

PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a publisher's version.

For additional information about this publication click this link.

<http://hdl.handle.net/2066/200823>

Please be advised that this information was generated on 2019-11-13 and may be subject to change.

CONSCIOUS PERCEPTION IN THE PREDICTIVE BRAIN



DONDERS
SERIES

ERIK L MEIJS

CONSCIOUS PERCEPTION IN THE PREDICTIVE BRAIN

Erik L. Meijs

CONSCIOUS PERCEPTION IN THE PREDICTIVE BRAIN

PhD thesis, Radboud University Nijmegen

© Erik L. Meijs, 2019

ISBN: 978-94-6284-184-0

Cover: Erik Meijs. Template brain image from SVGrepo.com (CC BY 4.0)

Print: Ipskamp printing BV

The research described in this thesis was carried out at the Donders Centre for Cognitive Neuroimaging. Data in chapter 2 was acquired at the University of Amsterdam. The PhD-project was funded by a TopTalent grant from the Donders Graduate School for Cognitive Neuroscience (Donders Centre for Medical Neuroscience / RadboudUMC) awarded to Erik Meijs. In addition, research was funded based on grants awarded to various co-authors by the Netherlands Organization for Scientific Research, the European Research Council and the James S. McDonnell foundation.

CONSCIOUS PERCEPTION IN THE PREDICTIVE BRAIN

Proefschrift

ter verkrijging van de graad van doctor

aan de Radboud Universiteit Nijmegen

op gezag van de rector magnificus prof. dr. J.H.J.M. van Krieken,

volgens besluit van het college van decanen

in het openbaar te verdedigen op maandag 25 maart 2019

om 14:30 uur precies

door

Erik Lambertus Meijs

geboren op 5 juni 1990

te Tilburg, Nederland

Promotor

Prof. dr. Floris P. de Lange

Copromotor

Dr. Simon van Gaal (Universiteit van Amsterdam)

Manuscriptcommissie

Prof. dr. W. Pieter Medendorp

Prof. dr. Victor A. F. Lamme (Universiteit van Amsterdam)

Dr. Manon G. J. Mulckhuyse (Vrije Universiteit Amsterdam)

TABLE OF CONTENTS

Chapter 1	Introduction	7
Chapter 2	Dynamic interactions between top-down expectations and conscious awareness	17
Chapter 3	Exploring the effect of expectations and stimulus relevance on stimulus-specific neural representations and conscious report	41
Chapter 4	Cue predictability does not modulate bottom-up attentional capture	67
Chapter 5	General discussion	89
Appendix	References	102
	Nederlandse samenvatting	110
	Publication list	113
	CV	114
	Acknowledgements	115
	Donders Graduate School for Cognitive Neuroscience	119

1

INTRODUCTION

In our everyday life, we normally have the experience that we have a rich perception of the world around us, driven by an accurate representation of it. However, it has been known for a long time that perception is not a passive process in which signals coming in through the senses are merely registered. Rather, perception is thought to be an active process in which the brain tries to make sense of the inputs from the environment it is in [1]. As a result, perception seems to be as much an interpretation as a representation of the outside world.

The idea that there is no one-to-one mapping between perception and the outside world can be illustrated with simple visual illusions. In these illusions, the perceptual system is misled, because assumptions the brain normally makes when interpreting (visual) input lead to an erroneous conclusion about specific stimuli. In the examples in **FIGURE 1.1**, the context in which a stimulus (a circle in A and a line in B) is presented has a big influence on how the stimulus is interpreted. But there are even more basic examples that demonstrate that perception deviates from sensory inputs. For example, think about the fact that we automatically fill in the input missing due to light falling onto the blind spot of the retina, leading us to have a “complete” visual experience even though part of the input is missing.

A. Ebbinghaus illusion



B. Müller-Lyer illusion




Figure 1.1: Visual illusions. (A) Ebbinghaus illusion (also called Titchener circles). Both centered circles are of equal size, but the surrounding circles make the center circle on the left look smaller than the center circle on the right. (B) Müller-Lyer illusion. Both lines are of equal length, but the positioning of the arrows heads makes the top line appear shorter and the bottom line appear longer.

PREDICTIVE BRAIN THEORIES

The hypothesis that the brain tries to understand its environment is also a prominent feature of recently developed neurocomputational theories like predictive coding [2,3] and adaptive resonance theory [4]. One central tenet in these theories is that the brain uses its knowledge about the environment to build expectations about upcoming inputs. Therefore, I will collectively refer to them as “predictive brain theories”. Expectations are typically thought to originate from higher-level brain regions, either

within the sensory cortex or in associative areas such as the (pre)frontal cortex, and they may guide information processing in lower-level sensory regions via top-down projections. These expectations can both help and harm perception [5]. Consider, for example, a situation where you are taking a walk in the park and hear a barking sound. The auditory system processes the sound and eventually you recognize the sound as that of a dog. This then generates a visual expectation about what you will see when we turn our head towards the sound source: a dog. When indeed there is a dog present, this may lead to an improved neural representation in visual regions while overall brain activity is reduced [6]. On the other hand, when you do not see a dog but a loudspeaker, a prediction error signal will trigger the brain to update its expectations about the source of the sound [7].

While a surge of scientific work now supports many of the central assumptions of predictive brain theories [8,9], there are still many open questions regarding specific aspects of these theories. In this thesis, I will focus on a number of these questions.

EXPECTATIONS AND CONSCIOUS AWARENESS

One unresolved issue relates to the relation between expectations and conscious awareness. Some prediction-based theories posit that expectations are intrinsic and automatic brain computations [9]. This suggests that they should take place irrespective of a person's behavioral state. Indeed, subjective perception feels like an effortless process that does not require much willful intervention. That is, perceptual expectations seem to arise automatically, in the absence of conscious awareness or focused attention. Still, to date there has been no conclusive empirical support for this hypothesis. A number of studies have shown that low-level invalidations of regularities can be processed outside of conscious awareness [10,11], but it is unclear whether these findings reflect true model-based comparisons of expectations to new input or more basic low-level adaption effects [12,13].

In addition, the notion of nonconscious expectations appears incompatible with many of the leading theories of conscious awareness (see **Box 1.1** for a general overview). Many predictive brain theories posit that expectations are implemented by feedback loops from other areas [3,4]. These feedback loops can be implemented both on a low, local level [14] for low-level expectations, by longer-range expectations originating in higher-order areas (e.g. frontal cortex) to implement more cognitive, conscious

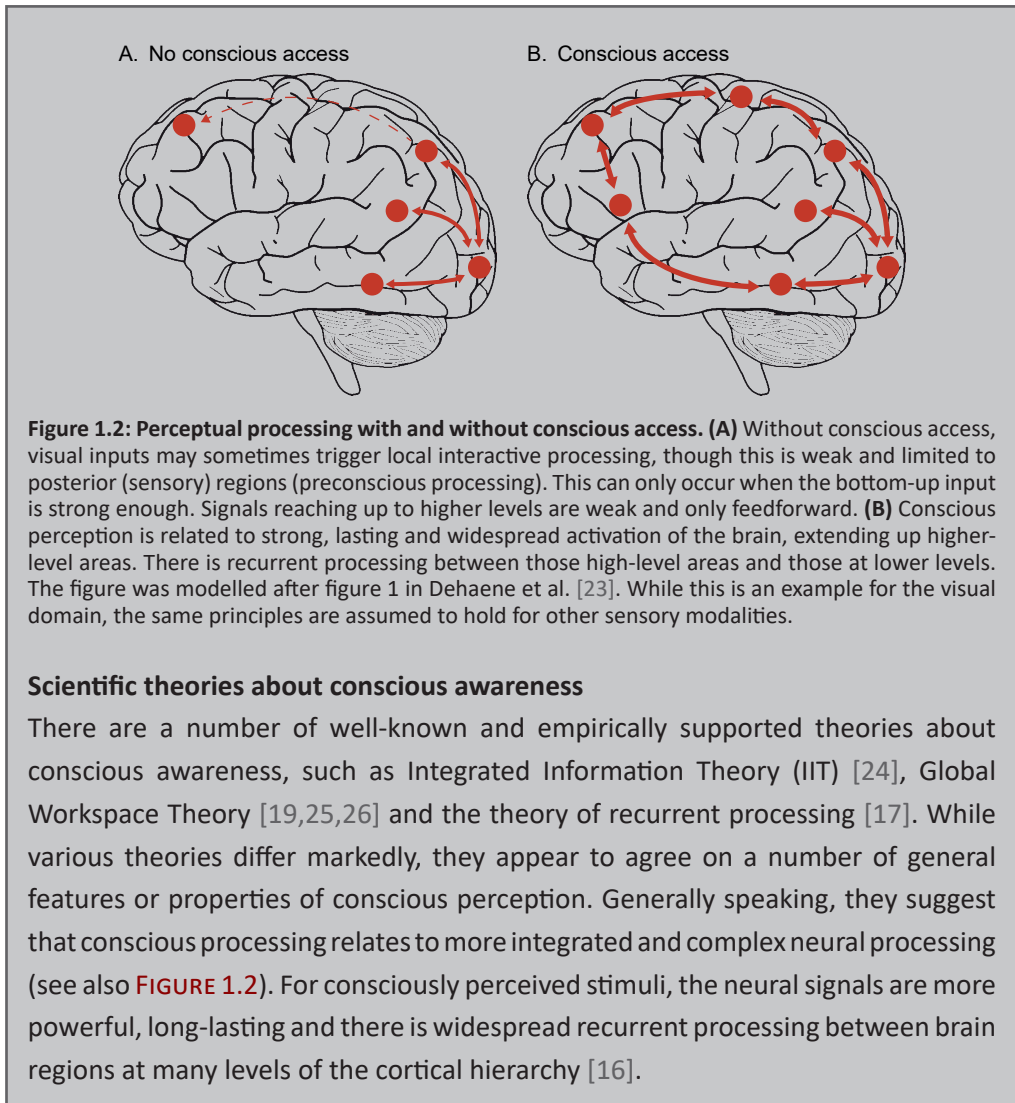
BOX 1.1: THE SCIENTIFIC STUDY OF CONSCIOUSNESS

Decades of scientific work have been aimed at explaining the properties of conscious perception (subjective experience) and its relation to the brain. This journey has often been characterized as the search for the neural correlates of consciousness (NCCs): the minimal set of neural mechanisms that are jointly sufficient for a conscious percept [18]. Below, I will give a short overview of how consciousness is usually studied and what the current leading theories about conscious awareness entail.

Studying consciousness

The techniques to manipulate the level of conscious access to sensory inputs can be roughly divided into two types of experiments. The first type of experiments degrades the stimulus input itself, for example by presenting it very shortly or embedding it within noise. Best known are masking paradigms, in which a weak stimulus is preceded and/or followed by much stronger masking stimuli. Often, this condition is compared to a condition where the stimulus input is stronger and masks are absent or weaker [19]. This results in differences in bottom-up input that can be confounded with the effect of conscious awareness. To prevent this potential confound, another type of experiments exploits more natural variations in conscious awareness, for example due to momentary lapses of attention. This is for example the case in the attentional blink paradigm [20] (used in **CHAPTER 2** and **CHAPTER 3**), in which conscious perception of the second of two target stimuli is impaired due to temporary unavailability of processing resources. A downside of methods like this is a lack of experimental control over conscious awareness.

The (un)consciousness of a percept can be determined based on either objective or subjective measures [21,22]. In the former, participants perform a (detection) task on the supposedly (non)conscious stimulus and their performance is compared to chance level. In a nonconscious condition, participants should then not detect stimuli better than chance. A more subjective way to determine conscious awareness of stimuli is to rely on participants' subjective estimates. Of course, the usefulness of this method relies on the reliability of subjective reports, which may be compromised by for example decision or response biases [22].



expectations [15]. Interestingly, leading theories of consciousness implicate similar top-down feedback mechanisms in the constitution of conscious access [16,17]. Information that cannot be consciously accessed is thought to trigger only feedforward activity or localized recurrent interactions within sensory brain areas. Therefore, based on consciousness theories, nonconscious effects of expectations are highly unlikely because the feedback mechanisms they rely on are not present. In contrast, predictive brain theories would suggest that the processing of expectations is intrinsic and essential to brain functioning, and should therefore be automatic and not dependent

on conscious awareness. These conflicting hypotheses will be tested against each other in Chapter 2. In this chapter, I report a series of experiments investigating the dynamic interaction between conscious awareness and perceptual expectations. In these experiments, participants learn conditional probabilities between two sets of target stimuli. Consequently, once they have seen the initial target, they can form an expectation about the most likely target stimulus to follow. I will study whether the use of this type of predictive information depends on conscious awareness of the initial predictive stimulus. Additionally, I will investigate to what extent a violation of expectations is processed when the second stimulus is not consciously perceived.

EXPECTATIONS AND CONSCIOUS ACCESS

Earlier in this introduction I have mentioned that perceptual experience can deviate substantially from the actual sensory inputs reaching the senses. Likewise, it has been shown that while perceptual experiences feel rich and complete, conscious perception is often quite limited and “obvious” inputs are not always consciously processed. That is, stimuli of which the bottom-up input in principle is strong enough to be perceived may nevertheless not be consciously accessible, for example due to momentary lapses of attention. Attention is often assumed to be required for conscious perception [19,23]. A well-known example of salient but unexpected stimuli being missed is the case of inattention blindness. In the original paradigm participants watch a video and count how often a group of people pass a ball around. Because they are focusing on this task, they fail to see (or at least report) a person in the video who walks around in a gorilla suit [27].

A. Mooney image



B. Original image



Figure 1.3: Expectations and stimulus detection. In (A) an example Mooney image is presented. While the stimulus in this figure may not be very easy, having prior knowledge of the stimulus to look for will improve the speed and accuracy of stimulus detection. (B) Original image of a young leopard from which the Mooney image was constructed.

Recently, it has been shown that stimulus detection can be improved when prior expectations match the stimulus input [e.g. 28,29]. This is in line with predictive brain

theories, which commonly assume that what we consciously see is strongly linked to the brain's expectations about the outside world [30]. This idea can be illustrated with the Mooney image [31] shown in **FIGURE 1.3**. In such an ambiguous image, it is often difficult to distinguish the stimulus figure from the background. However, having advance knowledge on the stimulus to expect makes it easier to consciously see it. It is still an open question to what extent expectations can also boost otherwise unseen stimuli into consciousness, thus enabling the switch from a nonconscious to a conscious stimulus representation. **CHAPTER 2** and **CHAPTER 3** will explore these questions using the attentional blink paradigm. In this task, participants often fail to detect target stimuli. I will investigate whether this limitation can be overcome when participants have a prior expectation about the upcoming target stimulus. Furthermore, I will investigate the neural mechanisms underlying the observed effects.

EXPECTATIONS AND BOTTOM-UP ATTENTION

Expectations may also affect the processing of stimuli in other ways, such as by modulating bottom-up attentional mechanisms. Bottom-up attention is an automatic, reflexive type of attention which is triggered by suddenly presented salient stimuli [32]. In a standard visual bottom-up attention task, a cue is flashed somewhere on the screen and then followed after some time by a target stimulus. If the target is presented in the same location as the cue, performance is improved because the cue already attracted attention [33,34].

Bottom-up allocation of attention can be conceptualized as a way to redistribute neural resources: attention is attracted to a different part of the sensory input (and hence disengaged from previously attended items). Empirical evidence suggests that unexpected events require more processing resources [35]. Thus, it is possible that unexpected inputs trigger a bottom-up attention mechanism in order to reallocate processing capacity to a part of the sensory input that is not yet “explained”. This would be a way for the perceptual system to remain sensitive to sensory inputs next to its ongoing perceptual processes. According to predictive coding theory, the brain may assume that strong signals have high precision and consequently a location where this signal occurs may be temporarily given higher gain [36]. Therefore, prediction errors for new inputs in the same location may be upregulated, eventually leading to stronger suppression by the perceptual hypothesis predicting the input. This in turn is then more likely to be the hypothesis that is also consciously experienced [30].

It is also possible that rather than strong inputs, predictable inputs are preferentially weighted because they are considered to be more reliable [37]. That is, attention may reflect gain modulations of expectations (in contrast to prediction errors) and thereby make expected inputs more salient. This would help to prevent distraction by unexpected or unexplained inputs. As may be clear, so far there is no consensus what the role of bottom-up attention in the predictive brain is. Subtle differences in assumptions about what the brain considers “reliable” or “useful” information can lead to directly opposing hypotheses. In **CHAPTER 4** of this thesis I will delve into this unresolved issue by investigating the role of spatial and temporal expectations in visual bottom-up attention affects.

MAIN RESEARCH AIMS AND THESIS OUTLINE

In this thesis, I report a series of seven experiments distributed over three experimental chapters. In the experiments, I used psychophysics, electroencephalography and magnetoencephalography (see **BOX 1.2**) to investigate a number of hypotheses relating to the predictive brain hypothesis and its relation to conscious awareness and bottom-up attention.

Specifically, in this thesis I aimed at answering the following questions:

1. Can expectations boost an otherwise unseen stimulus into conscious awareness (**CHAPTER 2**)?
2. Are expectations automatic brain computations, taking place irrespective of conscious access (**CHAPTER 2**)?
3. How do expectations affect the neural processing of visual stimuli and how does this relate to conscious access (**CHAPTER 2** and **CHAPTER 3**)?
4. To what extent is bottom-up attentional capture determined by perceptual expectations (**CHAPTER 4**)?

In **CHAPTER 5** of this thesis, I will give an overview of the main conclusions that can be drawn from the experiments I have conducted and suggest directions for future research.

Box 1.2: EEG AND MEG

Electroencephalography (EEG) and magnetoencephalography (MEG) are neuroimaging methods that can be used to measure indexes of neuronal activity. Extracellular voltage differences, generated mainly by synaptic activity, can be measured at the scalp level by the use of a set of electrodes (EEG) [38]. Every electrical signal generates a magnetic field. The tiny variations in magnetic field generated through neural activity can be measured from outside the head with MEG. Both EEG and MEG have a high temporal resolution (< 10 ms) but relatively low spatial resolution, though the large number of sensors in MEG allow for better spatial inference than is possible in EEG.

How do we interpret the EEG and MEG signals?

In this thesis, I use two different analyses types to interpret the signals measured with EEG or MEG. In the first analysis method, used in **CHAPTER 2** and **CHAPTER 3**, the signal per channel is averaged over trials relative to a relevant time-point (for example the onset of a visual stimulus). Because noise is supposedly not time-locked to this event it is cancelled out and only the signal related to the relevant event. The average signal that is left is called the event-related potential (ERP for EEG) or event-related field (ERF for MEG).

A second analysis method, which I used in **CHAPTER 3**, is decoding. This is a multivariate analysis that does not just look at the average signal per channel over trials, but rather takes into account the trial-by-trial signal over multiple channels. Simply put, in the decoding analysis a subset of all trials is given to a computer algorithm that learns to distinguish multiple categories (e.g. conditions in an experiment) based on the trial data. A new set of trials is then used to see if the algorithm can successfully label them. If this is the case, it can be assumed that the neural data contained information about the categories of interest.

2

DYNAMIC INTERACTIONS BETWEEN TOP-DOWN EXPECTATIONS AND CONSCIOUS AWARENESS

Adapted from:

Meijs EL, Slagter HA, de Lange FP, & van Gaal S. (2018). Dynamic Interactions between Top-Down Expectations and Conscious Awareness. *Journal of Neuroscience*, 38, 2318–2327.

ABSTRACT

It is well known that top-down expectations affect perceptual processes. Yet, remarkably little is known about the relationship between expectations and conscious awareness. We address three crucial questions that are outstanding: 1) How do expectations affect the likelihood of conscious stimulus perception?; 2) Does the brain register violations of expectations nonconsciously?; and 3) Do expectations need to be conscious to influence perceptual decisions? We performed three experiments, in human participants, in which we manipulated stimulus predictability within the attentional blink paradigm, while combining visual psychophysics with electrophysiological recordings. We found that valid stimulus expectations increase the likelihood of conscious access of stimuli. Furthermore, our findings suggest a clear dissociation in the interaction between expectations and consciousness: conscious awareness seems crucial for the implementation of top-down expectations, but not for the generation of bottom-up stimulus-evoked prediction errors. These results constrain and update influential theories about the role of consciousness in the predictive brain.

INTRODUCTION

A rapidly growing body of work indicates that sensory processing is strongly influenced by expectations that we have about likely states of the world. Such expectations are shaped by the context in which we are operating, but also by learning, past experience and our genetic makeup [3,8,9]. Expectations are typically thought to originate from higher-level brain regions, such as the (pre)frontal cortex, which may guide information processing in lower-level sensory regions via top-down projections. In this framework, what we consciously see is proposed to be strongly influenced by the brain's expectations about, or its best guess of, the outside world [30,39,40]. Initial studies support the idea that the brain uses information in the environment to build expectations of stimulus frequency or conditional probabilities to modify perceptual processing [5,6]. These ideas have been formalized in theoretical models, such as predictive coding and sequential sampling models [2,3,41]. Although these frameworks are attractive in their simplicity, how exactly expectations shape conscious perception, and to what extent awareness guides the formation of expectations, is still largely unknown.

At present, there are (at least) three issues that need to be resolved to further our understanding of the relationship between expectations and consciousness. The first issue relates to the effect that expectations may have on conscious awareness itself. It has been shown that valid expectations increase the speed of conscious access [28,42–44] and may help selecting or facilitating stimulus interpretation when (visual) input is ambiguous or noisy [15,29,39,45,46]. It is yet an open question whether expectations can boost an otherwise unseen stimulus into conscious awareness, thereby enabling the switch from a nonconscious to a conscious stimulus representation, instead of merely facilitating its cognitive interpretation or its speed of appearance in time.

Second, it is an open question to what extent prediction errors, arising in a situation of invalid expectations, can be registered outside of conscious awareness. It has been shown that “oddball” stimuli (e.g. simple violations in auditory tone sequences) elicit early mismatch responses in electrophysiological signals: the mismatch negativity (MMN) [47,48]. Interestingly, MMN's can even be observed when attention is distracted from the tone sequences [10] or in several reduced states of consciousness, such as sleep [49], anesthesia [50] and vegetative state [10]. This suggests that the MMN

reflects a pre-attentive nonconscious prediction error signal [11,51,52]. However, it remains debated whether these signals originate in model-based comparisons of expectations to new input or merely reflect passive low-level sensory adaptation to repeated inputs [12,53]. The one study in which these mechanisms were dissociated in a nonconscious state showed adaptation remains operative during sleep, whereas prediction error detection disappears [13], questioning the notion that prediction errors may be registered nonconsciously.

The final issue concerns the role of awareness in implementing expectations. Many expectation-based models assume that expectations are implemented via top-down neural activation. Interestingly, influential theories of consciousness suggest that conscious access requires similar top-down interactions between higher-level (e.g. prefrontal) and lower-level (e.g. visual) brain regions, referred to as feedback or recurrent processing [16,23]. Information that does not reach conscious access is thought to only trigger feedforward activity or local recurrent interactions between posterior brain regions. Therefore, it is unclear how nonconscious information, in the absence of feedback signals from higher-order cortical areas, could lead to the implementation of expectations.

METHODS

Participants

We tested 26 participants in Experiment 1 (21 females, age 19.5 ± 1.3 years), 34 participants in Experiment 2 (27 females, age 20.0 ± 1.1 years) and 85 participants in Experiment 3 (63 females, age 22.0 ± 3.2 years). All participants were right-handed and had normal or corrected-to-normal vision.

For all experiments, participants for whom the minimum number of observations in one or more conditions was lower than 10, were excluded from analysis. Additionally, for Experiment 2 (EEG), we had to exclude 2 participants due to problems with the reference electrodes. In the end, this resulted in the inclusion of 25 participants for Experiment 1 (20 females, age 19.5 ± 1.3 years), 29 participants for Experiment 2 (22 females, age 20.0 ± 1.1 years) and 67 participants for Experiment 3 (49 females, age 21.9 ± 3.0 years).

The studies were approved by the local ethics committee of the University of

Amsterdam and written informed consent was obtained from all participants according to the Declaration of Helsinki. Compensation was 20 Euros for Experiment 1, 30 euros for Experiment 2 and 25 Euros for Experiment 3, or equivalents in course credit.

Materials

All stimuli were generated using the Psychophysics Toolbox [54] within a MATLAB environment (MathWorks, Natick, MA, USA). Stimuli were displayed on an ASUS LCD monitor (1920 x 1080 pixels, 120Hz, 50.9x28.6 cm screen size, 46.3pixels/°) on a “black” (RGB: [0 0 0], $\pm 3\text{cd/m}^2$) background while participants were seated in a dimly lit room, approximately 70 cm away from the screen.

Procedure and Stimuli

Participants performed an attentional blink (AB) task [20], in which on every trial a rapid series of visual stimuli was presented consisting of a sequence of 17 uppercase letters drawn from the alphabet but excluding the letters I, L, O, Q, U, and V. Every letter appeared maximally once per trial. Letters were presented at fixation in a mono-spaced font (font size: 40; corresponding to a height of approximately 1.2°) for 92 ms each.

Experiment 1

Participants were instructed to detect target letters within the rapid serial visual presentation (RSVP). The first target (T1: G or H) was always presented at the fifth position of the RSVP. On most trials (80%) it was followed by a second target (T2: D or K) at lag 2, lag 4 or lag 10 (respectively 183, 367 or 917 ms later). Each lag was equally likely. T1 was presented in green (RGB: [0 255 0]), while T2 and the distractor letters were white (RGB: [255 255 255], $\pm 320\text{cd/m}^2$).

Crucially, there was a predictive relationship between the two targets (FIGURE 2.1A). Namely, in the 80% of trials where a T2 was presented, the identity of T1 (e.g. G) predicted which T2 was likely (75%, e.g. D) or unlikely (25%, e.g. K) to appear. On the 20% remaining trials without a T2 a random distractor letter was presented at the T2-timepoint (every distractor letter was presented maximally once per trial) The mapping of T1 and T2 was counterbalanced over participants, so that for half of the participants the most likely target combinations were G-D and H-K while for the other half G-K and H-D were most likely. To be able to distinguish different lags in the absence

of a T2 stimulus, 4 grey squares (RGB: [200 200 200], $\pm 188\text{cd/m}^2$; size: 0.35° ; midpoint of each square centered at 1.30° from fixation) were always presented around the stimulus (T2 or distractor) at the T2-timepoint. Participants were instructed to use the timing information this cue provided when making decisions about the presence of a T2 (only the letters “D” or “K”, all other letters where distractors).

Following a 150 ms blank period at the end of the RSVP, participants gave their responses. First, they indicated whether or not they had seen any T2 by pressing the left or right shift key on the keyboard. The mapping between the keys and the response options was randomized per trial to decouple participants’ responses from the decision they had to make. Then they were asked to make a forced choice judgment about the T2 letter (D or K) that was presented by typing in this letter. Finally, they made a similar response about the identity of T1 (G or H). We used long response timeout durations of 5s and participants were instructed to value accuracy over response speed. The inter-trial interval, as defined by the time between the last response and the onset of the stream, was 500-750 ms.

The experiment consisted of two one-hour sessions on separate days within one week. In the first session, participants received instructions about the task and subsequently performed the task for 6 blocks of 75 trials (total 300 trials). The goal of the training session was to familiarize participants with the task. Besides, since we did not instruct participants about the predictive relationship between T1 and T2, some practice on the task was required for them to (implicitly) learn this relationship. In the second session, participants first received a summary of the instructions, after which the actual experiment started. Participants performed 6 blocks of 90 trials (total of 540 trials) of the AB task. The first three participants performed 6 blocks of 105 trials (630 trials). In both sessions participants received summary feedback about their performance at the end of each block, followed by a short break.

Experiment 2 (EEG)

The task in the EEG experiment was the same as in Experiment 1, except that in Experiment 2, we only asked participant to give one response by typing in the target letters they observed. In addition, we only used two different lags: lag 3 (275 ms; 2/3rd of trials) and lag 10 (917 ms; 1/3rd of trials). To further increase the number of trials, the ITI range was reduced to 200-400 ms.

Again, the experiment consisted of two different sessions within one week. The first session (1 hour) consisted of instructions followed by extensive training (720 trials over 6 blocks) on the task. Participants were not explicitly informed about the predictive relationship between the targets. In the second session (2 hours) we first prepared the participant for the EEG measurements and gave brief instructions about the task. Then, participants performed 12 blocks of 120 trials (total 1440 trials) of the AB task.

Experiment 3

To investigate the importance of T1 detection for expectation effects on conscious access, we adjusted the task we used in Experiment 1 to decrease the visibility of T1 (FIGURE 2.3A). We now presented T1 in white instead of green, to make it stand out less among the other stimuli. Furthermore, T1 duration was staircased per participant such that participants could report T1 on roughly 75% of the trials. Starting in the second half of the training and continuing in the experimental session, after each block T1 duration was decreased by one frame (8 ms) if performance was higher than 85% and increased by one frame if performance was lower than 65%. To ensure T1 duration would not deviate too much from the duration of other stimuli, T1 duration was only allowed to be in the range of 42-142 ms (max. 50 ms different from other stimuli). The median duration of T1 in the second session was 125 ms. On 20% of trials no T1 was presented and a random distractor letter was presented instead. When both targets were present, T1 predicted which T2 was likely to follow with 75% accuracy.

We made a few changes to the task design to increase the efficiency of the design. The ITI was reduced to values between 300-500 ms. In addition, we only asked participants for one response. They were asked to type in any target letter they had seen during the trial and refrain from typing in a T1 and/or T2 letter when they did not see any. The response was confirmed by pressing the space bar on the keyboard or when a timeout of 4s had passed. To further increase the number of trials per condition, we decided to use only lag 3 (2/3rd of trials) and lag 10 (1/3rd of trials). Because T1 duration was staircased on an individual basis, the T1-T2 SOA differed between participants. On average, lag 3 corresponded to an SOA of 308 ms while lag 10 corresponded to an SOA of 950 ms.

Finally, we manipulated the instructions we gave to participants in order to see to what extent explicit knowledge of the relationship between T1 and T2 affected our results.

As in Experiment 1, we tested participants during two separate sessions within one week. The first group of the participants (N=25) did not receive any explicit instruction about this relationship, similar to Experiment 1. The second group of participants (N=19) received explicit instructions about the T1-T2 relationship at the start of the second session, and a third group of participants (N=23) received those instructions already at the start of their first session.

The first session (1 hour) was used for instructions and training the participants on the task (10 x 75 trials). The experimental session in which participants performed the AB task lasted 1.5 hour and contained 16 blocks of 75 trials (1200 trials).

Behavioral analyses

Preparatory steps were done with in-house MATLAB scripts. Statistical analyses (repeated measures ANOVAs and paired t-tests) were performed using JASP software [55]. In situations where a specifically tested hypothesis did not yield a significant result, we used a Bayesian equivalent of the same test to quantify the evidence for the null-hypothesis [56,57]. In those cases, using JASP's default Cauchy prior, Bayes Factors (BF) were computed for each effect. To increase the interpretability in analyses with multiple factors, we used Bayesian model averaging to get a single BF for each effect in ANOVA's. This BF is the change from prior to posterior inclusion odds, and can intuitively be understood as the amount of evidence the data gives for including an experimental factor in a model of the data. The BF will either converge to zero when the factor should not be included, or to infinity when it should be included in the model. Values close to one indicate that there is not enough evidence for either conclusion. We use the conventions from Jeffreys [58] to interpret the effect sizes of our Bayesian analyses.

Experiment 1

In our behavioral analyses we looked at the T2 detection performance, given that T1 was correctly identified. A response was considered to be correct when (1) the participant indicated no T2 was present when no T2 was presented or (2) the participant correctly indicated a T2 was present and subsequently reported the correct target letter. Since expectation is only a meaningful concept when a T2 target was presented, the T2 absent trials, on which a distractor letter was presented instead of a T2, were not taken into consideration for the main statistical analyses. Trials where

one of the responses was missing were deleted from all analyses. Percentage correct was used in a 3 x 2 repeated measures ANOVA with the factors lag (lag 2, lag 4, lag 10) and expectation (valid, invalid). In a control analysis, we repeated our analyses for Experiment 1 based on the T2 detection responses (ignoring the accuracy of the T2 identification) as dependent variable (see also **RESULTS**). Since the seen/miss response is orthogonal to the specific expectations about target letters, this analysis rules out simple response biases as a potential cause of our effects.

Experiment 2 (EEG)

The behavioral analyses for the EEG experiment were similar to those for Experiment 1. However, the factor lag had only 2 levels (lag 3, lag 10). Percentage correct T2 detection was computed as in Experiment 1 using only the trials on which T1 was correctly reported. A response was considered to be correct when the letter a participant entered was the letter that was presented or when a participant refrained from entering a letter when none was presented for the T2 absent trials. In addition, we computed a T2 detection measure to use in a control analysis: if a participant typed in any letter, we categorize the response as a “target seen” response, otherwise we call it a “target absent” response. This outcome measure was used in a control analysis.

Experiment 3

In this experiment, participants gave only one response by typing in the target they had perceived. Trials on which no response was given or on which an impossible response was given (e.g. two T1 targets reported) were excluded from analyses. For T1 and T2 separately, we assessed the accuracy of the responses. The definition of correct and incorrect responses was the same as in Experiment 2 and we also used the same T2 detection measure.

Subsequently, T2 percentage correct detection was used in a 2 x 2 x 2 x 3 mixed ANOVA with the within-subject factors lag (lag 3, lag 10), expectation (valid, invalid) and T1 visibility (T1 seen, T1 missed) and the between-subject factor instruction (none, start session 2, start session 1). As mentioned before, this between-subject factor was included to find out whether predictive effects would be modulated by explicit knowledge of the relation between T1 and T2. To investigate the effect of T1 visibility in more detail, we followed up the main analyses by other mixed ANOVAs in which we first split up the dataset based on T1 visibility. In situations where we

found interactions with the factor instruction, we compared the effects of lag and expectation separately per instruction condition using repeated measures ANOVAs and paired-samples t-tests.

Finally, to test for an interaction between expectation validity and lag, we combined behavioral data from all experiments in a post-hoc analysis. Only trials on which T1 was correctly identified were used. For Experiment 1 we averaged data for lag 2 and lag 4 to create an average “short lag” condition. Subsequently, these data were entered into a 2 x 2 x 3 mixed ANOVA with the within-subject factors lag (short, long) and expectation (valid, invalid) and the between-subject factor experiment (Experiment 1, Experiment 2, Experiment 3).

Electroencephalographic Measurements

EEG was recorded with a BioSemi ActiveTwo system and sampled at 512 Hz (BioSemi, Amsterdam, The Netherlands). Potentials were measured from sixty-four scalp electrodes, along with two reference electrodes on the earlobes and four electrodes measuring horizontal and vertical eye movements. After data acquisition, EEG data was pre-processed with the FieldTrip toolbox for MATLAB [59]. First, data were re-referenced to the linked earlobes, high-pass filtered at 0.01 Hz, and epoched from -0.750 to 1s surrounding the onset of T2. Data were visually inspected and trials and/or channels containing artifacts not related to eye blinks were manually removed, resulting in deletion of on average 9.1% ($\pm 3.9\%$) of trials and 2.0 (± 1.7) channels. Independent component analysis was used to identify components related to eye blinks or other artifacts that could easily be distinguished from other EEG signals. After the independent component analysis, previously deleted channels were reconstructed based on a nearest neighbor approach. Trials were baseline corrected to the average amplitude prior to T1 onset (-0.750 to -0.275 relative to T2). As a final step, we applied a 40Hz low-pass filter to the trial data, after which ERPs were created separately for each condition of interest.

Electroencephalographic analyses

All EEG analyses are based exclusively on trials where T2 appeared at lag 3 and T1 was correctly identified. We used a combination of Fieldtrip [59] and in-house MATLAB scripts to perform our analyses. As a first step, we performed cluster-based permutation tests [60] on the time-window 0-750 ms from stimulus onset to isolate significant ERP

events relating to expectation validity (valid, invalid; irrespective of T2 visibility) or T2 visibility (seen, missed; irrespective of validity) or the interaction between these factors. Next, we used an in-house built MATLAB script to isolate the significant events as clusters in time and space. For this purpose, we computed an average difference wave over all channels that were part of the cluster at any point in time. Subsequently, the onset and offset of a cluster were defined as the time period around the maximum difference where the difference did not drop below 50% of this maximum and where at least one channel showed a significant effect. We then selected the 10 channels that showed the largest effect in this time-window. One of the observed events reflected a mixture of the traditionally observed P3a and P3b components [61,62]. Therefore, we split the event into two clusters by manually selecting either the 32 most anterior or 32 most posterior EEG channels (from the central midline) before running the cluster selection procedure.

As an alternative way to establish potential interactions between T2 detection and validity, we inspected the clusters that were isolated in the previous step in more detail. This may be a more powerful (but also less sensitive) way to detect small effects, because data is averaged over more time-points and channels. Within each of the clusters, we performed a 2 x 2 repeated measures ANOVA (and its Bayesian equivalent, see also **BEHAVIORAL ANALYSIS**) with the factors T2 detection (seen, missed) and expectation validity (valid, invalid) on the cluster data averaged over channels and time. To prevent double dipping, in each analysis we only considered the effects orthogonal to the one that was used to define the cluster (e.g. not testing the effect of expectation in a cluster defined based on the expectation effect).

RESULTS

Experiment 1: (how) do expectations affect conscious access?

In the first experiment we addressed the question if expectations about the likelihood of stimulus identity modulate the likelihood of conscious access, and if so, in what direction. To do so, we used the attentional blink paradigm [20]. The attentional blink is an impairment in the conscious perception of the second of two target stimuli that are presented in rapid succession when the initial target was correctly perceived. Here we modified the paradigm in such a way that the first target (T1: the letter G or H, in green) predicted which of the second targets would most be likely to appear in case a T2 target was presented (T2: the letter D or K, predicted=75%, unpredicted=25%,

in white, **FIGURE 2.1A**). On 20% of trials we presented a random distractor letter instead of a T2 target. At the end of each stream of letters, participants gave three responses. First, they indicated whether or not they had seen any of the two T2 targets (“seen”/“unseen” response). Second, they were prompted to make a forced-choice judgment about the identity of T2 (whether the letter D or K was presented). Third, participants had to make a similar forced-choice decision about the identity of T1 (whether the letter G or H was presented) (see **METHODS** for details). Participants were not explicitly instructed about the predictive relationship between T1 and T2.

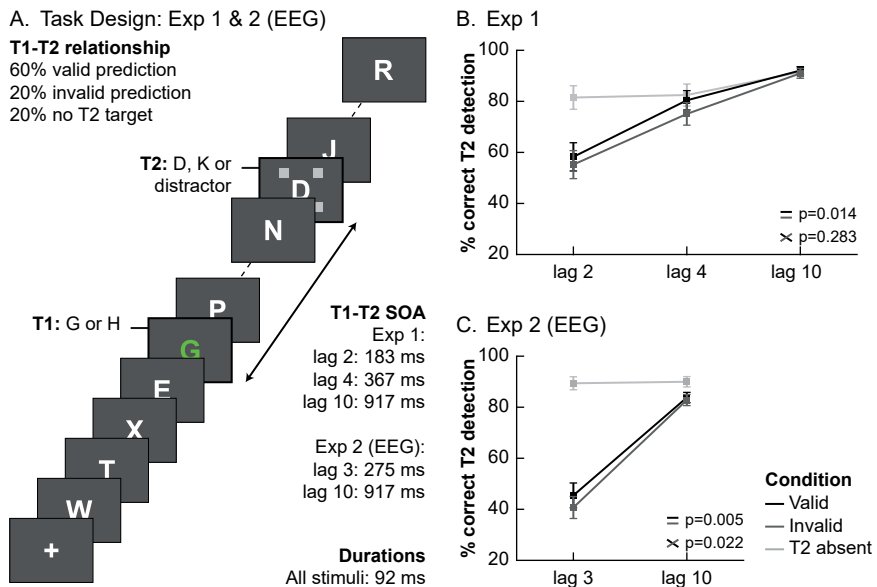


Figure 2.1: Task design and behavioral results of Experiment 1 and 2. (A) The trial structure of the attentional blink task used in Experiment 1 and 2. Each trial consisted of a stream of rapidly presented letters in which predefined target letters had to be detected and then reported at the end of the stream. The first target (T1: a green G or H) always appeared at the fifth position. The second target (T2: D or K), was presented at varying SOAs (lags) after the first one and was marked by placeholders. The identity of T1 predicted which of the T2 targets was most likely to appear, thereby introducing validly and invalidly predicted T2 targets. On 20% of the trials no second target was presented and a random distractor letter was presented instead. **(B)** Percentage correct T2 target detection at each of the T1-T2 lags, for valid expectations, invalid expectations and T2 absent trials in Experiment 1. Validly predicted T2’s were significantly more often perceived than invalidly predicted T2’s. **(C)** Percentage of T2 target detection at each of the T1-T2 lags, after a valid or invalid expectation or on a T2 absent trial for Experiment 2. Again, validly predicted T2’s were more often perceived, in particular at short lags. Error bars represent SEM.

In **FIGURE 2.1** we plot the percentage of trials in which T2 was correctly detected and T1 discrimination was also correct (average T1 accuracy was 94.20%, $sd=5.77\%$) for the three different lags (lag 2, 4 and 10). T2 was considered to be detected correctly when participants indicated to have seen it (based on the first response) and correctly

identified it (based on the second response). Overall, there was a clear attentional blink, as reflected by reduced T2 detection when the time (i.e. lag) between T1 and T2 was shorter (FIGURE 2.1B, main effect of lag: $F_{2,48}=48.15$, $p<0.001$). Importantly, expectations modulated T2 detection rate. T2 detection was significantly better when T1 validly predicted T2 (black lines) compared to when the expectation was invalid (gray lines, main effect of validity: $F_{1,24}=7.10$, $p=0.014$, no significant interaction between lag and validity: $F_{2,48}=1.30$, $p=0.283$). These results extend several previous studies [28,29,42,43,63] by showing that conscious perception is (partly) determined by the transitional probability of the input the brain receives.

While these data support the notion that valid expectations trigger access to consciousness, it has been recognized that such findings may not solely be due to changes in perception, but perhaps (also) due to changes in decision criteria or response biases [64–66]. To rule out the possibility that our effects could be explained by a response bias in which people simply report the target letter that they expected based on T1, irrespective of whether they consciously perceived T2, we performed an analysis with T2 detection (instead of T2 discrimination, see METHODS) as the dependent variable. This analysis takes into account only participants' first response (the “seen”/“unseen” response), regardless of whether subsequent T2 letter identification was correct or not. Crucially, this analysis cannot be influenced by any decision/response biases because the response was orthogonal to the participants' expectation. Information about the most likely letter to appear cannot predispose participants to better determine whether a target letter was presented at all. Still, we observed a qualitatively similar pattern of results (main effect of validity: $F_{1,24}=5.47$, $p=0.028$). This finding suggests that validity indeed boosted participants' conscious access of T2, instead of merely eliciting a shift in response bias.

Experiment 2: EEG markers of conscious and nonconscious expectation violations

Subsequently, we tested whether expectation violations can be elicited by nonconsciously processed unpredicted stimuli or whether conscious perception of a stimulus is a prerequisite for it to trigger neural expectation error responses. To test this, we measured subjects' brain activity with EEG while they performed a similar task as in Experiment 1. First, we replicated the behavioral effects of Experiment 1 (FIGURE 2.1C). Overall, T1 performance was high ($M=93.61\%$, $sd=7.31\%$) and T2 detection was higher at lag 10 than at lag 3 (main effect of lag: $F_{1,28}=128.72$, $p<0.001$), reflecting a

robust attentional blink. More importantly, validly predicted T2s were discriminated better than invalidly predicted T2s (main effect of validity: $F_{1,28}=9.49$, $p=0.005$). The effects were similar in a control analysis where we considered the percentage of T2 seen responses (regardless of the exact letter participants entered), making it less likely that our effect can be explained by a response bias (main effect of validity: $F_{1,28}=4.23$, $p=0.049$). In this experiment, the validity effect was significantly modulated by lag (validity \times lag: $F_{1,28}=5.86$, $p=0.022$), an effect that was numerically similar, but not significant in Experiment 1. Participants performed better for valid than invalid trials at lag 3, but there was no convincing evidence for an effect of expectations at lag 10 (lag 3 validity effect: $t_{28}=3.40$, $p=0.002$; lag 10 validity effect: $t_{28}=0.98$, $p=0.334$). Thus, expectation effects were larger at the time at which T2 more often goes unperceived.

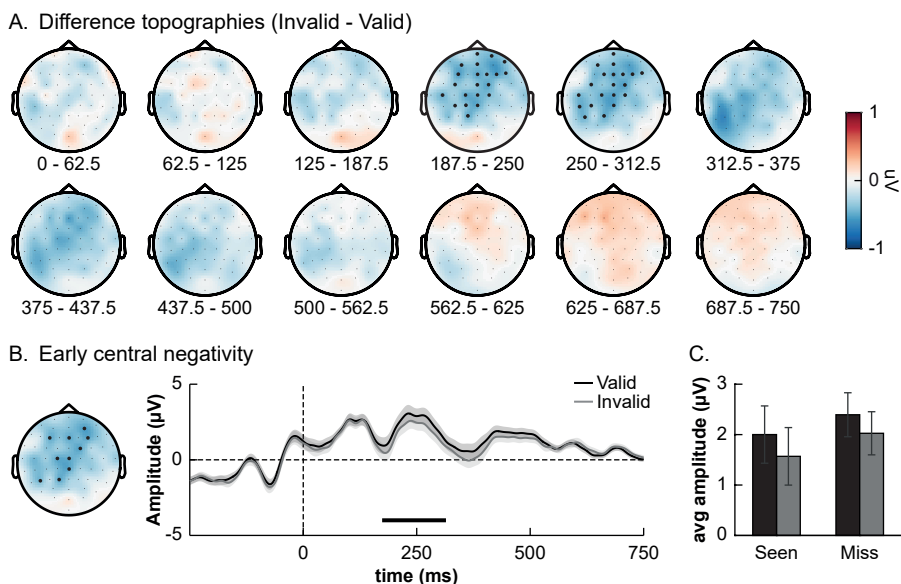


Figure 2.2: ERP effects related to T2 prediction validity. (A) Topographic maps of the difference between validly and invalidly predicted T2s over time (0 = T2 onset). Cluster-based permutation tests were used to isolate the significant events, while correcting for multiple comparisons across time and (electrode) space. On each head map, channels with a significant effect for at least 50% of its time window are highlighted. (B) The average ERP time-course of the 10 channels shown on the headmap on the left, shown separately for each validity condition. The significant time-window is marked by a black line above the x-axis. Invalidly predicted T2s were associated with greater fronto-central negativity than validly predicted T2s. (C) Bar graphs showing the average amplitude of the four conditions (visibility \times prediction) for the significant neural event shown in B. In all plots error bars represent SEM.

Next, we investigated potential differences in the neural processing of predicted and unpredicted stimuli, as a function of stimulus awareness. To this end, we contrasted invalidly and validly predicted T2s and tested this difference using cluster-based

permutation testing, correcting for multiple comparisons across both time (0-750 ms) and (electrode) space (see **FIGURE 2.2** and **METHODS**) [60]. We found one significant difference over fronto-central electrode channels, which reflected greater T2-elicited negativity for invalid compared to valid trials between 174-314 ms ($p=0.015$, **FIGURE 2.2B**), therefore potentially reflecting some type of mismatch response. We then further analyzed this event to test whether the difference was modulated by, or dependent on, conscious perception of T2. Crucially, the size of this fronto-central mismatch component was independent of T2 awareness ($F_{1,28}=0.04$, $p=0.850$, $BF=0.254$, **FIGURE 2.2C**), indicating that both seen and unseen T2's generated a fronto-central mismatch response.

Additionally, analyses of T2 visibility effects (irrespective of expectation validity) replicated previously reported findings [61,67,68]. In these analyses, we examined the difference in ERPs following seen and missed T2s using a cluster-based permutation test (**FIGURE 2.3**), revealing two significant events. First, a significant negative difference could be observed over (left) posterior electrodes from 170-355 ms after T2 onset ($p=0.010$; **FIGURE 2.3A**). This event was followed by a significant long-lasting positive event ($p<0.001$), reflecting a mixture of the P3a and P3b components, extending over frontal and central electrodes.

Subsequently, we had a closer look at the interactions between conscious access and expectation validity. Therefore, we analyzed the ERP events that were isolated in the previous step in more detail (**FIGURE 2.3B-G**). For this analysis we first isolated the traditionally observed AB-related P3a and P3b ERP components from the long-lasting positive ERP event that differentiated between seen and missed T2s [61]. Doing so resulted in an early positive P3a cluster (**FIGURE 2.3D**) over fronto-central channels that was significant between 395-586 ms and a somewhat later positive P3b cluster (**FIGURE 2.3F**) over more posterior parietal channels, which was significant between 445-611 ms. Within each of these clusters we performed repeated measures ANOVAs with the factors validity and T2 detection.

For none of the events we found evidence that the T2 detection effect was modulated by expectation validity (early left-posterior event: $F_{1,28}=0.29$, $p=0.597$, $BF=.260$; P3a: $F_{1,28}=1.56$, $p=0.222$, $BF=0.230$; P3b: $F_{1,28}=2.10$, $p=0.159$, $BF=0.296$), though the BF values suggest that the evidence for the absence of such interactions is moderate

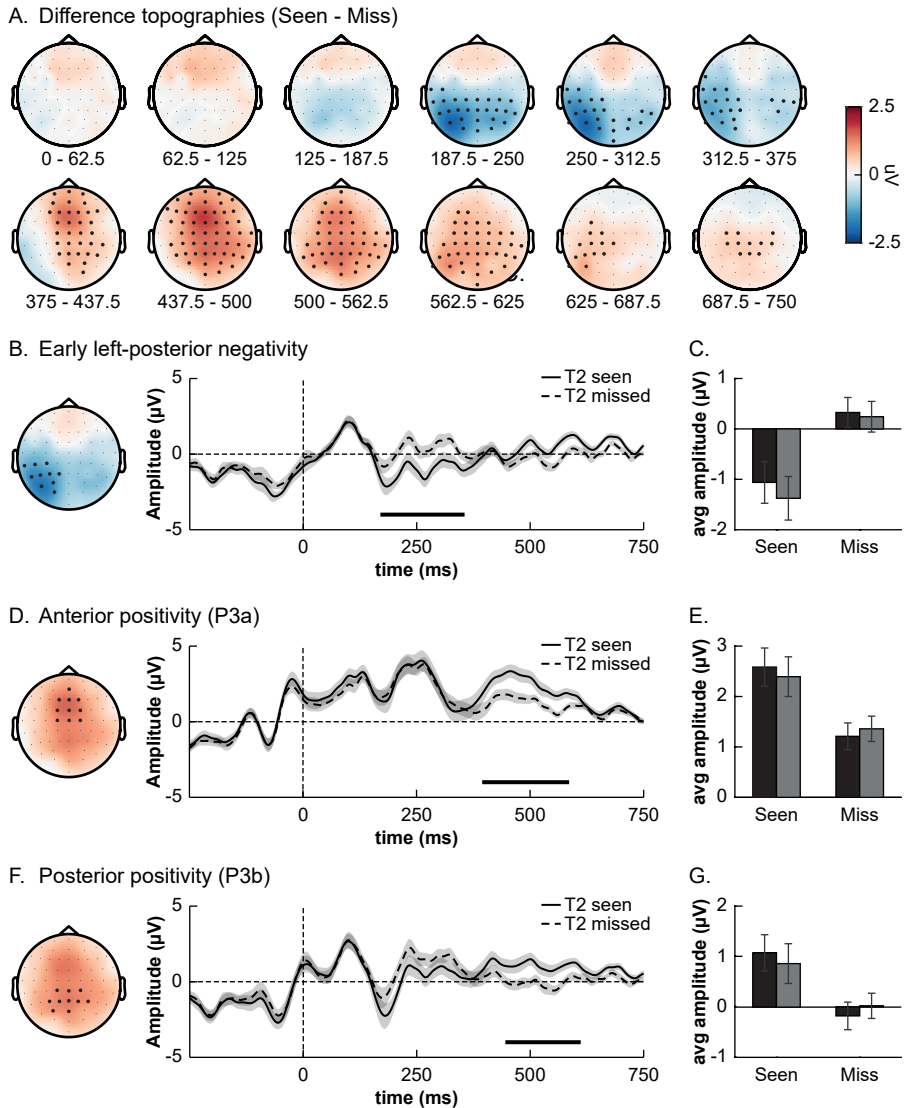


Figure 2.3: ERP effects related to T2 visibility analyses. (A) Topographic maps showing the difference between seen and missed T2s over time (0 = T2 onset). Cluster-based permutation tests were used to isolate the significant events while correcting for multiple comparisons across time and (electrode) space. On each head map, channels showing a significant difference for at least 50% of its time window are highlighted. Three events were isolated based on the permutation tests. (B,D,F) For each of the events individually, the average ERP time-course of the 10 channels shown on the headmap on the left, separately for T2 seen and T2 missed conditions is shown. The significant time-window is marked by a black line above the x-axis. (C,E,G) Bar graphs showing the average amplitude of the four conditions (visibility x prediction) for the significant neural events shown in (B,D,F). In all plots error bars represent SEM.

at best. This is somewhat surprising, because especially the late positive events have previously been related to conscious access [61,69] and metacognition [70].

However, recent investigations show it may also reflect even later more cognitive processes, merely arising as a consequence of becoming consciously aware of information [71,72]. We did not find evidence that the amplitude of these ERP events was modulated by expectation validity, which may suggest that once a stimulus has been perceived consciously, it is irrelevant whether or not the expectation was valid.

Finally, we directly tested for an interaction between conscious access and expectation by comparing the validity ERP effect (invalid-valid) for T2 seen and T2 missed trials in a cluster-based permutation test (this analysis takes into account the entire scalp topography). Again, no significant interactions between these factors were observed (all clusters $p > 0.10$).

Experiment 3: the role of conscious awareness in implementing top-down expectations

In our final experiment, we addressed the question whether expectation formation itself can unfold in the absence of awareness and subsequently influence conscious access (FIGURE 2.4). To address this question, we changed the color of T1 from green to white and for each subject staircased T1 duration in such a way that T1 was correctly identified on approximately 75% of the trials (actual T1 identification performance: $M=76.03\%$; $sd=8.65\%$). T1 duration did not differ between trials where T2 was seen and trials where T2 was missed (T2 detection: $t_{66}=0.31$, $p=0.752$; T2 seen: $M=117.42$ ms; T2 missed: $M=117.46$ ms), which indicates that T1 visibility was not determined by stimulus duration. Likely, internal fluctuations in the system (e.g. variability in attention) must be causing participants to sometimes see T1 and sometimes miss it. Moreover, on 20% of trials no T1 was presented (but replaced by a distractor). Further, to test to what extent explicit knowledge of the predictive relationships between stimuli would increase the validity effects, we varied the moment at which explicit information about the predictive relations between T1 and T2 was provided. The experiment consisted of a training session and a test session on separate days. A first group of subjects received no explicit instructions about the predictive relations in either session and had to learn them implicitly through experience with the task; the second group received explicit instructions about the T1-T2 relations in the test session only, but not in the first training session; and the third group received explicit instructions already from the start of the experiment.

T1 visibility strongly affected T2 detection. When T1 was seen, T2 detection was markedly lower than when T1 was missed (main effect of T1 awareness: $F_{1,64}=4.62$, $p=0.035$), in particular at short lags (T1 awareness x lag: $F_{1,64}=72.95$, $p<0.001$). Validly predicted targets were detected more often (main effect of validity: $F_{1,64}=33.39$, $p<0.001$). The effect of expectation validity on T2 detection varied as a function of T1 awareness and instructions (T1 awareness x validity: $F_{1,64}=40.55$, $p<0.001$; validity x instruction: $F_{1,64}=5.91$, $p=0.004$; T1 awareness x validity x instruction: $F_{2,64}=11.33$, $p<0.001$). When T1 was seen (FIGURE 2.4B), a clear attentional blink was observed (main effect of lag: $F_{1,64}=170.01$, $p<0.001$) and validly predicted targets were more often detected than invalidly predicted targets (main effect of validity: $F_{1,64}=64.97$, $p<0.001$) (as in Experiment 1 and 2). Like in the previous experiments, a control analysis considering only the percentage of T2 seen responses (regardless of the exact letter participants entered) also revealed a significant effect of validity (main effect of validity: $F_{1,64}=65.83$, $p<0.001$), making it unlikely that response biases are causing the effect. Interestingly, we also observed a significant attentional blink for missed T1's, reflecting a nonconsciously elicited AB (main effect of lag: $F_{1,64}=74.42$, $p<0.001$). This AB effect cannot be explained by an overall T2 detection performance benefit for targets that are presented later in the trial because the AB was larger for trials on which T1 was presented but missed compared to trials on which no T1 was presented in the trial (lag x T1 presence: $F_{1,66}=24.19$, $p<0.001$). However, although missed T1's triggered an AB, expectation validity did not affect T2 detection performance for missed T1's (main effect of validity: $F_{1,64}=0.35$, $p=0.554$), regardless of the type of instruction participants received about the predictive relation between T1 and T2 (validity x instruction: $F_{2,64}=0.64$, $p=0.533$). A Bayesian equivalent of the repeated-measures analysis strongly suggested validity should not be included in a model of the data ($BF=0.084$).

The above results highlight that only when T1 was seen, valid expectations facilitated T2 detection. A post-hoc analysis on T1-seen trials only revealed that this effect was modulated by how explicitly we instructed participants about the predictive relationship between T1 and T2 (validity x instruction: $F_{2,64}=14.83$, $p<0.001$). The validity effect, as defined by the difference between valid and invalid trials, averaged across the two lags, increased with more explicit instructions (group 1: 1.87%, group 2: 19.53%, group 3: 26.27%). These results indicate that, not only does the visibility of T1 define the predictive impact on T2 detection, but also the extent to which these predictive relations are (explicitly) known affects the impact of expectations on conscious access.

This may also explain why the validity effect appeared more pronounced in Experiment 3 compared to Experiments 1 and 2, because in Experiment 1 and 2 subjects were not explicitly instructed about the predictive relations between T1 and T2.

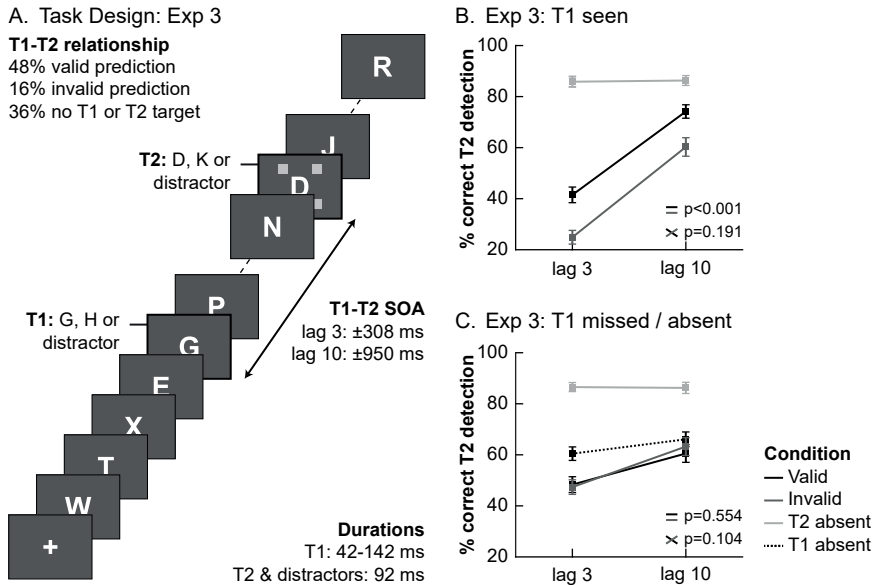


Figure 2.4: Task design and behavioral results of Experiment 3. (A) Trial structure of the task used in Experiment 3. T1 visibility was staircased at approximately 75% correct by manipulating its duration (on 20% of trials no T1 was presented). **(B)** Percentage of correct T2 target detection at each of the T1-T2 lags, after a valid or invalid expectation and on a T2 absent trials, for trials where T1 was correctly reported (T1 seen). As in Experiment 1 and 2, when T1 was seen, validly predicted T2’s were more often detected than invalidly predicted T2’s. **(C)** Solid lines show percentage of T2 target detection at each of the T1-T2 lags, after a valid or invalid expectation and on a T2 absent trials, for trials where T1 was presented but missed. In contrast to T1 seen trials **(B)**, when T1 was not seen, validity did not enhance T2 detection. However, a missed T1 still triggered a significant attentional blink, as compared to trials on which no T1 was presented (dotted line). Error bars represent SEM.

Finally, in contrast to Experiment 2, on T1 seen trials the validity effect was independent of lag (validity x lag: $F_{1,64}=1.750$, $p=0.191$). Since we anticipated stronger expectation effects at short lags, behavioral data from all three experiments was combined in a post-hoc analysis. Only trials on which T1 was correctly identified were used and for Experiment 1 we averaged data for lag 2 and lag 4 to create an average “short lag” condition. A significant interaction between validity and lag showed that across all experiments, the expectation effect was stronger at short lags compared to the long lags (validity x lag: $F_{1,118}=5.73$, $p=0.018$; no validity x lag x experiment interaction: $F_{2,118}=0.065$, $p=0.937$).

DISCUSSION

In this report we investigated three important questions regarding the intricate relationship between top-down expectations and conscious awareness. The first question that we addressed was how prior information about the identity of an upcoming stimulus influences the likelihood of that stimulus entering conscious awareness. Using a novel attentional blink paradigm in which the identity of T1 cued the likelihood of the identity of T2, we showed that stimuli that confirm our expectation have a higher likelihood of gaining access to conscious awareness than stimuli that violate our expectations, especially at short lags. The expectation effect was qualitatively similar across all three experiments, regardless of subtle experimental differences in task design and overall performance between those experiments. Furthermore, it could not be explained by simple shifts in response criterion, because it was also present for a dependent measure orthogonal to the expectation manipulation. Together, this suggests that valid expectations amplify the perceptual strength of a stimulus and therefore increase the chance of conscious access, possibly due to the sharpening of its neural representations [6]. This interpretation is supported by previous experiments that have shown varying effects of expectations on (subjective) perception, such as studies showing prior knowledge increases the speed [29,42–44] and accuracy [28] of stimulus detection. Furthermore, our findings complement recent studies showing that the AB can be reduced when there is knowledge about temporal statistics of the task [73,74] or when the latency of T2 targets is explicitly cued [75,76]. In addition, two recent reports have shown that, in some subjects, expecting the presentation of a stimulus can even elicit an illusory stimulus percept when no stimulus is presented, when attention is diverted [77,78]. Future experimentation is required to shed light on the generalizability of our effect to simpler tasks. Such experiments may also consider using other measures of subjective perception (e.g. perceptual awareness scale) [79].

The second question that we addressed was related to the extent to which nonconscious stimuli can trigger prediction error responses, as measured with EEG. Over the last 20 years, we and others have shown that nonconscious information processing is rather sophisticated [19,80], and that a diverse range of high-level cognitive processes can unfold nonconsciously [81–88]. Interestingly, in Experiment 2 we found that expectations that are violated by a nonconscious stimulus trigger a stronger negative fronto-central ERP component than expectations that are confirmed. This neural

event was similar for trials on which T2 was seen and on trials where T2 was missed, highlighting that conscious awareness of a stimulus is not a prerequisite for it to trigger a prediction error response [89,90]. This effect may reflect a mismatch signal, similar to the mismatch negativity [48], which is a negative deflection following oddball stimuli that develops 100-200 milliseconds after stimulus onset. Sometimes this effect lasts longer, in some experiments until ~400 milliseconds, depending on the specifics of the task and stimulus material [11,47,52]. While in terms of interpretation this effect is similar to a mismatch effect, its topography is slightly different than a typical visually evoked MMN, which generally peaks more posteriorly, although considerable variation in its topography has been reported [47]. Alternatively, it is possible that the higher activation for valid compared to invalid trials corresponds to the frontal selection positivity (FSP), which is a well-known marker of non-spatial attentional processes [91]. In our paradigm, this could be explained as improved attentional selection when expectations are confirmed. Although the exact nature of the observed component deserves future experimentation, the key finding is that the effect was independent of T2 perception and purely depends on the validity of the expectation. This is in line with studies that have shown context influences on nonconscious information processing [92–95], studies that have shown that the MMN can be observed when the expectation violations are unattended [10,11,52,96,97] and more generally evidence for relatively high-level processing of nonconscious stimuli [72,80,98]. Nevertheless, the absence of interactions in the ERP is also somewhat surprising (but see also Rutiku et al. [99]), because as noted earlier such interactions between expectation validity and conscious T2 detection were present in behavior. A neural basis for this effect should exist, but may be very subtle. Recently, a study by Aru et al. [45] found early (<100 ms) differences in signal amplitude over posterior channels that predicted the behavioral benefit of prior knowledge on the detection of stimuli presented at the threshold of perception. Another potentially interesting signature to investigate could be the onset of components related to conscious perception [42] and how they relate to expectations. Moreover, it is possible that instead of signal strength, it is the signal-to-noise ratio or sharpness of the representation that is improved [6]. Possibly, valid expectations do not modulate the amplitude of the neural response, but increase the specificity of the neural representation.

In the final experiment, we showed that conscious perception of T1, initiating the

expectation, is a prerequisite for influences on conscious access to occur. On the subset of trials where subjects did not see T1, there was no expectation effect of T1 on T2 detection performance. This result contrasts with findings from a recent study that suggested that some priors may operate nonconsciously [100]. Chang and colleagues presented participants with masked grey-scale natural scene images and found that the nonconscious processing of these images improved subsequent recognition of their degraded counterparts, so-called “Mooney images”, presented seconds later. One explanation for this difference is that the priors on which the effects of Chang et al. relied may be more automatic and hard-wired than the relatively arbitrary relationships that people have to learn and actively use in our experiments. It is possible that lower-level, automatic expectations are more easily processed outside of awareness compared to the more active ones studied here.

Further, it is also possible that with more training we would find nonconscious expectation effects. However, since subjects were already trained on the task on a separate day before performing the experimental session, this possibility seems unlikely. We did observe greater validity effects when subjects were made explicitly aware of the predictive nature of T1, suggesting that explicit knowledge of stimulus associations can facilitate the effects of stimulus-induced expectations. Finally, it should be noted that we did not test the full range of timing intervals between T1 and T2. It has been shown and proposed that the processing of nonconscious stimuli is relatively fleeting [23,101,102], so it is conceivable that the T1-T2 lags that we have used here may have been too long to observe expectation effects triggered by unseen T1’s. Further, a significant attentional blink was observed on trials on which T1 was missed, indicating that attention was still captured by a missed T1 at the T1-T2 lags used here. This latter result is in line with evidence showing that nonconscious stimuli are able to trigger attentional capture [103–105] and with a study showing lower T2 detection for T1’s that were missed compared to trials without a T1 (in that experiment this effect was independent of lag [106]).

In summary, three main conclusions can be drawn from the present series of studies. First, expectation confirmation, compared to violation, increases the likelihood of conscious awareness, suggesting that valid expectations amplify the perceptual strength of a stimulus. Second, nonconscious violations of conscious expectations are registered in the human brain. Third, however, expectations need to be implemented

consciously to subsequently modulate conscious access. These results suggest a differential role of conscious awareness in the hierarchy of predictive processing, in which the active implementation of top-down expectations requires conscious awareness, whereas a conscious expectation and a nonconscious stimulus can interact to generate prediction errors. How these nonconscious prediction errors are used for updating future behavior and shaping trial-by-trial learning is a matter for future experimentation.

3

EXPLORING THE EFFECT OF EXPECTATIONS AND STIMULUS RELEVANCE ON STIMULUS-SPECIFIC NEURAL REPRESENTATIONS AND CONSCIOUS REPORT

ABSTRACT

Subjective experience can be influenced by top-down factors, such as expectations and stimulus relevance. Recently, it has been shown that expectations can enhance the likelihood that a stimulus is consciously reported, but the neural mechanisms supporting this enhancement are still unclear. We manipulated stimulus expectations within the attentional blink (AB) paradigm using letters, and combined visual psychophysics with magnetoencephalographical (MEG) recordings to investigate whether prior expectations may enhance conscious access by sharpening stimulus-specific neural representations. We further explored how stimulus-specific neural activity patterns are affected by the factors expectation, stimulus relevance and conscious report. First, we show that valid expectations about the identity of an upcoming stimulus increases the likelihood that it is consciously reported. Second, using a series of multivariate decoding analyses, we show that the identity of letters presented in and out of the attentional blink can be reliably decoded from MEG data. Third, we show that early sensory stimulus-specific neural representations are similar for reported and missed target letters in the AB task (active report required) and an oddball task in which letter was clearly presented but its identity was task-irrelevant. However, later sustained and stable stimulus-specific representations were uniquely observed when target letters were consciously reported (decision-dependent signal). Fourth, we show that global pre-stimulus neural activity biases perceptual decisions for a “seen” response. Fifth and last, no evidence was obtained for the sharpening of sensory representations by top-down expectations. We discuss these findings in light of emerging models of perception and conscious report highlighting the role of expectations and stimulus relevance.

INTRODUCTION

What we perceive can be strongly influenced by top-down factors, such as our prior expectations about likely states of the world and the relevance of input for the task at hand [2,3,5,15]. According to a growing body of work, expectations, originating from past experience, can shape perception on both a neural and behavioral level [8]. When sensory input matches prior expectations, performance on tasks is higher [28,107] and neural activity is attenuated [6,108,109]. In such predictive brain frameworks, it is assumed that what we perceive consciously is strongly related to the brain's best guess about the current state of the outside world [30,39,40,110].

Indeed, the idea that our subjective experience is strongly influenced by top-down factors is supported by numerous behavioral studies that have shown beneficial effects of prior knowledge on subjective perception, such as on the accuracy [28] or speed [29,43] of detecting a stimulus. Likewise, it has been shown that prior knowledge increases the likelihood that a stimulus is consciously reported during the attentional blink [74,75,111]. Although expectations seem to affect conscious access, the neural underpinnings of these expectation-related modulations are still unclear. In a recent electroencephalographic (EEG) study, we did not find evidence that the amplitude of neural signals, as indexed by event-related potentials (ERPs), explained the effect of expectations on the likelihood of conscious report of the stimulus [111]. This finding may be explained by the fact that instead of modulating the strength of neural responses, expectations may improve the signal-to-noise ratio or “sharpness” of the representation of stimuli [6,112] by instantiating specific perceptual templates (e.g. orientation selectivity), even before stimulus presentation [113]. Because it has been shown previously that conscious perception may also be closely linked to the quality or variability of sensory representations [114,115], we here investigated whether similar enhancements of neural representations also underlie the effects that expectations have on the conscious accessibility of stimuli [111].

To this end, we used an attentional blink (AB) task in which each trial consisted of a sequence of rapidly presented letters in which one or two targets were to be detected and reported at the end of the stream (targets were marked by placeholders; **FIGURE 3.1**). Crucially, the first target stimulus (T1) could either validly, invalidly or not predict (neutral trials) the identity of the second target (T2). Then, multivariate decoding

analyses were used on magnetoencephalography (MEG) data to track the neural representations of target stimuli. Further, to explore the effects of task-relevance on stimulus-specific neural activity patterns subjects also performed an additional “oddball” task in which a similar rapid serial presentations of letters was presented but subjects were instructed to merely detect a contrast change of a stimulus that happened only on 10% of the trials (these oddball trials were not taken into account in the analyses). The combination of both tasks allowed us to trace the sensory processing of a letter stimulus in the absence of a target identity decision (because the identity of the letters was task-irrelevant). This decoding profile could then be compared to the decoding profile that we observed during the attentional blink (AB) task in which subjects did have to make a perceptual decision on the presented target stimuli (target identity was task-relevant). This approach allowed us to perform within-task decoding analyses (training and testing on the AB task using k-folding) and between-task decoding analyses (training on the oddball task and testing on the AB task). Both approaches are aimed at testing different hypotheses. Training the classifier on the oddball task will extract a relatively pure sensory signal (letter identity is task-irrelevant) and therefore between-task decoding will isolate those neural processes in the AB task related to sensory stages of information processing that may be associated with the effect of expectation on conscious report. On the other hand, within-task decoding in the AB will train and test on task-relevant stimuli that require a categorical decision and therefore may reveal decision-dependent processes that may underlie the role of expectation in conscious access. Thus, our approach allowed us to examine if expectations influence visual representations at sensory and/or decision-related stages and how this influences conscious access.

METHODS

Participants

We tested a total of 33 participants for this experiment. All participants had normal or corrected-to-normal vision. One participant was excluded because T1 identification performance was more than three standard deviations lower than the group average. Furthermore, four participants were excluded because their subjective estimates of T2 visibility were unreliable. Their indication that they perceived a target did not correspond to above chance performance in identifying the same target ($p > 0.05$ in a binomial test). Additionally, only in the MEG analyses we excluded participants for whom the number of observations in any of the relevant conditions was lower than

10, similar to Meijs et al [111]. As a result, we included 28 participants (18 females, age 22.5 ± 2.8 years) in the behavioral analyses and 19 participants (12 females, age 23.0 ± 2.8 years) in the MEG analyses.

The experiment was approved by the local ethics committee of the Radboud University (CMO Arnhem-Nijmegen; “Imaging Human Cognition”). Written informed consent was obtained from participants according to the Declaration of Helsinki. Compensation was either 36 Euros or course credit.

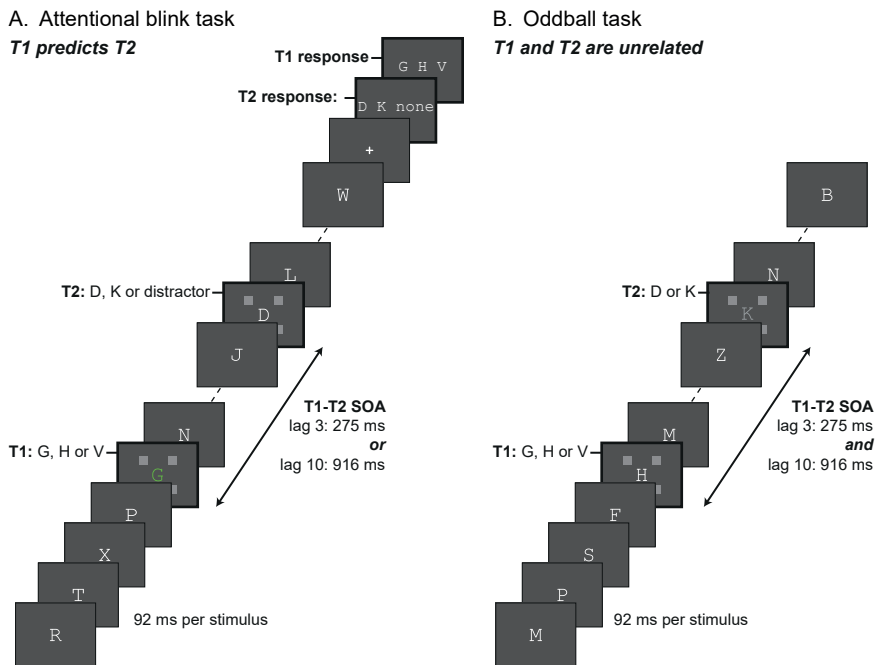


Figure 3.1: Experimental tasks. (A) The trial structure of the AB task. Each trial consisted of a sequence of rapidly presented letters in which targets were to be detected and reported at the end of the stream. Targets were marked by placeholders. The first target (T1: G, H or V) was always the fifth stimulus in the sequence. A second target (T2: D or K) was presented on 80% of trials, at varying lags. In a training session participants learned conditional probabilities between T1 and T2. One T1 stimulus was used as neutral condition and was thus followed equally often by each T2 stimulus. The other T1 stimuli predicted which T2 target was most likely to appear, thereby introducing valid and invalid expectations. **(B)** The oddball task was designed such that temporally it resembled the AB task as much as possible, with the most notable difference that on *every trial* three target stimuli were presented (T1, T2 at lag 3 and T2 at lag 10). Participants’ task was unrelated to the target identity: they had to respond to oddball stimuli that were present on 10% of trials at one of the target positions. In the example trial here, an oddball is presented at the T2 time point.

Materials

Stimuli were generated using the Psychophysics Toolbox [54] in a MATLAB (MathWorks,

Natick, MA, USA) environment. In the behavioral lab, stimuli were displayed on a 24" BENQ LED monitor (1920 x 1080 pixels; 120 Hz). A chinrest was used to control participants' distance from the screen (± 57 cm). In the MEG environment, a PROPixx projector (VPixx Technologies Inc., Saint-Bruno, Canada) located outside the magnetically shielded room projected the stimuli onto a screen approximately 80 cm in front of participants (1920 x 1080 pixels, 120 Hz). All visual input was presented on a "black" background (luminance: ± 3 cd/m²) and matched for visual luminance between the two labs.

Procedure and stimuli

The experiment consisted of two sessions that were completed within one week. The task of interest was a version of the attentional blink (AB) task [20], which participants performed in the second session while MEG was recorded. In the first session, participants performed an oddball task while MEG recordings were made and subsequently practiced the AB task outside of the MEG environment.

AB task

Participants had to detect targets within a sequence of rapidly presented distractors (92 ms per stimulus). Each stimulus in the sequence was an uppercase letter that was presented at fixation in a monospaced font ("Courier New"; letter size: $\pm 2.08^\circ$). The first target (T1: G, H or V) was presented in "green" at the fifth position of the sequence. In 80 % of trials a second target (T2: D or K) was presented as well, either at lag 3 (275 ms after T1 onset; 2/3rd of trials) or at lag 10 (917 ms after T1 onset; 1/3rd of trials). Each distractor letter (all alphabet letters excluding the targets) was presented maximally once per trial. T2 targets and distractors were presented in "white" (luminance: ± 230 cd/m²).

Crucially, the likelihood of each T2 target appearing was conditional on the identity of the T1 stimulus that was previously presented (FIGURE 3.1A). For every participant, one of the T1 stimuli (e.g., G) predicted that "D" was the most likely T2 target while another T1 (e.g., H) made the exact opposite prediction that "K" was the most likely T2 target. If a T2 was presented, the likely T2 stimulus was shown on 75 % of trials. A third T1 stimulus (e.g., V) had no predictive value (neutral condition; 20 % of trials), i.e., both T2 stimuli were equally likely to follow T1. All possible mappings of T1 and T2 were used across participants in a counterbalanced fashion, but mappings were fixed

within a given participant for the entire experiment. On 20 % of trials no T2 stimulus was presented but a random distractor letter was presented instead at either lag 3 or lag 10. Both at the T1-timepoint and the T2-timepoint of a trial, even when a T2 target was omitted, a placeholder consisting of 4 grey squares (luminance: ± 50 cd/m²; size: 0.62°; midpoint of each square centered at 2.34° horizontally and vertically from fixation) was presented around the target letter. This placeholder provided timing information, cueing participants which time points were relevant in a trial and thereby helping them decide which targets they saw during a trial.

Following a 150 ms blank period at the end of the letter sequence, participants gave their responses with the use of a (MEG-compatible) button-box, using the index, middle and ring finger of their right hand. First, they reported the T2 they had seen (3 response options: D, K or none). Subsequently they were asked to make a forced choice judgment about the T1 stimulus that was presented (3 response options: G, H or V). A long response timeout duration of 4 second was used and participants were explicitly instructed to value accuracy over response speed. The inter-trial interval was 800 - 1200ms.

Before starting the MEG recordings, participants were briefly reminded about the task instructions (given to them in a previous session, see next paragraph). Every participant completed approximately 750 trials (average: 748.7 ± 26.8 , minimum: 652; maximum: 803), with the exact number depending on the duration of MEG preparations and the number of breaks a participant needed. At the end of every block of 75 trials, participants received summary feedback about their performance and were provided with the opportunity to take a short break.

Behavioral training

At the end of the first session, every participant was behaviorally trained on the AB task. First, participants received on-screen instructions in which they were explicitly instructed about the predictive relationship between T1 and T2. Subsequently, they performed 6 blocks of 75 trials (total 300 trials) of the task. The goal of this training session was to familiarize participant with the task before the MEG session and to teach participants the predictive relationship between T1 and T2. We did not analyze the data from this training session.

Oddball task

At the start of the first session, we measured MEG while participants performed an oddball task that was highly similar to the AB task (FIGURE 3.1B). Again, every trial consisted of a sequence of letters in which letters that were targets in the AB task were presented. On every trial, a “T1” (the letter D, V, or G) was presented followed by a “T2” (the letter D or K) at both lag 3 and lag 10 and every of these targets was marked by placeholders. All stimuli, including targets, were “white”. Every combination of “T1” and “T2” stimuli was equally likely, meaning there was no predictive relationship between the targets. Importantly, participants were not made aware of the presence of the targets. Moreover, because they had not yet seen the AB task, they did not know about the existence of any specific target letters (or predictive relationship between those). The timing of the stimuli and trials was identical to that in the AB task, with the exception that at the end of the letter sequence the ITI (800 – 1200 ms) started immediately.

Participants were instructed to detect oddball stimuli that occurred on 10 % of the trials. An oddball was defined by its grey (luminance: ± 78 cd/m²) instead of white color. Once an oddball was detected, the participant had to press the index finger button on an MEG-compatible button box as quickly as possible while the task continued. Because, unbeknownst to participants, oddballs were always either a “T1” or “T2” target, they were accompanied by the same placeholders used in the AB task. To make sure participants would focus their attention on the relevant time points in the sequence, they were explicitly instructed that these squares marked the potential temporal positions of oddballs and would help them detect oddballs. Oddball trials were excluded from all analyses. Every participant completed 8 blocks of 96 trials (total 768 trials) of the task. Every block was followed by summary feedback and a short break.

Behavioral analyses: AB task

Behavioral data was preprocessed with in-house MATLAB scripts and subsequently analyzed using JASP software [55]. We focused on the effects of lag and expectation on percentage correct T2 visibility, given that the correct T1 target was reported. T2 responses were considered to be correct if a participant entered the target letter that was presented or reported not seeing a letter when none was presented on a T2 absent trials. Since expectations are undefined on T2 absent trials, these trials cannot be used

in the main statistical analyses. T2 percentage correct was used in a 2 x 2 repeated measures ANOVA with the factors expectation validity (valid, invalid, neutral) and lag (lag 3, lag 10). Subsequently, we performed post-hoc t-tests to directly compare the different levels of expectations within each of the lags.

MEG measurements and preprocessing

Whole-head magnetoencephalographic (MEG) recordings were acquired (sampling rate 1200 Hz) using a 275-channel MEG system with axial gradiometers (VSM/CTF Systems, Coquitlam, BC, Canada), which was located in a magnetically shielded room. Four channels (MLC11, MLC32, MLF62, MRF66) were disabled in all participants for technical reasons. Head position was monitored and corrected if required using three coils, placed on the nasion and on earplugs in both ears [116]. Importantly, at the start of the second MEG session an effort was made to reposition a participants' head as much as possible in the same location as during the first session by using a template head position saved in the first MEG session. In addition to the MEG, three sets of electrodes (+ground) were used to measure the electrocardiogram (ECG) and horizontal and vertical eye-movements (EOG). Finally, an Eyelink 1000 eyetracker (SR Research, Ottawa, Canada; sampling rate 1000 Hz) was used to measure pupil dilation and vertical and horizontal eye-movements.

For each session separately, we preprocessed the data with the FieldTrip toolbox for MATLAB [59]. Data were high-pass filtered at 0.01 Hz to remove slow signal drifts. Additionally, a set of notch filters was applied at 50, 100 and 150 Hz to remove line noise. Subsequently, we cut the data into epochs from -750 to 1500 ms relative to T1 stimulus onset. The data were visually inspected and trials and/or channels with artifacts were deleted (averages session 1: 7.0 % of trials, 1.2 channels; averages session 2: 5.4 % of trials, 1.1 channels). To remove noise originating from far away external sources, 3rd order gradient correction using the CTF reference sensors was applied. Independent component analysis was used to identify and remove data components related to eye blinks, eye movements or heartbeats. To get a reliable estimate of which components to delete, each of the components was correlated to the EEG and eyetracker channels. Following the independent component analysis, previously deleted channels were reconstructed using the average of neighboring channels. Finally, all trials were baseline corrected on the interval 500 ms prior to T1 onset (corresponding to -775 to -275 ms prior to T2 onset).

Decoding analyses and statistics

Prior to the decoding analyses, we applied a sliding window of 50 ms to the data, thereby smoothing the data in the temporal domain and improving the signal-to-noise ratio. Subsequently, we performed the decoding analyses by using linear discriminant analysis (LDA) decoders with the activity from all MEG-channels as features. In short, as outcome measure, the LDA decoder calculates the distance from a decision boundary on a trial-by-trial basis (full analysis details are available in [117]). This distance measure can be used as a quantitative measure of the evidence for a certain class in the decoder. In cases where a decoder was trained and tested on the same dataset, a 10-fold cross-validation procedure was implemented in which for each fold the LDA decoder was trained on 90% of the trials and tested on the remaining 10% of trials.

To be able to look at the stability of neural representations over time, all decoders were trained at one time point and then tested on all time points, resulting in a temporal generalization matrix (time range -250 to 1500 ms relative to T1; see also [118]). Cluster-based permutation tests with 1000 permutations were used to find the significant positive or negative clusters within the temporal generalization matrices of interest [60].

Using these analyses methods, we initially decoded T2 target identity at lag 3 within each of the two tasks. Further decoding analyses were done between-tasks, training the LDA decoder on one task and testing it on the other task. For the main analyses, we used the decoding that was trained on the oddball task and tested it on the AB task data (only trials where T1 was correctly identified), separately for valid and invalid trials. We did not include the neutral condition here because the number of observations in this condition was low (only presented on 20% of trials). A similar analysis was done to compare T2 identity decoding between T2 reported and T2 missed trials. As a control, we repeated all these decoding analyses for T2 targets presented at lag 10, the results of which are reported in the Supplementary material.

Next, we used the decoder trained on the oddball task and tested it on the attentional blink data (T1 correct trials) for T2-absent trials and trials with lag 10 together (in both conditions, no target stimulus was presented at lag 3). Instead of grouping trials based on T2 identity, we grouped trials based on the expected T2 stimulus so that we could investigate the neural representation of sensory expectations. In a final group

of analyses, decoders were trained and tested on the AB task (lag 3 & T1 correct trials only), allowing us to investigate the main effects of expectation validity and T2 visibility within the AB task. In addition, we did within-task T2 decoding analyses separately for T2 reported and T2 miss trials. Finally, to get a better view at the decoding results when training and testing on the same time point, we extracted the diagonal from the temporal generalization matrices of a number of analyses. We performed a paired-samples t-tests to each time-point of the data after T1 presentation (0 – 1500ms) and subsequently applied false discovery rate (FDR) correction.

RESULTS

Behavioral results: Expectations increase the likelihood a stimulus being consciously perceived

In **FIGURE 3.2A**, we show percentage correct T2 discrimination for trials on which T1 was correctly identified, separately for short and long lags (T1 accuracy T2 at lag 3: 95.63%, $sd=3.68\%$; T1 accuracy T2 at lag 10: 95.69%, $sd=3.95\%$). We observed a clear attentional blink, as indicated by a reduced performance when the time between T1 and T2 was short compared to long (lag 3: 34.57%; lag 10: 62.58%; $F_{1,27}=61.43$, $p<0.001$). Importantly, T2 stimuli were more often correctly reported when they were expected (collapsed across lag: expected: 59.44%, unexpected: 39.13%; $F_{2,54}=17.37$, $p<0.001$) and this effect of expectations was different at lag 3 than at lag 10 (interaction lag x validity: $F_{2,54}=5.97$, $p=0.005$).

To test whether valid expectations increase, or invalid expectations decrease, T2 reportability, we directly compared valid, invalid and neutral trials at each of the two lags. At lag 3, where T2 most often goes unreported, T2 discrimination performance was higher in the valid (black line; 47.67%) than in the neutral (dark grey line; 30.35%) and invalid (light grey line; 25.68%) conditions (valid-invalid: $t_{27}=5.49$, $p<0.001$; valid-neutral: $t_{27}=4.60$, $p<0.001$). In the invalid condition performance was lowest, although not significantly lower than performance in the neutral condition (neutral-invalid: $t_{27}=1.79$, $p=0.084$). At the long lag, performance was also higher in the valid than invalid condition (valid 71.21%; invalid 52.57%; valid-invalid: $t_{27}=3.88$, $p<0.001$) and performance in the neutral condition was in between (neutral: 63.97%; valid-neutral: $t_{27}=2.04$, $p=0.051$; neutral-invalid: $t_{27}=3.18$, $p=0.004$). Taken together, these results replicate earlier findings that conscious access can be affected by expectations about stimulus likelihood and that these effects are largest at short lags, inside the

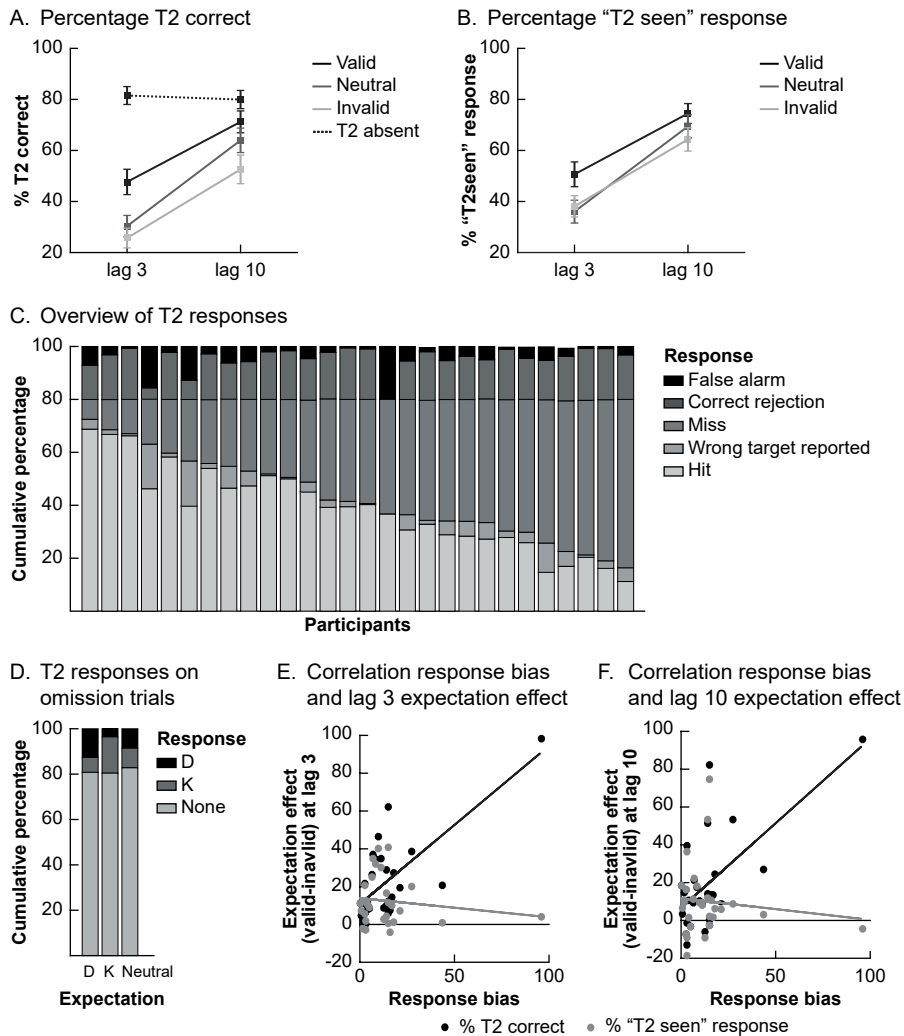


Figure 3.2: Behavioral results. (A) Percentage correct T2 (given that T1 was correctly identified) for each of the lags and conditions. Expectation validity significantly modulated the percentage of T2s that was consciously seen and this effect was different at the two lags. At lag 3, the expectation effect was mainly driven by the valid condition, while at lag 10 both invalid and valid expectations affected T2 detection as compared to neutral trials. (B) Results from a control analysis, using the percentage of trials on which participants reported seeing a T2 target (given that T1 was correctly identified) as dependent variable. Overall, the pattern of results was similar to that in (A). Error bars denote SEM. (C) A general overview of the response pattern over all trials for both lags together. The percentage of misses (no target reported when one was presented) was high in all participants. When no target was presented, participants usually correctly reported not seeing one. (D) The response pattern shows participants had a response bias: if they reported a target on an T2-absent trial, they more often guessed the expected target letter (when there was an expectation). Panels (E) and (F) depict the relationship between the response bias (% of T2-absent trials with a predictive T1 where the expected target was reported) and the behavioral effect of expectation validity (valid-invalid) at lag 3 (E) and lag 10 (F). In both, the relationship is shown for both outcome measures that were used to quantify behavioral performance: percentage correct (black) and percentage of "T2 seen" responses (regardless of the exact letter participants entered; grey).

AB [111]. Additionally, at this short lag performance was significantly improved when expectations were valid compared to neutral, while the negative effect of invalid expectations was relatively small. Indeed, the benefit of valid expectations was larger at lag 3 than at lag 10 (valid-neutral lag 3 vs. valid-neutral lag 10: $t_{27}=3.67$, $p=0.001$). This suggests that at short lags, when the stimulus is most often missed, the expectation effect we observe is most likely mainly driven by a higher likelihood of conscious access for valid expectations.

Overall, false-alarms were infrequent, on ~19% of T2-absent trials subjects reported a target (FIGURE 3.2CD), correct rejections for T2-absent trials: lag 3: 81.52%; lag 10: 79.94%) [77]. Further, if participants reported seeing a target on T2-absent trials, they most often reported the expected target ($t_{27}=2.52$, $p=0.018$), demonstrating that participants had indeed learned the predictive relationship between T1 and T2 and used this information for their decisions. Therefore, it may be that the behavioral effect of expectations we observed on T2-present trials is (partly) affected by this response bias. Indeed, a correlation between participants' response bias, as measured by the percentage of trials they reported the expected letter vs. another letter on T2-absent trials, and their expectation effect, as measured by the difference in performance on valid minus invalid trials, was observed (FIGURE 3.2EF, black dots; lag 3: spearman $r=0.49$, $p=0.008$; lag 10: spearman $r=0.34$, $p=0.078$). However, a control analysis where we considered the percentage of "T2 seen" responses (classifying a response as "seen" regardless of whether this letter was correct) showed comparable effects of expectation (validity effect: $F_{2,54}=6.06$, $p=0.004$; interaction lag x validity: $F_{2,54}=5.13$, $p=0.009$, FIGURE 3.2B). Note that, in this analysis, the dependent variable (seen vs. miss) is orthogonal to participants' expectations (D vs. K) and this outcome measure did not correlate with the response bias participants had on T2-absent trials (FIGURE 3.2EF; grey dots; lag 3: spearman $r=-0.09$, $p=0.649$; lag 10: spearman $r=-0.11$, $p=0.594$), which suggests that the observed expectation effects are present over and above the effects of response bias (note that the pattern of correlations was similar when the outlier subject with a response bias of nearly 100% was removed from the data, lag 3: spearman $r=-0.05$, $p=0.822$; lag 10: spearman $r=-0.05$, $p=0.807$).

Decoding of T2 identity is modulated by stimulus relevance

Multivariate analyses can provide insights in the dynamics of category-specific brain responses, as different neural and cognitive processes may be reflected in dissimilar

patterns of temporal generalization [118]. First, we tested the role of stimulus relevance in the processing of the letter stimuli. To do so, we focused on decoding the identity of T2 targets (D vs. K) at lag 3, irrespective of conscious report or expectations. **FIGURE 3.3AD** shows the temporal generalization profiles for T2 identity within each of the two tasks, using cross-validation procedures and corrected for multiple comparisons (see **METHODS** for details). As can be seen, T2 target identity could reliably be decoded from MEG data in both tasks, as reflected in a typically observed diagonal decoding pattern starting ~ 100 ms after T2 onset (Oddball task: $p=0.022$; AB task: $p<0.001$, permutation tests). This diagonal decoding pattern likely reflects a rapid sequence of distinct neural processes evolving over time [117–119].

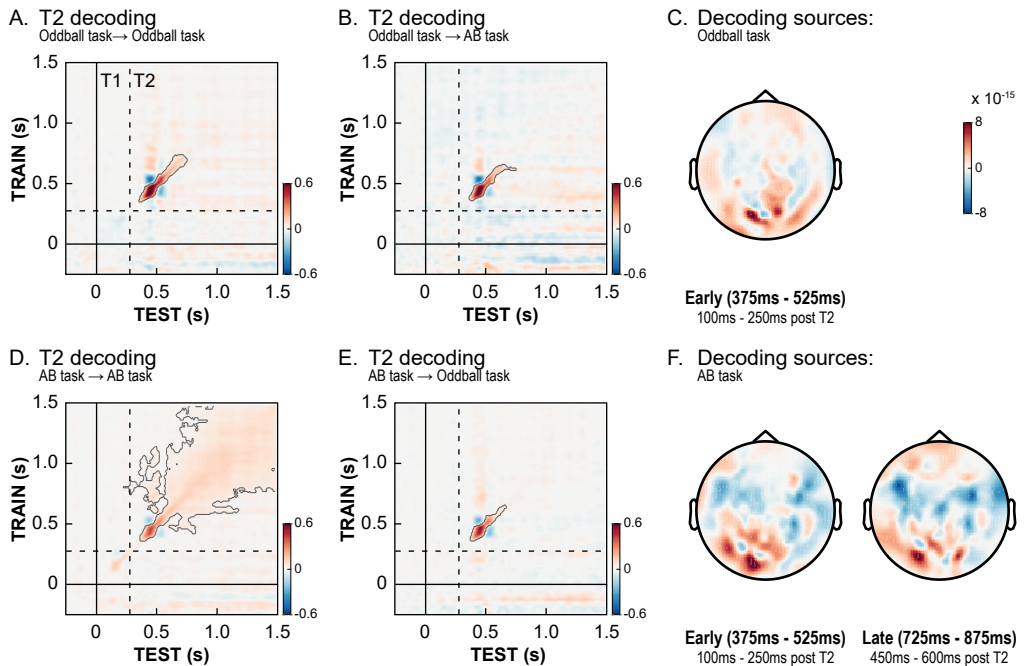


Figure 3.3: Decoding of T2 identity at lag 3. Temporal generalization matrices showing that T2 target identity (D vs K) could be reliably decoded from MEG data within both the oddball task (**A**) and the AB task (**D**). Decoding was longer and more widespread in the latter, suggesting neural representations were more stable over time. (**B,E**) Between-task decoding analyses (trained and tested on separate datasets) showed significant clusters along the diagonal, indicating that early neural sensory representations of T2 identity were highly similar between tasks. Timing of all panels is relative to T1 onset. Dashed lines indicate T2 stimulus onset. Contours of significant clusters are shown with a grey line (corrected for multiple comparisons). Decoding of T2 targets at lag 10 is shown in **SUPPLEMENTARY FIGURE 3.1**. To visualize the pattern of activity used by the decoding algorithm, we display topographic plots (**C,F**) showing the planar gradient ERF-difference between the two T2 targets. On the left, we show the topography in an early window (100-250ms post T2; including the peak in decoding) for both tasks. For the AB task only, a topographic plot in a later window (450-600ms post T2), in which there was no significant decoding for the oddball task, is shown.

Further, as expected, in the AB task only, decoding performance was not only strong on the diagonal of the temporal generalization matrix, but it's profile showed a mixture of a diagonal and a square shaped pattern and decoding was much more extended in time. Thus, we observed stimulus-specific stable and sustained activation patterns for task-relevant compared to task-irrelevant stimuli. Next, we trained decoders on each of the tasks and subsequently tested them on the data of the other task (between-task decoding) to test for neural similarities and differences between the tasks. This approach solely resulted in significant clusters along the diagonal of the temporal generalization matrices for both analyses, further highlighting a shared neural coding of target identity from ~100 to 375 ms post T2 onset (Oddball task → AB task: $p=0.013$, **FIGURE 3.3B**; AB task → Oddball task: $p=0.024$, **FIGURE 3.3E**). This conclusion is also supported by similar decoding results for T2 stimuli presented at lag 10 (**SUPPLEMENTARY FIGURE 3.1**). Overall, these results show that when stimuli are task-relevant (require a perceptual decision) additional processing steps can be picked up by the classifier and T2 identity is processed in more stable and temporally extended neural representations, as compared to when stimuli are task-irrelevant [117,118,120–122].

Conscious access is related to late and stable stimulus-specific neural representations

Next, we tested how category-specific neural representations relate to conscious report, or the absence thereof, in the AB task. We did so by comparing T2 decoding (D vs. K) at lag 3 for seen T2s vs. missed T2s. To rule out any effect caused by the identity of T1, only trials where T1 was correctly reported were incorporated in these analyses. **FIGURE 3.4** shows both between-task analyses (Oddball task → AB task) and within-task analyses (AB task → AB task; done separately per condition). In line with prevalent theoretical models of consciousness, early sensory processing of the stimulus was independent of conscious report, as reflected by similar decoding profiles along the diagonal of the temporal generalization matrix for reported and missed T2's (**FIGURE 3.4A-C**, Oddball task → AB task) [19,23,102,123]. **FIGURE 3.4** also shows within-task analyses (AB task → AB task) in which T2 identity was decoded based on solely T2 reported (**FIGURE 3.4D**) or T2 missed trials (**FIGURE 3.4E**). In both analyses, significant diagonal decoding of T2 identity was observed (T2 reported trials: $p<0.001$, miss trials: $p=0.042$), however, when T2 was reported, this decoding pattern was stronger, longer-lasting (**FIGURE 3.4F**) and had a broader spatial profile (more square-shaped instead of diagonal, T2 reported: 100 ms to 1200 ms post T2-onset; T2 missed: 100 ms to 400 ms post T2-

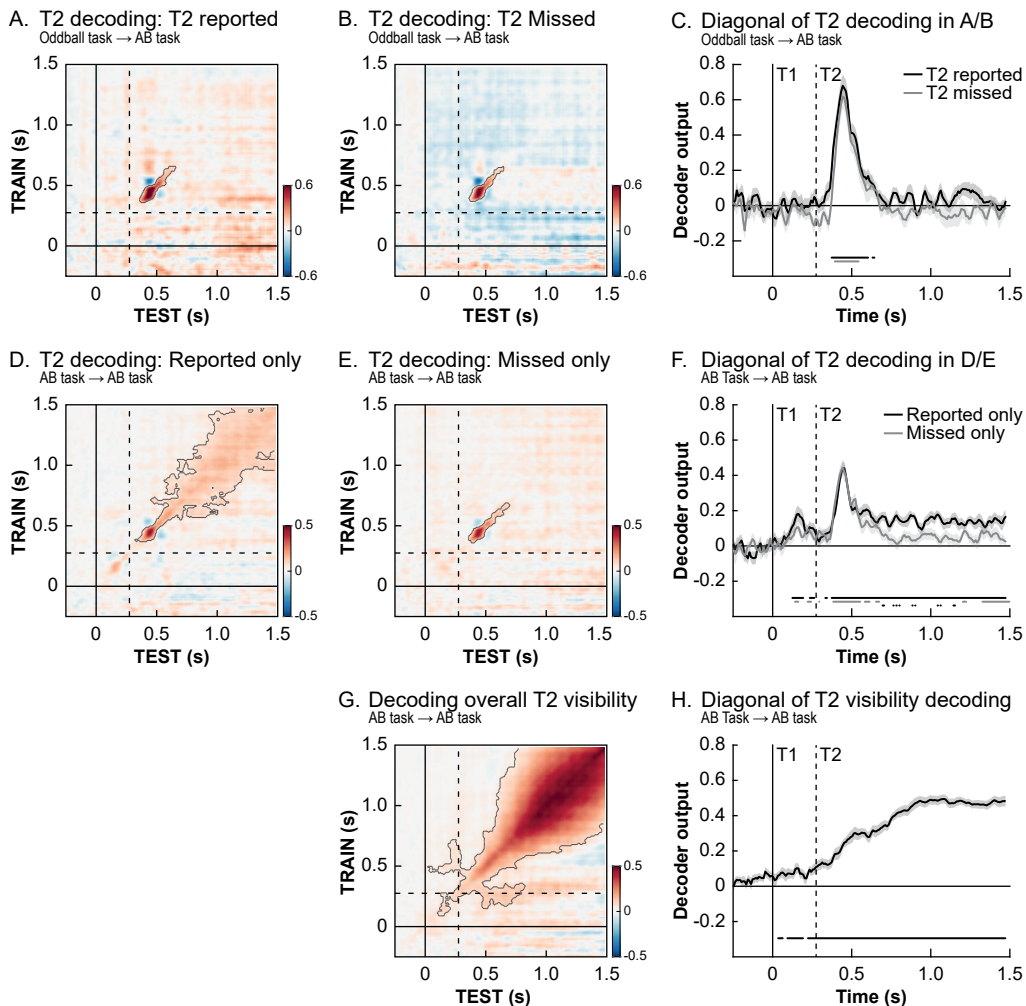


Figure 3.4: Conscious stimulus detection and its neural representation. Temporal generalization matrices for the decoding of T2 identity at lag 3, trained on the oddball task, for trials where T2 was consciously perceived (**A**) and trials where it was missed (**B**). T2 targets could be reliably decoded in both conditions and no significant differences between both conditions was observed along the diagonal (**C**). Within condition decoding analyses trained on the AB task show T2 decoding was more long-lasting (**F**) and stable on trials where T2 was consciously perceived (**D**) compared to where it was missed (**E**). (**G,H**) A direct investigation of the effect of T2 visibility (seen vs miss, irrespective of stimulus identity) showed widespread significant decoding starting before stimulus onset, indicating an early and long-lasting main effect of stimulus visibility. Timing of all panels is relative to T1 onset and in every panel only trials where T1 was correctly identified were used. Dashed lines indicate the onset of a T2 stimulus at lag 3. Significant clusters in the temporal generalization matrices are contoured by a grey line (corrected for multiple comparisons). Significant time-periods along the diagonals (**C,F,H**) are marked by horizontal lines at the bottom of the panel. Solid lines indicate significant time periods for conditions while the difference between conditions is indicated by a dotted line (**F**).

onset). This suggests that, compared to missed T2's, the signal was broadcast higher up the cortical hierarchy (for similar results see [117]). Please note that the apparent pre-stimulus decoding in [FIGURE 3.4F](#) should not necessarily be interpreted as evidence for decoding of stimulus specific T2 activations before the stimulus is presented. This effect most likely relates to decoding of T1, which in the AB task was correlated to T2.

Finally, we aimed to decode reported vs. missed T2's in the AB task, irrespective of the specific stimulus that was presented (classifier labels: reported vs. missed, target letter is not relevant, [FIGURE 3.4G](#), $p < 0.001$). Several interesting aspects of the obtained temporal generalization profile can be noted. First, along the diagonal, decoding was significant already before T2 presentation ([FIGURE 3.4H](#)). Neural activity before stimulus presentation thus predicted later stimulus report, in line with recent reports highlighting pre-stimulus fluctuations in neural activity that correlate with stimulus visibility (e.g. [123-125]). This analysis only included T1-correct trials, this pre-stimulus effect cannot easily be explained in terms of differences in T1-performance. Nevertheless, we cannot rule out that T1-related processes affect the results here, since others have shown that T1 processing may be different between blink and no-blink trials [127,128]. The second interesting aspect about this overall visibility effects was that the square shaped pattern was much stronger and clearer. This stimulus independent profile is likely related to an amplified P3-like component on T2 reported trials compared to T2 missed trials, which is picked up by the classifier algorithm (see [SUPPLEMENTARY FIGURE 3.4](#) for ERF results) [61,111].

Valid expectations do not enhance early sensory representations of stimulus identity

In our final set of decoding analyses, we aimed to explore the link between neural recordings, expectations and conscious report. Behaviorally, we observed that valid expectations increased the chance that a stimulus was subsequently reported and to test whether this may relate to the sharpening of sensory representation of T2 identity, decoding performance was compared between trials where participants had a valid ([FIGURE 3.5A](#)) or invalid ([FIGURE 3.5B](#)) expectation at lag 3 (only T1-correct trials were included). Importantly, only between-task analyses were performed in which the decoder was trained on the oddball task (no T1-T2 predictability) to prevent the decoding being biased by the predictive relationship between T1 and T2 in the first place, which is inherently present in our AB-task and difficult to circumvent. Training on the AB task would also incorporate undesirable factors (e.g. conscious access: valid

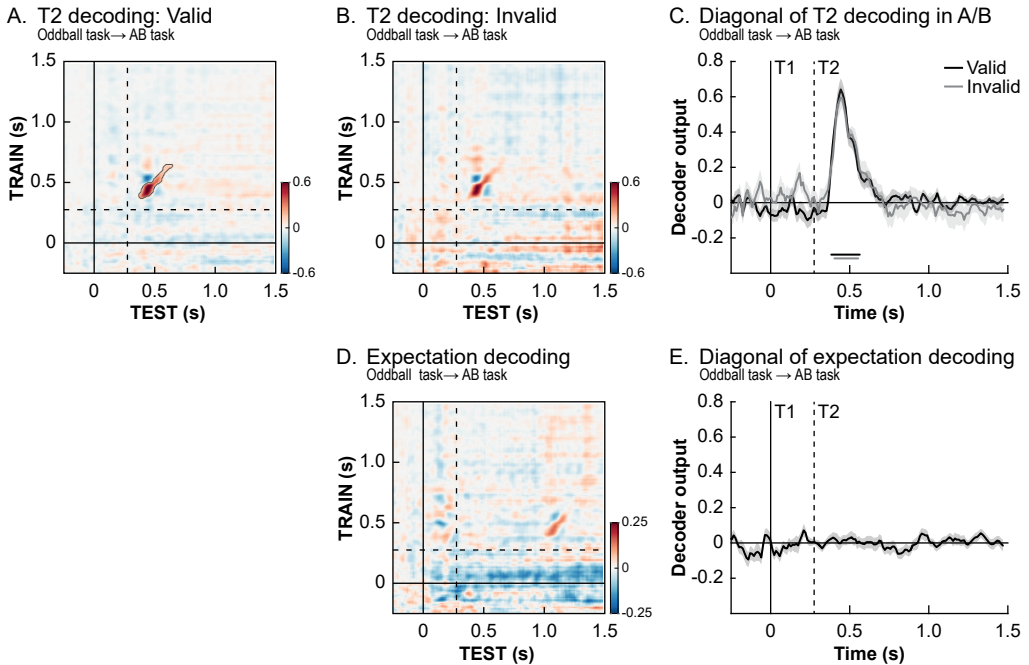


Figure 3.5: Expectation validity. Temporal generalization matrices for the decoding of T2 identity at lag 3 for valid trials (A) and invalid trials (B) separately. Only trials where T1 was correctly identified were used. Panel (C) shows that no significant differences (valid – invalid) between the two conditions were observed, indicating T2 identity could be decoded equally well on valid and invalid trials (significant time period indicated by horizontal lines at bottom of panels), thus suggesting early sensory representations were not affected by the validity of prior expectations. (D,E) It was not possible to decode participant’s T2 stimulus expectations based on a between-task decoder trained on the oddball task data and tested on the attentional blink data. Timing is relative to T1 onset, and the dashed lines indicate the onset of T2 stimuli at lag 3. We did not find any evidence for the existence of neural stimulus templates.

trials are reported more often than invalid trials) into the classifier training. No effect of prior knowledge on the sharpening of sensory representations was observed (T2 identity decoding was similar for valid and invalid trials, **FIGURE 3.5C**). Because several studies have shown that the effects of expectation may interact with (or even depend on) attention [129,130] we aimed to rule out that the absence of expectation-related sharpening could be attributed to a lack of attention for targets presented at lag 3 (i.e., during the “blink” period). However, similar results were obtained for stimuli presented at lag 10 and hence out of the blink period (**SUPPLEMENTARY FIGURE 3.3**).

The lack of any decoding difference between valid and invalid trials makes it unlikely that participants’ expectations induced neural “perceptual templates” in the current task design. On valid trials, when participants’ expectation about the upcoming T2

was equal to the T2 that was presented, a decoder trained to distinguish T2 identity in the oddball task should have been able to pick up the representation of the target (expectation) before target presentation (target and expectation matched). Similarly, on invalid trials one would have expected inverse decoding, because participant's expectation was opposite to the T2 stimulus that was later presented (target and expectation mismatched; see for example Kok et al [113]). To fully rule out that we may have missed expectation-induced sharpening we performed a final analysis in which we grouped trials based on the expected T2 stimulus instead of the actually presented T2 stimulus. Only T2-absent and lag 10 trials were used, so no T2 stimulus was ever presented at lag 3 in order to prevent any contamination of the results by actual stimuli. But again we did not observe the instigation of expectation-related perceptual templates (all $p > 0.221$, **FIGURE 3.5DE**).

DISCUSSION

In this report, we investigated the relationship between expectations and conscious access and the role of stimulus relevance on stimulus-specific neural representations. Using an attentional blink (AB) paradigm in which the identity of T1 predicted the likelihood of the T2, we found that T2 stimuli confirming T1-induced expectations are more likely to be consciously reported than T2 stimuli that violate these expectations. This is in line with studies that have shown that the speed [29,43,44] or accuracy [28,111] of stimulus detection is strongly influenced by expectations. Interestingly, by including neutral expectation trials we showed that the effect of expectations was mainly driven by a performance benefit for valid expectations, at least in the AB interval (at short lags) (see also [43]). Thus, at the time that targets are often missed, conscious access becomes more likely when expectations are confirmed (valid) compared to when they were absent or invalidated (**FIGURE 3.2A**).

We have also presented a series of multivariate decoding analyses, in which we first showed that it was possible to reliably decode the identity of a target letter from MEG data, both when the identity of this letter was task-irrelevant in an oddball task and when it was task-relevant and actively searched for in an AB task. Furthermore, the combination of both within- and between-task decoding analyses revealed that stimulus-specific neural representations underlying the early stages of information processing (on-diagonal generalization) were highly similar for both tasks and independent of stimulus relevance and the outcome of the perceptual decision

(whether T2's were reported or missed in the AB task) [102,117]. Similar early decoding profiles were observed for 1) task-irrelevant letters in the oddball task, 2) task-relevant, but missed T2's in the AB task, and 3) task-relevant and reported T2's in the AB task. These results suggest that early stimulus-specific neural processing stages are not modulated by either attention to the stimulus dimension (stimulus relevance) nor conscious access to the target identity (report), in line with previous findings [123]. It further shows that the representation of target letter identity is not "broadcast" to higher levels of the cortical hierarchy when this feature is task-irrelevant [119]. Usually, it is assumed that compared to nonconscious perception, conscious report is related to widespread neural activity (up to frontal areas) and more stable neural representations [16,19,23,114]. Indeed we observed that within the AB task, when targets were task-relevant, attended and reported, T2 identity could be decoded for a longer time frames and was more stable, reflected in a square shaped decoding profile [102,117,119,131–133]. Additionally, we observed a nonspecific (regardless of T2 identity or expectations) difference between T2 reported and missed trials, of which the late sustained part presumably corresponds to previously observed P3 modulations related to conscious report [61,67,111] and task-relevance [71,134]. We also observed that fluctuations in neural activity (or overall "brain state") before T2 presentation influenced whether or not the stimulus was later reported, as has been recently demonstrated by several others using at threshold or backward masking tasks [124–126,135–138]. These fluctuations may reflect differences in T1 processing [128].

Next, in order to explain the behavioral effects of expectations on conscious access, we looked at the neural representations of T2 identity conditioned on the validity of single-trial expectations. We hypothesized that the early neural representation of validly expected stimuli would be "sharpened", and that this improvement in improved stimulus representation would relate to enhanced stimulus report [6,107]. However, we did not observe this effect. Moreover, no evidence for expectation-induced pre-stimulus perceptual templates was observed, although these effects have been recently reported in a different task context [113]. One possible explanation for this discrepancy in findings is that it may take longer than 300ms to encode T1 and translate the T1-prediction into a sensory representation. Nevertheless, in our study, the influence of expectations on conscious report was not reflected in modulations of sensory representations prior to T2 presentation or at early stages in the processing hierarchy (for a similar conclusion in a different paradigm see also [139]). It could be

that expectations can influence perception at multiple levels of the cortical hierarchy, with different processes being affected depending on the type of target stimuli or expectations being involved. In our AB paradigm, participants learned conditional relationships between two letter stimuli. Consequently, the expectations in our task were most likely represented at a higher level of the cortical hierarchy (i.e., semantic) as compared to the rather low-level perceptual expectations in earlier studies [6,113]. As a consequence, the early sensory representation of the targets that we pick up based on the oddball task may not have been optimal for detecting expectation-based modulations of stimulus processing. Therefore, the “format” in which the expectations were represented may simply have been different (e.g. verbally or motor code) than that of the incoming visual information and hence interactive effects may not have been visible at the sensory level [139,140]. Future research would benefit from using tasks that can capture the full stimulus processing hierarchy involved in processing the relevant stimuli within the experimental task (sensory, semantic, decision).

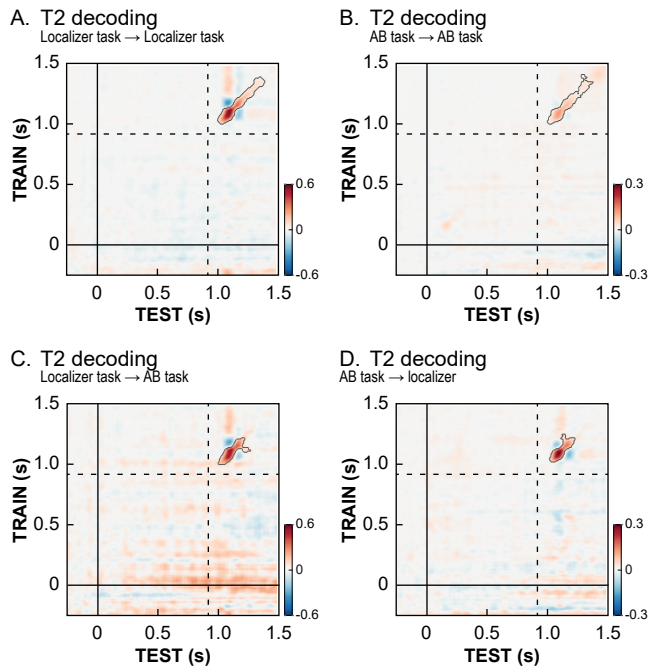
In line with a recent study [139], we did not find any reliable decoding effects that we could relate to the behavioral benefits of expectations, suggesting it may be quite subtle and hard to detect (see also [111]). One possible implementation of expectations would be a modulation of the onset of components related to conscious perception [42]. Other potentially interesting signatures are the power and/or phase of alpha oscillations because a number of studies have shown that low pre-stimulus alpha power is predictive of subsequent conscious stimulus perception [124–126,135,141–143]. Furthermore, it has been shown that expectations lead to changes in pre-stimulus alpha in a way that predicts stimulus visibility [144]. Future studies are required to unravel the effect of expectations on conscious access in the frequency domain.

A number of conclusions can be drawn from this study. First, we have shown that the neural representation of letters presented during rapid serial visual presentation can be reliably decoded from neural activity as measured by MEG. While early, sensory representations of letters could be decoded regardless of the behavioral state of the subject, later and more stable multivariate activity patterns were dependent on top-down modulations by task relevance, attention and conscious report. Second, we have shown that valid expectations enhance the likelihood of visual target detection but we did not find evidence that this was due to increased “sharpness” of the relevant sensory representations.

SUPPLEMENTARY MATERIAL

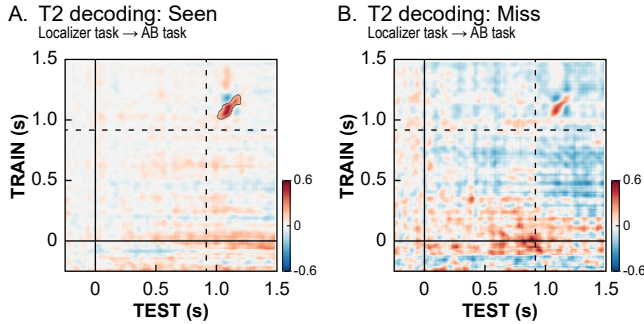
Decoding results for T2 stimuli presented at lag 10

T2 identity for targets presented at lag 10 could be reliably decoded (**SUPPLEMENTARY FIGURE 3.1**) from MEG data both within-task (Oddball task: $p=0.011$; AB task: $p=0.025$) and between-task (Oddball task \rightarrow AB task: $p=0.025$; AB task \rightarrow Oddball task: $p=0.020$). This further validates the conclusion that the initial representation of T2 targets was similar between the tasks.



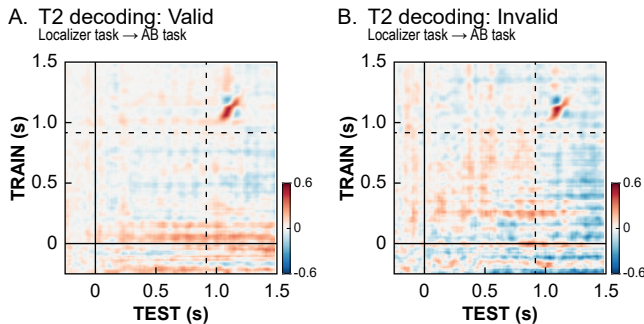
Supplementary Figure 3.1: Decoding of T2 targets at lag 10. We show the temporal generalization matrices for the decoding of T2 identity with different types of decoding analyses. Timing of all panels is relative to T1 onset. The dotted lines indicate the onset of T2 stimulus at lag 10. On the left, we show the decoding performance for cross-validated decoders that were trained and tested within either the oddball task (**A**) or attentional blink (**B**) task. On the right, we show the results for decoding analyses where the decoder was trained and tested on separate datasets. In (**C**) the decoder was trained on the oddball task and tested on the AB task, while in (**D**) the AB task was used for training and the oddball task was used for testing. The identity of T2 could be decoded in all the analyses. Contours of significant clusters are marked with a grey line (corrected for multiple comparisons).

The representation of T2 stimuli presented at lag 10 was similar for detected and missed T2s (**SUPPLEMENTARY FIGURE 3.2**; all $p>0.0785$). We should note that we did find effects for this contrast at lag 3, but these were on the off diagonal starting quite late, making it likely they simply “fall outside” the time axis in the temporal generalization matrix.



Supplementary Figure 3.2: Effects of conscious stimulus detection on the neural representation of T2s at lag 10. Temporal generalization matrices for the decoding of T2 identity at lag 10 for trials where the T2 was consciously perceived (A) and trials in which it was missed (B). Only trials where T1 was correctly identified were used. The timing of all panels is relative to T1 onset. Dashed lines indicate the onset of a T2 stimulus at lag 10.

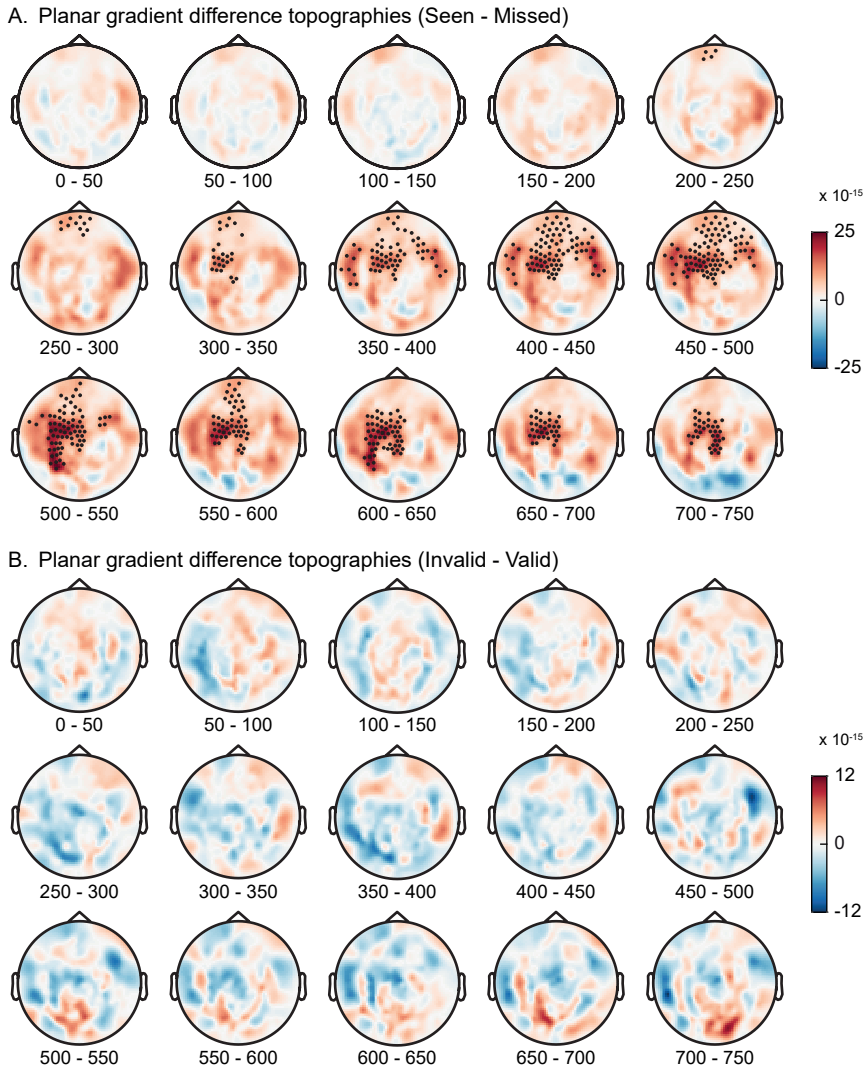
In line with the analyses performed for T2s presented at lag 3, no significant differences were observed between trials on which the stimulus was validly or invalidly expected (SUPPLEMENTARY FIGURE 3.3; all $p > 0.582$). These results show it is unlikely that the absence of an expectation validity effect at lag 3 was due to a lack of attention in early time-windows where attention “blinks” [129].



Supplementary Figure 3.3: Effects of expectation validity on the neural representation of T2s at lag 10. We show the temporal generalization matrices for the decoding of T2 identity at lag 10 for valid trials (A) and invalid trials (B) separately. Only trials where T1 was correctly identified were used. There were no significant differences (valid – invalid) between the two conditions. Timing of all panels is relative to T1 onset. Dashed lines indicate the onset of a T2 stimulus at lag 10.

ERF Methods

In order to be able to make a more direct comparison with our earlier study [111], we replicated the ERP analyses of main effects by performing two of ERF analyses in which we look at the effects of expectation validity and T2 visibility. Only lag 3 trials on which subject identified the correct T1 stimulus were used. Prior to the ERF analyses, the data was low-pass filtered at 40 Hz.



Supplementary Figure 3.4: ERF results. Topographic maps showing the difference between reported and missed T2s (**A**) and between valid and invalid expectations (**B**) over time (0=T2 onset) using planar gradients. Cluster-based permutation tests were used to isolate significant events. On every head map, channels belonging to a cluster that show significant difference for at least 50% of the time window are highlighted.

For each of the conditions in a comparison, the data were transformed to planar gradients to allow averaging over participants. Because this transformation is sensitive to the number of trials in a condition, we repeated this step 100 times for each condition, matching the number of trials between conditions that would be compared (by taking a random selection of trials from the largest condition) and averaging over these repetitions. Subsequently, we used cluster-based permutation tests [60] on

the time-window 750 ms after T2 onset to isolate significant ERF events relating to expectation validity (valid, invalid; irrespective of T2 visibility) or T2 visibility (seen, missed; irrespective of validity).

ERF Results

SUPPLEMENTARY FIGURE 3.4 shows the results of the ERF analyses. In line with previous results [61,111], conscious detection of a target was reflected in a long-lasting P3-like component ($p < 0.001$). However, we did not find an ERF effect analogous to the early negative posterior effect in the EEG results. Additionally, the comparison of the ERF on valid and invalid trials did not yield any significant differences (all $p > 0.332$). Future research is required to explain the discrepancies in results between these studies, that may result from differences in task design or general differences between EEG and MEG.

4

CUE PREDICTABILITY DOES NOT MODULATE BOTTOM-UP ATTENTIONAL CAPTURE

Adapted from:

Cue predictability does not modulate bottom-up attentional capture. *Royal Society Open Science*, 5, 180524.

ABSTRACT

Attention can be involuntarily captured by physically salient stimuli, a phenomenon known as bottom-up attention. Typically, these salient stimuli occur unpredictably in time and space. Therefore, in a series of three behavioral experiments, we investigated the extent to which such bottom-up attentional capture is a function of one's prior expectations. In the context of an exogenous cueing task, we systematically manipulated participants' spatial (Experiment 1) or temporal (Experiment 2 and 3) expectations about an uninformative cue, and examined the amount of attentional capture by the cue. We anticipated larger attentional capture for unexpected compared to expected cues. However, while we observed attentional capture, we did not find any evidence for a modulation of attentional capture by prior expectation. This suggests that bottom-up attentional capture does not appear modulated by the degree to which the cue is expected or surprising.

INTRODUCTION

When performing tasks in our everyday lives, we constantly have to battle potential distraction by task-irrelevant inputs. Even when we want to stay focused on the task at hand, it can be difficult to ignore other, often more salient, stimuli that capture our attention. Historically, there has been considerable debate on whether attentional capture is purely stimulus-driven [145,146], or also depends on top-down goals [147]. More recently it has been suggested that recent trial history [148] and associations with reward [149,150] may also modulate attentional capture. This implies that the presence and amount of attentional capture may be a complex function of both stimulus and internal variables [151].

Stimulus expectation is another factor that may modulate bottom-up capture. One of the studies providing evidence for this was carried out by Folk and Remington [152]. In a spatial cueing paradigm, they manipulated the frequency of salient but uninformative (i.e., not predictive of the target) cues. Their results indicated that these cues captured attention only when they were unlikely, regardless of the top-down task set participants were using (but see [153] for an alternative interpretation). In another study the proportion of distractors was systematically varied over blocks [154]. The distractors interfered more with target processing when they were presented in a block with fewer distractors, suggesting they captured attention more when they were more surprising. Similarly, it has been observed that novel stimuli are most potent in capturing attention [155] and also most robustly modulate the neural response in a macaque's V1 [156]. Taken together, these studies support the hypothesis that attentional capture by task-irrelevant stimuli may be modulated by perceptual expectations, and most notably by the violation of these expectations. This can be interpreted as evidence that surprising stimuli are more salient and therefore more attention-grabbing [157,158]. Additional evidence supporting this idea comes from studies on mismatch detection, in which it has been shown that unexpected deviant stimuli lead to larger mismatch responses in the EEG-signal and seem to subjectively "pop out" [159–161].

Besides influencing the amount of attentional capture by distracting stimuli, there is evidence that prior information about these cues can help participants to voluntarily diminish distraction [162,163]. For example, it has been shown that attentional capture

by unlikely distractors can be attenuated when the search task promotes suppressing features similar to those of the distractors stimuli [153]. Another study showed that it is easier to ignore regular sequences than irregular sequences [164]. Whether reducing the amount of distraction is caused by the inhibition of attentional capture, or by rapid disengagement at a later stage is still debated [148,165]. An electrophysiological study by Kiss and colleagues suggested that bottom-up capture can be inhibited, but that this only happens when task demands (i.e., timing) require it [166].

One interpretation of the empirical evidence above is that surprising stimuli are more salient and therefore more attention-grabbing [157,158]. Predictive coding theories have suggested that processing unexpected events requires more resources [6,8]. One may conceptualize bottom-up attention as a way of redistributing resources, for example towards processing unexpected events. This is in line with findings that bottom-up attention increases contrast sensitivity [167]. Nevertheless, many models based on predictive coding have actually suggested that regularity and predictability may attract attention. The idea behind this is that predictable inputs are more strongly weighted because they are more reliable. A number of studies have provided support for this idea [37], suggesting that regularities automatically attract attention (but see e.g. [168]).

It thus seems that the link between expectation and bottom-up attention is still far from clear [8,169]. Attention may be either drawn to surprising stimuli, or to regularly occurring ones. Therefore, we performed a series of three experiments using an exogenous cueing task [34], in which we explicitly looked at this relationship by manipulating participants' expectations about an otherwise uninformative (i.e., unrelated to target) cue stimulus. Specifically, we investigated to which extent prior expectations about the cue modulated bottom-up attentional capture. Based on the evidence listed earlier, we anticipated that unexpected cue stimuli will attract more attention and therefore result in larger cue-target validity effects (i.e., performance difference between validly and invalidly cued trials), whereas expected cue stimuli are followed by strongly reduced or even absent validity effects. To preview, in contrast to this hypothesis, we observed attentional capture in all experiments, but no direct modulation by prior knowledge about the cue stimulus in any experiment.

EXPERIMENT 1: DO SPATIAL EXPECTATIONS AFFECT BOTTOM-UP ATTENTION?

METHODS

Participants

We tested 120 participants in Experiment 1. This number was based on power analysis for a between-subjects design with 40 participants per group, with 80% power to detect medium-sized effects. All participants had normal or corrected-to-normal vision. We excluded two participants whose task performance was markedly (more than 3 standard deviations) worse than that of other subjects. As a result, we included 118 participants (87 females, age 22.7 ± 5.0 years). The experiment was approved by the local ethics committee of the Radboud University (CMO Arnhem-Nijmegen; 2014/288 “Imaging Human Cognition”). Written informed consent was obtained from all participants according to the Declaration of Helsinki.

Materials

Stimuli were presented using the Psychophysics Toolbox [54] within MATLAB (MathWorks, Natick, MA, USA), generated by a Dell T3500 Workstation and displayed on a 24” BENQ LED monitor (1920 x 1080 pixels; 60Hz; screen size 53.1cm x 29.9 cm). All presented stimuli were “black” (RGB: [0 0 0]; ± 0.3 cd/m²) on a grey (RGB: [150 150 150]; ± 103.9 cd/m²) background. A chinrest was used to control the distance participants were seated from the monitor (± 57 cm). Participants responded by means of two button boxes.

Procedure and stimuli

Participants performed an adjusted version of the exogenous cueing task ([34], **FIGURE 4.1A**). First, a cue (2° circular outline) was presented for 50 ms either 5 degrees above or below fixation. After the cue, a target was presented centered on either same (valid trials) or the opposite (invalid trials) screen location. Cue location and target location were unrelated, meaning that both target locations were equally likely throughout the experiment, regardless of where the cue was presented. Targets were small (0.48° wide and 0.60° high) arrows pointing either leftward or rightward. The participants’ task was to report the direction the arrow was pointing in (leftward or rightward) by pressing a button with either their left or right index finger, while maintaining fixation throughout the experiment.

In Experiment 1 we manipulated the likelihood that the cue would appear either above or below fixation. In two groups of participants the cue was most likely to appear respectively above or below fixation. Consequently, participants in these two groups (N=78) encountered both trials where the cue was in the expected location (80%) as well as trials where it was in the unexpected location (20%). In a third group of participants (N=40) both cue locations were equally likely, resulting in these participants experiencing only neutral trials.

The stimulus-onset asynchrony (SOA) between the cue and target was set to 117 ms. The target remained onscreen until a response was given or until 1000 ms after target onset had passed. Trials were separated by a variable inter-trial interval of 750-1500 ms. Participants responded to the arrows by pressing a button on a button box with their left or right index fingers, respectively for leftward and rightward pointing arrows.

In total the experiment lasted approximately one hour. Before starting with the main experimental task, participants received on-screen instructions and performed one practice block of 80 trials. Participants were told that the cue was irrelevant (i.e. non-predictive of the target) and could be ignored. During the practice block participants received on-screen feedback (response correct or incorrect) on a trial-by-trials basis. Subsequently, participants performed 800 trials of the main task divided over 10 blocks. After each block, participants received feedback about their task performance (overall percentage correct and number of late responses) and subsequently there was a short 20s break.

Behavioral analysis

Trials where the participants' reaction time exceeded three standard deviations from the participants' mean or was below 200 ms were discarded. Furthermore, trials on which no (relevant) response was given within 1 s from target onset were excluded from analyses. The remainder of the trials (98.3%) was exported to JASP [55] in order to perform statistical analyses. We analyzed the data using a combination of both frequentist statistics and their Bayesian equivalents [56,57].

For the statistical analyses of Experiment 1 we performed a 2 x 2 repeated-measures ANOVA with the factors Validity (valid, invalid) and Expectation (expected, unexpected) for the reaction times. This analysis only considered participants for whom the cue

was more likely to appear in one of the locations (N=78). Participants in the neutral condition were included as a reference group, to be able to interpret possible performance differences as either gains or losses in performance with respect to an expectation-neutral context. In two independent samples t-tests we compared the validity effect size between the neutral condition and the expected or unexpected condition in the test group participants. In addition to frequentist analyses, we also computed Bayes Factors for all relevant comparisons. As we were specifically interested in the interaction between expectations and the bottom-up validity effect, the Bayes Factor (BF) for a model with only the main effects was compared to the BF for a model with the main effects and the interaction. The ratio of the BF values then quantifies the evidence for including the interaction term in the model, and hence can be interpreted as evidence for or against the existence of an interaction between the two experimental factors. BF ratios will converge either to infinity when a model including the interaction explains the data better, or to zero when it explains the data worse than a model with only main effects. If the ratio is close to one, this indicates that both models are equally likely and that there is not enough evidence for either conclusion. We use the conventions from Jeffreys [58] to interpret the evidence in our Bayesian analyses.

RESULTS AND DISCUSSION

We examined whether expectations about the spatial location of a cue modulate the ensuing attentional capture by this cue. In **FIGURE 4.1B** we plot the reaction time results for Experiment 1 for trials on which participants gave the correct answer. Higher reaction times on invalid than valid trials indicate there was attentional capture by the cue (RT difference=29.31 ms, $F_{1,77}=216.53$, $p<0.001$, $\eta^2=0.738$). Importantly, however, this validity effect was not modulated by the spatial expectation participants had about the cue ($F_{1,77}=0.753$, $p=0.388$). The evidence against the existence of this modulation is moderate (BF₀₁=4.40). While spatial expectations did not modulate attentional capture, there was an overall reaction time benefit for trials where the cue was in the expected location compared to the unexpected location (RT difference=8.87 ms, $F_{1,77}=84.56$, $p<0.001$, $\eta^2=0.523$). Further post-hoc analyses showed that the validity effect in the neutral control group was not significantly different from that on either expected ($t_{116}=-0.96$, $p=0.338$) or unexpected ($t_{116}=-1.15$, $p=0.252$) trials for participants who had expectations about cues. Accuracy was close to ceiling level (98.21±1.21% correct) and is further reported and analyzed in the **SUPPLEMENTARY MATERIAL**.

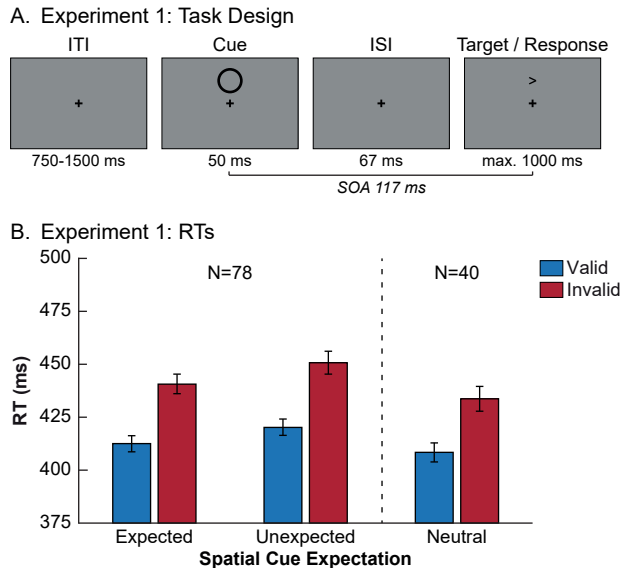


Figure 4.1: Task design and behavioral results of Experiment 1. (A) Trial structure of the exogenous cueing task used in Experiment 1. In every trial a cue (circular annulus) was presented for 50 ms, either above or below fixation. After an inter stimulus interval (ISI) of 67 ms (SOA 117 ms), a target (arrow) was presented in either the same (valid trials) or opposite (invalid trials) location. We manipulated spatial cue expectation by varying the likelihood the cue would appear in either location. In one group of participants the cue appeared equally often above and below fixation. In the two other groups the cue was more likely (80%) to appear in one of the locations. Target location was counterbalanced and unrelated to the cue location. Participants' task was to report the direction the arrow was pointing in. **(B)** Reaction time results for Experiment 1. Only trials on which the correct answer was given were used for the analysis. On the left we show results for participants that expected the cue either above or below fixation (N=78), meaning that it was sometimes presented in the expected location and sometimes in the unexpected location. Participants were faster on valid than on invalid trials, regardless of their spatial expectations about the cue. For reference, we present the results for participants in the neutral group (N=40) on the right. Error bars represent SEM.

Together, these results suggest that our manipulation of spatial expectations modulated overall behavioral speed and accuracy but did not result in a modulation of bottom-up attention capture. While this suggests that expectations do not interact with bottom-up attention, there are possible alternative explanations. Most notably, because targets and cues are often presented in the same location, and with only a short (117 ms) and predictable time delay between the two, suppressing processing at the cue location would possibly hamper target processing. Moreover, since the cue was temporally predictive of the target (fixed SOA), its presentation was informative about target onset and hence attending the cue at all times may have been useful for target perception. These considerations inspired Experiment 2, in which we examined whether cues still elicit bottom-up attention when they are temporally predictable but no longer predictive of when targets occur.

EXPERIMENT 2: THE EFFECT OF TEMPORAL EXPECTATIONS ON BOTTOM-UP ATTENTION

METHODS

Participants

We tested 67 participants in Experiment 2. Compared to Experiment 1, fewer participants were tested because a within-subjects design was used. All participants had normal or corrected-to-normal vision. We excluded four participants from Experiment 2 because button presses were not recorded properly. As a result, we included 63 participants (50 females, age 23.0 ± 3.5 years) in the final analysis. The experiment was approved by the local ethics committee of the Radboud University (CMO Arnhem-Nijmegen; 2014/288 “Imaging Human Cognition”). Written informed consent was obtained from all participants according to the Declaration of Helsinki.

Materials

We used the same materials as in Experiment 1, with as only exception that participants now used the computer keyboard (DELL KB522) to respond.

Procedure

Instead of manipulating spatial expectations about the cue, in Experiment 2 we manipulated temporal expectations. To remove standard temporal links between cues and targets, the task consisted of continuous blocks (duration 5 min) during which cues and targets were presented (FIGURE 4.2A). Temporal expectations were manipulated by varying the regularity of cue presentation onset between blocks. In the regular (expected condition) blocks, the cue would be presented every second (1 Hz presentation rate). In the irregular blocks (unexpected condition) cues were presented quasi-randomly every 0.5 – 1.5 seconds, with a uniform distribution over all possible intervals between two consecutive cues. This results in the participant having less (precise) information about the cue onset and the cue-target SOA. Targets lasted 200 ms and were presented quasi-randomly every 1.0 to 2.5 seconds. Importantly, the onsets of cues and targets were determined completely independently of each other. Because the SOA between two targets was longer than between cues, not every cue was followed by a target. Participants used the ‘z’ and ‘/’ (slash) buttons on the keyboard to respond to leftward and rightward pointing arrow targets respectively.

Participants completed 8 blocks of the task, switching between regular and irregular conditions after every two blocks. The initial condition was counterbalanced over participants. After every block there was a 20s break. If a condition switch occurred, this was explicitly mentioned on-screen at the end of the break. Participants were instructed that the cue was irrelevant to their task and could be ignored. Participants did not receive any feedback during the task. Together with the instructions and the practice block the experiment lasted approximately one hour.

A. Experiment 2: Task design




Figure 4.2: Task design of Experiment 2. (A) Example stimulus sequence for Experiment 2. The task consisted of continuous blocks (duration 300s) during which cues and targets were presented. In some blocks presentation was regular and hence expectations about cue onset were strong. In other blocks cue presentation was less expectable. Again, participants’ task was to report the direction the target arrow was pointing in.

Behavioral analysis

During the experiment, the task was presented to participants in continuous streams, without a clearly discernible trial structure. For analyses purposes, trials were defined post-hoc by isolating at all target presentations (on average 1371 trials per participant). Data preprocessing was similar to Experiment 1 and resulted in the inclusion of 96.0% of trials. Subsequently, for each target presentation, we identified the temporal distance between the target and the closest preceding cue stimulus. This generated a total of 90 bins, spanning a cue-target SOA between 0 and 1500 ms, in 17 ms steps (SOAs above 1000 ms were infrequent, <10% of trials). The first button press after target presentation was recorded as the participants’ response. If no response was made within a 1000 ms interval following the target, the trial was classified as a “miss” trial.

Because the number of trials per condition was low (due to the large amount of possible SOA bins; on average 20.6 trials per SOA bin in Experiment 2 and 21.0 trials per SOA bin in Experiment 3), we smoothed the data in the temporal domain by applying a sliding window over all SOA bins of interest (SOAs below 500 ms). In this procedure, for every relevant combination of conditions, the data we ascribe to an SOA are computed as a

weighted average from that SOA (SOA_0) and the two SOAs immediately preceding and following it ($0.1*SOA_{-2}+0.2*SOA_{-1}+0.4*SOA_0+0.2*SOA_{+1}+0.1*SOA_{+2}$). Trials with an SOA between 0 ms and 50 ms were excluded from the analyses because on those trials cue and target presentation overlapped in time.

Based on the smoothed data, we computed the overall validity effect for each SOA. We then identified two SOAs of interest based on these data: (1) an early maximally facilitatory validity effect (SOA with a maximally positive overall validity effect); and (2) a late maximally inhibitory validity effect (SOA with a maximally negative overall validity effect). We then averaged the (non-smoothed) data around these SOAs of interest, using the selected SOA and the two SOA's preceding or following it, to create an early and a late window of interest. Subsequently, we tested for an interaction between expectation and the bottom-up validity effect by performing a 2 x 2 repeated-measures ANOVA with the factors Validity (valid, invalid) and Expectation (expected, unexpected) in each of these windows. In addition, like in Experiment 1 the Bayesian equivalent of the ANOVA was performed. Please note that since cue onsets and target onsets were determined independently, not every participant had observations for every condition at every SOA. In case data was missing in one of the conditions of an analysis, we excluded the respective participant from that analysis.

RESULTS AND DISCUSSION

We temporally smoothed the data over the different SOA bins (FIGURE 4.3A). Subsequently, we identified an early and late window where bottom-up attention effects were most prominent (see also BEHAVIORAL ANALYSES). The windows can be interpreted as resulting from initial bottom-up attentional capture (early window), followed by inhibition of return (late window; [33]). In these windows, we tested for modulations by temporal expectations (FIGURE 4.3B and FIGURE 4.3C).

Like in Experiment 1, we observed strong evidence for bottom-up attentional capture as indexed by the validity effect (early window: RT difference=10.54 ms, $F_{1,60}=24.77$, $p<0.001$, $\eta^2=0.292$; late window: RT difference=-11.18 ms, $F_{1,59}=15.44$, $p<0.001$, $\eta^2=0.207$). Note that this effect is to be anticipated, since we chose our windows based on the validity effect size. Again, temporal expectations did not modulate bottom-up attentional capture, as shown by the absence of an influence of cue onset regularity on the validity effect (early window: $F_{1,60}=0.02$, $p=0.883$; late window: $F_{1,59}=2.85$,

$p=0.097$). The evidence against such modulations was moderate for the early window ($BF_{01} = 5.14$). For the late window, likely related to inhibition of return, there was only anecdotal evidence ($BF_{01} = 2.27$), suggesting the study's power was not sufficient to make any strong claims about effects in this time window. Unlike Experiment 1, there was no significant main effect of expectations on reaction time (early window: $F_{1,60}=0.45$, $p=0.504$; late window: $F_{1,59}=0.07$, $p=0.793$). Overall participants' task performance was close to ceiling ($94.00\pm 3.30\%$).

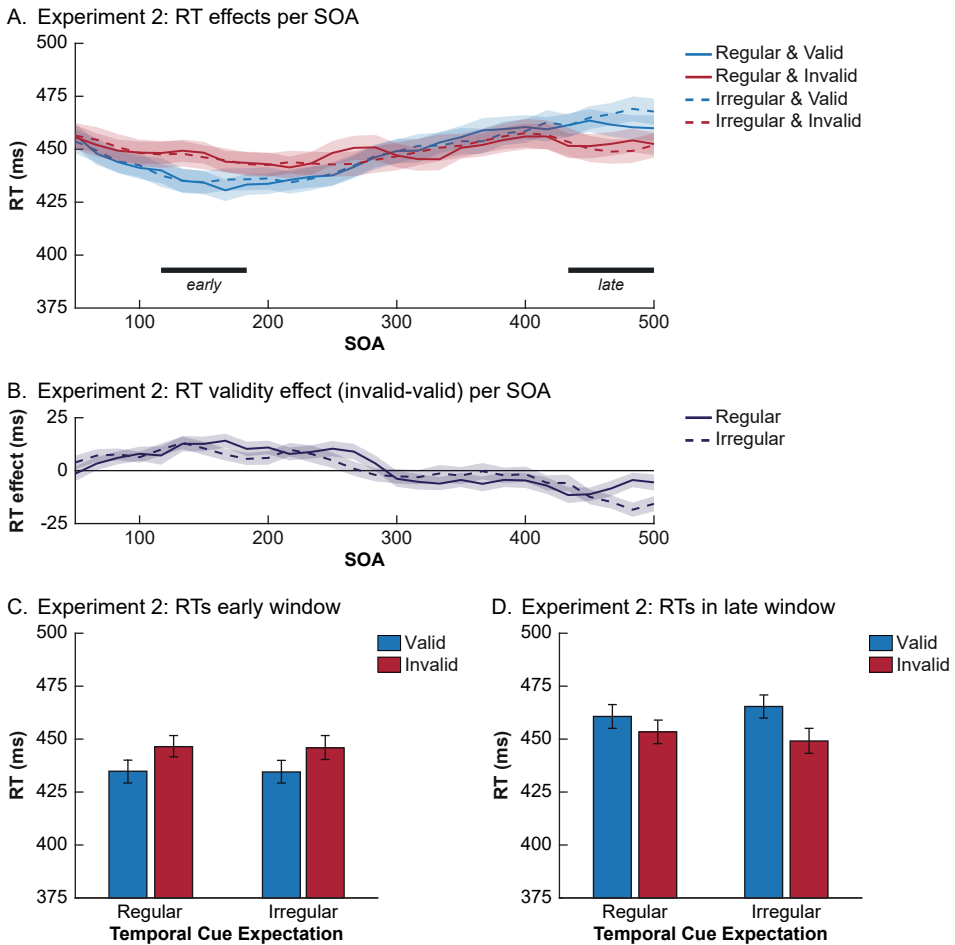


Figure 4.3: Behavioral results of Experiment 2. (A) The reaction time results and validity effect (invalid-valid) (temporally smoothed) for SOA bins between 50 ms and 500 ms are shown for each of the relevant conditions. Below the graph, the time-periods that were isolated as windows (around 150 ms and 467 ms) for further analyses are marked. (B) The validity effect (invalid-valid) results (temporally smoothed) after regularly and irregularly presented cues for SOA bins between 50 ms and 500 ms. The average reaction times per condition from the windows in (A) are presented in panel (C) for the early window and panel (D) for the late window. In both windows there was a significant validity effect that was not affected by temporal cue expectations. Error bars represent SEM.

Cues in the regular condition were temporally predictable but spatially unpredictable, i.e. they could equally likely occur above or below fixation. It is possible that cues need to be both temporally and spatially predictable for attentional capture to be reduced. This is what we set out to test in Experiment 3, in which we replicated Experiment 2 while keeping the cue location constant.

EXPERIMENT 3: THE ROLE OF TEMPORAL EXPECTATIONS IN BOTTOM-UP ATTENTION FOR SPATIALLY PREDICTABLE STIMULI

METHODS

Participants

We tested 61 participants in Experiment 3. All participants had normal or corrected-to-normal vision. Two participants were excluded because their performance markedly (more than 3 standard deviations) worse than that of other subjects. In addition, one participant was excluded because of a failure to respond on more than 20% of the trials. In the end, 58 participants were included in the analyses (42 females, age 22.7 ± 3.5 years). The experiment was approved by the local ethics committee of the Radboud University (CMO Arnhem-Nijmegen; 2014/288 “Imaging Human Cognition”) and written informed consent was obtained from all participants according to the Declaration of Helsinki.


Materials

The materials used were identical to those in Experiment 2.

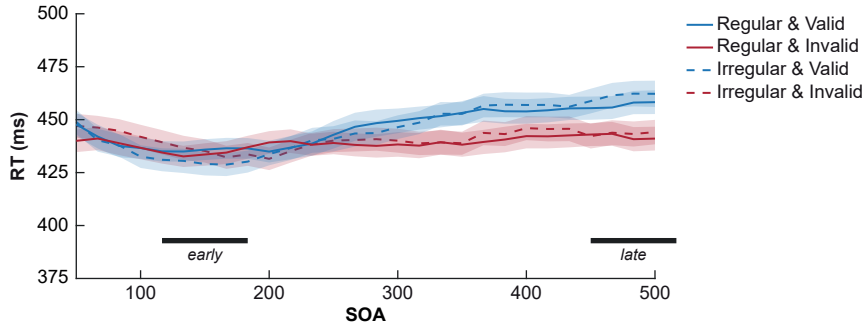
Procedure

The design of Experiment 3 was largely similar to that of Experiment 2. The most notable difference is that now for each participant the cue location was kept constant throughout the task (FIGURE 4.4A). This allowed us to investigate effects of temporal expectations in a context where the distractor is spatially fully predictable. The cue location was counterbalanced across participants. To limit possible carry-over effects, we only switched between the regular and irregular conditions halfway through the experiment. In addition, a short practice block was included at the start of each of the conditions to get participants used to the change in task structure.


A. Experiment 3: Task design




B. Experiment 3: RT effects per SOA



C. Experiment 3: RT validity effect (invalid-valid) per SOA



D. Experiment 3: RTs early window



E. Experiment 3: RTs late window

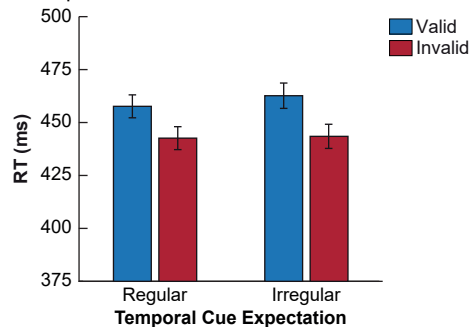


Figure 4.4: Task design and behavioral results of Experiment 3. (A) Example stimulus sequence for Experiment 3. The task that was used was highly similar to that in Experiment 2. The main difference was that for each participant the cue was now consistently presented in one location, making it completely spatially expected. (B) The reaction time results (temporally smoothed) for SOA bins between 50 and 500 ms are shown for each of the relevant conditions. (C) The validity effect (invalid-valid) results (temporally smoothed) after regularly and irregularly presented cues for SOA bins between 50 ms and 500 ms. At the bottom, the time-periods that were isolated as windows (around 150 ms and 483 ms) for further analyses are marked. In the early window (D) there was no significant bottom-up attentional capture, regardless of temporal expectations. The late window (E) did show a significant effect, but this inhibition of return effect was not modulated by temporal cue expectations. Error bars represent SEM.

Behavioral analysis

Data preprocessing (97.0% of trials included) and statistical analyses were identical to Experiment 2.

RESULTS AND DISCUSSION

In this final experiment, we created a condition in which cues were both spatially and temporally fully predictable. Within each participant the cue now always appeared in a single location and thus was always spatially expected. After defining the windows of interest (similar analysis pipeline as Experiment 2; **FIGURE 4.4B**), we again tested whether temporal predictability modulated the bottom-up attention effects (**FIGURE 4.4D** and **FIGURE 4.4E**). As in the previous experiments, no such modulations were found (early window: $F_{1,55}=1.40$, $p=0.241$; late window: $F_{1,57}=0.96$, $p=0.331$). The evidence against these interactions was moderate for both windows (early window: $BF_{01}=3.07$; late window $BF_{01}=3.91$). Similar to Experiment 2, there were no main effects of temporal cue expectations on reaction times (early window: $F_{1,55}=0.09$, $p=0.763$; late window: $F_{1,57}=1.21$, $p=0.276$). Surprisingly, while there was significant inhibition of return (late window validity effect: RT difference=-17.11ms, $F_{1,57}=30.38$, $p<0.001$), there was no overall attentional capture in this experiment (early window validity effect: $F_{1,55}=0.09$, $p=0.763$) and the evidence against the existence of such an effect was moderately strong ($BF_{01}=3.37$). As in the previous experiments, participants' overall task performance was near-ceiling ($95.50\pm 2.72\%$).

COMPARISONS BETWEEN EXPERIMENTS

Because the experiments differed markedly in the expectations participants had about cues and the context in which those cues were presented, we compared the validity effect sizes between experiments. Therefore, we performed two post-hoc analyses in which we directly compared reaction time validity effects between experiments. To ensure maximal comparability, for Experiment 2 and 3 we take the validity effect at an SOA of 117 ms (after smoothing), which is the SOA that was used in Experiment 1. First, we compared Experiment 1 (only neutral trials; **FIGURE 4.5A**) and Experiment 2 (**FIGURE 4.5B**). In terms of spatial expectations both are comparable (cue 50% in each location), but the experiments differ strongly in the temporal context in which stimuli are presented. Most notably, in Experiment 1 the SOA was fixed at 117 ms, while in Experiment 2 it was variable and unpredictable. Therefore, in Experiment 1 the cue was a good temporal predictor of target onset. A comparison of the validity effects by

means of an independent samples t-test shows a significantly smaller validity effect in Experiment 2 than in Experiment 1 ($t_{101}=4.47$, $p<.001$, $d=.904$). This is possibly explained by that fact that in Experiment 2 there was (1) less information about the onset of the target and (2) a lower likelihood (not after every cue) a target would appear.

Second, the absence of an initial bottom-up capture effect in Experiment 3 (FIGURE 4.5C) could be potentially explained by the perfect predictability of cue locations in the experiment. Because in contrast to Experiment 3 cue location was unpredictable in Experiment 2, the comparison between those experiments can be used to test this hypothesis. An independent samples t-test showed that there was no significant difference between the validity effects ($t_{119}=1.182$, $p=0.240$). Therefore, we did not find evidence that the bottom-up validity effect is modulated by spatial expectations about the cue. This conceptually replicates our findings of Experiment 1. However, it should be noted that there is only anecdotal evidence ($BF_{01}=2.75$) that the validity effects of Experiment 2 and 3 were of equal size and these results should thus be interpreted with caution.




Figure 4.5: Comparing overall validity effects between experiments. For each of the experiments, we display the average reaction times on valid and invalid trials for the selections of trials that were used to compare validity effects between the experiments. **(A)** Only the neutral condition of Experiment 1 (control group) was used. The SOA in this experiment was fixed at 117 ms. For experiment 2 **(B)** and 3 **(C)** we only used the data we obtained for the SOA of 117 ms after smoothing the SOA time courses. Error bars represent SEM.

DISCUSSION

In a series of three behavioral experiments we investigated whether bottom-up attentional capture is modulated by stimulus expectations. We did not observe empirical support for this hypothesis. On the contrary, in all experiments we found moderately strong evidence that the bottom-up validity effects were of comparable size when cue stimuli were expected, compared to when they were not (or less) expected.

The fact that participants were overall faster on expected compared to unexpected trials in Experiment 1 suggests that participants did form prior expectations, which had a sizeable influence on behavior. Based on studies showing that stimuli embedded in regular streams are better detected [37], we anticipated to see a similar main effect of expectation for the temporal paradigm in Experiment 2 and 3. Contrary to our predictions, we did not observe this effect, even though regular and irregular blocks were visibly different and every switch between conditions was explicitly marked in the participants' instructions. Other studies have used similar paradigms with streams of stimulation to investigate effects of temporal expectations and did find effects of temporal regularity on subsequent behavioral performance [164,170]. Nevertheless, temporal expectations about cues do not appear to influence the processing of subsequently presented target stimuli in situations where targets are salient and uncoupled from the cues.

Assuming that our manipulations of expectations actually instantiated priors in our subjects, the absence of any significant interaction between expectations and bottom-up attention in our experiments is surprising, because it contradicts the hypothesis that unexpected or surprising events capture attention (more) [155,157,158]. Interestingly, two recent studies observed that distractor predictability can modulate the amount of attentional capture in a visual search task, where target and distractor are presented concurrently [171,172]. The most notable difference in task design between their study and ours, is that in their paradigm distractors and targets were presented simultaneously on different locations of the screen, resulting in direct attentional competition between the stimuli. This competition can be biased by predictability. In contrast, in our studies we examine the consequences of presenting a salient cue stimulus in isolation on subsequent visual processing at that location. We find that the attention grabbing properties of such a cue are not modulated by predictability. In line with our findings, a recent study by Southwell and colleagues has also suggested that regular and random streams are equally salient [168]. Moreover, our results are in line with the stimulus-driven account of bottom-up attention [146,173], in which it is assumed that the initial capture of attention is automatic and independent from top-down factors. However, it must be noted that even within the stimulus-driven account there would have been room for expectations to suppress the effects of distracting inputs (i.e., the uninformative cues) at later processing stages [148,165].

One conceivable alternative explanation for the absence of effects of expectation on attentional capture is that the tasks we used were too simple. Participant's performance was close to ceiling (>90%) in all experiments. As a result, ignoring or suppressing the uninformative cues may not have been required in order to perform well. Indeed, recent studies [166,174] showed that attentional capture was only suppressed if task requirements were such that capture would interfere with target processing. Moreover, a study in macaques showed that modulations of V1 attention responses was larger for tasks that were more difficult [156].

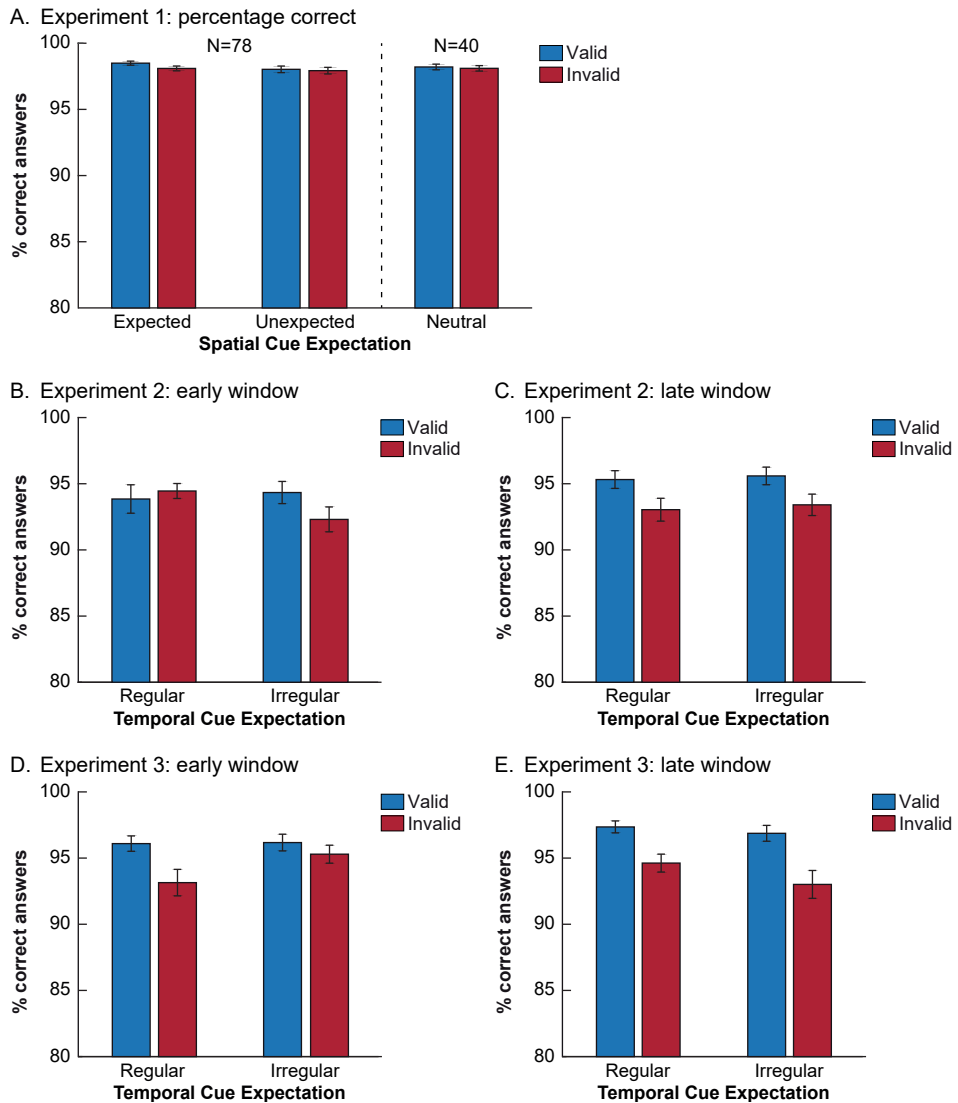
Furthermore, a recent study has suggested that only fully spatially predictable distractors may be suppressed at an early processing stage [175]. This hints at the possibility that our manipulation of expectations, especially in Experiment 1, was not potent enough to influence bottom-up capture. It is possible that with a stronger manipulation of expectations (i.e., making cues even more likely in one condition and less likely in the other) we would have observed a modulation of capture. Still, this cannot fully explain the absence of an effect in Experiment 1: even in Experiment 3 when cue stimuli were perfectly predictable in terms of timing, location and visual characteristics, no significant modulation of the validity effect was found.

It is noteworthy that ignoring or suppressing the uninformative cues in our tasks may generally not have been a useful strategy. It is conceivable that participants did not inhibit the cue location at any point in a trial, because a target would often (50%) be presented in the same location with only a short time-delay. As a consequence, ignoring one location systematically would be detrimental to target detection. Moreover, the fixed SOA in Experiment 1 resulted in cue onset being perfectly predictive of target onset time. Hence, paying attention to an informative cue, instead of ignoring it, was actually a viable strategy. Consequently, in all experiments the tasks we used may have had factors that made participants attend cues (instead of ignore them). In addition, cues and targets were both defined as abrupt onset stimuli, meaning cue features to some extent overlapped with the searched-for target feature [147]. It is conceivable that this overlap caused attention to be captured regardless of experimental condition. As a result, potential effects may have been obscured because task set did not optimally support suppression of uninformative cues [153].

We observed a significant difference in the amount of attentional capture between

Experiment 1 and Experiment 2. While it is difficult to directly compare both experiments because they differed on several dimensions, a likely candidate explanation for this difference is the fact that in Experiment 2 there was more (temporal) uncertainty about the onset of the targets, as well as an overall lower likelihood of target appearance. It is conceivable that participants deploy a different strategy in Experiment 2 compared to Experiment 1, in which they focus more on the targets and less on the cues (because those are less/not informative), which in turn leads to the cue having less influence on subsequent target processing. Future studies are required to test this idea.

In conclusion, we did not find evidence for modulations of bottom-up capture by spatial or temporal expectations about the cue. We therefore conclude that, at least in the exogenous cueing tasks we used, bottom-up attentional capture does not seem to be altered by prior knowledge about the location or time point of the distracting inputs. This calls into question perceptual surprise as an explanation for bottom-up attention. Future research may use more difficult tasks in which the relationship between targets and distractors can be more carefully controlled. In addition, we believe electrophysiological studies could possibly disentangle the effects of expectation and attention and precisely point at their interactions with high temporal resolution.

SUPPLEMENTARY MATERIAL: PERCENTAGE CORRECT RESULTS

Supplementary Figure 4.1: Percentage correct results. (A) Percentage correct responses in Experiment 1. On the left we show results for participants that expected the cue either above or below fixation (N=78), meaning that it was sometimes presented in the expected location and sometimes in the unexpected location. For reference, we present the results for participants in the neutral group (N=40) on the right. Below, we present the percentage correct results for Experiment 2 (BC) and Experiment 3 (DE), separately for the early and late time windows. Error bars represent SEM.

Experiment 1

Participants' performance was higher on expected than unexpected trials (SUPPLEMENTARY FIGURE 4.1A; $F_{1,77}=5.02$, $p=0.028$, $\eta^2=0.061$). However, there was no

validity effect ($F_{1,77}=2.10$, $p=0.152$), regardless of spatial expectations about the cue ($F_{1,77}=0.56$, $p=0.458$).

Experiment 2

We used the same time-windows as for the reaction time analyses. Percentage correct was not significantly affected by temporal expectations (early window: $F_{1,60}=1.22$, $p=0.275$; late window: $F_{1,61}=0.28$, $p=0.596$). Cue validity did impact on task performance, but only in the late time window (early window: $F_{1,60}=1.10$, $p=0.299$; late window: $F_{1,61}=12.93$, $p<0.001$, $\eta^2=0.175$). Moreover, the validity effect size was not significantly modulated by temporal expectations about the cue in either of the windows (early window: $F_{1,60}=3.29$, $p=0.075$; late window: $F_{1,61}=0.03$, $p=0.873$).

Experiment 3

We again used the same time-windows for the accuracy analyses as for the reaction time analyses. Expectations had a significant effect on participants' performance in the early window ($F_{1,55}=4.25$, $p=0.044$, $\eta^2=0.072$) but not in the late time window ($F_{1,57}=2.31$, $p=0.134$). In both windows, the validity of the cue significantly affected performance (early window: $F_{1,55}=5.36$, $p=0.024$, $\eta^2=0.089$; late window: $F_{1,57}=23.32$, $p<0.001$, $\eta^2=0.290$). Importantly though, in neither of the windows this validity effect was modulated by temporal expectations (early window: $F_{1,55}=2.19$, $p=0.145$; late window: $F_{1,57}=0.80$, $p=0.374$).

5

GENERAL DISCUSSION

In this thesis, I aimed to gain important insight in how perceptual expectations relate to conscious awareness and bottom-up attention. In particular, I wanted to answer the following questions:

1. Can expectations boost an otherwise unseen stimulus into conscious awareness?
2. Are expectations automatic brain computations, taking place irrespective of conscious access?
3. How do expectations affect the neural processing of visual stimuli and how does this relate to conscious access?
4. To what extent is bottom-up attentional capture determined by perceptual expectations?

Below, I will summarize the empirical findings reported in the previous chapters and interpret them within the context of different theories that postulate perception is an active process in which the brain tries to make sense of the outside world by forming prior expectations about it [8].

EXPECTATIONS BOOST CONSCIOUS ACCESS

In **CHAPTER 2** and **CHAPTER 3**, I reported a series of experiments in which the attentional blink paradigm was employed to investigate the influence of top-down expectations on conscious perception. In the attentional blink paradigm, the second of two targets is often missed when it is presented shortly after a first target that was detected [20]. I tested whether making the initial target (T1) predictive of the second target (T2) would help or harm participants overcome this limitation in conscious access. All those experiments, regardless of subtle design differences between them, supported the conclusion that stimuli confirming expectations are more likely to be consciously perceived than stimuli violating expectations. Moreover, in **CHAPTER 3** I showed that the expectation effect was mainly driven by a performance benefit for valid expectations. This facilitatory effect was strongest in the time window where conscious perception is usually most impaired in the attentional blink task [20]. Overall, this puts forward the idea that expectations are able to boost otherwise unseen stimuli into conscious awareness.

Together with other studies showing effects of top-down expectations on conscious stimulus detection, the results I reported are in line with claims derived from prominent predictive brain theories [2,3,30]. In these theories, the brain is thought to

actively interpret incoming evidence in light of expectations based on a “hypothesis” of what the outside world looks like. Conscious perception is assumed to reflect the hypothesis best fitting the incoming evidence at any given moment (while taking into account the precision of the model). At that time point, this is the model probably best representing the outside world [30,40]. In the example of the attentional blink task used here, when a “D” is presented and when it is a priori expected, it is more likely that the hypothesis “D” is the one best explaining the sensory input. On the other hand, if a “K” was expected and a “D” was presented, the overall evidence for the hypothesis that there is a “D” in the outside world is less strong, making it less likely that this hypothesis explains the sensory input well enough to be consciously perceived.

The fact that expectations had a stronger effect at short lags, when conscious target perception is least likely, can potentially be explained by a behavioral ceiling effect when T2 targets were presented late in a trial. In the attentional blink paradigm, overall performance for T2 targets presented at long lags is often high, leaving little room for improvement. Another interesting (not mutually exclusive) possibility is that the relative impact of expectations may be larger at short lags because the sensory input is weighted less. That is, at early lags, when attention “blinks”, sensory input is imprecise and therefore weighted less. As a result, prior expectations have a greater impact on what the “winning” perceptual hypothesis is and therefore on what one consciously perceives (and responds). This hypothesis can be tested by manipulating the attentional resources for the processing of T2 in a controlled way, for example by varying the salience of T1 and thereby changing the “blink rate” at a certain lag. In addition, one could compare this to a situation in which the bottom-up input for the relevant target is actually reduced, in which case one would also expect the effect of expectations to be larger [176].

To be able to better isolate the effect of expectation from other (e.g. attentional) effects, it could be interesting to change the predictive stimuli in the task to another stimulus than T1. In the current attentional blink paradigm, the T1 had two conflicting effects on perception of T2. On the one hand, seeing (and attending to) T1 of course impaired detection of T2 (standard attentional blink effect). However, seeing T1 also allowed forming an expectation about T2, actually facilitating its conscious perception if the expectation was valid. As a result, if for example the stimulus was not consciously

perceived, this could be due to either a strong perception of T1, drawing away attention from T2, or to a weak perception of T1, not allowing the forming of an expectation. Alternatively, adding manipulations of bottom-up stimulus strength could also help to improve understanding of the roles of bottom-up and top-down factors in conscious perception.

A note on response bias

Using target discrimination tasks to measure subjective perceptual experience, such as in [CHAPTER 2](#) and [CHAPTER 3](#), has often been criticized [177]. While the ability to report the correct target seems to imply conscious perception, part of the observed effect may be (partially) driven by decision or response bias effects. That is, participants may regularly report the target letter they expected to see, regardless of whether they have genuinely seen it. Some of the additional analyses reported in [CHAPTER 3](#) indicated that a part of the behavioral effect of expectations is likely driven by such biases. However, in both [CHAPTER 2](#) and [CHAPTER 3](#) I have also reported a bias-free outcome measure for which effects of expectations were qualitatively the same but only somewhat weaker. Therefore, it is not likely that biases can explain the full pattern of expectation-related results.

Nevertheless, I think future work could be done to better isolate the effect of expectations on conscious access from potential biases. In order to do this, it is important to use a response dimension that is entirely orthogonal to the manipulation of top-down expectations. One way to do this is simply asking participants to indicate whether or not they have seen the target stimulus. The manipulation of expectations about the features of upcoming visual stimuli cannot lead to any bias in determining whether a stimulus was presented at all. Another promising approach, similarly unbiased, would be to ask participants to rate the subjective strength of the targets, such as done with the Perceptual Awareness Scale (PAS) [79]. Still, it is advisable to always verify the trustworthiness of participants' responses against more objective measures, since they may not always reliably reflect their percept (i.e. participants rate an expected percept higher, regardless of actual perceptual strength).

AUTOMATICITY OF PERCEPTUAL EXPECTATIONS

Theories about the predictive brain often suggest that expectations are intrinsic and automatic brain computations, implemented by feedback loops originating in

higher-order brain areas [9]. However, this seems incompatible with consciousness theories (see **Box 1.1**), which implicate similar top-down feedback mechanisms in the establishment of conscious perception. In **CHAPTER 2** I have reported two experiments that directly investigated the interaction between consciousness and expectations. Specifically, we investigated to what extent predictive processing was dependent on conscious awareness of either the predictive stimulus (here: T1) or the target stimulus (here: T2).

In an EEG experiment employing the attentional blink task, expectation violations were related to a frontal ERP component. Interestingly, this component was similar for trials on which T2 was seen and trials where on which it was missed. This indicates that the brain recognizes and signals simple violations of expectations in the absence of conscious awareness. Such findings are in line with theories about the predictive brain that assume that expectation-related processes are intrinsic brain computations that happen automatically. In as far as the signals supporting this process are purely driven by a strong feedforward sweep, this could also be in line with theories about consciousness. More research is needed in order to find out how complex expectations can be while still being processed outside of conscious awareness.

Contrary to the findings in the previous paragraph, a behavioral study in **CHAPTER 2** showed that conscious perception of T1, initiating the expectation, was a prerequisite for later influences of this stimulus on conscious perception of T2. When T1 was not seen, the validity of the expectation did not modulate T2 perception. This shows that even after participants were explicitly made aware of the predictive relationship between T1 and T2, these expectations could not be implemented outside of conscious awareness. These findings are in line with theories about consciousness: it is likely that if T1 is not consciously perceived, a lack of recurrent processing between brain areas prevents implementation of any expectation [17,178]. On the contrary, the requirement of conscious perception for the implementation of expectations is more difficult to explain in light of predictive brain in which expectations are seen as intrinsic and automatic brain computations. Based on these assumptions, one would have expected that conscious perception of a stimulus is not required for it to influence predictive processing at a later time-point.

Nonetheless, even when one assumes expectations are intrinsic brain processes, it

is conceivable that the extent to which they are implemented is affected by many factors. For example, maybe only more basic, low-level expectations can be processed outside of awareness, because they rely less on long-range interactive processing in the brain. Indeed, there is some evidence that simpler expectations can be processed unconsciously [100]. As mentioned before, in our task, expectations were defined on a relatively high level for which it is more likely that conscious processing is required. Future work may focus on describing the limits of nonconscious predictive processing, and how this depends on the type, and complexity of the expectations. For example, it would be interesting to see if nonconscious expectation effects can be obtained in the attentional blink paradigm if more automatized (overlearned) expectations are employed. This could be achieved by training participants even more extensively than was currently done. Alternatively, one could take advantage of existing automatized conditional relationships, for example by presenting targets in alphabetical order.

Finally, it is likely that nonconscious stimuli generally instantiate weaker expectations, leading to them only being “available” on shorter time scales (e.g. [101]). In the task used in **CHAPTER 2**, the time between the T1 and T2 was minimally 275 ms. It is possible that nonconscious expectation effects would have been observed on a shorter time scale, because there is a lower chance the expectation has already “faded away”. A recent study in which shorter lags were used (92 ms) does not support this idea [Alilovic et al., in preparation]. Even at the shortest lag, there was no evidence that nonconsciously perceived T1 stimuli introduced effects of expectation validity on subsequent conscious perception of the T2 target. More research is needed to investigate what the time-course of nonconscious predictive processing is. Moreover, it would be interesting to track these processes in the brain using neuroimaging, to see to what extent they depend on feedforward or feedback neural processing.

NEURAL MECHANISMS UNDERLYING THE “BOOSTING” EFFECT OF EXPECTATIONS

In **CHAPTER 2**, I investigated the neural mechanisms underlying the “boosting” effect of expectations on conscious perception by investigating signal amplitude as measured by EEG. Surprisingly, no reliable neural interactions between expectations and conscious perception were observed. That is, in contrast to suggestions by other studies (e.g. [45]), expectation-related modulations of signal amplitude did not relate to changes in conscious access. In the discussion of this chapter, I hypothesized that it may not be

overall signal amplitude, but the “sharpness” of neural representations that is affected by expectations and leads to enhanced conscious perception.

This is what I set out to test in the MEG study reported in **CHAPTER 3**. Using a separate localizer task, it was possible to reliably decode the identity of targets presented within the attentional blink task. However, the representations we picked up were not modulated by expectations and hence cannot readily explain their behavioral impact. At a first glance, these findings appear to contradict intuitions based on predictive brain theories that expected signals should have improved neural representations (as in e.g. [6]). Nevertheless, this may not necessarily be the case. Participants’ expectations in the attentional blink task were defined on a “semantic level” rather than a “perceptual” level. As a result, it is possible that expectations in our task impacted on conscious perception at a relatively high higher cortical level. However, the decoding analyses in **CHAPTER 3** were only sensitive to relatively early sensory parts of the processing of target letters. When a classifier was trained on an independent oddball task, decoding in the attentional blink task was only possible roughly 100-375 ms after stimulus onset and was likely based on posterior MEG channels. This is probably due to the fleeting nature of stimulus representations in the oddball task, where the stimulus identity was not task-relevant. As a consequence, the hypothesis that expectations affect processing of targets at later levels could not be tested in the current dataset and warrants more research.

The suggestion that expectations may affect conscious perception at higher levels of the cortical hierarchy is supported by modelling work in which visibility decisions were related to processing at a late stages of the perceptual processing hierarchy [44,179]. They also make sense in the context of theories about consciousness, which usually implicate high cortical areas, such as prefrontal cortex, in conscious processing. This begs the question whether an area like the (pre)frontal cortex is necessarily at the top of the perceptual processing hierarchy. Predictive brain theories like predictive coding do not require a single specific brain region at the top of this hierarchy. Still, it is often assumed that conscious perception is supported by an optimization of perceptual hypotheses across the cortical hierarchy [30]. This implies optimization of hypotheses at all levels, including high-level hypotheses. Such a process is likely supported by a propagation of neural signals up to the higher cortical areas that process these high-level hypotheses, where neural representations are likely more stable [122]. While

this of course does not directly provide evidence for a causal link between frontal cortical activations and conscious processing, at least it suggests that, and explains why, conscious access may often go hand in hand with activations at the top of the cortical hierarchy, such as in (pre)frontal cortex.

Alternative mechanisms supporting the role of expectations in conscious perception

Obviously, it is not possible, based on empirical results of two experiments only, to rule out that perceptual expectations cause variations in signal amplitude or neural representations that support the facilitatory effect of expectations on conscious perception. In fact, there is some evidence that signal amplitude or signal latency can be modulated by previous experience while also being related to conscious perception [42,45]. Still, it is also possible that behavioral expectation effects are related to other neural factors. One interesting aspect of the neural data to look at, which may even be possible within the existing datasets, could be the variability (i.e. standard deviation) of the neural signal [114,115]. Maybe expectations increase the stability of neural responses, which may in turn lead to more reliable or stable percepts in both space and time.

A majority of recent studies has focused on the role of pre-stimulus oscillatory activity in conscious perception. In particular, the pre-stimulus phase and amplitude of alpha-oscillations have been related to the conscious perception of visual stimuli [124,125,141,180]. An exciting new line of research has shown that these neural markers may also be modulated by top-down factors like expectations [143,144,181]. More research is needed here to isolate the effects of prior expectations from that of more general top-down influences, such as attention (which is often related to alpha oscillations). As a side note, one should realize that the presentation rate of stimuli in the attentional blink paradigm (in my tasks around 11 items/second) lies in the alpha range, making it difficult to disentangle effects of intrinsic neural oscillations from mere entrainment by the 11Hz stimulus presentation protocol. Still, presenting stimuli in alternative frequency ranges has been shown to reduce the overall size of the attentional blink [182], and the blink phenomenon may in fact depend on entrainment to stimuli presentations in the alpha-band [183].

EXPECTATIONS DO NOT MODULATE BOTTOM-UP ATTENTIONAL CAPTURE

In **CHAPTER 4**, I investigated how bottom-up attentional capture interacts with

expectations. In a series of three behavioral experiments using modified exogenous cueing paradigms [33,34], I investigated to what extent temporal or spatial expectations about the cue would influence this cues' influence on targets presented shortly afterwards. Surprisingly, no modulations of attentional capture by cue expectations were observed in any of the experiments. Since our manipulation of expectations instantiated priors in our subjects, this suggests that bottom-up attentional capture was not altered by prior knowledge about the location or time point of potentially distracting inputs, highlighting the relatively automatic, involuntary nature of bottom-up attention [146,173]. The absence of an interactive effect between top-down expectations and attentional capture cannot be easily interpreted within the context of predictive brain theories. Based on those, one would expect that either predictable [37] or unpredictable [157] inputs would be “preferentially processed”, and therefore capture attention. Our evidence supports neither theory, and suggests that bottom-up attention and expectation may not be strongly related [168].

Recent studies have shown that disengagement from attentional capture by distracting stimuli is an active process [184]. It is possible that expectations affect both attentional capture and the active suppression of such influences, but under different conditions. Maybe in our task an expected stimulus initially attracted more attention, but was subsequently also better suppressed. In the experimental tasks used in **CHAPTER 4**, this would be hard to disentangle because only performance for the target stimulus is measured and therefore only an indirect measure of the processing of the cue is available. I would suggest that future studies employ neuroimaging methods, preferably with high temporal resolution, to be better able to tease apart effects on attentional capture itself, effects on suppression of this capture and subsequent effect on target processing in real-time.

In addition, it may be interesting to look at the type of expectations that are used in a task. Here, we looked at the temporal and spatial predictability of visual stimuli, thereby always having expectations that were independent of the feature dimension participants responded to. Maybe if we would have manipulated the expectations on a feature level, like in the other chapters in this thesis, we would have found effects of expectations on attentional capture because the expectations were more relevant. That is, if it is assumed that attentional capture is a way for the system to “snap out of active inference” [30], then it is likely that this would depend on the (estimated)

behavioral relevance of incoming information. In **CHAPTER 4**, the variability of expectedness of cues in different conditions may have been too low to be behaviorally relevant. Moreover, the tasks were all quite easy, making any active suppression of distracting cues more-or-less irrelevant and hence participants may not have invested in doing this.

While we found evidence that conscious processing was modulated by perceptual expectations in **CHAPTER 2** and **CHAPTER 3**, no such evidence was found for attentional processes in **CHAPTER 4**. When considering the idea that conscious processing is closely linked to attention (e.g. [23]), the findings reported here may seem contradictory. However, this is not necessarily the case because conscious processing is often assumed to be related to top-down attention and not bottom-up attention. The current findings may support the idea that bottom-up and top-down attention are independent and affect, or are affected by, expectations differently. That is, perceptual expectations could be correlated only to the top-down attentional processes that later affect conscious awareness, without at the same time affecting bottom-up attentional capture.

CONCLUSION

In summary, the studies I have reported in this thesis show that top-down expectations have a significant impact on conscious access. They have the potential to boost otherwise not perceived stimuli into consciousness. No evidence that these behavioral effects were supported by changes in signal amplitude or early neural representations was found, but preliminary evidence suggests that the boost in conscious access may relate to later decisional processes. While prediction errors can be generated in the absence of conscious awareness, conscious awareness may be essential for the implementation of at least some types of expectations. Finally, although expectations clearly affected conscious access, no evidence for a relationship between expectations and bottom-up attention was found, suggesting these are independent processes and furthermore underscoring the reflexive nature of bottom-up attention.

APPENDIX

REFERENCES

1. Helmholtz H. (1867). *Handbuch der physiologischen optik*. Leipzig: L. Voss.
2. Clark A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36, 181–204.
3. Friston K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360, 815–836.
4. Grossberg S. (2006). Adaptive Resonance Theory. In *Encyclopedia of Cognitive Science*, pp. 1–12. Chichester: John Wiley & Sons, Ltd.
5. Bar M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, 5, 617–629.
6. Kok P, Jehee JFM, de Lange FP. (2012). Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron*, 75, 265–270.
7. Den Ouden HEM, Kok P, de Lange FP. (2012). How prediction errors shape perception, attention, and motivation. *Frontiers in Psychology*, 3, 1–12.
8. Summerfield C, De Lange FP. (2014). Expectation in perceptual decision making: neural and computational mechanisms. *Nature Reviews Neuroscience*, 15, 745–756.
9. Bar M. (2009). Predictions: a universal principle in the operation of the human brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364, 1181–1182.
10. Bekinschtein TA, Dehaene S, Rohaut B, Tadel F, Cohen L, Naccache L. (2009). Neural signature of the conscious processing of auditory regularities. *Proceedings of the National Academy of Sciences*, 106, 1672–1677.
11. Kimura M, Takeda Y. (2015). Automatic prediction regarding the next state of a visual object: Electrophysiological indicators of prediction match and mismatch. *Brain Research*, 1626, 31–44.
12. Garrido MI, Kilner JM, Stephan KE, Friston KJ. (2009). The mismatch negativity: A review of underlying mechanisms. *Clinical Neurophysiology*, 120, 453–463.
13. Strauss M, Sitt JD, King J-R, Elbaz M, Azizi L, Buiatti M, Naccache L, van Wassenhove V, Dehaene S. (2015). Disruption of hierarchical predictive coding during sleep. *Proceedings of the National Academy of Sciences*, 112, E1353–E1362.
14. Rao RPN, Ballard DH. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2, 79–87.
15. Bar M et al. (2006). Top-down facilitation of visual recognition. *Proceedings of the National Academy of Sciences*, 103, 449–454.
16. Lamme VAF, Roelfsema PR. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, 23, 571–579.
17. Lamme VAF. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, 10, 494–501.
18. Crick F, Koch C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences*, 2, 263–275.
19. Dehaene S, Changeux J-P. (2011). Experimental and Theoretical Approaches to Conscious Processing. *Neuron*, 70, 200–227.
20. Raymond JE, Shapiro KL, Arnell KM. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*. 18, 849–860.
21. Seth AK, Dienes Z, Cleeremans A, Overgaard M, Pessoa L. (2008). Measuring consciousness: relating behavioural and neurophysiological approaches. *Trends in Cognitive Sciences*, 12, 314–321.
22. Persuh M. (2018). Measuring perceptual consciousness. *Frontiers in Psychology*, 8, 1–4.
23. 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, 204–211.
24. Tononi G, Koch C. (2008). The neural correlates of consciousness: An update. *Annals of the New*

- York Academy of Sciences, 1124, 239–261.
25. Dehaene S, Sergent C, Changeux J-P. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences*, 100, 8520–8525.
 26. Baars BJ. (2002). The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Sciences*, 6613, 47–52.
 27. Simons DJ, Chabris CF. (1999). Gorillas in Our Midst: Sustained Inattentive Blindness for Dynamic Events. *Perception*, 28, 1059–1074.
 28. Stein T, Peelen M V. (2015). Content-specific expectations enhance stimulus detectability by increasing perceptual sensitivity. *Journal of Experimental Psychology: General*, 144, 1089–1104.
 29. Chang AYC, Kanai R, Seth AK. (2015). Cross-modal prediction changes the timing of conscious access during the motion-induced blindness. *Consciousness and Cognition*, 31, 139–147.
 30. Hohwy J. (2012). Attention and conscious perception in the hypothesis testing brain. *Frontiers in Psychology*, 3, 96.
 31. Mooney CM. (1957). Age in the development of closure ability in children. *Canadian Journal of Psychology*, 11, 219.
 32. Corbetta M, Shulman GL. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3, 201–215.
 33. Posner MI, Cohen Y. (1984). Components of Visual Orienting. *Attention and performance X: Control of language processes*, 32, 531–556.
 34. Posner MI. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32, 3–25.
 35. de Lange FP, Heilbron M, Kok P. (2018). How Do Expectations Shape Perception? *Trends in Cognitive Sciences*, 22, 764–779.
 36. Feldman H, Friston KJ. (2010). Attention, Uncertainty, and Free-Energy. *Frontiers in Human Neuroscience*, 4.
 37. Zhao J, Al-Aidroos N, Turk-Browne NB. (2013). Attention is spontaneously biased toward regularities. *Psychological Science*, 24, 667–677.
 38. Buzsáki G, Anastassiou CA, Koch C. (2012). The origin of extracellular fields and currents-EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13, 407–420.
 39. Panichello MF, Cheung OS, Bar M. (2013). Predictive feedback and conscious visual experience. *Frontiers in Psychology*, 3, 620.
 40. Gregory RL. (1980). Perceptions as Hypotheses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 290, 181–197.
 41. Ratcliff R, McKoon G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, 20, 873–922.
 42. Melloni L, Schwiedrzik CM, Muller N, Rodriguez E, Singer W. (2011). Expectations Change the Signatures and Timing of Electrophysiological Correlates of Perceptual Awareness. *Journal of Neuroscience*, 31, 1386–1396.
 43. Pinto Y, van Gaal S, de Lange FP, Lamme VAF, Seth AK. (2015). Expectations accelerate entry of visual stimuli into awareness. *Journal of Vision*, 15, 13.
 44. De Loof E, Van Opstal F, Verguts T. (2016). Predictive information speeds up visual awareness in an individuation task by modulating threshold setting, not processing efficiency. *Vision Research*, 121, 104–112.
 45. Aru J, Rutiku R, Wibral M, Singer W, Melloni L. (2016). Early effects of previous experience on conscious perception. *Neuroscience of Consciousness*, niw004.
 46. Denison R, Piazza E, Silver M. (2011). Predictive context biases perceptual selection during binocular rivalry. *Journal of Vision*, 11, 312–312.
 47. Pazo-Alvarez P, Cadaveira F, Amenedo E. (2003). MMN in the visual modality: a review. *Biological Psychology*, 63, 199–236.
 48. Näätänen R, Paavilainen P, Rinne T, Alho K. (2007). The mismatch negativity (MMN) in basic research

- of central auditory processing: A review. *Clinical Neurophysiology*, 118, 2544–2590.
49. Ruby P, Caclin A, Boulet S, Delpuech C, Morlet D. (2008). Odd sound processing in the sleeping brain. *Journal of Cognitive neuroscience*, 20, 296–311.
 50. Koelsch S, Heinke W, Sammler D, Olthoff D. (2006). Auditory processing during deep propofol sedation and recovery from unconsciousness. *Clinical Neurophysiology*, 117, 1746–1759.
 51. Näätänen R, Tervaniemi M, Sussman E, Paavilainen P, Winkler I. (2001). ‘Primitive intelligence’ in the auditory cortex. *Trends in Neurosciences*, 24, 283–288.
 52. Stefanics G, Kimura M, Czigler I. (2011). Visual mismatch negativity reveals automatic detection of sequential regularity violation. *Frontiers in Human Neuroscience*, 5, 46.
 53. Stefanics G, Kremláček J, Czigler I. (2016). Mismatch negativity and neural adaptation: Two sides of the same coin. Response: Commentary: Visual mismatch negativity: a predictive coding view. *Frontiers in Human Neuroscience*, 10, 925–939.
 54. Brainard DH. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 433–436.
 55. Love J et al. (2015). JASP.
 56. Rouder JN, Morey RD, Speckman PL, Province JM. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56, 356–374.
 57. Rouder JN, Morey RD, Verhagen J, Swagman AR, Wagenmakers E-J. (2017). Bayesian analysis of factorial designs. *Psychological Methods*, 22, 304–321.
 58. Jeffreys H. (1967). *Theory of Probability*. 3rd edn. Oxford: Oxford University Press.
 59. Oostenveld R, Fries P, Maris E, Schoffelen JM. (2011). FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Intelligence and Neuroscience*, 2011.
 60. Maris E, Oostenveld R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 164, 177–190.
 61. Sergent C, Baillet S, Dehaene S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience*, 8, 1391–1400.
 62. Volpe U, Mucci A, Bucci P, Merlotti E, Galderisi S, Maj M. (2007). The cortical generators of P3a and P3b: A LORETA study. *Brain Research Bulletin*, 73, 220–230.
 63. Stein T, Kaiser D, Peelen M V. (2015). Interobject grouping facilitates visual awareness. *Journal of Vision*, 15, 10.
 64. Yang E, Brascamp J, Kang MS, Blake R. (2014). On the use of continuous flash suppression for the study of visual processing outside of awareness. *Frontiers in Psychology*, 5, 1–17.
 65. Gayet S, Van Der Stigchel S, Paffen CLE. (2014). Breaking continuous flash suppression: Competing for consciousness on the pre-semantic battlefield. *Frontiers in Psychology*, 5, 1–10.
 66. Attarha M, Moore CM. (2015). Onset rivalry: factors that succeed and fail to bias selection. *Attention, Perception, & Psychophysics*, 77, 520–35.
 67. Kranczioch C, Debener S, Engel AK. (2003). Event-related potential correlates of the attentional blink phenomenon. *Cognitive Brain Research*, 17, 177–187.
 68. Harris JA, McMahan AR, Woldorff MG. (2013). Disruption of Visual Awareness during the Attentional Blink Is Reflected by Selective Disruption of Late-stage Neural Processing. *Journal of Cognitive Neuroscience*, 25, 1863–1874.
 69. Rutiku R, Martin M, Bachmann T, Aru J. (2015). Does the P300 reflect conscious perception or its consequences? *Neuroscience*, 298, 180–189.
 70. Desender K, Van Opstal F, Hughes G, Van den Bussche E. (2016). The temporal dynamics of metacognition: Dissociating task-related activity from later metacognitive processes. *Neuropsychologia*, 82, 54–64.
 71. Pitts MA, Padwal J, Fennelly D, Martínez A, Hillyard SA. (2014). Gamma band activity and the P3 reflect post-perceptual processes, not visual awareness. *NeuroImage*, 101, 337–350.
 72. Silverstein BH, Snodgrass M, Shevrin H, Kushwaha R. (2015). P3b, consciousness, and complex unconscious processing. *Cortex*, 73, 216–227.

73. Lasaponara S, Dragone A, Lecce F, Di Russo F, Doricchi F. (2015). The 'serendipitous brain': Low expectancy and timing uncertainty of conscious events improve awareness of unconscious ones (evidence from the Attentional Blink). *Cortex*, 71, 15–33.
74. Visser TW, Ohan JL, Enns JT. (2015). Temporal cues derived from statistical patterns can overcome resource limitations in the attentional blink. *Attention, Perception, & Psychophysics*, 77, 1585–1595.
75. Martens S, Johnson A. (2005). Timing attention: Cuing target onset interval attenuates the attentional blink. *Memory & Cognition*, 33, 234–240.
76. Nieuwenstein MR, Chun MM, van der Lubbe RHJ, Hooge ITC. (2005). Delayed Attentional Engagement in the Attentional Blink. *Journal of Experimental Psychology: Human Perception and Performance*, 31, 1463–1475.
77. Aru J, Bachmann T. (2017). Expectation creates something out of nothing: The role of attention in iconic memory reconsidered. *Consciousness and Cognition*, 53, 203–210.
78. Mack A, Erol M, Clarke J, Bert J. (2016). No iconic memory without attention. *Consciousness and Cognition*, 40, 1–8.
79. Overgaard M, Rote J, Mouridsen K, Ramsøy TZ. (2006). Is conscious perception gradual or dichotomous? A comparison of report methodologies during a visual task. *Consciousness and Cognition*, 15, 700–708.
80. van Gaal S, Lamme VAF. (2012). Unconscious High-Level Information Processing. *The Neuroscientist*, 18, 287–301.
81. Custers R, Aarts H. (2005). Positive Affect as Implicit Motivator: On the Nonconscious Operation of Behavioral Goals. *Journal of Personality and Social Psychology*, 89, 129–142.
82. Pessiglione M, Schmidt L, Draganski B, Kalisch R, Lau H, Dolan RJ, Frith CD. (2007). How the Brain Translates Money into Force: A Neuroimaging Study of Subliminal Motivation. *Science*, 316, 904–906.
83. Pessiglione M, Petrovic P, Daunizeau J, Palminteri S, Dolan RJ, Frith CD. (2008). Subliminal Instrumental Conditioning Demonstrated in the Human Brain. *Neuron*, 59, 561–567.
84. Dehaene S, Naccache L, Cohen L, Bihan D Le, Mangin J-F, Poline J-B, Rivière D. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience*, 4, 752–758.
85. van Gaal S, Naccache L, Meuwese JDI, van Loon AM, Leighton AHL, Cohen L, Dehaene S. (2014). Can the meaning of multiple words be integrated unconsciously? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369, 20130212–20130212.
86. van Gaal S, Lamme VAF, Ridderinkhof KR. (2010). Unconsciously triggered conflict adaptation. *PLoS ONE*, 5, 2–7.
87. van Gaal S, de Lange FP, Cohen MX. (2012). The role of consciousness in cognitive control and decision making. *Frontiers in Human Neuroscience*, 6, 121.
88. Lau HC, Passingham RE. (2007). Unconscious Activation of the Cognitive Control System in the Human Prefrontal Cortex. *Journal of Neuroscience*, 27, 5805–5811.
89. Mathews Z, Cetnarski R, Verschure PFMJ. (2014). Visual anticipation biases conscious decision making but not bottom-up visual processing. *Frontiers in Psychology*, 5, 1–14.
90. Malekshahi R, Seth A, Papanikolaou A, Mathews Z, Birbaumer N, Verschure PFMJ, Caria A. (2016). Differential neural mechanisms for early and late prediction error detection. *Scientific Reports*, 6, 24350.
91. Kenemans JL, Kok A, Smulders FTY. (1993). Event-related potentials to conjunctions of spatial frequency and orientation as a function of stimulus parameters and response requirements. *Electroencephalography and Clinical Neurophysiology*, 88, 51–63.
92. Rohaut B, Alario F-X, Meadow J, Cohen L, Naccache L. (2016). Unconscious semantic processing of polysemous words is not automatic. *Neuroscience of Consciousness*, 2016, niw010.
93. Gayet S, Van der Stigchel S, Paffen CLE. (2014). Seeing is believing: Utilization of subliminal symbols requires a visible relevant context. *Attention, Perception, & Psychophysics*, 76, 489–507.

94. Van Opstal F, Calderon CB, Gevers W, Verguts T. (2011). Setting the stage subliminally: Unconscious context effects. *Consciousness and Cognition*, 20, 1860–1864.
95. Nakamura K, Dehaene S, Jobert A, Le Bihan D, Kouider S. (2007). Task-specific change of unconscious neural priming in the cerebral language network. *Proceedings of the National Academy of Sciences*, 104, 19643–19648.
96. King J-R et al. (2013). Single-trial decoding of auditory novelty responses facilitates the detection of residual consciousness. *NeuroImage*, 83, 726–738.
97. Dykstra AR, Gutschalk A. (2015). Does the mismatch negativity operate on a consciously accessible memory trace? *Science Advances*, 1, e1500677.
98. Luck SJ, Vogel EK, Shapiro KL. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature*, 383, 616–618.
99. Rutiku R, Tulver K, Aru J, Bachmann T. (2016). Visual masking with frontally applied pre-stimulus TMS and its subject-specific neural correlates. *Brain Research*, 1642, 136–145.
100. Chang R, Baria AT, Flounders MW, He BJ. (2016). Unconsciously elicited perceptual prior. *Neuroscience of Consciousness*, 2016, niw008.
101. Greenwald AG, Draine SC, Abrams RL. (1996). Three cognitive markers for unconscious semantic activation. *Science*, 173, 1699–1702.
102. King J-R, Pescetelli N, Dehaene S. (2016). Brain Mechanisms Underlying the Brief Maintenance of Seen and Unseen Sensory Information. *Neuron*, 92, 1122–1134.
103. Hsieh P-J, Colas JT, Kanwisher N. (2011). Pop-Out Without Awareness. *Psychological Science*, 22, 1220–1226.
104. Mulckhuysen M, Theeuwes J. (2010). Unconscious attentional orienting to exogenous cues: A review of the literature. *Acta Psychologica*, 134, 299–309.
105. Ansorge U, Kiss M, Eimer M. (2009). Goal-driven attentional capture by invisible colors: Evidence from event-related potentials. *Psychonomic Bulletin & Review*, 16, 648–653.
106. Nieuwenstein M, Van der Burg E, Theeuwes J, Wyble B, Potter M. (2009). Temporal constraints on conscious vision: On the ubiquitous nature of the attentional blink. *Journal of Vision*, 9, 18–18.
107. Cheadle S, Egnér T, Wyart V, Wu C, Summerfield C. (2015). Feature expectation heightens visual sensitivity during fine orientation discrimination. *Journal of Vision*, 15, 14.
108. Todorovic A, de Lange FP. (2012). Repetition Suppression and Expectation Suppression Are Dissociable in Time in Early Auditory Evoked Fields. *Journal of Neuroscience*, 32, 13389–13395.
109. Alink A, Schwiedrzik CM, Kohler A, Singer W, Muckli L. (2010). Stimulus Predictability Reduces Responses in Primary Visual Cortex. *Journal of Neuroscience*, 30, 2960–2966.
110. King J-R, Dehaene S. (2014). A model of subjective report and objective discrimination as categorical decisions in a vast representational space. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369, 20130204–20130204.
111. Meijs EL, Slagter HA, de Lange FP, van Gaal S. (2018). Dynamic Interactions between Top-Down Expectations and Conscious Awareness. *Journal of Neuroscience*, 38, 2318–2327.
112. Bell AH, Summerfield C, Morin EL, Malecek NJ, Ungerleider LG. (2016). Encoding of Stimulus Probability in Macaque Inferior Temporal Cortex. *Current Biology*, 26, 2280–90.
113. Kok P, Mostert P, de Lange FP. (2017). Prior expectations induce prestimulus sensory templates. *Proceedings of the National Academy of Sciences*, 114, 201705652.
114. Schurger A, Sarigiannidis I, Naccache L, Sitt JD, Dehaene S. (2015). Cortical activity is more stable when sensory stimuli are consciously perceived. *Proceedings of the National Academy of Sciences*, 112, E2083–E2092.
115. Schurger A, Pereira F, Treisman A, Cohen JD. (2010). Reproducibility Distinguishes Conscious from Nonconscious Neural Representations. *Science*, 327, 97–99.
116. Stolk A, Todorovic A, Schoffelen JM, Oostenveld R. (2013). Online and offline tools for head movement compensation in MEG. *NeuroImage*, 68, 39–48.
117. Mostert P, Kok P, de Lange FP. (2015). Dissociating sensory from decision processes in human

- perceptual decision making. *Scientific Reports*, 5, 18253.
118. King J-R, Dehaene S. (2014). Characterizing the dynamics of mental representations: The temporal generalization method. *Trends in Cognitive Sciences*, 18, 203–210.
 119. Marti S, Dehaene S. (2017). Discrete and continuous mechanisms of temporal selection in rapid visual streams. *Nature Communications*, 8.
 120. Stokes M, Thompson R, Nobre AC, Duncan J. (2009). Shape-specific preparatory activity mediates attention to targets in human visual cortex. *Proceedings of the National Academy of Sciences*, 106, 19569–19574.
 121. Myers NE, Stokes MG, Nobre AC. (2017). Prioritizing Information during Working Memory: Beyond Sustained Internal Attention. *Trends in Cognitive Sciences*, 21, 449–461.
 122. Kiebel SJ, Daunizeau J, Friston KJ. (2008). A hierarchy of time-scales and the brain. *PLoS Computational Biology*, 4.
 123. Fahrenfort JJ, van Leeuwen J, Olivers CNL, Hogendoorn H. (2017). Perceptual integration without conscious access. *Proceedings of the National Academy of Sciences*, 114, 3744–3749.
 124. Mathewson KE, Gratton G, Fabiani M, Beck DM, Ro T. (2009). To See or Not to See: Prestimulus Phase Predicts Visual Awareness. *Journal of Neuroscience*, 29, 2725–2732.
 125. Iemi L, Chaumon M, Crouzet SM, Busch NA. (2017). Spontaneous Neural Oscillations Bias Perception by Modulating Baseline Excitability. *Journal of Neuroscience*, 37, 807–819.
 126. Linkenkaer-Hansen K, Nikulin V V, Palva S, Ilmoniemi RJ, Palva JM. (2004). Prestimulus oscillations enhance psychophysical performance in humans. *Journal of Neuroscience*, 24, 10186–90.
 127. Ouimet C, Jolicoeur P. (2007). Beyond Task 1 difficulty: The duration of T1 encoding modulates the attentional blink. *Visual Cognition*, 15, 290–304.
 128. Slagter HA, Mazaheri A, Reteig LC, Smolders R, Figeer M, Mantione M, Schuurman PR, Denys D. (2017). Contributions of the Ventral Striatum to Conscious Perception: An Intracranial EEG Study of the Attentional Blink. *Journal of Neuroscience*, 37, 1081–1089.
 129. Jiang J, Summerfield C, Eger T. (2013). Attention Sharpens the Distinction between Expected and Unexpected Percepts in the Visual Brain. *Journal of Neuroscience*, 33, 18438–18447.
 130. Kok P, Rahnev D, Jehee JFM, Lau HC, De Lange FP. (2012). Attention reverses the effect of prediction in silencing sensory signals. *Cerebral Cortex*, 22, 2197–2206.
 131. Jehee JFM, Brady DK, Tong F. (2011). Attention Improves Encoding of Task-Relevant Features in the Human Visual Cortex. *Journal of Neuroscience*, 31, 8210–8219.
 132. Ress D, Backus BT, Heeger DJ. (2000). Activity in primary visual cortex predicts performance in a visual detection task. *Nature Neuroscience*, 3, 940–945.
 133. van Vugt B, Dagnino B, Vartak D, Safaai H, Panzeri S, Dehaene S, Roelfsema PR. (2018). The threshold for conscious report: Signal loss and response bias in visual and frontal cortex. *Science*, 7186, eaar7186.
 134. Pitts MA, Metzler S, Hillyard SA. (2014). Isolating neural correlates of conscious perception from neural correlates of reporting one's perception. *Frontiers in Psychology*, 5, 1–16.
 135. Busch NA, Dubois J, VanRullen R. (2009). The Phase of Ongoing EEG Oscillations Predicts Visual Perception. *Journal of Neuroscience*, 29, 7869–7876.
 136. Pincham HL, Szucs D. (2012). Conscious access is linked to ongoing brain state: Electrophysiological evidence from the attentional blink. *Cerebral Cortex*, 22, 2346–2353.
 137. Petro NM, Keil A. (2015). Pre-target oscillatory brain activity and the attentional blink. *Experimental Brain Research*, 233, 3583–3595.
 138. de Gee JW, Knapen T, Donner TH. (2014). Decision-related pupil dilation reflects upcoming choice and individual bias. *Proceedings of the National Academy of Sciences*, 111, E618–E625.
 139. Rungratsameetaweemana N, Itthipuripat S, Salazar A, Serences JT. (2018). Expectations do not alter early sensory processing during perceptual decision making. *Journal of Neuroscience*, 38, 5632–5648.
 140. Bang JW, Rahnev D. (2017). Stimulus expectation alters decision criterion but not sensory signal in

- perceptual decision making. *Scientific Reports*, 7, 1–12.
141. Benwell CSY, Tagliabue CF, Veniero D, Cecere R, Savazzi S, Thut G. (2017). Pre-stimulus EEG power predicts conscious awareness but not objective visual performance. *Eneuro*, 4.
 142. de Lange FP, Rahnev DA, Donner TH, Lau H. (2013). Prestimulus Oscillatory Activity over Motor Cortex Reflects Perceptual Expectations. *Journal of Neuroscience*, 33, 1400–1410.
 143. Sherman MT, Kanai R, Seth AK, VanRullen R. (2016). Rhythmic Influence of Top-Down Perceptual Priors in the Phase of Prestimulus Occipital Alpha Oscillations. *Journal of Cognitive Neuroscience*, 28, 1318–1330.
 144. Mayer A, Schwiedrzik CM, Wibral M, Singer W, Melloni L. (2016). Expecting to see a letter: Alpha oscillations as carriers of top-down sensory predictions. *Cerebral Cortex*, 26, 3146–3160.
 145. Theeuwes J, Reimann B, Mortier K. (2006). Visual search for featural singletons: No top-down modulation, only bottom-up priming. *Visual Cognition*, 14, 466–489.
 146. Theeuwes J. (2010). Top-down and bottom-up control of visual selection. *Acta Psychologica*, 135, 77–99.
 147. Folk CL, Remington RW, Johnston JC. (1992). Involuntary Covert Orienting Is Contingent On Attentional Control Settings. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 1030–1044.
 148. Belopolsky A V, Schreij D, Theeuwes J. (2010). What is top-down about contingent capture? *Attention, Perception, & Psychophysics*, 72, 326–341.
 149. Munneke J, Hoppenbrouwers SS, Theeuwes J. (2015). Reward can modulate attentional capture, independent of top-down set. *Attention, Perception, & Psychophysics*, 77, 2540–2548.
 150. Failing MF, Theeuwes J. (2014). Exogenous visual orienting by reward. *Journal of Vision*, 14, 6–6.
 151. Awh E, Belopolsky A V, Theeuwes J. (2012). Top-down versus bottom-up attentional control: a failed theoretical dichotomy. *Trends in Cognitive Sciences*, 16, 437–443.
 152. Folk CL, Remington RW. (2015). Unexpected Abrupt Onsets Can Override a Top-Down Set for Color. *Journal of Experimental Psychology: Human Perception and Performance*, 41, 1153–1165.
 153. Schönhammer JG, Kerzel D. (2018). Optimal task-sets override attentional capture by rare cues. *Journal of Experimental Psychology: Human Perception and Performance*, 44, 681–692.
 154. Müller HJ, Geyer T, Zehetleitner M, Krummenacher J. (2009). Attentional Capture by Salient Color Singleton Distractors Is Modulated by Top-Down Dimensional Set. *Journal of Experimental Psychology: Human Perception and Performance*, 35, 1–16.
 155. Becker SI, Horstmann G. (2011). Novelty and saliency in attentional capture by unannounced motion singletons. *Acta Psychologica*, 136, 290–299.
 156. Wang F, Chen M, Yan Y, Zhaoping L, Li W. (2015). Modulation of Neuronal Responses by Exogenous Attention in Macaque Primary Visual Cortex. *Journal of Neuroscience*, 35, 13419–13429.
 157. Horstmann G. (2015). The surprise-attention link: a review. *Annals of the New York Academy of Sciences*, 1339, 106–115.
 158. Itti L, Baldi P. (2009). Bayesian surprise attracts human attention. *Vision Research*, 49, 1295–1306.
 159. Kaya EM, Elhilali M. (2014). Investigating bottom-up auditory attention. *Frontiers in Human Neuroscience*, 8, 327.
 160. Nöstl A, Marsh JE, Sörqvist P. (2012). Expectations Modulate the Magnitude of Attentional Capture by Auditory Events. *PLoS ONE*, 7.
 161. Winkler I. (2007). Interpreting the Mismatch Negativity. *Federation of European Psychophysiology Societies*, 21, 147–163.
 162. Grubb MA, White AL, Heeger DJ, Carrasco M. (2015). Interactions between voluntary and involuntary attention modulate the quality and temporal dynamics of visual processing. *Psychonomic Bulletin & Review*, 22, 437–444.
 163. Gaspelin N, Luck SJ. (2018). The Role of Inhibition in Avoiding Distraction by Salient Stimuli. *Trends in Cognitive Sciences*, 22, 79–92.
 164. Andreou LV, Kashino M, Chait M. (2011). The role of temporal regularity in auditory segregation.

- Hearing Research*, 280, 228–235.
165. Sawaki R, Luck SJ. (2010). Capture versus suppression of attention by salient singletons: Electrophysiological evidence for an automatic attend-to-me signal. *Attention, Perception, & Psychophysics*, 72, 1455–1470.
 166. Kiss M, Grubert A, Petersen A, Eimer M. (2012). Attentional Capture by Salient Distractors during Visual Search Is Determined by Temporal Task Demands. *Journal of Cognitive Neuroscience*, 24, 749–759.
 167. Carrasco M, Ling S, Read S. (2004). Attention alters appearance. *Nature Neuroscience*, 7, 308–313.
 168. Southwell R, Baumann A, Gal C, Barascud N, Friston K, Chait M. (2017). Is predictability salient? A study of attentional capture by auditory patterns. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372, 20160105.
 169. Friston K. (2009). The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, 13, 293–301.
 170. Morillon B, Schroeder CE, Wyart V, Arnal LH. (2016). Temporal Prediction in lieu of Periodic Stimulation. *Journal of Neuroscience*, 36, 2342–2347.
 171. Wang B, Theeuwes J. (2018). Statistical regularities modulate attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, 44, 13–17.
 172. Wang B, Theeuwes J. (2018). How to inhibit a distractor location? Statistical learning versus active, top-down suppression. *Attention, Perception, & Psychophysics*, 80, 860–870.
 173. Theeuwes J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics*, 51, 599–606.
 174. Thomson DR, Willoughby K, Milliken B. (2014). Implicit learning modulates attention capture: evidence from an item-specific proportion congruency manipulation. *Frontiers in Psychology*, 5, 1–14.
 175. Noonan MP, Adamian N, Pike A, Printzlau F, Crittenden BM, Stokes MG. (2016). Distinct Mechanisms for Distractor Suppression and Target Facilitation. *Journal of Neuroscience*, 36, 1797–1807.
 176. Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN. (2011). Elapsed Decision Time Affects the Weighting of Prior Probability in a Perceptual Decision Task. *Journal of Neuroscience*, 31, 6339–6352.
 177. Schneider KA, Komlos M. (2008). Attention biases decisions but does not alter appearance. *Journal of Vision*, 8, 3.1-10.
 178. Dehaene S, Charles L, King J-R, Marti S. (2014). Toward a computational theory of conscious processing. *Current Opinion in Neurobiology*, 25, 76–84.
 179. Maniscalco B, Lau H. (2016). The signal processing architecture underlying subjective reports of sensory awareness. *Neuroscience of Consciousness*, 2016, 801–10.
 180. Samaha J, Bauer P, Cimaroli S, Postle BR. (2015). Top-down control of the phase of alpha-band oscillations as a mechanism for temporal prediction. *Proceedings of the National Academy of Sciences*, 112, 8439–8444.
 181. Samaha J, Boutonnet B, Postle BR, Lupyan G. (2018). Effects of meaningfulness on perception: Alpha-band oscillations carry perceptual expectations and influence early visual responses. *Scientific Reports*, 8, 6606.
 182. Shapiro KL, Hanslmayr S, Enns JT, Lleras A. (2017). Alpha, beta: The rhythm of the attentional blink. *Psychonomic Bulletin and Review*, 24, 1862–1869.
 183. Zauner A, Fellinger R, Gross J, Hanslmayr S, Shapiro K, Gruber W, Müller S, Klimesch W. (2012). Alpha entrainment is responsible for the attentional blink phenomenon. *NeuroImage*, 63, 674–686.
 184. Liesefeld HR, Liesefeld AM, Töllner T, Müller HJ. (2017). Attentional capture in visual search: Capture and post-capture dynamics revealed by EEG. *NeuroImage*, 156, 166–173.

NEDERLANDSE SAMENVATTING

In het dagelijks leven hebben we gewoonlijk de ervaring dat we de wereld om ons heen natuurgetrouw waarnemen. Deze subjectieve ervaring is echter onjuist. Het is al lang bekend dat onze waarneming geen passieve registratie van de buitenwereld is. Waarneming moet eerder gezien worden als een actief proces, waarbij het brein probeert binnenkomende zintuiglijke signalen te begrijpen. Onze waarneming is daardoor minstens zo veel een interpretatie als een representatie van de buitenwereld.

Recent is gesuggereerd dat het brein bij de interpretatie van zintuiglijke signalen gebruik maakt van eerder opgedane kennis. Mede op basis van deze kennis kan het brein verwachtingen vormen over mogelijke toekomstige zintuiglijke input. Na het horen van geblaf kan een verwachting worden gevormd dat je waarschijnlijk een hond gaat zien, en niet een kat. Zo kunnen nieuwe signalen snel en efficiënt worden verwerkt. Theorieën die dit aannemen spreken ook wel over “the predictive brain”, ofwel “het voorspellende brein”. Hoewel basisprincipes van deze theorieën worden ondersteund door veel wetenschappelijk onderzoek, zijn er ook nog veel aspecten van de theorieën waarover niet alles duidelijk is. Zo is nog niet veel bekend over de manier waarop verwachtingen in het “voorspellende brein” samenhangen met andere processen, zoals aandacht en bewustzijn. In dit proefschrift staan zeven experimenten waarin binnen de context van visuele waarneming onderzoek is gedaan naar deze verbanden.

VISUELE VERWACHTINGEN EN BEWUSTE WAARNEMING

In **HOOFDSTUK 2** en **HOOFDSTUK 3** van dit proefschrift wordt gekeken naar de relatie tussen visuele verwachtingen en bewuste waarneming. Daarbij wordt gebruik gemaakt van de zogeheten “attentional blink” taak. In deze taak zien proefpersonen telkens een stroom van letters die zeer snel na elkaar (ongeveer elke tiende van een seconde) worden gepresenteerd. Een aantal letters van het alfabet zijn vooraf aangewezen als zogenaamde targets (= doelstimuli). Proefpersonen moeten deze targets detecteren binnen de stroom van letters. Als er een enkel target is kunnen proefpersonen dit bijna altijd probleemloos vinden. Eerder onderzoek heeft echter laten zien dat als na zo’n target snel nog een ander target wordt gepresenteerd, dit tweede target vaak niet bewust wordt waargenomen.

Onze versie van de taak was dusdanig gemanipuleerd dat het eerste target voorspelde wat het tweede target waarschijnlijk zou zijn. Dit werd bereikt door elk eerste target te koppelen aan een tweede target dat vaak op dit eerste target volgde, dus een “G” werd bijvoorbeeld in 75% van de gevallen gevolgd door een “K”, terwijl een “H” vaak werd gevolgd door een “D” en juist bijna nooit door een “K”. Deze patronen zorgden ervoor dat proefpersonen na het zien van het eerste target een verwachting hadden wat het meest waarschijnlijke opvolgende target zou zijn. De vraag die hierbij werd gesteld was of deze verwachtingen invloed zouden hebben op het aantal targets dat proefpersonen bewust waar zouden nemen.

In een reeks aan experimenten lieten de resultaten keer op keer zien dat een juiste verwachting proefpersonen hielp om het tweede target bewust waar te nemen. We wilden ook graag weten welk proces in het brein hieraan ten grondslag lag. Onze hypothese was dat verwachtingen zorgen voor een efficiëntere verwerking van visuele signalen in gebieden van het brein die relatief vroeg betrokken zijn bij de verwerking van visuele signalen. Deze hypothese testten we met twee neuroimagingexperimenten, waarin we met EEG of MEG afgeleiden van elektrische activiteit in het brein meten. In tegenstelling tot onze hypothese, vonden we suggesties dat het gunstige effect van verwachtingen op bewuste waarneming gerelateerd was aan beslisprocessen in het brein die in een relatief laat stadium plaatsvinden.

In de taak kan het voorkomen dat een proefpersoon verwacht een bepaalde target letter (bijvoorbeeld “D”) te gaan zien, terwijl uiteindelijk een andere letter (bijvoorbeeld “K”) wordt gepresenteerd. Onverwachte gebeurtenissen, zoals het zien van een dergelijk “onverwacht” target, leiden vaak tot een “mismatch signaal” dat we met EEG kunnen meten. Dit signaal laat zien dat het brein heeft geconstateerd dat de verwachte en daadwerkelijke zintuiglijke input (zoals hier de letter) niet met elkaar overeenkwamen. Opvallend genoeg was dit signaal even sterk wanneer proefpersonen niet konden aangeven welke target letter ze gezien hadden. Bewuste waarneming van een target lijkt dus geen voorwaarde voor de vergelijking van input met eerder gemaakte verwachtingen. Het feit dat dit validatieproces onbewust kan plaatsvinden, suggereert dat het een relatief automatisch proces is.

In een laatste experiment zorgden we ervoor dat proefpersonen het eerste target, waarop ze hun verwachting moesten baseren, regelmatig niet bewust konden

waarnemen. Als proefpersonen het eerste target niet bewust hadden waargenomen, maakte de voorspellende waarde van dit target niet meer uit. Een visuele stimulus moet dus bewust zijn waargenomen om op basis van deze stimulus voorspellingen te kunnen maken welke toekomstige waarnemingen te verwachten zijn.

VISUELE VERWACHTINGEN EN AANDACHT

In een andere reeks van drie experimenten, in **HOOFDSTUK 4**, werd gekeken naar de relatie tussen visuele verwachtingen en “bottom-up” aandacht. Dit is aandacht die door externe gebeurtenissen wordt gestuurd. In een bekende “bottom-up” aandachtstaak wordt gekeken hoe het presenteren van een visuele stimulus op een bepaalde locatie in het visuele gezichtsveld (in experimenten vaak een plaats op het computerscherm) vervolgens leidt tot het richten van aandacht op die locatie. Stel je bijvoorbeeld een lamp voor die plotseling aangaat: welke factoren bepalen of je vervolgens extra aandacht hebt voor de locatie waar deze lamp zich bevindt. Onze hypothese was dat vooral onverwachte signalen (de verdeling van) aandacht zouden sturen. We vonden hier echter geen bewijs voor: elke stimulus trok de aandacht steeds evenveel, ongeacht of deze stimulus qua locatie of timing te verwachten was. Er lijkt dus geen relatie te zijn tussen visuele verwachtingen en “bottom-up” aandacht.

PUBLICATION LIST

PUBLICATIONS

- Francken JC, **Meijs EL**, Hagoort P, van Gaal S, & de Lange FP. (2016). Exploring the automaticity of language-perception interactions: Effects of attention and awareness. *Scientific Reports*, 5, 17725.
- Francken JC, **Meijs EL**, Ridderinkhof OM, Hagoort P, de Lange FP, & van Gaal, S. (2015). Manipulating word awareness dissociates feed-forward from feedback models of language-perception interactions. *Neuroscience of Consciousness*, 2015, niv003.
- Meijs EL**, Klaassen FH, Bokeria L, van Gaal S, & de Lange FP. (2018). Cue predictability does not modulate bottom-up attentional capture. *Royal Society Open Science*, 5, 180524.
- Meijs EL**, Slagter HA, de Lange FP, & van Gaal S. (2018). Dynamic Interactions between Top–Down Expectations and Conscious Awareness. *Journal of Neuroscience*, 38, 2318–2327.

FORTHCOMING

- Meijs EL**, Mostert P, Slagter HA, de Lange FP, & van Gaal S. (in preparation). Exploring the role of expectations and stimulus relevance on stimulus-specific neural representations and conscious report.
- Mongelli V, **Meijs EL**, van Gaal S, & Hagoort P. (in preparation). No language unification without neural feedback: How awareness modulates semantic combinatorial processes.

CURRICULUM VITAE

Erik Meijs was born on June 5th 1990 in Tilburg, the Netherlands. He grew up in Goirle, but after graduating from the Mill-Hill college in 2008, he moved to Nijmegen to study Psychology at the Radboud University. During the second and third year of this study, Erik participated in the Radboud Honours Academy, which offered extra courses and research opportunities. After receiving his Bachelor's degree with distinction, he started the research master Cognitive Neuroscience at the same university. As part of this master, Erik did a research internship under supervision of Jolien Francken, Simon van Gaal and Floris de Lange, studying the role of conscious awareness in language-perception interactions. This internship led to writing the grating proposal that funded Erik's PhD project to study the interactions between visual expectations and conscious awareness. During his studies and his PhD, Erik was involved in the teaching of multiple courses at the Bachelor level. In the final year of his PhD, he took an opportunity offered by the PhD Network Netherlands (PNN) to acquire work experience outside the scientific world, at the Dutch Inspectorate of Education. He later accepted a job at the Inspectorate of Education, where he has been employed as an analyst in the department of Higher Education since finishing his PhD manuscript.

ACKNOWLEDGEMENTS

This thesis would unquestionably not have been possible without the help and support of many others.

Floris and **Simon**, if it weren't for the two of you I would have never been able to start my PhD project in the first place. Thank you for taking me under your wings and supporting and encouraging me all along the way. Working with a dynamic duo like you guys was a great pleasure on both a scientific and a personal level. Your enthusiasm during our meetings was instrumental in keeping my motivation up when things were not going according to plan (which happened more often than one would hope). **Simon**, even though you were generally supervising me from a distance, I have always felt you were easy to reach when I needed advice. **Floris**, it was a great pleasure to be part of your lab. I really admire your ability to create a very friendly and social, but at the same time intellectually stimulating environment.

Roshan, even though we did not end up working together much, your willingness to support my PhD has meant a lot to me. Thank you for allowing me to follow my own path, even when this made me stray far from your own area of expertise. Furthermore, I want to thank you, **Floris** and **Simon** for giving me the freedom to explore other options than a scientific career, and for supporting my choice to leave science after finishing my PhD.

Thanks to all the past and present members of the **Predictive Brain lab** for their ideas, feedback and practical support. I have learned a great deal from all of you and being in a group with you was amazing. Thank you for all the times you agreed to be a buddy or (pilot) subject in one of my studies. I have fond memories of our meetings, retreats and the many social events (did someone say sushi?) we shared. I have also enjoyed our weekly group lunches a lot and I feel honored to have been your meme king for so many years. It is good to see the meme tradition is still continuing without me.

A number of people from the Predictive Brain Lab have been especially important in the realization of this thesis. **Felix** and **Levan**, it has been a great pleasure to supervise you while working on our bottom-up attention project. You guys managed to test a

ridiculous amount of participants. **Pim**, thank you for the countless times you have helped me with my study designs or (decoding) analyses. **Jolien**, thank you for your unlimited (scientific) enthusiasm. I learned so much from you during and after my master's internship, and I would have never started my PhD journey if it wasn't for you.

Pim and **Lieke**, thank you so much for agreeing to be my paranymphs and helping me get through the administrative and organizational hassle surrounding a PhD defense.

I was also lucky enough to be able to work with some amazing people from outside the Donders Institute. **Sjoerd**, **Doris** and **Thomas**, thank you for all your help in collecting data for my attentional blink experiments. You have made my life a lot easier. **Heleen**, thank you for all your insights and inputs, which have without a doubt greatly improved both my studies and papers. **Josipa**, I am happy to see somebody is continuing the attentional blink projects and to know I leave this line of research in good hands. I am curious to see what future experiments will show. **Valeria**, it was nice working with you on your N400 projects. I'm looking forward to seeing the resulting paper soon!

The **Donders Institute** is a great place to do a PhD, and I count myself lucky I was able to work there. Besides the expertise and intellectually stimulating environment, the institute also has a great support structure that makes a researchers' life lot easier. There are too many people to thank all of them personally, so I would like to thank everyone in the institute, especially in the **TG** and the **Administration**, for all their help and support.

Throughout the years, I had the pleasure to share my office at the Donders with a lot of incredibly nice people. Thank you **Jim**, **Daniel**, **Peter**, **Mirjam**, **Irati**, **Ann-Katrin**, **Ella**, **João** and **Martin** for keeping me company, for the useful input and for the many interesting and/or fun conversations we had.

Finishing a PhD would not have been possible without coffee. I really enjoyed drinking my coffee together with many nice colleagues, like **Erik**, **Pim**, **Matthias**, **Maarten** and **Lieke**. Especially in the final year of my PhD trajectory, "early" morning coffee appointments with **Lieke** and **Iris** were important in motivating me to get to work in time.

I also spend some time teaching students in the Psychology and Biology bachelor programs. I want to thank all my teaching colleagues for their help and support. In this context, I want to especially thank **Eric**, for giving me the opportunity to get some experience with all aspects of teaching at a university.

When a Swede, a Belgian, a Canadian, a Brit, a South-African and a Dutchman walk into a bar, this is not the start of a joke. It is quiz time for Fauxkapi! Long time quiz buddies **Daniel, Matthias, Peter, Amie** and **Ashley**, thank you for joining me for our weekly quiz nights throughout all these years, even though it has been obvious for a long time that we kind of suck at quizzing. Also thanks to many others, such as **Lorijn, Jeff, Sonny** and **Andrew**, that have, against all odds, at some point in time joined to team to help us claim a quizzing victory. With our quiz team, we have a long-running tradition to have dinner together before the quiz starts. **Daniel, Matthias, Peter** and **Amie**, thank you for all those great evenings with (not so) culinary adventures and lots of fun. Let's hope the tradition continues for a long time and will keep including the occasional Mario Kart or Goldeneye on the Nintendo64.

Rasmus and **Vincent**, there are just too many great memories from our student life together to even start listing them. Even though you guys moved to Berlin before I started my PhD, you have had a great impact my career. It has been great to share an interest in cognitive neuroscience with you guys. I'm happy we are still in touch and look forward to visiting Berlin to celebrate your PhD defenses in the (near) future!

Ik wil ook graag iedereen bij de **Onderwijsinspectie**, en met name in het Sectorbeelden team, erg bedanken voor het warme welkom tijdens mijn stage. Ik had en heb het erg naar mijn zin bij de inspectie en ben blij met de kansen die jullie me hebben geboden en het vertrouwen dat jullie in mij hebben getoond door mij na mijn promotietraject aan te nemen als analist. Daarnaast ben ik erg dankbaar dat jullie mij de ruimte hebben gegeven om ook mijn promotie nog tot een goed einde te brengen.

Verder wil ik nog heel graag mijn familie bedanken voor alle interesse, steun en vertrouwen die ik door de jaren heen van ze kreeg. **Inge** en **Marieke**, we zien elkaar dan wel niet meer zo vaak sinds ik in het verre Nijmegen woon, onze band is daar gelukkig niet slechter door geworden. Ik ben blij dat we zo goed met elkaar kunnen opschieten en regelmatig veel plezier samen kunnen hebben. Laten we Concert at Sea

2020 vast in de agenda's zetten! **Oma Corrie**, dit boekje zou niet compleet zijn als ik u niet had bedankt voor alles wat u vroeger voor ons hebt gedaan. Niemand als ik die zo'n omaatje heeft. **Papa en Mama**, jullie hebben me altijd vrijgelaten om mijn eigen weg te kiezen. Daarbij hebben jullie me immer gesteund en stonden jullie op de juiste momenten klaar om me met raad en daad vooruit te helpen. Ik wil jullie ontzettend bedanken voor jullie onuitputtelijke steun en vertrouwen. Ik ben ervan overtuigd dat ik dit alles nooit zonder jullie had kunnen bereiken.

Stephanie, ik kan nauwelijks onder woorden brengen hoe blij ik met je ben en hoe belangrijk je voor me bent. Ik ben zo ontzettend gelukkig dat jij al zoveel jaar aan mijn zijde staat om me te ondersteunen. Jij zorgt ervoor dat ik de best mogelijke versie van mijzelf ben. Het is bewonderenswaardig dat een gebrek aan interesse in de wetenschap jou nooit heeft weerhouden me te laten weten hoe trots je op me bent. Weet dat ik ook ongelofelijk trots op jou ben. Ik weet zeker dat we nog een ontzettend mooie toekomst tegemoet gaan. Ik hou van je!

DONDERS GRADUATE SCHOOL FOR COGNITIVE NEUROSCIENCE

For a successful research Institute, it is vital to train the next generation of young scientists. To achieve this goal, the Donders Institute for Brain, Cognition and Behaviour established the Donders Graduate School for Cognitive Neuroscience (DGCN), which was officially recognised as a national graduate school in 2009. The Graduate School covers training at both Master's and PhD level and provides an excellent educational context fully aligned with the research programme of the Donders Institute.

The school successfully attracts highly talented national and international students in biology, physics, psycholinguistics, psychology, behavioral science, medicine and related disciplines. Selective admission and assessment centers guarantee the enrolment of the best and most motivated students.

The DGCN tracks the career of PhD graduates carefully. More than 50% of PhD alumni show a continuation in academia with postdoc positions at top institutes worldwide, e.g. Stanford University, University of Oxford, University of Cambridge, UCL London, MPI Leipzig, Hanyang University in South Korea, NTNU Norway, University of Illinois, North Western University, Northeastern University in Boston, ETH Zürich, University of Vienna etc.. Positions outside academia spread among the following sectors: specialists in a medical environment, mainly in genetics, geriatrics, psychiatry and neurology. Specialists in a psychological environment, e.g. as specialist in neuropsychology, psychological diagnostics or therapy. Positions in higher education as coordinators or lecturers. A smaller percentage enters business as research consultants, analysts or head of research and development. Fewer graduates stay in a research environment as lab coordinators, technical support or policy advisors. Upcoming possibilities are positions in the IT sector and management position in pharmaceutical industry. In general, the PhDs graduates almost invariably continue with high-quality positions that play an important role in our knowledge economy.

For more information on the DGCN as well as past and upcoming defenses please visit: <http://www.ru.nl/donders/graduate-school/phd/>

DONDERS

I N S T I T U T E

ISBN 978-94-6284-184-0

Radboud University



Radboudumc