an all or none phenomenon. Rather, prolonged periods of continuous flash suppression might enable distinct stages of suppression strengths, which, in turn, might vary in their susceptibility to experimental manipulations. Within this frame of reference, an important mediating factor might be the point in time (i.e., the strength of suppression) at which an experimental manipulation affected suppression durations in a b-CFS experiment. Recently, stimuli under continuous flash suppression were shown to only elicit high-level behavioral priming effects under conditions of partial consciousness (Gelbard-Sagiv, Faivre, Mudrik, & Koch, 2016). Similarly, a recent study demonstrated that manipulations of attention only affected dominance durations in a binocular rivalry paradigm around the time of perceptual transitions (Dieter, Melnick, & Tadin, 2015). Although interocular competition depends on low-level inhibitory interactions, high-level excitatory modulations also affect suppression durations (Tong, Meng, & Blake, 2006). Indeed, interocular competition is modulated at different levels throughout the visual processing hierarchy (Blake, & Logothetis, 2002). We propose that (local) reductions in suppression strength (which by definition occur in a b-CFS task) allow for higher-level visual areas to modulate suppression durations of initially fully suppressed visual input. In the present case (i.e., chapters, 5, 6, and 7), a priori elevated activity levels in LOC and/or IPS representing the relevant feature, enhance the response to initially suppressed visual stimuli. This enhancement, however, occurs only once a period of reduced suppressive strength is attained, such that a sufficient portion of the visual input perspires to these higher level visual areas. In these higher level visual areas, the content of visual working memory has the potency to modulate the visual input. In case the visual input remains under conditions of full interocular suppression (and therefore its neural response is restricted to early visual areas), however, the visual input does not interact with the concurrent content of visual working memory.

Implications for the field

The rationale behind the b-CFS paradigm is that, if one acknowledges that some manipulation influenced *when* a stimulus reaches conscious, it is implied that said manipulation affected the processing of the stimulus *before* it reached consciousness. The inhomogeneity of continuous flash suppression, eliciting periods of partial consciousness, has been proposed as an explanation for b-CFS findings before (chapter 3; Peremen, & Lamy, 2014; Stein, Hebart, & Sterzer, 2011). Because of this, it has been argued that, while b-CFS provides the means to measure *access to consciousness*, it does not allow for making inferences about *non-conscious processing* per se (Stein, & Sterzer, 2014). I will make the case that this inhomogeniety does pose a problem, but a different one. For this, it is important to first have a look at the nature of the manipulations that produce an effect in a b-CFS paradigm.

In the review article presented in chapter 3, we made a distinction between two types of experimental manipulations that are used in the b-CFS paradigm to modulate the time it takes for stimuli to reach consciousness. The first type comprises manipulations of stimulus *content* (i.e., comparing suppression durations between different classes of stimuli), and the second type comprises manipulations of stimulus *context* (comparing suppression durations for the same stimulus between different contextual settings). The most robust manipulation of stimulus *content* in the b-CFS literature is the reduction in suppression durations for upright compared to inverted faces (e.g., Jiang, Costello, & He, 2007; Stein et al., 2011). One of the most extensively studied manipulations of stimulus *context*, is arguably the modulation of suppression duration depending on the contingency between the target and the concurrent content of visual working memory (originally observed in chapter 5). Indeed, this effect has been replicated in at least 11 different experiments, over 5 different studies using color, shape and face stimuli (chapters 5, 6, and 7; Moorselaar, Olivers, & Theeuwes, 2015; Pan, Lin, Zhao, & Soto, 2014). Aside from its robustness, this effect is also characterized by its low demands on complex cognitive functions. The relative simplicity of this effect stands in contrast to other effects observed with the same paradigm, which in some cases require complex cognitive functions, including semantic extraction, arithmetic problem solving, and scene integration (reviewed in chapter 3; Hassin, 2013; Hesselmann, & Moors, 2015). It is striking that this robust phenomenon, for which simple cognitive functions suffice, apparently only arises under conditions of shallow suppression (as occurs in the b-CFS paradigm) but not under full suppression (as in our second fMRI experiment). This observation brings us to the problem that I mentioned above: different manipulations in b-CFS paradigm, require different cognitive processing of varying complexity, and are therefore also likely to occur at different depths of suppression. While the problem used to be that b-CFS did not provide an unequivocal measure of non-conscious perception, the problem now has increased in magnitude. That is, the graded nature of non-conscious perception, as deduced from the different levels of non-conscious cognitive functioning, indicates that there is no such thing as non-conscious processing per se.

Indeed, the present observations have important implications for the field of consciousness, that go beyond validity issues of experimental paradigms. To make this case it is important to emphasize that the modulation of visual working memory affected *when* a stimulus reached consciousness. That is, observers were not able to report the occurrence, identity or location of said stimulus prior to a certain point in time, after which (a part of) the stimulus was accessible to consciousness and they provided a response. The potency of experimental manipulations to affect the (still) interocularly suppressed visual input, however, was not homogenous over time. Depending on the depth of suppression, experimental manipulations may or may not affect the processing of the interocularly suppressed stimulus. As we suggested above, manipulations that need to tap into increasingly complex cognitive functions in order to affect suppression durations, require shallower suppression. As such, despite the binary nature of observer's conscious perception of a stimulus property (e.g., its location), it appears that non-conscious processing itself is not unitary. Rather, different levels of non-conscious processing allow neural responses to non-conscious stimuli to travel up to different levels of the (visual) processing hierarchy, thereby enabling different cognitive functions (e.g., Kaniwsher, 2001; Bar et al., 2001; Grill-Spector, Kushnir, Hendler, & Malach, 2000; Sergent, & Dehaene, 2004; but see, Dehaene et al., 2001; Supèr, Spekreijse, & Lamme, 2001).

In the introduction we considered the point of view that consciousness cannot be separated from the cognitive functions that are associated with it (Cohen, & Dennett, 2011; Dennett, 1991). Given the continuity in the levels of complexity of cognitive functions that are associated with conscious and non-conscious states, it seems likely that consciousness is a graded phenomenon. Indeed, some extremely complex cognitive functions are reserved to the conscious realms, whereas other extremely simple cognitive functions can safely reside within the non-conscious realms. In between the two, however, there is a gray area in which the neural response to visual information is not entirely restricted to early sensory areas, but also does not give rise to a widespread distributed pattern of neural activity. As a result of this, visual information within this gray area might be determined to be conscious (depending on how it is asked) and certain cognitive functions might be applied to this visual information, whereas others may not. In the end, complex cognitive functions seem to require consciousness, because they require the distributed activity pattern that characterizes conscious processing. Thus, by making the claim that consciousness is an all-or-none phenomenon, while agreeing that the patterns of activity associated with conscious and non-conscious events are graded, one dissociates consciousness from said activity patterns. Similarly, by making the claim that consciousness is an all-or-none phenomenon, while agreeing that the functions that are associated with consciousness are graded, one dissociates consciousness from the functions that it accompanies. Dissociating consciousness from the functions that it accompanies, and from the patterns of activity that are associated with it, simultaneously removes consciousness from the realm of empiricism, and degrades it to an epiphenomenon. Similarly, attributing complex cognitive functions to the non-conscious mind, is like attributing a bird's ability to fly to its spine, while asserting that its wings have emerged as an epiphenomenon.

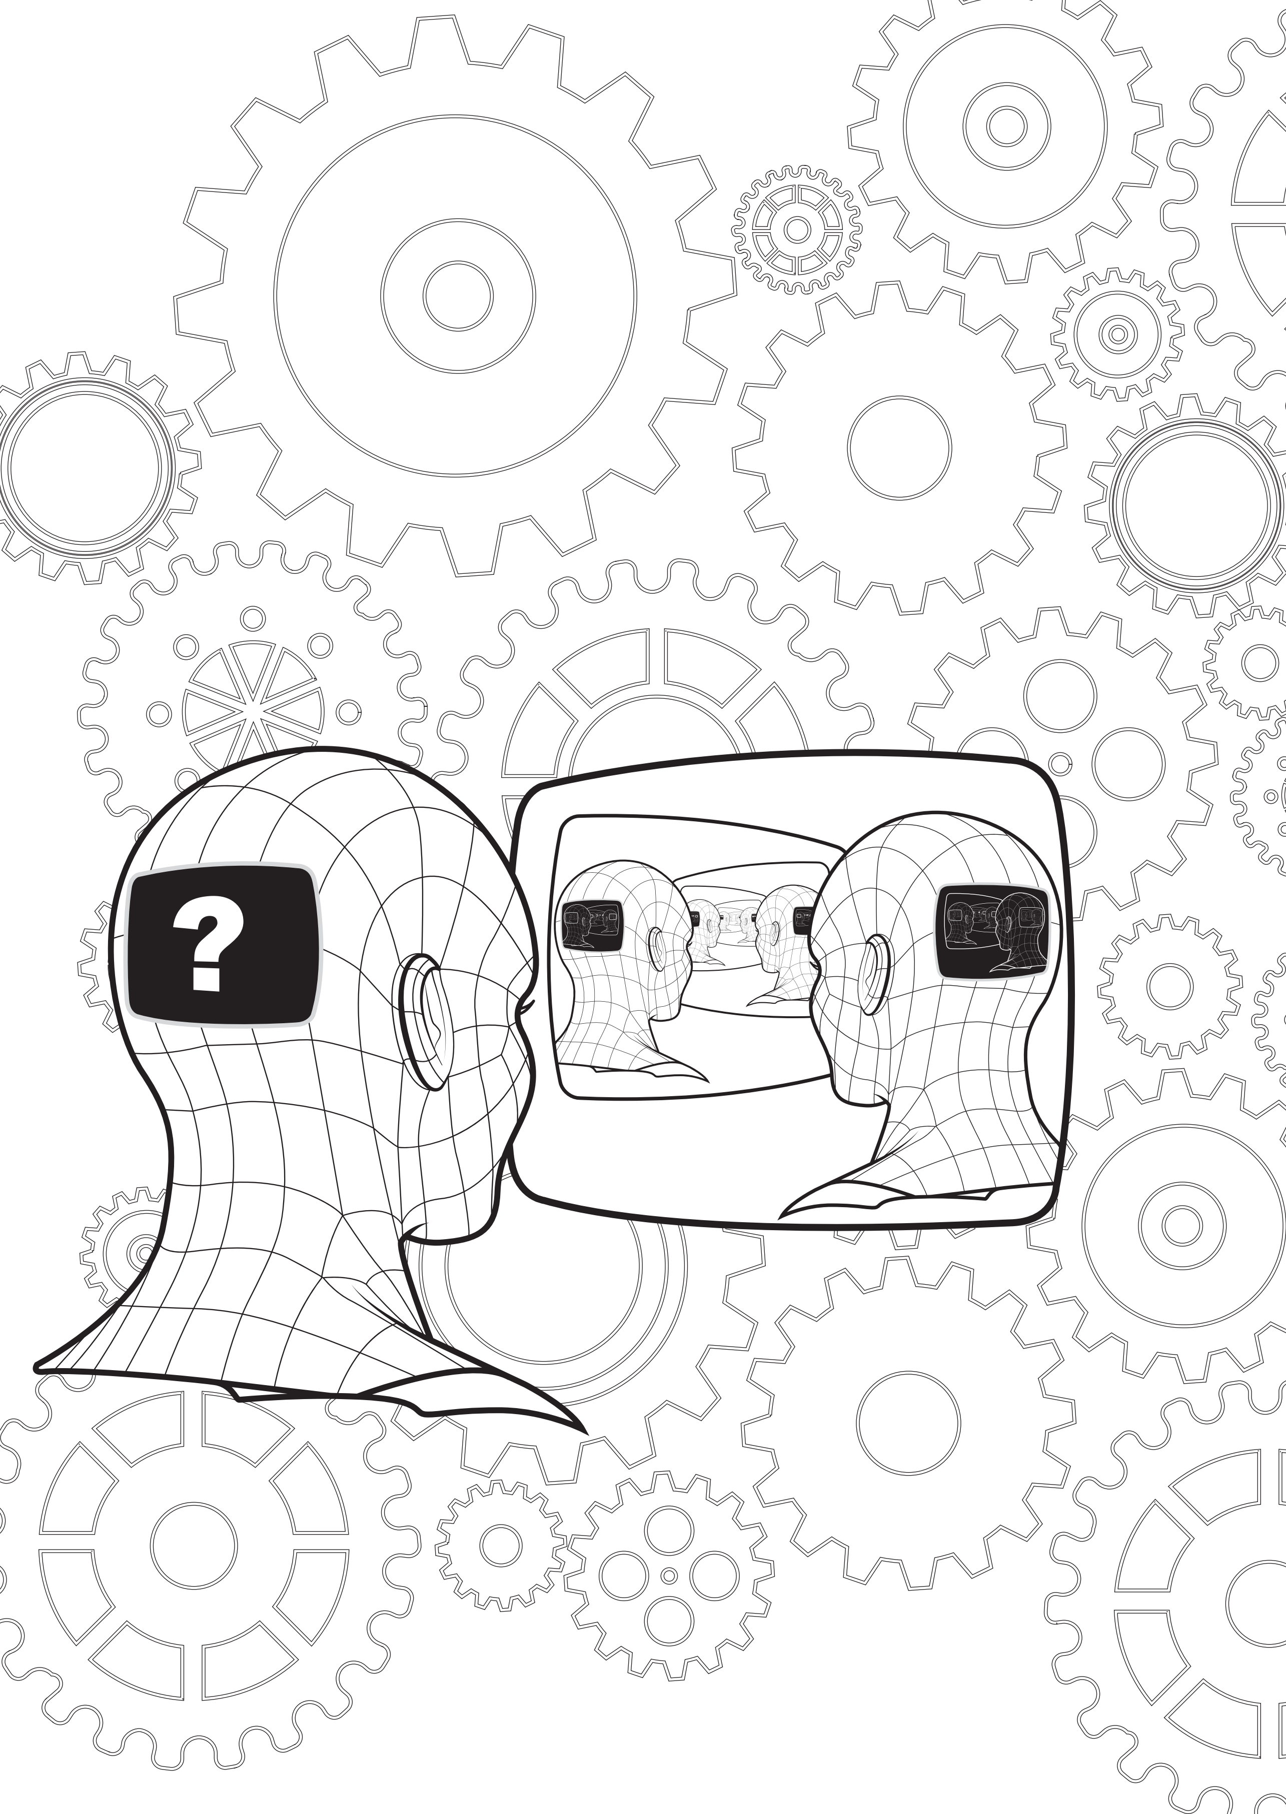

References

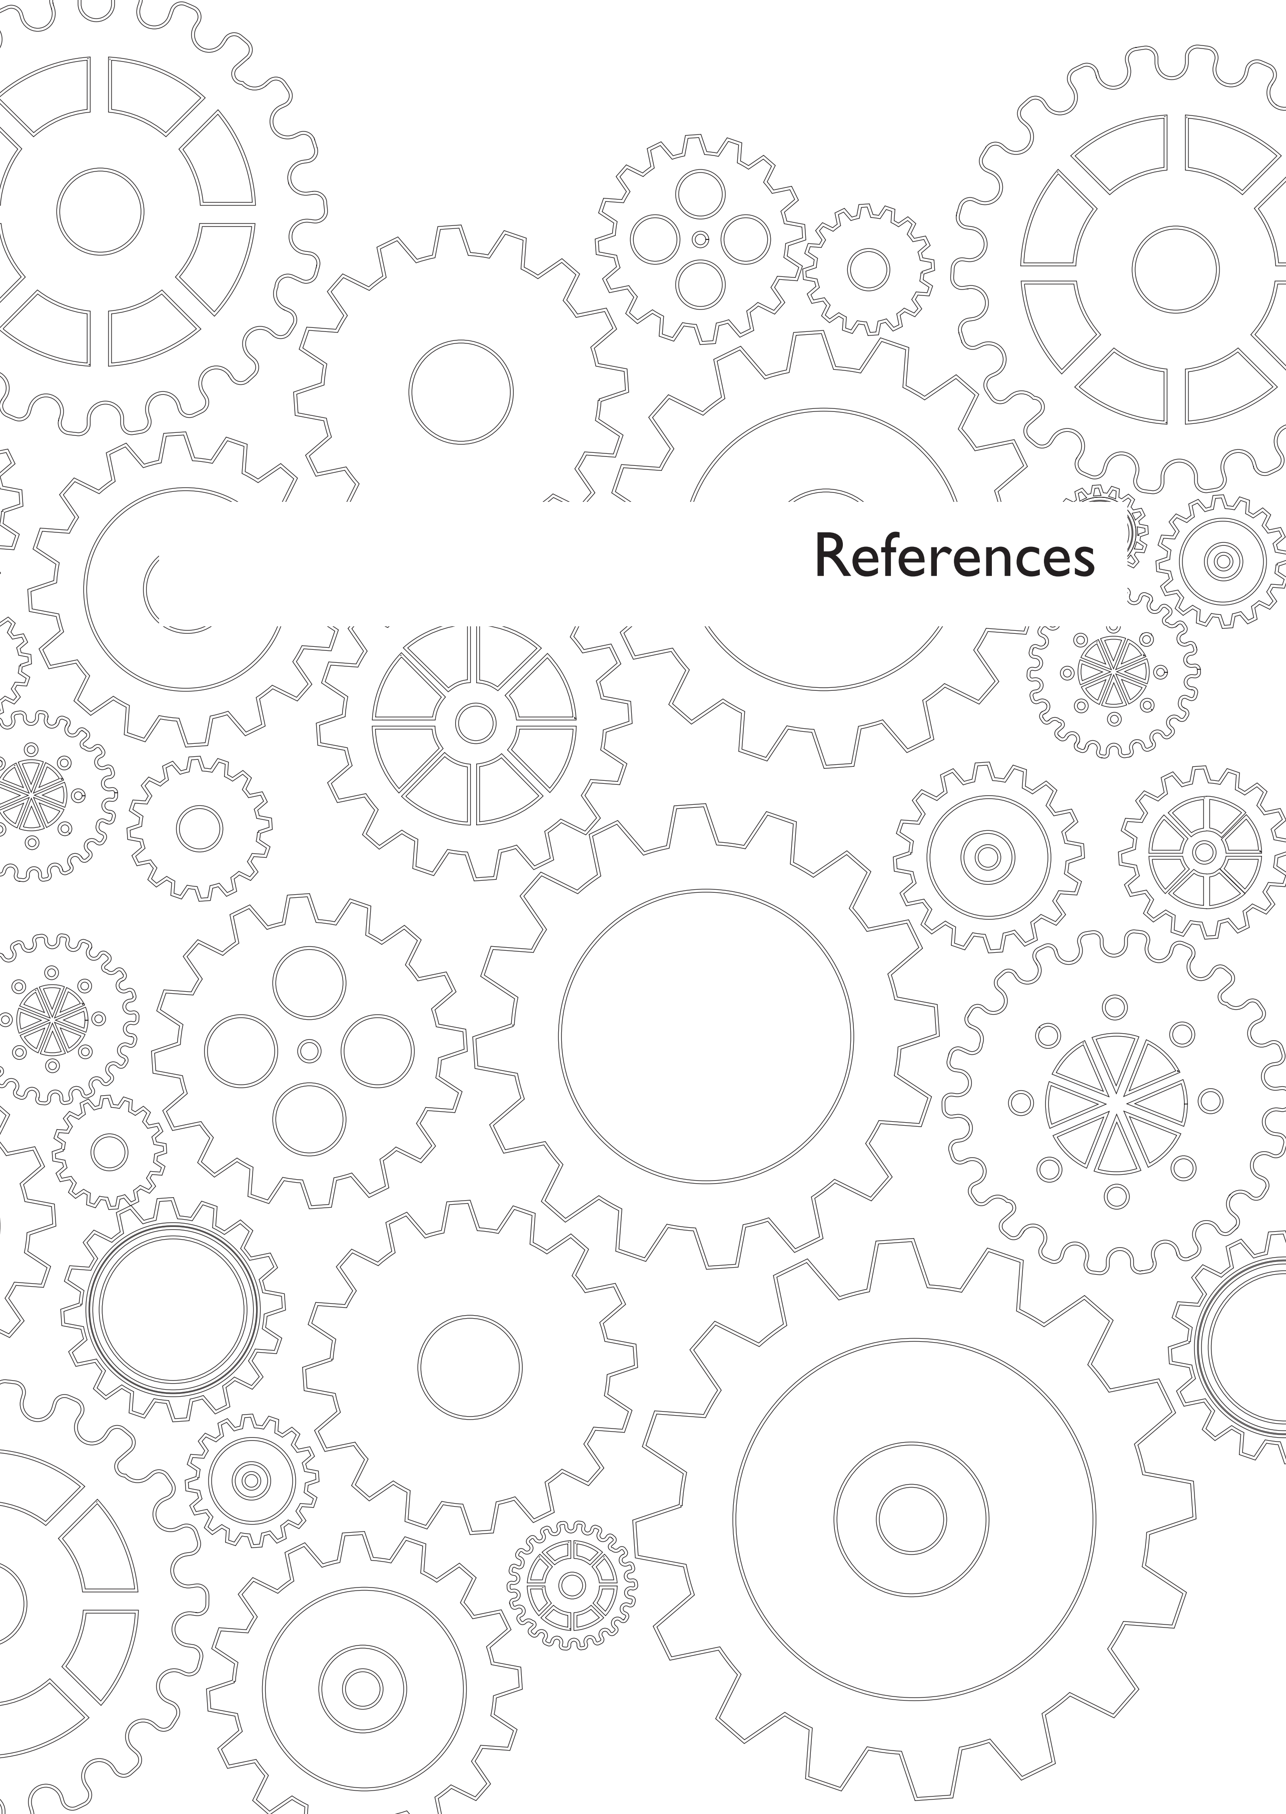

Abrams, R. L., & Greenwald, A. G. (2000). Parts outweigh the whole (word) in unconscious analysis of meaning. *Psychological Science*, *11*(2), 118-124.

Alais, D., & Blake, R. (Eds.). (2005). *Binocular rivalry.* Cambridge: MA: MIT Press.

Alais, D., & Parker, A. (2006). Independent binocular rivalry processes for motion and form. *Neuron*, *52*(5), 911-920.

Albers, A. M., Kok, P., Toni, I., Dijkerman, H. C., & de Lange, F. P. (2013). Shared representations for working memory and mental imagery in early visual cortex. *Current Biology*, *23*(15), 1427-1431.

Alpers, G. W., & Gerdes, A. (2007). Here is looking at you: emotional faces predominate in binocular rivalry. *Emotion*, *7*(3), 495-506.

Alpers, G. W., Ruhleder, M., Walz, N., Mühlberger, A., & Pauli, P. (2005). Binocular rivalry between emotional and neutral stimuli: A validation using fear conditioning and EEG. *International Journal of Psychophysiology*, *57*(1), 25-32.

Alsius, A., & Munhall, K. G. (2013). Detection of audiovisual speech correspondences without visual awareness. *Psychological Science*, *24*(4), 423-431.

Amaral, D. G., & Price, J. L. (1984). Amygdalo‐ cortical projections in the monkey (Macaca fascicularis). *Journal of Comparative Neurology*, *230*(4), 465-496.

Amaral, D. G., Behniea, H., & Kelly, J. L. (2003). Topographic organization of projections from the amygdala to the visual cortex in the macaque monkey. *Neuroscience*, *118*(4), 1099-1120.

Anders, R., Alario, F. X., & Van Maanen, L. (in press). The shifted Wald distribution for response time data analysis. *Psychological Methods*.

Andrews, T. J., & Blakemore, C. (1999). Form and motion have independent access to consciousness. *Nature Neuroscience*, *2*(5), 405-406.

Ansorge, U., & Heumann, M. (2006). Shifts of visuospatial attention to invisible (metacontrast-masked) singletons: Clues from reaction times and event-related potentials. *Advances in Cognitive Psychology, 2*, 61-76.

Ansorge, U., & Neumann, O. (2005). Intentions determine the effect of invisible metacontrast-masked primes: Evidence for top-down contingencies in a peripheral cueing task. *Journal of Experimental Psychology: Human Perception and Performance, 31*, 762-777.

Ansorge, U., Horstmann, G., & Worschech, F. (2010). Attentional capture by masked colour singletons. *Vision Research*, *50*(19), 2015-2027.

Ansorge, U., Kunde, W., & Kiefer, M. (2014). Unconscious vision and executive control: How unconscious processing and conscious action control interact. *Consciousness and Cognition*, *27*, 268-287.

Armony, J. L., & Dolan, R. J. (2002). Modulation of spatial attention by fear-conditioned stimuli: an event-related fMRI study. *Neuropsychologia*, *40*(7), 817-826.

Baars, B. J. (1993). *A cognitive theory of consciousness*. Cambridge: Cambridge University Press.

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

Baars, B. J. (1997b). In the theatre of consciousness: Global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies*, *4*, 292-309.

Baars, B. J. (1997c). Some essential differences between consciousness and attention, perception, and working memory. *Consciousness and Cognition*, *6*(2), 363-371.

Baars, B. J. (1998). Metaphors of consciousness and attention in the brain. *Trends in neurosciences*, *21*(2), 58-62.

Baars, B. J. (2002). The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Sciences*, *6*(1), 47-52.

Baars, B. J. (2003). Working memory requires conscious processes, not vice versa: A global workspace account. In N. Osaka (Ed.), *Neural basis of consciousness* (pp. 11-26). Amsterdam, The Netherlands: John Benjamins.

Baars, B. J. (2005). Global workspace theory of consciousness: Toward a cognitive neuroscience of human experience. *Progress in Brain Research*, *150*, 45-53.

Baars, B. J., & Franklin, S. (2003). How conscious experience and working memory interact. *Trends in Cognitive Sciences*, *7*(4), 166-172.

Baddeley, A. (2003). Working memory: Looking back and looking forward. *Nature Reviews Neuroscience*, *4*, 829-839.

Badre, D., & Wagner, A. D. (2007). Left ventrolateral prefrontal cortex and the cognitive control of memory. *Neuropsychologia*, *45*(13), 2883-2901.

Badre, D., Poldrack, R. A., Paré-Blagoev, E. J., Insler, R. Z., & Wagner, A. D. (2005). Dissociable controlled retrieval and generalized selection mechanisms in ventrolateral prefrontal cortex. *Neuron*, *47*(6), 907-918.

Baker, D. H., & Graf, E. W. (2009). Natural images dominate in binocular rivalry. *Proceedings of the National Academy of Sciences*, *106*(13), 5436-5441.

Balcetis, E., Dunning, D., & Granot, Y. (2012). Subjective value determines initial dominance in binocular rivalry. *Journal of Experimental Social Psychology*, *48*(1), 122-129.

Bannerman, R. L., Milders, M., De Gelder, B., & Sahraie, A. (2008). Influence of emotional facial expressions on binocular rivalry. *Ophthalmic and Physiological Optics*, *28*(4), 317-326.

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

Barbot, A., & Kouider, S. (2012). Longer is not better: nonconscious overstimulation reverses priming influences under interocular suppression. *Attention, Perception and Psychophysics, 74*, 174-184.

Bargh, J. A., & Morsella, E. (2008). The unconscious mind. *Perspectives on Psychological Science, 3*, 73-79.

Bettencourt, K. C., & Xu, Y. (2016). Decoding the content of visual short-term memory under distraction in occipital and parietal areas. *Nature Neuroscience*, *19*(1), 150-157.

Blake, R. (1989). A neural theory of binocular rivalry. *Psychological Review, 96*, 145-167.

Blake, R. (2001). A primer on binocular rivalry, including current controversies. *Brain and Mind 2*(1), 5-38.

Blake, R., & Fox, R. (1974). Adaptation to invisible gratings and the site of binocular rivalry suppression. *Nature, 249*(5456).

Blake, R., & Logothetis, N. K. (2002). Visual competition. *Nature Reviews Neuroscience*, *3*(1), 13-21.

Blake, R., Ahlström, U., & Alais, D. (1999). Perceptual priming by invisible motion. *Psychological Science*, *10*(2), 145-150.

Blake, R., O'Shea, R. P., & Mueller, T. J. (1992). Spatial zones of binocular rivalry in central and peripheral vision. *Visual Neuroscience*, *8*(05), 469-478.

Blake, R., Tadin, D., Sobel, K. V., Raissian, T. A., & Chong, S. C. (2006). Strength of early visual adaptation depends on visual awareness. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(12), 4783-4788.

Blake, R., Westendorf, D., & Fox, R. (1990). Temporal perturbations of binocular rivalry. *Perception & Psychophysics, 48*(6), 593-602.

Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, *15*(12), 567-575.

Bodner, G. E., & Dypvik, A. T. (2005). Masked priming of number judgments depends on prime validity and task. *Memory & Cognition, 33*, 29-47.

Boehler, C. N., Schoenfeld, M. A., Heinze, H. J., & Hopf, J. M. (2008). Rapid recurrent processing gates awareness in primary visual cortex. *Proceedings of the National Academy of Sciences, 105*(25), 8742-8747.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*, 433-436.

Brascamp, J. W., Van Ee, R., Noest, A. J., Jacobs, R. H., & van den Berg, A. V. (2006). The time course of binocular rivalry reveals a fundamental role of noise. *Journal of vision*, *6*(11), 8.

Breitmeyer, B. G. (2007). Visual masking: past accomplishments, present status, future developments. *Advances in Cognitive Psychology, 3*, 9-20.

Brouwer, G. J., & Heeger, D. J. (2009). Decoding and reconstructing color from responses in human visual cortex. *The Journal of Neuroscience*, *29*(44), 13992-14003.

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*, 153-178.

Carney, T., Shadlen, M., & Switkes, E. (1987). Parallel processing of motion and colour information. *Nature*, *328*, 647-649.

Carter, O., & Cavanagh, P. (2007). Onset rivalry: Brief presentation isolates an early independent phase of perceptual competition. *PLoS one, 2*(4): e343.

Cave, C. B., Blake, R., & McNamara, T. P. (1998). Binocular rivalry disrupts visual priming. *Psychological Science*, *9*(4), 299-302. doi:10.1111/1467-9280.00059

Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of consciousness studies*, *2*(3), 200-219.

Chen, Y. C., &Yeh, S. L. (2012). Look into my eyes and I will see you: unconscious processing of human gaze. *Consciousness and Cognition*, *21*(4), 1703-1710.

Chong, S. C., & Blake, R. (2006). Exogenous attention and endogenous attention influence initial dominance in binocular rivalry. *Vision Research*, *46*(11), 1794-1803.

Chong, S. C., Tadin, D., & Blake, R. (2005). Endogenous attention prolongs dominance durations in binocular rivalry. *Journal of Vision*, *5*(11), 6.

Chopin, A., & Mamassian, P. (2010). Task usefulness affects perception of rivalrous images. *Psychological Science*, *21*(12), 1886-1893.

Christophel, T. B., Hebart, M. N., & Haynes, J. D. (2012). Decoding the contents of visual short-term memory from human visual and parietal cortex. *The Journal of Neuroscience, 32*(38), 12983-12989.

Cichy, R. M., Chen, Y., & Haynes, J. D. (2011). Encoding the identity and location of objects in human LOC. *Neuroimage*, *54*(3), 2297-2307.

Clerc, M. (2010). *Particle swarm optimization*. Vol. 93. Hoboken: John Wiley & Sons.

Cohen, J. (1992). A power primer. *Quantitative Methods in Psychology, 112*, 155-159.

Cohen, M. A., & Dennett, D. C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, *15*(8), 358-364.

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

Coren, S., & Russell, J. A. (1992). The relative dominance of different facial expressions of emotion under conditions of perceptual ambiguity. *Cognition & Emotion*, *6*(5), 339-356.

Costello, P., Jiang, Y., Baartman, B., McGlennen, K., & He, S. (2009). Semantic and subword priming during binocular suppression. *Consciousness and Cognition, 18*(2), 375-382. doi:10.1016/j.concog.2009.02.003

Cowan, N. (2001). Metatheory of storage capacity limits. *Behavioral and brain sciences*, *24*(01), 154-176.

Creed, R. S. (1935). Observations on binocular fusion and rivalry. *The Journal of Physiology*, *84*(4), 381-392.

Crick, F., & Koch, C. (2003). A framework for consciousness. *Nature neuroscience*, *6*(2), 119-126.

Damian, M. F. (2001). Congruity effects evoked by subliminally presented primes: automaticity rather than semantic processing. *Journal of Experimental Psychology: Human Perception and Performance*, *27*(1), 154.

de Lange, F. P., Rahnev, D. A., Donner, T. H., & Lau, H. (2013). Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. *The Journal of Neuroscience*, *33*(4), 1400-1410.

De Loof, E., Poppe, L., Cleeremans, A., Gevers, W., & Van Opstal, F. (2015). Different effects of executive and visuospatial working memory on visual consciousness. *Attention, Perception, & Psychophysics*, *77*(8), 2523-2528.

De Loof, E., Verguts, T., Fias, W., & Van Opstal, F. (2013). Opposite effects of working memory on subjective visibility and priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(6), 1959-1965.

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

Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, *10*(5), 204-211. doi:10.1016/j.tics.2006.03.007

Dehaene, S., Kerszberg, M., & Changeux, J. P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences*, *95*(24), 14529-14534.

Dehaene, S., Naccache, L., Cohen, L., Le Bihan, D., Mangin, J. F., Poline, J. B., & Rivière, D. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature neuroscience*, *4*(7), 752-758.

Dennett, D. C. (1991). *Consciousness explained*. Boston, MA: Little, Brown.

D'Esposito, M. (2007). From cognitive to neural models of working memory. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1481), 761-772.

Di Lollo, V., Enns, J. T., & Rensink, R. A. (2000). Competition for consciousness among visual events: The psychophysics of reentrant visual processes. *Journal of Experimental Psychology: General, 129*, 481-507.

Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psycholology, 5*: 781.

Dieter, K. C., & Tadin, D. (2011). Understanding attentional modulation of binocular rivalry: A framework based on biased competition. *Frontiers in Human Neuroscience*, *5*, 155.

Dieter, K. C., Melnick, M. D., & Tadin, D. (2015). When can attention influence binocular rivalry?. *Attention, Perception, & Psychophysics*, *77*(6), 1908-1918.

Edelman, G. M. (1989). *The remembered present: A biological theory of consciousness*. New York, NY: Basic Books.

Edelman, G. M., & Tononi, G. (2000). *A universe of consciousness: how matter becomes imagination.* New York, NY: Basic Books.

Efron, R. (1967). Biology without consciousness - and its consequences. *Perspectives in biology and medicine*, *11*(1), 9-36.

Eger, E., Ashburner, J., Haynes, J. D., Dolan, R. J., & Rees, G. (2008). fMRI activity patterns in human LOC carry information about object exemplars within category. *Journal of Cognitive Neuroscience*, *20*(2), 356-370.

Egeth, H.E., & Yantis, S. (1997). Visual attention: Control, representation, and time course. *Annual Review of Psychology. 48*, 269-297.

Eimer, M., & Schlaghecken, F. (1998). Effects of masked stimuli on motor activation: Behavioral and electrophysiological evidence. *Journal of Experimental Psychology: Human Perception and Performance, 24*, 1737-1747.

Eimer, M., & Schlaghecken, F. (2002). Links between conscious awareness and response inhibition: Evidence from masked priming. *Psychonomic Bulletin & Review, 9*, 514-520.

Engel, E. (1956). The role of content in binocular resolution. *The American Journal of Psychology, 69*(1), 87-91.

Ester, E. F., Sprague, T. C., & Serences, J. T. (2015). Parietal and frontal cortex encode stimulus-specific mnemonic representations during visual working memory. *Neuron*, *87*(4), 893-905.

Fahrenfort, J. J., & Lamme, V. A. F. (2012). A true science of consciousness explains phenomenology: comment on Cohen and Dennett. *Trends in Cognitive Sciences*, *16*(3):138.

Faivre, N., & Koch, C. (2014). Temporal structure coding with and without awareness. *Cognition*, *131*(3), 404-414.

Fang, F. He, S. (2005). Cortical responses to invisible objects in the human dorsal and ventral pathways. *Nature Neuroscience*, 8, 1380-1385.

Finkbeiner, M., & Al-Janabi, S. (2012). Effective processing of masked eye gaze requires volitional control. *Experimental Brain Research, 216*, 433-443.

Folk C. L., Remington RW., & Johnston J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human. Perception & Performance, 18*, 1030-1044.

Forstmann, B. U., Brown, S., Dutilh, G., Neumann, J., & Wagenmakers, E.-J. (2010). The neural substrate of prior information in perceptual decision making: a model-based analysis. *Frontiers in Human Neuroscience, 4*: 40.

Forstmann, B. U., Ratcliff, R., & Wagenmakers, E.-J. (2016). Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annual Review of Psychology*, *67*(24), 1-26.

Friesen, C. K., Ristic, J., & Kingstone, A. (2004). Attentional effects of counterpredictive gaze and arrow cues. *Journal of Experimental Psychology: Human Perception and Performance, 30*, 319-329.

Frischen, A., & Tipper, S. P. (2004). Orienting attention via observed gaze shift evokes longer term inhibitory effects: Implications for social interactions, attention, and memory. *Journal of Experimental Psychology: General, 133*, 516-533.

Frischen, A., Smilek, D., Eastwood, J.D., & Tipper, S. P. (2007). Inhibition of return in response to gaze cues: The roles of time course and fixation cue. *Visual Cognition, 15*, 881-895.

Gabrieli, J. D., Poldrack, R. A., & Desmond, J. E. (1998). The role of left prefrontal cortex in language and memory. *Proceedings of the National Academy of Sciences*, *95*(3), 906-913.

Gaillard, R., Del Cul, A., Naccache, L., Vinckier, F., Cohen, L., & Dehaene, S. (2006). Nonconscious semantic processing of emotional words modulates conscious access. *Proceedings of the National Academy of Sciences, USA, 103,* 7524-7529.

Gayet, S., Brascamp, J.W., Van der Stigchel, S., & Paffen, C. L. (2015). Cogito ergo video: task-relevant information is involuntarily boosted into awareness. *Journal of Vision*, *15*(5):3.

Gayet, S., Paffen, C. L. E., & Van der Stigchel, S. (2013). Information matching the content of visual working memory is prioritized for conscious access. *Psychological Science*, *24*(12), 2472-2480.

Gayet, S., Paffen, C. L., E., Belopolsky, A. V., Theeuwes, J., & Van der Stigchel, S. (2016). Visual input signaling threat gains preferential access to awareness in a breaking continuous flash suppression paradigm. *Cognition*, *149*, 77-83.

Gayet, S., Van der Stigchel, S., & Paffen, C. L. (2014b). Breaking continuous flash suppression: competing for consciousness on the pre-semantic battlefield. *Frontiers in Psychology*, *5*: 460.

Gayet, S., Van der Stigchel, S., & Paffen, C. L. (2014a). Seeing is believing: Utilization of subliminal symbols requires a visible relevant context. *Attention, Perception, & Psychophysics*, *76*(2), 489-507.

Gazzaniga, M. S. (2005). Forty-five years of split-brain research and still going strong. *Nature Reviews Neuroscience*, *6*(8), 653-659.

Gelbard-Sagiv, H., Faivre, N., Mudrik, L., & Koch, C. (2016). Low-level awareness accompanies "unconscious" high-level processing during continuous flash suppression. *Journal of Vision*, *16*(1): 3.

Gobbini, M. I., Gors, J. D., Halchenko, Y. O., Hughes, H. C., & Cipolli, C. (2013). Processing of invisible social cues. *Consciousness and Cognition*, *22*(3), 765-770.

Gobbini, M. I., Gors, J. D., Halchenko, Y. O., Rogers, C., Guntupalli, J. S., Hughes, H., & Cipolli, C. (2013). Prioritized Detection of Personally Familiar Faces. *PloS one*, *8*(6), e66620.

Godijn, R., & Theeuwes, J. (2004). The relationship between inhibition of return and saccade trajectory deviations. *Journal of Experimental Psychology: Human Perception and Performance, 30*, 538-554.

Gray, K. L., Adams, W. J., & Garner, M. (2009). The influence of anxiety on the initial selection of emotional faces presented in binocular rivalry. *Cognition*, *113*(1), 105-110.

Gray, K. L., Adams, W. J., Hedger, N., Newton, K. E., & Garner, M. (2013). Faces and awareness: low-level, not emotional factors determine perceptual dominance. *Emotion*, *13*(3), 537.

Grill-Spector, K., Kourtzi, Z., & Kanwisher, N. (2001). The lateral occipital complex and its role in object recognition. *Vision Research*, *41*(10), 1409-1422.

Grill-Spector, K., Kushnir, T., Edelman, S., Avidan, G., Itzchak, Y., & Malach, R. (1999). Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron*, *24*(1), 187-203.

Grill-Spector, K., Kushnir, T., Hendler, T., & Malach, R. (2000). The dynamics of object-selective activation correlate with recognition performance in humans. *Nature Neuroscience*, *3*(8), 837-843.

Guggenmos, M., Thoma, V., Cichy, R. M., Haynes, J. D., Sterzer, P., & Richardson-Klavehn, A. (2015). Non-holistic coding of objects in lateral occipital complex with and without attention. *Neuroimage*, *107*, 356-363.

Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*(7238), 632-635.

Harvey, B. M., & Dumoulin, S. O. (2011). The relationship between cortical magnification factor and population receptive field size in human visual cortex: constancies in cortical architecture. *The Journal of Neuroscience*, *31*(38), 13604-13612.

Hassin, R. R. (2013). Yes it can on the functional abilities of the human unconscious. *Perspectives on Psychological Science*, *8*(2), 195-207.

Haynes, J. D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews Neuroscience*, *7*(7), 523-534.

Haynes, J. D., Deichmann, R., & Rees, G. (2005). Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus. *Nature, 438*, 496-499. doi:10.1038/nature04169

Heathcote, A. (2004). Fitting wald and ex-Wald distributions to response time data: an example using functions for the S-PLUS package. *Behavior Research Methods, Instruments and Computers*, *36*(4), 678-694.

Hebart, M. N., Görgen, K., & Haynes, J. D. (2014). The Decoding Toolbox (TDT): a versatile software package for multivariate analyses of functional imaging data. *Frontiers in Neuroinformatics*, *8*: 88.

Heitland, I., Groenink, L., Bijlsma, E.Y., Oosting, R. S., & Baas, J. M. (2012). Human fear acquisition deficits in relation to genetic variants of the corticotropin releasing hormone receptor 1 and the serotonin transporter. *PloS one*, *8*(5), e63772.

Hesselmann, G., & Moors, P. (2015). Definitely maybe: can unconscious processes perform the same functions as conscious processes?. *Frontiers in Psychology*, *6*: 584.

Hietanen, J. K., Leppänen, J. M., Nummenmaa, L., & Astikainen, P. (2008). Visuospatial attention shifts by gaze and arrow cues: An ERP study. *Brain Research, 1215*, 123-136.

Ho, T., Brown, S., van Maanen, L., Forstmann, B. U., Wagenmakers, E. J., & Serences, J. T. (2012). The optimality of sensory processing during the speed-accuracy tradeoff. *The Journal of Neuroscience*, *32*(23), 7992-8003.

Hollingworth, A., & Luck, S. J. (2009). The role of visual working memory (VWM) in the control of gaze during visual search. *Attention, Perception, & Psychophysics*, *71*(4), 936-949.

Hollingworth, A., Matsukura, M., & Luck, S. J. (2013). Visual working memory modulates rapid eye movements to simple onset targets. *Psychological Science*, *24*(5), 790-796.

Hommel, B. (2007). Consciousness and control: not identical twins. *Journal of Consciousness Studies*, *14*(1-2), 155-176.

Hommel, B., Pratt, J., Colzato, L., & Godijn, R. (2001). Symbolic control of visual attention. *Psychological Science, 12*, 360-365.

Hong, S. W., & Shevell, S. K. (2006). Resolution of binocular rivalry: Perceptual misbinding of color. *Visual Neuroscience*, *23*(3-4), 561-566.

Humphreys, G. W., Besner, D., & Quinlan, P.T. (1988). Event perception and the word repetition effect. *Journal of Experimental Psychology: General, 117*, 51-67.

Ishihara, S. (1917). Test for Colour-Blindness. *Tokyo: Hongo Harukicho*.

James, W. (1904). Does consciousness exist?. *The Journal of philosophy, psychology and scientific methods*, *1*(18), 477-491.

Jaskowski, P., Skalska, B., & Verleger, R. (2003). How the self controls its "automatic pilot" when processing subliminal information. *Journal of Cognitive Neuroscience, 15*, 911-920.

Jeffreys, H. (1961). *The Theory of Probability*. Oxford: Oxford University Press.

Jiang, Y., & He, S. (2006). Cortical responses to invisible faces: dissociating subsystems for facial-information processing. *Current Biology*, *16*(20), 2023-2029.

Jiang, Y., Costello, P., & He, S. (2007). Processing of invisible stimuli: Advantage of upright faces and recognizable words in overcoming interocular suppression. *Psychological Science*, *18*(4), 349-355.

Jiang, Y., Costello, P., Fang, F., Huang, M., He, S. (2006). A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Procedures of the National Academy of Science, 103*, 17048-17052.

Jonides, J. (1981). Voluntary versus automatic control over the mind's eye's movement. In J. B. Long & A. D. Baddeley (Eds.), *Attention and performance IX* (pp. 187-203). Hillsdale, NJ: Erlbaum.

Kaiser, P. K., & Comerford, J. P. (1975). Flicker photometry of equally bright lights. *Vision Research*, *15*, 1399-1402.

Kang, M. S., Blake, R., & Woodman, G. (2011). Semantic analysis does not occur in the absence of awareness induced by interocular suppression. *Journal of Neuroscience, 31*, 13535-13545.

Kang, P., & Shevell, S. K. (2008). The role of luminance edges in misbinding of color to form. *Vision Research*, *48*(25), 2495-2500.

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

Kaufman, L. (1963). On the spread of suppression and binocular rivalry. *Vision Research*, *3*(9), 401-415.

Kaunitz, L., Fracasso, A., Lingnau, A., & Melcher, D. (2013). Non-conscious processing of motion coherence can boost conscious access. *PloS one*, *8*(4), e60787.

Kavey, N. B., Whyte, J., Resor, S. R., & Gidro-Frank, S. (1990). Somnambulism in adults. *Neurology*, *40*(5), 749-749.

Kiefer, M., & Brendel, D. (2006). Attentional modulation of unconscious "attentional" processes: Evidence from event-related potentials in masked priming. *Journal of Cognitive Neuroscience, 18*, 184-98.

Kiefer, M., & Martens, U. (2010). Attentional sensitization of unconscious cognition: Task sets modulate subsequent masked semantic priming. *Journal of Experimental Psychology: General, 139*, 464-489.

Kiefer, M., & Spitzer, M. (2000). Time course of conscious and unconscious semantic brain activation. *Neuroreport, 11*, 2401-2407.

Kiesel, A., Kunde, W., & Hoffmann, J. (2007). Mechanisms of subliminal response priming. *Advances in Cognitive Psychology, 3*, 307-315.

Kinoshita, S., Forster K. I., & Mozer, M. C. (2008). Unconscious cognition isn't that smart: modulation of masked repetition priming effect in the word naming task. *Cognition, 107*, 623-649.

Kinoshita, S., Mozer, M. C., & Forster, K. I. (2011). Dynamic adaptation to history of trial difficulty explains the effect of congruency proportion on masked priming. *Journal of Experimental Psychology: General, 140*, 622-636.

Klein, R.M. (2000). Inhibition of return. *Trends in Cognitive Sciences, 4*, 138-147.

Klotz, W., & Neumann, O. (1999). Motor activation without conscious discrimination in metacontrast masking. *Journal of Experimental Psychology: Human Perception and Performance, 25*, 976-992.

Koch, C. (2004). The quest for consciousness: a neurobiological approach. New York, NY: Roberts & Company Publishers.

Koch, C., & Tsuchiya, N. (2007). Phenomenology without conscious access is a form of consciousness without top-down attention. *Behavioral and Brain Sciences*, *30*(5-6), 509-510.

Koechlin, E., Naccache, L., Block, E., & Dehaene, S. (1999). Primed numbers: Exploring the modularity of numerical representations with masked and un- masked semantic priming. *Journal of Experimental Psychology: Human Perception and Performance, 25*, 1882-1905.

Konstantinou, N., & Lavie, N. (2013). Dissociable roles of different types of working memory load in visual detection. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(4), 919-924.

Kouider, S., & Dehaene, S. (2007). Levels of processing during non-conscious perception: A critical review of visual masking. *Philosophical Transactions of the Royal Society B, 362*, 857-875.

Kunde, W., Kiesel, A., & Hoffmann, J. (2003). Conscious control over the content of unconscious cognition. *Cognition*, *88*(2), 223-242.

Lack, L. C. (1978). Selective attention and the control over binocular rivalry. *Perception & Psychophysics*, *15*, 193-200.

Lamme, V. A. F. (2001). Blindsight: the role of feedforward and feedback corticocortical connections. *Acta Psychologica, 107*, 209-228.

Lamme, V. A. F. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, *7*(1), 12-18.

Lamme, V. A. F. (2004). Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness. *Neural networks*, *17*(5), 861-872.

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

Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences*, *15*, 365-373.

Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences, 15,* 365-373.

Leopold, D. A., & Logothetis, N. K. (1999). Multistable phenomena: changing views in perception. *Trends in Cognitive Sciences*, *3*(7), 254-264.

Levelt, W. J. M. (1965). *On binocular rivalry.* (Doctoral dissertation). Institute for Perception RVO-TNO, Soesterberg, the Netherlands.

Lidell, B. J., Williams, L. M., Rathjen, J., Shevrin, H., & Gordon, E. (2004). A temporal dissociation of subliminal versus supraliminal fear perception: An event-related potential study. *Journal of Cognitive Neuroscience, 16*, 479-486.

Lim, S. L., Padmala, S., & Pessoa, L. (2009). Segregating the significant from the mundane on a moment-to-moment basis via direct and indirect amygdala contributions. *Proceedings of the National Academy of Sciences*, *106*(39), 16841-16846.

Lin, Z., & He, S. (2009). Seeing the invisible: The scope and limits of unconscious processing in binocular rivalry. *Progress in Neurobiology, 87*(4), 195-211.

Lipp, O.V., Kempnich, C., Jee, S. H., & Arnold, D. H. (2014). Fear conditioning to subliminal fear relevant and non fear relevant stimuli. *PloS one*, *9*(9), e99332.

Liu, D., Wang, L., & Jiang, Y. (2013). Working memory modulates unconscious visual processing. *Journal of Vision, 13*(9), 802. doi:10.1167/13.9.802

Logothetis, N. K. (1998). Single units and conscious vision. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *353*(1377), 1801-1818.

Love, J., Selker, R., Marsman, M., Jamil, T., Dropmann, D., Verhagen, A. J., Ly, A., Gronau, Q. F., Smira, M., Epskamp, S., Matzke, D., Wild, A., Rouder, J. N., Morey, R. D. & Wagenmakers, E.-J. (2015). JASP (Version 0.7). [Computer software]

Luck, S. J., & Hillyard, S. A. (1994). Spatial filtering during visual search: Evidence from human electrophysiology. *Journal of Experimental Psychology: Human Perception and Performance, 20*, 1000-1014.

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*(6657), 279-281.

Lupyan, G., & Ward, E. J. (2013). Language can boost otherwise unseen objects into visual awareness. *Proceedings of the National Academy of Sciences, 110*(35), 14196-14201.

Mackintosh, N. J. (1983). *Conditioning and associative learning*. Oxford: Clarendon Press.

Macknik, S., & Martinez-Conde, S. (2004). Dichoptic visual masking reveals that early binocular neurons exhibit weak interocular suppression: Implications for binocular vision and visual awareness. *Journal of Cognitive Neuroscience, 16*, 1049-59.

Malek, N., Mendoza-Halliday, D., & Martinez-Trujillo, J. (2012). Binocular rivalry of spiral and linear moving random dot patterns in human observers. *Journal of Vision*, *12*(10), 16.

Maniscalco, B., & Lau, H. (2015). Manipulation of working memory contents selectively impairs metacognitive sensitivity in a concurrent visual discrimination task. *Neuroscience of Consciousness*, *2015*(1), 1-13.

Maruya, K., & Blake, R. (2009). Spatial spread of interocular suppression is guided by stimulus configuration. *Perception*, *38*(2), 215-231.

Marx, S., & Einhäuser, W. (2015). Reward modulates perception in binocular rivalry. *Journal of Vision*, *15*(1), 11.

Mather, M., & Sutherland, M. R. (2011). Arousal-biased competition in perception and memory. *Perspectives on Psychological Science*, *6*(2), 114-133.

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

Meng, M., & Tong, F. (2004). Can attention selectively bias bistable perception? Differences between binocular rivalry and ambiguous figures. *Journal of Vision*, *4*(7), 2.

Meng, Q., Cui, D., Zhou, K., Chen, L., & Ma, Y. (2012). Advantage of hole stimulus in rivalry competition. *PloS one*, *7*(3), e33053.

Mitchell, J. F., Stoner, G. R., & Reynolds, J. H. (2004). Object-based attention determines dominance in binocular rivalry. *Nature*, *429*(6990), 410-413.

Morris, J. S., Öhman, A., & Dolan, R. J. (1998). Conscious and unconscious emotional learning in the human amygdala. *Nature*, *393*(6684), 467-470.

Morris, S. B., & DeShon, R. P. (2002). Combining effect size estimates in meta-analysis with repeated measures and independent-groups designs. *Psychological Methods*, *7*(1), 105-125.

Mudrik, L., & Koch, C. (2013). Differential processing of invisible congruent and incongruent scenes: A case for unconscious integration. *Journal of Vision*, *13*(13), 24.

Mudrik, L., Breska, A., Lamy, D., & Deouell, L. Y. (2011). Integration Without Awareness Expanding the Limits of Unconscious Processing. *Psychological Science*, *22*(6), 764-770.

Mudrik, L., Faivre, N., & Koch, C. (2014). Information integration without awareness. *Trends in Cognitive Sciences, 18*(9), 488-496.

Mudrik, L., Gelbard-Sagiv, H., Faivre, N., & Koch, C. (2013). Knowing where without knowing what: partial awareness and high-level processing in continuous flash suppression. *Journal of Vision*, *13*(9): 1103.

Mulckhuyse, M., & Theeuwes, J. (2010). Unconscious attentional orienting to exogenous cues: a review of the literature. *Acta psychologica*, *134*(3), 299-309.

Mulckhuyse, M., Crombez, G., & Van der Stigchel, S. (2013). Conditioned fear modulates visual selection. *Emotion*, *13*(3), 529-536.

Mulckhuyse, M., Talsma, D., & Theeuwes, J. (2007). Grabbing attention without knowing: Automatic capture of attention by subliminal spatial cues. *Visual Cognition, 15*, 779-788.

Nagel, T. (1974). What is it like to be a bat?. *The philosophical review*, *83*(4), 435-450.

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

Nummenmaa, L., & Hietanen, J.K. (2006). Gaze distractors influence saccadic curvature: evidence for the role of the oculomotor system in gaze-cued orienting. *Visual Research, 46*, 3674-3680.

O'Shea, R. P., & Crassini, B. (1981). The sensitivity of binocular rivalry suppression to changes in orientation assessed by reaction-time and forced-choice techniques. *Perception, 10*, 283-293.

O'Shea, R. P., Sims, A. J., & Govan, D. G. (1997). The effect of spatial frequency and field size on the spread of exclusive visibility in binocular rivalry. *Vision Research*, *37*(2), 175-183. doi:10.1016/S0042-6989(96)00113-7

Öhman, A., Flykt, A., & Esteves, F. (2001). Emotion drives attention: detecting the snake in the grass. *Journal of experimental psychology: general*, *130*(3), 466-478.

Olivers, C. N. L. (2009). What drives memory-driven attentional capture? The effects of memory type, display type, and search type. *Journal of Experimental Psychology: Human Perception and Performance, 35,* 1275-1291.

Olivers, C. N. L., Meijer, F., & Theeuwes, J. (2006). Feature-based memory-driven attentional capture: Visual working memory content affects visual attention. *Journal of Experimental Psychology: Human Perception and Performance, 32*(5), 1243-1265.

Ooi, T. L., & He, Z. J. (1999). Binocular rivalry and visual awareness: the role of attention. *Perception-London*, *28*(5), 551-574.

Padmala, S., & Pessoa, L. (2008). Affective learning enhances visual detection and responses in primary visual cortex. *The Journal of Neuroscience*, *28*(24), 6202-6210.

Paffen, C. L. E., & Alais, D. (2011). Attentional modulation of binocular rivalry. *Frontiers in Human Neuroscience*, *5*, 105-105.

Paffen, C. L. E., Alais, D., & Verstraten, F. A. J. (2006). Attention speeds binocular rivalry. *Psychological Science, 17*, 752-756.

Paffen, C. L. E., Hessels, R. S., & Van der Stigchel, S. (2012). *Attention, Perception & Psychophysics, 74*, 251-256.

Paffen, C. L. E., Naber, M., & Verstraten, F. A. (2008). The spatial origin of a perceptual transition in binocular rivalry. *PloS One*, *3*(6), e2311.

Palmer, S., & Mattler, U., (2013). Masked stimuli modulate endogenous shifts of spatial attention. *Consciousness and Cognition, 22*, 486-503.

Pan, Y., & Luo, Q. Y. (2012). Working memory modulates the perception of time. *Psychonomic Bulletin & Review*, *19*(1), 46-51.

Pan, Y., Cheng, Q. P., & Luo, Q. Y. (2012). Working memory can enhance unconscious visual perception. *Psychonomic Bulletin & Review*, *19*(3), 477-482.

Pan, Y., Lin, B., Zhao, Y., & Soto, D. (2014). Working memory biasing of visual perception without awareness. *Attention, Perception, & Psychophysics*, *76*(7), 2051-2062.

Parker, A., & Alais, D. (2007). A bias for looming stimuli to predominate in binocular rivalry. *Vision Research*, *47*(20), 2661-2674.

Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin, 116,* 220-244.

Pasley, B. N., Mayes, L. C., & Schultz, R. T. (2004). Subcortical discrimination of unperceived objects during binocular rivalry. *Neuron*, *42*(1), 163-172.

Pavlov, I. P. (1927). *Conditioned reflexes: An investigation of the physiological activities of the cerebral cortex*. New York: Dover.

Pelli, D. G. (1997). The Video Toolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437-442.

Peremen, Z., & Lamy, D. (2014). Comparing unconscious processing during continuous flash suppression and meta-contrast masking just under the limen of consciousness. *Frontiers in Psychology*, *5*:969.

Pessoa, L. (2005). To what extent are emotional visual stimuli processed without attention and awareness?. *Current Opinion in Neurobiology*, *15*(2), 188-196.

Pessoa, L., & Adolphs, R. (2010). Emotion processing and the amygdala: from a'low road'to'many roads' of evaluating biological significance. *Nature Reviews Neuroscience*, *11*(11), 773-783.

Pessoa, L., Gutierrez, E., Bandettini, P.A., & Ungerleider, L. G. (2002). Neural correlates of visual working memory: fMRI amplitude predicts task performance. *Neuron*, *35*(5), 975-987.

Phelps, E. A., Ling, S., & Carrasco, M. (2006). Emotion facilitates perception and potentiates the perceptual benefits of attention. *Psychological Science*, *17*(4), 292-299.

Plazzi, G., Vetrugno, R., Provini, F., & Montagna, P. (2005). Sleepwalking and other ambulatory behaviours during sleep. *Neurological Sciences*, *26*, s193-s198.

Polonsky, A., Blake, R., Braun, J., & Heeger, D. J. (2000). Neuronal activity in human primary visual cortex correlates with perception during binocular rivalry. *Nature Neuroscience, 3*, 1153-1159.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32A*, 3-25.

Posner, M. I., & Cohen, Y. P. C. (1984). Components of visual orienting. In H. Bouma & D. Bouwhuis (Eds.), *Attention and performance X* (pp. 531-556). London: Erlbaum.

Postle, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience*, *139*(1), 23-38.

Pratt, J., & Hommel, B. (2003). Symbolic control of visual attention: The role of working memory and attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance, 29*, 835-845.

Rafal, R., Calabresi, P., Brennan, C., & Sciolto, T. (1989). Saccade preparation inhibits reorienting to recently attended locations. *Journal of Experimental Psychology: Human Perception and Performance, 15*, 673-685.

Raffone, A., & Wolters, G. (2001). A cortical mechanism for binding in visual working memory. *Journal of Cognitive Neuroscience*, *13*(6), 766-785.

Raio, C. M., Carmel, D., Carrasco, M., & Phelps, E. A. (2012). Nonconscious fear is quickly acquired but swiftly forgotten. *Current Biology*, *22*(12), R477-R479.

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

Ratcliff, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin*, *86*, 446-461.

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, *111*(2), 333-367.

Rees, G., Kreiman, G., & Koch, C. (2002). Neural correlates of consciousness in humans. *Nature Reviews Neuroscience*, *3*(4), 261-270.

Ress, D., & Heeger, D. J. (2003). Neuronal correlates of perception in early visual cortex. *Nature Neuroscience, 6,* 414-420.

Reuss, H., Pohl, C., Kiesel, A., & Kunde, W. (2011). Follow the sign! Top-down contingent attentional capture of masked arrow cues. *Advances in Cognitive Psychology, 7*, 82-91.

Riggall, A. C., & Postle, B. R. (2012). The relationship between working memory storage and elevated activity as measured with functional magnetic resonance imaging. *The Journal of Neuroscience*, *32*(38), 12990-12998.

Ristic, J., Friesen, C. K., & Kingstone, A. (2002). Are eyes special? It depends on how you look at it. *Psychonomic Bulletin & Review, 9*, 507—513.

Ristic, J., Wright, A., & Kingstone, A. (2007). Attentional control and reflexive orienting to gaze and arrow cues. *Psychonomic Bulletin & Review, 14*, 964-969.

Rock, I. (1974). The perception of disoriented figures. *Scientific American, 230*, 78-85.

Roseboom, W., & Arnold, D. H. (2011). Learning to reach for 'invisible' visual input. *Current Biology, 21*, 493-494.

Rothkirch, M., Stein, T., Sekutowicz, M., & Sterzer, P. (2012). A direct oculomotor correlate of unconscious visual processing. *Current Biology, 22*, 514-515.

Salomon, R., Lim, M., Herbelin, B., Hesselmann, G., & Blanke, O. (2013). Posing for awareness: proprioception modulates access to visual consciousness in a continuous flash suppression task. *Journal of vision*, *13*(7): 2.

Schall, J. D., Nawrot, M., Blake, R., & Yu, K. (1993). Visually guided attention is neutralized when informative cues are visible but unperceived. *Vision Research, 33*, 2057- 2064.

Schenck, C. H., & Mahowald, M. W. (1995). A polysomnographically documented case of adult somnambulism with long-distance automobile driving and frequent nocturnal violence: parasomnia with continuing danger as a noninsane automatism?. *Sleep*, *18*(9), 765-772.

Schlaghecken, F., & Eimer, M. (2004). Masked prime stimuli can bias "free" choices between response alternatives. *Psychonomic Bulletin & Review, 11*, 463-468.

Schmack, K., Burk, J., Haynes, J. D., & Sterzer, P. (2015). Predicting subjective affective salience from cortical responses to invisible object stimuli. *Cerebral Cortex*, bhv174.

Schmidt, L. J., Belopolsky, A. V., & Theeuwes, J. (2015a). Attentional capture by signals of threat. *Cognition and emotion*, *29*(4), 687-694.

Schmidt, L. J., Belopolsky, A. V., & Theeuwes, J. (2015b). Potential threat attracts attention and interferes with voluntary saccades. *Emotion*, *15*(3), 329-338.

Schneegans, S., Spencer, J. P., Schöner, G., Hwang, S., & Hollingworth, A. (2014). Dynamic interactions between visual working memory and saccade target selection. *Journal of Vision, 14*(11), 9.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, *84*(1), 1.

Scocchia, L., Valsecchi, M., Gegenfurtner, K. R., & Triesch, J. (2013). Visual working memory contents bias ambiguous structure from motion perception. *PloS one*, *8*(3), e59217.

Scocchia, L., Valsecchi, M., Gegenfurtner, K. R., & Triesch, J. (2014). Differential effects of visual attention and working memory on binocular rivalry. *Journal of Vision*, *14*(5), 13.

Self, M. W., & Roelfsema, P. R. (2010). A monocular, unconscious form of visual attention. *Journal of Vision, 10*: 17.

Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychological Science*, *20*(2), 207-214.

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

Sergent, C., & Dehaene, S. (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychological Science 15*(11): 720-728.

Shore, D. I., & Klein, R. M. (2000). The effects of scene inversion on change blindness. *The Journal of General Psychology*, *127*(1), 27-43.

Silvis, J. D., & Van der Stigchel, S. (2014). How memory mechanisms are a key component in the guidance of our eye movements: Evidence from the global effect. *Psychonomic Bulletin & Review*, *21*(2), 357-362.

Sklar, A. Y., Levy, N., Goldstein, A., Mandel, R., Maril, A., & Hassin, R. R. (2012). Reading and doing arithmetic nonconsciously. *Proceedings of the National Academy of Sciences, 109,* 19614-19619.

Sligte, I. G., van Moorselaar, D., & Vandenbroucke, A. R. (2013). Decoding the contents of visual working memory: evidence for process-based and content-based working memory areas?. *The Journal of Neuroscience*, *33*(4), 1293-1294.

Sommer, W., Leuthold, H., & Schubert, T. (2001). Multiple bottlenecks in information processing? An electrophysiological examination. *Psychonomic Bulletin & Review, 8*, 81-88.

Song, J. H., & Jiang, Y. (2006). Visual working memory for simple and complex features: An fMRI study. *Neuroimage*, *30*(3), 963-972.

Soto, D., & Silvanto, J. (2014). Reappraising the relationship between working memory and conscious awareness. *Trends in Cognitive Sciences*, *18*(10), 520-525.

Soto, D., Heinke, D., Humphreys, G. W., & Blanco, M. J. (2005). Early, involuntary top-down guidance of attention from working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(2), 248-261.

Soto, D., Mäntylä, T., & Silvanto, J. (2011). Working memory without consciousness. *Current Biology, 21,* R912-R913. doi:10.1016/j.cub.2011.09.049

Sperry, R. (1984). Consciousness, personal identity and the divided brain. *Neuropsychologia*, *22*(6), 661-673.

Srinivasan, R., Russell, D. P., Edelman, G. M., & Tononi, G. (1999). Increased synchronization of neuromagnetic responses during conscious perception. *The Journal of Neuroscience*, *19*(13), 5435-5448.

Stein, T. (2012). *Visual processing of social information during interocular suppression* (Doctoral dissertation). Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät II.

Stein, T., & Sterzer, P. (2012). Not just another face in the crowd: Detecting emotional schematic faces during continuous flash suppression. *Emotion*, *12*(5), 988-996.

Stein, T., & Sterzer, P. (2014). Unconscious processing under interocular suppression: getting the right measure. *Frontiers in psychology*, *5*, 387.

Stein, T., Hebart, M. N., & Sterzer, P. (2011). Breaking continuous flash suppression: a new measure of unconscious processing during interocular suppression?. *Frontiers in human neuroscience*, *5*: 167.

Stein, T., Kaiser, D., & Hesselmann, G. (2016). Can working memory be non-conscious? *Neuroscience of Consciousness*, 2016, 1-3.

Stein, T., Kaiser, D., & Peelen, M. V. (2015). Interobject grouping facilitates visual awareness. *Journal of Vision*, *15*(8), 10-10.

Stein, T., Peelen, M. V., & Sterzer, P. (2011). Adults' awareness of faces follows newborns' looking preferences. *PLoS one, 6*, e29361.

Stein, T., Senju, A., Peelen, M. V., & Sterzer, P. (2011). Eye contact facilitates awareness of faces during interocular suppression. *Cognition*, *119*(2), 307-311.

Stein, T., Seymour, K., Hebart, M. N., & Sterzer, P. (2014). Rapid fear detection relies on high spatial frequencies. *Psychological Science, 25(2), 566-574.*

Stein, T., Sterzer, P., & Peelen, M. V. (2012). Privileged detection of conspecifics: Evidence from inversion effects during continuous flash suppression. *Cognition, 125*, 64-79.

Sterzer, P., Haynes, J. D., & Rees, G. (2008). Fine-scale activity patterns in high-level visual areas encode the category of invisible objects. *Journal of Vision*, *8*(15), 10.

Sterzer, P., Stein, T., Ludwig, K., Rothkirch, M., & Hesselmann, G. (2014). Neural processing of visual information under interocular suppression: a critical review. *Frontiers in psychology*, *5*: 453.

Stuit, S. M., Paffen, C. L. E., van der Smagt, M. J., & Verstraten, F. A. J. (2011). Suppressed images selectively affect the dominant percept during rivalry. *Journal of Vision, 11*: 10.

Stuit, S. M., Verstraten, F. A., & Paffen, C. L. (2010). Saliency in a suppressed image affects the spatial origin of perceptual alternations during binocular rivalry. *Vision Research*, *50*(19), 1913-1921.

Supèr, H., Spekreijse, H., & Lamme, V. A. F. (2001). A neural correlate of working memory in the monkey primary visual cortex. *Science*, *293*(5527), 120-124.

Supèr, H., Spekreijse, H., & Lamme, V. A. F. (2001). Two distinct modes of sensory processing observed in monkey primary visual cortex (V1). *Nature Neuroscience*, *4*(3), 304-310.

Sweeny, T. D., Grabowecky, M., & Suzuki, S. (2011). Awareness becomes necessary between adaptive pattern coding of open and closed curvatures. *Psychological Science, 22*, 943-950.

Thomas, L. (1997). Retrospective power analysis. *Conservation Biology, 11*, 276-280.

Thompson-Schill, S. L., D'Esposito, M., & Kan, I. P. (1999). Effects of repetition and competition on activity in left prefrontal cortex during word generation. *Neuron*, *23*(3), 513-522.

Tipples, J. (2002). Eye gaze is not unique: Automatic orienting in response to uninformative arrows. *Psychonomic Bulletin & Review, 9*, 314—318.

Tong, F. (2001). Competing theories of binocular rivalry: a possible resolution. *Brain and Mind, 2*, 55-83.

Tong, F., Meng, M., & Blake, R. (2006). Neural bases of binocular rivalry. *Trends in Cognitive Sciences, 10*, 502-511.

Tong, F., Nakayama, K., Vaughan, J. T., & Kanwisher, N. (1998). Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron*, *21*(4), 753-759.

Tononi, G. (2008). Consciousness as integrated information: a provisional manifesto. *The Biological Bulletin*, *215*(3), 216-242.

Tononi, G., & Edelman, G. M. (1998). Consciousness and complexity. *Science*, *282*(5395), 1846-1851.

Tononi, G., & Koch, C. (2015). Consciousness: here, there and everywhere?. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *370*(1668), 20140167.

Tononi, G., Srinivasan, R., Russell, D. P., & Edelman, G. M. (1998). Investigating neural correlates of conscious perception by frequency-tagged neuromagnetic responses. *Proceedings of the National Academy of Sciences*, *95*(6), 3198-3203.

Treisman, A.M. Consciousness and perceptual binding. (2003). Cleeremans, A. (Ed.), *The Unity of Consciousness: Binding, Integration, and Dissociation*, pp. 95-113, Oxford University Press

Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience*, *8*(8), 1096-1101.

Tsuchiya, N., Koch, C., Gilroy, L. A., & Blake, R. (2006). Depth of interocular suppression associated with continuous flash suppression, flash suppression, and binocular rivalry. *Journal of Vision, 6*, 1068-1078.

Turk-Browne, N. B., Scholl, B. J., Chun, M. M., & Johnson , M. K. (2008). Neural evidence of statistical learning: Efficient detection of visual regularities without awareness. *Journal of Cognitive Neuroscience, 21*, 1934-1945.

Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M., (2010). Implicit perceptual anticipation triggered by statistical learning. *The Journal of Neuroscience, 30*, 11177-11187.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., ... & Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage*, *15*(1), 273-289.

van Ee, R., Van Dam, L. C. J., & Brouwer, G. J. (2005). Voluntary control and the dynamics of perceptual bi-stability. *Vision Research*, *45*(1), 41-55.

van Gaal, S., De Lange, F. P., & Cohen, M. X. (2012). The role of consciousness in cognitive control and decision making. *Frontiers in Human Neuroscience*, *6*: 121.

van Gaal, S., Lamme, V. A. F., Fahrenfort, J. J., & Ridderinkhof, R. K. (2011). Dissociable brain mechanisms underlying the conscious and unconscious control of behavior. *Journal of Cognitive Neuroscience, 23*, 91-105.

van Maanen, L., Forstmann, B. U., Keuken, M. C., Wagenmakers, E.-J., & Heathcote, A. (2016). The impact of MRI scanner environment on perceptual decision making. *Behavior Research Methods, 48*(1), 184-200.

van Moorselaar, D., Theeuwes, J., & Olivers, C. N. L. (2014). In competition for the attentional template: can multiple items within visual working memory guide attention? *Journal of Experimental Psychology: Human Perception and Performance, 40*(4), 1450-1464.

van Moorselaar, D., Theeuwes, J., & Olivers, C. N. L. (2015). The limitations of visual working memory in prioritizing visual stimuli for conscious access. *Journal of Vision*, *15*(12): 869.

Vecera, S. P., & Johnson, M. H. (1995). Gaze detection and the cortical processing of faces: evidence from infants and adults. *Visual Cognition*, *2*(1), 59-87.

Vincent, S. B. (1912). *The functions of the vibrissae in the behavior of the white rat.* (Vol. 1, No. 5). University of Chicago.

Wade, N. J., & Wenderoth, P. (1978). The influence of colour and contour rivalry on the magnitude of the tilt after-effect, *Vision Research, 18*, 827-836.

Wagenmakers, E. J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic bulletin & review*, *14*(5), 779-804.

Wagner, G., & Boynton, M. (1972). Comparison of four methods of heterochromatic photometry. *Journal of the Optical Society of America*, *62*, 1508-1515.

Walraven, J. (1972). *TNO test for stereoscopic vision (9th edition)*. Utrecht, the Netherlands: Lameris Instrumenten.

Wang, L., Weng, X., & He, S. (2012). Perceptual grouping without awareness: superiority of Kanizsa triangle in breaking interocular suppression. *PloS one*, 7(6), e40106.

Weiskrantz, L. (1986). *Blindsight*. Oxford, England: Oxford University Press.

Weiskrantz, L. (1997). *Consciousness lost and found: a neuropsychological exploration*. New York, NY: Oxford University Press.

Whalen, P. J., Rauch, S. L., Etcoff, N. L., McInerney, S. C., Lee, M. B., & Jenike, M. A. (1998). Masked presentations of emotional facial expressions modulate amygdala activity without explicit knowledge. *The Journal of Neuroscience*, *18*(1), 411-418.

Wheatstone, C. (1838). Contributions to the Physiology of Vision.--Part the First. On Some Remarkable, and Hitherto Unobserved, Phenomena of Binocular Vision. *Philosophical transactions of the Royal Society of London*, *128*, 371-394.

White, K.D., Petry, H.M., Riggs, L.A., Miller, J., 1978. Binocular interactions during establishment of McCollough effects. *Vision Research, 18*, 1201-1215.

Wilbertz, G., van Slooten, J., & Sterzer, P. (2014). Reinforcement of perceptual inference: reward and punishment alter conscious visual perception during binocular rivalry. *Frontiers in psychology*, *5*: 1377.

Wilke, M., Logothetis, N. K., & Leopold, D. A. (2006). Local field potential reflects perceptual suppression in monkey visual cortex. *Proceedings of the National Academy of Sciences, 103*, 17507-17512.

Williams, M. A., Morris, A. P., McGlone, F., Abbott, D. F., & Mattingley, J. B. (2004). Amygdala responses to fearful and happy facial expressions under conditions of binocular suppression. *The Journal of Neuroscience*, *24*(12), 2898-2904.

Wilson, H. R., Blake, R., & Lee, S. H. (2001). Dynamics of travelling waves in visual perception. *Nature*, *412*(6850), 907-910.

Wolfe, J. M. (1984). Reversing ocular dominance and suppression in a single flash. *Vision Research, 24*, 471-478.

Wyart, V., & Tallon-Baudry, C. (2009). How ongoing fluctuations in human visual cortex predict perceptual awareness: baseline shift versus decision bias. *The Journal of Neuroscience*, *29*(27), 8715-8725.

Xu, Y., & Chun, M. M. (2006). Dissociable neural mechanisms supporting visual short-term memory for objects. *Nature*, *440*(7080), 91-95.

Xu, Y., & Chun, M. M. (2009). Selecting and perceiving multiple visual objects. *Trends in Cognitive Sciences*, *13*(4), 167-174.

Yang, E., & Blake, R. (2011). Deconstructing continuous flash suppression. *Journal of Vision*, *12*(3), 8.

Yang, E., Zald, D. H., & Blake, R. (2007). Fearful expressions gain preferential access to awareness during continuous flash suppression. *Emotion*, *7*(4), 882-886.

Yang, Y. H., & Yeh, S. L. (2011). Accessing the meaning of invisible words. *Consciousness and cognition*, *20*(2), 223-233. doi:10.1016/j.concog.2010.07.005

Yang, Y. H., & Yeh, S. L. (2014). Unmasking the dichoptic mask by sound: spatial congruency matters. *Experimental Brain Research*, 1-8. doi:10.1007/s00221-014-3820-5

Yantis, S., Jonides, J., 1990. Abrupt visual onsets and selective attention: Voluntary versus automatic allocation. *Journal of Experimental Psychology: Human Perception and Performance, 16*, 121-134.

Yiend, J. (2010). The effects of emotion on attention: A review of attentional processing of emotional information. *Cognition and Emotion*, *24*(1), 3-47.

Yuval-Greenberg, S., & Heeger, D. J. (2013). Continuous flash suppression modulates cortical activity in early visual cortex. *The Journal of Neuroscience*, *33*(23), 9635-9643.

Zadbood, A., Lee, S. H., & Blake, R. (2011). Stimulus fractionation by interocular suppression. *Frontiers in Human Neuroscience, 5*: 135.

Zhang, P., Jiang, Y., & He, S. (2012). Voluntary attention modulates processing of eye-specific visual information. *Psychological Science, 23*, 254-260.

Zhao, J., Al-Aidroos, N., & Turk-Browne, N. B. (2013). Attention is spontaneously biased towards regularities. *Psychological Science, 24*, 667-677.

Zhaoping, L. (2008). Attention capture by eye of origin singletons even without awareness: A hallmark of a bottom-up saliency map in the primary visual cortex. *Journal of Vision, 8*: 1.

Zhou, G., Zhang, L., Liu, J., Yang, J., & Qu, Z. (2010). Specificity of face processing without awareness. *Consciousness and Cognition*, *19*(1), 408-412. doi: 10.1016/j.concog.2009.12.009

Zhou, W., Jiang, Y., He, S., & Chen, D. (2010). Olfaction modulates visual perception in binocular rivalry. *Current Biology*, *20*(15), 1356-1358. doi:10.1016/j.cub.2010.05.059

Zimba, L. D., & Blake, R. (1983). Binocular rivalry and semantic processing: out of sight, out of mind. *Journal of Experimental Psychology: Human Perception and Performance*, *9*(5), 807. doi:10.1037/0096-1523.9.5.807

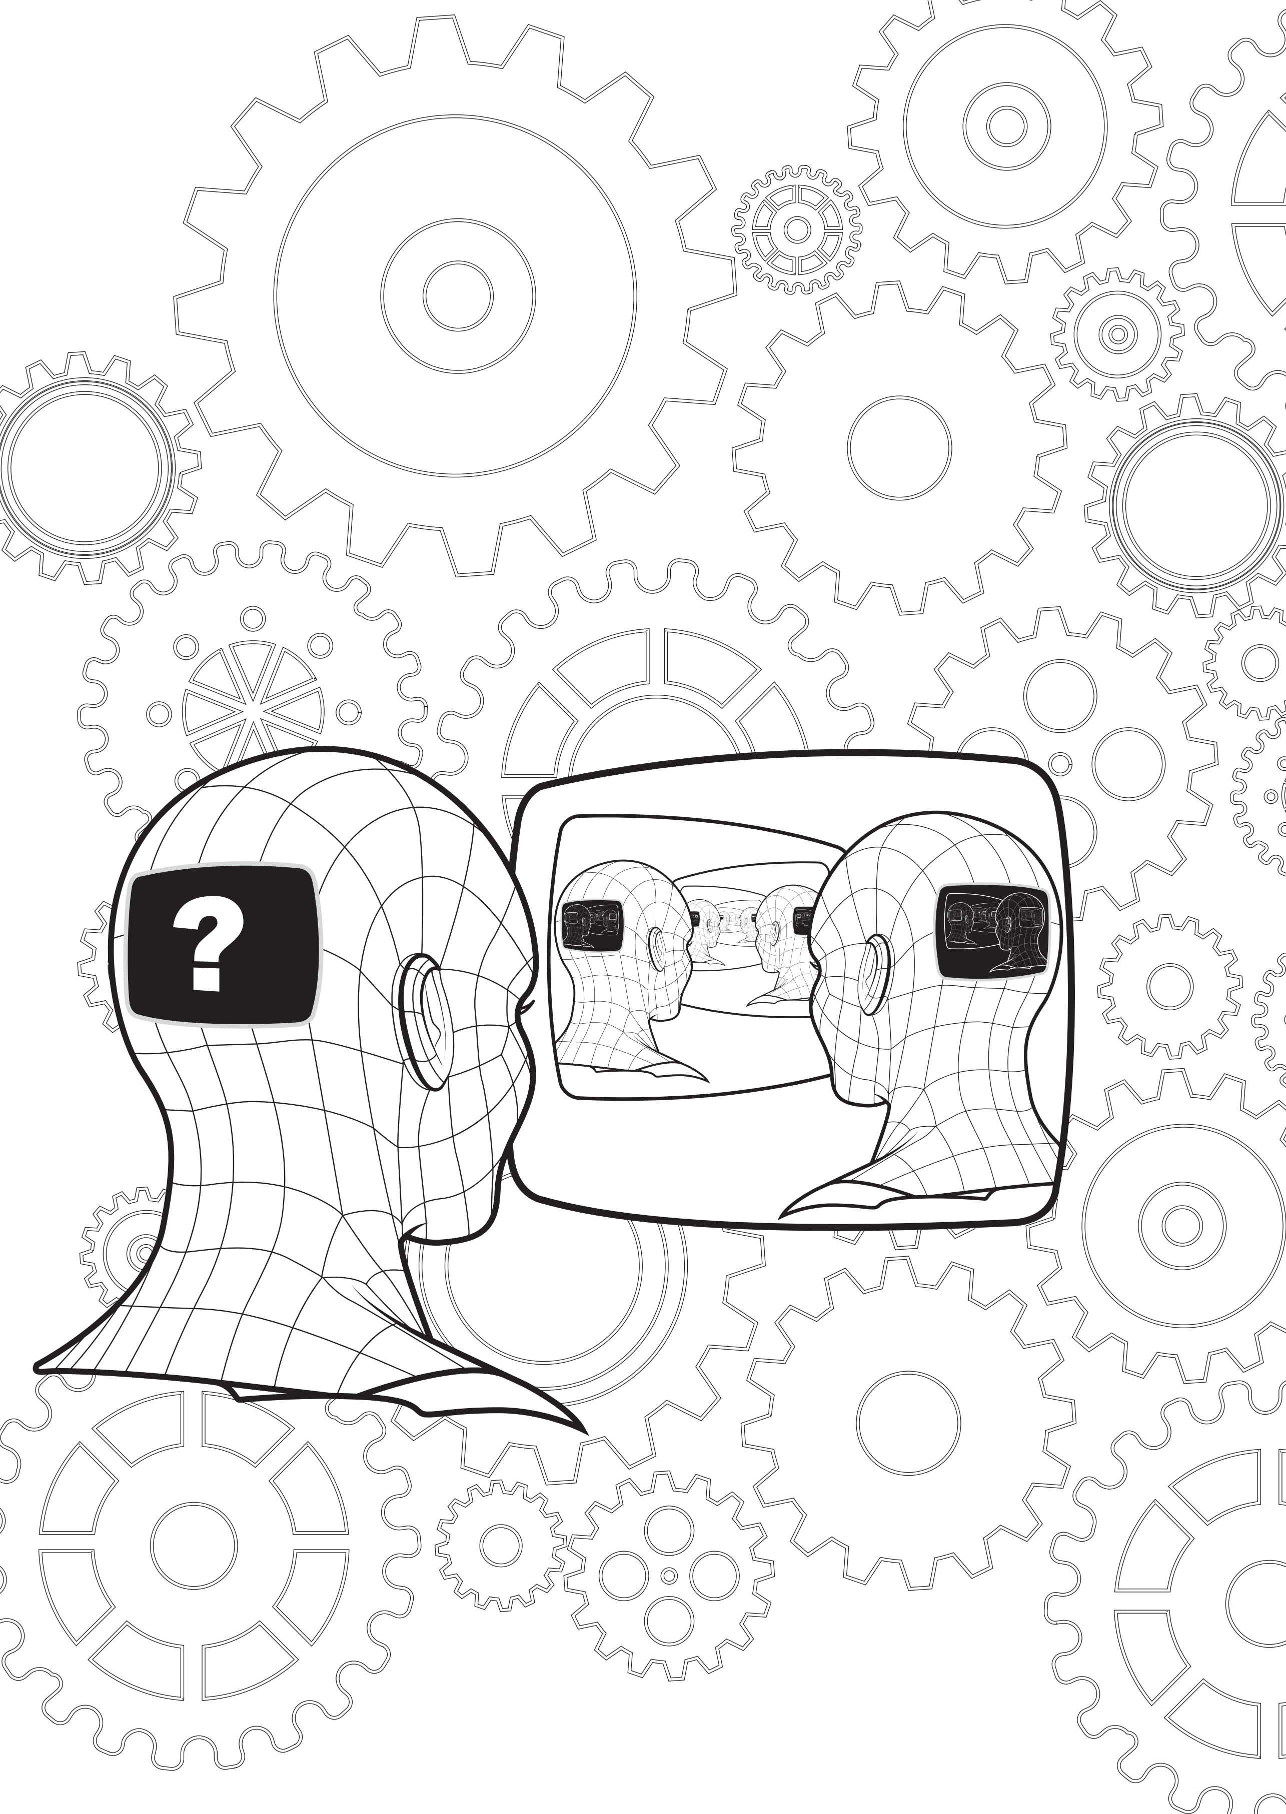

Supplementary materials

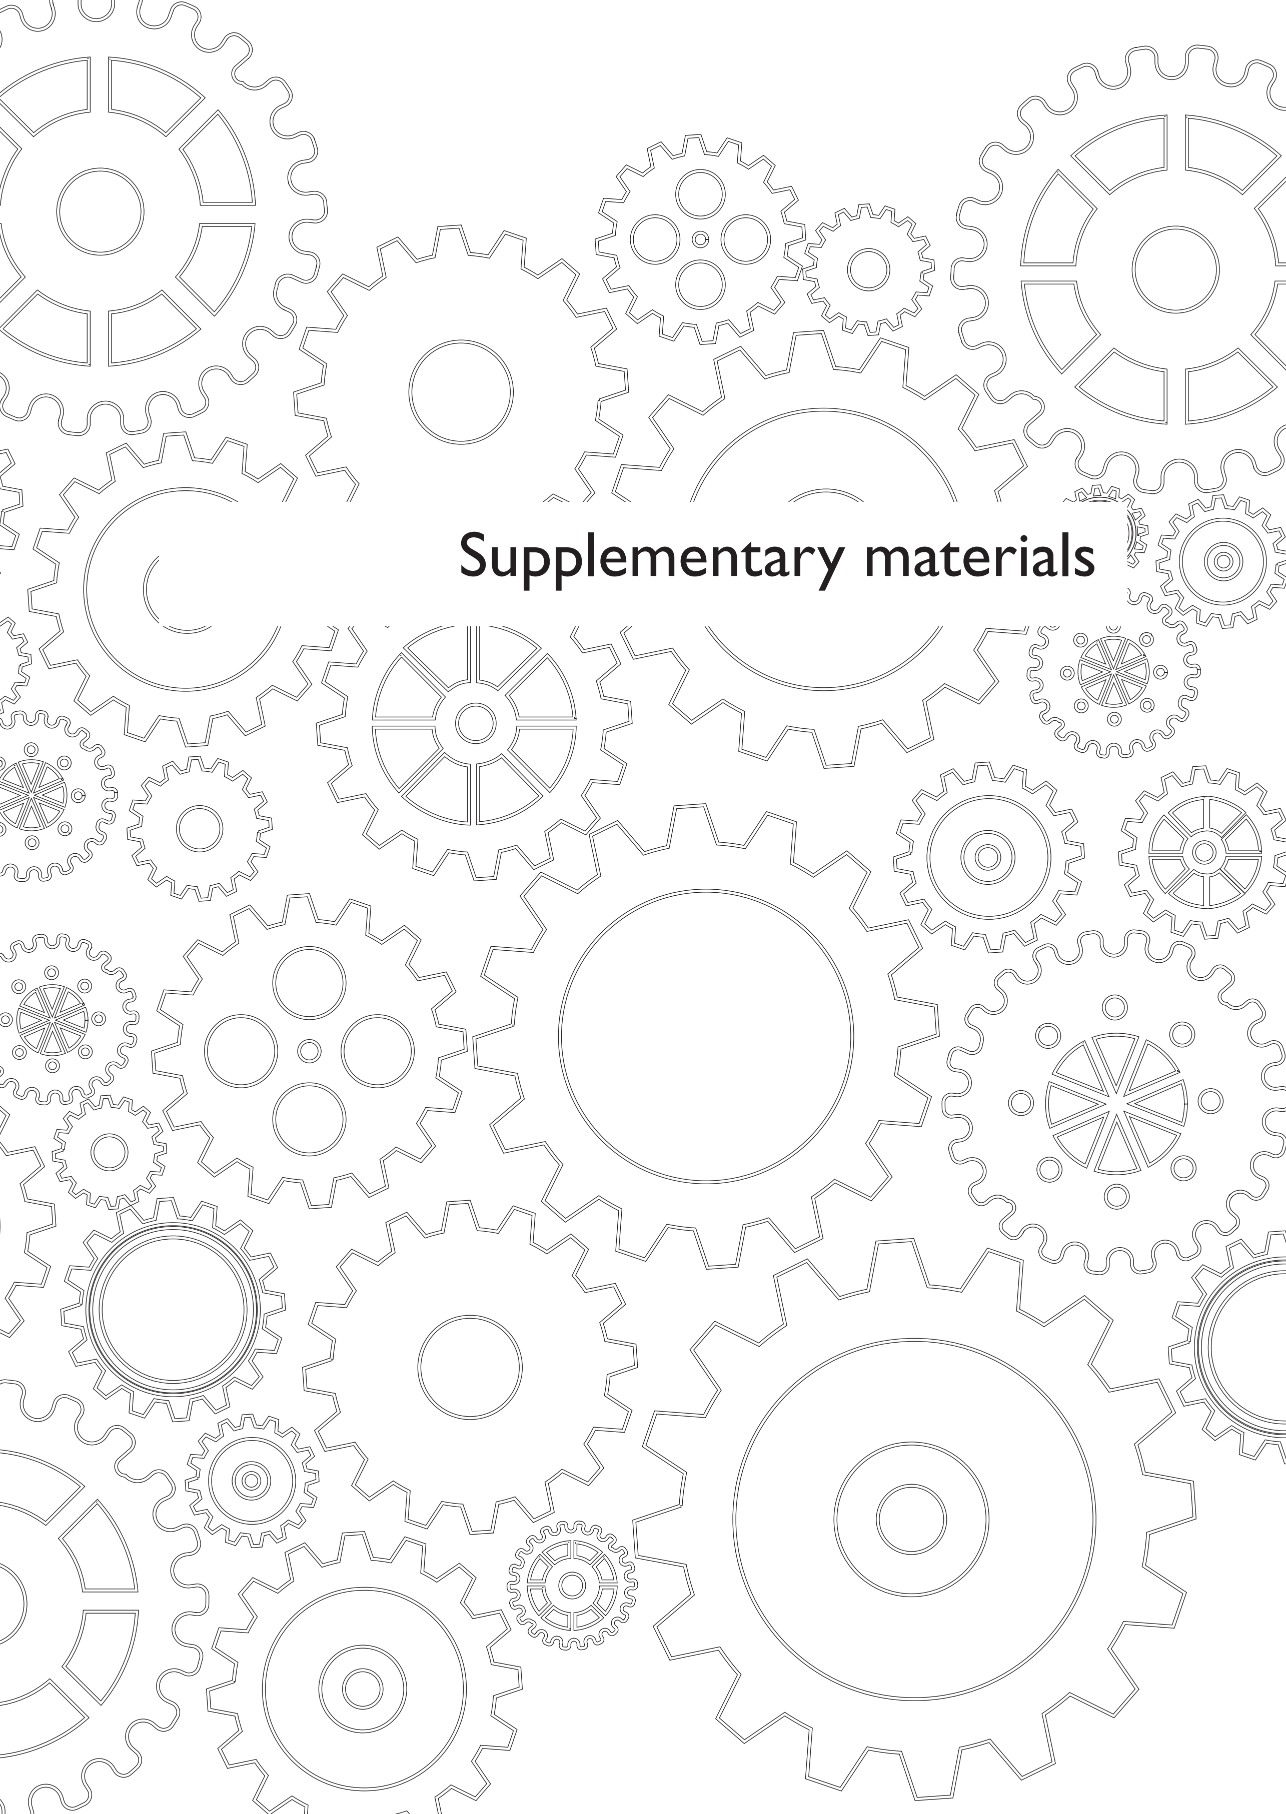

Chapter 1

S1: Mixed versus blocked control

In the experimental conditions of Experiment 1 and 2 supraliminal and subliminal trials were intermixed, whereas only flash-suppressed cues were presented in the control tasks. Possibly, the intermixed visible arrow cues facilitated breakthrough of the suppressed arrow cues in Experiment 1 and 2 (target detection task), as a result of which the control task (cue detection task) would have underestimated the visibility of the suppressed arrow cues. To test for this possibility we ran an additional control experiment (N = 8) in which participants were instructed to report the direction of the arrow as soon as they could discern it. One experimental session was comprised of 120 flash-suppressed trials (blocked condition), and was identical to the control condition in Experiment 1 and 2. In the other experimental session visible trials were intermixed (intermixed condition) in such a way that the stimulus presentation was identical to the target-detection task of Experiment 1.

Overall, 85% of the subliminal arrow cues were detected within 6 seconds, after which a new trial was initiated if no response was given. These trials yielded an average detection time of 1.7 seconds ($SD = 1.1$). Importantly, there was no difference between the blocked and mixed presentation on detection times of arrow directions; neither mean detection time of arrow cues that were perceived within 6 seconds (all SOA's: $p > 0.5$), nor the number of masked arrows that were not seen within 6 seconds ($p = 0.08$) differed between presentation conditions. However, there was a trend that more (masked) arrows were perceived in the blocked design, as compared with the mixed design. This suggests that, by using a blocked design in the cue detection task (the control task) we might have overestimated the number of unsuccessfully masked arrow cues in the (mixed) target detection task of Experiment 1 and 2. Hence, our initial approach appears somewhat conservative in detecting non-conscious cueing effects.

Finally, this supplementary control experiment was used to explicitly test our assumption that supraliminally presented arrow cues were detected faster than subliminally presented cues. This assumption was confirmed in a repeated measures ANOVA with the factors SOA and Cue Visibility by a main effect of Cue Visibility, $F(1, 7) = 141.89$, $p < 0.001$, $\eta^2 = 0.95$. Specifically, subsequent paired-samples t-tests revealed that in the SOA 100 condition, subliminal cues were detected 1191 ms

(*SD* = 264) later than supraliminal cues, $t(7) = 12.77$, $p < 0.001$. Similarly, in the SOA 500 condition, subliminal cues were detected 1241 ms (*SD* = 370) later than supraliminal cues, $t(7) = 9.47$, $p < 0.001$.

S2: Analyses of subliminal trials in Experiment 2

In the result section of Experiment 2 only the supraliminal conditions were analyzed. Firstly, because the main objective was to test whether subliminally presented statistical context would affect the utilization of intermixed supraliminal cues. Secondly, because we mainly provide analyses of cueing effects (congruent RTs subtracted from incongruent RTs) and there were no incongruent trials in the subliminal predictive condition. While investigating cueing effects on subliminal trials was not the main purpose of Experiment 2, our general conclusions lead us to assume that the predictive value of the cues would have no effect on RTs in subliminal trials. To test this assumption we compared the RTs of the three subliminal conditions of Experiment 2 as a supplementary analysis: predictive congruent, non-predictive congruent and non-predictive incongruent. Since an SOA of 100 ms led to longer RTs (*M* = 379 ms, *SD* = 42.5) than an SOA of 500 ms (*M* = 350 ms, *SD* = 40.8), $F(1, 12) = 39.90$, $p < 0.001$, $\eta^2 = 0.77$, we analyzed both SOA conditions separately.

Within the SOA 100 condition none of the three pairwise comparisons were significant: Congruent RTs in the predictive condition did not differ from congruent RTs in the non-predictive condition (*M* = 2.44 ms, *SD* = 27.9), $t(12) = 0.31$, $p = 0.759$, $d = 0.08$; Congruent RTs in the non-predictive condition did not differ from incongruent RTs in the non-predictive condition (*M* = .77 ms, *SD* = 8.71), $t(12) = 0.32$, $p = 0.755$, $d = 0.09$; Congruent RTs in the predictive condition did not differ from incongruent RTs in the non-predictive condition (*M* = 3.21 ms, *SD* = 25.5), $t(12) = 0.454$, $p = 0.658$, $d = 0.13$.

Next, we compared RTs in the three subliminal conditions (predictive congruent, non-predictive congruent and non-predictive incongruent) for trials with an SOA of 500 ms. Again, none of the pairwise comparisons yielded significant results: Congruent RTs in the predictive condition did not differ from congruent RTs in the non-predictive condition (*M* = 1.34 ms, *SD* = 19.2), $t(12) = 0.25$, $p = 0.806$, $d = 0.07$; Congruent RTs in the non-predictive condition did not differ from incongruent RTs in the non-predictive condition (*M* = -0.79 ms, *SD* = 6.14), $t(12) = 0.46$, $p = 0.651$, $d = 0.13$; Congruent RTs in the predictive condition did not differ from incongruent

RTs in the non-predictive condition ($M = 0.55$ ms, $SD = 18.3$), $t(12) = 0.11$, $p = 0.916$, $d = 0.03$.

In sum, there was no effect of predictive value on subliminal trials, nor was there an effect of cue-target congruency. This is in line with our general conclusion that visible statistical relevance is required for subliminal cue utilization to occur.

S3: Supplementary tables

Table S1. *Mean RTs in ms (SD) and percentage of errors (SD) for all conditions in Experiment 1.*

	100 ms SOA							
	Supraliminal				Subliminal			
	Non-predictive		Predictive		Non-predictive		Predictive	
	Congr.	Incon.	Congr.	Incon.	Congr.	Incon.	Congr.	Incon.
Errors	0.90	0.83	0.28	3.85	0.77	0.96	0.96	0.96
(SD)	(1.20)	(1.02)	(0.34)	(4.41)	(1.58)	(2.40)	(1.63)	(1.63)
R.T.	390	400	390	422	396	400	406	410
(SD)	(62)	(62)	(48)	(49)	(61)	(62)	(45)	(44)
	500 ms SOA							
Errors	0.96	1.67	0.48	3.53	0.77	1.35	0.77	0.58
(SD)	(0.89)	(2.41)	(0.99)	(4.02)	(1.58)	(1.65)	(1.58)	(1.10)
R.T.	352	384	337	420	371	378	378	394
(SD)	(50)	(49)	(38)	(42)	(46)	(45)	(53)	(54)

Table S2. *Mean RTs in ms (SD) and percentage of errors (SD) for all conditions in Experiment 2.*

	100 ms SOA							
	Supraliminal				Subliminal			
	Non-predictive		Predictive		Non-predictive		Predictive	
	Congr.	Incon.	Congr.	Incon.	Congr.	Incon.	Congr.	Incon.
Errors	0.77	0.77	1.15	1.35	1.09	1.09	0.58	n.a.
(SD)	(2.14)	(1.58)	(2.42)	(1.94)	(1.54)	(1.84)	(0.79)	
R.T.	370	385	377	391	379	378	381	n.a.
(SD)	(48)	(47)	(44)	(43)	(45)	(45)	(43)	
	500 ms SOA							
Errors	0.96	2.88	0.96	4.23	1.47	1.79	1.22	n.a.
(SD)	(2.17)	(4.98)	(2.17)	(6.57)	(2.41)	(2.38)	(1.39)	
R.T.	334	364	336	363	349	350	351	n.a.
(SD)	(48)	(41)	(44)	(38)	(40)	(41)	(44)	

Chapter 4

S1: Analyses with alternative inclusion criteria

To assert the robustness of the present findings, and to demonstrate that these findings are not the result of selective data or participant inclusion, the same set of analyses was conducted with more conservative inclusion criteria. The four participants of which only one Eye condition allowed for computing medians were excluded from these additional analyses. As such the present analyses are based on 12 participants (6 males; 22 years of age, $SD = 3.5$), for seven of which the red annulus was associated with electric shocks in the acquisition phase of the experiment. The group results are presented in Figure S1 and the results for individual participants are presented in Figure S2.

To examine whether classical fear conditioning affected the time it took for conditioned stimuli to break through continuous flash suppression, we conducted a 2 * 2 within-subject analysis of variance, with the factors Conditioning (CS+ and CS-) and Suppression (interocular suppression or monocular presentation) and the between-subjects factor Color (the red or blue annulus was paired with electric shocks). This analysis revealed a main effect of Conditioning, $F(1, 10) = 11.57$, $p = 0.007$, $\eta^2 = 0.54$, showing that, irrespective of the Suppression condition, trials in which a CS+ annulus was presented ($M = 1625$ ms, $SD = 229$) yielded faster response times than trials in which a CS- annulus was presented ($M = 1738$, $SD = 381$). The absence of a main effect of Suppression ($p > 0.9$) indicated that we successfully matched the response times of trials in which the targets were interocularly suppressed ($M = 1625$ ms, $SD = 398$), and trials in which they were not ($M = 1656$ ms, $SD = 249$). None of the interactions with the between-subjects factor Color reached significance (all p's > 0.2), indicating that none of the effects reported here depended on which specific annulus color was associated with shocks. Finally, there was an interaction between Suppression and Conditioning on response times, $F(1, 10) = 22.15$, $p < 0.001$, $\eta^2 = 0.69$, revealing that the effect of fear conditioning on response times was dependent on whether the conditioned annuli were interocularly suppressed or not.

To further investigate the nature of this interaction we subsequently conducted paired-samples t-tests. This revealed that when targets were interocularly suppressed, discrimination of the oriented gratings was faster when they were surrounded by a CS+ annulus (*M* = 1440 ms, *SD* = 317) than a CS- annulus (*M* = 1810 ms, *SD* = 524), *t*(11) = 3.83, *p* = 0.003, Cohen's *d* = 1.37. When the grating and annuli were not interocularly suppressed (i.e. in the monocular condition), however, response times did not differ between CS+ trials (*M* = 1648 ms, *SD* = 246) and CS- trials (M = 1665 ms, *SD* = 273; *p* > 0.7). These findings demonstrate that stimuli that were previously paired with aversive stimulation, and therefore signal threat, are released from interocular suppression (370 ms) faster than stimuli that were not paired with a shock.

The magnitude of the conditioning effect (response times on CS- minus response times on CS+) in trials with interocular suppression correlated with participants' overall average response times, *r*(10) = 0.72, *p* = 0.008. This means that participants with slower response times showed a larger difference between Conditioning conditions. In order to remove the between-subject variability in response speed from the effect of conditioning, we computed the normalized difference between response times on CS+ trials and on CS- trials for every subject (CS- minus CS+ divided by CS- plus CS+). This was done for both Suppression conditions independently. One-sample t-tests with the normalized data revealed that CS+ trials were indeed responded to 20.5% (*SD* = 18.1) faster than CS- trials, when the annuli were interocularly suppressed, *t*(11) = 3.89, *p* = 0.003, *d* = 1.12, but not when they were presented to the same eye as the masks (*M* = 0.4%, *SD* = 4.5, *p* > 0.7). A paired-samples t-test between the normalized data of suppression trials and monocular trials further substantiated the interaction observed in the main analyses, *t*(11) = 4.84, *p* < 0.001, *d* = 1.74.

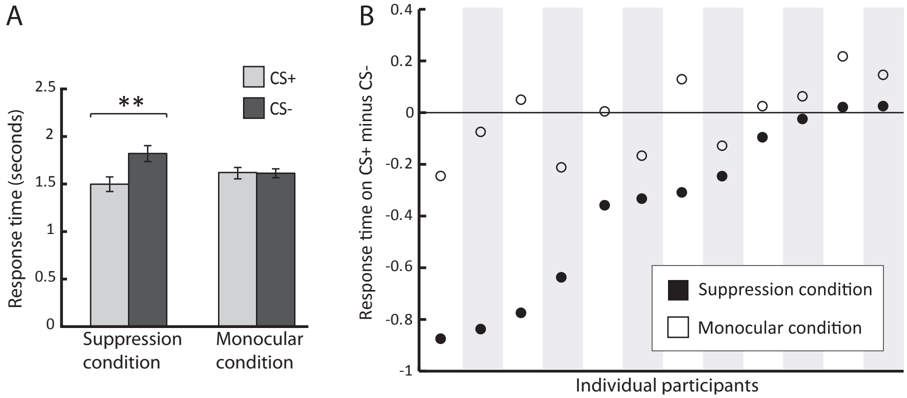

Figure S1. Group results (A) and individual results (B) for 12 included participants. In panel A, the y-axis represents the response time as a function of presentation condition (labeled on the x-axis) and fear conditioning (shades of gray). In panel B, the y-axis represents the difference in response time between CS+ trials and CS- trials for each participant, depicted along the x-axis. ***p* < 0.001

S2: Linear mixed-effects model analysis

In light of the relatively small amount of trials per condition, we conducted a linear mixed-effects model analysis to further substantiate the robustness of the present findings. We entered the raw response times of all trials that yielded a correct response within the 4 seconds deadline. In the Suppression condition, this included 13.2 (*SD* = 3.2, range 8:16) CS+ trials and 13.0 (*SD* = 2.8, range 8-16) CS- trials, and in the Monocular condition, this included 14.3 (*SD* = 2.2, range 10-16) CS+ trials and 14.1 (*SD* = 2.3, range 11-16) CS- trials per condition. Participants were entered as a random factor, whereas the within-subject factors Conditioning and Suppression were entered as fixed factors. This analysis revealed a main effect of Conditioning, $F(1,944) = 11.166$, $p = 0.001$, no main effect of Suppression, $F(1,944) = 2.674$, $p = 0.102$, and an interaction between Conditioning and Suppression, $F(1,944) = 6.060$, $p = 0.014$. To further investigate the nature of this investigation we conducted subsequent linear mixed-effects model analyses for the two levels of the factor Suppression separately. This revealed an effect of Conditioning in the Suppression condition, $F(1,497) = 8.694$, $p = 0.003$, but not in the Monocular control condition, $F(1,564) = 0.874$, $p = 0.350$, which corroborates our findings of the main analyses. Finally, to account for the skew in the response time distributions, we conducted the same analysis on log-transformed response times (for a similar approach, see Stein, Kaiser, & Peelen, 2015; Peremen, & Lamy, 2014). This analysis also revealed a main effect of Conditioning, $F(1,944) = 10.672$, $p = 0.001$, no main effect of Suppression, $F(1,944) = 1.114$, $p = 0.291$, and an interaction between Conditioning and Suppression, $F(1,944) = 6.027$, $p = 0.014$.

Chapter 7

S1: Decoding accuracies per region of interest

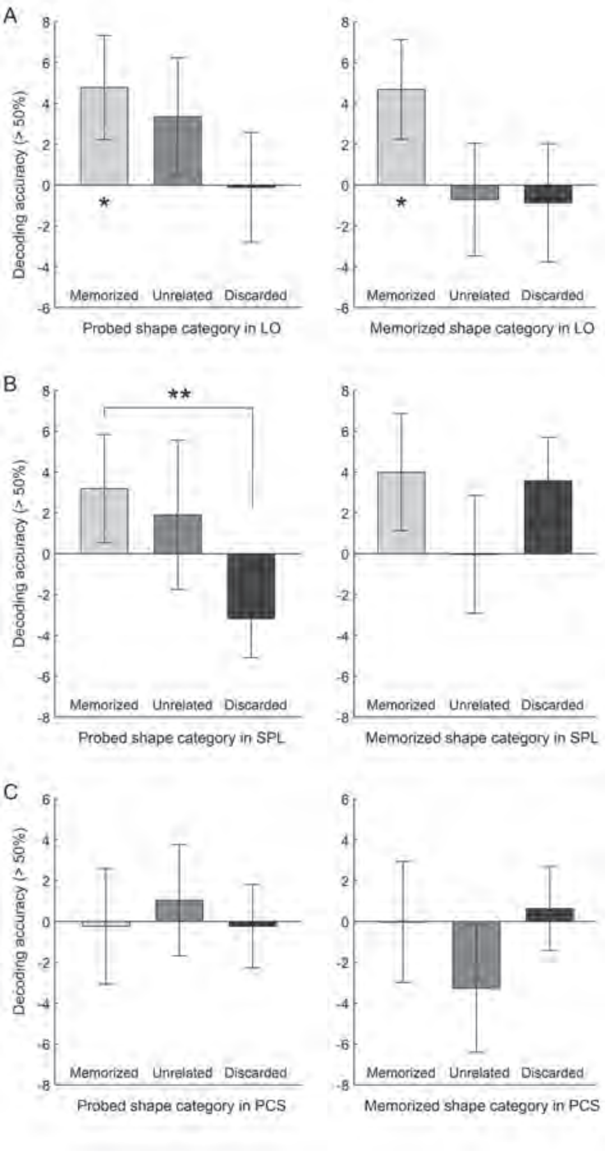

Figure S1. Classifier performance for decoding the shape category of the probe (left) and the shape category maintained in working memory (right) in Experiment 1. Decoding was done separately for three different ROI's, comprised of the significant voxels in the Stimulus > Baseline contrast of the functional localizer run in the lateral occipital cortex (panel A), the superior parietal lobule (panel B) and the precentral sulcus (panel C). These ROI's correspond to the ROI's used in the univariate analyses, as depicted in Figure 3C. $^*p < 0.1$, $^{**}p < 0.05$

S2: Behavioral experiment

Introduction

In Experiment 2 the content of VWM did not modulate the neural response to visual input that was rendered invisible by continuous flash suppression (CFS; Tsuchiya, & Koch, 2005). The data that we present here is aimed at validating the current stimulus set. To ensure that the present stimulus-set allows for finding a faster release from CFS for probes that match rather than mismatch the content of VWM, we replicated the original paradigm of Gayet et al. (2013; who used red, green, and blue color categories) with the stimulus set used in Experiments 1 and 2 (i.e., triangle, ellipse and rectangle shape categories).

Methods

A group of 19 participants took part in the behavioral experiment. They were drawn from the same pool of participants as those who participated in the fMRI experiments. As in Experiment 2, participants were instructed to memorize the cued prime for subsequent recall, and a CFS suppressed probe was presented during the delay interval. The memory task was identical to that of the fMRI experiments. In this experiment, however, participants were additionally instructed to report the location of the probe (left or right of fixation) by means of a key press, as soon as they could discern its location. We expected that probes that match the shape category of the cued prime (Memorized condition), would be released from suppression faster, and therefore elicit faster response times, than probes that either matched the shape category of the uncued prime (Discarded condition) or matched neither shape category (Unrelated condition). A schematic depiction of a trial is provided in Figure S2.

Participants were positioned in front of a dichoptic mirror set-up mounted on a chin-rest, ensuring a fixed viewing distance (70cm) and independent stimulation of both eyes. A Brownian noise frame that surrounded the presentation area was presented to the center of view of each eye, such as to facilitate binocular fusion. The behavioral experiment followed the exact same procedure, stimulus presentation and stimulus timing as the fMRI experiments, except for two differences. First, all shape stimuli (primes, probes and test-stimuli) had a height and width of about 0.8 degrees of visual angle, and the probes were presented at an eccentricity of 1.3 degrees of visual angle from the fixation dot. Second, the probe was presented on the screen between 0.3 and 0.6 seconds ($M = 0.45$) after the onset of the CFS, and ramped up from zero to full intensity (about 30% Michelson contrast) in one second, after which the stimuli remained on screen for five seconds, or until the participant had made a response.

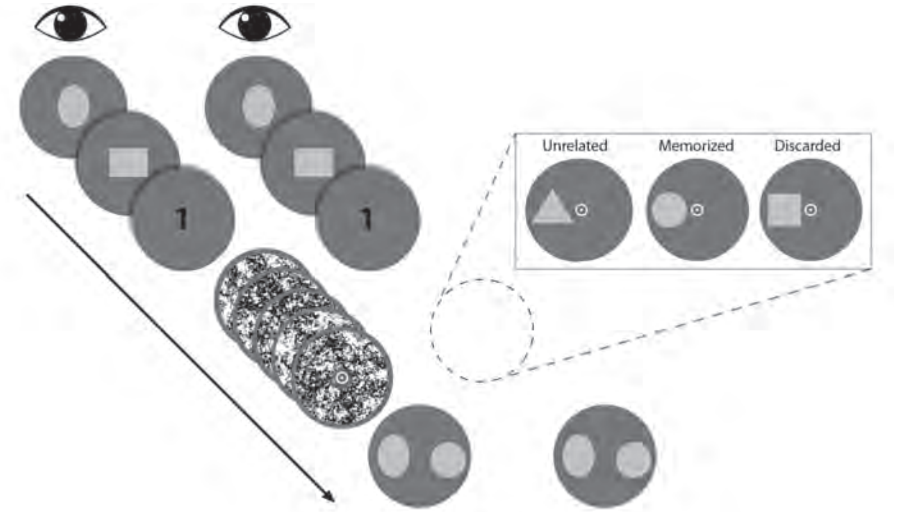

Figure S2. Schematic depiction of a trial in the behavioral experiment. Similarly to the fMRI experiments, participants were presented with two different prime shapes and a cue, which indicated which of the two primes they should memorize for subsequent recall. Again, during the delay interval a probe was presented that could match the cued (i.e., Memorized) prime, the uncued (i.e., Discarded) prime or that matched neither (i.e., Unrelated to the memory task). After the onset of the CFS masks, the probe was ramped up from zero to full intensity in one seconded, and remained on the screen until five seconds had elapsed or until participants reported whether it was presented left or right of fixation.

Results and discussion

Participants were on average 65.1% (*SD* = 8.0) correct on the memory task, and indicated the location of the probe correctly on 99.1% (*SD* = 2.0) of the trials. All trials in which participants correctly reported the location of the probe were included in the analyses of response times. We first conducted a repeated measures ANOVA with the factor Congruence (Memorized, Discarded, Unrelated) as a three level within-subject factor. This revealed a main effect of Congruence, $F(2, 17) = 5.31$, $p = 0.016$, which demonstrates that the response times to the probe were modulated by its match with the concurrent content of VWM. To further investigate the nature of this main effect between the different levels of the factor Congruence, we normalized the data, to account for the between subject variability in response times (for a similar approach, see Stein, 2012; Gayet, Paffen, Belopolsky, Theeuwes, & Van der Stigchel, 2016). The normalized difference between condition A and B was obtained as follows: A versus B = (A-B)/(A+B). One-sample t-tests against no difference revealed that probes in the Memorized condition (*M* = 1657 ms, *SD* = 510) were reported on average 11% (*SD* = 11) faster than probes in the Discarded condition (*M* = 1766 ms, *SD* = 616), $t(18) = 4.42$, $p < 0.001$, and 12% (*SD* = 12) faster than probes in the Unrelated condition, $t(18) = 2.23$, $p = 0.039$. There was no difference in response times for probes in the Unrelated condition compared to probes in the Discarded condition (*M* = 2%, *SD* = 12, $p > 0.7$). These results are depicted in Figure S3. Numerically, response times were faster in the Memorized condition than in the Discarded condition for 18 out of 19 participants. As such, the present stimulus set allows for modulating the behavioral response to CFS suppressed probes by the content of VWM. This demonstrates that not only colors (Gayet et al., 2013) and faces (Pan et al., 2014) but also shapes gain prioritized access to awareness when they match the content of VWM. More importantly for the present purpose, this finding validates the stimulus set used in Experiment 2.

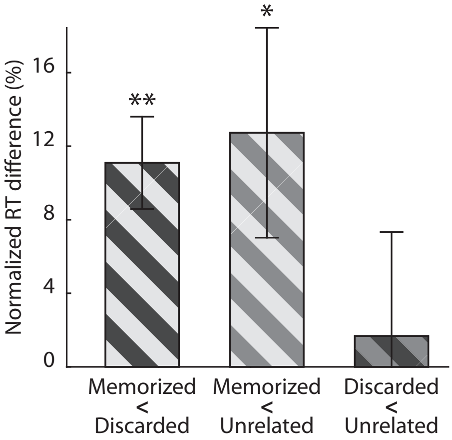

Figure S3. Normalized response time difference in the behavioral experiment. Positive values indicate a decrease in response times for the conditions at the top relative to the bottom conditions. Error bars depict the standard error of the mean. *$p < 0.05$, **$p < 0.001$

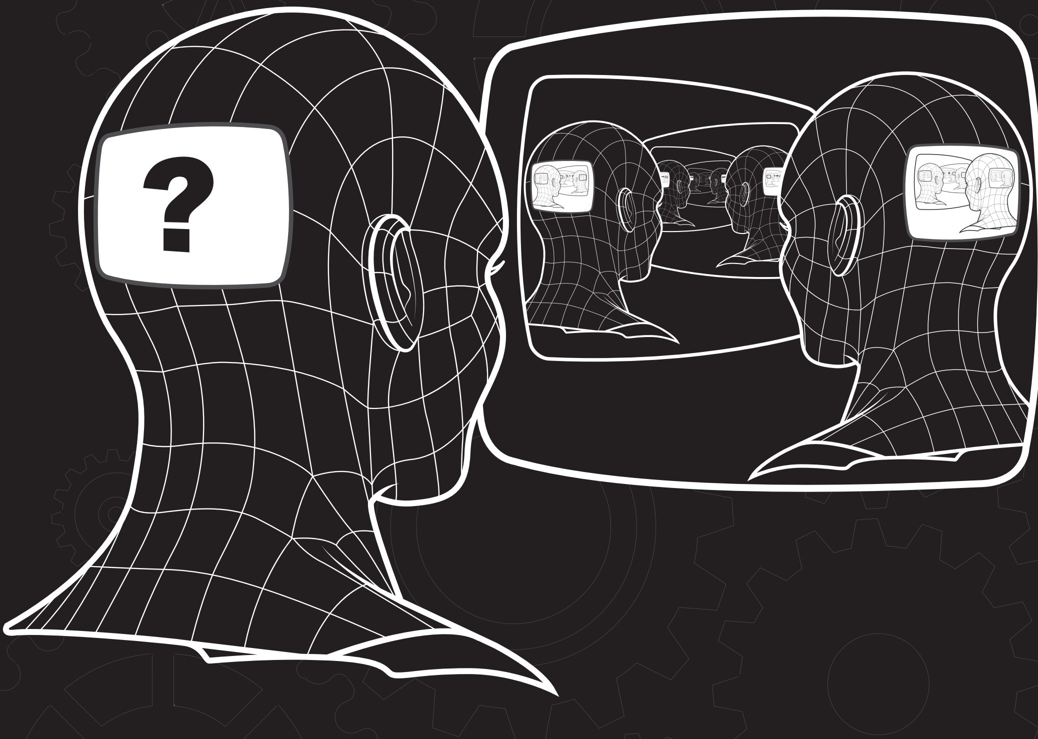

Appendix

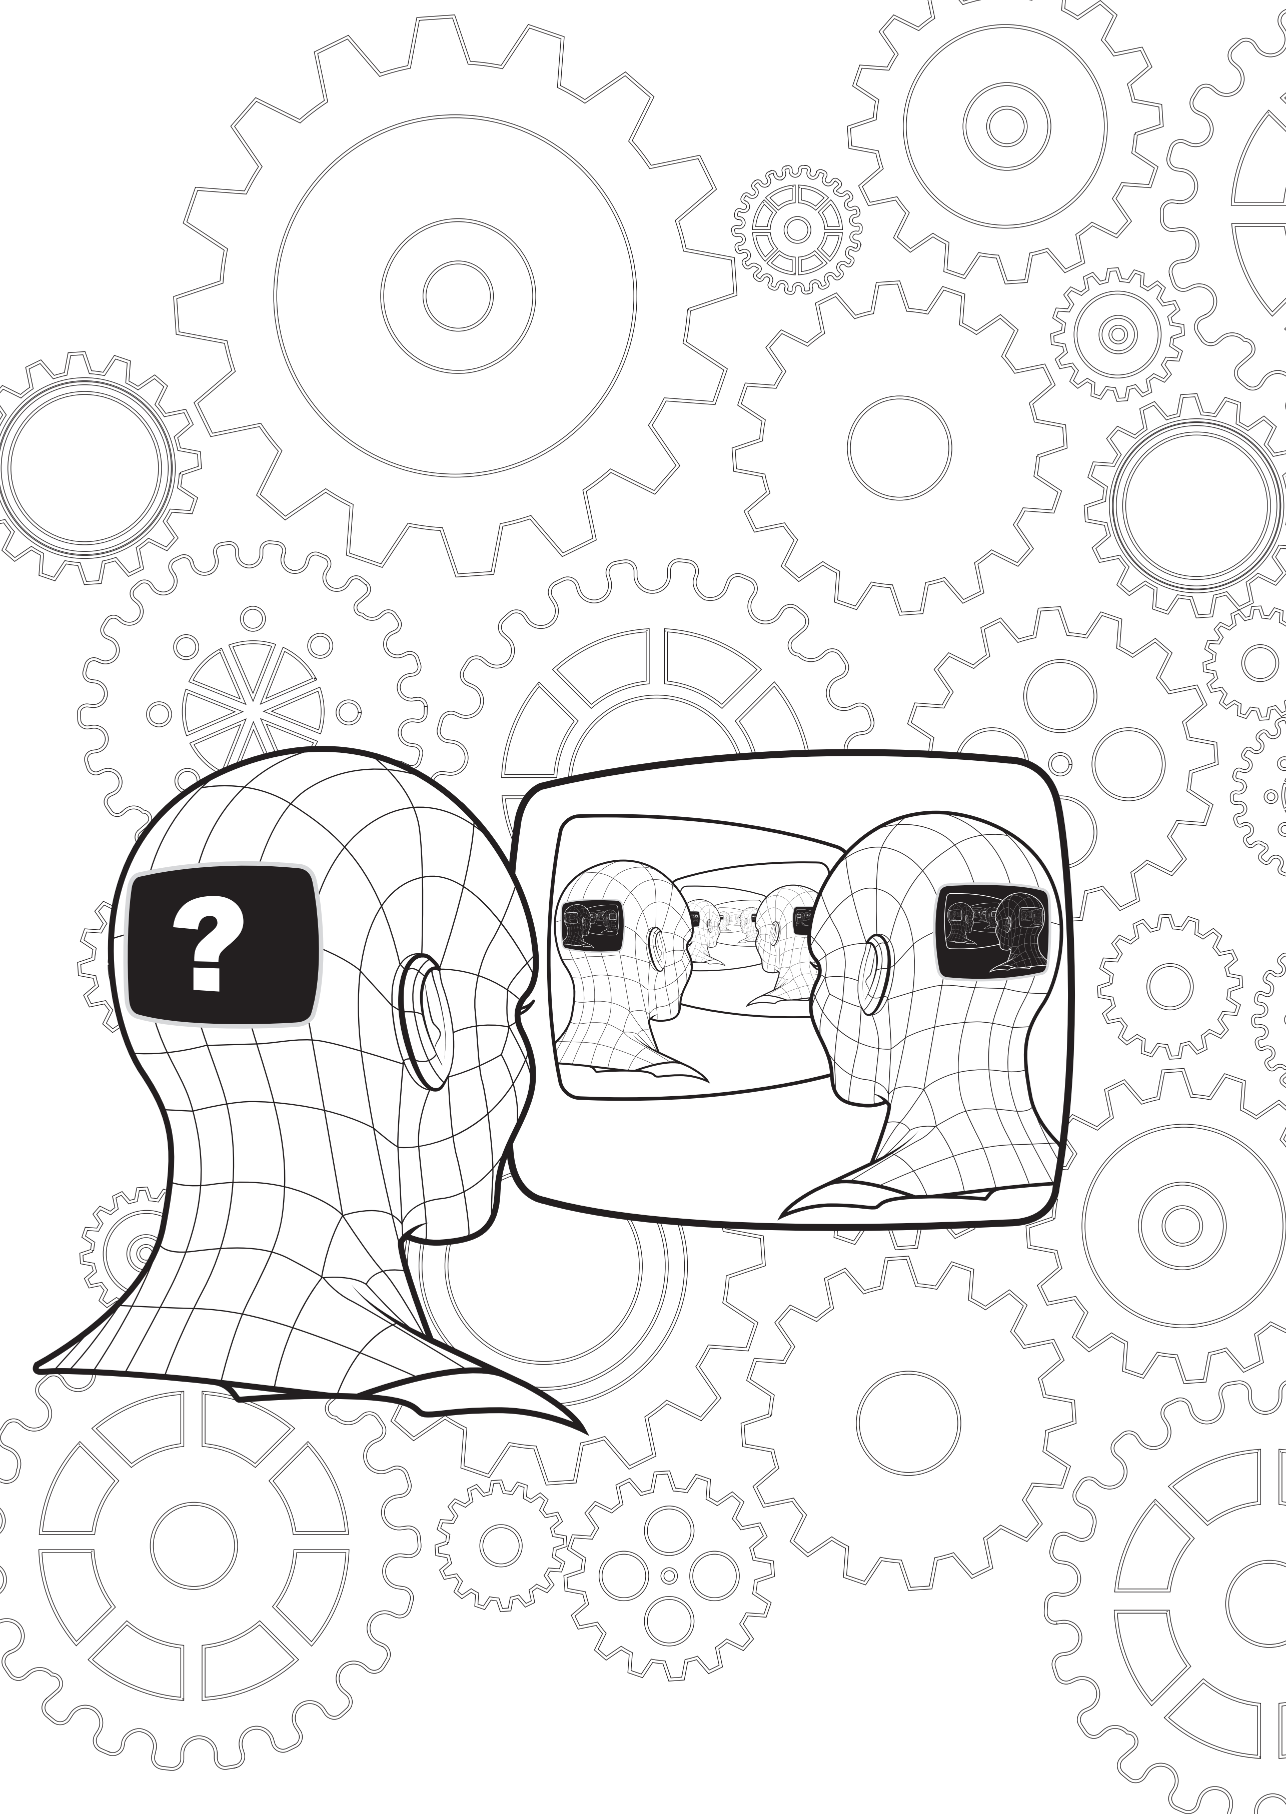

Nederlandse samenvatting
(Summary in Dutch)

Van blind oog tot geestesoog: Hoe de relevantie van visuele informatie tot bewuste waarneming leidt

Introductie

Elke seconde landen er tussen de 10^7 en 10^{13} fotonen (lichtdeeltjes) op elke vierkante millimeter van ons netvlies. Het aantal verschillende configuraties waarin deze fotonen op ons netvlies kunnen landen overschrijdt het aantal atomen in het zichtbare heelal. Als we al deze visuele informatie stap voor stap moesten interpreteren, zou ons leven te kort zijn om alleen al deze enkele zin te kunnen lezen. Gelukkig is het menselijk brein bewapend met een visueel systeem dat een deel van deze visuele informatie selecteert, en samenvat in een enkele coherente visuele ervaring: het visueel bewustzijn. Visuele informatie die bewust wordt waargenomen wordt gekenmerkt door een veel wijdverspreider patroon van hersenactiviteit dan dezelfde visuele informatie die niet bewust wordt waargenomen. Dit heeft twee belangrijke gevolgen. Ten eerste berust het merendeel van complexe mentale handelingen (zogeheten cognitieve functies) op samenwerkingen tussen verschillende hersengebieden, waardoor we deze alleen maar kunnen toepassen op visuele informatie die bewust waarneembaar is (een telefoonnummer onthouden, mentale rotatie, etc.). Ten tweede heeft de wijdverspreide hersenactiviteit als consequentie dat er maar ruimte is voor één bewuste ervaring op elk moment. Hieruit vloeit direct een cruciale vraag voort: welk deel van onze visuele wereld wordt bewust waarneembaar? In dit proefschrift heb ik getracht te achterhalen welk type informatie deze bevoorrechte positie krijgt toebedeeld, en hoe het visuele systeem dit klaarspeelt[4]. De hypothese die ten grondslag ligt aan dit proefschrift stelt dat visuele informatie die relevant is voor het gedrag van de waarnemer meer kans zou moeten maken om tot een bewuste ervaring te leiden dan visuele informatie die irrelevant is voor de waarnemer.

Hoofdstuk 1: De invloed van onbewuste visuele informatie op gedrag

Veel informatie die ons netvlies bereikt leidt nooit tot een bewuste ervaring, maar in sommige gevallen kan deze informatie wel ons gedrag beïnvloeden. In het eerste hoofdstuk van dit proefschrift heb ik onderzocht of de mate waarin dergelijke onbewuste visuele informatie ons gedrag beïnvloedt, afhangt van hoe

---

[4] Deze Nederlandstalige samenvatting dient geïnterpreteerd te worden als een versimpelde beschrijving van de belangrijkste bevindingen uit dit proefschrift op een theoretisch niveau. De samenvatting is bedoeld voor leken op het gebied van cognitieve psychologie, en derhalve zijn alle zaken die niet cruciaal zijn voor het theoretisch begrip van de bevindingen hier weggelaten.

relevant deze visuele informatie voor ons gedrag is. Hiertoe heb ik een experiment gebruikt waarbij participanten in het midden van het scherm fixeerden, en verzocht werden zo snel mogelijk aan te geven of een doelwit (een blauwe stip) links of rechts van het midden verscheen. Dit doelwit werd meestal voorafgegaan door een bewust waarneembare pijl, die ofwel in de richting wees waar het doelwit zou verschijnen ofwel in de andere richting. In sommige gevallen werd het doelwit echter voorafgegaan door een pijl die niet bewust waarneembaar was. Bij bewust waarneembare pijlen is het bekend dat participanten de locatie van een doelwit sneller kunnen rapporteren als de pijl in de richting van het doelwit wijst. Uit dit experiment bleek dat dit ook het geval was met de niet waarneembare pijlen, maar wel onder een voorwaarde. Als de bewust waarneembare pijlen relevant voor de taak waren (d.w.z. als ze in 80% van de gevallen naar de locatie van het doelwit wezen) beïnvloedden de niet bewust waarneembare pijlen ook de responssnelheid van participanten. Als de bewust waarneembare pijlen niet relevant waren voor de taak (d.w.z. als ze in 50% van de gevallen naar de locatie van het doelwit wezen) was er geen invloed van de niet-bewust waarneembare pijlen op de responssnelheid van participanten. Hieruit kan geconcludeerd worden dat onbewuste visuele informatie alleen het gedrag beïnvloedt, wanneer uit de context blijkt dat deze informatie relevant is voor het gedrag van de waarnemer.

Hoofdstuk 2: Relevantie beïnvloedt bewuste waarneming

Wanneer een waarnemer in elk oog een andere afbeelding aangeboden krijgt, bijvoorbeeld een blauw patroon in het linkeroog en een rood patroon in het rechteroog, fluctueert de bewuste waarneming stochastisch tussen de twee afbeeldingen: eerst is gedurende een paar seconden alleen het rode patroon zichtbaar, vervolgens het blauwe patroon, dan weer het rode patroon, enzovoort. Dit fenomeen wordt binoculaire rivaliteit genoemd (ofwel competitie tussen de ogen). Dit fenomeen is een waardevol instrument in het bewustzijnsonderzoek omdat de veranderende bewuste waarneming gemeten kan worden (bijvoorbeeld door de participant met drukknoppen te laten bijhouden welke afbeelding op dat moment wordt waargenomen) terwijl de visuele informatie die op het netvlies landt constant blijft. In de experimenten die in dit hoofdstuk worden beschreven werden participanten verzocht om telkens een kleur (bijvoorbeeld een bepaald type rood) te onthouden, en een tiental seconden later aan te geven welk van twee vergelijkbare kleuren identiek is aan de te-onthouden kleur. Nadat participanten de kleur hadden onthouden, maar voordat ze de kleur moesten rapporteren, kregen ze een binoculaire rivaliteitstaak, waarbij ze verzocht werden doorlopend aan te geven welk van de twee afbeeldingen ze waarnamen (met behulp van drukknoppen). De afbeelding in het ene oog bestond telkens uit de kleur die relevant was voor de gelijktijdig uit te voeren geheugentaak (bijvoorbeeld rood), terwijl de afbeelding in het andere oog uit een kleur bestond die niet relevant was voor de geheugentaak. Uit de resultaten bleek dat een kleur die relevant is voor een gelijktijdig uit te voeren geheugentaak een groter deel van de tijd bewust waarneembaar is dan een kleur die niet relevant is voor de geheugentaak. Hieruit kan geconcludeerd worden dat visuele informatie die relevant is voor het gedrag van de waarnemer een grotere kans heeft om bewust te worden waargenomen.

Hoofdstuk 3: Toegang tot bewuste waarneming meten

Uit de bevindingen van het vorige hoofdstuk bleek dat visuele informatie die relevant is voor het gedrag van de waarnemer een groter deel van de tijd bewust waarneembaar is dan niet relevante visuele informatie. Uit deze bevindingen valt echter niet op te maken of dit veroorzaakt wordt doordat (A) relevante visuele informatie eerder toegang krijgt tot het visueel bewustzijn (en dus voorrang krijgt in het visuele systeem), of dat (B) relevante visuele informatie langer bewust waarneembaar blijft wanneer het reeds bewust waarneembaar is (bijvoorbeeld door de irrelevante visuele informatie opzettelijk te onderdrukken). Om dit onderscheid te kunnen maken is er een experimentele methode ontwikkelt die voortbouwt op binoculaire rivaliteit. Deze methode (genaamd *breaking continuous flash suppression*, afgekort b-CFS) bestaat uit het aanbieden van een dynamisch patroon aan het ene oog, waardoor een afbeelding in het andere oog in eerste instantie niet bewust waarneembaar wordt (het is gemaskeerd). Na een aantal seconden neemt de maskerende werking van het dynamisch patroon af en wordt de oorspronkelijk gemaskeerde afbeelding bewust waarneembaar. Deze methode maakt het mogelijk om te vergelijken hoe snel verschillende categorieën afbeeldingen (die in eerste instantie niet bewust waarneembaar zijn) bewust waarneembaar worden: het meet de toegang tot het visueel bewustzijn. De redenering die ten grondslag ligt aan deze b-CFS methode is de volgende: als een experimentele manipulatie meebepaalt *wanneer* een afbeelding zichtbaar wordt, dan werd de verwerking van de afbeelding beïnvloed nog *voor* dat deze bewust waarneembaar was.

Hoofdstuk 4: Visuele informatie die gevaar aankondigt dringt sneller door tot het visueel bewustzijn

Een van de hoofdvragen die ik in dit proefschrift probeer te beantwoorden is of visuele informatie die relevant is voor de waarnemer eerder bewust waarneembaar wordt dan irrelevante visuele informatie. Nu we bewapend zijn met een methode om te meten hoe snel iets bewust waarneembaar wordt (b-CFS), rest mij alleen nog een methode te vinden om visuele informatie relevant te maken voor de waarnemer. Het toppunt van relevantie voor een waarnemer is, vanuit een evolutionair perspectief, visuele informatie die naderend gevaar aankondigt. In dit hoofdstuk heb ik de mate waarin neutrale visuele informatie (blauwe en rode ringen) naderend gevaar aankondigt gemanipuleerd door middel van een klassieke angstconditioneringsprocedure. Deze procedure hield in dat participanten na het zien van (bijvoorbeeld) de blauwe ring met hoge waarschijnlijkheid een vervelende stroomstoot kregen toegediend, terwijl na het zien van (bijvoorbeeld) de rode ring dit nooit gebeurde. Hierdoor associeerden participanten de kleur van één van de ringen met een naderend gevaar, en de ander niet. Na de angstconditioneringsprocedure werden er geen schokken meer gegeven, maar voerden participanten een b-CFS taak uit met blauwe en rode ringen, zodat ik kon meten welk van de twee kleuren ringen gemiddeld sneller bewust waarneembaar werd. De hoofdbevinding van dit hoofdstuk is dat visuele informatie die een naderend gevaar aankondigt tot wel 20% sneller doordringt tot het visueel bewustzijn dan soortgelijke visuele informatie die geen naderend gevaar aankondigt.

Hoofdstuk 5: Visuele informatie die relevant is voor een neventaak dringt sneller door tot het visueel bewustzijn

Uit de bevindingen van hoofdstuk twee bleek al dat visuele informatie die relevant is voor een gelijktijdig uitgevoerde geheugentaak een groter deel van de tijd bewust waarneembaar is dan visuele informatie die niet relevant is voor deze neventaak. In dit hoofdstuk toon ik aan dat visuele informatie die relevant is voor de gelijktijdig uitgevoerde geheugentaak ook daadwerkelijk (15%) sneller doordringt tot het visueel bewustzijn dan visuele informatie die niet relevant is voor deze neventaak. Een belangrijk kenmerk van deze geheugentaak is dat de relevante informatie (die nodig is om uiteindelijk te kunnen rapporteren welke specifieke kleur vooraf werd onthouden) door de participant wordt vastgehouden in het zogenaamde visuele werkgeheugen. Het visuele werkgeheugen is een hersensysteem waarmee mensen actief visuele informatie in het geestesoog vasthouden dat relevant is voor aanstaand gedrag. Wanneer visuele informatie niet (meer) vanuit de buitenwereld op ons netvlies landt zijn wij dankzij het visueel werkgeheugen toch nog in staat om in ons hoofd een beeld te creëren; een interne visuele representatie. In de volgende twee hoofstukken heb ik de relatie tussen het visueel werkgeheugen en het visueel bewustzijn verder onderzocht, om te achterhalen hoe ons visueel systeem in staat is om informatie die relevant is (voor een geheugentaak) sneller te laten doordringen tot het visuele bewustzijn.

Hoofdstuk 6: Een valse start in de visuele verwerking van relevante informatie

Evenals in de voorgaande hoofdstukken, wordt de versnelde toegang tot het visueel bewustzijn van visuele informatie die relevant is (voor de geheugentaak) gemeten als een verschil in responstijden. In deze responstijden schuilen wellicht patronen die iets meer kunnen vertellen over de hersenprocessen die ten grondslag liggen aan de responstijdverschillen tussen relevante en irrelevante afbeeldingen. Zo kan men bijvoorbeeld kijken naar de snelste responstijden, de minder snelle responstijden, of juist naar de meest langzame responstijden, en onderzoeken wanneer het responstijdverschil tussen relevante en irrelevante afbeeldingen ontstaat. Een eerste hypothese is bijvoorbeeld dat relevante afbeeldingen efficiënter worden verwerkt door het visuele systeem, en daarom sneller doordringen tot het visueel bewustzijn. In dat geval is de verwachting dat het responstijdverschil tussen relevante en irrelevante afbeeldingen op het moment van presentatie nihil is, en vervolgens toeneemt naarmate de responstijden langzamer worden. Een tweede hypothese is dat het visuele systeem een a priori voorkeur heeft voor de relevante afbeelding, waardoor de drempel om toegang te krijgen tot het visueel bewustzijn lager ligt voor relevante afbeeldingen. In dit geval is de verwachting dat het responstijdverschil al aanwezig is bij de snelste responstijden, maar vervolgens gelijk blijft voor de langzamere responstijden. Door deze twee hypothetische responstijdpatronen te vergelijken met het patroon van de geobserveerde responstijden, bleek het vijfendertigduizend keer waarschijnlijker te zijn dat de geobserveerde responstijd verklaard werd door een a priori voorkeur voor de relevante afbeelding. Hieruit kan dus geconcludeerd worden dat het responstijdverschil (en dus het verschil in toegang tot bewustzijn) tussen relevante en irrelevante visuele informatie gestoeld is op een a priori voorkeur in het visueel systeem voor de relevante visuele informatie.

Hoofdstuk 7: Sterkere neurale respons voor relevante visuele informatie

Een mogelijke interpretatie van de observatie uit hoofdstuk 6 is dat visuele informatie in dezelfde hersengebieden verwerkt wordt, ongeacht of deze voortkomt uit externe bronnen (waarneming via de ogen) of interne bronnen (bijvoorbeeld opgeroepen door dromen, visuele illusies of, zoals in dit geval, het visueel werkgeheugen). Het idee is dat elk type visuele informatie, ongeacht de bron, geassocieerd is met een kenmerkend patroon van hersenactiviteit in de visuele verwerkingsgebieden van onze hersenen: dit noemen we een visuele representatie. Wanneer een bepaalde representatie reeds geactiveerd is door het visueel werkgeheugen (de interne bron), zal een representatie die voortkomt uit visuele waarneming (de externe bron) sterker zijn als deze overeenkomt met de representatie die voortkomt uit het werkgeheugen, dan als deze niet overeenkomt.

In dit laatste hoofdstuk heb ik onderzocht of de representatie van een waargenomen afbeelding inderdaad een sterkere respons in visuele verwerkingsgebieden teweeg brengt wanneer deze overeenkomt met de inhoud van het werkgeheugen. Hiervoor heb ik een geheugentaak gebruikt die vergelijkbaar was met die van de hoofdstukken 5 en 6. De afbeelding die tijdens de geheugentaak werd aangeboden was ditmaal echter bewust waarneembaar, en participanten hoefden alleen maar naar deze afbeelding te kijken (zonder een respons te geven). Wederom kon de waargenomen afbeelding relevant zijn voor (overeenkomen met de afbeelding van) de geheugentaak of niet. Ditmaal lagen de participanten tijdens het experiment echter in een MRI-scanner. Aan de hand van functionele MRI kan de hersenactiviteit van participanten worden herleidt uit de veranderingen in de verhouding tussen zuurstofrijke en zuurstofarme hemoglobine (een eiwit dat voornamelijk in de rode bloedcellen te vinden is) die veroorzaakt worden door hersenactiviteit.

De resultaten van dit onderzoek laten zien dat de visuele hersengebieden die geactiveerd worden door een bepaald type waargenomen afbeeldingen, een sterkere respons geven wanneer een waargenomen afbeelding overeenkomt met de inhoud van het werkgeheugen dan wanneer deze afbeelding niet overeenkomt. Hieruit kan geconcludeerd worden dat visuele informatie die relevant is voor het gedrag van de waarnemer een sterkere neurale respons teweegbrengt dan irrelevante visuele informatie.

Conclusie

In de experimentele hoofdstukken van dit onderzoek heb ik aangetoond dat visuele informatie die relevant is voor de waarnemer (A) gedrag kan beïnvloeden zonder dat deze informatie bewust waarneembaar wordt en (B) sneller bewust waarneembaar wordt en langer bewust waarneembaar blijft dan visuele informatie die niet relevant is voor het gedrag van de waarnemer. Het onderscheid tussen relevante en irrelevante visuele informatie (C) berust op een a priori verschil in visuele verwerking, die er voor zorgt dat (D) visuele informatie die relevant is voor de waarnemer een sterker patroon van hersenactiviteit teweegbrengt dan visuele informatie die niet relevant is voor de waarnemer. De relevantie van visuele informatie is op verschillende manieren experimenteel gemanipuleerd. Zo manipuleerden we of de informatie relevant was voor de huidige taak (hoofdstuk 1), voor een parallelle geheugentaak (hoofdstukken 2, 5, 6, 7) of voor het detecteren van gevaar (hoofdstuk 4). De convergerende evidentie van verschillende manipulaties van relevantie en meetmethodes ondersteunen de algemene conclusie dat visuele informatie die relevant is voor de waarnemer een voorkeursbehandeling krijgt in het visuele systeem. Door dit a priori verschil zal relevante visuele informatie aan een zwak binnenkomend visueel signaal genoeg hebben om de competitie te winnen van andere (wellicht sterkere) binnenkomende visuele signalen, en zo eerder tot bewuste waarneming leiden.

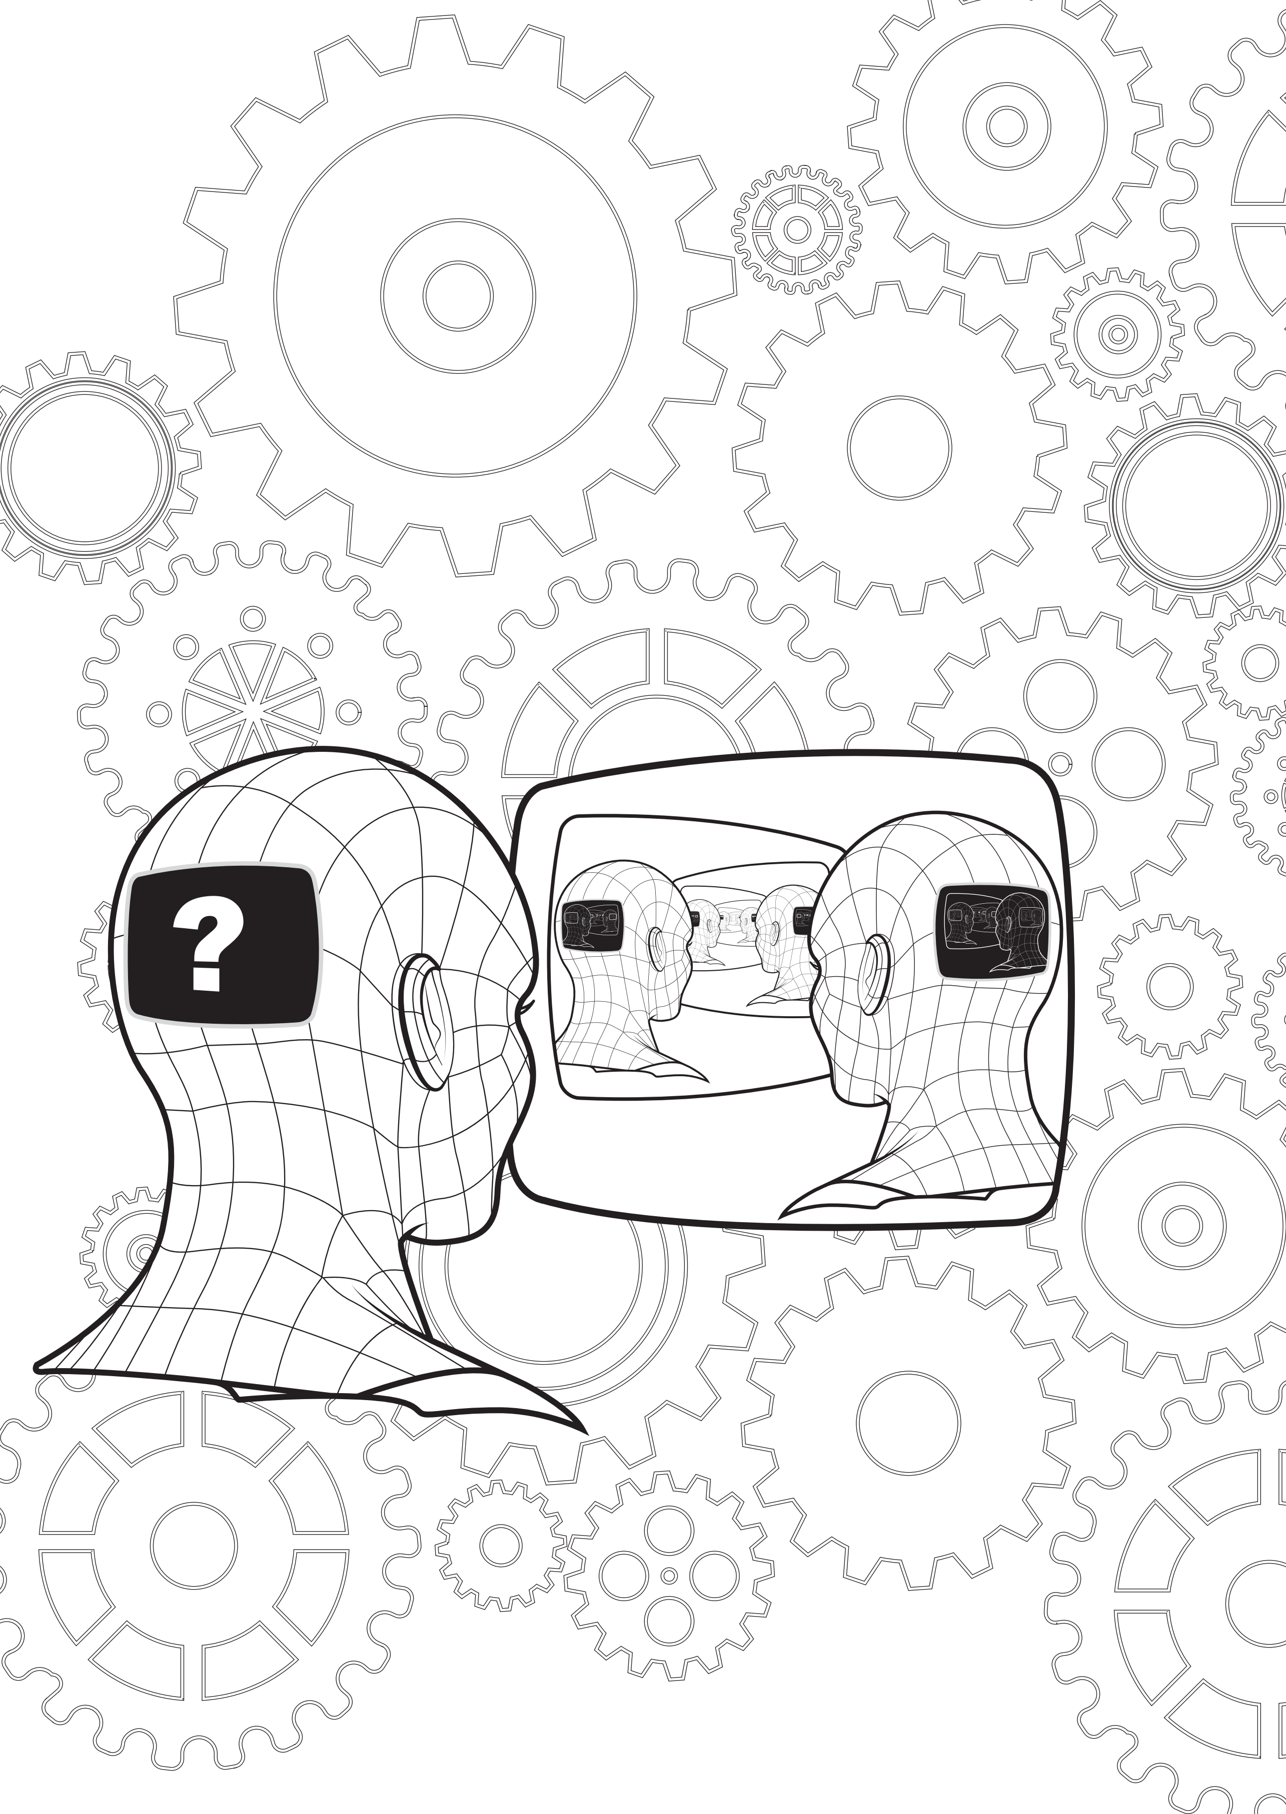

Dankwoord (Acknowledgments)

Toen ik bezig was met mijn promotieonderzoek was ik ervan overtuigd dat het schrijven van het dankwoord het makkelijkste deel van het proefschrift zou zijn. Het onderzoek is gedaan, de hoofdstukken zijn geschreven, het proefschrift is goedgekeurd, wat kan er nog fout gaan? Het gevaar is dat je met zo een dankwoord eigenlijk alleen maar mensen boos kan maken die je vergeet te noemen. Ter compensatie ga ik daarom zo nu en dan ook wat onaardigs zeggen over de mensen die ik wel benoem, zodat het relatief minder erg is als ik je vergeet.

Allereerst wil ik natuurlijk mijn twee copromotoren bedanken. **Chris** en **Stefan**, in het eerste jaar van mijn promotietraject stuurde ik jullie wel eens een mail met de aanhef "He stelletje copromotors", waarop jullie de zorg uitspraken of ik jullie wel serieus nam als begeleiders. De waarheid is dat ik geweldige begeleiding heb gehad. Ik wil jullie dan ook bedanken voor alles wat ik van jullie heb geleerd, maar ook voor jullie relaxte houding en het vertrouwen dat jullie in mij hebben getoond. Waar ik in het eerste jaar regelmatig bij jullie aanklopte ("2.5 of 3 standaarddeviaties?"), ben ik in de laatste jaren van mijn promotietraject steeds zelfdstandiger geworden ("jij begeleidt jezelf wel, he?"), en dat is grotendeels jullie verdienste. Het vertrouwen dat jullie gaven was soms wellicht wat onterecht. Als Stefan zei "dat paper schrijf je wel in een weekendje, joh" was ik 2 maanden later nog steeds met de eerste paragraaf van de inleiding bezig. Maar voor de paar keer dat ik echt iets doms deed (bijvoorbeeld een te laat opgemerkte fout in een analyse script) hebben jullie een sfeer weten te creëren waarin het geen probleem voor mij was om dat toe te geven. Meestal lachte Chris me dan gewoon keihard uit. En vreemd genoeg werkt dat heel relativerend. Waar menig AIO klaagt over nalatige begeleiders die er te lang over doen om hun manuscripten te lezen, hielden jullie regelmatig een wedstrijdje wie als eerst feedback op mijn manuscript had teruggestuurd. Dat was vaak nog dezelfde dag. Ik kijk er naar uit om komende tijd met jullie verder te werken aan alle lopende projecten.

Ook wil ik natuurlijk mijn promotoren bedanken. **Frans** en **Albert**, jullie zijn slechts zijdelings betrokken geweest bij mijn project. Desondanks waren jullie altijd bereid om vragen te beantwoorden, aanbevelingsbrieven te schrijven, of tips te geven. Ik heb altijd het idee gehad dat ik op jullie input kon rekenen wanneer ik dat kon gebruiken.

Waar is een promovendus zonder zijn paranymfen? **Inge** (Pling Plong, Pingel) en **Ruben** (Robian, Rups), jullie zijn waarschijnlijk de niet-collega's tegen wie ik het meest heb gezeurd over mijn onderzoek. Tegen Rups heb ik hele tirades gehouden over fMRI onderzoek en diffusiemodellen, terwijl we eigenlijk een optreden moesten voorbereiden. Ook hebben we onder het genot van een fles whisky tot diep in de nacht een supersnelle sinusfunctie geschreven, terwijl we ook aan een liedje hadden kunnen werken. En Inge, die heeft de laatste weken van mijn promotietraject vrijwel dagelijks mijn gestressede schrijfhoofd moeten aanschouwen. Ik woonde practisch bij jou op kantoor. Ook wil ik je bedanken voor alle botte opmerkingen en sappige roddels die we in de loop der jaren hebben gedeeld. Pling en Rups, ik ben heel dankbaar dat jullie mijn paranymfen willen zijn.

Eens in de duizend jaar zou er een legendarische Super Sayian worden geboren, wiens potentie die van al zijn soortgenoten met afstand ontstijgt. Aldus Dragon Ball Z. Als een dergelijke mythe ook in de academische wereld van toepassing was, dan zou het over **Micha** (a.k.a Mr. Adaptive Staircase) gaan. In de zes maanden dat je meewerkte aan mijn onderzoek (eerst als student, toen als onderzoeksassistent), ben je niet alleen van grote waarde geweest voor mij, maar heb je ook vele anderen op de afdeling geholpen met de kennis die je in zo korte tijd hebt opgedaan. Het is mij vooralsnog onduidelijk hoe een bachelor student van amper 20 jaar oud bekend is met de difussiemodellen van Ratcliff en de predictive coding modellen van Friston. "Het stelt niks voor hoor, ik heb alleen maar de bekendste boeken erover gelezen" was daarop je antwoord. Op die leeftijd las ik de Donald Duck.

"Ik zat in tram vijf en m'n lul stond stijf, want naast me zat een lekker wijf". De klassieke teksten van muzikant en beeldend kunstenaar **Pascal** (Def P) knalden uit de oordopjes van mijn walkman toen ik op de naschools opvang zat. Inmiddels een kwart-eeuw verder, ben ik dankbaar dat ik een aantal liedjes met mijn jeugdheld heb mogen maken, en zelfs een aantal keer met hem het podium heb mogen delen. De laatste tijd ontaarden de hiphopavonden waar we gezamenlijk aanwezig waren echter steeds vaker in discussies over het menselijk bewustzijn, de vrije wil en lichaam-geest dualisme. Die discussies zijn nog niet over, maar voortaan kom ik wel bewapend met een proefschrift, waarvan jij de geweldige cover hebt ontworpen. Dank je wel.

Wetenschappelijke ontwikkeling meemaken lukt niet in je eentje. Dat lukt alleen met inspirerende mensen om je heen. Ik heb dan ook veel geluk gehad met mijn kamergenoten op de afdeling Experimentele Psychologie. Stuk voor stuk waren jullie ideale sparringpartners om mee te klankborden. Allereerst, gewaardeerde kamergenoot en mede-Pafje **Jim** (a.k.a. Bonobo Jim, a.k.a. Doe-je-koplampen-aan, a.k.a. Mr. Ringloos), van onze gesprekken laat ik citaties achterwege. Ik kan me namelijk niet voorstellen dat er drukkers zijn die akkoord zouden gaan met de grofheid van de grappen die we op dagelijkse basis door het kantoor heen smijten. Versace Versace Versace (dat is ghetto voor "mijn dank is groot"). Een van de hoogtepunten van mijn jonge academische carriere was het tripje naar Miami Beach en de daarop aansluitende cruise naar de Bahama's met **Sjoerd** (a.k.a. Steward Steward, a.k.a. Mr. Pocahontas). In de reflectie van de kajuit kwam Sjoerd er achter dat hij, tot zijn verbazing, weldegelijk groter was dan ik. Onvergetelijke momenten. Leve de Credit Card! **Wietske** (a.k.a. Wortels, a.k.a. Ms. Barfight), wat is er beter dan een goeie zuipsessie met Tante Wortels. Of een vers gembertheetje natuurlijk. Ik zal maar niet te vaak vragen hoe het met je gaat, want het antwoord is waarschijnlijk een immer kenmerkende "Pfffff". **Carlijn**, bedankt dat je mij wilde babysitten in de eerste 2 maanden van mijn poging tot academisch pruttelen. Ik heb de laatste maanden overigens een flinke slag geslagen in onze dagelijkse lunchwedstrijd, als ik zo doorga heb ik je in 2038 ingehaald. **Martijn**, **Vivian**, ook jullie heel erg bedankt voor al het klankborden en academisch sparren de afgelopen jaren.

Een dankwoord zou incompleet zijn zonder **de hele Vision Group** te benoemen. Elk project waar ik aan heb gewerkt is beter geworden dankzij de feedback tijdens de lab meetings, en de discussies op de gang. Met name **Maarten** en **Susan** hebben het moeten ontgelden wat betreft mijn vragen over chromaticiteit en equiluminantie. **Marnix** is steeds de spreekwoordelijke sjaak als het om pupillen gaat. **Ignace** en **Casper** wil ik bedanken voor de kritische blik tijdens de labmeetings, het belang van expliciete formulering, en het model-denken in het bijzonder. **Jan** heeft meermaals als wandelende encyclopedie gefungeerd als ik weer eens een binoculaire rivaliteitsvermoeden had dat gegrond moest worden in empirische literatuur (en anders moest **Maartje** het wel ontgelden). Ook wil ik je natuurlijk bedanken voor een mooie samenwerking. **Serge**, **Alessio**, **Ben**, **Barrie**, **Wietske** en **Martijn**, bedankt voor het beantwoorden van al mijn naïeve fMRI vragen. Mede dankzij jullie input heb ik het idee dat ik me tijdens de verdediging staande zal houden. **Hinze**, jouw bewonderenswaardige snelheid van denken zou onwaarschijnlijk irritant zijn geweest, als je het niet zou inzetten om anderen te helpen. Tijdens elke lab meeting weet je weer de gedachtegang van spreker en publiek helder (en objectief) te parafraseren, in het belang van iedereen. **Jelmer**, het is jammer dat we je niet meer langs horen sloffen op de gang, moedeloos neerploffend in een vrije stoel om te klagen over een pilot. 's Avonds is het een stuk saaier op de afdeling sinds je (ongeveer 79 jaar geleden) bent vertrokken naar verre oorden. **Siarhei**, als ik vraag om een stip op een beeldscherm, kom je een dag later terug met een bionische-plasma-robot. Maar het werkt wel. Bedankt voor alle technische ondersteuning en uitleg (en thee).

De Stigcheltjes zijn de sterkst groeiende populatie binnen de afdeling Experimentele Psychologie. **Rudmer** (Ruddybuddy), de prevalentie van slechte woordgrappen op de afdeling is na jouw vertrek gehalveerd. Dankzij jouw turbowerkweek heeft mijn snelheidsrecord proefschriftschrijven welgeteld 1 week stand gehouden. Top. **Nathan** (Mr. From-the-Sidewalk), bedankt dat je altijd open staat voor vragen, inzichten, en commentaar. Bij jou kan ik altijd al mijn onderzoeksideeën testen, tenzij het gesprek per ongeluk ontaard in een religieus debat. Ook ben ik blij dat jij de rol van binnenploffen-en-klagen-over-pilotresultaten recentelijk hebt overgenomen van Jelmer. **Paul**, omdat je als gelukszoeker ook gewoon Nederlands moet leren spreek ik jou ook aan in het Nederlands. Ik heb er veel vertrouwen in dat jij een mooi project tot stand kan brengen met Stefan, en ik ben er enthousiast over om er een bijdrage aan te kunnen leveren. **Jasper** (Bonenkweker, Tuinboon) en **Martijn** (hoezo heb jij geen bijnaam?), bedankt voor het aanraden van Rick & Morty. Of eigenlijk wil ik jullie vooral bedanken dat jullie mij deze gouden tip pas na de deadline van mijn proefschrift hebben gegeven. Ik verheug me op de volgende lunches met grappen die door geen enkele spreekwoordelijke dan wel letterlijke beugel kunnen. **Erik**, **Mijke**, **Jessica**, ook jullie bedankt voor de Stigchelende momenten.

Natuurlijke lopen er veel meer figuren rond op de afdeling waar ik ontzettend veel (al dan niet academische, maar altijd) goede gesprekken mee kan voeren. **Estrella** (Esther-Elle-Ja, Tante Es), om de een of andere reden ben jij een uiterst geschikte kandidaat om roddels mee te delen (meestal net iets te hard in een net iets te stille treincoupé). Ook was ik er stiekem erg van gecharmeerd dat tante Wortels en jij, alsmede jullie respectievelijke aanhang, als toeschouwers aanwezig waren toen ik mijn academische pet tijdelijk had verruild voor een rapperspetje. PS: Amsterdam is de mooiste stad van de wereld. **Miranda** (tante Mier), ook voor jou geldt: treinreis = roddeltijd. Daarbij geldt: gangtijd = met een gedicideerde tred doormarcheren. PS: de Indische buurt is de mooiste buurt van de wereld. **Manasa** (the SC girl), fijn om iemand op de afdeling te hebben die daadwerkelijk mijn naam kan uitspreken. Ik wacht nog op de spelavond! **Ivo**, jij hebt de meest Duitse snor van de afdeling (en ook de enige). Volgend biertje bij NVP of bij TEAP? **David**, bedankt voor het helpen met de SPM probleempjes. Het heugt mij dat de toewijding die jij toont op je werk ook terug te vinden is in de toewijding die jij toont in de kroeg. **Peter**, ook jij bedankt. Ik weet niet precies waarvoor, maar je bent gewoon een mooi figuur, en ik vond je in dit rijtje thuishoren. En **Anouk**, hoe zit het met dat etentje in de meest psychologische tapasbar van Amersfoort? Voorts, zou het werken op de afdeling Experimentele Psychologie ondenkbaar zijn zonder de ondersteunende krachten van **Ria** (ik kom zo een dropje jatten), **Dennis** en **Eveline**. Vanzelfsprekend wil ik ook alle andere mensen op de afdeling bedanken die ik niet expliciet (maar bij dezen wel impliciet) heb benoemd.

Omdat er natuurlijk maar één echte stad in Nederland is kan ik niet ontkomen aan een dankwoord naar mijn collega's uit het immer pittoreske Amsterdam. **Ilja**, **Yaïr**, **Anouk**, **Tomas**, **Annelinde**, **Simon**, jullie laten elke keer weer zien dat feestjes op VSS, NVP en ASSC een extra vleugje aanheid krijgen met een UvA state-of-mind. Daarbij is het ook niet geheel onbelangrijk dat jullie tijdens het middagprogramma van voorgenoemde congressen ook stuk voor stuk indrukwekkend werk laten zien. Bedankt voor de inspiratie en de goede sfeer. Hetzelfde geldt voor mijn gewaardeerde (ex) VU-ers, zoals **Sebastiaan**, **Lotje**, **Dirk**, **Wouter**, **Jaap** en **Joanne** van de jongere garde, en natuurlijk ook **Chris**, **Artem** en **Jan** van de gevestigde orde. Ondanks dat jullie bij de NVP pubquiz het podium net niet hebben gehaald, weten jullie mij op (alcoholisch en) academisch vlak wél te inspireren. We hebben een aantal mooie samenwerkingen op poten gezet, en ik hoop dat we samenwerkingen ook in de toekomst niet uit de weg gaan. **Romke**, bedankt voor de ruimte die je me hebt gegeven tijdens mijn masterthesis, en bedankt voor de enthousiaste begeleiding en het mooie cijfer. Ik weet zeker dat ik zonder dat project nooit zo een goede promotieplek had kunnen krijgen. **Leendert**, bedankt voor je bijdrage aan dit proefschrift, en ook voor je prettige manier van werken en geduldige uitleg. Ik heb veel van onze samenwerking geleerd. Ook wil ik de rest van de VU/UvA-ers bedankt voor jullie bijdrage aan een altijd goede sfeer als ik op bezoek kom.

If I had to list the multiple highlights of my PhD project, the lab visit to the Charité Universitätsmedizin in Berlin would definitely be somewhere at the top. **Philipp**, I am incredibly thankful that you gave me the opportunity to join your lab for a couple of months. The collaboration that we worked on resulted in a great chapter for this dissertation, as well as a number of papers that I should be writing up now, instead of writing these acknowledgements. Your experimental insights and scientific vigilance are all the more impressive considering that you divide your attention between experimental and clinical work at the psychiatry clinic. The Perception lab is lucky to have such an inspiring principal investigator. **Matthias**, you basically taught me everything that I know about fMRI. I feel quite confident now in setting up, running, and analyzing a new (but arguably rather simple) fMRI experiment all by myself, and that is fully thanks to your fMRI crash course during my stay in Berlin. Thanks for all the time you took in teaching me, in eating cookies at the scanner, and in answering my incredibly lengthy emails. **Thomas**, thank you as well for all the technical and theoretical insights revolving around fMRI experimentation, as well as for answering equally lengthy emails. And let's not forget the many beers at VSS and TEAP. **John**, although we only met on few occasions, your input in our project was crucial for its success. Aside from my brilliant collaborators, I am also grateful for all the scientific and social moments that I shared with the rest of the Perception lab. **Marcus** and **Apoorva** thanks for helping me out with both fMRI and eye-tracking issues on multiple occasions. **Karin** and **Guido**, thanks for a lot of inspiring conversations, I always enjoy your scientific rigor and theoretical insights. **James**, thanks for the Rauchbier! **Gregor**, **Veith**, and the rest of the lab, you all contributed to my enthusiasm vis-à-vis the Perception lab. Last, but not least (although not part of the lab anymore), thank you **Timo** for all our extensive email discussions on b-CFS, monocular control conditions, and measuring the content of awareness. I am looking forward to collaborating in the near future (but then who will review it?).

The Summer School in Visual Neuroscience, in Rauischolzhausen, was definitely one of the highlights of my PhD curriculum as well. I want to express my gratitude towards all my peers, for the great time that I had in Karl's castle as well as at the VSS reunion. Above all, of course, I want to take this opportunity to thank **Karl**, **Jochen**, **Wolfgang**, and all the lecturers, for making this summer school possible. It was an invaluable experience that I recommend to any PhD candidate or postdoctoral researcher in vision science. The only downside: We were told that, by the end of the summer school, we would be able to pronounce the name "Rauischolzhausen".

Unfortunately, I am quite confident that (at least) all international attendees miserably failed at achieving this goal.

Als laatste wil ik nog iedereen bedanken die er de afgelopen jaren op persoonlijk vlak voor mij is geweest. De lijst van vrienden is te lang om op te noemen (zeker als ik de bijnamen er ook bij vermeld), maar jullie weten zelf wel wie ik bedoel. Bedankt voor alles! Een aantal mensen wil ik ook in het bijzonder bedanken. Niet omdat hen per se meer dank ten deel valt dan de rest, maar omdat ze wellicht op meer regelmatige basis last hebben gehad van mijn wetenschappelijk geneuzel. **Arjen** (de Buffel, Arie, Knut), je kent ongetwijfeld het gezegde "een schoon huis vergt twee opruimdriftige huisgenoten". Nee? Ik ook niet. Laten we dan snel maar chillen in één van onze huidige en uiterst volwassen schone onderkomens. **Kim**, ondanks dat je mij misschien liever als flierefluitende troubadour had gezien, ben ik je heel erg dankbaar in het grenzeloze vertrouwen dat je altijd in mij toont. **Vera** (Verski, Noppie) en **Lotte** (Mullie), bedankt voor de oneindigheid aan fantastische roddelsessies. Waar moeten we de komende roddelverhalen vandaan halen nu we allemaal bezet zijn? **Syl** (Grilly, Nadia, . . .), jij bent de personificatie van opgewektheid, enthousiasme en positivisme. Ik verheug me nu al op elk volgend avondje drinken, dat ofwel eindigt met een pseudo-vriendschappelijke knokpartij, ofwel met een substantieel financieel verlies voor een van ons, middels een onzinnige weddenschap. **Anouk** (Noukie), jij bent waarschijnlijk degene die het meest heeft meegekregen van het wel en wee van mijn promotietraject. Bedankt voor het fijne thuiskomen na een dag hard werken, en bedankt dat ik mijn hart kon luchten als er weer van alles misgegaan was tijdens mijn experimenten. **Manu**, bedankt voor het keuren van mijn posterpresentaties, voor het legoën, en voor het geleidelijk aan steeds vaker op normale tijdstippen wakker worden. **Patrice** (Patje Vano), jij bent waarschijnlijk de meest eerlijke persoon die ik ken. Bedankt voor je altijd oprechte en ongezouten mening. **Barend** (Dopaganda, Blubberman), **Edgar** (Greg Kerkwerper, etc.), **Gwendal** (Kalie, Darmpje), **Rachid** (Crackbek, Frits van Dijk), **Ronald** (Huussuh), en al mijn andere muzikale metgezellen plus aanhang, bedankt voor de vele muzikale, alcoholische en intellectuele intermezzo's. **Floor** (Frenkel), jij bent een soort van sociale gezelligheidsbom waar elk mens gelukkig van wordt. Wanneer gaan we weer whisky drinken? Ik had **Inge** al bedankt in haar rol als paranimf, maar laten we vooral de rest van ons neurotische nerdclubje niet vergeten. **Amber**, **Oscar**, **Mats** (ik laat de bijnamen hier even achterwege in verband met de drukkosten voor extra pagina's), bedankt voor alle geweldige gesprekken, onder het genot van overheerlijke maaltijden. Ook bedankt voor het aanhoren van mijn gewauwel over Bayesiaanse statistiek om drie uur 's nachts. **Alden** (Dino the Don, Aarslikker D), jij bent een onmisbaar figuur in deze wereld. Daar moeten er meer van zijn (maar ook niet te veel). **Vera** (Kribbie, Subsidieblondje), bedankt voor de leuke tijd samen. Jij hebt de pech gehad dat je het laatste stuk eindstress van mijn promotietraject van dichtbij hebt moeten ervaren. Gelukkig gaf je aan er zo goed als niets van te merken (mijn gemiddelde stressniveau tijdens het koken ligt hoger). Al met al val je dus met je spreekwoordelijke neus in de spreekwoordelijke boter, en heb je juist de perfecte timing gekozen om samen het eind van mijn promotietraject te vieren! Ik geniet van de tijd die ik met jou doorbreng, en hoop in de toekomst nog veel samen mee te maken.

Bien-sûr je ne peux pas laisser passer l'occasion de remercier ma Grand-Mère. **Grand-Mère** (mémé pour l'un, Marie-Thé pour l'autre), merci d'être toujours si contente de m'entendre au téléphone quand je t'appelle (bien que pas assez souvent). Et merci d'être toujours si fier de me prouesses académique, depuis le début de mes études, et même avant. Je suis sûr que, maintenant que ton petit-fils a reçu son doctorat, tu vas raconter ça à tous tes collègues aux Hespérides. J'espère pouvoir te rendre visite bientôt!

Uiteraard wil ik boven alles ook **mijn ouders** bedanken, die altijd ongelofelijk enthousiast zijn over alles wat ik doe, en mij steunen in alle keuzes die ik maak. Ze waren enthousiast toen ik regelmatig op het podium stond, ze waren enthousiast toen ik psychologie studeerde, ze waren enthousiast toen ik bouwkunde studeerde, en ze waren zelfs enthousiast toen ik gestopt was met bouwkunde en in een winkeltje rondvaart-tickets verkocht (of ze deden alsof ze enthousiast waren, maar dat is net zo bewonderenswaardig). Ook wil ik jullie bedanken voor de sporadische financiële injecties, waarmee ik in geval van nood toch de gaten kon dichten die ik in vlagen van opportuun positivisme in mijn portemonnee had geslagen. Zonder jullie was dit allemaal een onmogelijke opgave geweest.

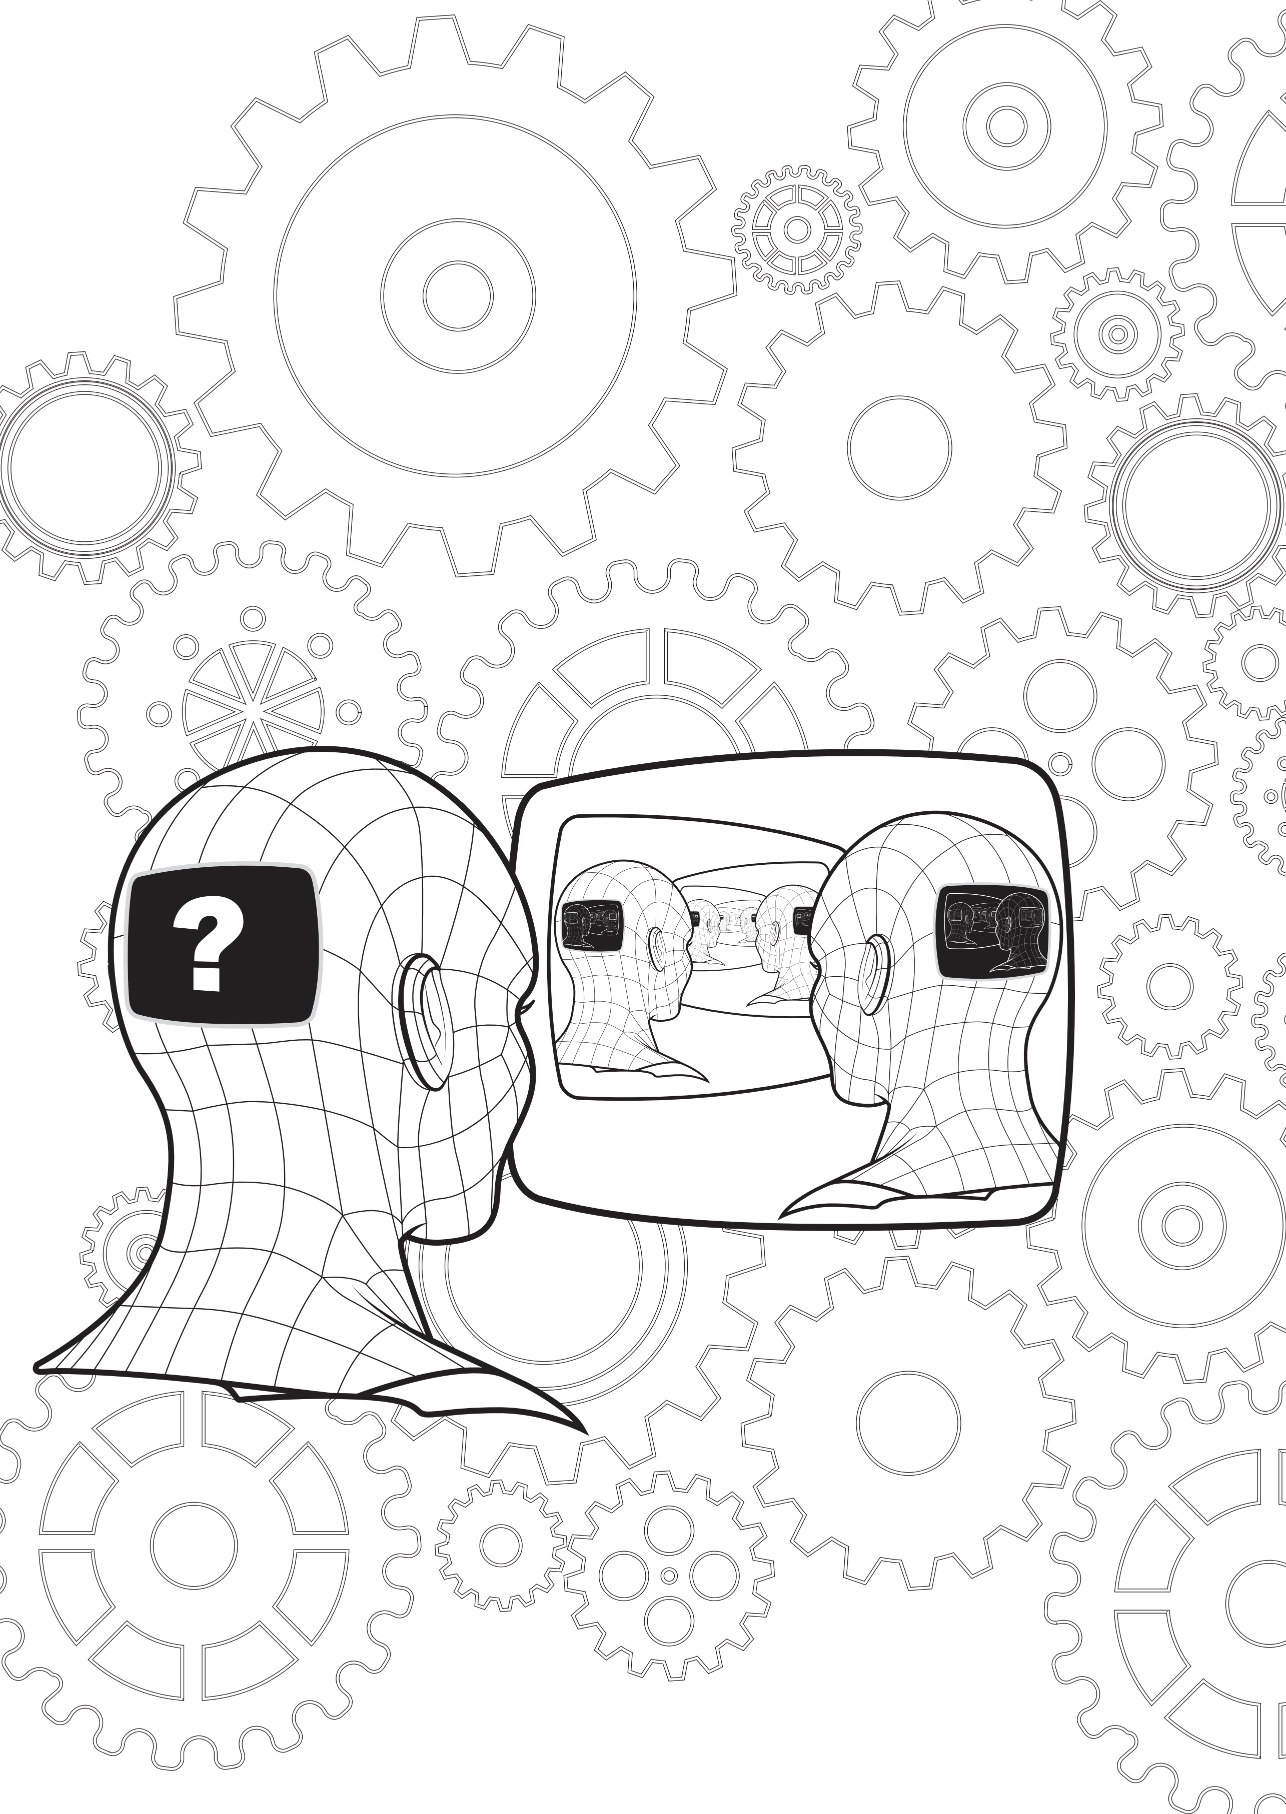

Curriculum vitae

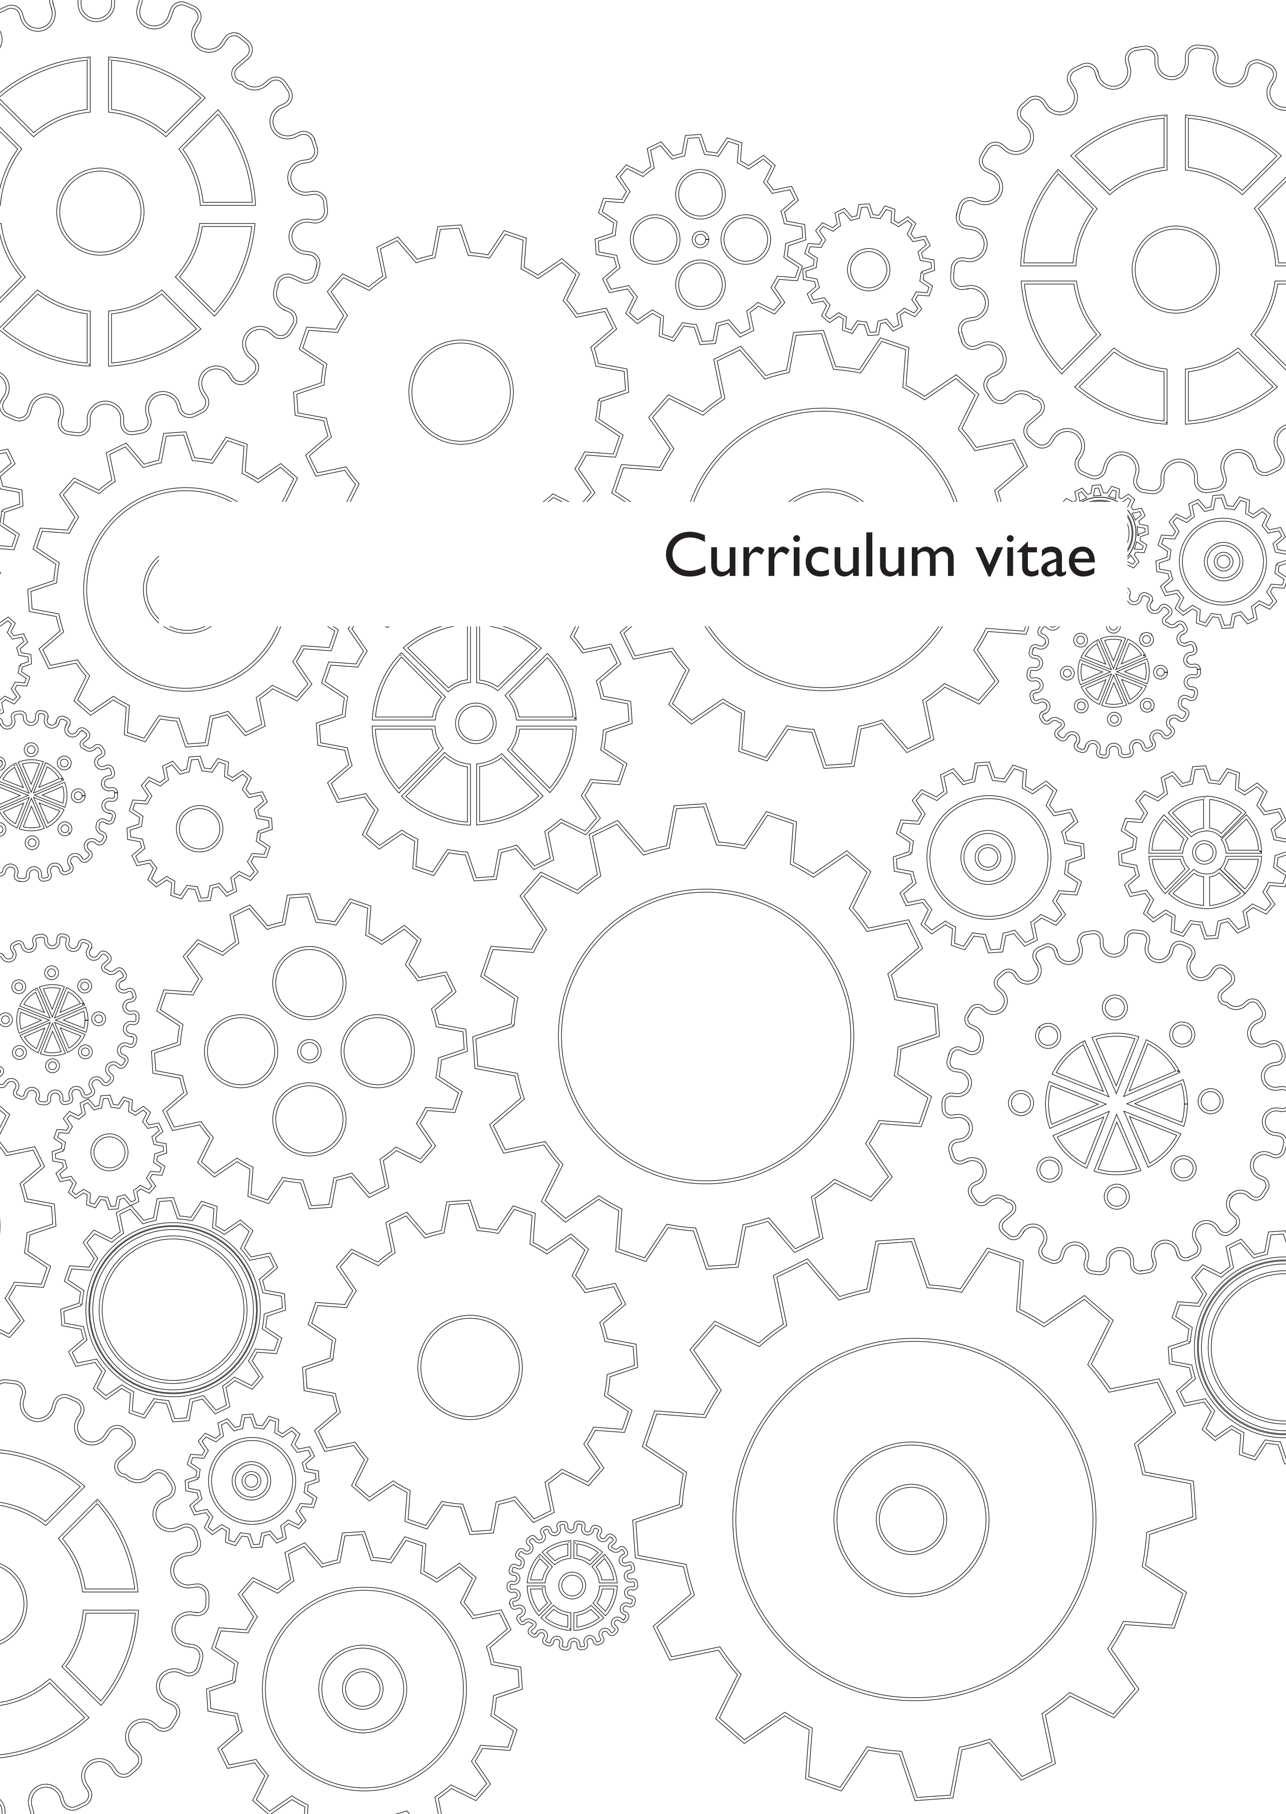

Biography

Surya Gayet was born on May 16th 1984 in Amsterdam, The Netherlands. In 2002 Surya graduated from secondary school 'het Vossius Gymnasium'. After a detour studying Architectural Engineering (Technical University of Delft), he obtained his Bachelor of Science degree in Psychology at the University of Amsterdam, where he subsequently obtained his Master of Science degree (*with honors*) in Cognitive Psychology in 2011. While he was in University, Surya was also active as a musician. Within the hip-hop subculture his reputation mostly revolved around his achievements in *rap battles* (verbal improvisation competitions), winning most major Dutch tournaments between 2005 and 2010. However, he also released two studio albums and two records that were well received by critics. During that period, Surya performed in hundreds of venues and festival stages throughout The Netherlands and Belgium. Between 2012 and 2016 Surya did his PhD at Utrecht University, under the direct supervision of dr. Chris Paffen and dr. Stefan van der Stigchel. During his PhD curriculum, Surya spent three months in Prof. dr. Philipp Sterzer's Perception lab, at the Charité University Hospital in Berlin, Germany. Currently, Surya is a postdoctoral researcher in dr. Stefan van der Stigchel's Attention Lab at Utrecht University, where he conducts behavioral and neuroimaging research in the fields of human consciousness, visual perception and visual working memory.

Biografie

Surya Gayet werd geboren op 16 mei 1984 te Amsterdam, Nederland. In 2002 behaalde Surya zijn Gymnasium diploma op het Vossius Gymnasium. Na een omweg via de studie Bouwkunde (Technische Universiteit Delft) behaalde Surya zijn Bachelor of Science in Psychologie aan de Universiteit van Amsterdam, alwaar hij in 2011 *cum laude* zijn Master of Science in Cognitieve Psychologie behaalde. Tijdens zijn studie was Surya ook actief als muzikant. Surya was binnen de hip-hop subcultuur voornamelijk bekend om zijn *rap battles* (verbale improvisatie competities), waarin hij tussen 2005 en 2010 vrijwel alle belangrijke toernooien van Nederland op zijn naam wist te schrijven. Daarnaast heeft Surya echter ook twee studioalbums en twee singles uitgebracht, die lovend werden ontvangen door recensenten. In die periode heeft Surya op honderden podia in Nederland en België zijn muziek mogen presenteren. Van 2012 tot 2016 promoveerde Surya aan de Universiteit Utrecht, onder begeleiding van zijn copromotoren dr. Chris Paffen en dr. Stefan van der Stigchel. Tijdens zijn promotietraject werkte Surya gedurende drie maanden in het Perceptie lab van Prof. dr. Philipp Sterzer, in het Charité Academisch Ziekenhuis te Berlijn, Duitsland. Momenteel is Surya werkzaam als postdoctoraal onderzoeker in het Attention Lab van dr. Stefan van der Stigchel, waar hij neuroimaging en gedragsexperimenten uitvoert naar de werking van het menselijk bewustzijn, visuele waarneming en visueel werkgeheugen.

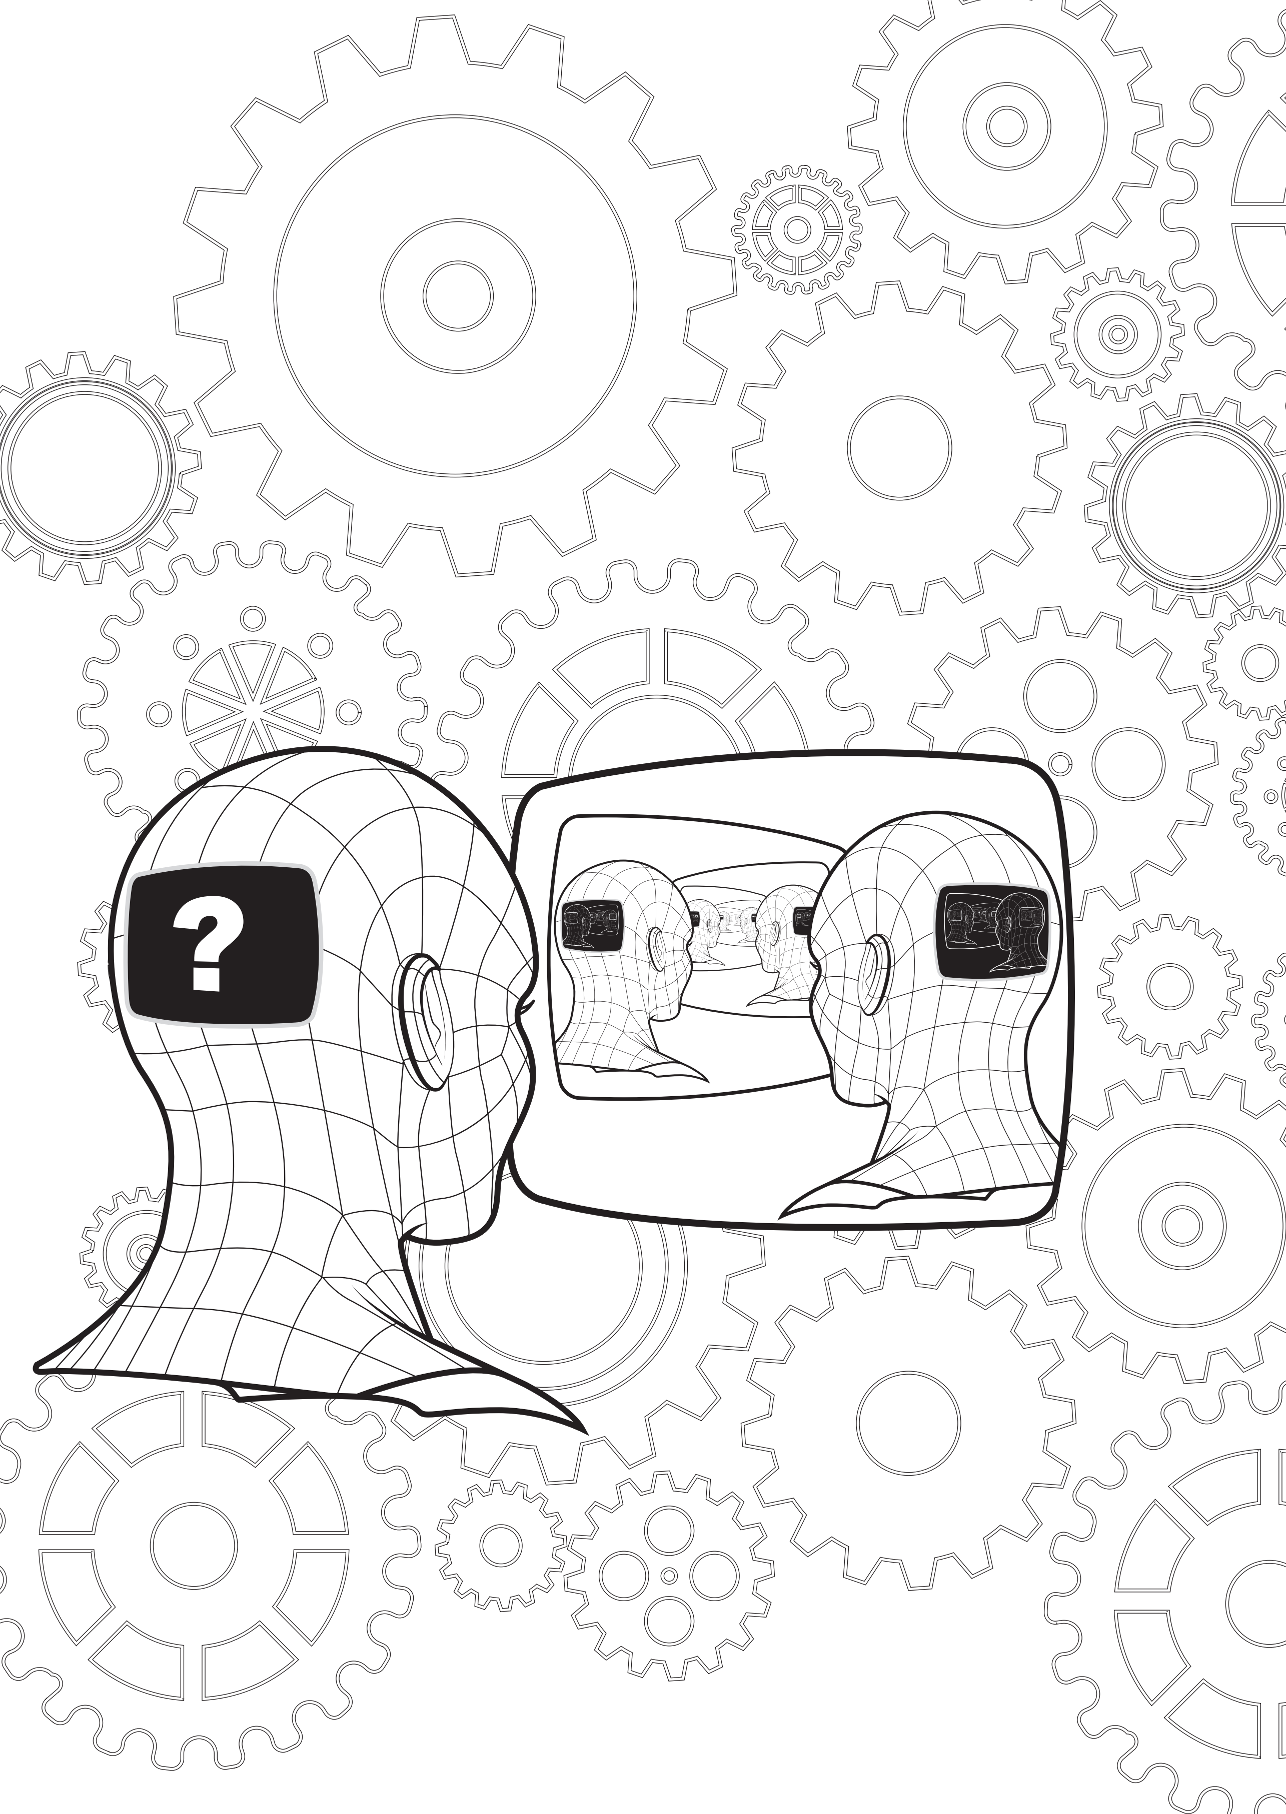

Bibliography

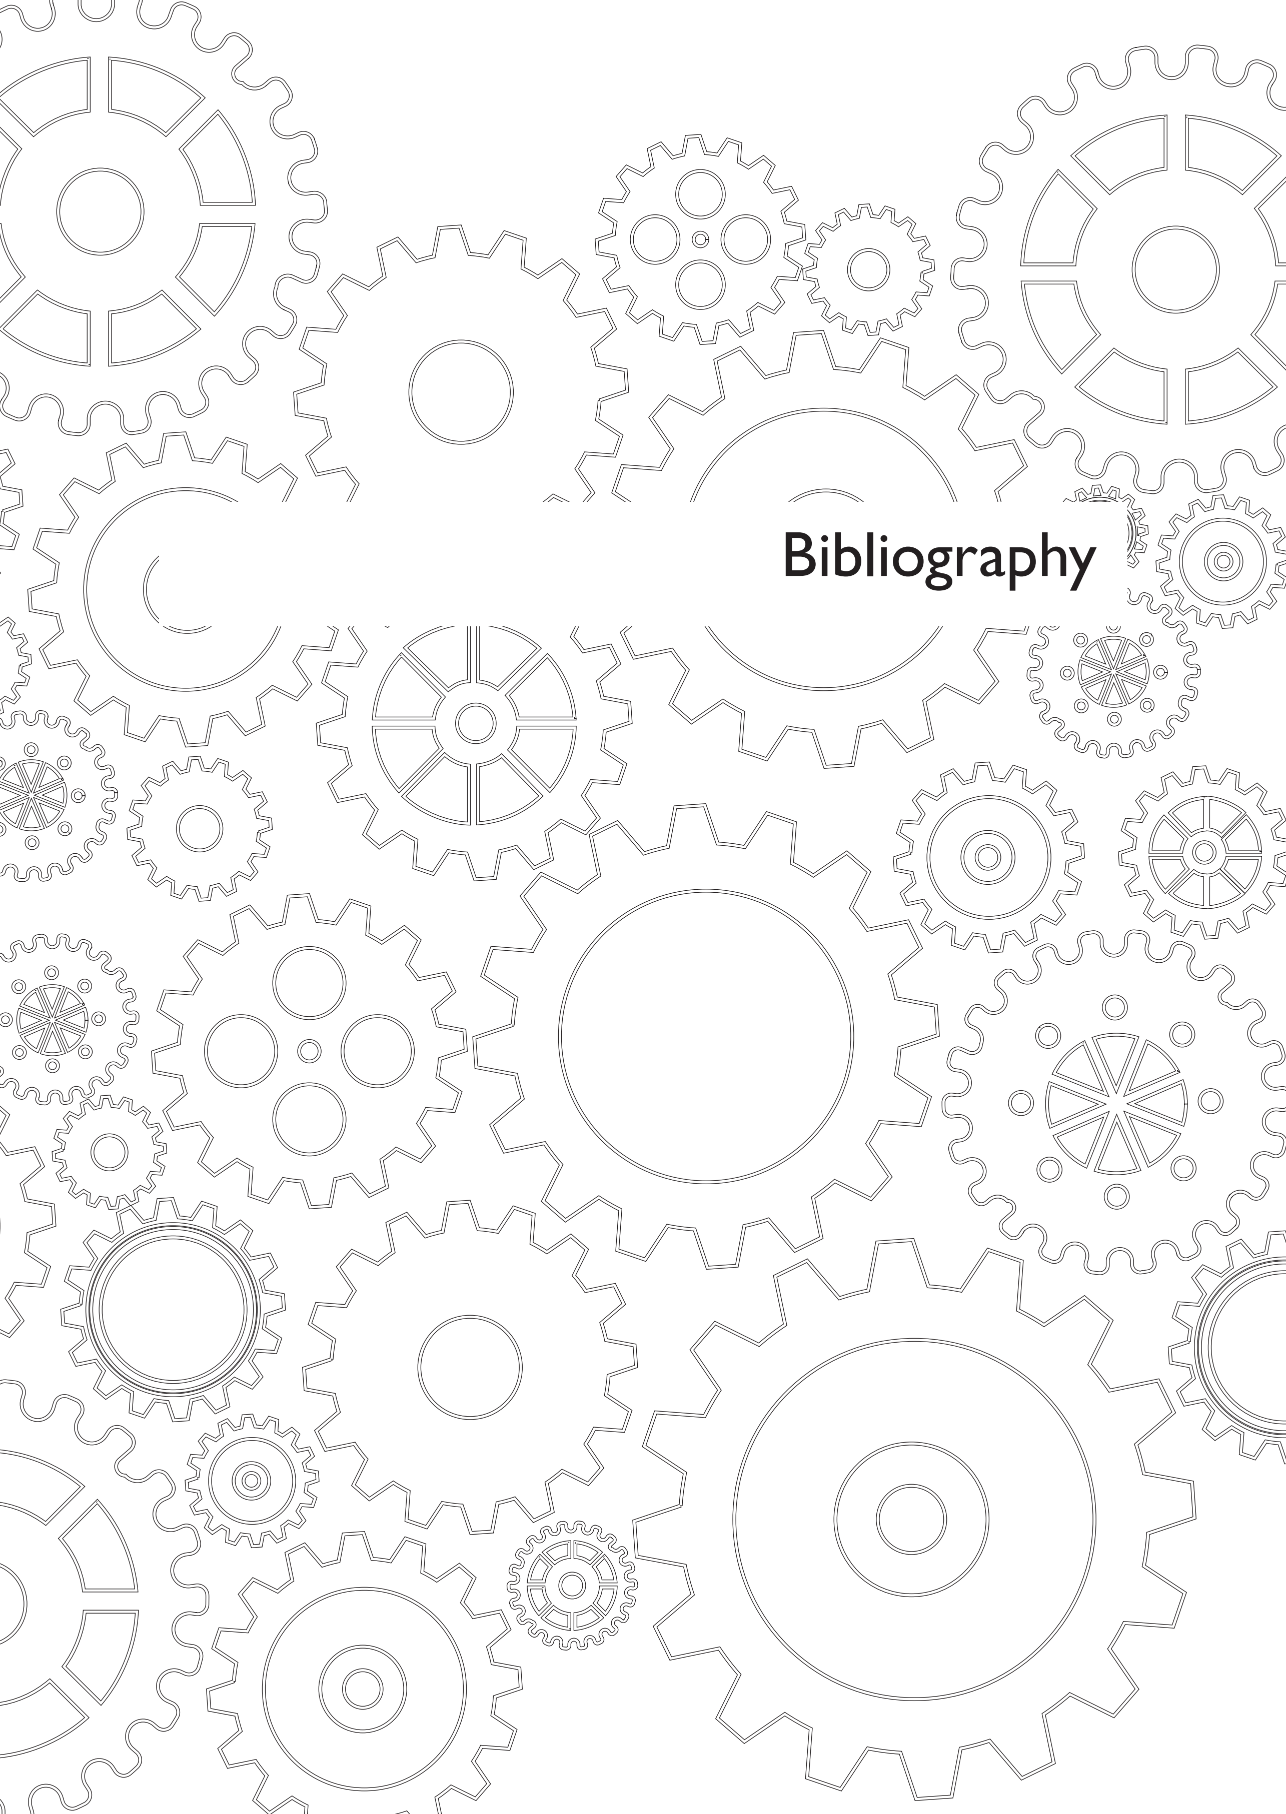

Published peer-reviewed articles

Gayet, S., Paffen, C. L. E., Belopolsky, A. V., Theeuwes, J., & Van der Stigchel, S. (2016). Fear conditioned visual input gains preferential access to awareness in a breaking continuous flash suppression paradigm. *Cognition, 149*, 77-83. doi:10.1016/j.cognition.2016.01.009

Gayet. S., Brascamp, J. W., Van der Stigchel, S., & Paffen, C. L. E. (2015). Cogito ergo video: Task relevant information is involuntarily boosted into awareness. *Journal of Vision, 15*(5)3, 1-18. doi:10.1167/15.5.3

Gayet. S., Van der Stigchel, S., & Paffen, C. L. E. (2014). Breaking continuous flash suppression: Competing for consciousness on the pre-semantic battlefield. *Frontiers in Psychology: Consciousness Research, 5*: 460. doi:10.3389/fpsyg.2014.00460

Gayet. S., Van der Stigchel, S., & Paffen, C. L. E. (2014). Seeing is believing: Utilization of subliminal symbols requires a visible relevant context. *Attention, Perception & Psychophysics, 76, 489-507*. doi:10.3758/s13414-013-0580-4

Gayet. S., Paffen, C. L. E., & Van der Stigchel, S. (2013). Information matching the content of visual working memory is prioritized for conscious access. *Psychological Science, 24(12), 2472-2480*. doi:10.1177/0956797613495882

Articles in preparation (data used in this dissertation)

Gayet. S., van Maanen, L., Heilbron, M., Paffen, C. L. E., Van der Stigchel, S. (under review). Visual working memory reduces the threshold for matching visual input to reach conscious access: evidence from sequential sampling models.

Gayet, S., Guggenmos, M., Paffen, C. L. E., Van der Stigchel, S., Christophel, T. B., Haynes, J. D., & Sterzer, P. (in preparation). Increased BOLD response and higher decoding accuracy for visual input that matches the content of visual working memory.

Gayet, S., Guggenmos, M., Paffen, C. L. E., Van der Stigchel, S., Christophel, T. B., Haynes, J. D., & Sterzer, P. (in preparation). The content of visual working memory does not modulate the BOLD response to interocularly suppressed visual input.

Gayet, S., Paffen, C. L. E., Guggenmos, M., Sterzer, P., & Van der Stigchel, S. (in preparation). Pupillary constriction accompanies the synergetic effect of visual working memory on pre-conscious visual input.

Hofstadter’s Law: It always takes longer than you expect,
even when you take into account Hofstadter’s Law.
Douglas Hofstadter