Temporal but not Spatial Expectation Modulates Bottom-Up Attention

Master thesis

Author: Felix H. Klaassen (s4126599)

Supervision: Erik Meijs, Simon van Gaal, & Floris de Lange

Donders Institute for Brain, Cognition & Behavior, Centre for Cognitive Neuroimaging

Radboud University Nijmegen, The Netherlands

August 10th, 2016

Abstract

In the attentional literature there has been considerable debate about whether bottom-up attentional processes are purely stimulus driven or dependent on top-down sets and goals. One possible top-down factor that might influence bottom-up attentional processing is expectation. In this study we investigate the relationship between bottom-up attention and expectation, specifically asking whether bottom-up attention can be explained by a prediction-error. We hypothesize that unexpected stimuli generate a larger prediction error than expected stimuli, consequently leading to more attentional capture. In two experiments we used an exogenous cueing paradigm to investigate if spatial and temporal expectations about a distracting cue modulate the amount of attentional capture. In the first experiment we investigated the effect of spatial expectation on bottom-up attention, but found no evidence for an interactive relationship. In the second experiment we focused on the temporal predictability of the cue, and found a modulation of bottom-up attention by cue predictability. Specifically, unpredictable cues lead to more attentional capture compared to predictable cues, though only for a specific range of cue-target stimulus onset asynchronies. With these findings, we provide a first direct indication that expectation influences bottom-up attention, although the exact mechanism underlying the modulation is not yet clear. We propose a modulatory account of the interaction between expectation and bottom-up attention, and suggest that bottom-up attention is in principle stimulus driven but can be modulated by expectation if the temporal relationship between the cue and the target is optimal.

Keywords: bottom-up attention, expectation, predictive coding, prediction error

In order to properly perform everyday tasks and successfully navigate through a complex and ever changing environment, it is crucial to be able to ignore sensory information that is task-irrelevant, and to emphasize information that is task-relevant. We are capable of doing this by means of our attentional system which enables us to, as von Helmholtz (1867) already stated, shine a mental spotlight on that which we deem interesting, consequently facilitating efficient processing of the attended information (Baluch & Itti, 2011; Posner, Snyder, & Davidson, 1984). This voluntarily deployed attention (also called top-down or endogenous attention) can be contrasted against its involuntary counterpart: bottom-up or exogenous attention (Posner, 1980). Bottom-up attention is, contrary to top-down attention, not oriented as a consequence of a deliberate decision to do so. It is a reflexive orientation of attention caused by a sudden salient event in the environment (think for example of a loud noise or a bright light; Corbetta & Shulman, 2002). The functional role of bottom-up attention presumably lies within making the observer aware that something important might be happening in the environment, by biasing attention towards salient information.

1.1 Bottom-up attention

Studies on bottom-up attention show its practical implications for our behavior. If a distracting stimulus (cue or distractor) suddenly appears and captures our attention, it can have striking consequences for task performance. Typically, valid cues (i.e., cues that appear at the same spatial location as a task-relevant stimulus or target) result in faster reaction times (RTs) to the target than invalid cues. This *attentional capture* effect of validity is due to the cue biasing attention either towards or away from the target, subsequently making it easier or harder to respond quickly to it (Posner, 1980; Posner, et al., 1984).

The mechanism underlying bottom-up attention has been subject of a large debate in which so far no real consensus has been reached. On the one hand, bottom-up attention is thought to be purely stimulus driven and independent of top-down intentions and goals (e.g., Belopolsky, Schreij, & Theeuwes, 2010; Pinto, Van der Leij, Sligte, Lamme, & Scholte, 2013; Theeuwes 1992; Theeuwes, 2004; Theeuwes & Godijn, 2002). On the other hand, there have been several researchers who proposed that bottom-up attention is not independent from top-down mechanisms and that it can be modulated or even suppressed by task goals and top-down attentional sets (e.g., Connor, Egeth, & Yantis, 2004; Folk, Remington, & Johnston, 1992; Gibson & Kelsey, 1998). While both view-points have their merits, an alternative way to approach this dichotomy is perhaps to reconcile the two theories. It may be that top-down signals are able to influence bottom-up attentional processes, but are not capable of completely diminishing them (Theeuwes & van der Burg, 2007). Expectation is one top-down mechanism that might in particular play such a role.

1.2 Expectation

With the rise of predictive coding theory, the idea that the brain is a prediction machine has become increasingly popular. The model assumes that the brain continuously compares primary expectations of the state of the environment (priors; E_p) with the actual state of the environment (E_a), as observed through the sensory systems. Whenever E_p does not match E_a a surprise signal called a prediction error (PE) is generated, signaling that the prior should be adjusted (den Ouden, Kok, & de Lange, 2012). Careful monitoring of the state of the prior (i.e., adjusting it to new input if necessary) makes sensory processing more efficient by suppressing what is expected and enhancing what is not (Summerfield, et al., 2006).

Indeed, predictive mechanisms have been associated with a wide range of cognitive domains such as auditory and visual perception, reward processing and memory, and have been found to be widely implemented in the brain (e.g., see Baldeweg, 2006; Friston & Kiebel, 2009; Hohwy, Roepstorff, & Friston, 2008; Kumaran & Maguire, 2009; O'Doherty, et al., 2006; Rao & Ballard, 1999; Todorovic, van Ede, Maris, & de Lange, 2011). For example, in electrophysiological methods PEs are characterized by the mismatch negativity (MMN, den Ouden, Kok, & de Lange, 2012; Summerfield & Koechlin, 2008; Wacongne, Changeux, & Dehaene, 2012) and P300 (Polich, 2007) components. Additionally, in functional magnetic resonance imaging (fMRI) prediction errors have been associated with an increase of the blood oxygenation level dependent (BOLD) response (den Ouden, Friston, Daw, Mcintosh, & Stephan, 2009). Exactly what the brain responses that are associated with PEs represent is still subject of debate. It has long been thought that the suppression of expected information in low-level sensory areas reflects the brain 'explaining away' expected stimulus information. However, it has been shown that the reduced response amplitude corresponds to a sharpening of the stimulus representation, enhancing the specificity by which neurons respond to incoming sensory information (Kok, Jehee, & de Lange, 2012), thereby facilitating sensory processing. It might be that expectation can influence bottom-up attentional processing in a similar manner by calling for more attentional resources if a stimulus does not fall within the sharpened representation. Still, this is yet to be investigated as the exact relationship between attention and expectation is largely unclear.

1.3 Expectation & attention

At first sight, predictive and attentional mechanisms may seem comparable with each other: They both enhance information that might be relevant (i.e., surprising or interesting) and suppress information that might not. Indeed, both increased attention (Boynton, 2009) and violations of expectations (Rao & Ballard, 1999) are associated with an increased brain response in sensory areas. Based on this, one might suspect that the two mechanisms are additive. However, attention has been found to be capable of reversing predictive suppression of sensory processing (Kok, Rahnev, Jehee, Lau, & de Lange, 2012), suggesting an interaction between the two mechanisms rather than an additive relationship.

Some research has indeed proposed an interactive relationship between expectation and attention, which has also been referred to as the predictive coding account of attention. For example, Feldman and Friston (2010) formed a theoretical framework of this account in which they argued that attention can be seen as performing an inference on the level of uncertainty of perceptual information (see also Friston, 2009). Empirical support for this model however only focused on top-down attention (Feldman & Friston, 2010). Hohwy (2012) did briefly address bottom-up attention in the context of this framework, proposing that salient stimuli result in a high level of sensory input, consequently increasing the gain of associated PE units. In accordance with this, Spratling (2012) also proposed PEs to signal salience, redefining salience not only as the result of low-level stimulus characteristics (e.g., Li, 2002) but also as the extent to which a stimulus is unexpected.

Evidently, there have been some ideas on the interaction between expectation and attention, but it seems that so far the relationship itself has mostly been explored theoretically and mainly focused on top-down attention. A recent study did shine some light on the interaction between expectation and bottom-up attention, showing that infrequent but not frequent onsets capture attention and override top-down attentional sets (Folk & Remington, 2015). However, this effect depended greatly on the specific task-related properties of the cue and target, making it hard to distinguish between the contributions of different top-down factors (e.g., expectation vs. task set). A direct assessment of expectation as a modulator of bottom-up attention has thus not yet been done.

In the current study we aim to find out exactly how expectation and bottom-up attention relate to each other, assessing directly whether bottom-up attention is indeed purely stimulus driven or if it is instead caused by how (un)expected a stimulus is. Concretely, we investigate whether bottom-up attention is the result of a prediction error. The proposed mechanism entails that whenever a stimulus appears unexpectedly a PE will be generated, signaling that something surprising and therefore potentially relevant might be happening, as in accordance with Hohwy's (2012) and Spratling's (2012) ideas. Consequently, in order to properly deal with a possible interesting event, attentional resources will quickly be allocated towards the stimulus, enhancing target detection after valid cues and decreasing it after invalid cues. In order to test this we conduct two experiments, one looking at spatial and the other at temporal expectation by using an adapted version of Posner's exogenous cueing paradigm to measure bottom-up attention (Posner, 1980). In short, we hypothesize that unexpected cues lead to increased attentional capture and expected cues lead to decreased attentional capture. If our hypotheses are confirmed it would suggest that exogenous attention is dependent on prior expectations, providing support for the predictive coding account of attention. Insight into how exactly prior expectations are involved in exogenous orienting of attention could potentially lead to a re-evaluation of the dichotomy between bottom-up and top-down information processing, by reconciling both mechanisms and bringing us further towards a more complete model of the interplay between top-down knowledge on the one hand and bottom-up information processing on the other.

Experiment I: Spatial Prediction

Here, we investigate the influence of spatial predictability of a distracting cue on the amount of attentional capture. While it may seem intuitive that being able to predict where a distraction will appear will make it easier to ignore it, concrete empirical evidence has been lacking. We will investigate this by using an adapted version of a well-established exogenous cueing task (Posner, 1980). To address our research question we adjust this paradigm such that one cue position is more likely than the other. Additionally, we aim to isolate the effect of expectation and rule out the influence of other top-down factors such as task goals. To achieve this, the cues are made completely unrelated to the task. This way, the expectation of the cue location will be based on the pure statistical likelihood of where a cue will appear, and violations of this expectation will thus generate a PE that is purely based on the location of the cue. We predict that cues appearing on the expected location will generate a small PE, causing less attentional capture (i.e., a decrease of the cue validity effect), and that unexpected cues will generate a large PE, causing more attentional capture (i.e., an increase of the cue validity effect).

2. Methods

2.1 Participants

In the experiment 120 healthy Dutch speaking participants (89 females, age; $M \pm SD$, 22.6 ± 4.95) participated in return for money or course credit. All participants had normal or corrected-to-normal vision. Participants were not aware of the purpose of the study. All participants gave written informed consent prior to experiment initiation according to the institutional guidelines of the local ethics committee (CMO-Commissie Mensgebonden Onderzoek region Arnhem-Nijmegen, the Netherlands). Participants were randomly appointed to one of three groups, resulting in 40 participants per group (two experimental groups and one control group).

2.2 Material

The experiment was programmed and executed in MATLAB® using the Psychophysics Toolbox extensions (R2014b & R2011b, The MathWorks, Natick, Massachusetts; Brainard, 1997, Kleiner, Brainard, & Pelli, 2007). Computer hardware during data acquisition consisted of a Dell T3500 Workstation desktop running GNU/Linux (Cent OS 6.2), and a BENQ XL2420T 24" (1920 × 1080 pixels) 60 Hz LED TFT monitor. Two separate button boxes were used to record participants' responses. Additionally a chinrest was used to control the distance from which participants were viewing the monitor.

2.3 Design

2.3.1 Paradigm & stimuli. To study bottom-up attention effects, an adjusted version of the Posner exogenous cueing task was used (Posner, 1980). In our version, at the start of each trial a cue was flashed for 50 ms either above or below the fixation cross (subtending .5° × .5° visual angle). The cue consisted of an annulus with a distance from fixation of 5°, a diameter of 2° and line thickness of .2° visual angle. After a stimulus onset asynchrony (SOA) of 117 ms the target was presented. We chose this particular SOA because it has been shown that bottom-up facilitation effects are largest around SOA's of 100 ms (Dukewich & Klein, 2015; Klein, 2000).

To prevent backward masking of the cue on the target (see for example Kouider & Dehaene, 2007) we decided to take a slightly longer SOA than usual.

After the inter-stimulus interval (67 ms), the target was presented either above or below fixation (same distance from fixation as the cue). The target stimulus was an arrow pointing either leftward or rightward (subtending a w × h of .48° × .6° visual angle, a line thickness of .15° visual angle, and a 60° angle between the upper and bottom arrow lines). The target remained onscreen for 1000 ms or until a response was given. If the participant failed to respond in time the trial was labeled as a miss. After a response a variable inter-trial interval (ITI; 750 – 1500 ms) followed. Participants' task was to indicate via button presses whether the arrow was pointing leftward (left button) or rightward (right button) with the index fingers of both hands. All stimuli were black on a grey (RGB: 150, 150, 150) background (Fig. 1).

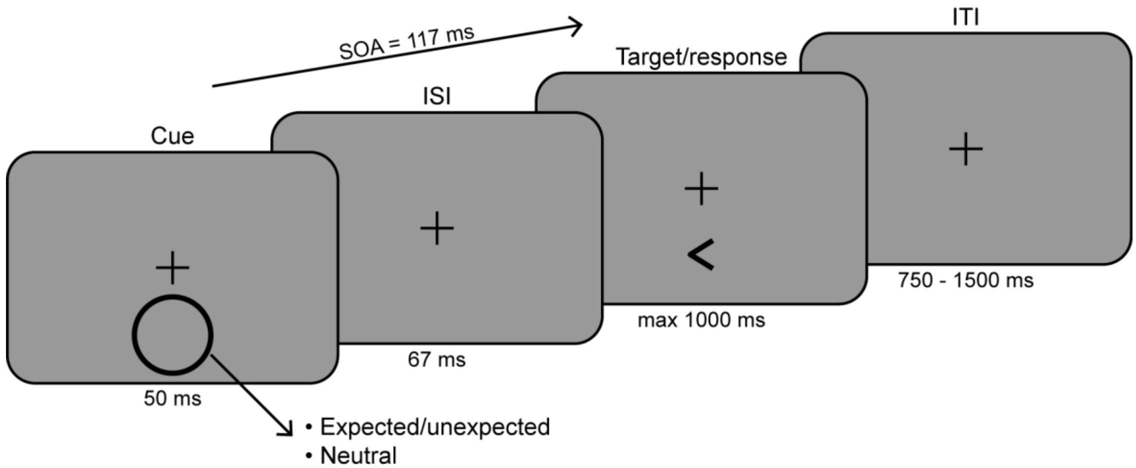

Figure 1. Experimental paradigm. Each trial started with an uninformative cue below or above fixation that remained onscreen for 50 milliseconds. The probability of a cue appearing at each of the locations was manipulated, depending on the experimental group participants were in. After an SOA of 117 ms the target was presented at either the same location (valid trial) or the opposite location (invalid trial) as the cue. The task was to discriminate between a leftward and rightward pointing arrow. In the figure an example of a valid trial is shown.

2.3.2 Factors & manipulations. The overall experiment consisted of a within-subjects design and featured the independent factors Validity (valid vs. invalid cues) and Expectation (expected vs. unexpected cues). Validity was manipulated within-subjects and was given to all participants. That is, in each trial the cue and target location could be either the same (valid trial) or opposite (invalid trial), with both types of trials occurring with equal probabilities (50%). The factor Expectation was only included for the two experimental groups. That is, for those participants we manipulated the location where the cue was most likely to appear. For one group the cue location was expected to be above fixation (i.e., the cue appeared above fixation for 80% of the trials and below for 20% of the trials) while for the other group the cue was expected to appear below fixation. In the control group the cue appeared at both locations with equal probability (50% of the trials for each location), therefore inducing no expectation of cue location.

Importantly, in all groups the target always appeared on either location 50% of the trials, and the amount of valid and invalid trials was always equal. The cue had thus no predictive relationship with either target identity or location. Dependent variables were reaction time (RT, relative to target onset) and error rate (ER). The total experiment consisted of 11 blocks of 80 trials including 1 practice block, and lasted approximately 40 minutes.

2.4 Procedure

All experiments were conducted in a dimly lit behavioral lab. In the instructions participants were fully informed on all the contingencies of the task and were made aware that the cue did not predict the target location and direction in any way. Participants were also asked to always keep their eyes focused on the fixation cross, centered on the screen. After instruction the participants were positioned into the chinrests (distance of 57 cm from the screen). Subsequently participants started with a block of 80 practice trials during which the door of the cubicle remained open, while the experiment leader waited outside. In this block the participants received trial-by-trial feedback about their performance. After the practice block the actual experiment started and the experiment leader closed the door. During the main task, participants received feedback about their performance after each block and subsequently a break of 20 seconds was initiated. In this period the participant was also sometimes asked to open the door to let the experiment leader in.

2.5 Data analysis

2.5.1 Preprocessing. RTs below 200 ms and measurements deviating more than 3 standard deviations from the mean were removed from the dataset. This amounted to an average removal of 11.15 trials per participant (1.4% of data set). ERs were calculated as the proportion of trials in which an incorrect response was given. Participants with an average ER of more than 20% would be excluded from the dataset, but no participants reached that threshold. For the statistical analysis only trials with correct responses were included. Preprocessing was done using MATLAB® (R2014b, The MathWorks, Natick, Massachusetts).

2.5.2 Statistical testing. To test the effect of Expectation on the Validity effect the data of the experimental groups (N = 80) were jointly put in a 2x2 Repeated Measures (RM) ANOVA with within-subjects variables Validity (valid vs. invalid) and Expectation (expected vs. unexpected) and dependent variables RT and ER. Additional (post-hoc) analyses to test for the Validity effect in the control group (N = 40) and to compare the experimental and control groups with each other were done with paired- and independent-sample t-tests. Finally, a Bayesian RM ANOVA was done to test the amount of evidence for including different factor models (i.e., additive vs. interactive factor effects). The probability of the prior for all models was $P(\mathrm{M}) = 0.2$, which is the default of the analysis software (JASP, 2016; but see also Love, et al., 2015; Morey &

Rouder, 2015; Rouder, Morey, Speckman, & Province, 2012). Data was analyzed with MATLAB® (R2014b, The MathWorks, Natick, Massachusetts) and IBM SPSS statistics 21.

3. Results

3.1. Reaction Times

3.1.1 Main analysis. We found significantly faster RTs for expected (M = 424, SD = 4.48) compared to unexpected (M = 433, SD = 4.69) cues (F (1, 79) = 89.231, p < .001, η^2_p = .530), and significantly faster RTs for valid (M = 414, SD = 4.05) compared to invalid (M = 443, SD = 5.17) cues (F (1, 79) = 217.149, p < .001, η^2_p = .733). Importantly, the RT difference between valid and invalid trials was equal for expected and unexpected cues, as shown by the absence of a significant Validity x Expectation interaction effect (p = .356, Fig. 2, left).

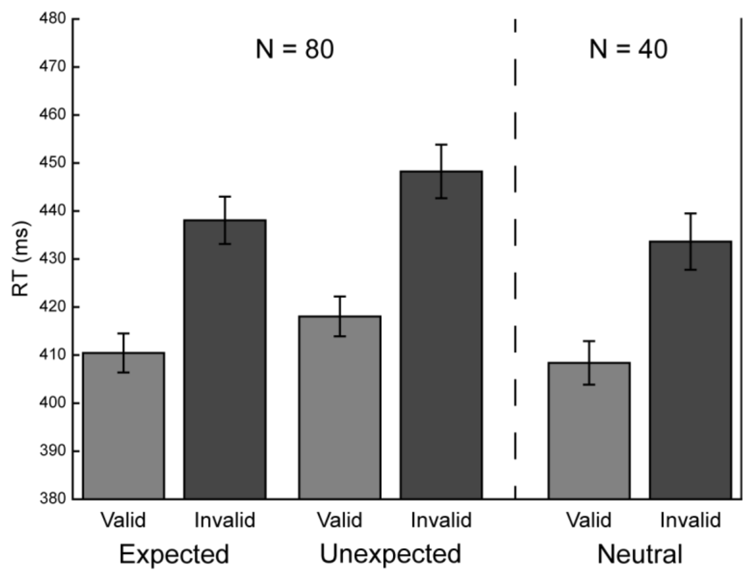

Figure 2. Mean Reaction Time (RT) as a function of cue Validity for the experimental groups (left of dashed line) and the neutral control group (right). **Left;** We found faster RTs for valid compared to invalid (p < .001) cues, and faster RTs for expected compared to unexpected (p < .001) cues. These effects did not interact (p = .356). **Right;** In the control group we found faster RTs for valid compared to invalid cues (p < .001). Error bars reflect 1 standard error of the mean (SEM).

3.1.2 Control group and post-hoc analyses. In the control group, we observed significantly faster RTs for valid (*M* = 408, *SD* = 28.63) compared to invalid (*M* = 433, *SD* = 37.04) cues (t (39) = 10.536, $p < .001$, Cohen's *d* = 6.23; Fig. 2, right).

To confirm the absence of the Validity x Expectation interaction in the main analysis, two additional analyses were done. Firstly, we performed a 2x2 Bayesian RM ANOVA using the exact same factors and levels as the original analysis. By interpreting the associated Bayes Factor (BF) we found substantial evidence against a model that assumes an interaction ($BF_{01} \approx 4$, Jeffreys, 1961; Wetzels & Wagenmakers, 2012). In other words, the additive model assuming no interaction is approximately four times more likely than the interaction model, given the data. Secondly, we performed post-hoc independent samples t-tests to compare the mean Validity effect (invalid – valid) of the expected and unexpected conditions separately to the mean Validity effect of the neutral control condition. The mean Validity effect is effectively a way of quantifying the amount of bottom-up attentional capture. The comparisons showed no significant difference of the Validity effect between neutral and expected ($p = .419$) or neutral and unexpected conditions ($p = .198$). In short, the evidence against an interaction model and the absence of a significant difference of the Validity effect between the experimental and neutral groups indicate that the Validity and Expectation effects do not interact.

3.2 Error rates

A 2x2 RM ANOVA with independent factors Expectation (expected vs. unexpected) and Validity (valid vs. invalid) was done to test for differences in error rates. We observed only a main effect of Expectation (F (1, 79) = 6.293, $p = .014$, $\eta^2_p = .074$), showing that participants made significantly more errors for unexpected (*M* = 2.4, *SD* = 0.22) than for expected cues

(*M* = 2.1, *SD* = 2.08). Error rates did not differ significantly between the neutral and expected ($p = .993$) or neutral and unexpected ($p = .121$) conditions.

4. Discussion experiment I: Spatial Prediction

In this experiment we investigated the influence of spatial expectation of cue location on the amount of bottom-up attentional capture. We replicated the traditional bottom-up attentional capture effect, showing significantly faster RTs for valid compared to invalid cues (Posner, 1980). Additionally, we observed significantly slower RTs for unexpected than for expected cues. This indicates that participants learned the cue probabilities and that the manipulation of cue expectation was successful. We propose that the effect itself can be explained as a startle response: While distractive stimuli enhance RTs if presented simultaneously with a target (Valls-Solé, Kumru, & Kofler, 2008), distractions presented prior to a target slow RTs (Corneil & Munoz, 1996; Theeuwes, 1992). The latter effect has been suggested to be due to noisier perceptual decisions (Garrido, Dolan, & Sahani, 2011). If indeed unexpected onsets cause noisier perceptual decision making, we might also expect a higher proportion of errors for those stimuli. This is indeed what we observe.

Finally, and importantly, no interaction was found between the Validity and Expectation effects, indicating that the amount of attentional capture did not change as a function of cue expectation. This is intriguing, especially since the main effect of Expectation indicates that the manipulation of expectation did have an effect on behavior. Two explanations can be considered. First, it is possible that indeed expected distracting stimuli are not more easily ignored than neutral or unexpected ones. However, this directly contradicts the theories of how predictive mechanisms work. According to those, expected information should be suppressed and evoke smaller sensory responses (den Ouden, Kok, & de Lange, 2012; Kok, Jehee, & de Lange, 2012). It might be that suppression of sensory responses does not necessarily correspond to suppression of behavior, but this is not likely as behavior and neural activity (especially in low-level sensory areas) have often been shown to correlate (Ruff & Cohen, 2013; Nienborg, Cohen, & Cumming, 2012).

Secondly, an alternative explanation of this finding is that expectation does influence attentional capture, but that the current paradigm did not allow for it show. In our paradigm, the cue position, regardless of expectation, was often congruent with the subsequent location of the target (i.e., in valid trials). Therefore, it may not be beneficial for task performance to completely ignore the expected cue location, as this would mean that targets would also potentially be missed. It is thus possible that because of the necessity to always pay attention to both locations attentional capture could not be modulated by expectation. It has indeed been implied that successful suppression of a distractor only occurs when its location is 100% predictable (Noonan, et al., 2016) and furthermore has no potential spatial overlap with the target (Grubb, White, & Heeger, 2015). In that case we would have to conclude that the current paradigm is not appropriate for testing spatial expectation and that a different approach where the dimension of the expectation does not interfere with the task might be more successful. An additional factor that might render the expectation ineffective is that there is a tight temporal relationship between the cue and the target: The target always followed the cue shortly in time. This might make it more difficult to suppress the expected cue, as suppression of information on the predicted location would have to be very specific and brief in time to be successful.

In conclusion, while the results indicate that there is no interaction between expectation and bottom-up attention, it might be that the nature of the paradigm was not fit to investigate spatial expectation. Further, the tight temporal relationship between the cue and the target might be detrimental for effective suppression of expected cues. Therefore, we suspect that we might gain a better insight into the relationship between expectation and bottom-up attention if we correct for these factors.

Experiment II: Temporal Prediction

To further investigate the relationship between expectation and attentional capture, we conduct a second experiment which focuses on temporal expectation. Although we observed no interaction between expectation and bottom-up attention in the first experiment, the nature of the paradigm prevented us from drawing firm conclusions on the exact relationship between predictive and attentional mechanisms. Hence, in this second experiment we make sure that expectation is manipulated such that the specific type of expectation induced by the paradigm (i.e., the temporal predictability of the cue) does not interfere with the task. We decouple both the spatial and the temporal relationship between the cue and the target by making the cue location neutral, and implementing a temporal independency between the cue and target onsets.

2. Method

2.1 Participants

A total of 31 healthy Dutch speaking participants (25 females; age $M \pm SD$, 21.94 ± 2.72) joined the experiment. All participants had normal or corrected-to-normal vision and gave written informed consent prior to experiment initiation where the same ethical approval applied as previously. All participants were unaware of the purpose of the experiment. Data from four participants were excluded because the participants did not properly perform the task (i.e., they used incorrect response buttons), leaving 27 participants for data analysis.

2.2 Material

We used the same experimental setup as in experiment 1, with the exception of the use of the computer keyboard for responses instead of using separate button boxes. For this they used the index fingers of both their hands. Also, the experiment was run on Windows 7 instead of Linux.

2.3 Design

2.3.1 Paradigm & stimuli. We used the same stimulus configurations as in experiment 1. However, the design of the paradigm was changed in order to manipulate the predictability of the cue in the temporal domain. Instead of a traditional trial-based paradigm, for this experiment people performed the task in continuous blocks of 5 minutes. Cues and targets were presented continuously throughout the block, and remained on screen for 50 (cue) and 200 (target) milliseconds. The regularity of cue onset timing was manipulated to induce temporal expectations of the cue (see section 2.3.2 of this experiment). All cue locations and target locations and directions were equally likely. Like in experiment 1, participants' task was to indicate the direction of the arrow with the index fingers of both hands using the 'z' key for a "left" response and the '/' (slash) key for a "right" response. RTs and ERs were measured.

2.3.2 Factors & manipulations. The experiment consisted of a 2x2 within-subjects design, with independent factors Validity (valid vs. invalid cue) and Predictability (predictable vs. unpredictable cue) and dependent factors RT and ER. Validity was defined in terms of the spatial relationship between the target and the last cue prior to target onset. Concretely, the cue location was either the same as (valid trial) or opposite from (invalid trial) the target location.

Temporal predictability of the cue was manipulated such that cues appeared either with a frequency of 1 Hz (predictable) or quasi-randomly every 0.5 – 1.5 seconds (unpredictable), with a temporal resolution of 1 frame (16.7 ms). In the unpredictable condition the probability distribution of the time intervals between two cues was uniform, rendering all cue-to-cue intervals equally likely. Targets were presented quasi-randomly every 1 – 2.5 seconds. The target-to-target intervals were also uniformly distributed and had the same temporal resolution as the cue-to-cue intervals. Target onset timing was determined independently from cue presentation, ensuring that the participants' expectation was only on the cues' timing and not also the targets'. Predictable and unpredictable conditions were blocked and the block order was counterbalanced across participants.

2.4 Procedure

Via the instructions participants were made aware that the cue could (in blocks) either appear exactly every second or randomly, but also that the cue did not in any way have a predictive relationship with the target location or direction and that they should try to ignore it. Halfway through the instruction there was a short practice block of 25 s in which only the sequence of targets was presented, to make the participants familiar with the stimuli and task. Participants were instructed to keep their eyes focused on the fixation cross throughout the stimulus sequence.

After the instructions participants started with the main experiment. The experiment consisted of 8 blocks with the Predictability conditions alternating every 2 blocks. The experiment was counterbalanced such that half of the participants received the conditions in an AABBAABB fashion and the other half the other way round. After each block-transition participants were informed about the nature of the timing of the cue for the next two blocks. Each block was followed by a short break of 20 seconds. In total the main task took approximately 45 minutes. The participants received no feedback on their performance.

2.5 Data analysis

2.5.1 Reaction Times

Preprocessing. Because the experiment consisted of a continuous stream of stimuli, trials were defined post acquisition. For every presented target we determined the closest cue prior to target onset, and based on that defined the associated SOA and validity conditions. If there was no target response the trial was labeled as a miss. RTs below 200 ms or measurements deviating more than 3 standard deviations from the mean were removed from the dataset. Trials in which there was no cue preceding the target (i.e., the first target of a block) and trials during the first 60 seconds of each transition block (i.e., blocks where the cue predictability had changed) were removed from the dataset, to account for any carry-over effects between conditions. This amounted to an average removal of 162 trials per participant (11.2% of data set). For the RT analysis only trials with correct responses were included.

SOA's. As a result of the irregular onset timing of the target, and in the unpredictable condition also the cue, there was a wide range of possible SOA's. The individual SOA bins were defined in steps of frames, which is a result of the cue and target onsets also being determined in terms of frames. The bins ranged from 17 ms to 483 ms, with a final bin that contained all SOA's of 500 ms and longer, resulting in a total of 30 bins. Initially we planned to only look at the SOA bins of interest for bottom-up attention (50 – 200 ms; Dukewich & Klein, 2015; Klein, 2000), however there appeared to be large fluctuations in the effects per SOA. Therefore we decided to look at the individual SOA's rather than average over them.

After splitting the data into the SOA bins, there was a very low amount of trials per cell, rendering any effects potentially unreliable. To compensate for this we applied a moving average on the data to average out random noise. Specifically this means that for every bin in the analysis (and the plot), we used the average of that bin combined with the average of the preceding and following bin. This way we managed to gather a respectable amount of trials per bin which allows for more reliable statistical testing.

Statistical testing. A disadvantage of the moving average is that the same data is used multiple times over all bins, making a standard Validity x Predictability x SOA interaction analysis no longer valid. For that reason, we did a 2x2 (Validity x Predictability) RM ANOVA for each SOA bin separately. Consequently we cannot directly compare effects between the different SOA bins.

2.5.2 Error Rates. Error rates were calculated in the same way as in experiment 1, and the same exclusion threshold of an ER of 20% was maintained. Apart from the four participants who were excluded due to incorrect performance of the task, no participants reached that threshold. For the statistical analysis we performed a 2x2 (Validity x Predictability) RM ANOVA collapsed over all SOA's, as it was not within the interest of the main research question to apply a moving average on the ERs as well.

3. Results

3.1 Reaction Times

In several SOA bins significant effects were found. Firstly, we observed significant main effects of Validity in the SOA bins of 117 ms (F (1,26) = 8.2615, p = .008), 133 ms (F (1,26) = 6.5131, p = 0.17) and 150 ms (F (1,26) = 14.1258, p < .001) . More specifically, in these bins RTs were significantly faster for valid compared to invalid cues. Furthermore, in the later SOA bins (417 ms and higher) we found a reversed Validity effect where RTs were faster for invalid than for valid cues (for detailed statistics per bin see Table 1). Overall, no main effects of Predictability were found. Secondly, we found a significant interaction between Validity and Predictability in the 133 ms bin (F (1,26) = 4.9325, p = .035). At this SOA, there was a significantly larger Validity effect for unpredictable cues compared to predictable cues (Fig. 3).

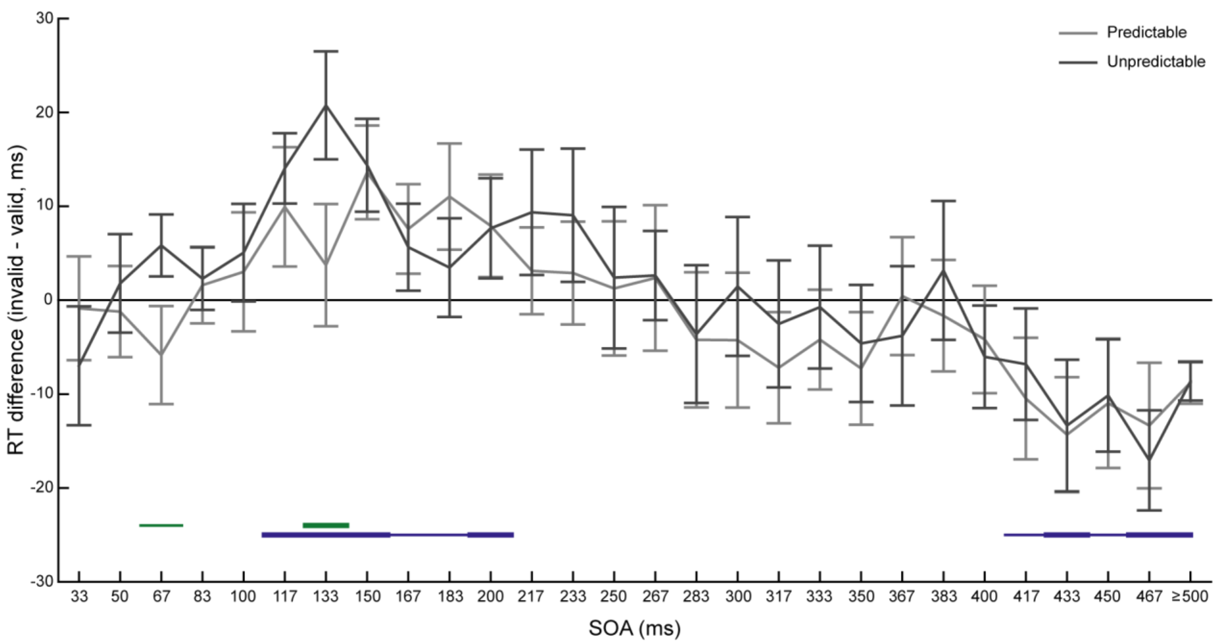

Figure 3. Moving average data. For both predictable and unpredictable conditions the mean Validity effect (invalid – valid) in ms is plotted per SOA bin. Values above zero indicate faster RTs for valid compared to invalid cues, and values below zero indicate faster RTs for invalid compared to valid cues. Lines in the bottom of the figure indicate significance levels of Validity (blue) and Validity x Predictability (green) effects (thick line p < .05; thin line p < .1). Error bars reflect 1 SEM.

Table 1

Repeated Measures ANOVA output per SOA bin[a]

		F (df = 1, 26)			p (η^2_p)	
SOA bin	Val.	Pred.	Val. x Pred.	Val.	Pred.	Val. x Pred.
67 ms	0.0000	0.0004	3.3276	.998	.984	.080
117 ms	8.2615	0.0002	0.4280	*.008* (.24)	.989	.519
133 ms	6.5131	0.1242	4.9325	*.017* (.20)	.727	*.035* (.16)
150 ms	14.1258	0.6060	0.0132	< *.001* (.35)	.443	.909
167 ms	3.6717	0.6428	0.0925	.066	.430	.763
183 ms	3.4576	1.6543	0.9783	.074	.210	.332
200 ms	4.6126	0.0415	0.0011	.041	.840	.973
417 ms	3.3460	0.0056	0.2068	.079	.941	.653
433 ms	6.7053	0.7124	0.0160	*.016* (.21)	.406	.900
450 ms	3.7858	0.1819	0.0139	.063	.673	.907
467 ms	9.3116	0.0044	0.2941	*.005* (.26)	.948	.592
≥ 500 ms	30.7063	0.0088	0.0021	< *.001* (.54)	.926	.964

Note. [a]Only the SOA bins where the *p*-value of at least one effect was below .1 are displayed (df: degrees of freedom; Val.: Validity; Pred.: Predictability; Val. x Pred.: two-way interaction).

3.2 Error Rates

A 2x2 RM ANOVA with within-subjects factors Validity (valid vs. invalid) and Predictability (predictable vs. unpredictable) showed a significant main effect of Validity (F (1, 26) = 11.6252, p = .0021, η^2_p = .309). Participants made more errors for invalid compared to valid cues.

4. Discussion experiment II: Temporal Prediction

In the second experiment we investigated the effect of temporal predictability of a distracting cue on the amount of bottom-up attentional capture. After application of a moving average on the different SOA bins, we found a range of SOA's (around 117 – 150 ms) where there was a significant effect of Validity, indicating an area of effective attentional capture. The finding is in accordance with a vast amount of literature showing that bottom-up attention is dependent on the exact timing relationship between the cue and the target, demonstrating that bottom-up attention is facilitated only at certain SOA intervals (e.g., Dukewich & Klein, 2015; Klein, 2000). While it is common that bottom-up attention also occurs at SOA's below 100 ms (Posner & Cohen, 1984), we suspect that in our experiment this was hindered by a masking effect of the cue on the target, consequently slowing RTs in valid trials (Kouider & Dehaene, 2007). The increased amount of errors for invalid compared to valid cues further indicates that the reaction time effect is not the result of a speed-accuracy trade-off. Furthermore, a significant reversed Validity effect was found in the late SOA bins (approx. 417 ms and higher), showing faster RTs for invalid than valid trials. This effect has been found frequently in other bottom-up attention studies and is often referred to as inhibition of return (IOR). According to the IOR principle, locations that have recently been the subject of attention are subsequently being prevented from being attended to again (Itti & Koch, 2001; Klein, 2000). Consequently, attentional capture will impair performance on valid trials while simultaneously boosting performance on invalid trials. In short, our data are in line with already existing literature on bottom-up attention, providing further confirmation of basic bottom-up attention principles.

We also manipulated the temporal predictability of the cue onset. While we did not find a main effect of Predictability, we did observe a significant interaction between Predictability and Validity. More specifically, around an SOA of 133 ms there was a significantly larger Validity effect for unpredictable than for predictable cues. Though the effect is only local in time, this finding may confirm our hypothesis that unpredictable cues enlarge the bottom-up attentional capture relative to predictable cues. We suspect that because of the unexpected nature of the cue onset, a larger PE was generated resulting in increased allocation of attention towards the cue, subsequently enlarging the difference between valid and invalid trials. Predictable cues in contrast would have only generated small PEs, leading to less attentional allocation towards the cue and subsequently a smaller cue validity effect. However, our behavioral results cannot provide concrete evidence for this proposition. Further studies using neuroimaging methods that are capable of directly measuring PEs (e.g., fMRI and magnetoencephalography; MEG) are needed to confirm our hypothesis (den Ouden, Kok, & de Lange, 2012).

In conclusion, we found evidence that, if the cue-target SOA falls within a certain range, expectation modulates bottom-up attention such that attentional capture is increased for unexpected compared to expected stimuli. Admittedly, the modulation of the Validity effect is only present in one SOA bin. Implications of our findings will further be discussed and we will subsequently propose a potential mechanism underlying the interaction between expectation and bottom-up attention.

General discussion

In this study we conducted two experiments to investigate if bottom-up attention is modulated by expectation. In the first experiment we focused on spatial expectation and replicated the classical bottom-up attention effect (Posner, 1980; Posner, Snyder, & Davidson, 1984). However, we did not find evidence that supports a predictive coding account of (bottom-up) attention, as attentional capture was the same for expected and unexpected cues. In the second experiment we investigated the influence of temporal predictability on attentional capture and found a range of cue-target SOA's with effective bottom-up attention, as is in accordance with the literature (Dukewich & Klein, 2015; Klein, 2000). Additionally, we observed a temporally localized interaction between Validity and Predictability. Around an SOA of 133 ms, unpredictable cues resulted in significantly more attentional capture than predictable cues. This finding elaborates further on the recent work done by Folk & Remington (2015), and provides, to our knowledge, a first direct indication of an interactive relationship between bottom-up attention and expectation. Further, in a broader perspective this finding provides support for a predictive coding account of attention (Feldman & Friston, 2010; Friston, 2009; Hohwy, 2012).

However, not all is yet clear. For example, while we found a modulation of bottom-up attention by expectation in experiment 2, no such modulation was found in the first experiment. This suggests a dissociation between different types of expectation, where perhaps spatial expectations are less powerful than temporal expectations. However, there seems to be no concrete evidence that both types of expectation differentially influence attention and that they rather have additive effects on both behavior and brain activity (Doherty, Rao, Mesulam, & Nobre, 2005). It is difficult to claim otherwise on the basis of the present study, as our paradigm has its limitations with regards to interpreting the influence of the spatial dimension of expectation (as reviewed in the Discussion of experiment 1). It would therefore be interesting to investigate how, if at all, bottom-up attention is altered if spatial and temporal expectations are jointly manipulated.

Furthermore, it is unclear exactly what the mechanistic relationship is between expectation and bottom-up attention. While it seems evident that bottom-up attention is vulnerable to influences of prior expectations, the effect that we found only occurred in a very specific cue-target SOA range (around 133 ms). This raises the question why only in this particular time period expectation is able to modulate attentional capture. While this has to be investigated further to able to develop a complete framework, we will briefly speculate on possible underlying mechanisms that might cause the pattern we observed in the data of our second experiment.

A first important consideration is that predictive mechanisms are top-down processes, meaning that prior information in higher-order brain areas is fed back to low-level areas, which is a process that takes time (De-Wit, Machilsen, & Putzeys, 2010; den Ouden, Daunizeau, Roiser, Friston, & Stephan, 2010; Friston, 2005; Rao & Ballard, 1999). Thus it might be that prior expectations are not able to influence bottom-up attentional processing for short SOA's (e.g., < 100 ms) because there is not sufficient time to compare the prediction (being fed back from the top of the processing hierarchy) with the sensory evidence before the target appears. Reward PE-neurons have for example indeed been shown to fire approximately 100 ms after reward presentation (Tobler, Fiorillo, & Schultz, 2005). If a resembling mechanism underlies perceptual PEs, this would imply that at short SOA's bottom-up attention is in principle stimulus driven and not modulated by expectation. Then as the SOA gets longer (e.g., towards 133 ms), predictions have more time to be fed-back to sensory areas, compute the prediction error, and from there on modulate attention.

Another open question is why the modulatory influence of expectation on attentional capture diminishes rather quickly and does not last for SOA's of 150 to 200 milliseconds. One possible explanation is that modulation of bottom-up attention relies on the presence of a bottom-up attention effect. In literature on the relationship between SOA's and bottom-up attention there is consensus that attentional capture tends to disappear at longer SOA's (Dukewich & Klein, 2015; Klein, 2000). It might thus be that at longer SOA's predictive mechanisms are ineffective because there is no bottom-up attention present to be modulated. Another factor that might play a role here is that with increased time after cue onset, the attentional system may be able to endogenously reorient attention quickly towards the target and is hence less dependent on bottom-up processing, making the prediction irrelevant. Still, you would presumably expect a gradual decrease of the influence of expectation rather than a sudden cut-off.

A more speculative consideration involves that to act on a dynamic environment, it is crucial for a predictive brain to always have the most up-to-date ‘template’ of the world available (Summerfield, et al., 2006). However, as time progresses already fed-back information becomes less relevant for the current state of the environment as it was based on relatively old sensory information. To compensate for this, it might be the case that predictive signals lose weight over time, causing them to exert less influence on bottom-up attentional processing. Another way this might be implemented is that new predictions overwrite the old predictions, rendering the latter ineffective. In this context, it would be interesting for further research to investigate at what rate old predictions get overwritten by novel information.

All in all, we propose a modulatory account of the interaction between bottom-up attention and expectation. Specifically, we pose that attentional capture is in principle stimulus driven, but becomes vulnerable to predictive influences as the temporal conditions allow for effective comparison of prior expectations with incoming sensory information. More research is needed to elucidate exactly what mechanism underlies the susceptibility of bottom-up attention to expectation. We suggest using temporally high-resolution neuroimaging methods such as MEG to find out what happens in the brain when bottom-up attention and expectation interact, in particular in the context of their temporal dependencies.

References

Baldeweg, T. (2006). Repetition effects to sounds: evidence for predictive coding in the auditory system. *Trends in Cognitive Sciences*, *10*(3), 93-94.

Baluch, F., & Itti, L. (2011). Mechanisms of top-down attention. *Trends in Neuroscience*, *34*(4), 210-224.

Belopolsky, A.V., Schreij D., & Theeuwes, J. (2010). What is top-down about contingent capture? *Attention, Perception & Psychophysics*, *72*(2), 326-341.

Boynton, G.M. (2009). A framework for describing the effects of attention on visual responses. *Vision Research*, *49*, 1129-1143.

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

Connor, C.E., Egeth, H.E., & Yantis, S. (2004). Visual Attention: Bottom-Up Versus Top-Down. *Current Biology*, *14,* R850-R852.

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

Corneil, B.D., & Munoz, D.P. (1996). The Influence of Auditory and Visual Distractors on Human Orienting Gaze Shifts. *The Journal of Neuroscience*, *16*(24), 8193-8207.

De-Wit, L., Machilsen, B., & Putzeys, T. (2010). Predictive Coding and the Neural Response to Predictable Stimuli. *The Journal of Neuroscience*, *30*(26), 8702-8703.

Den Ouden, H.E.M., Daunizeau, J., Roiser, J., Friston, K.J., & Stephan, K.E. (2010). Striatal prediction error modulates cortical coupling. *The Journal of Neuroscience*, *30*, 3210-3219.

Den Ouden, H.E.M., Kok, P., & de Lange, F.P. (2012). How prediction errors shape perception, attention, and motivation. *Frontiers in Psychology*, *3*:548.

Den Ouden, H.E.M., Friston, K.J., Daw, N.D., Mcinosh, A.R., & Stephan, K.E. (2009). A dual role for prediction error in associative learning. *Cerebral Cortex*, *19*, 1175-1185.

Doherty, J.R., Rao, A., Mesulam, M.M., & Nobre, A.C. (2005). Synergestic Effect of Combined Temporal and Spatial Expectations on Visual Attention. *The Journal of Neuroscience*, *25*(36), 8259-8266.

Dukewich, K.R., & Klein, M.R. (2015). Inhibition of return: a phenomenon in search of a definition and a theoretical framework. *Attention, Perception, & Psychophysics*, *77*(5), 1647-1658.

Folk, C.L., & Remington, R.W. (2015). Unexpected Abrupt Onsets Can Override a Top-Down Set for Color. *Journal of Experimental Psychology*, *41*(4), 1153-1165.

Folk, C.L., Remington, R.W., & Johnston, J.C. (1992). Involuntary Covert Orienting Is Contingent on Attentional Control Settings. *Journal of Experimental Psychology*, *18*(4), 1030-1044.

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society*, *360*, 815-836.

Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends of Cognitive Science, 13*, 293-301.

Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1521), 1211-1221.

Garrido, M., Dolan, R., & Sahani, M. (2011). Surprise leads to noisier perceptual decisions. *I-Perception*, *2*(2), 112-120.

Gibson, B.S., & Keylsey, E.M. (1998). Stimulus-driven attentional capture is contingent on attentional set for displaywide visual features. *Journal of Experimental Psychology: Human Perception & Performance, 24*, 699-706.

Grubb, M.A., White, A.L., & Heeger, D.J. (2015). Interactions between voluntary and involuntary attention modulate the quality and temporal dynamics of visual processing. *Psychonomic Bulletin Review*, *22*, 437-444.

Helmholtz, H. von. (1867/1925). *Treatise on physiological optics* (Translated from the Third German edition). Dover, NY: Constable and Company.

Hohwy, J. (2012). Attention and conscious perception in the hypothesis testing brain. *Frontiers in Psychology*, *3*:96.

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

Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience*, *2*(3), 194-203.

JASP Team (2016). JASP (Version 0.7.5.5) [Computer software].

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

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

Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? *Perception*, *36*, Abstract Supplement.

Kok, P., Jehee, J.F., & de Lange, F.P. (2012). Less is more: expectation sharpens representations in the primary visual cortex. *Neuron*, *75*(2), 265-270.

Kok, P., Rahnev, D., Jehee, J.F.M., Lau, C.H., & de Lange, F.P. (2012). Attention Reverses the Effect of Prediction in Silencing Sensory Signals. *Cerebral Cortex*, *22*(9), 2197-2206.

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

Kumaran, D., & Maguire, E.A. (2009). Novelty signals: a window into hippocampal information processing. *Trends in Cognitive Science*, *13*, 47-54.

Li, Z. (2002). A saliency map in primary visual cortex. *Trends in Cognitive Sciences, 6*(1), 9-16.

Love, J., Selker, R., Marsman, M., Jamil, T., Dropmann, D., Verhagen, A. J., Ly, A., Gronau, Q. F., … Wagenmakers, E.J. (2015). JASP (Version 0.7.5) [Computer software].

MATLAB Release 2011b, 2014b, The MathWorks, Inc., Natick, Massachusetts, United States.

Morey, R. D., & Rouder, J. N. (2015). BayesFactor (Version 0.9.11-3) [Computer software].

Nienborg, H., Cohen, M.R., & Cumming, B.G. (2012). Decision-related activity in sensory neurons: correlations among neurons and with behavior. *Annual Review Neuroscience*, *35*, 463-83.

Noonan, M.P., Adamian, N., Pike, A., Printzlau, F., Crittenden, B.M., & Stokes, M.G. (2016). Distinct Mechanisms for Distractor Suppression and Target Facilitation. *The Journal of Neuroscience*, *36*(6), 1797-1807.

O'Doherty, J.P., Buchanan, T.W., Seymour, B., & Dolan, R.J. (2006). Predictive neural coding of reward preference involves dissociable responses in human ventral midbrain and striatum. *Neuron, 49*(1), 157-166.

Pinto, Y., van der Leij, A.R., Sligte, I.G., Lamme, V.A.F., & Scholte, H.S. (2013). Bottom-up and top-down attention are independent. *Journal of Vision*, *13*(3):16.

Polich, J. (2007). Updating P300: An integrative theory of P3a and P3b. *Clinical Neurophysiology*, *118*, 2128-2148.

Posner, M.I. (1980). Orienting of attention. *The Quarterly Journal of Experimental Psychology*, *32*(1), 3-25.

Posner, M.I., & Cohen, Y. (1984). *Components of visual orienting. In Attention and Performance Vol. X.* (Bouma, H., and Bouwhuis, D., eds), pp. 531-556, Erlbaum.

Posner, M.I., Snyder, C.R.R., & Davidson, B.J. (1984). Attention and the detection of signals. *Journal of Experimental Psychology: General*, *109*, 160-174.

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

Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, *56*, 356-374.

Ruff, D.A., & Cohen, M.R. (2013). Pursuing the Link between Neurons and Behavior. *Neuron*, *79*, 6-9.

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.

Summerfield, C., & Koechlin, E. (2008). A neural representation of prior information during perceptual inference. *Neuron*, *59*, 336-347.

Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics*, *51*, 599-606.

Theeuwes, J. (2004). Top-down search strategies cannot override attentional capture. *Psychonomic Bulletin & Review*, *11*, 65-70.

Theeuwes, J., & Godijn, R. (2002). Irrelevant singletons capture attention: Evidence from inhibition of return. *Perception & Psychophysics, 64*(5), 764-770.

Theeuwes, J., & van der Burg, E. (2007). The role of spatial and nonspatial information in visual selection. *Journal of Experimental Psychology: Human Perception & Performance*, *33,* 1335-1351.

Tobler, P.N., Fiorillo, C.D., & Schultz, W. (2005). Adaptive Coding of Reward Value by Dopamine Neurons. *Science*, *307*(5715), 1642-1645.

Todorovic, A., van Ede, F., Maris, E., & de Lange, F.P. (2011). Prior expectation mediates neural adaptation to repeated sounds in the auditory cortex: an MEG study. *The Journal of Neuroscience, 31*, 9118-9123.

Valls-Solé, J., Kumru, H., & Kofler, M. (2008). Interaction between startle and voluntary reactions in humans. *Experimental Brain Research*, *187*, 497-507.

Wacongne, C., Changeux, J.P., & Dehaene, S. (2012). A Neuronal Model of Predictive Coding Accounting for the Mismatch Negativity. *The Journal of Neuroscience*, *32*(11), 3665-3678.

Wetzels, R., & Wagenmakers, E.J. (2012). A default Bayesian hypothesis test for correlations and partial correlations. *Psychonomic Bulletin & Review*, *19*(6), 1057-1064.