

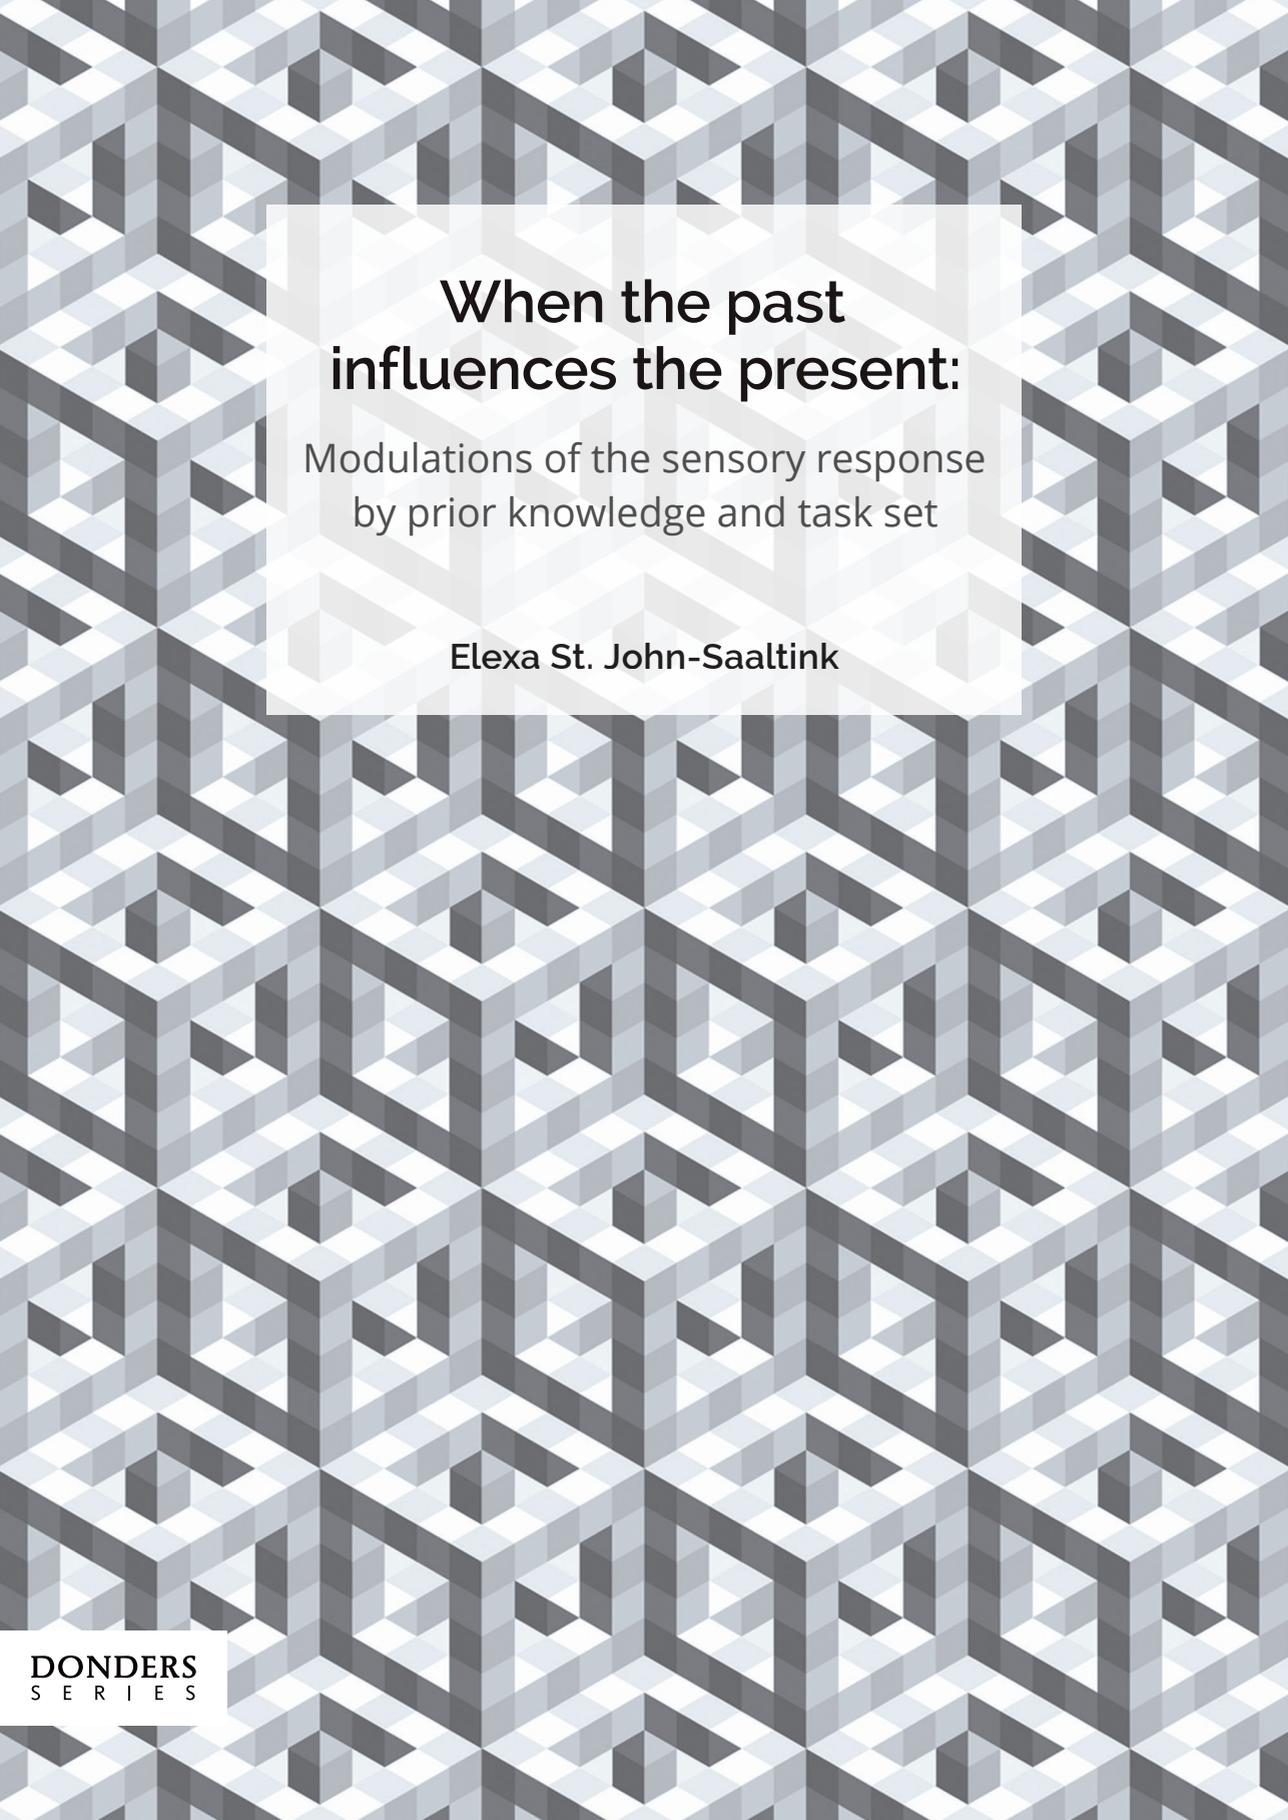
PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a publisher's version.

For additional information about this publication click this link.

<http://hdl.handle.net/2066/160914>

Please be advised that this information was generated on 2019-08-21 and may be subject to change.



When the past influences the present:

Modulations of the sensory response
by prior knowledge and task set

Elexa St. John-Saaltink

The work described in this thesis was carried out at the Donders Institute for Brain Cognition and Behaviour, Radboud University, with financial support from the Netherlands Organization for Scientific Research (NWO 404-10-037) awarded to Dr. Hakwan Lau and Dr. Floris de Lange.

ISBN

978-94-6284-083-6

Design/lay-out

Promotie In Zicht, Arnhem

Print

Ipskamp Printing, Enschede

© Elexa St. John-Saaltink, 2016

No part of this thesis may be reproduced, stored in a retrieval system or transmitted in any form or by any means without permission from the author or, when appropriate, from the publishers of the publications.

When the past influences the present:

Modulations of the sensory response
by prior knowledge and task set

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 5 december 2016
om 12.30 uur precies

door

Elexa St. John-Saaltink
geboren op 5 mei 1985
te Toronto, Canada

Promotor

Prof. dr. P. Hagoort

Copromotor

Dr. F.P. de Lange

Manuscriptcommissie

Prof. dr. R. Cools

Dr. H. Slagter (Universiteit van Amsterdam)

Prof. dr. L. Muckli (University of Glasgow, Verenigd Koninkrijk)

When the past influences the present:

Modulations of the sensory response
by prior knowledge and task set

Doctoral Thesis

to obtain the degree of doctor
from Radboud University Nijmegen
on the authority of the Rector Magnificus prof. dr. J.H.J.M. van Krieken,
according to the decision of the Council of Deans
to be defended in public on Monday, December 5, 2016
at 12.30 hours

by

Elexa St. John-Saaltink

Born on May 5, 1985
in Toronto, Canada

Supervisor

Prof. dr. P. Hagoort

Co-supervisor

Dr. F.P. de Lange

Doctoral Thesis Committee

Prof. dr. R. Cools

Dr. H. Slagter (University of Amsterdam)

Prof. dr. L. Muckli (University of Glasgow, United Kingdom)

Contents

Chapter 1	Introduction	9
Chapter 2	Serial dependence in perceptual decisions is reflected in activity patterns in primary visual cortex	19
Chapter 3	Expectation suppression in early visual cortex depends on task set	37
Chapter 4	Graded expectations: How does the brain weigh prior expectation with task-relevant sensory input?	55
Chapter 5	Does the weighting of prior expectations with sensory input depend on task?	73
Chapter 6	Discussion	89
	References	103
	Nederlandse samenvatting	109
	Acknowledgements	113
	List of Publications	119
	Biography	121
	Donders Graduate School for Cognitive Neuroscience	122

1

Introduction

Sight is an extremely important sense for humans. Relative to most other mammals, we rely on vision much more than on our other senses, such as smell and touch, and accordingly a large part of our brain is devoted to visual processing. In vision, the light that enters the eyes is converted into electrical signals by the retina, and then sent to the visual regions of the brain. Here, these electrical signals are processed by a hierarchy of brain regions that extract increasingly complex information, ultimately resulting in awareness of the visual scene. For decades, this sequence of events was thought to entirely capture the processes giving rise to vision: perception was seen as a passive process, in which the contents of perception were completely determined by a one-way stream that began with the incoming sensory signal. So long as the incoming sensory information was of sufficient strength, it would result in perception, and the contents of perception directly reflected the input to the eyes. However, visual illusions reveal that what we perceive is *not* always a veridical reflection of the world. To experience this, look at the centre blocks on the front and top faces of the cube in Figure 1.1A.

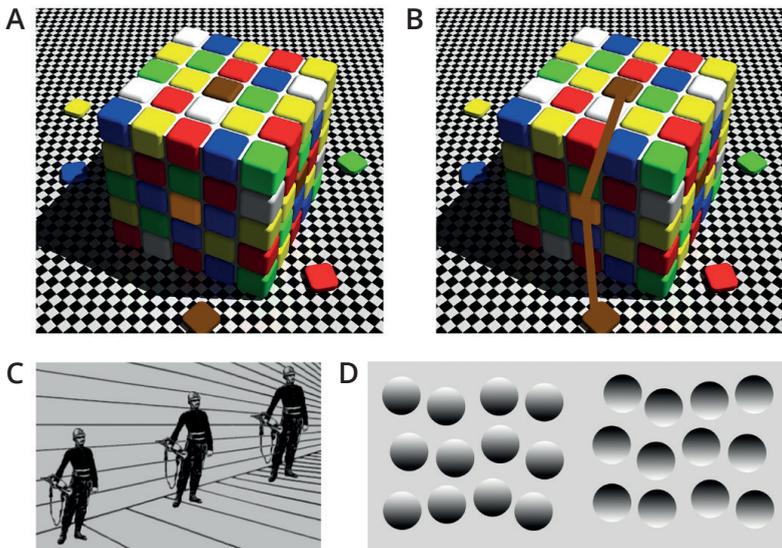


Figure 1.1. Perception is more than meets the eye.

A) Patches with an identical hue are perceived differently depending on inferred lighting conditions. **B)** Connecting the patches breaks the illusion. **C)** Simple depth cues cause identical objects to appear to differ in size. A prior that distance reduces an object's size causes the rightmost soldier to appear larger than his counterparts. **D)** A prior that light shines from above determines whether the circles are seen as protruding (convex) or receding (concave).

The square in the centre of the front face looks to be an orange colour, whereas the square in the centre of the top face looks to be a brown colour. In fact, these two squares are exactly the same colour, meaning that the wavelengths of light reflected by those squares are identical when they enter your eye. Since colour is highly dependent on lighting, mechanisms in your visual system correct for brightness context (Xing et al., 2015) in order to determine the true colour of the object. While this usually results in more accurate colour perception, this mechanism can at times lead to incorrect inferences about the true colour of an object. In this case, the illusion is so strong that visual proof – the squares connected by lines of the same colour (Figure 1.1B) – is required to accept it!

Modulations of perception by prior knowledge

Already over a century ago, Helmholtz questioned the idea of passive perception, casting it instead as an intrinsically inferential process wherein past experiences are used to make sense of incoming sensory information. Indeed, as just demonstrated in Figure 1.1A, our perception can often differ markedly from the information contained in the sensory signal alone. The knowledge that the brain has learnt about the world through past experience is thought to be contained in a collection of priors – ‘prior’ in the sense that in the absence of any sensory input, certain things are more likely to occur and therefore *a priori* expected. For example, we have learnt that distance causes things to appear smaller than they are up close, and due to this prior knowledge, our interpretation of size can be profoundly influenced by depth cues. In Figure 1.1C, you may experience the illusion that the rightmost soldier is larger than his companions – however all three soldiers are identical in size. The rightwards converging lines in the background of Figure 1.1C suggest that the rightmost soldier is farther away, which makes him appear larger. In another example, the circles on the left of Figure 1.1D probably appear as bumps that protrude out of the page towards you, whereas the circles on the right likely appear as dimples that recede into the page, away from you. However, the circles on the left are just 180° rotations of the circles on the right – if you turn this book upside down, the dimples will turn to bumps and vice versa! From experience in the world, we know that the most likely source of light is from above (i.e., the sun), and therefore developed a ‘light from above’ prior (Adams et al., 2004a) that gives rise to the illusion in Figure 1.1D: if light is shining from the top of the page downwards, then a convex circle would have a shadow at its top edge and a convex circle would have a reflection.

It is interesting to observe that knowledge of the ‘true’ visual input does not release the viewer from the illusory percepts displayed in Figure 1.1. You may find it

difficult, or impossible, to perceive the dots on the right of Figure 1.1D as bumps instead of dimples or to perceive all three solders as equally sized, despite knowing that they are! This demonstrates just how strong the influence of some priors on our perception can be. These examples demonstrate that perception is not a passive process, but instead that it is shaped by what we have learnt about the world.

Forming Expectations

While these examples may leave you with the impression that our brains are open to deception or trickery, applying prior knowledge to interpret incoming sensory information conveys large advantages. By constraining the hypothesis space, it enables the brain to form more efficient and accurate representations of the outside world. For example, when an object is presented in its usual context (i.e. a bicycle on a sidewalk) it is recognised more accurately and more quickly than when this context is scrambled (Biederman, 1972), and placing objects in unusual configurations or size relations to each other has the same impeding effect upon behaviour (Biederman et al., 1982). Making predictions about our environment improves processing efficiency and enables us to quickly infer properties that are ambiguous from the two-dimensional sensory input alone, as size, distance, and colour often are. Furthermore, it improves performance in the face of noisy or weak signals, such as the ability to navigate during a rainstorm or make our way through a darkened room, and allows us to infer the presence of objects that are difficult to detect or are occluded from view (and thus imperceptible from sensory information alone). Such inference is possible because past experience creates expectations about what sorts of things are likely to occur in the environment, which then can be used to guide and bias the interpretation of subsequent sensory input.

There are a number of ways in which the brain is able to form expectations about what is likely in the environment, all of which involve statistical learning at their core. Certain objects are more frequently encountered than others, and therefore more probable. If someone is pointing at the sky, it is more likely they are trying to draw your attention to an airplane or a bird than to a blimp or a bat. Other objects tend to come together, such that one predicts the presence of the other. When you hear a siren, you often see a vehicle with flashing lights moving quickly through the traffic. Finally, there is a high level of constancy in our visual scene, such that objects and features do not rapidly vary from one moment to the next. If you look up from this page to your surroundings, you will likely find that the majority of the scene has remained the same as the last time you looked at it – and

if you did find that an object had suddenly appeared or disappeared, this would be quite alarming! Based on the high degree of temporal autocorrelation in the environment, the brain forms a prior that the world is generally stable over short timescales. I investigate the perceptual and neural effects of this prior in **Chapter 2**.

Once these expectations have been formed, they modulate neural responses to sensory information at very early stages of processing. Neuroimaging experiments have demonstrated that expected items are processed with fewer resources than unexpected items (Murray et al., 2002; Alink et al., 2010; Meyer and Olson, 2011; Kok et al., 2012a). This reduction in neural activity to the same stimulus when it is expected compared to when it is unexpected is commonly referred to as *expectation suppression* (ES; Todorovic and De Lange, 2012). The reduced response to expected stimuli occurs whether expectations were induced via base rate occurrence (Summerfield et al., 2008; Larsson and Smith, 2012) or conditional probability (Den Ouden et al., 2009; Arnal et al., 2011; Kok et al., 2012a). However, whether or not such expectation suppression occurs depends on several factors, which I outline below.

Which factors influence whether the brain uses predictive information?

The brain is constantly forming hypotheses about the causes of its sensory input, and weighing the evidence for the different hypotheses against each other. Imagine that you are entering your bedroom late one evening, and discover that the light switch by the door does not work. As you go to turn on the bedside lamp, you startle – there is a cat curled up on your bed, but you have no pets! In the light, the ‘cat’ is revealed to be your discarded pyjamas, fallen in just the right way to resemble a curled body and tail. This type of hallucination is common, especially when light is low. The folds and bumps in the pyjamas and their position on the bed temporarily had more in common with previous experiences with cats than with clothing, and therefore offered more evidence for the ‘cat’ hypothesis, reflected in a brief hallucination. Like in the case of the illusions pictured in Figure 1.1, hallucinations result from the priors we use to make sense of the world. Importantly, the input supported this hypothesis only because of the poor lighting conditions – as soon as the light came on, the true cause of the input was revealed. This is in line with the idea that perception is most strongly influenced by prior knowledge when sensory input is noisy or ambiguous. In other words, when sensory input is clear, there is less need to rely on previous experience with the world in order to tell what we are looking at. However, when input is noisy or ambiguous, sensory expectations facilitate perceptual inference (Bar, 2004). Indeed, many theories of perception

hold that sensory input should be weighted according to its reliability (Mumford, 1992; Rao and Ballard, 1997; Knill and Pouget, 2004; Friston, 2009a), such that priors have a larger influence when sensory input is weak or ambiguous, but little or no influence when sensory input is strong and clear. In **Chapter 4**, I manipulated the strength of visual input in order to investigate whether there was evidence for an increased reliance on prior information when sensory information was weak, compared to when it was strong.

Apart from differences in the sensory signals themselves, the influence of priors may also depend on differences in the cognitive state of the observer. For example, when the stimulus/object that you are able to form expectations about (e.g., a disk flying through the air) is relevant to the task you are doing (e.g., playing Frisbee), the brain may capitalize on its ability to predict what is about to happen. Contrast this to the situation where you are having a picnic on the grass: here the brain may only make use of the ability to predict the irrelevant stimulus when there is a specific benefit of doing so (such as a need to suppress distracting information, or to duck when the disk comes flying at you). In general, most research on the effects of expectation on neural processing tends to focus on task-relevant stimuli, such that predictive cues signal features of the stimuli that participants need to respond to. I investigate whether an observer's specific requirements change how the brain responds to expected or predictable stimuli in the environment in **Chapters 3 and 5**.

Overview of this thesis

The aim of the work described in this thesis was to investigate the effects of expectation on sensory processing. I approached this question by probing the response of visual areas of the brain using functional magnetic resonance imaging (fMRI), a neuroimaging technique that provides an indirect measure of neural activity within different brain areas with high spatial specificity.

In **Chapter 2**, I examined how recent perceptual history influences how subsequent stimuli are perceived and processed by visual brain regions. Earlier research had demonstrated that perception is biased towards recently seen stimuli, such that perception is not fully determined by the currently presented stimulus, but also influenced by the stimuli that preceded it (Fischer and Whitney, 2014). This effect can be understood to result from a natural prior that the environment is generally stable over short timescales, given the high amount of autocorrelation commonly displayed by objects in the world. In a recent demonstration of this, Fischer & Whitney (2014) judge orientations of line segments, and found that participants perceived the orientations to be more similar to the orientation of

the previous stimulus than they really were. What is happening in the brain when perception is influenced by recently seen stimuli? Does the stimulus from the previous trial influence how the stimulus on the current trial is represented in visual areas of the brain? We addressed this question by measuring neural activity in visual cortex using fMRI while participants performed an orientation-judgment task. Our results suggest a direct neural correlate of stimulus history effects in visual cortex – to our knowledge, the first reported neural effect of perceptual history.

Chapter 2 focused on a natural prior – one that is learnt through normal experience with the autocorrelation present in people’s environment over the course of their lives. These long term, natural priors are still flexible, and can be adjusted through exposure to different environment properties (Adams et al., 2004b; Sotiropoulos et al., 2011). However, it is also possible to form new priors relatively quickly and flexibly, and **Chapters 3-5** investigated priors that were obtained in the laboratory. In these experiments, I induced simple, artificial priors by pairing the orientation of visual stimuli to auditory tones of specific frequencies. Participants saw stimuli containing oriented lines of one of two orientations, leftward tilted or rightward tilted. In **Chapter 3**, each visual stimulus was preceded by an auditory tone that either cued the orientation of the upcoming stimulus with 100% validity, or provided no information about the orientation of the upcoming stimulus (50% validity). This allowed me to compare neural activity in response to stimuli that were fully predicted to stimuli for which a precise expectation of orientation could not be formed. In **Chapters 4 and 5**, we made a slight adjustment to how we cued participants that allowed us to investigate what happens when predictions are violated: cues were 75% valid, meaning that a quarter of the time participants’ orientation expectation was invalid.

In **Chapter 3**, I investigated the effect of task-set on the neural response to predictable and non-predictable stimuli, to evaluate the automaticity of expectation effects. We used three tasks: one in which the cued stimulus was relevant, and two that diverted attention away from the cued stimulus by either a perceptual or working memory task on stimuli presented elsewhere on the screen. We found a reduced response to stimuli during predictable compared to non-predictable blocks during the distracting perceptual task, but not during the other two tasks. These results suggest that expectation suppression is not an automatic phenomenon, but dependent on attentional state and type of available cognitive resources.

In **Chapter 4**, I investigated the hypothesis that priors have a larger influence on sensory processing when sensory input is weak, as described in the example of the cat pyjamas, above. Here we manipulated the strength of the cued sensory input to investigate whether there is a larger effect of prior expectation on the

neural response in visual cortex when the strength of the sensory input is reduced. Our results supported this conclusion: expectation suppression was stronger for low contrast stimuli than high contrast stimuli. In other words, the effect of the prior was increased when the sensory input was weak, as would be expected when priors and sensory inputs are weighted by their relative reliability (Mumford, 1992; Rao and Ballard, 1997; Knill and Pouget, 2004; Friston, 2009a).

In Chapter 4, the cued stimuli were task-relevant. In **Chapter 5**, I studied the automaticity of the effect that was found in the previous chapter – that expectation had more weight when the sensory input was weak. The same participants completed a second fMRI session with the same cues and stimuli, however here they performed a distracting task and the cued stimuli were irrelevant. In this context, the task-irrelevant stimuli evoked much less activity in visual cortex, and this activity was not modulated by expectation. These results suggest that the modulation of sensory processing by expectation may not always occur automatically, but may be influenced by task demands.

Finally, in **Chapter 6**, I discuss how the empirical findings presented in this thesis contribute to our knowledge of how prior knowledge shapes the way the brain responds to its sensory environment.

2

Serial dependence in perceptual decisions is reflected in activity patterns in primary visual cortex

This chapter has been published as: St. John-Saaltink E, Kok P, Lau HC & De Lange FP. (2016) Serial Dependence in Perceptual Decisions is Reflected in Activity Patterns in Primary Visual Cortex. *Journal of Neuroscience* 36(23): 6186-92.

Abstract

Sensory signals are highly structured in both space and time. These regularities allow expectations about future stimulation to be formed, thereby facilitating decisions about upcoming visual features and objects. One such regularity is that the world is generally stable over short time scales. This feature of the world is exploited by the brain, leading to a bias in perception called serial dependence: previously seen stimuli bias the perception of subsequent stimuli, making them appear more similar to previous input than they really are. What are the neural processes that may underlie this bias in perceptual choice? Does serial dependence arise only in higher-level areas involved in perceptual decision-making, or does such a bias occur at the earliest levels of sensory processing? In this study, human subjects made decisions about the orientation of grating stimuli presented in the left or right visual field while activity patterns in their visual cortex were recorded using fMRI. In line with previous behavioral reports, reported orientation on the current trial was consistently biased towards the previously reported orientation. We found that the orientation signal in V1 was similarly biased towards the orientation presented on the previous trial. Both the perceptual decision and neural effects were spatially specific, such that the perceptual decision and neural representations on the current trial were only influenced by previous stimuli at the same location. These results suggest that biases in perceptual decisions induced by previous stimuli may result from neural biases in sensory cortex induced by recent perceptual history.

Introduction

The visual input we receive about the world is constantly interrupted by eye movements, blinks, and the occlusion of objects within our visual field. However, we perceive objects as continuous and the visual scene as stable. How is this stability obtained? One candidate mechanism for deriving stable representations from fluctuating noisy signals is temporal smoothing, i.e. the brain may generate a weighted average of current input with previously obtained input. This may be a beneficial strategy, given that the world is stable over short time scales (Dong and Atick, 1995). Indeed, perceptual judgments are known to be influenced by previous trial history (Gao et al., 2009; De Lange et al., 2013).

Historically, priming is a classic example of how a previously seen stimulus can alter the response to a subsequent stimulus. When stimuli are physically (or conceptually) repeated, the behavioural response is facilitated. A set of recent studies also demonstrated strong serial dependence of perception between temporally adjacent stimuli, even for reliable (suprathreshold) visual stimuli that varied randomly over time (Cicchini et al., 2014; Fischer and Whitney, 2014; Liberman et al., 2014; Rahnev et al., 2015). In particular, remarkably, these recent studies show that previous stimuli can change – in other words, *distort* – perception.

Here, we sought to clarify the neural mechanisms underlying this perceptual effect of recent stimulus history. On the one hand, it is conceivable that recent sensory input may change the sensitivity of sensory neurons, for example by increasing the sensitivity of neurons tuned to the previous input for a period of time following stimulus presentation (Fischer and Whitney, 2014). On the other hand, biases may only occur in downstream areas, at the stage of evidence accumulation and integration into a perceptual decision (Gold and Shadlen, 2007; Law and Gold, 2008; Hanks et al., 2011), leaving sensory processing unaffected. Here we examined whether serial dependence in visual perception is already manifest at the level of early sensory representations using an fMRI dataset in which subjects performed a perceptual decision task on the orientation of briefly presented grating stimuli. This allowed us to determine the influence of the previous stimulus on sensory representations in early visual cortex and perceptual report on the current trial. To preview, we found that sensory representations in early visual cortex were biased by the perceptual choice on the previous trial, in a spatially specific fashion. This suggests a potential sensory mechanism for serial dependence in visual perception.

Materials and Methods

Participants

Twenty-seven healthy right-handed individuals (17 females, age 22 ± 2 , mean \pm standard deviation (SD)) with normal or corrected-to-normal vision gave written informed consent to participate in this study. Three participants did not complete the full fMRI session due to poor task performance or poor fixation ability, therefore data from twenty-four participants were used for analyses. Experimental procedures were approved by the local ethics committee (Commissie Mensgebonden Onderzoek region Arnhem-Nijmegen, the Netherlands).

Stimuli

Stimuli consisted of two circular sinusoidal gratings (7° visual angle) presented for 200 ms at 5° along the horizontal meridian to the left and right of a central fixation point. Gratings were oriented at 45° and 135° independently of each other, and had the same orientation on 50% of trials. Grating orientations were pseudo-randomized such that all four orientation combinations (clockwise (CW)/CW; counter-clockwise (CCW)/CCW; CW/CCW; CCW/CW) occurred equally often. To increase task difficulty, grating stimuli were embedded in random noise at 80% contrast. The contrast of the grating within the stimulus was presented at two levels; the same contrast was used for both gratings within a trial. The stimuli (grating + noise) were normalized such that overall contrast and luminance were constant for high and low grating contrast values (all stimuli were 80% contrast). This meant that more of the overall stimulus contrast was driven by the grating for high contrast than for low contrast stimuli (mean grating contrast was $\sim 5\%$). The central fixation point was displayed on a gray background throughout the experiment. Stimuli were generated using MATLAB (MathWorks, Natick, MA, US) in conjunction with Psychophysics Toolbox (Brainard, 1997). In the fMRI session, stimuli were displayed on a rear projection screen using a luminance-calibrated EIKI projector (60 Hz refresh rate, 1024×768 resolution) which participants viewed through a mirror. Stimuli were displayed on a LCD monitor (60 Hz refresh rate, 1024×768 resolution) during the behavioural session.

Experimental Design

Stimuli were presented in an event-related design, with 5-7 s between trials. On each trial, a stimulus display of two gratings was briefly presented between two 500 ms periods of fixation (see Figure 2.1A). 700 ms after the onset of the stimulus display, two small chevrons pointing to the left or the right were presented for 1000 ms on either side of the fixation, which cued participants to respond to the grating that had been presented on that side of the screen. Participants performed a two-

alternative forced choice (2AFC) task on the orientation of the grating specified by the response cue, using two buttons (45°/135°) on an MR-compatible button box. Response contingency was counterbalanced across participants, and participants switched response hands halfway through the experiment.

Each trial also contained a pre-cue that consisted of two additional, smaller chevrons on either side of fixation that pointed in the same direction as the post-cue on 75% of trials. The pre-cue remained on screen during the initial fixation display and the stimulus display (i.e., 700 ms). On trials in which an orientation response was given, two small 'C's appeared on either side of the fixation point 500 ms after the offset of the response cue, to prompt participants to rate their perceptual confidence on a scale from one to four (using four buttons in their other hand). These aspects of the task were included to answer a different research question than the one addressed in the current paper. Therefore, in order to maximize the reliability of the estimation of the neural response, we collapsed over the congruency between the pre- and post-cues in all analyses presented here and for simplicity we display only the stimuli relevant to the current analysis in Figure 2.1.

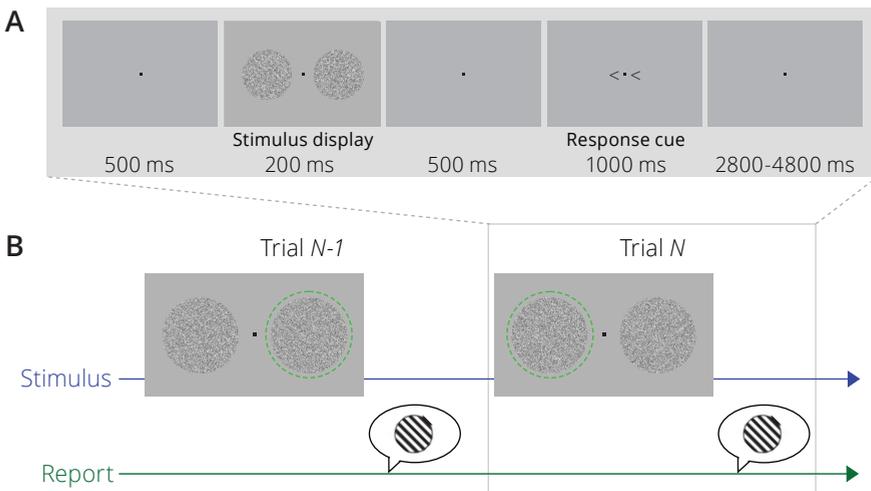


Figure 2.1. Experimental paradigm.

- A)** Every 5-7 s, two noisy gratings independently oriented at either 45 or 135 degrees were presented. Participants reported the orientation of the grating indicated by a response cue.
- B)** Analysis focuses on features of the current and previous trial.

Subjects participated in a behavioral session outside the scanner in the week before the fMRI session to familiarize subjects with the task and to titrate performance in the different conditions to 75% using a Quest staircase procedure (Watson and Pelli, 1983). During scanning, half of both congruent and incongruent trials were presented at low grating contrast, and the other half at high grating contrast, and contrast values were updated from the staircases after each block instead of after every trial. As mentioned above, we collapsed across cue-congruency and the different grating contrast levels in all analyses presented here.

The task was split into two runs of four blocks each, for a total of 512 trials. Between the two task runs, participants practiced remapping their response hand (the hand for grating and confidence response were switched) during the anatomical scan. Two additional scans were carried out after the main experiment. A functional localizer was collected to enable identification of voxels that were maximally responsive to the grating stimuli in the left and right hemifield, and a retinotopy scan to allow delineation of early visual cortices. The localizer consisted of full contrast gratings that were identical in size and position to those in the main experiment. Gratings were flickered at one of the stimulus locations per trial, alternating between left and right hemifield, at 2 Hz for 23.4 s. Each orientation (45° or 135°) was presented four times per location in a pseudo-random order. To ensure fixation, participants' task was to detect two letters ('X', 'Z') in a stream of letters within the fixation bull's-eye. During the retinotopy scan, a flashing black-and-white checkerboard pattern (3 Hz) in a 90 degree wedge rotated on a black background in 30 degree steps (1 position per TR). Participants' task was to detect unpredictable changes in the colour of the central fixation point (white to black). Nine cycles of clockwise and counterclockwise rotation were presented. During both additional scans, participants responded to target events with a button press.

fMRI Acquisition Parameters

Functional images were acquired using a 3T Trio MRI system (Siemens, Erlangen, Germany) using a 32-channel head coil, with a T2*-weighted gradient-echo EPI sequence (TR 1.95 s, 31 transversal slices, 3 x 3 x 3 mm in-plane resolution, TE 30 ms, field of view 192 mm x 192 mm, flip angle of 80). A high resolution anatomical image was collected using a T1-weighted MP-RAGE sequence (TR 2.3 s, TE 3.03 ms, 1 x 1 x 1 mm in-plane resolution, GRAPPA acceleration factor of 2).

fMRI Data Preprocessing

Data were pre-processed using SPM8 (<http://www.fil.ion.ucl.ac.uk/spm>; Wellcome Trust Centre for Neuroimaging, London, UK). The first three volumes of each task run were discarded to allow for time to achieve initial equilibrium. Functional images from the two task runs, the localizer, and retinotopy runs were spatially

realigned to the mean image, and temporally realigned to the first slice of each volume. The motion parameters resulting from spatial realignment were included as nuisance regressors in the general linear models. The structural image was coregistered with the functional volumes.

fMRI Data Analysis

An initial analysis of the functional localizer data was performed using SPM8, with regressors for left hemifield stimulation, right hemifield stimulation, and the motion parameters resulting from spatial realignment. A 128 s high-pass filter removed low-frequency signal components. Subtraction of the response to left and right hemifield stimulation was used to select stimulus-responsive voxels in each hemisphere for further analysis. In a separate analysis, Freesurfer (surfer.nmr.mgh.harvard.edu/) was used to inflate the cortical surface of each participant's T1-weighted structural image and to analyze the functional data from the retinotopy session. Polar-angle maps were generated using Fourier-based methods and projected onto the surface of the inflated cortex according to established methods (Serenio et al., 1995), allowing retinotopic areas within early visual cortex to be visually identified and delineated. Freesurfer and SPM functions were used to convert the retinotopic labels from surface to volume space and to transform them into regions of interest (ROIs).

Within retinotopic ROIs V1, V2, and V3 of each hemisphere, we restricted our analyses to the 50 voxels that were the most responsive to the localizer. To remove slow-drifts, preprocessed data from the localizer and task were linearly detrended. To estimate the response amplitude of each of these voxels to each single trial during the task, we applied the Least-Squares - Separate (LS-S) method outlined in Mumford et al. (2012) to the preprocessed data (Kok et al., 2013; Schoenmakers et al., 2014; Schlichting et al., 2015). This method consists of running a separate GLM for every trial, such that each trial is modeled once as a regressor of interest, with all other trials combined into a single nuisance regressor. This method has been shown to improve the estimation of single-trial BOLD response, compared to a GLM with one regressor for each trial (Mumford et al., 2012). In addition to these regressors, we included separate regressors for break and end of run screens, as well as the motion parameters resulting from spatial realignment, their derivatives, and the square of these derivatives (i.e., 18 motion parameters in total). The data from the functional localizer were analyzed similarly using the LS-S method, with one GLM performed per trial. The resulting task and localizer beta weights were normalized by z-scoring the values for each voxel.

For the main analyses, we first computed an orientation-specific signal for each trial by training a support vector machine (SVM) on the localizer data per hemisphere, and applying these SVMs to the task data in order to produce an SVM

decision value for each task trial. We used the SVM decision value as a proxy for orientation-signal strength. To maximize the strength of this orientation signal, we determined the optimal number of voxels for each participant and each ROI. To do so, we calculated the mean orientation signal over all task trials for different numbers of voxels (5 to 50, in steps of 5) and selected the number of voxels at which the mean orientation signal peaked. We applied Platt Scaling (Platt, 2000) to transform SVM outputs to probabilities by passing them through a sigmoid (Niculescu-Mizil and Caruana, 2005; Charles et al., 2014).

Serial Dependence Analyses (behaviour and fMRI)

For all following analyses, orientation was re-coded such that 45° was positive and 135° was negative. First, we constructed four regressors that captured, for each trial, the stimulus orientation at the responded to and non-responded to locations, on both trial N and trial $N-1$ (refer to Figure 2.2A). We applied logistic regression to participants' binary perceptual choice to characterize the impact of current and previous stimuli on current perceptual choice. Parameter estimates indicate the extent to which the perceptual report on trial N is influenced by the stimulus orientations presented at each location on the current and previous trial. To correspondingly characterize the impact of current and previous stimuli on the orientation signal in early visual areas, we applied linear regression to the orientation-specific BOLD signal (i.e., SVM output) in primary visual cortex (V1).

To investigate whether the serial effect is dependent upon the previous stimulus or the previous *percept* of that stimulus, we separately modeled the previous stimulus (the light blue bar in Figure 2.2B-C) depending on response: 'correct' when the previous percept was congruent with the previous stimulus; 'incorrect' when the previous percept was incongruent with the previous stimulus; and 'non-responded to' when the response was made to the stimulus at the other location. Using these three regressors in combination with the two regressors that captured the stimuli on the current trial, we again applied logistic regression to participants' binary perceptual choice, and linear regression to the orientation signal in V1.

For all analyses, we used simple t -tests at the group level to determine the robustness of each regressor's influence. To assess the location-specificity of the effects, we used paired-sample t -tests at the group level to compare regressors for reported and non-reported locations, separately for current and previous trial. Likewise, paired-sample t -tests at the group level were used to compare the strength of the bias following correct and non-responded to stimuli, and correct and incorrect stimuli. Finally, to assess whether higher order extrastriate cortex displayed comparable serial dependence, we ran the same linear regression analyses performed in V1, separately for V2 and V3 ROIs.

Control Analyses (fMRI)

Because we used an event-related design, with an average ITI of 6 s, a potential concern is that the serial dependence we find in the fMRI data could be attributed to BOLD from the previous trial that is yet to return to baseline. Although this concern is partially mitigated by the fact that all trials were modelled in the context of the GLM, which attributes only the unique variance to each regressor, we performed a conservative control analysis in which we explicitly modelled variations in hemodynamic effects of the previous trial during single trial beta estimation (i.e., the LS-S single trial GLMs, see above). Specifically, we captured any variation in the onset and duration of the BOLD response to the previous trial by modelling it with three complementary regressors: a canonical HRF and its first and second order derivatives. If the serial dependence effects in visual cortex reflect residual BOLD activity evoked by the previous stimulus, then the trial history effects should no longer be present when this signal is removed from the single trial estimates. We therefore repeated the two linear regression analyses presented above on these single trial estimates.

Persistence of serial dependence

In order to evaluate the temporal limit of serial dependence on the perceptual report, we investigated the influence of the preceding four trials on the current behavioural response. We constructed regressors capturing the stimulus orientation at the responded to and non-responded to locations on trial $N-2$, $N-3$, and $N-4$, and added them to the regressors for trials N and $N-1$ (refer to Figure 2.2A). We then repeated the logistic regression on participants' binary perceptual choice.

Results

We investigated an fMRI dataset collected while participants reported the orientation one of two visual grating stimuli briefly presented to the left and right of a central fixation point. On average, participants were 78% correct ($\pm 5\%$) and responded after 595 ms (± 119 ms), indicating that participants followed task instructions.

To characterize the impact of current as well as previous stimuli on perceived orientation, we applied logistic regression to participants' binary perceptual choice. Similarly, in order to investigate the impact of current and previous stimuli on the orientation signal in early visual cortical areas, we extracted orientation specific BOLD signals from visual cortex on every trial, and applied linear regression to the orientation-specific BOLD signal.

Perceived orientation was consistently biased towards the orientation of the preceding stimulus ($t_{(23)}=5.41$, $p = 1.7e-05$; Figure 2.2B). This serial dependence

effect was spatially specific: perceptual decisions were more strongly influenced by previous stimuli at the same location than by stimuli at a different location ($t_{(23)}=5.34$, $p = 2.0e-05$). In fact, perceived orientation on the current trial was slightly repelled away from the orientation of the previous stimulus at the other location ($t_{(23)}=-2.18$, $p = 0.040$).

Strikingly, the orientation signal in V1 was similarly biased towards the orientation presented on the previous trial ($t_{(23)}=3.21$, $p = 0.0039$; Figure 2.2C), which suggests

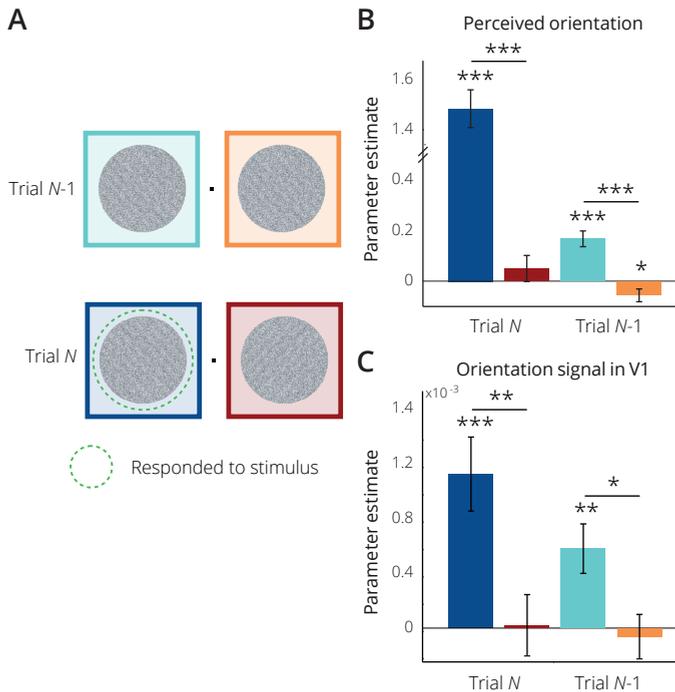


Figure 2.2. Serial dependence effects on perceptual choice and orientation signal in V1.

A) For each trial, four regressors captured the orientation of the stimuli at the responded to (blue) and non-responded to (red) locations, on both trial N and trial $N-1$. **B-C)** Parameter estimates for the four conditions indicated in **A**. Parameter estimates indicate how strongly perceptual report on trial N is influenced by the stimulus orientations presented at each location on the current and previous trial. **B)** Perceptual report on trial N is influenced by the orientation of the stimulus cued for report, but also by the previous stimulus presented at that location. **C)** The same as **B** but with the orientation signal in V1 (classifier output) as the dependent variable.

that recently seen stimuli alter the low-level sensory representations of subsequent stimuli. Again, this effect was retinotopically specific, with a stronger influence by previous stimuli at the same location than by stimuli at the other location ($t_{(23)}=2.63$, $p = 0.015$). V2 and V3 displayed the same pattern: the orientation signal was biased towards the orientation presented on the previous trial at the same location (V2: $t_{(23)}=2.08$, $p = 0.049$; V3: $t_{(23)}=2.88$, $p = 0.0085$). This effect was not present at the opposite location (V2: $t_{(23)}=0.65$, $p = 0.52$; V3: $t_{(23)}=0.26$, $p = 0.79$), although the difference between the two locations was not statistically significant in V2 ($t_{(23)}=0.63$, $p = 0.54$) and only approached significance in V3 ($t_{(23)}=2.02$, $p = 0.056$).

Is the serial effect on perceptual choice and neural representation dependent upon the previous *stimulus*, or instead upon the previous *percept*? Most often, perception follows the stimulus, precluding such an analysis. However, trials in which perception diverges from stimulus input offer an opportunity to tease these factors apart. In the following, we will only consider effects of stimuli presented on previous trials at the same spatial location as the currently responded to stimulus, given the spatial specificity of the serial dependence effect demonstrated above. To investigate this question, we separately modeled the previous stimulus on the basis of the on perceptual report, resulting in regressors for 'correct', 'incorrect', and 'non-responded to' trials (see Methods). As can be seen in Figure 2.3B, the perceptual decision on the current trial was consistently biased towards the previous stimulus when it was correctly perceived ($t_{(23)}=6.52$, $p = 1.2e-06$), but biased away from the physically presented stimulus towards the *perceived* stimulus when stimulus and perceptual choice diverged on the previous trial ($t_{(23)}=-4.98$, $p = 4.9e-05$), resulting in a significant difference between these conditions ($t_{(23)}=6.92$, $p = 4.7e-07$). When no explicit perceptual decision was made on the previous stimulus (e.g. on the current trial the left grating was responded to, but on the previous trial the right grating was responded to), there was still a strong bias towards the previous stimulus ($t_{(23)}=4.91$, $p = 5.8e-05$), ruling out an explanation of response bias. The bias towards the previous stimulus was stronger when it was correctly reported than when it was not responded to ($t_{(23)}=2.48$, $p = 0.021$).

The orientation signal decoded from BOLD activity in V1 displayed a similar profile, with a significant bias of orientation signals towards the previous stimulus when the previous stimulus was correctly perceived ($t_{(23)}=2.38$, $p = 0.026$) or passively viewed ($t_{(23)}=2.30$, $p = 0.031$) but no reliable effect when stimulus and choice diverged ($t_{(23)}=0.40$, $p = 0.69$). This however did not culminate in a significant difference between correctly versus incorrectly responded trials ($t_{(23)}=0.51$, $p = 0.61$), potentially due to the higher variability of the neural orientation signal compared to the behavioural report, in combination with the relatively low number of error trials (23% on average). There was no difference in the strength of the bias following correctly reported compared to non-reported stimuli ($t_{(23)}=0.71$, $p = 0.48$). Again, a

similar pattern of results was found in V2 and V3, albeit non-significantly in V2. Orientation signals were biased towards previous stimuli when they were correctly perceived (V2: $t_{(23)}=1.05$, $p = 0.31$; V3: $t_{(23)}=2.37$, $p = 0.027$) or passively viewed (V2: $t_{(23)}=1.51$, $p = 0.14$; V3: $t_{(23)}=2.45$, $p = 0.023$), with no reliable effect when stimulus and choice could be dissociated (V2: $t_{(23)}=0.51$, $p = 0.62$; V3: $t_{(23)}=1.54$, $p = 0.14$).

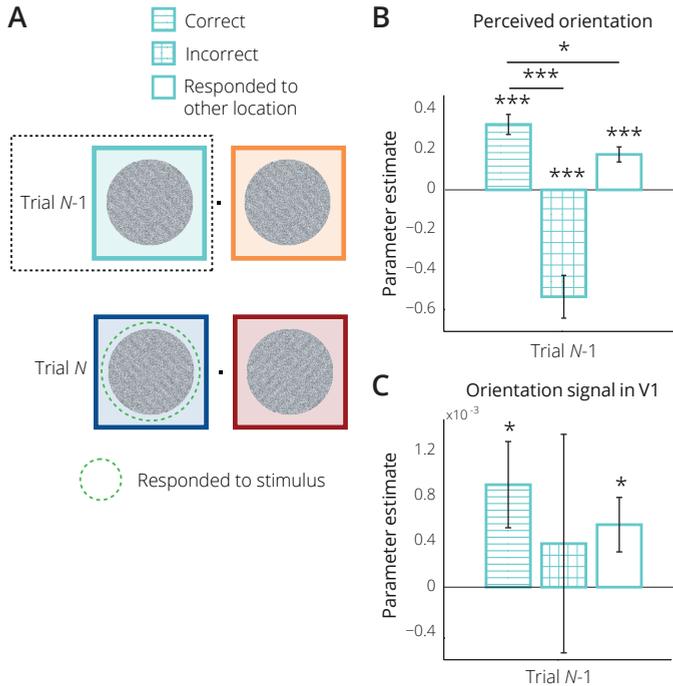


Figure 2.3. Serial dependence is governed by previous percept rather than stimulus or response.

A) The regressor for the stimulus on trial $N-1$ at the responded location on trial N (light blue in Figure 2.2B-C) was sub-divided on the basis of the perceptual decision on trial $N-1$: correct response (horizontal stripes), incorrect response (grid), responded to other location (no fill). **B-C)** Parameter estimates for the three conditions outlined in **A**. **B)** Previous stimuli that were correctly reported or not responded to exert a positive bias on current perception. For incorrect trials, perception is biased towards previous percept rather than stimulus. **C)** The same as **B**, but with the orientation signal in V1 (classifier output) as the dependent variable. Here only the positive biases are present.

To ensure that our neural results were not dependent on residual BOLD signal from the evoked response to the previous stimulus, we modeled the neural response to the previous trial during the single trial beta estimation. If the serial dependence effects in visual cortex reflect left over BOLD activity evoked by the previous stimulus, then the trial history effects should no longer be present when this signal is removed from the single trial estimates. However, there was still a reliable effect of the previous stimulus on the current neural response in V1 ($t_{(23)}=2.38, p = 0.026$).

To investigate how long the influence of previous trials persists, we looked at the influence of stimuli from the four preceding trials on the current perceptual report. There is a consistent positive bias towards the orientation of preceding stimuli from the preceding three trials ($N-2: t_{(23)}=5.10, p = 3.7e-05; N-3: t_{(23)}=5.04, p = 4.2e-05$), specific to stimuli at the location that is responded to on trial N ($N-2: t_{(23)}=5.19, p = 2.9e-05; N-3: t_{(23)}=3.90, p = 0.00072$; Figure 2.4). This bias drops off after three trials ($N-4: t_{(23)}=1.41, p = 0.17$). Serial dependence of the orientation signal in visual cortex does not persist beyond the directly preceding trial ($N-2: t_{(23)}=0.77, p = 0.45$), possibly due to the fact that the fMRI orientation signals were generally noisier than behaviour.

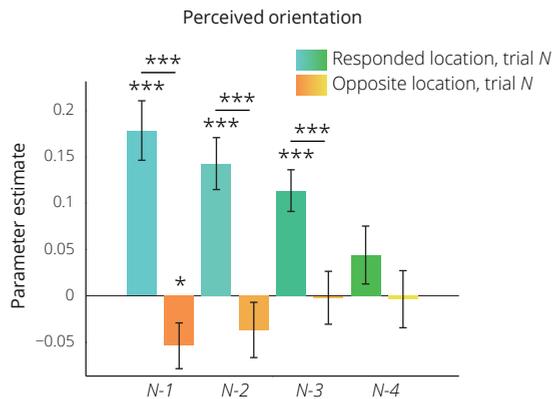


Figure 2.4. The influence of previous trials on perceptual choice persists for three trials.

For this analysis, the orientation of the stimuli at the responded to (blue-green) and non-responded to (orange-yellow) locations on trial $N-2, N-3$ and $N-4$ were added to the regressors in Figure 2.2A. The positive bias of previous stimuli on the current perceptual report persisted for three trials.

Discussion

Our perception at any given moment is influenced by both current and previous sensory signals. In this study, we investigated the neural mechanism underlying this serial dependence in perceptual decisions. Through extracting orientation-specific signals from visual cortex, we determined that the attractive bias exerted by the previous percept is present at the level of early sensory representations. This result sheds light on the mechanism behind the serial dependencies that have been reported in the literature, such as that choice on a current trial is influenced by the directly preceding stimulus (Gao et al., 2009; De Lange et al., 2013; Fischer and Whitney, 2014).

What specifically is carried over from one trial to the next? We distinguish three possibilities. First, serial dependence could be driven by the previous stimulus, via the bottom-up signal. Another option is that the perceptual choice made *on the basis of* the previous stimulus is carried over to the next trial. The crucial distinction according to this explanation is that serial dependence is more attributable to the percept of the previous stimulus than the previous stimulus *per se*. The third possibility is that serial dependence relies on the behavioural response that is coupled to the perceptual choice. In this explanation, the serial effect is motoric (as opposed to sensory or perceptual) in nature.

If serial dependence is driven by the bottom-up input, then there should be a consistent influence of previous stimuli based on their orientation. This is indeed what we found: there was a reliable bias towards the orientation of previous stimulus presented at the same location. However, perceptual choice is highly correlated to sensory input, i.e., generally our perceptual experience of the world is a good reflection of the stimulation received by our sensory cortices. This makes it difficult to determine whether it is the bottom-up signal, or the perceptual choice, that carries over across trials. Because the orientation task used here was deliberately difficult (the contrast of the orientation signal within the noisy stimulus patches was titrated such that participants were 75% correct; the gratings were on average 5% contrast in 80% contrast noise), it resulted in a proportion of trials in which stimulus and choice diverged. These trials allow us to dissociate perceptual choice from the sensory input. On these trials, subsequent perceptual decision was biased towards the previous (incorrect) perceptual choice, instead of the stimulus that was presented (see Figure 2.3B). Furthermore, the attractive bias elicited by the previous trial was stronger following a correct perceptual choice than when the stimulus was not responded to. This suggests that *perception* – as opposed to sensory input *per se* – is what is carried over across time, exerting a positive bias on subsequent perceptual decisions.

If, on the other hand, the behavioural response (i.e., the button that was pressed) is carried over between trials, one would expect serial dependence to be location *unspecific* – i.e., if pressing button 1 would prime a subsequent button 1 response, this would not be localized to one visual hemifield, but would instead transfer across stimulus locations (as response location was randomized over trials). However, all effects of previous stimuli on both perceptual report and orientation signal in this study were location specific (see Figure 2.2B), which rules out that simple motor response biases were responsible for our results. For the same reason, this rules out non-specific decision-biases, such as a predisposition to report the same orientation repeatedly, and instead suggests a spatially specific perceptual carry-over.

The current results are in line with other recent reports suggesting a perceptual nature of serial dependence (Burr and Cicchini, 2014; Cicchini et al., 2014; Fischer and Whitney, 2014; Liberman et al., 2014). It should be noted that perceptual decisions in this experiment were measured using a 2AFC task, which can be contrasted with the more continuous measures of perception used in these previous reports. Given that it is necessarily less fine-grained, a binary measure of perception may be less sensitive to subtle perceptual biases. However, because the oriented gratings in our stimuli were low contrast, embedded in high contrast white noise which contains signal for all possible orientations, it is feasible that small (serial dependence) biases would lead to the false perception of the orthogonal grating orientation (Pajani et al., 2015). Since in the current study participants were presented with only two (orthogonal) orientations, the biases reported here likely reflect carry-over of perceptual decisions about whether the grating was oriented clockwise (45 degrees) or counterclockwise (135 degrees), rather than subtle perceptual biases on the order of a few degrees such as those previously described in continuous-report designs (Fischer and Whitney, 2014). A paradigm which combines a continuous measure of perception with a neuro-imaging measure of stimulus representations in low-level sensory cortices, such as one in which neural correlates of subtle perceptual biases have previously been measured (Kok et al., 2013), could be a promising avenue for future investigation of serial dependence.

What may be the neural mechanisms underlying the serial dependence of perceptual choice? One hypothesis is that the sensitivity of sensory neurons tuned to the previous percept may be increased for a brief period following stimulus presentation, thereby influencing current perception (Fischer and Whitney, 2014). In line with this idea, we found that the orientation of the previous perceptual decision biases the representation of stimuli in early visual cortex. Interestingly, similar effects of the decision variable on sensory responses have been observed during the period in which a perceptual decision unfolds (Nienborg and Cumming,

2009; Wimmer et al., 2015). We speculate that this biasing of sensory responses due to the decision may persist, thereby biasing subsequent sensory processing.

Notably, we found that non-reported gratings also have an influence on the subsequent perceptual decision and orientation signal in visual cortex. While this result could be seen as evidence for the stimulus-driven account of serial dependence, it should be borne in mind that non-reported gratings were still attentively perceived by the participants. Namely, subjects were only informed about which stimulus to report *after* the stimuli had been removed from the screen, therefore necessitating an implicit perceptual decision about both stimuli.

The effects in primary visual cortex and in perceptual report were spatially specific: perceptual decisions and neural representations on the current trial were only influenced by previous stimuli at the same location. This was equally true of the influence that stimuli from two and three trials back had on the perceptual report. This location-specificity may appear at odds with a previous report (Fischer and Whitney, 2014), in which serial dependence transferred across spatial locations (i.e., that serial dependence smoothes across time *and* space). One explanation for this could be due to differences in participants' attentional state between the two designs: when serial dependence was found to transfer across locations, only one stimulus was attended on each trial, whereas in the current study, both stimuli required a certain level of attention since either stimulus could be cued for report. Future designs that manipulate both the number of attended stimuli and total number of stimuli may elucidate the cause of this discrepancy.

A potential limitation of the present study is the possibility that the serial dependence we measured in visual cortex is a result of a residual BOLD response to the previous stimulus. However, there are several reasons why we believe that our results are not the result of autocorrelation in the BOLD signal. First, in order to increase the separability of single trials, the interval between trials was jittered such that trials were presented every 5-7 s. Second, each trial was independently modeled using a technique which maximizes extraction of the signal unique to each trial (Mumford et al., 2012). This approach to estimating the trial-specific BOLD signal, in combination with the jitter between trials, should allow the signal from individual trials to be dissociated from neighbouring trials. Furthermore, the serial dependence in visual cortex persisted even after we applied a conservative approach to regress out the BOLD response to the previous trial. Therefore, it appears plausible that the bias in sensory cortex is generated at the moment of bottom-up stimulation, rather than reflecting a spill-over of activity to the previous stimulus.

Our perception of the world is partly determined by our (often implicit) priors about the statistical regularities in the environment (Yuille and Kersten, 2006; Chalk et al., 2010; Kok et al., 2013). Serial dependence, such as reported here, can

be understood as one such prior – that the world is stable over short time scales. Interestingly, serial dependence is not restricted to low-level stimuli such as used here, but also extends to complex and naturalistic stimuli such as faces (Lieberman et al., 2014) and numerosity (Cicchini et al., 2014). Other sequential effects such as repetition suppression may similarly be cast as the result of this same prior (Summerfield et al., 2008; Todorovic et al., 2011; Henson, 2015): if the world is generally stable, objects are more likely to repeat than change. Given the stability of the sensory world, such a prior could make visual processing more robust by filtering out temporal noise.

3

Expectation suppression in early visual cortex depends on task set

This chapter has been published as: St. John-Saaltink E, Utzerath C, Kok P, Lau HC, & De Lange FP (2015) Expectation suppression in early visual cortex depends on task set, *PLoS ONE* 10(6), e0131172.

Abstract

Stimulus expectation can modulate neural responses in early sensory cortical regions, with expected stimuli often leading to a reduced neural response. However, it is unclear whether this expectation suppression is an automatic phenomenon or is instead dependent on the type of task a subject is engaged in. To investigate this, human subjects were presented with visual grating stimuli in the periphery that were either predictable or non-predictable while they performed three tasks that differently engaged cognitive resources. In two of the tasks, the predictable stimulus was task-irrelevant and spatial attention was engaged at fixation, with a high load on either perceptual or working memory resources. In the third task, the predictable stimulus was task-relevant, and therefore spatially attended. We observed that expectation suppression is dependent on the cognitive resources engaged by a subjects' current task. When the grating was task-irrelevant, expectation suppression for predictable items was visible in retinotopically specific areas of early visual cortex (V1-V3) during the perceptual task, but it was abolished when working memory was loaded. When the grating was task-relevant and spatially attended, there was no significant effect of expectation in early visual cortex. These results suggest that expectation suppression is not an automatic phenomenon, but dependent on attentional state and type of available cognitive resources.

Introduction

Stimulus expectation can modulate neural responses in early sensory cortical regions, with expected stimuli often leading to a reduced neural response (Murray et al., 2002; Meyer and Olson, 2011). This effect has been found in visual (Summerfield et al., 2008; Meyer and Olson, 2011; Kok et al., 2012a) and auditory (Arnal et al., 2011; Todorovic et al., 2011) cortices, and in both electrophysiological (Meyer and Olson, 2011; Todorovic et al., 2011) and haemodynamic (Summerfield et al., 2008; Den Ouden et al., 2009; Kok et al., 2012a) measurements.

Is expectation suppression an automatic process that happens outside the focus of attention? Several studies suggest that this is the case. A reduced neural response for predictable stimuli has been found during passive viewing (Alink et al., 2010), as well as when stimuli are fully task irrelevant (Den Ouden et al., 2009), supporting the idea that suppression occurs automatically, whenever sensory input is predictable. In contrast to this notion however, other authors found no effect of expectation on sensory activity when stimuli were unattended (Larsson and Smith, 2012), suggesting that expected background stimuli are not automatically suppressed.

One potential explanation for these conflicting results could be that the specific task set a subject is engaged in, and thereby the available resources, may determine whether stimulus expectations modulate the sensory response. For example, load theory (Lavie et al., 2004) states that the processing of task-irrelevant stimuli is determined by the type and level of resource load posed by a given task. In line with this, previous research has shown that task set can have a profound effect on the extent to which items in the visual background are processed (Yi et al., 2004; Lavie, 2005): irrelevant background stimuli tend to be suppressed during tasks that load perceptual resources, whereas there is no suppression for background stimuli when working memory is taxed. Therefore, expectation suppression of background stimuli might be especially pronounced during tasks with a high perceptual (but not working memory) load. Critically, in support of this idea, the study which found that predictable background stimuli were not suppressed by expectation used a paradigm that loaded working memory resources (Larsson and Smith, 2012).

Therefore, in the current study we asked whether expectation suppression for background stimuli depends on the type of available resources. To investigate this, we compared the neural response to predictable and non-predictable visual background stimuli during tasks that placed a higher load on either perceptual or working memory resources. If task set interacts with the expectation effect, this would indicate that expectation suppression is dependent on how processing resources are engaged by the task at hand. Conversely, if no such interaction is present, this would indicate that expectation suppression is independent of the

type of available resources. To compare the effect of spatial attention, we included a task that made the background stimulus task-relevant. Due to previous reports of expectation suppression for task-relevant, spatially-attended stimuli, we hypothesized that predictable stimuli would be suppressed relative to non-predictable stimuli during this task. With this design we aimed to elucidate the conditions for which sensory input is suppressed by expectation.

Materials and Methods

Participants

Thirty-five healthy right-handed individuals (25 females, age 22 ± 4 , mean \pm standard deviation (SD)) with normal or corrected-to-normal vision gave written informed consent to participate in this study. Experimental procedures were approved by the local ethics committee (Commissie Mensgebonden Onderzoek region Arnhem-Nijmegen, the Netherlands). Data from two subjects were excluded due to chance level performance on one or more of the tasks.

Stimuli

Stimuli were generated using MATLAB (MathWorks, Natick, MA, US) in conjunction with Psychophysics Toolbox (Brainard, 1997). In the behavioural session, stimuli were displayed on a Samsung SynchMaster 940BF monitor (60 Hz refresh rate, 1280 X 1024 resolution). In the fMRI session, stimuli were displayed on a rear projection screen using a luminance-calibrated EIKI projector (60 Hz refresh rate, 1024 X 768 resolution) which participants viewed through a mirror. A fixation “bull’s-eye” (outer ring 0.8° of visual angle) was presented at the centre of a gray background throughout each task. On each trial, a grating annulus (outer diameter: 15° of visual angle; inner diameter: 2°) of luminance-defined sinusoids at 80% contrast was displayed around the fixation bull’s-eye (200 ms; 3.31A). Gratings were oriented at either 45° or 135° , with a phase randomly selected from 10 possibilities, evenly spaced between π and 2π . To mitigate afterimages, the phase of the grating was inverted halfway through stimulus presentation. On each trial, gratings had one of two possible spatial frequencies (mean: 1.5 cpd), with the specific spatial frequency values set for each individual by a staircasing procedure (see below). Simultaneously with the grating presentation, coloured letters were presented in the centre of the fixation bull’s-eye, with noise of the same colour superimposed (Figure 3.1B). Six letters (A, H, N, R, T, Z) and six colours (red, blue, green, cyan, yellow, magenta) were used. The number of coloured pixels that degraded the letters was similarly set by a staircasing procedure. The titration of these stimulus parameters for each subject ensured that tasks were matched on

difficulty. Auditory cues that were played before each stimulus consisted of four pure tones (329, 440, 493 and 659 Hz) that were played for 200 ms.

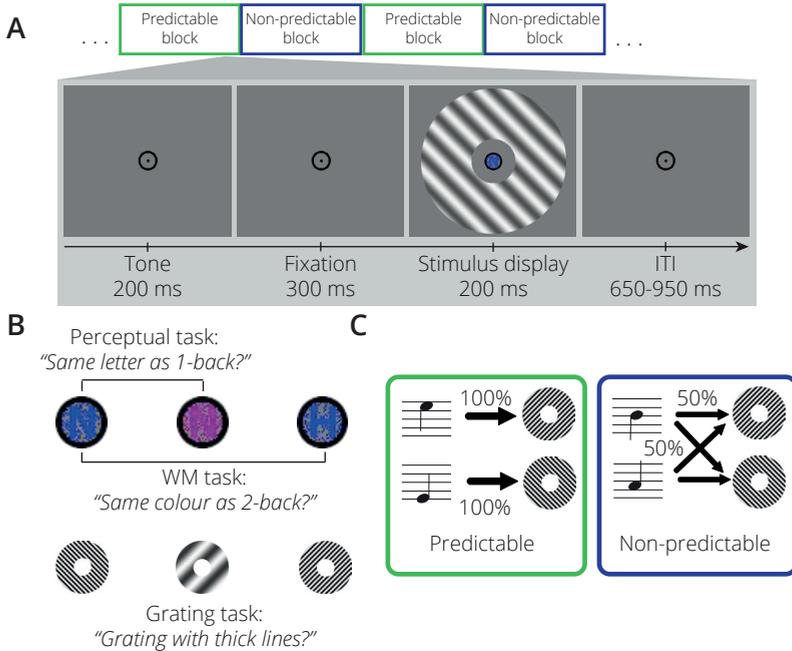


Figure 3.1. Experimental paradigm.

A) During each of the three tasks, stimuli were presented in predictable and non-predictable blocks, alternating every 12 trials. Each trial started with an auditory tone that either perfectly predicted the orientation of the subsequent grating stimulus (45° or 135°), or carried no orientation information. **B)** Participants performed two tasks on the stimuli in the fixation bull's-eye. During the perceptual task, targets were 1-back letter repetitions (the magenta 'N' is a target) that were difficult to perceive due to added noise. During the working memory task, targets were 2-back colour repetitions (the blue 'H' is a target) that were easy to perceive (as the whole inner ring has the same colour) but taxed the working memory system more strongly than the 1-back task. During the grating task, participants responded to the spatial-frequency of the grating stimuli. Targets had lower spatial-frequency than non-targets. **C)** Pairing between auditory tones and grating orientations. In predictable blocks, tones predicted grating orientation with 100% accuracy. In non-predictable blocks, the tones provided no orientation information. ITI = inter-trial interval.

Experimental Design

To maximize sensitivity to the effects of predictability, stimuli were presented in a block design. On each trial, an auditory tone was presented 500 ms before the onset of the visual stimulus display. These auditory tones contained predictive information about upcoming visual stimuli, similar to previous studies (Den Ouden et al., 2009; Den Ouden et al., 2010; Kok et al., 2012a; Kok et al., 2013; Kok et al., 2014). During predictable blocks, the auditory tones cued the orientation of the upcoming grating with 100% validity. During non-predictable blocks, the auditory tones provided no orientation information (50% validity; Figure 3.1C). Then, the stimulus display was presented, which consisted of both the grating annulus and the degraded coloured letter at fixation. The stimulus display remained on screen for 200 ms (see Figure 3.1A), followed by an inter-trial interval jittered between 650 and 950 ms. We reasoned that including this jitter between trials would increase the pairing between tones and gratings *within* trials (which had a fixed temporal distance of 500 ms), relative to stimuli *across* trials. Hereby we aimed to increase the salience of the predictable statistical structure of the tone-orientation pairings. Cue validity alternated every 12 trials (18 s), with one trial presented on average every 1.5 s (range 1.35-1.65 s). Using cross-modal cues to manipulate conditional probabilities has advantages over inducing expectations by manipulating base rate occurrence of stimuli, since the latter could cause differences in stimulus-specific adaptation (Krekelberg et al., 2006; Solomon and Kohn, 2014).

We manipulated task set while diverting spatial attention away from the grating stimuli by loading either perceptual resources or working memory resources at the fovea (see Figure 3.1B). Following the study of Yi et al. (2004), we induced a high working memory load by a 2-back task on easy to perceive stimuli, whereas a high perceptual resource load was induced by a 1-back task on visually degraded stimuli. Instead of manipulating the level of resource load, we elected to use only high load conditions with carefully titrated stimuli to match performance, so as to facilitate a direct comparison between the different types of load (rather than the amount of load). During the perceptual task, participants responded to the degraded letters within the fixation bull's-eye: a letter was a target if it was the same as during the previous stimulus display (1-back task). This task was perceptually challenging because the letters were occluded by noise of the same colour, making them difficult to identify. During the working memory task, participants responded to the colour of the letters and noise pixels within the fixation bull's-eye: a colour was a target if it was the same as the stimulus display presented two trials ago (2-back task). The 2-back task required participants to encode, maintain and compare more items in working memory than the 1-back task, and thereby it placed a higher load on working memory. We also included a task in which attention was directed to the grating. During this task, participants responded to the spatial-frequency of

the grating stimuli: low-spatial frequency gratings were targets. Therefore, the predictable grating stimuli could be either unattended with a load on perceptual resources (letter task); unattended with a load on working memory resources (colour task); or attended and task-relevant (grating task). For all tasks, participants were instructed to respond after every trial by indicating 'target' (button 1) or non-target (button 2). Each task had the same proportion of targets (33%), and participants were instructed to maintain fixation during all tasks.

Per scanner run of 5.4 min, participants performed 14 blocks of one of the tasks, with four fixation blocks interspersed. During fixation blocks, only the bull's eye was presented and participants were instructed to maintain fixation. Stimulus sequences were repeated for three runs such that each task was performed on the same stimuli. Participants cycled through the tasks until each had been performed three times, yielding 504 trials per task. Task order was counterbalanced across participants. Feedback (target accuracy (%), number of timeouts to targets, non-target accuracy (%), and number of timeouts to non-targets) was displayed for 2 s at the end of every run.

After the main experiment, we carried out two additional scans: one to identify voxels that were maximally responsive to the grating stimulus and one to retinotopically delineate early visual cortices. During the grating localizer, full contrast gratings of the same size and position as the main experiment were presented. Gratings were flickered at 2 Hz for 14.4 s at eight orientations (22.5, 45.0, 67.5, 90.0, 112.5, 135.0, and 157.5 degrees). Each orientation was presented four times in a pseudo-random order. To ensure fixation, participants' task was to detect two letters ('X', 'Z') in a stream of letters within the fixation bull's-eye. During the retinotopy scan, a 90 degree wedge stimulus consisting of a flashing black-and-white checkerboard pattern (3 Hz) rotated in 30 degree steps (1 position per TR) on a black background. Nine cycles of clockwise and counter clockwise rotation were presented. Participants' task was to detect unpredictable changes in the colour of the central fixation point (white to black), which occurred four to eight times per 36 s block. During both scans, participants responded to target events with a button press.

Behavioural Session

To familiarize participants with the tasks and to ensure all tasks were equally difficult, stimuli were calibrated prior to the scanning session. Participants were explicitly taught the relationships between the informative tones and the grating orientations, and their knowledge of these relationships was explicitly tested at the beginning and the end of the training session as well as directly preceding the fMRI session. Participants were trained on the working memory task until they could achieve better than 60% target accuracy on one run of this task (using identical

presentation parameters to the fMRI session). Participants who could not achieve this performance were excluded from further participation. The stimuli from the perceptual and grating tasks were adjusted using an adaptive staircase-procedure (Watson and Pelli, 1983) set to the accuracy that the participant attained on the working memory task. First, accuracy on the perceptual task was manipulated by adjusting the number of coloured noise pixels surrounding the letter. Then, accuracy on the grating task was manipulated by adjusting the difference between high and low spatial-frequency gratings (around a mean of 1.5 cpd). During both tasks, stimulus timing was identical to during the fMRI experiment, except that participants had the opportunity to pause every four blocks while stimulus parameters were updated according to the staircase value. This continued until the target accuracy was reached. This procedure matched performance on the different tasks within participants, while allowing differences in overall performance between participants. The behavioural session was held within seven days prior to the scanning session. Just before entering the scanner, participants were exposed to the relationship between the informative tones and the grating orientations with 24 practice trials.

To determine whether expectation or task had an effect on behaviour during the fMRI session, we performed a repeated measures ANOVA (rmANOVA) with factors expectation (predictable; non-predictable) and task (perceptual; WM; grating) on accuracy and RT. A main effect of task on RT was further investigated using post-hoc t-tests. Finally, to probe whether expectation had an effect when the predictable stimulus was attended and task-relevant, we compared accuracy and RT during predictable versus non-predictable blocks for the grating task using post-hoc t-tests.

fMRI Acquisition and Analysis

Functional images were acquired using a 3T Trio MRI system (Siemens, Erlangen, Germany) using a 32-channel head coil, with a 3D EPI sequence (TR 1.8 s, 64 transversal slices, 2 x 2 x 2 mm in-plane resolution, TE 25 ms, field of view 224 mm x 224 mm, GRAPPA acceleration factor of 2). A high resolution anatomical image was collected using a T1-weighted MP-RAGE sequence (TR 2.3 s, TE 3.03 ms, 1 x 1 x 1 mm in-plane resolution, GRAPPA acceleration factor of 2).

Data were pre-processed using SPM8 (<http://www.fil.ion.ucl.ac.uk/spm>; Wellcome Trust Centre for Neuroimaging, London, UK). The first four volumes of each task run were discarded to allow for time to achieve initial equilibrium. Functional images from all sessions were spatially realigned to the mean image, and the resulting movement parameters, their first order derivatives and the square of these derivatives were included as nuisance regressors in the general linear model (GLM). The structural image was coregistered to the functional volumes.

Functional data from each subject were modelled using a block design approach, within the context of the General Linear Model (GLM) that included the data from all nine task runs. A 128 s high-pass filter removed low-frequency oscillations. Regressors for the task conditions were specified per run (180 scans), and convolved with SPM8's canonical hemodynamic response function, which is comprised of the sum of two gamma functions. In addition to task regressors (predictable, non-predictable and fixation), the motion parameters as described above and a regressor to capture adaptation effects were included. The resulting beta-weights for predictable, non-predictable, and fixation for each run were analyzed in Matlab using in-house code. For each run, the fixation beta weight was used to normalize the beta-weights for predictable and non-predictable, and then the normalized betas were averaged across the three runs of each task. A separate design matrix was constructed for the grating localizer data, with regressors for stimulation, fixation, and the motion parameters (with derivatives as discussed above).

Freesurfer (surfer.nmr.mgh.harvard.edu/) was used to inflate the cortical surface of each participant's T1-weighted structural image and to analyze the functional data from the retinotopy session. Polar-angle maps were generated using Fourier-based methods and projected onto the surface of the inflated cortex according to established methods (Serenó et al., 1995), allowing retinotopic areas within early visual cortex to be visually identified and delineated. Freesurfer and SPM functions were used to convert the retinotopic labels from surface to volume space and to transform them into regions of interest (ROIs).

Within these retinotopic ROIs of V1-V3, we averaged the task-related activity of the 150 voxels that were most responsive to the grating stimulus during the localizer. To test whether the effect of expectation was influenced by task demands, we performed an rmANOVA with factors expectation (predictable; non-predictable) and task (perceptual; WM; grating) on the normalized beta weights averaged per task. This rmANOVA was performed separately for each ROI. A main effect of task and an interaction between task and expectation were found. Because we wanted to investigate whether expectation suppression for background stimuli depends on the type of available resources, we compared the two fixation tasks directly, using an rmANOVA with factors expectation (predictable; non-predictable) and task (perceptual; WM) on the normalized beta weights averaged per task. To verify that the expectation effect also differed between the perceptual and grating tasks, we compared the expectation effects during these tasks with an rmANOVA with factors expectation (predictable; non-predictable) and task (perceptual; grating).

To further examine the interactions, we performed post-hoc t-tests on the expectation effects for each task separately, within each ROI. These expectation effects were calculated by subtracting predictable from non-predictable beta

weights, such that positive values would indicate a lower response to predictable than non-predictable gratings (expectation suppression). We statistically evaluated the robustness of the effect using 150 voxels. Additionally, to assess whether the effect was stable across different voxel selection criteria, we calculated the expectation effect and the corresponding standard error of the mean (SEM) over a range of included voxels (50 to 300 voxels, in steps of 50 voxels).

Additionally, to investigate neural activity modulations in areas outside early visual cortex, we performed a separate, whole-brain analysis. Preprocessing followed a similar pipeline, except that subjects' T1 scans were normalized to MNI space. Functional images were brought into MNI space using the anatomical normalization parameters, and then smoothed with 8 mm kernel. The same GLM outlined above for the non-normalized images was applied to these data at the subject-level. The resulting beta weights were taken to a second (between-subjects) level, and tested using one-sample t-tests over linear combinations of the beta weights across subjects. To specifically test for areas showing the same resource-dependence of expectation suppression as found in the early visual cortex during the tasks at fixation, we specified a contrast to evaluate where there was greater expectation suppression during the perceptual task than the working memory task. Statistical inference was performed at the group-level using a cluster-level statistical test to assess clusters of significant activation (Friston et al., 1996). We used a familywise error (FWE) corrected cluster threshold of $p < .05$, with the spatial extent of clusters defined by a voxel threshold of $p < .001$ at the whole-brain level.

Results

Expectation suppression for unattended stimuli

Behaviour

Participants performed equally well on all three tasks (accuracy 85.6%, 87.8%, and 87.9% for perceptual, WM, and grating tasks, respectively, rmANOVA over all tasks: $F_{2,64}=1.34$, $p = 0.27$), demonstrating that participants followed task instructions and task difficulty was matched. There was no influence of grating predictability on either accuracy (rmANOVA over all tasks: $F_{1,32}=2.48$, $p = 0.13$) or RT (rmANOVA over all tasks: $F_{1,32}=2.13$, $p = 0.15$). There was a significant effect of task on RT (rmANOVA over all tasks: $F_{2,64}=23.45$, $p < 0.001$). Due to the additional time necessary for evidence accumulation of perceptually challenging stimuli, responses during the perceptual task (mean RT: 564 ms) were slower than during the grating task (mean RT: 524ms; $t_{32}=-4.21$ $p = 0.002$), and grating task responses were slower than during the working memory task (mean RT: 504 ms, $t_{32}=2.22$, $p = 0.035$). While valid expectations about task-relevant stimulus features often lead to behavioural

improvements, post-hoc t-tests on behavioural data from the grating task confirmed that there was no behavioural benefit of predictability on accuracy ($t_{32}=-1.03$, $p = 0.31$) or RT ($t_{32}=-0.20$, $p = 0.84$). For the grating task, expectations about stimulus orientation are not directly relevant to the task at hand (a spatial-frequency task) and therefore may not provide a behavioral benefit.

Visual cortex activity

To probe whether the effect of expectation depended on the task participants engaged in, we investigated the BOLD response in voxels in primary visual cortex that responded to the grating stimuli (see Methods). An interaction between expectation suppression and task would indicate that the effect of expectation is influenced by participants' task set. This is indeed what we found: the effect of expectation depended on task (rmANOVA over all tasks, V1: $F_{2,64}=4.15$, $p = 0.020$; V2: $F_{2,64}=3.44$, $p = 0.038$). We were particularly interested to know whether the expectation effect differed between the two tasks in which the predictable stimulus was task-irrelevant, since these tasks differed only in the relative load they placed on perceptual and working memory resources, but were matched on the locus of spatial-attention and the task-irrelevance of the grating. Indeed, the expectation effect depended on how processing resources were constrained by task demands: there was greater expectation suppression during the perceptual task than during the working memory task (rmANOVA over fixation tasks, V1: $F_{1,32}=5.08$, $p = 0.0312$; V2: $F_{1,32}=5.95$, $p = 0.020$; Figure 3.2A). Expectation suppression was also greater during the perceptual task than during the grating task in V1 (rmANOVA over perceptual and grating tasks, V1: $F_{1,32}=7.02$, $p = 0.012$; V2: rmANOVA, $F_{1,32}=2.03$, $p = 0.17$). In fact, there was a significantly reduced neural response to predictable stimuli only during the perceptual task (V1: $t_{32}=2.90$, $p = 0.0068$; V2: $t_{32}=2.54$, $p = 0.016$), but not during the working memory task (V1: $t_{32}=-0.94$, $p = 0.35$; V2: $t_{32}=-1.17$, $p = 0.25$) or the grating task (V1: $t_{32}=-1.52$, $p = 0.14$; V2: $t_{32}=0.26$, $p = 0.79$).

V3 displayed the same pattern of results, but the interaction between task and expectation did not reach significance across any of the task combinations (rmANOVA over all tasks: $F_{2,64}=2.80$, $p = 0.068$; rmANOVA over fixation tasks: V3: $F_{1,32}=3.90$, $p = 0.057$; rmANOVA over perceptual and grating tasks, $F_{1,32}=2.40$, $p = 0.13$).

These results were obtained on the basis of the 150 most grating responsive voxels (as determined by an independent functional localizer, see Methods) per region of interest, but the effects were largely independent of the number of voxels included. Irrespective of exact voxel selection criteria, the expectation effect (activity for non-predictable – predictable blocks) was different from zero for the perceptual task, but overlapping with zero for the working memory and grating tasks (see Figure 3.2B).

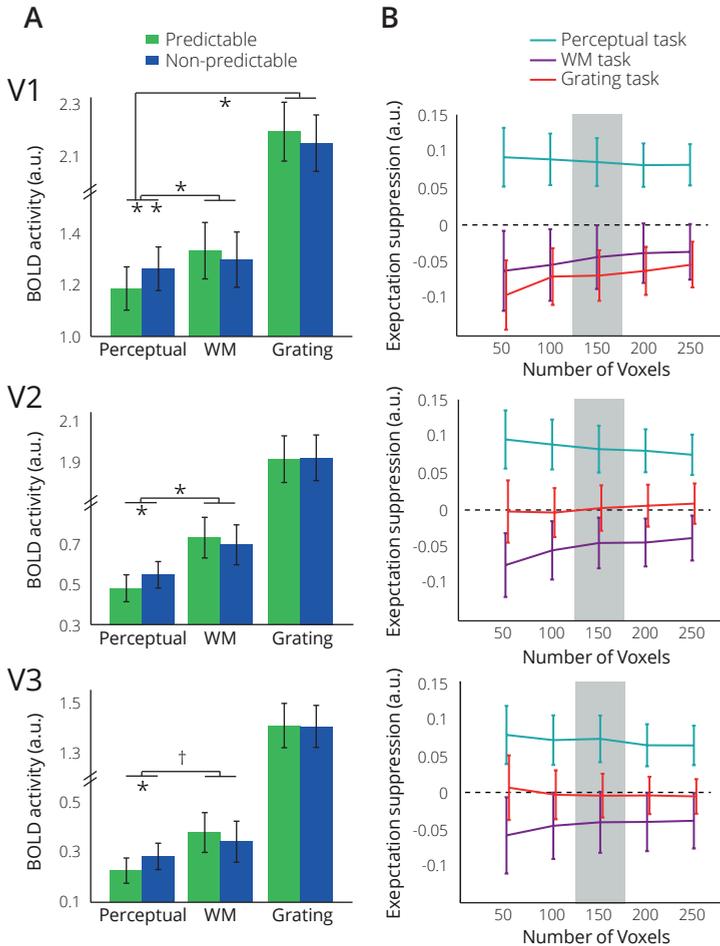


Figure 3.2. Expectation suppression for unattended stimuli depends on task set.

A) Amplitude of BOLD response (in arbitrary units; “a.u.”) within the 150 most grating-responsive voxels for predictable (green) and non-predictable (blue) stimuli, per visual area, per task. In each area, there was a reduced neural response to predictable stimuli within the perceptual task, but not during the working memory task or the grating task. Error bars reflect unbiased within-subjects corrected SEM (Cousineau, 2005; Morey, 2008). **B)** Expectation effects per number of included voxels, defined as the neural response to non-predictable minus the response to predictable gratings. Positive values indicate a lower response to predictable than non-predictable gratings (expectation suppression). Values close to zero indicate that expectation does not have an effect. Error bars reflect SEM. Grey region indicates the number of voxels depicted in **A**. Significance reported (if any) on the basis of t-tests for effect $\neq 0$. $**p < 0.01$; $*p < 0.05$; $\dagger 0.1 > p > 0.05$

In addition to the interaction between task and expectation, task itself had a strong effect on the BOLD response of voxels in the primary visual cortex (main effect of task, rmANOVA over all tasks, V1: $F_{2,64}=21.51$, $p < .001$; V2: $F_{2,64}=52.80$, $p < .001$; V3: $F_{2,64}=60.47$; $p < .001$; see Figure 3.2A). Neural activity within grating-responsive voxels in early visual cortex was more than twice as high during the grating task (i.e., when the grating stimulus was task-relevant) than during the fixation tasks (when the grating was a task-irrelevant background stimulus: $t_{32}=7.94$, $p < 0.0001$), indicating that this task successfully manipulated spatial attention. Additionally, in V2, the comparison between fixation tasks revealed a trend of overall stronger suppression of the grating during the perceptual load task than

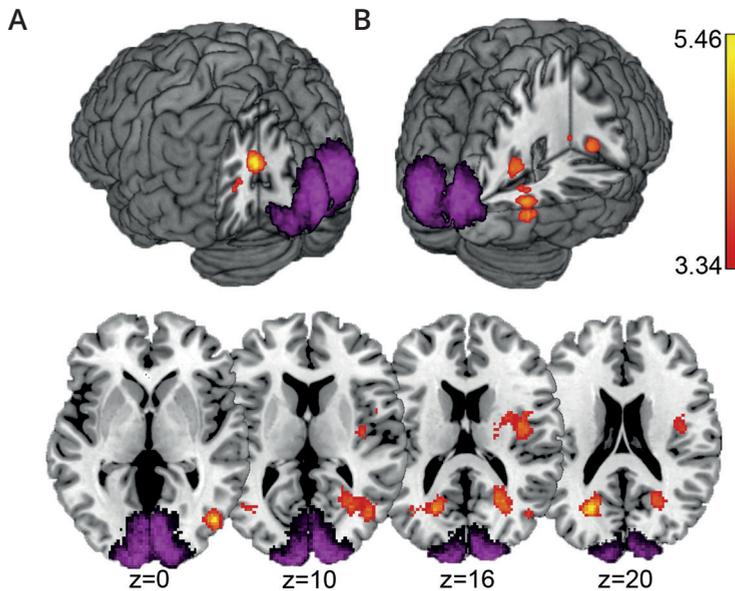


Figure 3.3. Interaction between task and expectation within bi-lateral cuneus and right insula.

Clusters within bilateral cuneus and right insular cortex show significantly stronger expectation suppression during the perceptual task than the working memory task (red/yellow). Cuneus activity is anterior and lateral to early visual cortex (anatomical contours in magenta; saturation indicates degree of overlap between subjects). **A)** Results displayed on a 3D MNI brain. **B)** The same results displayed on axial slices showing the bilateral cuneus clusters and right insular cortex activation. Z-coordinates displayed under each slice in mm. Cluster extents determined by $p < 0.001$ uncorrected. Scale indicates t values and applies to both panels.

during the working memory load task (main effect of task, rmANOVA over fixation tasks: $F_{1,32}=3.23, p = 0.082$). This trend was not present in V1 ($F_{1,32}=0.41, p = 0.53$) or V3 ($F_{1,32}=1.36, p = 0.25$). Because tasks were manipulated between runs, the absence of a robust task effect could therefore be caused by the increased variance in the task comparisons.

Whole-brain activity

Next, we probed whether other brain regions showed a similar interaction between the relative load on perceptual and working memory resources and expectation suppression. This interaction was present in bilateral cuneus (right: $p_{FWE} < .001$; left: $p_{FWE} = 0.013$) as well as right insular cortex ($p_{FWE} = 0.014$, Figure 3.3). In other words, bilateral occipital cortex and right insula showed stronger expectation suppression in the perceptual than in the working memory task, similar to what we observed in V1-V3.

Discussion

Expected stimuli often evoke a reduced sensory response compared to unexpected stimuli. Here we investigated whether the expectation suppression of unattended and irrelevant stimuli in the periphery depends on the type of task people perform centrally. Specifically, we compared the neural response evoked by predictable and non-predictable stimuli during tasks that loaded either perceptual resources or working memory resources. We observed that expectation significantly suppressed peripheral stimuli only during the perceptual task, in which perceptual resources were heavily loaded and working memory resources were largely available. On the other hand, when working memory resources were heavily taxed and perceptual resources were largely available, expectation suppression was abolished.

What could be the functional role of sensory suppression? Previous studies have found that sensory expectations facilitate perceptual inference, especially when input is complex and noisy or ambiguous (Bar, 2004). Anticipating sensory input naturally allocates processing resources in an efficient manner, with expected stimuli (that are thus low in informative value) receiving relatively fewer resources than surprising (and thus more informative) inputs (Rao and Ballard, 1999; Friston, 2005; Summerfield et al., 2008; Den Ouden et al., 2009; Friston, 2009b; Summerfield and Egner, 2009; Alink et al., 2010; Todorovic et al., 2011; Kok et al., 2012a; Kok et al., 2012b; Todorovic and De Lange, 2012). Predicting sensory input consequently minimizes both sensory uncertainty and processing resources consumed. For these reasons, it appears advantageous for the system to suppress predictable sensory inputs whenever possible.

During the task that placed a high load on working memory resources, there was no evidence of expectation suppression of the peripheral stimuli. Why would expectation suppression be abolished when working memory resources are strongly loaded by a concurrent task? The answer may lie in the fact that both expectation and working memory rely on a common process: bringing online stimulus templates in visual cortex. Recent studies show that the working memory maintenance of visual items leads to a cortical reinstatement of the maintained material in early sensory areas (Harrison and Tong, 2009; Serences et al., 2009; Riggall and Postle, 2012; Albers et al., 2013; Bosch et al., 2014), suggesting that the same neurons driven by bottom-up stimulus input become active during recall. Interestingly, stimulus expectation may employ this same process of cortical reinstatement (Albright, 2012; Summerfield and de Lange, 2014), with early sensory neurons representing a template of the expected stimulus, in the absence of bottom-up input (Kok et al., 2014). Once online, such stimulus templates can then be compared with incoming sensory information, facilitating perception and behaviour when there is a match (Lee and Mumford, 2003; Bar, 2004; Friston, 2005; Friston, 2009b; Kok et al., 2012a). We therefore speculate that expectation suppression may be abolished during the working memory task because this task places a high demand on the same resources needed for cortical reinstatement of expectation templates of the peripheral, irrelevant stimuli. Consequently, this inability to bring online templates of expected stimuli abolishes expectation suppression. Given that the task-relevant stimuli and the gratings were processed in non-overlapping regions of V1, the conflict may not be at the level of sensory cortex, but rather at the level of the frontal and parietal regions that are involved both in working memory (Fuster, 1973; Curtis and D'Esposito, 2003) and the generation of sensory predictions (Summerfield et al., 2006; Rahnev et al., 2011). The prefrontal cortex is known to be involved in top-down driven distractor suppression (Gazzaley and Nobre, 2012; Marini, 2014), and taxing this region with a working memory task may thereby hinder effective top-down suppression. Supporting this link between working memory and top-down suppression mechanisms, it has been found that during a challenging cognitive control task, individuals with high working memory capacity display relatively more distractor suppression than individuals with low working memory capacity (Gulbinaite et al., 2014).

Alternatively, the lack of expectation suppression during the working memory task may not solely be due to the high load on working memory resources, but caused by the relative absence of load on perceptual resources posed by this task. In other words, expectation suppression may take place only when predictable stimuli compete for task-relevant perceptual resources, as is the case in our perceptual load condition. During the (perceptually simple) working memory task,

perceptual resources are untaxed; therefore there is no need to suppress the irrelevant peripheral grating stimuli. In line with this, there was a trend towards an overall reduced sensory response to the peripheral grating when perceptual resources were centrally loaded (i.e., the grating evoked the least neural activity while participants performed the perceptual task compared to the other two tasks). This explanation could also explain the lack of expectation suppression during the grating task, since in this case the predictable stimuli are task-relevant, and so do not induce perceptual competition. However, previous research has shown that expectation suppression is also present in the absence of perceptual competition (Summerfield et al., 2008; Alink et al., 2010; Kok et al., 2012a). We cannot distinguish between the working memory and the perceptual competition explanations of our results on the basis of the current study.

This study manipulated cognitive load type by including different types of tasks and matching them for overall difficulty using an adaptive staircase procedure. Future studies may examine the role of different types of cognitive load more fully by manipulating the level of load within tasks (i.e., high vs. low load for each task). This may allow for more specific conclusions to be drawn about the factors relevant to expectation suppression, and distinguish between the two explanations of our results presented above.

Our results suggest that expectation suppression does not require attention to the predictable stimuli, but instead depends on the type of available cognitive resources. However, one might argue that the working memory task was more demanding than the perceptual task, and therefore left less attention for the gratings. In other words, the gratings might have been somewhat attended during the perceptual task, but not during the working memory task, culminating in expectation suppression for the former task but not the latter. In contrast to this notion however, we find an overall relatively *stronger* response to the irrelevant gratings during the working memory task, opposite to what would be expected with increased attentional load at the fovea. Additionally, accuracy was matched between tasks, suggesting that the complexity of the tasks were not markedly different. Both of these points argue against the possibility of reduced spatial attention for the grating during the working memory task.

Whole brain analyses revealed that bilateral cuneus and right insula showed the same interaction pattern between expectation and task that was present in early visual cortex. These areas have been previously reported to show sensitivity to statistical regularities in the environment (Turk-Browne et al., 2009).

We also included a task where participants responded to the grating stimuli, causing the predictable stimuli to become attended. Here we did not find an effect of expectation on the neural response. This result is perhaps unexpected, given that several previous studies show expectation suppression also for attended

stimuli (Summerfield et al., 2008; Todorovic et al., 2011; Kok et al., 2012a). However, the interaction between attention and expectation is complex and multifaceted (Summerfield and Egnér, 2009; Summerfield and de Lange, 2014). Recent studies suggest that attention interacts with expectation and may even reverse its effect (Bendixen et al., 2012; Kok et al., 2012b), such that expectation enhances the response when stimuli are attended (Doherty et al., 2005; Chaumon et al., 2008). Indeed, many studies on attention employ cues that concurrently signal what is likely and what is relevant (Posner and Petersen, 1990), thereby conflating expectation and attention. These studies generally observe increased activity for expected/attended sensory events (Maunsell and Treue, 2006). In light of this, we speculate that the absence of an expectation effect for attended stimuli in the current study could potentially be explained by an interaction between the opposing effects of attention and expectation. Additionally, previous research has generally focused on the contrast between confirmed and violated expectations (i.e., expected vs. *unexpected*), whereas here we contrasted situations with and without expectation cues (i.e., expected vs. *non-expected*). It is possible that the expectation suppression is particularly visible when contrasted against stimuli for which the expectation is violated, rather than for stimuli for which no strong expectation is present.

To summarize, our results provide insight into the resource dependency of expectation suppression. We show that expectation suppresses unattended, predictable stimuli when perceptual resources are loaded, but that this suppression is absent when working memory resources are loaded. This suggests that differences in the type of cognitive load might explain the conflict in the literature regarding the presence or absence of expectation suppression. Furthermore, our findings indicate that predictable stimuli do not need to be attended in order to be suppressed by expectation, pointing to resource availability as a crucial factor instead.

4

Graded expectations:
How does the brain weigh prior
expectation with task-relevant
sensory input?

Introduction

Imagine that you are out for a run in your neighbourhood in the evening. You notice a dark shape on the path ahead. A person bent over to pick something up, you think. But as you draw nearer, you realize it is just a pair of garbage bags, and what you saw as a head and a bent knee are just some bulges in the bags, catching the streetlight. We have all experienced this type of mistaken perception: when lighting is poor, or visibility is otherwise impaired (due to rain, mist, or poor eyesight), chance configurations of visual input are often interpreted as objects. This is because – far from being a passive process which gives rise to a veridical representation of external stimuli – perception is an inferential process wherein previous experience is used to make sense of current sensory input (Helmholtz, 1867; Summerfield and de Lange, 2014). The sensory inputs caused by the configuration of the garbage bags fit with the hypothesis of a crouching person, thus this percept reflected the brain's 'best guess' of the world at that moment.

We are more prone to mistaken perceptual inference when sensory input is noisy or ambiguous – we probably wouldn't have mistaken the garbage bags for a person in broad daylight. This is in line with the idea that perception is more influenced by prior knowledge when sensory input is noisy or ambiguous. Indeed, many theories of perception hold that sensory input should be weighted according to its reliability (Mumford, 1992; Rao and Ballard, 1997; Knill and Pouget, 2004; Friston, 2009a), with the consequence that priors should be up-weighted when sensory input is less reliable.

Empirical evidence supports this view. For example, low speeds are more frequent in the natural environment, and accordingly human perception is governed by a low speed prior (Weiss et al., 2002). Interestingly, this prior has more influence on perception when the sensory input has low contrast as compared to high contrast (Thompson, 1982; Stone and Thompson, 1992; Stocker and Simoncelli, 2006) – in other words, there was an increased effect of the prior when the stimulus strength was weak. Similarly, a study that manipulated stimulus strength through presentation duration found that a bias in perceived orientation was stronger when stimuli were weak (i.e., short duration) than when they were strong (i.e., long duration; (Wei and Stocker, 2015). Furthermore, a recent study showed that the expectation that a specific stimulus would be presented increased the sensitivity to weak stimulus-like sensory signals more so than to strong signals (Wyart et al., 2012).

In addition to this behavioral evidence, there is also some neural data that support an increased role for prior expectation when sensory input is weak. The brain is organized hierarchically, with sensory information fed forwards, up the hierarchy, and top-down signals such as prior knowledge or expectations sent

backwards, down the hierarchy. Feedback from higher-order areas can modulate the neural response in lower-order areas, amplifying or suppressing it. Hupé et al. (1998) demonstrated that the neural activity in a lower-order area was most strongly influenced by feedback signals from a higher-order area when the sensory input was weak. This is in line with an increased role of prior knowledge (sent via feedback signals) when sensory input is weak.

One well known influence of prior knowledge on the neural response is expectation suppression: a stimulus evokes a reduced response when it is expected compared to when it is unexpected (Summerfield et al., 2008; Den Ouden et al., 2009; Alink et al., 2010; Todorovic et al., 2011). For example, when the orientation of a visual grating stimulus is validly cued by a preceding auditory tone, the response in voxels that code for the stimulus in primary visual cortex is reduced relative to when that stimulus is invalidly cued (Kok et al., 2012a). Interestingly, this reduction of the neural activity is accompanied by an improved stimulus representation, quantified as the ability of a classifier to detect orientation information in the neural activity (Kok et al., 2012a). Neuroimaging paradigms to investigate the effects of expectation typically use stimuli that evoke a strong neural response. Visual stimuli, for example, are typically high contrast and otherwise easy to perceive (large, ample duration, etc.) – in the case of Kok and colleagues (2012a), grating stimuli were presented at 80% contrast for 500 ms.

Given that expectation suppression reflects a modulation of the sensory response based on prior knowledge, and the evidence that priors have a larger influence on perception when sensory input is weak, we formulated two hypotheses. First: the amount of expectation suppression should depend on the quality of the sensory input, and be (relatively) stronger when sensory input is weak. Second: the amount of sharpening in the sensory representation should likewise be relatively increased when sensory input is weak.

To test these hypotheses, we modified the paradigm of Kok et al. (2012a). Following Stocker & Simoncelli (2006), we manipulated the strength of the sensory input through stimulus contrast: 80% contrast for strong sensory input, and 10% contrast for weak sensory input.

Materials and Methods

Participants

Twenty-seven healthy right-handed individuals (19 females, age 23.4 ± 3.1 , mean \pm standard deviation (SD)) with normal or corrected-to-normal vision gave written informed consent to participate in this study. The study consisted of one behavioural session and two fMRI sessions, the first of which is described here.

The second session is described in Chapter 5. Experimental procedures were approved by the local ethics committee (Commissie Mensgebonden Onderzoek region Arnhem-Nijmegen, the Netherlands). Two participants were excluded due to excessive head movement (>2 mm in several task sessions), and one participant was excluded due to poor task performance (>2 SD below average task performance; 60% correct). Therefore, 24 participants were included in the analyses.

Stimuli

A small central fixation point was displayed at full contrast on a mean luminance gray background throughout the experiment. Grating stimuli were annuli (outer diameter: 15° of visual angle; inner diameter: 1.5°) containing luminance-defined sinusoids that remained on screen for 500 ms. On each trial, gratings were presented at one of two main orientations (45° or 135°) and one of two contrast values (10% (Figure 4.1C) or 80% (Figure 4.1B)), and a random spatial phase was selected from 10 values evenly spaced between 0 and 2π . Within each contrast level, 50% of trials were rotated by a small amount from the diagonals (25% clockwise (CW) and 25% counter clockwise (CCW); mean rotation: 7.2°), and 50% of trials were presented along the diagonal. The amount of rotation was determined per condition in a preceding behavioural session (see *Behavioural Session* below), and was updated online using adaptive staircases during the fMRI session (one for each condition; 1up-1down staircase with step-size ratio 0.2845, intended to achieve approximately 78% accuracy (Garcia-Perez, 1998)). A subtle flicker of the fixation dot, which was included for purposes of the second fMRI session (described in Chapter 5) occurred 0-3 times per trial for a total of 64 times a block. Auditory stimuli consisted of two pure tones (450 and 1000 Hz) with 200 ms duration. Stimuli were generated using MATLAB (MathWorks, Natick, MA, US) in conjunction with Psychophysics Toolbox (Brainard, 1997). In the fMRI session, stimuli were displayed on a rear projection screen using a luminance-calibrated EIKI projector (60 Hz refresh rate, 1024×768 resolution) which participants viewed through a mirror. In the behavioural session, stimuli were displayed on a Samsung SynchMaster 940BF monitor (60 Hz refresh rate, 1024×768 resolution).

Experimental Design

Stimuli were presented in an event-related design, with 4.6-6.6 s between trials. Each trial consisted of an auditory tone and a grating stimulus. The auditory tones preceded the gratings by 750 ms, and predicted the orientation of the grating that would appear on that trial (45° or 135°) with 75% validity (Figure 4.1A). The timing and duration of the tone and grating stimuli is based on Kok et al. (2012a). Participants' task was to report whether the grating was oriented along the diagonal (i.e., 45° and 135°) or was slightly rotated in either direction from the diagonal

(Figure 4.1B-C). Participants had 1.5 s from grating onset to respond using an MR-compatible button box (button 1='on diagonal'; button 2='off diagonal'). To encourage participants to respond to every trial, the fixation dot would turn red for 100 ms following a missed response. Feedback was given between blocks (mean accuracy, mean RT, and the number of missed trials for that block). Each block contained 64 trials, and participants completed 8 blocks (2 blocks per scanner run) for a total of 512 trials. A screen displaying the predictive relationship between the tones and gratings was displayed for 2 s at the beginning of each scanner run. The contingencies between tone and orientation were counterbalanced across participants.

Two additional scans were carried out after the main experiment: one to identify voxels that were maximally responsive to the grating stimulus, and one to enable the segmentation of early visual cortex into retinotopic areas. During the localizer, full contrast grating stimuli with the same size and position as the main experiment were flickered at 2 Hz for 14.4 s. Gratings were presented at eight orientations (22.5°, 45.0°, 67.5°, 90.0°, 112.5°, 135.0°, and 157.5°), with each

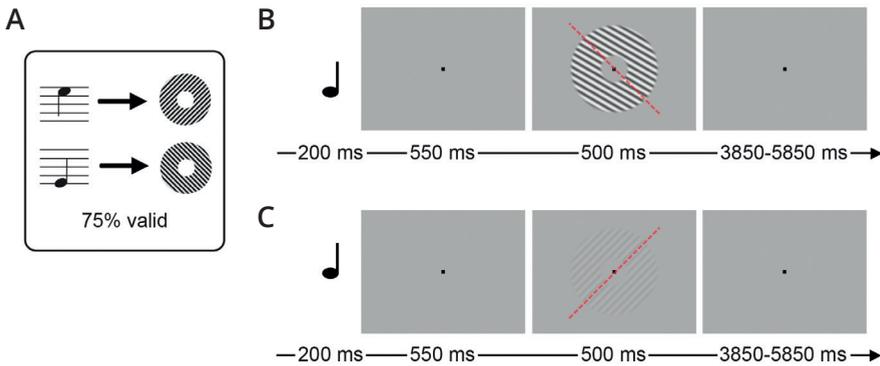


Figure 4.1. Experimental paradigm.

A) Auditory cues predicted grating orientation (~45° or ~135°) with 75% validity. **B-C)** On every trial, a grating stimulus was presented 550 ms after the auditory cue. Grating stimuli were either perfectly aligned to a diagonal or rotated slightly clockwise or counterclockwise from the diagonal. Participants' task was to judge, based on an internal reference of the diagonals, whether the grating was rotated or aligned with respect to the diagonal. The red dashed lines were not presented to participants. **B)** A high contrast trial. Here the overall orientation of the grating is expected, and the grating is slightly off-diagonal. **C)** A low contrast trial. Here the grating is unexpected, and on the diagonal. High and low contrast trials were randomly intermixed within blocks.

orientation displayed eight times in a pseudo-random order, interspersed with eight fixation blocks. Thus, each cycle of eight grating blocks and one fixation block lasted 129.6s, and this was repeated 8 times, with a break halfway through. To ensure fixation, participants' task was to detect two letters ('X', 'Z') in a stream of letters within a 0.8° fixation bull's-eye. During the retinotopy scan, a flashing black-and-white checkerboard pattern (3 Hz) in a 90° wedge rotated on a black background in 30° steps (1 position per TR). Participants' task was to detect unpredictable changes in the colour of the central fixation point (white to black), which occurred four to eight times per 36 s block. Nine cycles of CW and CCW rotation were presented. During both additional scans, participants responded to target events with a button press.

Behavioural Session

Within a time frame of approximately one week prior to the fMRI session, participants took part in a behavioral session outside the scanner. Participants were gradually introduced to the task, and trained on the contingencies between the tones and grating orientations. After this, performance was titrated to 75% using a Quest staircase procedure (Watson and Pelli, 1983), which adjusted the rotation angle of off-diagonal trials. Because stimuli with high contrast and/or an expected orientation might increase sensitivity to small rotation deviations, and therefore induce performance differences between the conditions, we used separate staircases for each contrast and expectation value (i.e., four staircases presented in an interleaved fashion). To limit the differences between bottom-up input between the conditions, the maximal rotation angle of off-diagonal trials could not exceed 12° in either (CW or CCW) direction. During this procedure, participants completed blocks of 64 trials until rotation offsets were stable (approximately 320 trials). Once this was completed, participants were introduced and staircased on a second task (described in Chapter 5). The behavioural session lasted between 45-90 minutes.

fMRI Acquisition

Functional images were acquired on a 3T Trio MRI system (Siemens, Erlangen, Germany) using a 32-channel head coil, with a multi-band EPI sequence (TR 1.8 s, 84 transversal slices, 2 x 2 x 2 mm in-plane resolution, TE 28 ms, field of view 210 mm x 210 mm, flip angle 73°, multi-band acceleration factor of 3, GRAPPA acceleration factor of 2). A high resolution anatomical image was collected using a T1-weighted MP-RAGE sequence (TR 2.3 s, TE 3.03 ms, 1 x 1 x 1 mm in-plane resolution, GRAPPA acceleration factor of 2).

fMRI Data Preprocessing

Data were pre-processed using SPM8 (<http://www.fil.ion.ucl.ac.uk/spm>; Wellcome Trust Centre for Neuroimaging, London, UK). The first four volumes of each task run were discarded to allow for time to achieve initial equilibrium. Functional images were spatially realigned to the mean image, and the resulting movement parameters, their first order derivatives and the square of these derivatives were included as nuisance regressors in all general linear models (GLMs). The structural image was coregistered to the mean functional image.

fMRI Data Analysis

Data were modelled using an event-related approach and analyzed within the framework of a General Linear Model (GLM). Regressors for each task condition were constructed by convolving delta functions at stimulus onsets with SPM8's canonical hemodynamic response function and its temporal and dispersion derivatives (Friston et al., 1998). Temporal and dispersion derivatives were included to capture possible deviations from the canonical HRF in terms of latency and duration of the BOLD response. In addition to task regressors (expected-high contrast; unexpected-high contrast; expected-low contrast; unexpected-low contrast), the model included nuisance regressors for the screens displaying tone-grating contingencies, feedback, and breaks, as well as the motion parameters as described above (see *fMRI Data Preprocessing*). Regressors were specified per run, and all task runs (from both fMRI sessions) were included in one GLM. A 128 s high-pass filter removed low-frequency signal drifts.

Data from the grating localizer were modelled in a separate design matrix using a blocked-design approach, with regressors for stimulation, fixation, and the motion parameters (with derivatives as discussed above). Due to the slower stimulus frequency during the localizer (each cycle of grating and fixation blocks took 129.6 s), we employed a high-pass filter of 256 s for the localizer scan to remove low-frequency signal components. In a separate analysis, *Freesurfer* (surfer.nmr.mgh.harvard.edu/) was used to inflate the cortical surface of each participant's T1-weighted structural image and to automatically identify the anatomical boundaries of V1. This was used as the region of interest (ROI) for our subsequent analyses.

As a measure of BOLD amplitude, the resulting canonical HRF beta-weights for the conditions of interest were averaged across runs of the grating task. To perform our main analysis, we selected 200 voxels that were most responsive to the grating stimulus during the localizer within the V1 ROI. For each condition, we averaged the beta weights of these voxels. To test whether the effect of expectation was influenced by the contrast of the grating stimulus, we performed a repeated measures ANOVA (rmANOVA) with factors expectation (expected; unexpected)

and contrast (high; low). Significant interactions were followed up with post-hoc *t*-tests on the expectation effect at each level of stimulus contrast. Expectation effects were calculated by subtracting the beta for expected gratings from the beta for unexpected gratings, such that positive values indicate a lower response to expected than unexpected gratings (expectation suppression). We statistically evaluated the robustness of the effect using 200 voxels. To verify whether our results were stable across different voxel selection criteria, we calculated the expectation effect (per contrast level) and the corresponding standard error of the mean (SEM) over a range of included voxels (50 to 200 voxels, in steps of 25 voxels).

It is possible that expectation may influence the latency and/or the duration of the neural response, and therefore the timing of the BOLD response. Therefore, to visualise the time course of the fitted BOLD response for each condition, we multiplied each HRF basis function (i.e., canonical, temporal, dispersion) by its fitted beta weight, and summed the three resulting timecourses.

Forward Modelling

To perform the forward modelling analysis, the above task regressors were defined separately for each orientation. Additionally, we split expected trials over three regressors so that the same number of trials would be used during the estimation of the expected and unexpected conditions. Again, information, feedback, and break screens, as well as the motion parameters as described above (see *fMRI Data Preprocessing*), were included in the model as nuisance regressors, and regressors were specified per run, with all task runs included in a single GLM. A separate GLM was specified for the grating localizer in which each trial was modeled by a separate regressor, with the motion parameters included as nuisance regressors. All GLMs for the multivariate analyses were performed in Matlab on the preprocessed data (high-pass filtered as per the amplitude analysis). To compensate for overall amplitude differences between runs, parameter estimates were normalized to z-scores for the localizer run and the experimental runs separately.

To probe stimulus representations in the visual cortex, we used a forward modelling approach to reconstruct the orientation angle from the BOLD signal (Brouwer and Heeger, 2009, 2011). This model assumes that each voxel is comprised of a large number of orientation-selective neurons with different orientation preferences, and that these preferences are fixed across sensory stimulation. We characterized the orientation selectivity of each voxel as a weighted sum of six hypothetical channels, each with an idealized direction tuning curve (or basis function). Model parameters are specified according to Brouwer & Heeger (2011). Each basis function was a half-wave sinusoid raised to the fifth power. Raising to the fifth power made the tuning curves narrower and thereby comparable to physiological findings. The six basis functions were evenly spaced within the 180°

orientation space, such that one channel responded maximally to the 135° stimulus and another channel responded maximally to the 45° stimulus, with two channels separating them on either side. The shape of the resulting channels approximate observed tuning curves of neurons in early visual cortex (Heeger, 1992).

In the first stage of the analysis, we used parameter estimates obtained from the localizer run to estimate the weights on the six hypothetical channels separately for each voxel, using linear regression. The details of this analysis have been reported previously (Brouwer and Heeger, 2011; Kok et al., 2013). Briefly, we estimated a weight matrix which related the matrix of estimated response amplitudes for the eight orientations presented during the localizer to the matrix of hypothetical channel outputs. Weights therefore reflected the relative contribution of the six hypothetical channels in the forward model to the observed response amplitude of each voxel. Using these weights, the second stage of analysis reconstructed the channel outputs associated with the pattern of activity across voxels evoked by the stimuli in the main experiment, again using linear regression. This step transformed each vector of n voxel responses (parameter estimates) to each trial into a vector of six channel responses. We reconstructed the population response separately for each condition (expected-high contrast; unexpected-high contrast; expected-low contrast; unexpected-low contrast) within the 200 voxels that were most responsive to the grating stimulus during the localizer within the V1 ROI (the same voxels as used in the amplitude analysis). Note that each expected condition was modelled by three regressors, each containing 1/3 of the expected trials – the forward model was run on each of these subsets and the resulting channel responses were averaged. Then, we averaged the channel responses to each condition over runs of the same task. Channel responses were zero-centred such that presented orientations (i.e., 135° and 45°) were set to 0°, and neighbouring channels were re-labelled relative to the presented channel (in other words, the 0° to 180° space became a -90° to +90° space).

To quantify the amount of orientation information in the channel responses, for each participant we subtracted the response of the channel orthogonal to the presented orientation from the response of the channel at the presented orientation (i.e., for trials with 135° gratings, the response of the 45° channel was subtracted from the 135° channel, and vice versa for trials with 45° gratings) separately for each condition. To test whether our conditions influenced the amount of orientation signal present in the population response, we performed an rmANOVA with factors expectation (expected; unexpected) and contrast (high; low).

To determine whether non-significant t -values reflected support for the null hypotheses, we used the online Bayes Factor Calculator (Rouder et al., 2009) to convert t -values to Scaled Jeffrey-Zellner-Siow (JZS) Bayes Factors. These factors

are directly interpretable: they provide the odds ratio in favour of one hypothesis over the other. For each test we used a JZS Prior (Cauchy distribution on effect size), without scaling the prior on the effect size ($r = 1$; default value) (Rouder et al., 2009). Smaller values of r can be used when effect sizes are expected to be small, however the effect size of the expectation effect in Kok et al. (2012a) is large (0.83), obviating the need for scaling.

Results

Behaviour

Participants were presented with high and low contrast gratings at orientations that were validly cued by a preceding auditory tone on 75% of trials. The task was to indicate whether grating were oriented precisely along one of the two diagonals, or slightly rotated from the diagonals: therefore gratings were relevant to the task. Using the two factors of contrast and expectation, we hoped to elucidate whether prior expectation are weighed differently depending on the strength of the sensory input.

On average, participants were 79% correct ($\pm 7\%$) and responded after 780 ms (± 88 ms), indicating that participants followed task instructions and were able to perform the task. Because we implemented a staircasing procedure to match task performance on the different conditions, there should be no benefit of high

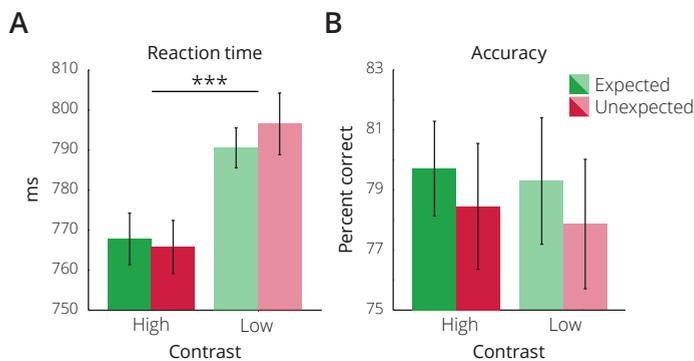


Figure 4.2. Behavioural results.

- A)** Faster reaction times to high contrast trials; however no significant benefit of expectation.
B) The staircases successfully calibrated the stimuli such that there were no accuracy differences between the conditions. *** $p < 0.0001$.

contrast or valid expectation on accuracy. Indeed, accuracy was not influenced by either factor (contrast: $F_{1,23}=0.27$, $p = 0.61$; expectation: $F_{1,23}=1.55$, $p = 0.23$), nor did they interact ($F_{1,23}=0.01$, $p = 0.94$; Figure 4.2B). The average rotation offset for off-diagonal trials was $7.2^\circ (\pm 2.6^\circ)$ CW or CCW. Though one might expect high contrast trials to have a smaller rotation offset, conditions did not differ in the average amount of rotation (mean high contrast: 7.3° ; mean low contrast: 7.1° , ($t_{23}=1.26$, $p = 0.22$)).

As expected, responses were faster on high contrast trials (mean low contrast: 794 ms, mean high contrast: 767 ms; $F_{1,23}=29.70$, $p = 1.54e-5$; Figure 4.2A). There was no effect of expectation on response time ($F_{1,23}=0.35$, $p = 0.56$). There was also no interaction between contrast and expectation ($F_{1,23}=1.08$, $p = 0.31$) on response time.

V1 amplitude results

To investigate whether expectation and contrast influenced how stimuli were processed by primary visual cortex, we selected voxels with the largest response to the grating in an independent localizer scan.

The amplitude of the stimulus-evoked BOLD response in V1 was strongly driven by grating contrast ($F_{1,23}=59.27$, $p = 8.23e-8$; Figure 4.3A). There was an interaction between contrast and expectation ($F_{1,23}=6.43$, $p = 0.0185$), such that the response to expected gratings was suppressed only on low contrast trials (low contrast: $t_{23}=3.27$, $p = 0.00339$, high contrast: $t_{23}=-1.08$, $p = 0.29$). Irrespective of exact voxel selection criteria, the expectation effect (activity for unexpected minus expected gratings) was different from zero for low contrast gratings, but overlapping with zero for high contrast gratings (see Figure 4.3B). To further investigate the null effect of expectation on high contrast gratings, we converted the t -value to a JZS Bayes factor (Rouder et al., 2009). The resulting value of 3.7 suggests some moderate support for the null hypothesis (i.e., the null hypothesis is approximately 3.7 times as likely as the alternative hypothesis), as opposed to a lack of power to adjudicate between the null and alternative hypotheses. For illustration purposes, we visualised the timecourse of the fitted BOLD response to the different conditions (Figure 4.3C). The timecourses are highly similar for expected and unexpected stimuli: expectation does not change the onset or peak latency of the BOLD response.

V1 stimulus representation results

Next, in the same set of voxels, we investigated how the pattern of neural activity representing the orientation of the stimulus was influenced by contrast and expectation. We quantified the amount of orientation information by subtracting the response of the orientation channel tuned to the orthogonal orientation from

the channel tuned to the presented orientation. High contrast increased the amount of orientation information ($F_{1,23}=63.10$, $p = 4.84e-8$; Figure 4.4) in the population response. There was no effect of expectation ($F_{1,23}=0.22$, $p = 0.64$) and

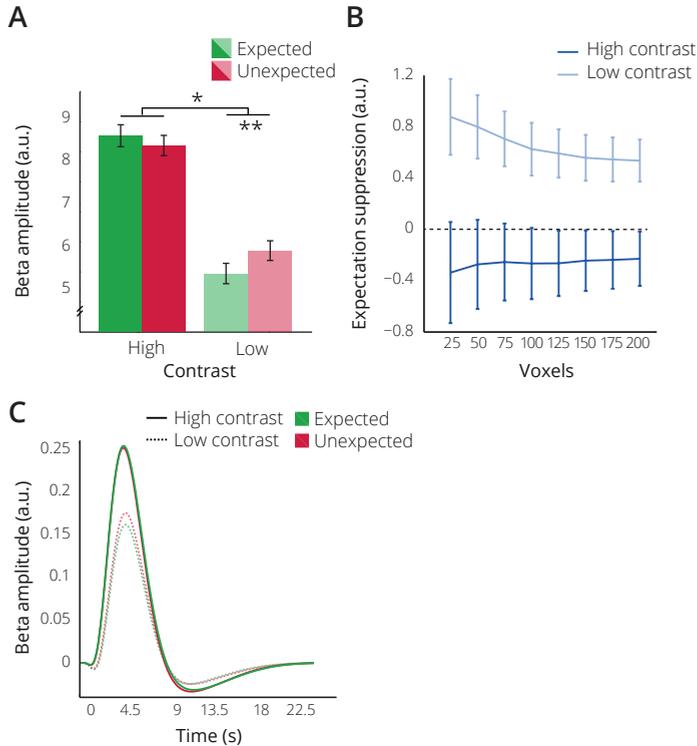


Figure 4.3. Stimulus activity in V1.

A) Amplitude of BOLD response (in arbitrary units; “a.u.”) within the 200 most grating-responsive voxels for expected (green) and unexpected (red) stimuli for each contrast level. High contrast gratings drive a larger BOLD response than low contrast gratings. Expectation interacts with contrast: low contrast stimuli are suppressed when they are expected compared to unexpected, however, there is no modulation of the response to high contrast gratings by expectation. Error bars reflect unbiased within-subjects corrected SEM (Cousineau, 2005; Morey, 2008). **B)** Expectation effects per number of included voxels, defined as the neural response to unexpected minus the response to expected gratings. Positive values indicate a lower response to expected than unexpected gratings (expectation suppression). Values close to zero indicate that expectation does not have an effect. Error bars reflect SEM. **C)** Timecourses of the fitted BOLD response. Significance reported (if any) on the basis of t-tests for effect $\neq 0$. $**p < 0.01$; $*p < 0.05$.

these factors did not interact ($F_{1,23}=0.36$, $p = 0.55$). The JZS Bayes factor for the effect of expectation is 5.7, which can be interpreted as some support for the null hypothesis (the odds are about 5.7 to 1 in favour of the alternative hypothesis).

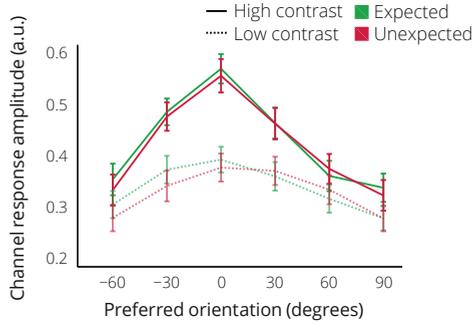


Figure 4.4. Orientation information in V1.

The channel response (in arbitrary units; “a.u.”) reconstructed from the 200 most grating-responsive voxels for expected (green) and unexpected (red) stimuli for high (solid lines) and low (dotted lines) contrast. Channels are ordered such that 0° maps onto the nearest diagonal of the stimulus orientation presented on that trial (i.e., 45° becomes 0° for $\sim 45^\circ$ trials), with 90° representing the orthogonal orientation. Error bars reflect unbiased within-subjects corrected SEM (Cousineau, 2005; Morey, 2008).

Discussion

In this experiment we investigated the effect of stimulus contrast on effects of expectation. In line with our first hypothesis, we found that the effect of the prior on the amplitude of the response was modulated by the quality of the sensory input: expectation suppression was stronger for low contrast stimuli than high contrast stimuli.

This result is in line with previous studies that have reported both a stronger influence of prior knowledge on perception (Thompson, 1982; Stone and Thompson, 1992; Stocker and Simoncelli, 2006; Wyart et al., 2012) and a larger modulation of sensory processing via feedback signals (Hupé et al., 1998) when sensory input is weak. These results suggest that there is a push-pull relationship between the weight given to bottom-up and top-down sources of information, with sensory input weighted according to its strength or quality. Indeed, when integrating sensory input from multiple sources, both humans (Ernst and Banks, 2002) and monkeys (Fetsch et al., 2012) proficiently weigh each source of information

according to its reliability. For instance, sensory information acquired through vision is more precise than sensory information acquired through touch, and as such vision tends to dominate the integrated percept – however, when visual input is noisy and less reliable, haptic cues are given more weight (Ernst and Banks, 2002).

The interaction between expectation and contrast was driven by the presence of a suppression effect only for low contrast stimuli: counter to our expectations, there was no difference in the response to expected and unexpected stimuli at high contrast. This null effect does not appear to be driven by a lack of power to adjudicate between the null and alternative hypotheses; however nor is the support for the null hypothesis strong. Given that our task was based on the paradigm used in Kok et al., (2012a), which reported robust expectation suppression at the identical contrast value, what might account for the fact that we find no modulation of the BOLD response by expectation in our design?

One possibility is the subtle differences to the task participants performed in this study: a ‘target/non-target’ task based on rotation from an internal template of the diagonals, compared to an orientation discrimination task between two consecutively presented gratings. Indeed, we have previously demonstrated that subtle differences in task set have profound effects on expectation suppression (St. John-Saaltink et al., 2015). However, an explanation purely based on task differences seems unlikely, given that we found a robust effect of expectation at low contrast using the same paradigm. In other words, the task used here does not preclude expectation suppression per se (although it may have attenuated it, thereby rendering it detectable only when the stimulus was weak). A second possibility is that the interleaved presentation of gratings at high and low contrast altered the context of stimulus processing. In the present study, the mean contrast over the experiment was 45%: perhaps this context emphasized the reliability of the high contrast gratings (which were substantially above the mean contrast), resulting in less use of prior knowledge. Put differently, what is considered a ‘strong’ or ‘weak’ sensory signal may depend on the expected strength of the sensory signal in the given context (Friston, 2009a). Alternatively, the fact that the contrast of the gratings changed randomly from trial to trial may have made the gratings appearance less predictable, lessening the effect of the predictive cues. Future studies could test these hypotheses by comparing the effects of a contrast manipulation in trial-by-trial (as here) and block-wise designs.

Our second hypothesis was in regards to the orientation information present in the neural response to the grating stimuli. Through training a forward model on the neural responses to an independent localizer, we were able to reconstruct the orientation signal present in the neural response to each stimulus during the task. Sensibly, this signal peaked at the orientation of the presented stimulus, and it was

stronger for high contrast gratings than for low contrast gratings. However, we did not find evidence for the sharpening account of expectation at either contrast: expectation did not influence the amount of orientation information present in the neural response. This is at odds with previous evidence for 'sharpened' neural representations of expected stimuli (Kok et al., 2012a). Future research is therefore necessary to investigate whether differences in task underlie whether or not expected stimuli benefit from a sharpened neural representation.

In conclusion, sensory expectations facilitate perceptual inference, especially when input is noisy or ambiguous (Bar, 2004). This study demonstrates that, in a context where the strength of the sensory signal varies, expectations modulate the response specifically when the signal strength is low. This is in line with a probabilistic view of perceptual inference, whereby perception depends more strongly on previous experience with the world when sensory input is ambiguous or incomplete.

5

Does the weighting of prior expectations with sensory input depend on task?

Many theories of perceptual inference posit that sensory input should be weighted according to its reliability (Mumford, 1992; Rao and Ballard, 1997; Knill and Pouget, 2004; Friston, 2009a), with the consequence that prior knowledge of the sensory statistics should be given more weight when sensory input is less reliable. In line with this, studies have found that prior knowledge has an increased effect on perception when signal strength is low compared to high (Thompson, 1982; Stone and Thompson, 1992; Stocker and Simoncelli, 2006; Wyart et al., 2012). In the previous chapter, we demonstrated that this also held for the effect of prior knowledge on the neural response to sensory input: valid expectations reduced the neural response when the stimulus signal was weak, but not when it was strong. However, in all of the aforementioned studies have a feature in common: the stimuli were attended by participants, because they were relevant to the task participants were performing. Therefore, it is unknown whether this reduction in neural activity to expected stimuli is the result of an automatic computation of the visual system, or instead critically depends on the requirements of the observer.

Previous research on this topic has been inconclusive. There is substantial evidence that when stimuli are attended, a valid expectation about which stimulus to expect suppresses the neural response (Summerfield et al., 2008; Todorovic et al., 2011). This is true even when participants' task is orthogonal to the predictive information, or when cues provide information about a different dimension of the stimulus than is relevant to their task (Kok et al., 2012a) – however, for a counter-example, see St. John-Saaltink et al. (2015). Evidence is more divided on the question of whether expectation can influence stimulus processing in the absence of (spatial) attention. When stimuli are task-irrelevant but spatial attention is not actively diverted from them, studies report a reduced neural response to predictable stimuli (Den Ouden et al., 2009; Alink et al., 2010). On the other hand, when attention *is* actively diverted from the stimuli, there is evidence that the *type* of task that is used to divert spatial attention can have important consequences. For instance, expectation suppression of a task-irrelevant stimulus has been reported when the distracting task loaded perceptual resources (St. John-Saaltink et al., 2015), but not when it loaded working memory resources (Larsson and Smith, 2012; St. John-Saaltink et al., 2015).

In the current study we examined whether prior expectation is weighed as a function of the strength of the sensory evidence independently of task demands, or instead depends on attentional set. To this end, we invited the participants who completed the fMRI session discussed in Chapter 4 for a second fMRI session. We presented grating stimuli with orientations that were either expected or unexpected, and manipulated the strength of the sensory input through stimulus contrast (Chapter 4; Stocker and Simoncelli, 2006). Crucially, during this session we rendered the peripheral grating stimuli task-irrelevant by having participants

perform a perceptual task at fixation. This allowed us to compare the response characteristics of expectation when stimuli are attended and task-relevant to the situation where they are fully irrelevant and outside the focus of spatial attention, with an aim to elucidate whether expectation suppression is automatic or task-dependent.

Materials and Methods

Participants

Approximately one week after participants performed the fMRI session described in Chapter 4, the same individuals completed the second fMRI session described here. Analyses are performed on the same 24 participants that were reported in the previous chapter (16 females, age 23.3 ± 3.2 , mean \pm standard deviation (SD)). All participants had normal or corrected-to-normal vision and gave written informed consent to participate in this study. Experimental procedures were approved by the local ethics committee (Commissie Mensgebonden Onderzoek region Arnhem-Nijmegen, the Netherlands).

Stimuli

A black central fixation point (luminance: 0.7 cd/m^2) was displayed on a mean luminance gray background (luminance: 530 cd/m^2) throughout the experiment. Stimuli were identical to those used in Chapter 4, with the exception of which stimuli were adjusted according to task performance (here, the fixation dot contrast 'dims', as opposed to the rotation angle of off-diagonal gratings in Chapter 4). As per the first fMRI session, a subtle 'flicker' of the otherwise full contrast fixation dot was implemented through reducing the contrast of the dot (i.e., black to grey) for 100 ms, 64 times in each block. In this session, however, the difference between full contrast (i.e., black) and 'dimmed' fixation dots was updated online using an adaptive staircase (1up-1down with step-size ratio 0.2845), intended to achieve approximately 78% accuracy (Garcia-Perez, 1998). The initial contrast value of the 'dimmed' fixation dots was determined by a preceding behavioural session (see *Behavioural Session* below). Across participants, the mean luminance value of the fixation dot 'dims' during the fMRI session was $159 (\pm 64) \text{ cd/m}^2$. Grating stimuli were annuli (outer diameter: 15° of visual angle; inner diameter: 1.5°) containing luminance-defined sinusoids that remained on screen for 500 ms. On each trial, gratings were presented at one of two average orientations (45° or 135°) and one of two contrast values (10% (Figure 5.1C) or 80% (Figure 5.1B)), with a random spatial phase selected from 10 values evenly spaced between 0 and 2π . Within each contrast level, 50% of trials were rotated by a small amount from the diagonals

(25% clockwise (CW) and 25% counter clockwise (CCW), and 50% of trials were presented along the diagonal (see Figure 5.1B-C). In order to match the sensory input between the first and second fMRI sessions as much as possible, for each participant, off-diagonal trials were rotated by the mean rotation angle across all off-diagonal trials presented during the first fMRI session (mean rotation: 7.2° , $\pm 2.5^\circ$). Auditory stimuli consisted of two pure tones (450 and 1000 Hz) with 200 ms duration. Stimuli were generated using MATLAB (MathWorks, Natick, MA, US) in conjunction with Psychophysics Toolbox (Brainard, 1997). In the fMRI session, stimuli were displayed on a rear projection screen using a luminance-calibrated EIKI projector (60 Hz refresh rate, 1024×768 resolution) which participants viewed through a mirror. In the behavioural session, stimuli were displayed on a Samsung SynchMaster 940BF monitor (60 Hz refresh rate, 1024×768 resolution).

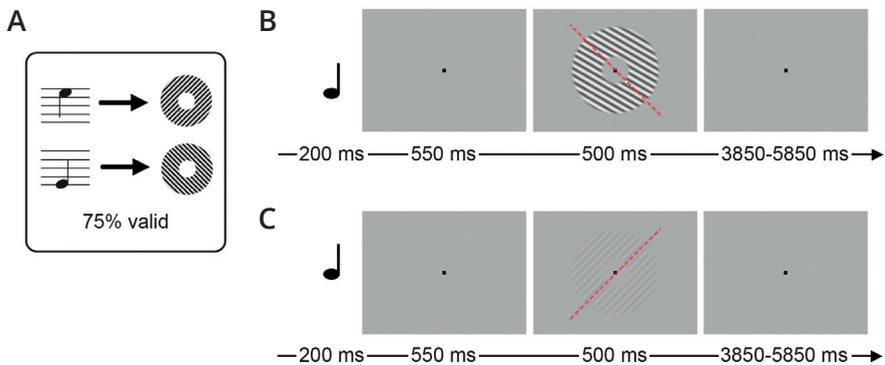


Figure 5.1. Experimental paradigm.

A) Auditory cues predicted grating orientation ($\sim 45^\circ$ or $\sim 135^\circ$) with 75% validity. **B-C)** On every trial, a grating stimulus was presented 550 ms after the auditory cue. Grating stimuli were either perfectly aligned to a diagonal or rotated slightly clockwise or counterclockwise from the diagonal (red dashed line). These stimuli were irrelevant to participants' task, which was to respond to subtle changes in the contrast of the fixation dot (black to grey). Red dashed lines were not presented to participants. **B)** A high contrast trial. Here the overall orientation of the grating is expected, and the grating is slightly off-diagonal. **C)** A low contrast trial. Here the grating is unexpected, and on the diagonal. High and low contrast trials were interleaved within blocks.

Experimental Design

Stimuli were presented in an event-related design, with 5.1-7.1 s between trials. Each trial consisted of an auditory tone, and 750 ms later, a grating stimulus (Figure 5.1B-C). Changes to the fixation contrast occurred independently of the tone and grating stimuli, with 0-3 fixation 'dims' per trial. The auditory tones predicted the orientation of the grating that would appear on that trial (CW/CCW) with 75% validity (Figure 5.1A) using the same tone-orientation contingencies used in the training and first fMRI session, however these stimuli were now completely irrelevant to the task that participants were performing. Participants were instructed to press a button on an MR-compatible button box as quickly as possible every time they saw the fixation dot briefly dim. Feedback was given between blocks (hit rate, mean RT, and the number of missed trials for that block). Participants completed eight blocks of 64 trials (two blocks per scanner run), for a total of 512 trials. A screen displaying the predictive relationship between the tones and gratings was displayed for 2 s at the beginning of each scanner run, so before every second block.

Behavioural Session

Participants were introduced to both tasks during the behavioural session that was held outside the scanner approximately one week prior to the first fMRI session (approximately one to two weeks prior to the second fMRI session). After participants completed the section on the grating task (see *Methods*, Chapter 4), they were introduced to the fixation task. Once they understood and practiced the task, performance was titrated to 75% using a Quest staircase procedure (Watson and Pelli, 1983), which adjusted the contrast of the dimmed fixation dot. During this procedure, participants completed blocks of 64 trials until the fixation contrast was stable. We opted for a Quest staircase during the behavioural session because this allowed the evolution of the staircase to be relatively independent of the starting value; however we opted for a 1up-1down staircase in the fMRI session – starting at the final Quest value – because this ensured a lower level of trial-to-trial fluctuations in the online staircase. The behavioural session lasted between 45-90 minutes in total.

fMRI Acquisition

Functional images were acquired on a 3T Trio MRI system (Siemens, Erlangen, Germany) using a 32-channel head coil, with a multi-band EPI sequence (TR 1.8 s, 84 transversal slices, 2 x 2 x 2 mm in-plane resolution, TE 28 ms, field of view 210 mm x 210 mm, flip angle 73°, multi-band acceleration factor of 3, GRAPPA acceleration factor of 2).

fMRI Data Preprocessing

Data were pre-processed using SPM8 (<http://www.fil.ion.ucl.ac.uk/spm>; Wellcome Trust Centre for Neuroimaging, London, UK). Functional images were aligned together with the functional images of the experiment described in Chapter 4 to facilitate comparisons between the two fMRI sessions. For full details about preprocessing steps, please see Chapter 4, *fMRI Data Preprocessing*.

fMRI Data Analysis

Data were modelled using an event-related approach and analyzed within the framework of a General Linear Model (GLM), according to the procedures described in Chapter 4. Regressors for each task condition were constructed by convolving delta functions at stimulus onsets with SPM8's canonical hemodynamic response function and its temporal and dispersion derivatives (Friston et al., 1998). Temporal and dispersion derivatives were included to capture possible deviations from the canonical HRF in terms of latency and duration of the BOLD response. In addition to task regressors (expectation x stimulus strength: expected-high contrast; unexpected-high contrast; expected-low contrast; unexpected-low contrast), the model included nuisance regressors for the screens displaying tone-grating contingencies, feedback, and breaks, as well as the motion parameters as described above (see *fMRI Data Preprocessing*). Regressors were specified per run, and all task runs (from both fMRI sessions) were included in one GLM. A 128 s high-pass filter removed low-frequency signal drifts.

As a measure of BOLD amplitude, the resulting canonical HRF beta-weights for the conditions of interest were averaged across runs of each task separately. We restricted our analyses to 200 voxels within the V1 ROI that were most responsive to the grating stimulus during the localizer (for details about the localizer task and analysis, and construction of V1 ROIs, please see *Experimental Design* and *fMRI Data Analysis*, Chapter 4). To compare the response to the grating between the two tasks, we averaged the beta weights of these voxels separately for each condition, for each task. To test for an effect of task on the interaction between expectation and the contrast of the grating stimulus, we performed a three-way repeated measures ANOVA (rmANOVA) in SPSS with factors task (grating; fixation), expectation (expected; unexpected), and contrast (high; low). To further examine the interactions, we performed post-hoc *t*-tests on the expectation effects at each level of stimulus contrast. Expectation effects were calculated by subtracting the beta for expected gratings from the beta for unexpected gratings, such that positive values indicate a lower response to expected than unexpected gratings (expectation suppression). We statistically evaluated the robustness of the effect using 200 voxels. To verify whether our results were stable across different voxel selection criteria, we calculated the expectation effect (per contrast level) and the

corresponding standard error of the mean (SEM) over a range of included voxels (50 to 200 voxels, in steps of 25 voxels).

It is possible that expectation may influence the latency and/or the duration of the neural response, and therefore the timing of the BOLD response. Therefore, to visualise the time course of the fitted BOLD response for each condition, we multiplied each HRF basis function (i.e., canonical, temporal, dispersion) by its fitted beta weight, and summed the three resulting timecourses.

Forward Modelling

To probe stimulus representations in the visual cortex, we used a forward modelling approach to reconstruct the orientation angle from the BOLD signal (Brouwer and Heeger, 2009, 2011). To perform the forward modelling analysis, task data were modelled with a separate GLM. The task regressors explained above were defined separately for each orientation, and additionally, because expected trials were three times as likely as unexpected trials, we split expected trials over three regressors so that the same number of trials would be used during the estimation of the expected and unexpected conditions. Again, information, feedback, and break screens, as well as the motion parameters, their first order derivatives and the square of these derivatives, were included in the model as nuisance regressors. Regressors were specified per run, with all task runs included in a single GLM. All GLMs for the multivariate analyses were performed in Matlab on the preprocessed data (high-pass filtered as per the amplitude analysis). To compensate for overall amplitude differences between runs, parameter estimates were normalized to z-scores for the localizer run and the experimental runs separately.

For a detailed description of the forward modelling approach, please see *Forward Modelling*, Chapter 4. Briefly, we characterized the orientation selectivity of each voxel as a weighted sum of six hypothetical channels, evenly spaced within the 180° orientation space. In the first stage of the analysis, we used parameter estimates obtained from the localizer run to estimate a weight matrix which related the matrix of estimated response amplitudes for the eight orientations presented during the localizer to the matrix of hypothetical channel outputs. Weights therefore reflected the relative contribution of the six hypothetical channels in the forward model to the observed response amplitude of each voxel. Using these weights, the second stage of analysis reconstructed the channel outputs associated with the pattern of activity across voxels evoked by the stimuli in the main experiment, again using linear regression. We reconstructed the population response separately for each condition (expected-high contrast; unexpected-high contrast; expected-low contrast; unexpected-low contrast), within the 200 voxels that were most responsive to the grating stimulus during the localizer within the V1 ROI (the same voxels as used in the amplitude analysis). We averaged the channel

responses to each condition over runs of the same task. Channel responses were zero-centred such that presented orientations (i.e., 135° and 45°) were set to 0°, and neighbouring channels were re-labelled relative to the presented channel (in other words, the 0° to 180° space became a -60° to +90° space – see Figure. 5.3).

To quantify the amount of orientation information in the channel response, for each subject, we subtracted the response of the channel orthogonal to the presented orientation from the response of the channel at the presented orientation (i.e., for trials with 135° gratings, the response of the 45° channel was subtracted from the 135° channel, and vice versa for trials with 45° gratings) separately for each condition per task. To test whether our conditions influenced the amount of orientation signal present in the population response, we performed an rmANOVA with factors expectation (expected; unexpected) and contrast (high; low).

To determine whether non-significant *t*-values reflect support for the null hypotheses, we used the online Bayes Factor Calculator (Rouder et al., 2009) to convert *t*-values to Scaled Jeffrey-Zellner-Siow (JZS) Bayes factors. These factors are directly interpretable: they provide the odds ratio in favour of one hypothesis over the other. For each test we used a JZS Prior (Cauchy distribution on effect size), without scaling the prior on the effect size ($r = 1$; default value) (Rouder et al., 2009). Smaller values of *r* can be used when effect sizes are expected to be small, however the effect size of the expectation effect in Kok et al. (2012a) is large (0.83), obviating the need for scaling.

Results

Behaviour

During this session, the statistical regularities between the tones and gratings were irrelevant to the task, which was to respond to subtle decreases in the contrast of the fixation dot (black to grey). Participants performed the fixation task well: average hit rate was 75% ($\pm 1.5\%$) and response time was 546 ms (± 36 ms), indicating that participants followed task instructions. The tight clustering of hit rates around 75% indicates that the online staircase functioned as intended.

Activity profile in primary visual cortex

To investigate whether spatial attention modulated the interaction between expectation and stimulus strength on neural activity in primary visual cortex, we compared the neural response to the grating stimulus when it was task-irrelevant (session 2) to when it was task-relevant (session 1). For this analysis, we selected voxels that displayed the largest response to the grating in an independent localizer

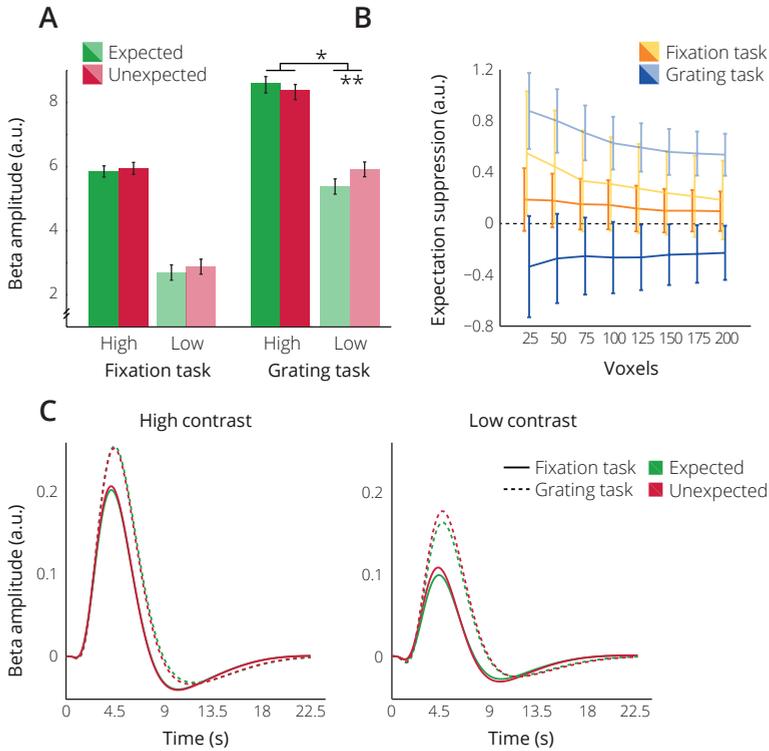


Figure 5.2. Stimulus activity in V1.

A) Amplitude of BOLD response (in arbitrary units; “a.u.”) within the 200 most grating-responsive voxels for expected (green) and un-expected (red) stimuli for each contrast level for each task. There is a larger BOLD response to gratings when they are task-relevant, and independent of relevance, high contrast gratings (bright colours) drive a larger BOLD response than low contrast gratings (faded colours). Error bars reflect unbiased within-subjects corrected SEM (Cousineau, 2005; Morey, 2008). **B)** Expectation effects per number of included voxels, defined as the neural response to unexpected minus the response to expected gratings (fixation task, orange/yellow (high/low contrast); grating task, blue/cyan (high/low contrast)). Positive values indicate a lower response to expected than unexpected gratings (expectation suppression). Values close to zero indicate that expectation does not have an effect. Error bars reflect SEM. **C)** Timecourses of the fitted BOLD response for expected (green) and unexpected (red) stimuli for the fixation (solid lines) and grating (dashed lines) tasks, separately for high and low contrast. Significance reported on the basis of t-tests for effect $\neq 0$. $**p < .01$; $*p < .05$.

scan. As in the previous chapter, we found that the effect of expectation was modulated by stimulus strength ($F_{1,23}=4.47$, $p = 0.045$; Figure 5.2A); expectation suppressed the response to weak (low contrast), but not to strong (high contrast) stimuli. This effect was not significantly modulated by task-relevance ($F_{1,23}=1.71$, $p = 0.20$). However, the effect did seem to be driven mainly by the task-relevant condition; when considering the fixation task in isolation, there was no evidence for a modulation of expectation by stimulus strength ($F_{1,23}=0.06$, $p = 0.8$). To investigate whether the data supported the null hypothesis that there is no expectation suppression for task-irrelevant stimuli (i.e., during the fixation task), we converted the corresponding t -values for the expectation effect at each contrast level to JZS Bayes factors (Rouder et al., 2009). There was moderate support for the null hypothesis (i.e. no effect of expectation) for both contrast levels (5.3 and 5.4 times more likely than the alternative hypotheses, respectively). Therefore, while there was a clear effect of expectation when the grating stimuli were task-relevant, there was moderate support for the absence of such an effect when the grating stimuli were task-irrelevant. This suggests that the absence of a task effect on the interaction between stimulus strength and expectation does not, in this case, provide strong evidence that task-relevance is unimportant. Rather, the JZS Bayes factor corresponding to this effect suggests that the results are inconclusive (2.9; i.e., about 2.9 to 1 in favour of the null hypothesis). This indicates that these data lack the power to discern an effect of task-relevance on expectation suppression.

However, this is not for lack of an effect of task relevance overall. Task relevance strongly modulated the amplitude of the BOLD activity evoked by the grating stimuli ($F_{1,23}=12.44$, $p = 0.0018$): there was a reduced neural response to the gratings when participants performed the fixation task. This suggests that our task manipulation was successful at diverting spatial attention away from the gratings. Similarly, stimulus strength strongly modulated the BOLD amplitude ($F_{1,23}=106.68$, $p = 4.14e-10$). The strength of this modulation did not differ between tasks ($F_{1,23}=0.77$, $p = 0.39$).

To evaluate whether the effect of expectation was consistent over different voxel selection criteria and to compare it to the effects during the grating task, we calculated an expectation suppression index by subtracting the beta for expected gratings from the beta for unexpected gratings over a range of included voxels. The absence of expectation suppression was independent of how many voxels were included in the analysis: the expectation effect for both low and high contrast gratings consistently overlapped with zero during the fixation task (orange/yellow (high/low contrast) lines in Figure 5. 2B).

For illustration purposes, we visualised the timecourse of the fitted BOLD response to the different conditions and tasks (Figure 5.2C). Both task and stimulus strength modulated the peak of the BOLD response: peak amplitudes were higher

during the grating task, and across both tasks, peak amplitudes were higher for high contrast gratings. However, the timecourses were highly similar for expected and unexpected stimuli: expectation did not change the onset or peak latency of the BOLD response.

Sensory representations in primary visual cortex

Next, in the same set of voxels, we investigated how the pattern of neural activity representing the stimulus was influenced by task, stimulus strength and expectation. We quantified the amount of orientation information by subtracting the response of the orientation channel tuned to the orthogonal diagonal orientation ($45^\circ/135^\circ$) from the channel tuned to the presented diagonal orientation ($45^\circ/135^\circ$). Consistent with the BOLD amplitude results, the amount of orientation information in the population response was modulated by both task-relevance ($F_{1,23}=10.51$, $p = 0.0036$) and stimulus strength ($F_{1,23}=44.14$, $p = 8.89\text{e-}7$; Figure 5.3). Expectation did not significantly modulate the orientation signal ($F_{1,23}=1.54$, $p = 0.23$), nor did it interact with either or both of the other factors (all $p > 0.10$). However, task did modulate the effect of stimulus strength ($F_{1,23}=14.63$, $p = 0.00087$), such that

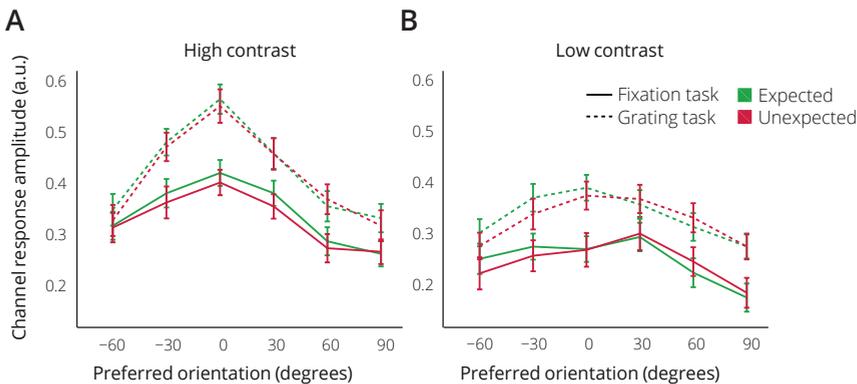


Figure 5.3. Orientation information in V1.

The channel response (in arbitrary units; "a.u.") reconstructed from the 200 most grating-responsive voxels for expected (green) and un-expected (red) stimuli for the fixation (solid lines) and grating (dashed lines) tasks, separately for high and low contrast. Channels are ordered such that 0° maps onto the nearest diagonal of the stimulus orientation presented on that trial (i.e., 45° becomes 0° for $\sim 45^\circ$ trials), with 90° representing the orthogonal orientation. Error bars reflect unbiased within-subjects corrected SEM (Cousineau, 2005; Morey, 2008).

there was a task-relevance boosted the increase of the orientation signal for high contrast than for low contrast gratings (i.e., multiplicative gain of the contrast effect when gratings were spatially attended).

Discussion

In this experiment, we investigated whether modulations of sensory processing by expectation constitute an automatic computation within the visual system. We compared the response modulation by expectation when stimuli were task-relevant and attended to the situation where they were task-irrelevant and outside the focus of spatial attention.

While there was a clear suppressive effect of expectation for task-relevant stimuli presented at low contrast (Chapter 4), this effect was absent for the same stimuli when they were irrelevant to the task. Furthermore, the Bayes factors for the effect of expectation during the fixation task suggest moderate support for an absence of an expectation effect when attention is diverted. While it then appears that the expectation effect is abolished when attention is diverted from the grating stimuli, the absence of a significant difference in the expectation effect between tasks precludes this strong conclusion. In fact, the Bayes factor for this statistical test suggests that our data does not strongly favour either the alternative or null hypothesis, possibly due to a lack of power. Stimulus contrast and task-relevance both strongly modulated the BOLD activity in the primary visual cortex, suggesting that our experimental manipulations of these factors were successful.

In the present study, there was no effect of expectation on the neural response in V1 to the grating stimuli when participants were performing a task at fixation. Put differently, when grating stimuli were task-irrelevant, the ability to predict a specific grating did not appear to culminate in a differential neural response to the stimulus. The lack of expectation suppression during the fixation task appears to be at odds with previous studies that do report expectation suppression for task-irrelevant stimuli (Den Ouden et al., 2009; Alink et al., 2010; St. John-Saaltink et al., 2015). The discrepancy between the current findings and those of one of our own previous studies (St. John-Saaltink et al., 2015) seems particularly striking, given that many aspects of the experimental design were similar. Specifically, in both studies the orientation of grating stimuli was predicted by preceding auditory cues, stimuli were presented at the same spatial location, and attention was diverted from the gratings by a perceptual task at fixation. While at first glance these results appear to be contradictory, there are two important differences between the current study and our previous one that may provide an explanation for the distinct results, which we outline below.

First, the previous study compared a fully predictable condition (100% cue validity) with a non-predictable condition (50% cue validity; no orientation information). In the current design, all cues were 75% valid, and we compared validly and invalidly expected stimuli. In other words, in the previous study we compared situations with and without predictions, whereas here there was always an expectation that was then violated in 25% of trials. A recent study has shown that this may be an important distinction (Hsu et al., 2015): in the context of auditory tone sequences, these authors report reduced activity when predictions were confirmed compared to violated, yet the least activity for tones about which there was no specific expectation in the first place. While this issue is not yet fully understood, these findings do suggest that the brain responds differently to sensory input that violates an expectation compared to input for which no specific expectation could be formed.

Second, while our previous study reported expectation suppression while subjects performed a perceptual task at fixation, the task used here (detection of subtle changes in the contrast of the fixation point) was less perceptually demanding than the task used in the previous study (identification of letters distorted by high levels of visual noise). Therefore, the amount of perceptual competition between the task-relevant stimuli at fixation and the irrelevant grating stimuli was likely much higher in the previous study. While this may seem like an inconsequential difference, the level of load posed by a task can have important consequences on the processing of task-irrelevant stimuli (Lavie et al., 2004; Lavie, 2005). It is possible that perceptual competition may be one of the factors which drives expectation suppression, such that predictable stimuli are suppressed only when they compete for task-relevant perceptual resources (St. John-Saaltink et al., 2015). Indeed, the previous study found no expectation suppression when attention was diverted from the gratings by a (perceptually simple) task at fixation that instead loaded working memory resources (St. John-Saaltink et al., 2015).

In addition to increasing the overall BOLD amplitude, we found that task-relevance also increased the orientation-specific BOLD signal. Furthermore, this increase was modulated by stimulus strength; there was a larger effect of attention on the orientation signal for high contrast than for low contrast stimuli, in line with a multiplicative gain effect of attention (Reynolds and Heeger, 2009). Notably, this modulation of attention by stimulus strength was opposite in direction to the modulation of expectation suppression by stimulus strength, which was rather strongest for low contrast stimuli. We found no evidence of an effect of expectation on the orientation signal.

A potential limitation to the current study is that all participants performed the grating task during the first session and the fixation task during the second fMRI session, making it difficult to rule out possible contributions of increased training

or stimulus familiarity on our results. However, we do not believe this would be a likely explanation for the effects reported here, as increased exposure to the statistical relationships between the tones and grating orientations should have only reinforced the predictive nature of the stimuli. Furthermore, participants were extensively trained on the stimuli in the behavioural session, so they had a high level of familiarity with the experimental stimuli and tasks before either fMRI session.

In conclusion, the issue of whether expectation suppression is automatic or task-dependent is complex, and has not yet been resolved fully (Schröger et al., 2015). Some studies have concluded that expectation effects require attention (Larsson and Smith, 2012), while other results provide evidence for attention independent effects (Den Ouden et al., 2009; Alink et al., 2010; St. John-Saaltink et al., 2015). It is likely that this partly depends on the exact task used to divert attention, such as what type of resources (e.g. perceptual or working memory) are loaded (St. John-Saaltink et al., 2015), and whether the task induces perceptual competition between the predictable (irrelevant) stimuli and the task-relevant stimuli. Future studies may be able to shed more light on this issue, by orthogonally manipulating these factors within the same experimental design.

6

Discussion

Vision is profoundly important to how we perceive and interact with the external world. The information we collect about the world through vision is rich and complex: a typical scene contains numerous objects which may be at different distances, under different lighting conditions, and often partially occluded by other elements of the scene. In order to process such dense and complex visual information efficiently, the brain makes predictions about the likely causes of sensory input. These predictions depend on prior experiences with the world, through which we learn, for instance, that distance causes objects to appear smaller than they are. Under certain circumstances, priors can lead to visual illusions – for instance, by inflating perceived object size. This is why the moon appears to be larger than usual when it is close to the horizon, and why the rightmost soldier in Figure 1.1C appears to be larger than the other figures.

In this thesis, I have examined how prior experience influences the way in which the brain responds to the sensory environment, and consequently shapes how we perceive the world. Specifically, I have investigated how visual information is processed by the brain depending upon the circumstance in which it arrives: what input has preceded it, whether specific features of the visual input could be predicted prior to the arrival of the stimulus, and which cognitive resources were available.

Chapter 2 investigated a natural prior about object stability across time. From moment to moment, the majority of visual stimulation does not change or changes only subtly. For example, when you are looking at the houses that line the other side of your street and a large truck momentarily occludes them from view, the row of houses will reappear unchanged once the truck has passed. Even when you are in a more volatile environment, such as walking in a forest during a storm, wind and rain can cause local variations in the appearance of objects, but they do not change the relevant aspects of the environment needed to navigate to shelter, such as position and size. Indeed, faster variations in input often occur as the result of noise, and performance (here, navigation) can be improved when these variations are disregarded. One way to accomplish this is to ‘smooth’ visual input over time, such that perception depends partly on recent visual history. Humans appear to use a prior which acts as precisely such a temporal filter: perception is biased towards recently-seen items, such that objects appear to be more similar to the preceding object than they truly are. In the experiment described in Chapter 2, I provided the first evidence that this bias, or serial dependence, is reflected in stimulus representations in the primary visual cortex. This suggests that our prior knowledge of object stability affects the processing of visual input already at the earliest cortical stages, rather than being something we take into account after visual processing is completed.

In **Chapter 3**, I compared how the ability to generate precise sensory predictions about visual stimuli influenced how visual areas of the brain responded to those stimuli, and how this effect was altered by the type of task that participants were engaged in. Here, I contrasted blocks in which the orientation of stimuli displayed in the visual surround could be precisely predicted to blocks in which a precise prediction about orientation could not be formed, while participants performed three tasks that loaded distinct cognitive resources. When the cued stimuli in the surround were task-irrelevant and perceptual resources were engaged by a challenging task centrally, predictable stimuli were suppressed relative to nonpredictable stimuli. However, when instead working memory resources were loaded, the task-irrelevant stimulus in the surround was not modulated by predictability. The fact that a distracting working memory task (but not a perceptual task) prevented prediction effects may suggest that visual prediction shares some cognitive resources with working memory – for instance, the ability to reinstate stimuli from memory into sensory cortex. Surprisingly, when participants performed a task in which the cued stimuli were actually relevant, there was no effect of expectation on the neural response. I return to this finding below.

In **Chapters 4-5**, I manipulated participants' expectations about sensory input while simultaneously manipulating the strength of the sensory input. Current theories of perceptual processing would predict that expectations influence perception more strongly when sensory inputs are weak. In **Chapter 4**, participants performed a task on the stimuli that were being cued. The effect of expectation on the sensory response in primary visual cortex interacted with the strength of sensory input: expectation suppression was found for the low contrast stimuli, while expectation did not influence the response to high contrast stimuli. As stated above, this is in line with current theories of perceptual inference: you should rely on prior knowledge more when sensory inputs are weak than when they are strong. In **Chapter 5** we found that when participants performed a distracting task at fixation, while being exposed to identical predictions and stimuli as in Chapter 4, this interaction was not present. In fact, here there was no evidence of an effect of expectation on the neural response for either contrast level.

In sum, the studies in this thesis provide evidence that the brain's response to the environment is modulated by prior experience, but that this effect is highly dependent on internal and external processing conditions.

Expectation evokes stimulus templates

A central question that can be asked about the effects described in this thesis is whether they rely upon the same neural mechanism. To answer this, I first discuss how, mechanistically, expectations might influence the way in which the brain responds to sensory input. When, for example, a tone predicts a specific stimulus – say a grating oriented at 45 degrees – what happens in the brain after the predictive tone sounds? A proposal gaining increasing popularity is that a ‘template’ of the expected stimulus becomes activated in sensory cortex (Wyart et al., 2012; SanMiguel et al., 2013; Kok et al., 2014). This may be achieved via subthreshold activations of the neurons that code for the stimulus (Wyart et al., 2012). This template can then be compared to the sensory input when it arrives, facilitating stimulus processing when the input and template match. Conversely, a mismatch between the template and sensory input often results in increased activity (Den Ouden et al., 2009; Todorovic et al., 2011; Kok et al., 2012a), possibly because both the population that codes for the template as well as the population coding for the unexpected input become active.

Support for the notion that expectations activate stimulus templates is provided by studies that induce the expectation of a particular stimulus, but then do not present any stimuli for the remainder of the trial. When expected stimuli are omitted, sensory cortex still responds at the time the stimulus was expected to occur, and this response resembles the expected stimulus (Todorovic et al., 2011; Kok et al., 2012a; SanMiguel et al., 2013). Importantly, since there is no bottom-up input to drive the activity, this cortical response must be internally generated. A template can therefore be understood as a neural representation of a stimulus that is representationally similar to the neural response elicited when the stimulus is presented (however, significantly less strong). This result has been demonstrated in both visual and auditory cortices.

The idea that expectation activates weak or subthreshold stimulus templates provides a compelling account of the interaction between expectation and the strength of the sensory input: the influence of a weak template compared to the (strong) sensory input would be relatively much larger when the sensory signal is weak, leading to a larger expectation effect. In other words, the same subthreshold template will have a larger net effect when it is integrated with weak sensory input than with strong sensory input (Wyart et al., 2012). This is precisely what I found in **Chapter 4**: expectation suppression was larger for low contrast than high contrast gratings.

Once active, expectation templates have been shown to bias the perception of ambiguous, weak or noisy sensory input (Sterzer et al., 2008; Chalk et al., 2010). Perhaps you have previously encountered an ambiguous stimulus such as the

Necker Cube (Figure 6.1) – a three-dimensional line drawing famous due the perceptual ambiguity it induces as to which face is ‘closest’ to the viewer. If you stare at Figure 6.1 long enough, your percept will switch to the other ‘solution’, of the input (i.e., you will see the other face as ‘closest’). To demonstrate the power of priors on the perception of ambiguous stimuli, one experiment led participants to believe that a certain colour of eyeglasses lens would lead to a particular solution (i.e., percept) of an ambiguous stimulus (Sterzer et al., 2008). However during this training, participants were actually presented with non-ambiguous stimuli displaying one of the two perceptual solutions. Tested afterwards on *truly* ambiguous stimuli, participants reported their percepts in accordance to the colour of eyeglasses lens they were wearing, despite the fact that the glasses had no influence on the ambiguity of the visual input. In this case, the colour of eyeglasses lens could be understood as activating a template of a particular perceptual solution, which then biased perception of the incoming, ambiguous stimulation. Similar results have been found when stimulation is unambiguous but noisy or weak (Chalk et al., 2010). For example, through manipulating the motion coherence of clouds of moving dots, the overall direction of motion can be very difficult to perceive. An experiment that induced expectations about motion directions using auditory cues (in a manner similar to **Chapters 3-5**) found that both perceptual reports and neural representations in visual cortex were biased towards the expected motion direction (Kok et al., 2013). While subtle – the attractive bias was on the order of a few degrees – the effect was highly reliable and furthermore correlated to the strength of the representational bias in primary visual cortex.

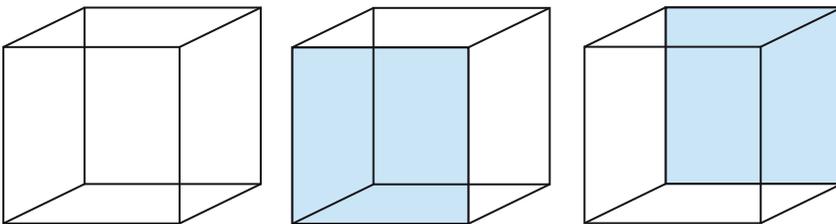


Figure 6.1. The Necker Cube.

This two-dimensional image is perceived as a three-dimensional cube; however this image is ambiguous as to which face of the cube is closest to the viewer and therefore perception oscillates between the two possibilities. The middle and right cube depict the two ways in which the cube can be perceived (with the shaded face ‘closest’).

Can stimulus templates also account for the effects of more implicit or automatic priors such as the serial dependence reported in **Chapter 2**? This is indeed a possibility. However, in this case the templates may originate from a different source, either from within visual cortex or from the fronto-parietal regions implicated in perceptual decision-making (Gold and Shadlen, 2007). Given that serial dependence was driven by the previous perceptual *decision* more than by the previous *input*, the latter option appears more likely. This is in line with previous research that demonstrated that primary visual cortex first shows activity corresponding to the sensory input and later to the perceptual decision made on the basis of that input (Nienborg and Cumming, 2009).

Effects of expectation depend on task

Why do we not *always* see a modulation of the neural response by expectation? Many studies (including those in this thesis) have found that participants' task set influences the effect of expectation on neural processing (Larsson and Smith, 2012; Schröger et al., 2015). For instance, in **Chapter 3**, I found that expectation had a suppressive effect when participants performed a task that loaded perceptual resources, but showed no such effect when the task instead loaded working memory resources. What makes this particularly striking is that during both tasks, the sensory stimulation was identical, and the cued stimuli were entirely irrelevant to participants' tasks. Indeed, the only aspect that differed was whether participants were performing a 1-back task on hard to identify letters, or a 2-back task on the (easy to perceive) colour of those same letters. One explanation of this result is that working memory and perceptual predictions share some neural resources, and that therefore perceptual predictions can be disrupted by high working memory loads. For instance, the template activation that occurs during stimulus expectation has an equivalent concept within the memory literature: cortical reinstatement (Danker and Anderson, 2010). When a stimulus is recalled during a working memory task, the same neurons that were active during stimulus processing become reactivated during the recall period. In other words, both stimulus expectation and stimulus recall appear to recruit the population of neurons that process the sensory input, when this input is either expected (Kok et al., 2014) or recalled (Harrison and Tong, 2009; Albers et al., 2013; Bosch et al., 2014). Hence, it is possible that when a distracting task occupies a lot of the resources needed to form expectation templates, no modulation of neural responses by expectation occurs.

A special case of this may occur when the task itself requires instantiating stimulus templates. For instance, one study found that the neural response to sensory input was suppressed by expectation *only* when that specific sensory

input was not already being prepared for as a result of the task instructions (Todorovic, 2015, Chapter 5). Specifically, when the task involved preparing for a particular tone (tone A), the neural response to this tone was no longer modulated by expectation. However, when a different tone was prepared or the visual modality was attended, the predictive cue that tone A would occur *did* suppress the neural response to tone A (i.e., typical expectation suppression). One explanation of this result is that preparing for a particular stimulus involves forming and maintaining a template of that stimulus, which may interfere the formation of expectation templates, therefore abolishing any neural effects of expectation. This explanation also offers a possible account of the absence of expectation suppression when the cued stimuli in **Chapter 3** were task relevant. The task involved detecting low spatial frequency gratings, thus the instantiation of a target (low spatial frequency grating) template in order to solve this task may have similarly prevented the formation of expectation templates.

Interference of the ability to generate expectation templates may occur either because the sensory neurons that code for the stimulus are already 'occupied' by the stimulus template induced by the current task, or because the higher-level resources required for generating stimulus templates (e.g. in frontoparietal cortex (Curtis and D'Esposito, 2003; D'Esposito and Postle, 2015) or hippocampus (Bosch et al., 2014; Gordon et al., 2014)) are otherwise engaged. The results of **Chapter 3** could be interpreted as evidence of a dependence on working memory resources, given that the expectation effect in **Chapter 3** was abolished during the working memory task. However, a resource competition account provides an alternative interpretation of these results: the perceptual task may have induced perceptual competition between the task-relevant degraded letters and the cued grating stimuli in the surround, creating an incentive to suppress the distracting gratings as much as possible (and predictive cues allowed this to be done more effectively). Thus, according to this explanation, the relative lack of perceptual competition during the working memory task (and not the working memory load itself) was responsible for the lack of expectation suppression during the working memory load. The question of which resources are required for the generation of expectation templates could be resolved by an experiment that uses a different modality for the distracting tasks – for example, a target detection task and a 2-back task on auditory tones – while expectations about visual grating stimuli are induced by coloured cues at fixation. If working memory resources are necessary for generating expectation templates, there should be no difference in the response to expected and unexpected gratings in visual cortex when participants are performing the auditory 2-back task, but there should be expectation suppression during the auditory target detection task. However, if the alternative, perceptual competition account is correct, then expectation suppression should occur during *neither* of the

distracting auditory tasks, since the cued and task-relevant stimuli no longer compete for perceptual resources in visual cortex, given that they are presented in different sensory modalities. Insight into which resources are involved in the ability to generate predictions about upcoming sensory input will greatly enhance our understanding of the nature of expectation, and how it relates to other concepts like cortical reinstatement during memory tasks.

I have suggested that a form of expectation templates may also drive the serial dependence effect found in **Chapter 2**. This proposal could be tested using a paradigm that has participants perform two distracting tasks, one loading perceptual resources and one loading working memory resources, while leveraging neuroimaging techniques to evaluate the serial dependence between the representations of subsequent *ignored* grating stimuli in visual cortex. Firstly, this could determine whether serial dependence occurs for stimuli on which no explicit perceptual report is made. Is the neural representation of an ignored stimulus biased by previous stimulation, or does this bias require a perceptual decision on the current stimulation? This question is hard to answer using psychophysics alone, since in such studies a behavioural report is necessarily required in order to have a measure of stimulus processing. Therefore, neuroimaging techniques would provide a unique opportunity to investigate this issue. Additionally, there is the question of whether serial dependence, like other types of perceptual expectations, requires cognitive resources to generate stimulus templates, or whether it is more automatic. This question can be answered by comparing serial dependence effects between the two distracting tasks; if serial dependence is a fully automatic process, it should occur during both tasks, while it should be abolished during the distracting working memory task in case it requires the generation of stimulus templates (see **Chapter 3**).

Concluding statement

In closing, considered together these findings point to a complex interaction between expectation and task set which is currently insufficiently understood. However, it is clear that task set plays a less straightforward role than previously thought: it is not simply a matter of whether stimuli that are task relevant or irrelevant, but rather the exact requirements of the task that determines whether prior expectation modulates the sensory response. Future research should try to more precisely determine the interplay between expectation and other cognitive processes, such as attention, working memory, and imagery. One particularly intriguing outstanding question is whether expectation, working memory, and imagery all use the same cortical reinstatement process (Pearson and Westbrook, 2015). Resolving these issues will hopefully shed more light on the neural mechanisms underlying our perception of the world.

References

Nederlandse samenvatting

Acknowledgements

List of Publications

Biography

Donders Graduate School
for Cognitive Neuroscience

References

- Adams WJ, Graf EW, Ernst MO (2004a) Experience can change the 'light-from-above' prior. *Nat Neurosci* 7:1057-1058.
- Adams WJ, Graf EW, Ernst MO (2004b) Experience can change the 'light-from-above' prior. *Nat Neurosci* 7:1057-1058.
- Albers AM, Kok P, Toni I, Dijkerman HC, De Lange FP (2013) Shared representations for working memory and mental imagery in early visual cortex. *Curr Biol* 23:1427-1431.
- Albright TD (2012) On the perception of probable things: neural substrates of associative memory, imagery, and perception. *Neuron* 74:227-245.
- Alink A, Schwiedrzik CM, Kohler A, Singer W, Muckli L (2010) Stimulus predictability reduces responses in primary visual cortex. *J Neurosci* 30:2960-2966.
- Arnal LH, Wyart V, Giraud A (2011) Transitions in neural oscillations reflect prediction errors generated in audiovisual speech. *Nat Neurosci* 14:797-801.
- Bar M (2004) Visual objects in context. *Nat Rev Neurosci* 5:617-629.
- Bendixen A, SanMiguel I, Schroger E (2012) Early electrophysiological indicators for predictive processing in audition: a review. *Int J Psychophysiol* 83:120-131.
- Biederman I (1972) Perceiving real world scenes. *Science* 177:77-80.
- Biederman I, Mezzanotte R, Rabinowitz J (1982) Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology* 14:143-177.
- Bosch SE, Jehee JF, Fernandez G, Doeller CF (2014) Reinstatement of associative memories in early visual cortex is signaled by the hippocampus. *J Neurosci* 34:7493-7500.
- Brainard DH (1997) The Psychophysics Toolbox. *Spatial Vision* 10:433-436.
- Brouwer GJ, Heeger DJ (2009) Decoding and reconstructing color from responses in human visual cortex. *J Neurosci* 29:13992-14003.
- Brouwer GJ, Heeger DJ (2011) Cross-orientation suppression in human visual cortex. *J Neurophysiol* 106:2108-2119.
- Burr D, Cicchini GM (2014) Vision: Efficient adaptive coding. *Curr Biol* 24:R1096-R1098.
- Chalk M, Seitz AR, Seriès P (2010) Rapidly learned stimulus expectations alter perception of motion. *J Vision* 10:1-18.
- Charles L, King JR, Dehaene S (2014) Decoding the dynamics of action, intention, and error detection for conscious and subliminal stimuli. *J Neurosci* 34:1158-1170.
- Chaumon M, Drouet V, Tallon-Baudry C (2008) Unconscious associative memory affects visual processing before 100 ms. *Journal of vision* 8:1-10.
- Cicchini GM, Anobile G, Burr DC (2014) Compressive mapping of number to space reflects dynamic encoding mechanisms, not static logarithmic transform. *Proc Natl Acad Sci USA* 111:7867-7872.
- Cousineau D (2005) Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorial in Quantitative Methods for Psychology* 1:42-45.
- Curtis CE, D'Esposito M (2003) Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences* 7:415-423.
- D'Esposito M, Postle BR (2015) The cognitive neuroscience of working memory. *Annu Rev Psychol* 66:115-142.
- Danker JF, Anderson JR (2010) The ghosts of brain states past: remembering reactivates the brain regions engaged during encoding. *Psychol Bull* 136:87-102.
- De Lange FP, Rahnev DA, Donner TH, Lau H (2013) Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. *J Neurosci* 33:1400-1410.
- Den Ouden HEM, Friston KJ, Daw ND, McIntosh AR, Stephan KE (2009) A dual role for prediction error in associative learning. *Cereb Cortex* 19:1175-1185.
- Den Ouden HEM, Daunizeau J, Roiser J, Friston KJ, Stephan KE (2010) Striatal prediction error modulates cortical coupling. *J Neurosci* 30:3210-3219.
- Doherty JR, Rao A, Mesulam MM, Nobre AC (2005) Synergistic effect of combined temporal and spatial expectations on visual attention. *J Neurosci* 25:8259-8266.

REFERENCES

- Dong DW, Atick JJ (1995) Statistics of natural time-varying images. *Network: Computation in Neural Systems* 6:345-358.
- Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415:429-433.
- Fetsch CR, Pouget A, DeAngelis GC, Angelaki DE (2012) Neural correlates of reliability-based cue weighting during multisensory integration. *Nat Neurosci* 15:146-156.
- Fischer J, Whitney D (2014) Serial dependence in visual perception. *Nat Neurosci* 17:738-743.
- Friston (2005) A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci* 360:815-836.
- Friston (2009a) The free-energy principle: a rough guide to the brain? *Trends Cogn Sci* 13:293-301.
- Friston, Holmes A, Poline JB, Price CJ, Frith CD (1996) Detecting Activations in PET and fMRI: Levels of Inference and Power. *Neuroimage* 4:223-235.
- Friston K (2009b) The free-energy principle: a rough guide to the brain? *Trends Cogn Sci* 13:293-301.
- Friston KJ, Josephs O, Rees G, Turner R (1998) Nonlinear event-related responses in fMRI. *Magn Reson Med* 39:41-52.
- Fuster JM (1973) Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *J Neurophysiol* 36:61-78.
- Gao J, Wong-Lin K, Holmes P, Simen P, Cohen JD (2009) Sequential effects in two-choice reaction time tasks: Decomposition and synthesis of mechanisms. *Neural Comput* 21:29.
- Garcia-Perez M (1998) Forced-choice staircases with fixed step sizes: asymptotic and small-sample properties. *Vision Res* 38:1861-1881.
- Gazzaley A, Nobre AC (2012) Top-down modulation: bridging selective attention and working memory. *Trends Cogn Sci* 16:129-135.
- Gold JJ, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535-574.
- Gordon AM, Rissman J, Kiani R, Wagner AD (2014) Cortical reinstatement mediates the relationship between content-specific encoding activity and subsequent recollection decisions. *Cereb Cortex* 24:3350-3364.
- Gulbainate R, Johnson A, de Jong R, Morey CC, van Rijn H (2014) Dissociable mechanisms underlying individual differences in visual working memory capacity. *Neuroimage* 99:197-206.
- 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. *J Neurosci* 31:6339-6352.
- Harrison SA, Tong F (2009) Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458:632-635.
- Heeger DJ (1992) Half-squaring in responses of cat striate cells. *Visual Neurosci* 9:427-443.
- Helmholtz H (1867) *Handbuch der physiologischen optik*. Leipzig: L. Voss.
- Henson RN (2015) Repetition suppression to faces in the fusiform face area: A personal and dynamic journey. *Cortex*:In Press.
- Hsu YF, Le Bars S, Hamalainen JA, Waszak F (2015) Distinctive Representation of Mispredicted and Unpredicted Prediction Errors in Human Electroencephalography. *J Neurosci* 35:14653-14660.
- Hupé JM, James AC, Payne BR, Lomber SG, Girard P, Bullier J (1998) Cortical feedback improves discrimination between figure and background by V1, V2 and V3 neurons. *Nature* 394:784-787.
- Knill DC, Pouget A (2004) The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci* 27:712-719.
- Kok P, Jehee JFM, De Lange FP (2012a) Less is more: Expectation sharpens representations in the primary visual cortex. *Neuron* 75:265-270.
- Kok P, Failing MF, De Lange FP (2014) Prior expectations evoke stimulus templates in the primary visual cortex. *J Cognitive Neurosci*.
- Kok P, Brouwer GJ, Van Gerven MAJ, De Lange FP (2013) Prior expectations bias sensory representations in visual cortex. *J Neurosci* 33:16275-16284.
- Kok P, Rahnev D, Jehee JFM, Lau HC, De Lange FP (2012b) Attention reverses the effect of prediction silencing sensory signals. *Cereb Cortex* 22:2197-2206.
- Krekelberg B, Boynton GM, van Wezel RJ (2006) Adaptation: from single cells to BOLD signals. *Trends Neurosci* 29:250-256.

- Larsson J, Smith AT (2012) fMRI repetition suppression: neuronal adaptation or stimulus expectation? *Cereb Cortex* 22:567-576.
- Lavie N (2005) Distracted and confused?: selective attention under load. *Trends Cogn Sci* 9:75-82.
- Lavie N, Hirst A, de Fockert JW, Viding E (2004) Load theory of selective attention and cognitive control. *Journal of experimental psychology General* 133:339-354.
- Law CT, Gold JI (2008) Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat Neurosci* 11:505-513.
- Lee TS, Mumford D (2003) Hierarchical Bayesian inference in the visual cortex. *J Opt Soc Am A Opt Image Sci Vis* 20:1434-1448.
- Liberman A, Fischer J, Whitney D (2014) Serial dependence in the perception of faces. *Curr Biol* 24:2569-2574.
- Marini F, Demeter, E., Roberts, K. C., Chelazzi, L., Woldorff, M. G. (2014) The context-dependent interplay between proactive and reactive mechanisms for distraction filtering in the human brain. Program No. 835.20. In: Neuroscience Meeting Planner. Washington, DC: Society for Neuroscience, 2014. Online., p 1373. Washington, DC.
- Maunsell JHR, Treue S (2006) Feature-based attention in visual cortex. *Trends Neurosci* 29:317-322.
- Meyer T, Olson CR (2011) Statistical learning of visual transitions in monkey inferotemporal cortex. *Proc Natl Acad Sci USA* 108:19401-19406.
- Morey RD (2008) Confidence intervals from normalized data: A correction to Cousineau (2005). *Tutorial in Quantitative Methods for Psychology* 4:61-64.
- Mumford D (1992) On the computational architecture of the neocortex. II. The role of cortico-cortical loops. *Biol Cybern* 66:241-251.
- Mumford JA, Turner BO, Ashby FG, Poldrack RA (2012) Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses. *NeuroImage* 59:2636-2643.
- Murray MM, Wylie GR, Hidding BA, Javitt DC, Schroeder CE, Foxe JJ (2002) The spatiotemporal dynamics of illusory contour processing: Combined high-density electrical mapping, source analysis, and functional magnetic resonance imaging. *J Neurosci* 22:5055-5073.
- Niculescu-Mizil A, Caruana R (2005) Predicting good probabilities with supervised learning. In: Proceedings of the 22nd International Conference on Machine Learning, pp 625-632. Bonn, Germany.
- Nienborg H, Cumming BG (2009) Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* 459:89-92.
- Pajani A, Kok P, Kouider S, de Lange FP (2015) Spontaneous Activity Patterns in Primary Visual Cortex Predispose to Visual Hallucinations. *J Neurosci* 35:12947-12953.
- Pearson J, Westbrook F (2015) Phantom perception: voluntary and involuntary nonretinal vision. *Trends Cogn Sci* 19:278-284.
- Platt JC (2000) Probabilities for SV Machines. In: *Advances in Large Margin Classifiers* (Smola AJ, Bartlett P, Schölkopf B, Schuurmans D, eds), pp 61-74. Cambridge, Massachusetts: The MIT Press.
- Posner MI, Petersen SE (1990) The attention system of the human brain. *Annu Rev Neurosci* 13:25-42.
- Rahnev D, Lau H, De Lange FP (2011) Prior expectation modulates the interaction between sensory and prefrontal regions in the human brain. *J Neurosci* 31:10741-10748.
- Rahnev D, Koizumi A, McCurdy LY, D'Esposito M, Lau H (2015) Confidence Leak in Perceptual Decision Making. *Psychol Sci* 26:1664-1680.
- Rao RPN, Ballard DH (1997) Dynamic model of visual recognition predicts neural response properties in the visual cortex. *Neural Comput* 9:721-763.
- Rao RPN, Ballard DH (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci* 2:79-87.
- Reynolds JH, Heeger DJ (2009) The Normalization Model of Attention. *Neuron* 61:168-185.
- Riggall AC, Postle BR (2012) The relationship between working memory storage and elevated activity as measured with functional magnetic resonance imaging. *J Neurosci* 32:12990-12998.
- Rouder JN, Speckman PL, Sun D, Morey RD, Iverson G (2009) Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic bulletin & review* 16:225-237.

REFERENCES

- SanMiguel I, Widmann A, Bendixen A, Trujillo-Barreto N, Schröger E (2013) Hearing silences: human auditory processing relies on preactivation of sound-specific brain activity patterns. *J Neurosci* 33:8633-8639.
- Schlichting ML, Mumford JA, Preston AR (2015) Learning-related representational changes reveal dissociable integration and separation signatures in the hippocampus and prefrontal cortex. *Nature communications* 6:8151.
- Schoenmakers S, Guclu U, van Gerven M, Heskes T (2014) Gaussian mixture models and semantic gating improve reconstructions from human brain activity. *Front Comput Neurosci* 8:173.
- Schröger E, Marzecova A, SanMiguel I (2015) Attention and prediction in human audition: a lesson from cognitive psychophysiology. *The European journal of neuroscience* 41:641-664.
- Serences JT, Ester EF, Vogel EK, Awh E (2009) Stimulus-specific delay activity in human primary visual cortex. *Psychological Science: A Journal of the American Psychological Society / APS* 20:207-214.
- Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB (1995) Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268:889-893.
- Solomon SG, Kohn A (2014) Moving sensory adaptation beyond suppressive effects in single neurons. *Curr Biol* 24:R1012-1022.
- Sotiropoulos G, Seitz AR, Series P (2011) Changing expectations about speed alters perceived motion direction. *Current biology : CB* 21:R883-884.
- St. John-Saaltink E, Utzerath C, Kok P, Lau H, De Lange F (2015) Expectation suppression in early visual cortex depends on task set. *PLoS ONE* 10:1-14.
- Sterzer P, Frith C, Petrovic P (2008) Believing is seeing: expectations alter visual awareness. *Curr Biol* 18:R697-698.
- Stocker AA, Simoncelli EP (2006) Noise characteristics and prior expectations in human visual speed perception. *Nat Neurosci* 9:578-585.
- Stone LS, Thompson P (1992) Human speed perception is contrast dependent. *Vision Res* 32:1535-1549.
- Summerfield C, Egner T (2009) Expectation (and attention) in visual cognition. *Trends Cogn Sci* 13:403-409.
- Summerfield C, de Lange FP (2014) Expectation in perceptual decision making: neural and computational mechanisms. *Nat Rev Neurosci* 15:745-756.
- Summerfield C, Trittschuh EH, Monti JM, Mesulam MM, Egner T (2008) Neural repetition suppression reflects fulfilled perceptual expectations. *Nat Neurosci* 11:1004-1006.
- Summerfield C, Egner T, Greene M, Koechlin E, Mangels J, Hirsch J (2006) Predictive codes for forthcoming perception in the frontal cortex. *Science* 314:1311-1314.
- Thompson P (1982) Perceived rate of movement depends on contrast. *Vision Res* 22:377-380.
- Todorovic A (2015) Predictive adaptation in early auditory processing. In, p 139. Enschede, the Netherlands: Radboud University.
- Todorovic A, De Lange FP (2012) Repetition suppression and expectation suppression are dissociable in time in early auditory evoked fields. *J Neurosci* 32:13389-13395.
- Todorovic A, Van Ede F, Maris E, De Lange FP (2011) Prior expectation mediates neural adaptation to repeated sounds in the auditory cortex: an MEG study. *J Neurosci* 31:9118-9123.
- Turk-Browne NB, Scholl BJ, Chun MM, Johnson MK (2009) Neural evidence of statistical learning: efficient detection of visual regularities without awareness. *J Cognitive Neurosci* 21:1934-1945.
- Watson AB, Pelli DG (1983) QUEST: a Bayesian adaptive psychometric method. *Percept Psychophys* 33:113-120.
- Wei XX, Stocker AA (2015) A Bayesian observer model constrained by efficient coding can explain 'anti-Bayesian' percepts. *Nat Neurosci* 18:1509-1517.
- Weiss Y, Simoncelli EP, Adelson EH (2002) Motion illusions as optimal percepts. *Nat Neurosci* 5:598-604.
- Wimmer K, Compte A, Roxin A, Peixoto D, Renart A, de la Rocha J (2015) Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area MT. *Nature communications* 6:6177.
- Wyart V, Nobre AC, Summerfield C (2012) Dissociable prior influences of signal probability and relevance on visual contrast sensitivity. *Proc Natl Acad Sci USA* 109:3593-3598.

- Xing D, Ouni A, Chen S, Sahmoud H, Gordon J, Shapley R (2015) Brightness-color interactions in human early visual cortex. *J Neurosci* 35:2226-2232.
- Yi DJ, Woodman GF, Widders D, Marois R, Chun MM (2004) Neural fate of ignored stimuli: dissociable effects of perceptual and working memory load. *Nat Neurosci* 7:992-996.
- Yuille A, Kersten D (2006) Vision as Bayesian inference: analysis by synthesis? *Trends Cogn Sci* 10:301-308.

Nederlandse samenvatting

Heb je ooit een slok genomen van iets waarvan je dacht dat het water was, maar wat koffie bleek te zijn? Als je ooit een dergelijke ervaring hebt gehad, dan weet je dat het iets wat je normaal gesproken lekker vindt heel onprettig kan maken. Dit simpele voorbeeld illustreert dat verwachtingen je sensorische systeem kunnen voorbereiden op een bepaalde invoer, en dat hoe je iets waarneemt ervan afhangt of het verwacht of verrassend was.

Een ander voorbeeld van hoe verwachtingen je waarneming kunnen beïnvloeden zijn visuele illusies. Kijk maar eens naar de afbeelding hieronder. De maan aan de hemel is groter dan de maan tussen de treinrails, toch? In werkelijkheid zijn ze even groot, maar het perspectief aanwezig in de afbeelding leidt ertoe dat ons visueel systeem een andere conclusie trekt. Dit komt doordat we door onze uitgebreide ervaring met de wereld hebben geleerd dat waarnemen op afstand ervoor zorgt dat dingen er kleiner uitzien dan wanneer we ze van dichtbij bekijken. Doordat het perspectief van de afbeelding ervoor zorgt dat de maan in de hemel veel verder weg lijkt dan de maan tussen de rails, nemen we hem als groter waar.



Image credit: NASA

Figuur 1. Geloof je ogen niet.

Perspectief kan ervoor zorgen dat identieke voorwerpen verschillende afmetingen lijken te hebben. Onze kennis dat afstand een voorwerp kleiner doet lijken zorgt ervoor dat de maan aan de hemel veel groter lijkt dan de maan tussen de treinrails.

Visuele informatie kan heel complex zijn, met meerdere voorwerpen op verschillende afstanden, elk met andere belichting, waarbij het ene voorwerp ook nog eens deels het zicht op het ander kan ontnemen. Het gebruik van onze kennis van de wereld om deze wirwar te kunnen ontrafelen zorgt ervoor dat onze waarneming soms onderhevig is aan illusies, zoals hierboven beschreven, maar dit kleine nadeel weegt niet op tegen het grote voordeel dat we hierdoor complexe en vaak ambiguë visuele informatie efficiënt (en gewoonlijk correct!) kunnen verwerken.

Een ander soort verwachting dat we over onze omgeving hebben is dat de meeste dingen van het een op het andere moment grotendeels hetzelfde blijven: een fiets verandert niet opeens in een auto. Stel je voor dat je de skyline van New York bewondert vanaf de oever van de Hudson rivier en dat je even wegstapt om je camera erbij te pakken: wanneer je weer opkijkt zal de skyline er nog net zo uitzien als vlak daarvoor. Dit soort verwachting kan er zelfs voor zorgen dat dingen meer lijken op wat je net gezien hebt dan dat ze eigenlijk waren! Dit lijkt misschien raar, maar dit heeft een duidelijk voordeel: heel snelle veranderingen in de wereld zijn meestal het gevolg van ruis (stel je voor dat je naar de skyline kijkt tijdens een regenstorm), en daardoor kun je voorwerpen beter herkennen wanneer je deze snelle veranderingen weg filtert. In Hoofdstuk 2 heb ik dit effect onderzocht, en vond ik dat mensen hetzelfde voorwerp anders waarnemen afhankelijk van wat ze net daarvoor hadden gezien. En dat niet alleen, dit effect was ook zichtbaar in de manier waarop het voorwerp door de hersenen werd verwerkt, zelfs al in de vroegste stadia! Dit suggereert dat de verwachtingen die we hebben over onze omgeving op basis van onze kennis van de wereld kan beïnvloeden hoe sensorische informatie “eruit ziet” volgens ons brein, misschien zelfs al vanaf het moment dat het binnenkomt.

Ondanks dat de meeste van dit soort invloeden op onze waarneming onbewust plaatsvinden, en daardoor automatisch lijken, zou het kunnen dat ze afhangen van hoe afgeleid we zijn door andere dingen die op hetzelfde moment gebeuren. Een manier om dit te onderzoeken is door te kijken hoe je brein reageert op voorspelbare informatie wanneer je afgeleid wordt door taken die verschillende cognitieve systemen bezighouden. Stel je bijvoorbeeld voor dat je twee dingen tegelijkertijd doet, zoals het volgen van een gesprek terwijl je i) een telefoonnummer onthoudt tot het je lukt om het te draaien, of ii) door een zware regenstorm rijdt. Taak i) houdt je werkgeheugen bezig, terwijl taak ii) je perceptuele systemen bezighoudt. In Hoofdstuk 3 liet ik proefpersonen dit soort taken doen, en heb ik onderzocht of dat beïnvloedde hoe de hersenen reageerden op (irrelevante) voorspelbare informatie in de omgeving. De resultaten van dit onderzoek lieten zien dat het brein gebruik maakt van voorspelbare informatie, ook wanneer je afgeleid wordt door een andere taak, maar alleen als deze taak je werkgeheugen niet belast. Dit suggereert dat het vormen en toepassen van verwachtingen waarschijnlijk een bepaalde mate van werkgeheugen vereist.

In conclusie, in dit proefschrift heb ik onderzocht hoe onze verwachtingen beïnvloeden hoe ons brein reageert op sensorische informatie, en hoe dit vervolgens onze waarneming vormt. Om precies te zijn heb ik onderzocht hoe visuele invoer verwerkt wordt door ons brein, afhankelijk van de omstandigheden waarin ze binnenkomt: wat ervóór kwam, wat je verwacht te zien, en welke cognitieve systemen je tot je beschikking hebt.

Acknowledgements

This thesis would have absolutely not been possible without the incredible support of a large number of people.

Floris, thank you for being an excellent role model and supervisor. I greatly admire your approach to scientific research, which combines a genuine enjoyment of your work, an unwavering engagement with the subject matter, and a sharp mind that would be intimidating if not for your sociable nature that makes working with you so enjoyable. You have always been encouraging, understanding, and supportive, and you have challenged me to grow both as a researcher and a person. I appreciate that you are never too busy to express your interest in how people are doing, however tight your schedule is!

You have brought together a great group of people who, like you, are inquisitive, clever, and engaged, and great fun to work with. To all members of the predatt group, past and present, who have shared time with me in the lab: Ana, Freek, Sasha, Peter, Jolien, Loek, Simon, Claudia, Matthias E., Anke Marit, Erik W., Pim, Erik M., Christian, Lieke, Matthias F., Marisha, Yulia, Jasper, Patricia, Felix, Paige, Poppy, Bronagh, Simone, Auréliane, Nathalie, Vanessa, Larissa, Thomas, Heidi, Marius, Lies, Michel, and Remco, thank you for sharing your ideas, constructive feedback, collaboration and baked goods! Christian, it was a valuable and enjoyable experience to collaborate with you on ‘impexp’ – thank you for your hard work and dedication. Paige, I learnt a lot from being involved in the rapid data collection for your rapid TR project, and I enjoyed collaborating with you. Claudia and Anke Marit, it was fun organizing one of the retreats together. Creating a human predictive coding model was a wild, but educational, idea!

Hakwan, we only ended up seeing each other in person a handful of times, but your input and the counterpoint your ideas provided to my discussions with Floris was very helpful to me – I learnt a lot through the email exchanges between the three of us! Your lightning-fast responses to emails consistently impressed me, as did your in-depth responses to my questions (especially about the VR model during the design for the first project). I always valued your words of encouragement and support that you offered through almost every communication.

The Donders is a phenomenal place to do research, and a stimulating, challenging, and well-supported environment to be a PhD candidate. Thank you to Peter and David for setting up the centre and being its co-captains and visionaries! I think that the collaborative and constructively-critical atmosphere that characterizes the

ACKNOWLEDGEMENTS

DCCN can in part be attributed to the PPM policy. I appreciate that you keep an open door and mind to feedback and suggestions about how the centre could be improved. An extra thank you to Peter for agreeing to be my promotor.

Tildie, Arthur, Sandra, Nicole, Ayse, Chantal, and Berend, I am aware that there were innumerable times that you made my life as a PhD and employee substantially easier; and this holds true for practically every person who has spent time in the centre. Tildie, thank you so much for fielding all my questions, completing all my credit card forms, and always having a minute to exchange stories from our latest (or upcoming) vacation. Nicole, Sandra, and Ayse, thank you for your friendly, helpful, and understanding responses to all participant and administration-related issues.

Paul, thanks for your dedication to keeping the MR lab running smoothly and efficiently, and for maintaining your relaxed manner when researchers get stressed. I enjoyed being on the team of 'super certified' users, despite how often things seemed to go awry when I was covering for you.

To the TG, Erik, Marek, Uriël, Rene, Edward, Mike, Jessica, Hong, and Sander, a huge thanks to you for all your hard work. The technical support at the DCCN is the finest I have encountered, and your proactive approach allows the rest of us to focus on our research.

Lucia, thank you for your dedication to keeping the DCCN looking so good! I am impressed by your eye for detail, and your cheerful and friendly disposition.

Roel, I came to you with many doubts, and your sincere and pragmatic advice helped greatly.

Clauda & Lieke, thank you for agreeing to be my paranymphs and for supporting me on this day!

Thanks to the organizing committees of various DCCN events that have increased the friendly and social atmosphere of the centre. A special thanks to the fellow Canadians Lauren, Nietzsche, and Ian who helped organize the Christmas dinner back in 2012, and who celebrated Thanksgivings together in the years following. Thanks to Ian as well for explaining the intricacies of US immigration and visa policy to me!

To my officemates in 2.24, 170 and 2.269, Bart, Madelon, Kim, Ilke, Mieke, Sander, Dan, Mahur, and Diana, thanks for sharing your tips and stories, answering my questions, and taking care of my plants when I was away.

To everyone I have exchanged stories, interesting facts, debates and frustrations with, Lauren, Ian, Tom, Sander, Ruben, Hanneke, Mathilde, Guillaume, Irati, Maria, Anna, Catalina, Silvy, Eelke, Ruud, Jim, Til, Johanna, Linda, Isabella, Markus, Jenni, Rasim, Jose, Tobias, Eelco, Verena, Lisa, Linda, Rene, Iska, Alina, Atsuko, Rene, Flora, Jeanette, Nietzsche, Tim, Jeroen, Richard, Lieneke, Monja, Vincent, Stephen, Payam, Andrea, Zahra, Frauke, Branka, Iris, Martin, David, and everyone else I have forgotten to mention here, thank you!

Heleen, I was so lucky to walk into your office to inquire about your internship positions: your supervision and mentorship started me on the path to becoming a researcher. Thank you for your encouragement and support! Tom, it was great to work with you in Heleen's lab, and I have you to thank for putting Nijmegen on my map! Also thanks to you and Iska as well for hosting me during my first Day Out. Anouk, your dedication to and enthusiasm for your research were very inspiring to me during my master's thesis project with you. Thank you for being my advocate.

Anna, it was great to get to know you and exchange ideas about predictive coding at various conferences.

To the girls in Amsterdam, Daniela, Eftychia, Natalie, Vanessa and Astrid: the master's program is what started this all! I have had such great times with you over the years; you became my new roots in the Netherlands.

Max, your solid friendship has seen me through so much. I'm so glad we've stayed close despite the distance.

Coby and Lammert, Klaas and Anke, Saskia and Roel, thank you for welcoming me into your families, and for all of your interest and kindness. This holds for the rest of the Veenstra and Kok family as well! To the Saaltinks, Elmie, Dale, Christopher, Lia, Rik, Brenda, Robyn, Adam, Heath, Clara, Emma, Mark, Jane, Philip and Rebecca, thank you for all of your support, interest, and being so much fun to hang out with! Grootmoeder and Grootvader, you are such role models for me, and inspirations for how to live life to the fullest. Thank you for all your encouragement. Colleen and Jon, I'm so grateful for your support, particularly during my early years. Kelly and Andre, it is amazing to have a home in Toronto where I'm welcomed and loved and fully feel at home. You have been a huge support to me. Knowing that Owen is well

ACKNOWLEDGEMENTS

loved and supported has made it easier for me to be so far away. Thanks for all the fruit salads Dee!

O, I'm so glad we've continued to share things we love, made us laugh, and made us angry. It's always great when hanging out doesn't involve talking through computers and telephones. Thanks for always listening. Papa, thank you for all your support and encouragement. You fostered my analytical thinking, and developed my writing through editing most of my undergrad essays! Mom, you fostered my creativity, determination, and love of food. I know you'd be so proud.

Peter, to say I couldn't have done it without you would be a huge understatement. You have supported the creation of this thesis in about every way imaginable. You are the most patient, understanding, supportive, and clever person I know. Thank you for your unwavering belief in me. I'm so glad that we're on this adventure together.

List of Publications

St. John-Saaltink E, Kok P, Lau HC, De Lange FP (2016) Serial Dependence in Perceptual Decisions is Reflected in Activity Patterns in Primary Visual Cortex. *Journal of Neuroscience*, 36(23): 618609192. doi:10.1523/JNEUROSCI.4390-15.2016

St. John-Saaltink E, Utzerath, C, Kok P, Lau HC, De Lange FP (2015) Expectation suppression in early visual cortex depends on task set. *PLoS ONE* 10(6): e0131172. doi:10.1371/journal.pone.0131172

Van Loon AM, Knapen T, Scholte HS, St. John-Saaltink E, Donner TH, Lamme VA (2013) GABA Shapes the Dynamics of Bistable Perception. *Current Biology* 23(9):823-7. doi: 10.1016/j.cub.2013.03.067

In Preparation

Utzerath C, St. John-Saaltink E, Buitelaar J, de Lange FP (submitted) Separable effects of expectation and repetition of objects at different levels of the visual processing hierarchy.

Scalf, PE, St. John-Saaltink E, Barth M, Lau HC, de Lange FP (submitted) Time-resolved fMRI reveals serial shifting of attention divided across multiple locations.

Biography

Elexa St. John-Saaltink was born on May 5, 1985 in Toronto, Canada. In 2008, she completed a Bachelor of Science at University of Toronto, with majors in Psychology and Philosophy (with distinction). Driven in part by a desire to connect to her Dutch heritage, she spent two semesters abroad at the University of Amsterdam during the last two years of her undergraduate studies. Through a bit of determination she was allowed to participate in some of Masters' level courses (offered in English), and it was here that she first experienced the excitement of reading research articles first hand, and became eager to get involved with research herself. In 2009, she moved to the Netherlands to return to the University of Amsterdam as a student in the Psychology Research Master programme, which she completed in 2011 (cum laude). She began research in Floris de Lange's lab in January 2012 on a co-supervised project with Hakwan Lau about the neural processes involved in the top-down factors that influence visual perception. Currently, she is extending her research background into the field of user experience research and design, and learning web development.

When the past influences the present:

Modulations of the sensory response by prior knowledge and task set

© Elexa St. John-Saaltink

ISBN 978-94-6284-083-6

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/donders-graduate/>

DONDERS

I N S T I T U T E

ISBN 978-94-6284-083-6

Radboud University  Radboudumc