Pinging the brain to reveal the hidden attentional priority map

Dock Duncan[1,2] , Dirk van Moorselaar[1,2] & Jan Theeuwes[1,2,3]

[1]Vrije Universiteit Amsterdam, Amsterdam, the Netherlands

[2] Institute Brain and Behavior Amsterdam, the Netherlands (iBBA)

[3]William James Center for Research, ISPA-Instituto Universitario, Lisbon, Portugal

Author notes

This research was supported by a European Research Council (ERC) advanced grant 833029 – [LEARNATTEND]. Correspondence should be addressed to Docky Duncan, Department of Experimental and Applied Psychology, Vrije Universiteit Amsterdam, Van der Boechorststraat 7, 1081 BT Amsterdam, The Netherlands. Email: D.H.Duncan@vu.nl. All scripts and data will be made publicly available on GitHub and OSF upon publication.

Abstract

Priority maps are winner-take-all neural mechanisms that are thought to guide the allocation of attention. Within these maps, attentional priority is coded as weights on a topographic representation of physical space. Research has shown that these weights are not just influenced by the current stimuli presented but also by selection episodes that occurred in the past – called selection history. The current study demonstrates, for the first time, that these latent, implicitly-learned spatial biases can be visualized by 'pinging' the brain (a technique previously used to study activity-silent working memory). By briefly flashing high-contrast, task-irrelevant visual stimuli in the inter-stimulus interval between searches, participants' current history-trained latent attentional priority could be decoded via scalp encephalography. Our findings not only offer a novel method of visualizing this latent attentional priority map, but also shed light on the underlying mechanisms allowing our past experiences to influence future behaviour.

Introduction

The immense complexity of our visual surroundings presents a fundamental challenge to our finite minds. Fortunately, our brains are able to take advantage of two general principles of the world to greatly simplify this perceptual problem: first, the world is highly repetitive, and therefore predictable; and second, much of what we perceive at any moment can be safely ignored. The brain's propensity to automatically learn environmental regularities is often referred to as "Statistical Learning" (Frost et al., 2019). Recently it has been claimed that uncovering the underlying correlational structure of perceptual input may be a key method by which the brain simplifies the perceptual problem space by sharpening percepts around robust predictions, thereby reducing computational costs (Friston, 2010; Y. Huang & Rao, 2011; Kok et al., 2012; Srinivasan et al., 1982). Selective attention on the other hand, is the process whereby relevant information is prioritized while task irrelevant information is suppressed. Attention plays a crucial role in structuring perception as the majority of what we initially perceive is filtered out by this system, thereby greatly reducing the redundancy of perceptual experience (Broadbent, 1958; Rensink et al., 1997; Treisman & Gelade, 1980).

Given the importance of these two cognitive mechanisms in shaping our perceptual experience, it may come as no surprise that these mechanisms are deeply interconnected. For instance; while driving, it is important to direct attention to some stimuli (e.g., crossing pedestrians) but little-to-no, attention to others (e.g. irrelevant billboards). When driving in unfamiliar roads, sorting relevant from irrelevant stimuli is an attentionally demanding task; however, when driving on familiar roads the experience is quite different. Through our past experiences the visual system can be tuned in space and time to expect relevant stimuli to appear in certain places (e.g. familiar pedestrian crossings) as well as where distractors are likely to be (e.g. familiar billboards), thereby combining the cognitive tools of attention and statistical learning into an integrated system which sharpens perception to maximize performance.

The influence of past experience on current behavior in attention is known as selection history - a category which encompasses statistical learning processes as well as other history based effects such as value-driven attentional capture or intertrial priming (Anderson et al., 2021; Awh et al., 2012; Theeuwes, 2018; Theeuwes et al., 2022). Together with top-down (goal driven) and bottom-up (saliency driven) attentional mechanisms, selection history is thought to converge in what is known as the attentional priority map – a real time representation of the behavioral relevance and saliency of the stimulus present in the visual field (Fecteau & Munoz, 2006; Ferrante et al., 2018; Theeuwes, 2018; Theeuwes et al., 2022). While these maps can exist in any feature space, they are often studied in the context of spatial features, where attentional priority is coded as weights on a topographic representation of physical space, and attentional selection is ultimately awarded in a winner-take-all fashion to the region with the highest activity (Itti & Koch, 2001; Zelinsky & Bisley, 2015). Selection history can then be thought of as a layer in this system, continuously up-regulating weights to locations which contained relevant information in the past, and down-regulating locations that frequently contain distracting information (Theeuwes et al., 2022; Theeuwes & Failing, 2020).

While the concept of a priority map that drives attention selection has been a prominent notion in many theories on attentional selection (for a selection, see Desimone & Duncan, 1995; Fecteau & Munoz, 2006; Itti & Koch, 2001; Serences & Yantis, 2006; Theeuwes, 2010; Treisman & Gelade, 1980; Wolfe, 1994), neuroimaging research on priority maps have notably focused almost exclusively on top-down or bottom-up influences (Ptak, 2012; Silver et al., 2005; Sprague et al., 2018; Sprague & Serences, 2013). Preliminary studies investigating the

influence of selection history on the attentional priority map have noted that, unlike top-down attention, there is little evidence that the neural processes underlying the effects of selection history can be described in terms of sustained neural processes (Adam & Serences, 2021; Barbosa et al., 2020; Moorselaar & Slagter, 2019; van Moorselaar et al., 2021; see Wang et al., 2019 for contrary evidence). Instead, it has been suggested that selection history effects exist on the network level, where processes of synaptic plasticity strengthen or weaken neural connections in an emergent manner (Adam & Serences, 2021; Ferrante et al., 2021; van Moorselaar et al., 2020; see Zelinsky & Bisley, 2015 for discussion on possible neural mecahnisms). Due to this latent characteristic, history-based influences on attentional selection may fall into the category of 'activity-silent' cognitive mechanisms; named such because these effects are invisible to common neural imaging techniques which measure the downstream traces of action potentials (i.e., ERP's or BOLD-responses). As a result, the study of selection history effects on the priority map has generally been restricted to the study of differences in evoked responses to predictable and unpredictable stimuli as a proxy for latent expectations in the brain (van Moorselaar et al., 2020; Wang et al., 2019; Zhang et al., 2022), and their study in their pure form has remained out of reach.

Recently, in the field of working memory a novel approach has been proposed to visualize activity-silent neural structures; synaptic theories of working memory propose that working memory may be partially (or fully) mediated by modifying synaptic weights such that remembered information is primed for reactivation in the near future (Barak & Tsodyks, 2014; Mongillo et al., 2008). In this energetically efficient model, memories can be retrieved in subsequent sweeps of network activity in the brain, where the stored memory will be primed to reactivate (D'Esposito & Postle, 2015; Sugase-Miyamoto et al., 2008). Critically, it has recently been demonstrated that these latent memories can also be incidentally mis-activated simply by flooding the perceptual system with sudden input; often in the form of high-contrast visual 'pings' (Wolff et al., 2015, 2017 see Rose et al., 2016 for a demonstration using TMS). This salient input causes neurons to fire in discernable patterns associated with currently retained memory items, possibly due to memory related primed neurons incidentally mis-firing at a higher rate than other neurons (see discussion), thereby leading to systematic neural activity patters which can be used to decode contents of otherwise activity-silent memory (Buonomano & Maass, 2009; Stokes et al., 2013).

While statistically learned attentional priority and synaptic working memory are clearly distinct processes with unique underlying cognitive mechanisms, they arguably share the feature of being latent neural structures mediated by features of network connectivity at the synaptic level (see discussion for further debate). As there is no a priori reason to believe that memory-relevant primed neural structures should be vulnerable to an impulse perturbation ("pinging") while similar latent structures related to the spatial priority map should not be, we postulate that the 'ping' technique may be appropriated to visualize learned attentional priority in a behaviorally independent manner. Such a finding would both inform the neural mechanisms underlying history influences on attentional selection as well as represent the first time the latent attentional priority map was imaged in a neutral way.

In the current study we employed the additional singleton task (Theeuwes, 1992) with imbalanced target distributions, a paradigm that can be used to implicitly train participants to expect relevant information (i.e., targets) to appear in certain regions of space (Ferrante et al., 2018; Gao & Theeuwes, 2020; C. Huang et al., 2022). We then showed that while the ongoing EEG signal did not contain information regarding the current high probability location (indicating this learned bias was in an activity-silent state), we could robustly decode this

high probability target location based on the ping's evoked response. Control analyses showed that this decoding could not be attributed to temporal confounds or eye movements, suggesting instead that the ping-evoked neural responses revealed a latent, implicitly learned spatial bias operating as a silent layer of the attentional priority map.

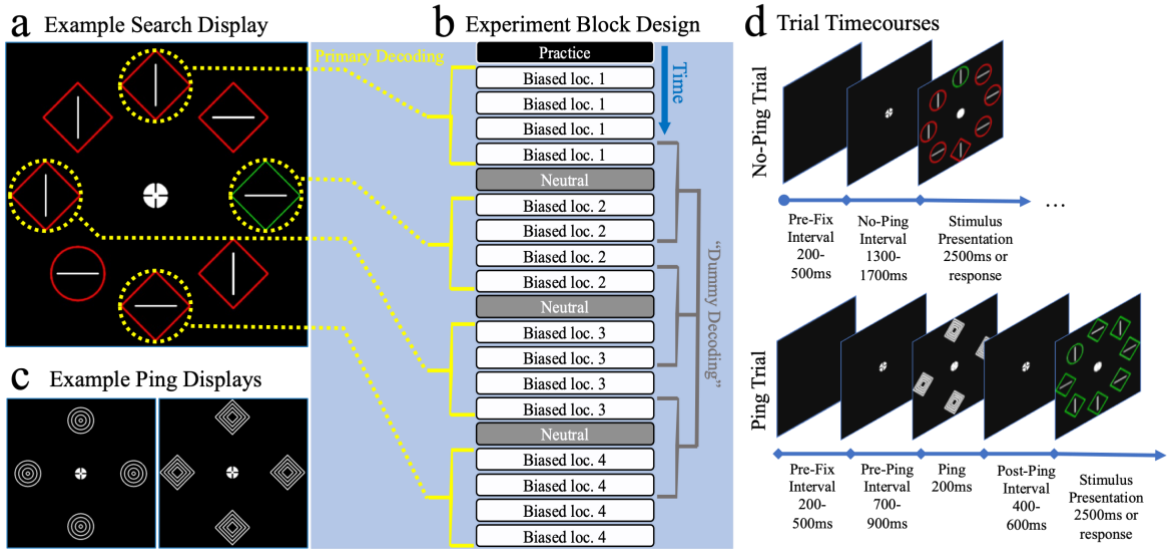

FIGURE 1 – Trial timecourse and experiment design: **A)** Example of a stimulus display. Note that the yellow markings were not present in the actual experiment. Participants' task was to find a uniquely shaped shape singleton and report the line within by pressing the 'z' key if it was horizontally oriented or the '/' key if it was vertical. Colors and shapes swapped identities randomly between trials such that participants did not know what color or shape the target would be on each trial. A distractor was present on 2/3 of trials in the form of a salient color singleton. **B)** Experiment block order and decoding regime. The experiment consisted of biased blocks and neutral blocks. Biased blocks contained one location in space which was significantly more likely to contain a target than any other location (37.5% of trials). Neutral blocks had no high-probability location and instead targets appeared equally frequently in all eight locations. Participants would encounter four biased blocks followed by a neutral block. During the next set of four biased blocks, a different high-probability locations would be used, cycled among four locations (top, bottom, left and right locations in space – circled in yellow in A). Through the course of the experiment each participant encountered each of the four high-probability locations. The order in which participants encountered the high-probability locations was counterbalanced across participants (the order shown in the figure is just one example). The rightmost panel shows the 'dummy decoding' regime – as a control analysis, decoders were passed fake category labels whith represented random stretches of time in the experiment in which no consistent high-probability location was used. These labels consisted of four blocks – beginning with the last block of one high-probability condition and ending with the second block of another high-probability phase. These dummy labels preserved temporal features of the EEG signal for the included trials while removing categorical features of a shared high-probabiity location in space. **C)** Examples of the two types of pings used in this experiment. Both were made up of the same shape stimuli as were used in the task (circles and diamonds) but instead of being filled with a line, the shapes were filled with increasingly smaller shapes. **D)** Timecourses for no-ping(top) and ping trials(bottom). Pre-stimulus periods were time matched for both trial types, ranging from 1500-2200ms. Trials began with a black screen, followed by the presentation of the fixation dot. On half of trials, a salient ping was presented next for 200ms. Trials ended with the stimulus presentation in which participants needed to find a unique shape and report the orientation of the line within. The stimulus was presented until a reponse was provided or 2500ms had passed. If participants did not provide a response, the trial was counted as incorrect. On trials in which no ping was presented, a trigger event was still recorded at a random timepoint in which a ping could occur. This trigger event was later used for baselining in the no-ping decoding analysis.

Results

The experimental design and all analysis methods were preregistered on Open Science Framework. Preregistration can be found at https://osf.io/xdvz4/.

Behavior: Flexible learned prioritization at high probability target locations in a changing environment

Participants (N = 24) performed a variant of the additional singleton task (Theeuwes, 1992, see figure 1A), with sequences of biased blocks where the target appeared with a higher probability at specific locations (i.e., four out of eight possible target locations served as high probability target location throughout the experiment; see Figure

1A) intermixed with neutral blocks without a spatial target location imbalance serving to reset the learned priority landscape (see Figure 1B; see methods for full details). Critically, on half of all trials, participants would encounter task irrelevant high-contrast visual 'pings' in the intertrial period in between visual searches (see figure 1C & D).

As visualized in Figures 2A-C behavioral results indicated that participants were sensitive to the distributional properties of the targets across the various high probability locations throughout the experiment, as indicated by faster responses when targets appeared in high probability target locations relative to all other locations (t (23) = 10.62, $p < 0.001$, $d_Z = 2.17$; Figure 2A). Furthermore, a reliable intertrial target location effect was observed when the target location repeated from one trial to the next (t (23) = 5.956, $p < 0.001$, $d_Z = 1.22$; Figure 2B). Controlling for this intertrial priming effect, the observed speed up of RT's at high probability locations remained highly reliable after excluding all trials where the target repeated from one trial to the next (t (23) = 9.882, $p < 0.001$, $d_Z = 2.02$; 18% of all trials excluded, Figure 2 A & C exclude target repetitions). Furthermore, in line with selective changes in attentional priority as a function of the introduced statistical regularities, distractor interference was reliably more pronounced at high than low probability distractor locations (t (23) = 3.442, $p = 0.002$, $d_Z = 0.7$; Figure 2C).

To examine whether the learning effect differed across experimental phases, a repeated measure analysis of variance (ANOVA) with within subjects' factors Target location (high probability vs. low probability) and experiment phase (high probability location 1-4, see figure 1B) was conducted. This yielded two main effects (all F's > 8.7, all p's < 0.001, all $\eta^2 > 0.033$), but no interaction ($F(3, 69) < 1$; $BF_{10} = 0.073$), indicating that the observed speed up at high probability locations was invariable throughout the experiment (all t's > 2.2, all p's < 0.036, all $d_Z > 0.46$; see Supplementary Figure 1). To further test whether the neutral blocks led to an extinction of previously acquired attentional bias, mean reaction times were compared in the neutral blocks between trials in which the target appeared at the high-probability location of the previous block with all seven other locations. No significant difference was found between these trial types, indicating a successful extinction of bias (t (23) = 1.444 $p = 0.162$, $d_Z = 0.30$; displayed per-block in Supplementary Figure 1).

Out of the 24 participants, 11 indicated that they noticed targets were presented more frequently at one location than any other. Seven out of these 11 participants were also able to correctly identify the HP location of the preceding block. Both behavioral and decoding analyses (see below) did not differ when excluding these 11 participants (t (12) = 6.751, $p < 0.001$, $d_Z = 1.9$; see Supplementary Figure 2). While this does not rule out the possibility that participants had some explicit knowledge regarding the underlying manipulation (Vadillo et al., 2020; Vicente-Conesa et al., 2021), it does appear that both behavioral and decoding results were consistent regardless of reported awareness levels. These results further support recent findings that participants utilize statistically learned attentional strategies even regardless of explicit awareness (Gao & Theeuwes, 2022).

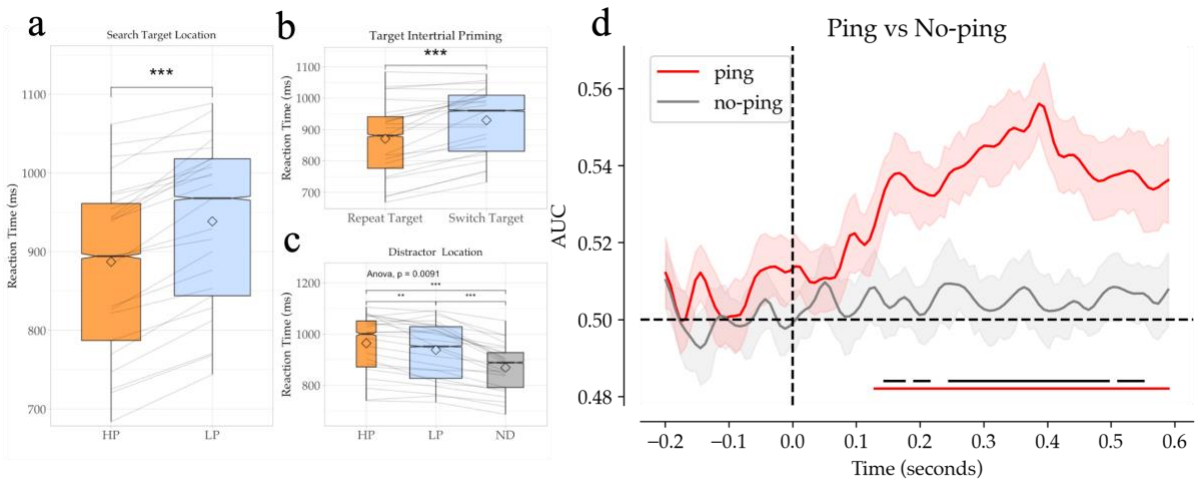

FIGURE ERROR! UNKNOWN SWITCH ARGUMENT. **–** Behavioral and decoding results (In box plots: grey lines indicate individual participant results; diamonds indicate condition mean; notches indicate confidence interval of median; plot size indicates data volume; * = p < 0.05; ** = p < 0.01; *** = p < 0.001). **A)** Participants were faster to respond to targets at high-probability (HP) locations than at low-probability (LP) locations. **B)** Participants were faster when targets were presented at the same location on sequential trials (repeat trials) than when they switched to a new location (switch target). **C)** Participants were slower when distractors were present than when they were absent. Participants were, additionally, especially slow when distractors were present at the HP target location. **D)** Decoder results comparing ping and no-ping trials. Data was baselined in the -200 to 0ms window pre-ping onset. Shaded areas represent standard error of participant means. Lines are smoothed using the scipy function gaussian_filter with an alpha of 1.5. Lower red bars represent significant clusters identified wherein decoding of ping trials was above chance. Black bars represent clusters identified in which ping and no-ping decoding differed significantly.

Decoding: Visual pings reveal the priority landscape in anticipation of search display onset.

After having validated that locations with a higher target probability were prioritized over other locations via reaction time measurements, we set out next to examine whether this prioritization could be revealed in anticipation of search display onset, and, critically, whether such decoding was dependent on the presentation of high-contrast visual 'pings'. For this purpose, we trained a classifier on the response pattern of 64 electrodes using all four high probability target location as classes (see Methods for full detail) separately for ping and no-ping trials. As visualized in Figure 2D, we observed robust above chance decoding of the high probability target location, but critically only from the evoked activity elicited by visual pings. By contrast, in the absence of a visual impulse, ongoing EEG signals appeared to contain little to no information about the current high probability target location. Indeed, cluster-based permutation tests across time confirmed that decoding in ping trials not only reliably differed from chance, but also from no-ping trials. The observed high probability target location decoding following visual pings is in line with the idea that learned attentional priority is encoded in a latent layer of the attentional priority map, mediated by dynamic changes of synaptic weights underlying spatial attention networks leading to priming of learned responses in preparation for new sensory input. The current results show that it is then possible to decode the weights of this otherwise activity-silent spatial priority map via a salient, task-irrelevant ping. However, before reaching this conclusion, several pre-planned and unplanned alternative explanations were investigated.

Decoding effects cannot be explained by temporal correlation

A concern in the present design is the blocked nature of the high probability distractor location, such that classification labels not only signaled a unique spatial high probability location, but also temporally separate phases in the experiment (see Figure 1B). While the absence of reliable decoding in the no-ping trials rules out that decoding is driven by oscillatory temporal artefacts, it is nevertheless possible that the evoked response elicited by the ping varied over time resulting in spurious decoding (Z. Huang et al., 2015; Linkenkaer-Hansen et al., 2001; Palazzo et al., 2020). To explore whether the ping evoked response produced meaningful decoding, in an unplanned control analysis we repeated the main analysis, but rather than classifying high probability target locations, the model was trained to decode the position of the target in the preceding search display. Although less pronounced, this analysis again resulted in selective decoding following ping onset, and critically this decoding could not be explained by temporal artefacts as trials were now randomly sampled from various timepoints (Figure 3A). To further rule out temporal structure as a confound, in another unplanned control analysis, we explored whether it was possible to decode distinct temporal phases in the experiment. That is, we artificially split the experiment in three separate phases, with each phase containing the final block in a high probability location sequence, the subsequent neutral block and the first two blocks of the following regularity ('dummy decoding' in Figure 1B). We chose these blocks specifically to minimize overlap with the high-probability conditions while matching trial volume as closely as possible to the original decoding (224 trials). When these 'dummy' labels were passed to the exact same decoding pipeline as in the preceding analyses decoding collapsed, never deviating significantly from chance level (Supplementary Figure 3A). Together, these analyses suggest that the observed decoding was indeed driven by learned latent attentional biases in response to the high probability location manipulation.

Decoding was driven both by statistical learning and inter-trial effects

It is well known that in visual search response times are affected by intertrial priming in a similar way as statistical learning affects response times (Maljkovic & Nakayama, 1994; Theeuwes & Van der Burg, 2011). Indeed, consistent with previous studies on intertrial priming, the current experiment also shows that response times were significantly faster when the target location repeated from one trial to the next (Figure 2B). Critically, the ping which was presented during the intertrial interval of the current trial was able to also retrieve which location contained the target location on the preceding trial (Figure 3A). Therefore, due to our experimental design in which there is a higher proportion of ping trials preceded by targets at the high probability location, it is possible that the observed decoding does not reflect learning across longer time scales, but instead solely reflects intertrial priming effects. To determine whether this was the case, in an unplanned analysis we repeated the main analysis after excluding all trials in which the preceding trial contained a target at the current high-probability (HP) location (37.5% of trials, called HP-exclusionary). Under these conditions, ping trials continued to show significant decoding of the high-probability location (Figure 3B). A further exploratory analysis was done to test whether decoding would remain high when decoders were trained *only* on pings following trials in which the target was at the current high-probability location (called HP-only decoding). Decoding continued to remain high despite a significantly lower trial count than the HP-exclusion analysis, indicating that the decoding contribution was approximately equal between HP trailing and non-trailing trials despite the trial imbalance (Figure 3B). Importantly, while intertrial effects can be separated from statistical learning effect via trial exclusions, statistical

learning effects cannot be separated from intertrial effects as they are present in all trials. Decoding in the HP-exclusionary condition of Figure 3B should thus be thought of as a combination of statistically learned enhancement and serial biases, while Figure 3A better reflects the pure intertrial effect. Overall, these results match the behavioral results, where significant effects of inter-trial priming were found on top of statistically learned spatial enhancement, and indicate that the ping is able to retrieve the priority landscape both induced by serial priming from the previous trial and statistical learning across longer timescales.

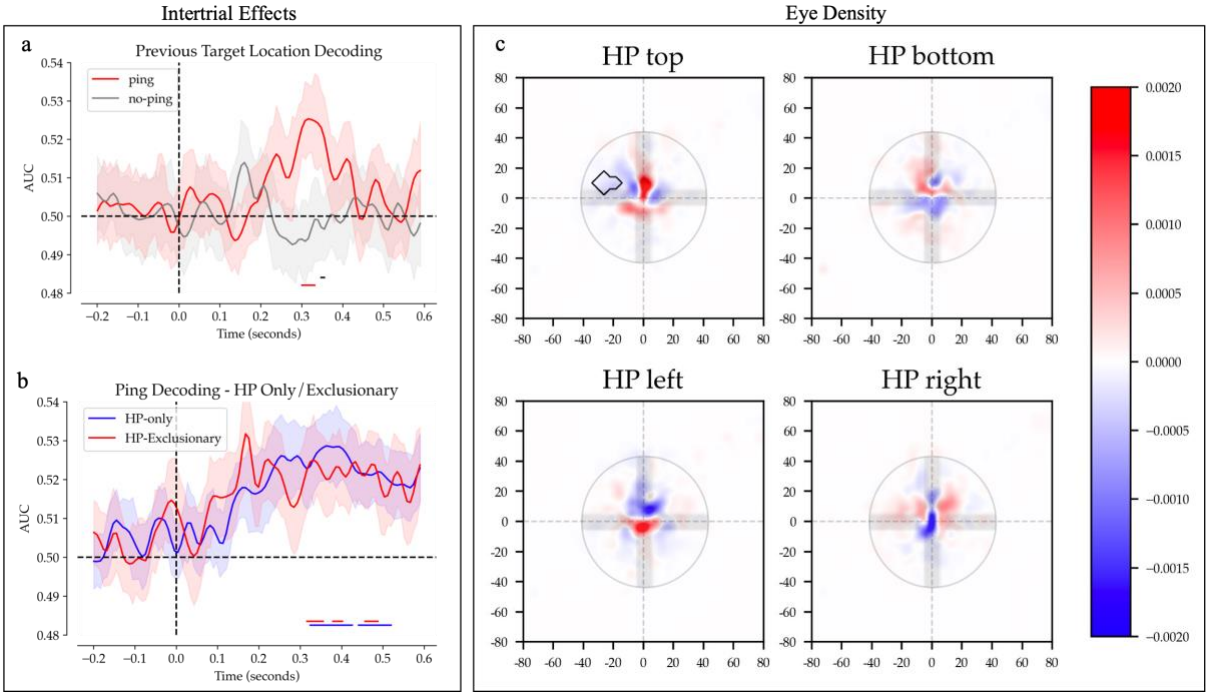

FIGURE 3 – Control Analyses. Solid colored bars below decoding results indicate significant clusters as identified via permutation tests (see Methods). All results were smoothed using a gaussian filter. Shaded areas represent standard error of mean. **A)** Decoder results when passed preceding trial target location as labels instead of current high-probability target location. **B)** Decoding results when decoders are trained only on ping epochs which followed a trial in which the target was at a HP location (blue) or excluding all such HP trailing epochs (red). Note that due to low trial counts, no trial averaging was possible for HP-only decoding (See Methods). **C)** Eye-density plots showing average eye position in the 600ms window following pings. Red represents regions with higher densities than combined average of other three conditions. Blue represents regions with lower densities. Black outlines indicate significant clusters identified using permutation tests (see methods). X and y-axis ticks indicate pixel distance. Color bar indicates proportion out of 1. The fixation dot has been superimposed over the densities for reference.

Decoding effects cannot be explained by systematic shifts in gaze position

A common concern in multi-variate EEG analyses is the influence of eye movements on decoding (Mostert et al., 2018). For example, the position of the eye can influence neural decoding when participants make (small) eye movements, affecting how the image is projected to early visual cortex and thus possibly leading to distinct patterns of activity (Mostert et al., 2018; Tse et al., 2010). There is some reason to believe that if eye movements were significantly influencing decoding, then they should have also influenced decoding in no-ping trials as the fixation cross remained present on the screen, thus leading to patterns of activity which should have been picked up by the decoders. This assertion is further supported by the fact that foveal receptive fields are much denser parafoveal fields, meaning minute movements should have evoked an outsized response to the foveal fixation

point in comparison to the peripherally presented ping stimuli (Harvey & Dumoulin, 2011; Van Essen et al., 1984). Despite this reasoning, and to address the issue of eye movements, we systematically investigated the extent to which the decoding results could be attributed to eye-related artefacts. First, we examined whether gaze was shifted relative to the high probability target locations. To do this, we calculated average gaze position for each of the HP trials across participants to form density maps of gaze positions per condition. As is evident from the heat maps of gaze density, this analysis yielded no evidence that gaze was systematically shifted towards the high probability target location (Figure 3B). The only significant cluster found was a negative cluster to the left of the fixation dot when the high-probability location was the top of the screen (cluster-based permutation test, $p < 0.05$). Excluding this condition from the main analysis did not change the overall pattern of results.

While the heatmaps did not suggest any systematic eye movements across participants in relation to the high-probability location, there is still the possibility that eye movements differed between high probability target locations, but in a way that does not generalize across participants. To test this with an unplanned analysis, we adopted the same procedure as in the preceding analyses, but instead entered the horizontal and vertical gaze position as features. As visualized in Supplementary Figure 4A there was reliable above chance decoding, independent of ping presence. When decoder results from the eye-tracker-trained decoders was compared to the EEG-trained decoders, no significant correlation was found, indicating high decoder accuracy in the eye tracker decoders did not predict high decoder accuracy in the post ping window in the EEG-trained decoders (Supplementary Figure 4B). Because eye-movements seemed to differ systematically on a participant level, but not at the group level, it was plausible that these drifts represented temporal artifacts rather than meaningful evoked movements. To examine this possibility, the same control analysis was done for the eye tracking decoding as for the EEG decoding to test for temporal confounds. Eye tracker data was passed to a decoder trained on 'dummy' labels with no meaningful overlap with experimental conditions (Supplementary Figure 3B) and produced the same pattern of results as when passed meaningful labels, a result that was not found for the EEG analysis. This indicates that the decoding observed in the EEG analysis was not the result of ping-related eye confounds, but rather that above chance eye tracker decoding likely represents temporal noise in the eye tracking data.

General Discussion

How does our past experience influence future behavior, and how can we study this latent bias? We sought to answer this question on the level of attentional deployment by testing whether the ping technique (a method used previously to reveal latent working memory content) could be used to visualize the layer of the attentional priority map maintained by selection history. Consistent with an activity-silent model of learned attentional bias, here we demonstrate that when observers learned to prioritize a given location in space (as indexed by faster target selection and increased distractor interference at that location), the ongoing EEG signal did not contain information about the current high probability location. Critically however, this otherwise hidden latent priority landscape could be revealed by inserting neutral visual ping displays in between search displays. This above chance decoding could not be attributed to temporal artefacts, nor could it be explained by systematic biases in gaze. Instead, high probability location decoding indicates that the ping technique can be used to visualize not only latent memory representations but also our latent attentional biases.

The current evidence suggests that updating of local priorities across visual searches accrues extremely fast (Valsecchi & Turatto, 2021), which makes spatial probability learning very flexible (Theeuwes et al., 2022). Indeed, in the current study where the high probability location was not static but systematically varied across time, learned priority was quickly adjusted in response to a location change. One cannot exclusively attribute this observed benefit to statistical learning however, as such statistical learning is naturally conflated with intertrial priming effects (Goschy et al., 2014; Jiang et al., 2013). Interestingly, it was recently observed in monkeys and humans that intertrial effects are mediated by activity silent mechanisms similar to those proposed for working memory (Barbosa et al., 2020). Consistent with this notion, we found that the target location on the preceding trial could also be decoded in response to the onset of a visual ping. Critically, our control analyses demonstrated that the observed high probability decoding could not exclusively be attributed to such intertrial priming effects, as decoding showed more or less identical time courses when high probability location decoding was limited to trials with a target on that location or after excluding those same trials. We thus conclude that the visual ping is able to both envision selection history effects on a very short time scale (i.e., intertrial priming) and effects that arise across longer time scales (i.e., statistical learning). These results support the notion that intertrial effects are mediated by synaptic mechanisms, as well as provide for the first time, neural evidence that statistical learning may also be mediated by such network level mechanisms in the brain. Furthermore, these findings suggest that 'pinging' may become a useful tool for the study of both statistical learning as well as intertrial effects, and may be expanded to the study of other selection history effects (e.g. reward based history effects; Failing & Theeuwes, 2018).

Dynamic shifts in synaptic weights have previously been proposed as a viable mechanism for spatial attention (Desimone & Duncan, 1995; Zelinsky & Bisley, 2015), and pinging techniques have been proposed as candidates for visualizing these network level attentional changes in the past (Rose et al., 2016). However, the enthusiasm for the study of these network level influences has been blunted by the fact that top-down attention does, in fact, produce ongoing, measurable neural activity (Foster et al., 2017; Samaha et al., 2016; Worden et al., 2000). Selection history effects, however, have recently been shown not to be driven by similar active neural mechanisms (Adam & Serences, 2021; Barbosa et al., 2020; Moorselaar & Slagter, 2019; van Moorselaar et al., 2021). The current results build upon these findings by demonstrating that the ping technique can in fact be used to reveal activity-silent history modulated attentional bias. While these findings seem to suggest that selection history influences exert themselves at the level of synaptic weight changes, some caution is called for in this interpretation as several open questions remain over the neural mechanism's underlying ping-revealed structures as well as why these structures are usefully revealed via these transient pings.

As decoding is by itself uninformative about the underlying neural representation (Naselaris & Kay, 2015; van Moorselaar & Slagter, 2020), the nature of these activity-silent states that the ping succeeds in visualizing remains an open question. In the field of working memory, the debate over why pings reveal otherwise hidden memory content has focused on the question of whether the decoded memories are encoded in truly latent networks, mediated by neuroplasticity, and which are then reactivated by the ping, (Rademaker & Serences, 2017; Stokes, 2015; Wolff et al., 2017, 2021), or alternatively whether these memories are simply mediated by ongoing neural activity below a certain detection threshold case, the ping would simply serve to reduce signal variance such that these states can be visualized (Bae & Luck, 2019; Barbosa et al., 2021; Schneegans & Bays, 2017). While these questions are ultimately better answered via neurophysiology, in the current case of ping-evoked

decoding of learned attentional preparation the existing literature favors the former interpretation, as selection history, whether in a form of statistical learning or intertrial priming, has generally exhibited none of the characteristics of active neural processes (Adam & Serences, 2021; Barbosa et al., 2020; Moorselaar & Slagter, 2019; van Moorselaar et al., 2020). Furthermore, while active traces, may in principle, be a plausible neural mechanism for driving intertrial priming, such active mechanism seem unlikely to drive statistically learned spatial biases as these effects are variously known to persist over very long time periods ranging from minutes (Britton & Anderson, 2020; Duncan & Theeuwes, 2020; Valsecchi & Turatto, 2021) to weeks (Chun & Jiang, 2003; Jiang et al., 2013; Turatto et al., 2018).

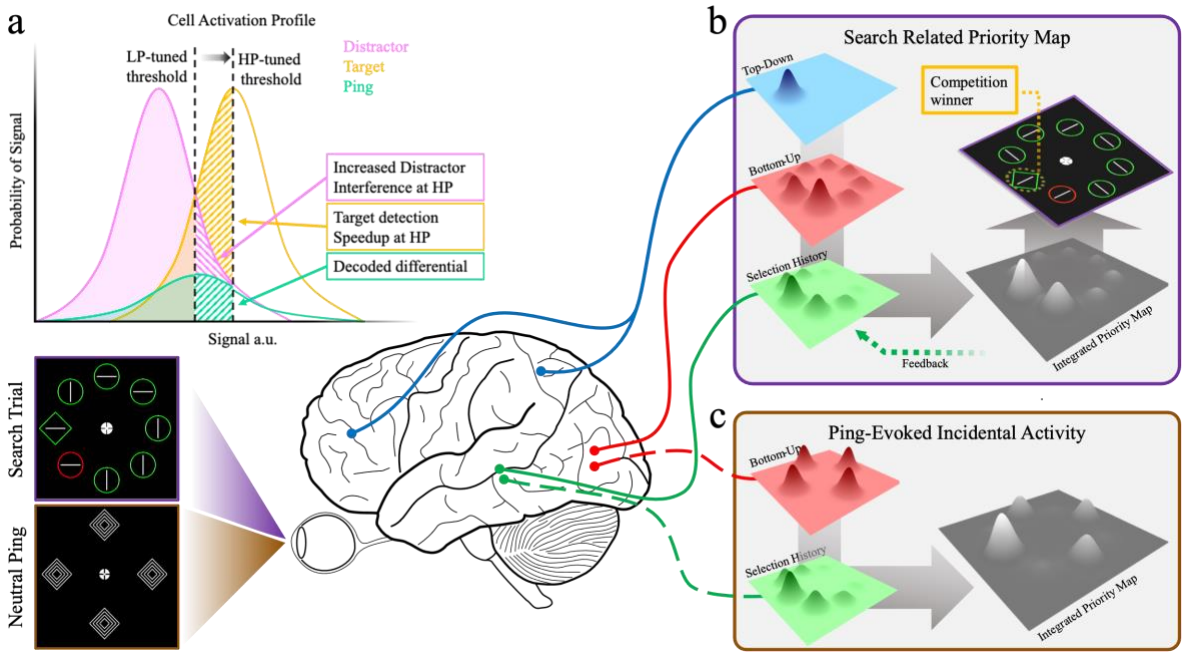

Figure 4 – Possible mechanisms underlying search priority and ping evoked activity **A)** Proposed tuning profiles of spatial cells in priority map. Cells tuned to HP locations are proposed to adopt a more liberal activity threshold, thereby leading to more correct hits (faster correct RT's) as well as more false alarms (greater distractor capture) in relation to stimuli presented in that location. Importantly, this shifted bias also leads to more incidental activity when pings are presented. This differential in incidental activity is proposed as driving ping decoding. **B)** Integrated activity from top-down, bottom-up and selection history effects driving attentional selection in an example search task. Attention is directed to region with highest activity, in this case the far-left location (note: the brain schematic is based on figure used in Theeuwes et al., 2022). **C)** Ping-evoked incidental activity integrating bottom-up and selection history effects. Top-down effects are absent as no task goals were present in relation to pings.

A further important unanswered question remains: why exactly do these otherwise hidden structures produce discernable activity pattern in response to a neutral, salient ping? In the case of working memory, it has been proposed that the observed reactivation mirrors the way that active sonar uses reflections of audio pulses to reveal underwater structures (Rademaker & Serences, 2017; Stokes, 2015; Wolff et al., 2015, 2017). While a useful tool for conceptualizing the pinging technique, the reflections of audio pulses underwater is clearly mechanically very different than the interaction of visual pulses with hidden states in the brain (Rademaker & Serences, 2017). Alternatively, it has been proposed that a latent history-dependent filter may drive decoding in working memory versions of ping experiments (Sugase-Miyamoto et al., 2008; Wolff et al., 2017), an explanation which is loosely valid in our attention variant of the paradigm and which has support in various template theories of attention (Desimone & Duncan, 1995; van Moorselaar et al., 2014; Wolfe, 1994). However, a filter match

model does not necessarily imply irrelevant stimuli should also activate these templates; thus, further work must be done to justify this filter correlated activity. We propose that the most parsimonious explanation for this decoding is that the sudden, salient stimulation of the ping drives systematic mis-activation patterns by dynamically task primed neurons (Figure 4). Firstly, the influence of selection history is proposed to lower activation thresholds of predicted stimulus tuned neurons, illustrated in Figure 4A in a shifting of the activation threshold. This liberalization would then account both for speeded responses when targets were correctly predicted, but also increased distractor interference when present at expected target locations (illustrated by larger areas under their respective curves, see figure 4B for a hypothetical trial). Secondly, this shifted activation profile should also lead to increased incidental misfiring of primed neurons when presented with the task irrelevant ping, illustrated by the turquoise curve in Figure 4A. Incidental misfiring of neurons to the presentation of non-tuned stimuli is a familiar phenomenon in research using single unit recordings (see, for instance, Quiroga et al., 2005). If latent learned attentional priority is mediated by neural structures entering a 'primed' neural state via processes of neural plasticity (Figure 4A), then the passing of irrelevant but high-contrast pings through the visual cortex may incidentally misfire these primed neurons at a high rate. This rate of misfire, then, would be the weak signal to which the decoders are sensitive (Figure 4C).

Our findings also inform the debate whether learned expectations exert their influence already in advance (Fiser et al., 2016) or, alternatively, only become apparent in response to sensory stimulation (Bar et al., 2006; Y. Huang & Rao, 2011). Based on behavioral studies that randomly intermix search displays with probe displays, where the latter allows to take a peek at selection priorities immediately prior to search display onset, it has been argued that learned attentional biases are already evident pre-stimulus (C. Huang et al., 2021, 2022). Consistent with such proactive enhancement in the spatial priority map, here the ping evoked decoding indicates preparatory spatial tuning towards the high probability location before the onset of the actual search display. At the same time however, this proactive tuning was not evident in ongoing oscillatory activity and only became apparent in response to sensory stimulation by the visual pings. Therefore, it appears that while statistical learning proactively adjusts the spatial priority map, this priority landscape only becomes apparent after integration of bottom-up sensory input, such as a probe display, or in this case visual pings (but see Wang et al., 2019).

A number of influential new theories of attention have highlighted the need to integrate selection history into models of dynamic attentional priority (Awh et al., 2012; Chelazzi et al., 2019; Liesefeld & Müller, 2019; Pearson et al., 2022; Theeuwes et al., 2022). The current results advance this project by providing a novel method of visualizing the history-mediated layer of the attentional priority map while also suggesting the neural mechanisms underlying such latent biases. While priority maps have typically been studied in their influence on spatial attention, it has also been noted that they can easily be thought of as a general cognitive tool useful in the conceptualization of any goal directed behaviour (Zelinsky & Bisley, 2015). Under this view, the current findings may represent just the first step in expanding the application of the ping technique to reveal latent biases in the brain. Fields of study that have previously been excluded from neuroimaging research should now reconsider the possibility of undertaking study of latent neural states using similar pinging approaches, as much future research will undoubtedly reveal the extent of the ability of the pinging technique to reveal previously hidden structured in the brain.

Methods

Participants

Based on previous pinging studies on working memory (Wolff et al., 2015, 2017a, 2020) we included 24 participants (17 female, mean age 24) in our final dataset after replacement of eight participants. Two participants were replaced because of accuracies 2.5 standard deviations (SD's) below the group average on the behavioral task; two as a result of overall response times slower than 2.5 SD's from the group mean; two for producing low quality encephalography as revealed through visual inspection, and two for failing to maintain fixation during the task as revealed through the electro-oculogram and eye-tracking analysis (saccades detected on more than 30% of trials, see eye-tracing acquisition and preprocessing below). The study was conducted at Vrije Universiteit Amsterdam, all participants indicated normal or corrected-to-normal vision and reported no history of cognitive impairments. All participants provided written informed consent prior to participation and were compensated with 25 euros or course credits. The study conformed to the Declaration of Helsinki and was approved by the Ethical Review Committee of the Faculty

Apparatus and Task

Participants were seated in a dimly lit sound-attenuated room with a chinrest 60cm away from a 23.8 inch, 1920x1080 pixel ASUS ROG STRIX XG248 LED monitor with a 240hz refresh rate upon which all stimuli were presented. The behavioral task was programmed using OpenSesame (Mathôt et al., 2012) and utilized functions from the Psychopy library of psychophysical tools (Peirce, 2007). EEG data were recorded with default Biosemi settings at a sampling rate of 512 Hz using a 64-electrode cap with electrodes placed according to the international 10–10 system (Biosemi ActiveTwo system; biosemi.com) with two earlobe electrodes used as offline reference. Vertical (VEOG) and horizontal EOG (HEOG) were recorded via external electrodes placed ~2 cm above and below the right eye, and ~1 cm lateral to the left and right lateral canthus. Eye-tracking data were collected using an Eyelink 1000 (SR research) eye tracker tracking both eyes. Participants used a headrest for stability positioned 60cm from the screen. Sampling differed between participants between 500, 1000 and 2000hz due to the different EEG labs used having different versions of the Eyelink 1000 with variable sampling limits. All participant data was later standardized to 500hz during preprocessing. Calibration was done before the first block and at the halfway point for all participants. For some participants, additional calibrations occurred due to subtle changes in resting position in the chinrest or other factors which caused noticeable drift in the calibrated signal.

The task was a modified version of the additional singleton task with underlying stimulus regularities known to provoke statistical learning (Theeuwes, 1992; Wang & Theeuwes, 2018). All stimuli were presented on a black background. Each trial started with a blank screen randomly jittered between 200-500ms in length. A circular fixation point (40 px diameter, ~1.1°) was then presented alone between 1300-1700ms at the center of the screen. The fixation point design was taken from Thaler et al. (2013) who showed their fixation dot design provided the most stable fixation results from a range of possible designs. Participants were informed to maintain fixation throughout the trial and not to saccade towards any peripheral stimuli. Participants received feedback when the eye tracker detected a fixation deviation above 1.5° away from fixation in the form of a low volume audio beep.

During the fixation period, on 50% of trials a high-salience visual 'ping', which appeared at a time point randomly jittered between 700-900ms after fixation onset, was presented for 200ms (Figure 1C). Pings were

comprised of four white (RGB) high contrast shapes, either all diamonds (diagonal length 116 pixels; ~3°) or all circles (diameter 90 pixels; ~2.4°), presented in the four cardinal directions (up, down, left and right) on the same locations where subsequent search shapes appeared (see below). The ping shapes comprised of a large outer shape with three smaller shapes recursively embedded within the outer shape. Following ping offset, the fixation point remained on screen for a randomly jittered duration of 400-600ms. To contrast ping and no-ping trials, a ping trigger was created on both trial types, but on no-ping trials these triggers did not correspond with an actual ping presentation; instead, the screen remained blank except for a fixation dot until the next trial began.

An additional singleton search display (Theeuwes, 1992) was next presented, which remained on screen for 2000ms or until participants provided a response (Figure 1A). Each search display contained eight equidistantly placed shapes on an imaginary circle (radius 185 pixels; ~4.8°) centered on fixation. The shapes, which could either be circles (diameter 90px, ~2.4°) or diamonds (82x82 px or ~ 2.1°x2.1° square rotated 45 degrees) and could either be colored red (rgb 255,0,1) or green (rgb 0,128,0), all contained a white (rgb 255,255,255) line (70px; ~1.8°) oriented horizontally or vertically bar. On each display, one shape was unique from the rest (either one circle among seven diamonds or vice versa) and participants were instructed to report the orientation of the line inside this target shape via button press on a standard keyboard ('z' for horizontal lines; '/' for vertical lines). On a subset of trials (70%) one of the homogenous shapes was assigned a unique color (e.g., if the shapes were green, the distractor was red, or vice versa) rendering it a colored singleton distractor. The sole purpose of these distractors was to make the task more challenging, thereby making the target probability manipulation less apparent (see below).

Critically, target locations were not selected at random, but instead, the experiment was structured such that for several experiment blocks in a sequence, one location would be more likely to contain the target than the other locations (see Figure 1A). Specifically, in high probability target blocks, one location was disproportionately likely to contain a target, with targets appearing at this high-probability location on 37.5% of trials, making them 4.2 times more likely to contain a target than any of the other, low-probability locations. To be able to decode the high probability target location, high probability locations did not remain static for an entire experimental session, but instead periodically shifted across the cardinal locations (up, down, left and right). Specifically, every four blocks the high probability location changed to a new location (order counterbalanced across participants), with every change being preceded by a neutral block, in which targets appeared in each of the eight locations with equal probability (see Figure 1B). These neutral blocks served to, at least partly, unlearn the acquired spatial priority settings such that observers entered a new sequence of learning blocks without robust lingering attentional biases (Britton & Anderson, 2020; Duncan & Theeuwes, 2020; Valsecchi & Turatto, 2021) [1]..

Each of the 19 experiment blocks and the preceding practice block consisted of 56 trials. Breaks were offered in-between experimental blocks where participants were informed of their progress, accuracy, average reaction time and were encouraged to rest their eyes. Following the completion of the experiment, the participant was asked to answer an additional four debrief questions (for three participants the debrief was collected verbally). Firstly, they were asked if they noticed the target tended to appear in certain locations more frequently than others. Secondly, they were asked to indicate where they believed the target was most frequently present on the final experiment block which they had just completed. If they were not sure they were instructed to provide a best

[1] Studies which have shown the time course of the extinction of learned attentional bias in the past have exclusively used high-probability *distractor* locations, thus our design was the first to examine the time course of learned *target* enhancement extinction.

guess. Thirdly they were asked if they felt any of the other seven locations at some point in the experiment was more likely to hold a target. Finally, the participant was asked if they had ever performed a task similar to the one they had just finished.

Analysis

All analyses involving EEG and eye tracking data were done using custom Python scripts. All behavioral analyses were programmed using R. All preprocessing and analysis scripts are available online and can be downloaded at https://github.com/dvanmoorselaar/DvM. Functions from the open-source MNE analysis package were heavily used and central to our analyses (Gramfort et al., 2014). All analysis steps followed those listed in our preregistration unless otherwise stated.

Behavioral Preprocessing

If a participant's overall accuracy was 2.5 SD's below group mean, then that participant was excluded from our dataset and replaced. Average accuracy after exclusions was 89% (range 81-97%). Next, participant data was restricted to correct trials and mean and SD based on each individual participant's RT data were then calculated. Trials 2.5 SD's faster and slower than each participant's average RT were excluded, resulting in an average of 1.1% of abnormally slow/fast trials excluded per participant. The group average reaction time (RT) was then calculated and participants 2.5 SD's away from this mean were excluded and replaced. Combined exclusion of incorrect responses (~11%) and data trimming (~1%) resulted in an overall loss of approximately 12% of total trials

EEG acquisition and preprocessing

Following re-referencing of all EEG data to the average of the two earlobe electrodes, the data was high-pass filtered at 0.01Hz using Hamming windowed FIR filter to remove slow signal drifts. Continuous EEG was then epoched from -700 to 1100ms relative to ping display onset, or relative to matching timepoints on no ping trials, with trial rejection procedures being limited to a smaller time window (i.e., -200 – 600ms). Prior to subsequent trial rejection and artefact correction, malfunctioning electrodes marked as bad by the experimenter during recording were temporarily removed. First, EMG contaminated epochs were identified with an adapted version of an automated trial-rejection procedure as implemented in Fieldtrip (Oostenveld et al., 2011). To specifically capture muscle activity we used a 110-140 Hz band-lass filter and allowed for variable z-score thresholds per subject based on within-subject variance of z-scores (Vries et al., 2017). Moreover, to reduce the number of false alarms, rather than immediate removal of epochs exceeding the z-score threshold, the algorithm first identified the five electrodes that contributed most to accumulated z-score within the time period containing the marked EMG artefact[2]. Then in an iterative procedure, the worst five electrodes per marked epoch were interpolated using spherical splines (Perrin et al., 1989) one by one, checking after each interpolation whether that epoch still exceeded the determined z-score threshold. Epochs were selectively dropped when after this iterative interpolation procedure the z-score threshold was still exceeded. Second, Independent Component Analysis (ICA) as implemented in the MNE (using the 'picard' method) was fitted on 1 Hz high pass filtered epoched data to remove

[2] We chose to deviate from our preregistered method here as the new interpolated method allowed us to preserve more data. Results were also analyzed using the preregistered exclusion criteria and did not differ in the pattern of results.

eye-blink components from the cleaned data. Third, manually marked malfunctioning electrodes were interpolated using spherical splines.

Lastly, we excluded trials in which a saccade was detected by the eye tracker, or significant drift was recorded through the course of the trial. To detect drift in our eye tracking data, we first baselined our eye tracking data on the 300ms pre-ping window. We then observed the maximum deviation from zero measured in segments of data 40ms long for each trial. If this maximum exceeded 1° then the trial was marked for exclusion (~8% of data per participant). If the number of trials excluded due to saccades or drift exceeded 30% of the total trials, their eye tracking data was further examined. If it was revealed that eye tracking data was exceptionally low quality, then their eye tracking data was ignored and their data was cleaned based on the HEOG recordings instead (one participant had their data treated in this way). If their eye tracking data was found to be of high quality, then the participant was excluded from the analysis and their data replaced with a new participant (two participants were excluded in this way). For participants with no reliable eye tracking, we identified trials with sudden jumps in the recorded EOG facial electrodes using a step method with a window of 200ms, a step size of 10ms and a threshold of 20 μV, and excluded these trials.

Eye-tracking acquisition and preprocessing

One participants eye data was excluded due to poor quality tracking (see above). For the remaining 23 datasets, eye data was first converted from the native EDF format to the ASC format before further processing. The eye-tracking data was then epoched around the ping onset event, or relative to matching timepoints on no ping trials, including -200ms before onset and 600ms after onset to match EEG data epochs. Blinks were then identified in the eye-tracking signal and linearly interpolated using custom scripts. A 200ms pad was applied before and after any identified blink time window to ensure movement artifacts before and after blinks were not accidentally preserved. All eye data was then baselined relative to the 200ms window prior to ping onset when participants held fixation. If the sampling rate of the participant was above 500hz, their data was next down-sampled to 500hz using scipy resample function. Eye tracking data was then trial matched to the processed EEG data such that only trials included in the EEG analysis were also included in any eye tracking analysis.

Eye-density Calculation

For each participant, eye data was first separated between the four HP conditions and concatenated. Next, each datapoint was rounded to the nearest quarter of a pixel and centered around zero. Density was then calculated for every quarter pixel by dividing the count of the number of timepoints at each (x,y) location by the total number of timepoints. Next, density differences were calculated by cycling through each of the four conditions and subtracting every cell of the current density matrix by the averaged density matrix of the other three conditions. The group average density figure was then generated by averaging across the 24 individual density matrixes for each HP condition. Finally, a cluster permutation test was done for each of the four locations, across the 24 datasets (see statistics).

Decoding Analysis

To test whether the ping evoked discernable activity based on the different high probability locations used across various blocks, we applied multivariate pattern analysis (MVPA) using a cross-validated linear discriminant

analysis (LDA) on ping and no-ping trials separately with all 64 electrode channels as features and block high probability locations as classes. EEG data was first baselined in a window from -200 to 0ms prior to ping onset trigger (for no-ping trials, no stimuli was actually presented but instead the screen remained blank except for a fixation point), and subsequently down sampled to 128 Hz. While not preregistered, to further increase the signal-to-noise ratio in our data (Grootswagers et al., 2017), we adopted a trial-averaging approach and subsequently transformed our data using principal component analysis (PCA). Specifically, we averaged over three trials of the same exemplars, after which the data was PCA transformed in each cross-validation run (see below) with a model fitted on the training data only. The main decoding results without these additional signal-to-noise boosting effects are included as Supplementary Figure 5.

The data was next randomly split into 10 equally sized subsets where each class (i.e. the four possible high probability locations) were selected equally often in each condition so that training would not be biased towards any one class. Next, a leave-one-out procedure was used such that that each classifier was trained on nine folds and tested on the one excluded fold until each fold was tested once; thus, ensuring that training and testing never occurred using the same data. Classifier performance was then averaged across the ten folds. We used an Area Under the Curve (AUC) approach to rate classifier performance, which is an approach considered sensitive, nonparametric and criterion free measure of classification performance (Hand & Till, 2001). Using AUC, a rating of 0.5 is considered chance level classification. An analysis score was collected for every time point, showing how decoding performance fluctuated across time. We chose to decode over the time period from -200ms before ping onset until 600ms after ping onset (or 400ms after ping *offset*, the minimum amount of time that could pass before a search display appeared, see Figure 1D).

Statistics

RT and accuracy were analyzed using simple two-sided t-tests and analyses of variance (ANOVA's). In case of insignificant results, where applicable, we run Bayesian equivalents of the above specified tests. Decoding scores were evaluated for both ping and no-ping trials via cluster-based permutation tests and paired sampled two-sided t-tests with cluster correction ($p = 0.05$ and 1024 iterations) using MNE functionality (Gramfort et al., 2014). A similar permutation test was used to compare whether the difference between decoding on ping and no-ping trials was significant across time.

Supplementary Materials

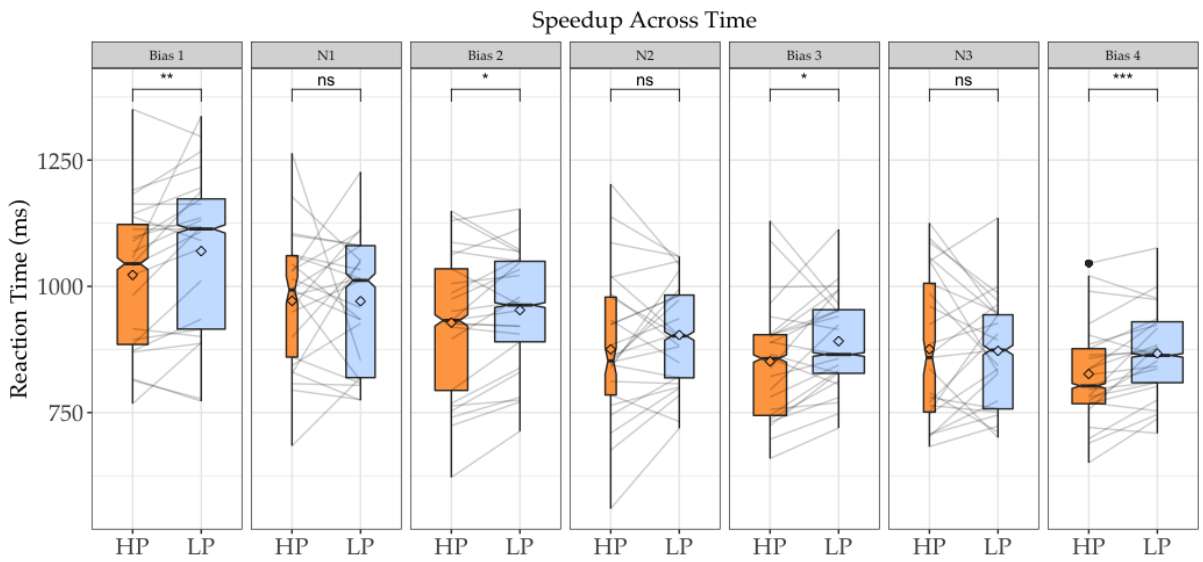

SUPPLEMENTARY FIGURE 1 - RT effects across experiment phases when targets were present at high-probability (orange) versus low-probability (blue) locations in space (see Figure 2 for additional figure nomenclature). Between each set of four bias blocks was a single neutral block (N#). For neutral blocks, there was no HP target location; comparison was thus done using the previous blocks HP location in orange to measure lingering spatial bias from the previous blocks. Note that these averages were calculated on very small trial counts (as illustrated by the narrow bar widths). Note that each bias phase contained a different high-probability location, counterbalanced across participants.

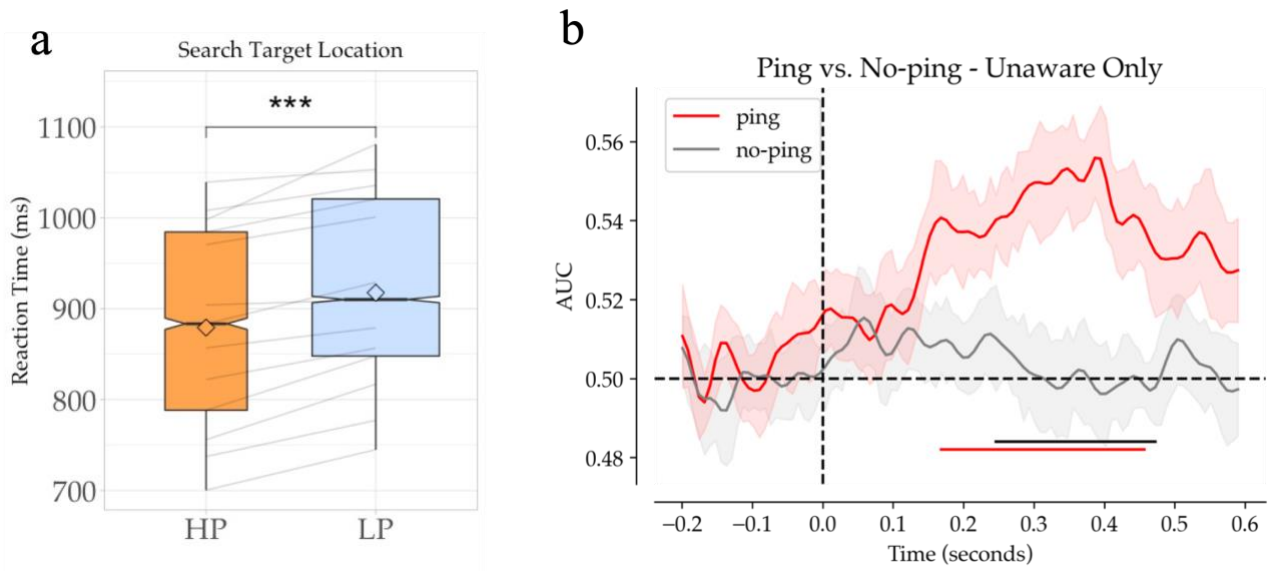

SUPPLEMENTARY FIGURE 2 - Excluding Reported-Aware Participant (N=13). Shown is the behavioral speedup effect (**A**) and decoding results (**B**) of the 13 participants that did not agree in a debrief survey that they noticed some pattern of target spatial distributions. Note that these figures also exclude those who agreed that they notice some pattern of target presentation but were unable to identify where the most recent high-probability location was.

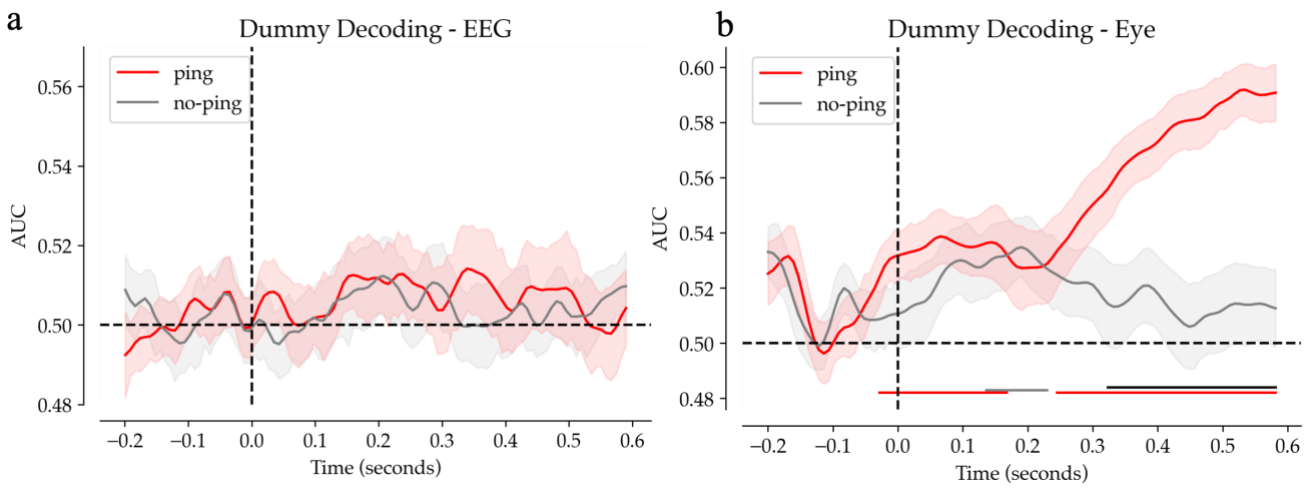

SUPPLEMENTARY FIGURE 3 – Dummy decoding. A) decoder performance when trained on EEG data and passed dummy labels (see figure 1B for illustration of dummy decoding). Note: y-axis is matched to results in Figure 2D for comparison. **B)** decoder performance when trained on eye tracking data and passed dummy labels.

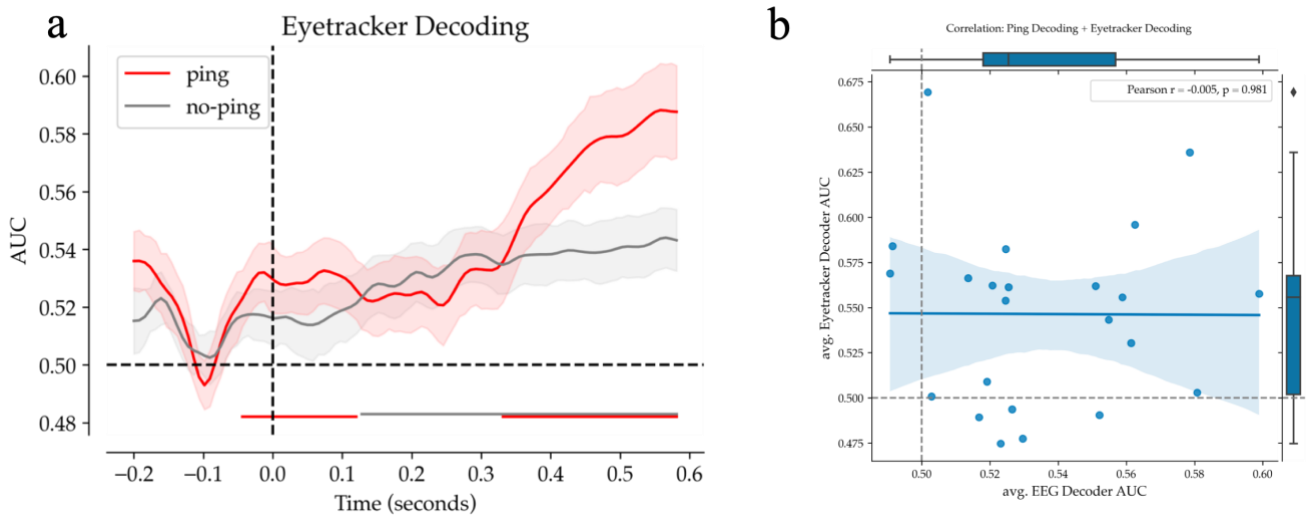

SUPPLEMENTARY FIGURE 4 – Decoders trained on eye tracker data and the correlation between eye tracker and EEG decoding. A) Shown is the decoder results when trained on eye tracking data corresponding to the four high-probability locations. Note the similarity of results to those in Figure 3D, when the same decoders were trained on meaningless labels. B) Correlation between average decoding for EEG and eye tracking decoders.

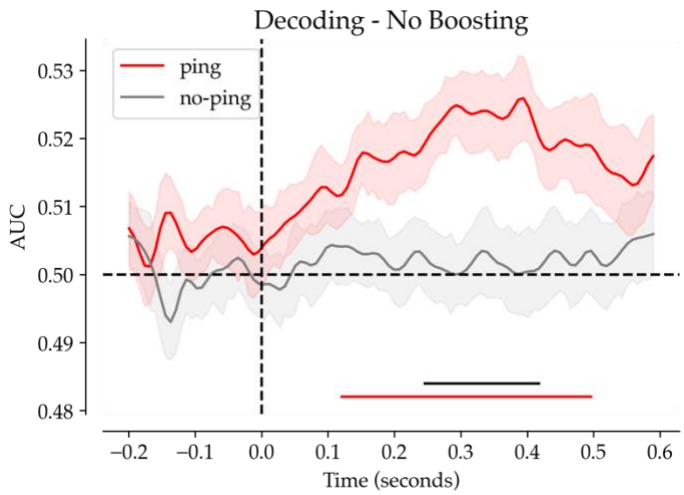

SUPPLEMENTARY FIGURE 5 – Decoding without signal-to-noise boosting. In order to increase the signal-to-noise ratio, techniques were borrowed from Grootswagers et al. (2017). These techniques were originally not planned for in the preregistration. The originally planned decoding without these boosting additions, in accord with the preregistered decoding pipeline, are shown in this figure. Note that the general pattern of results does not differ greatly in the timecourse of the overall pattern of decoding from that shown in Figure 2.

References

Adam, K. C., & Serences, J. T. (2021). History modulates early sensory processing of salient distractors. *Journal of Neuroscience*, *41*(38), 8007–8022.

Anderson, B. A., Kim, H., Kim, A. J., Liao, M.-R., Mrkonja, L., Clement, A., & Grégoire, L. (2021). The past, present, and future of selection history. *Neuroscience & Biobehavioral Reviews*, *130*, 326–350.

Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, *16*(8), 437–443.

Bae, G.-Y., & Luck, S. J. (2019). Reactivation of previous experiences in a working memory task. *Psychological Science*, *30*(4), 587–595.

Bar, M., Kassam, K. S., Ghuman, A. S., Boshyan, J., Schmid, A. M., Dale, A. M., Hämäläinen, M. S., Marinkovic, K., Schacter, D. L., & Rosen, B. R. (2006). Top-down facilitation of visual recognition. *Proceedings of the National Academy of Sciences*, *103*(2), 449–454.

Barak, O., & Tsodyks, M. (2014). Working models of working memory. *Current Opinion in Neurobiology*, *25*, 20–24.

Barbosa, J., Lozano-Soldevilla, D., & Compte, A. (2021). Pinging the brain with visual impulses reveals electrically active, not activity-silent, working memories. *PLoS Biology*, *19*(10), e3001436.

Barbosa, J., Stein, H., Martinez, R. L., Galan-Gadea, A., Li, S., Dalmau, J., Adam, K., Valls-Solé, J., Constantinidis, C., & Compte, A. (2020). Interplay between persistent activity and activity-silent dynamics in the prefrontal cortex underlies serial biases in working memory. *Nature Neuroscience*, *23*(8), 1016–1024.

Britton, M. K., & Anderson, B. A. (2020). Specificity and persistence of statistical learning in distractor suppression. *Journal of Experimental Psychology: Human Perception and Performance*, *46*(3), 324.

Broadbent, D. E. (1958). *The effects of noise on behaviour*.

Buonomano, D. V., & Maass, W. (2009). State-dependent computations: Spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, *10*(2), 113–125.

Chelazzi, L., Marini, F., Pascucci, D., & Turatto, M. (2019). Getting rid of visual distractors: The why, when, how, and where. *Current Opinion in Psychology*, *29*, 135–147.

Chun, M. M., & Jiang, Y. (2003). Implicit, long-term spatial contextual memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(2), 224.

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*(1), 193–222.

D'Esposito, M., & Postle, B. R. (2015). The cognitive neuroscience of working memory. *Annual Review of Psychology*, *66*, 115–142.

Duncan, D., & Theeuwes, J. (2020). Statistical learning in the absence of explicit top-down attention. *Cortex*, *131*, 54–65.

Failing, M., & Theeuwes, J. (2018). Selection history: How reward modulates selectivity of visual attention. *Psychonomic Bulletin & Review*, *25*(2), 514–538.

Fecteau, J. H., & Munoz, D. P. (2006). Salience, relevance, and firing: A priority map for target selection. *Trends in Cognitive Sciences*, *10*(8), 382–390.

Ferrante, O., Chelazzi, L., & Santandrea, E. (2021). *Statistical learning of target and distractor spatial probability shape a common attentional priority computation*. PsyArXiv. https://doi.org/10.31234/osf.io/p3sq2

Ferrante, O., Patacca, A., Di Caro, V., Della Libera, C., Santandrea, E., & Chelazzi, L. (2018). Altering spatial priority maps via statistical learning of target selection and distractor filtering. *Cortex*, *102*, 67–95.

Fiser, A., Mahringer, D., Oyibo, H. K., Petersen, A. V., Leinweber, M., & Keller, G. B. (2016). Experience-dependent spatial expectations in mouse visual cortex. *Nature Neuroscience*, *19*(12), 1658–1664.

Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K., & Awh, E. (2017). Alpha-band oscillations enable spatially and temporally resolved tracking of covert spatial attention. *Psychological Science*, *28*(7), 929–941.

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*(2), 127–138.

Frost, R., Armstrong, B. C., & Christiansen, M. H. (2019). Statistical learning research: A critical review and possible new directions. *Psychological Bulletin*, *145*(12), 1128.

Gao, Y., & Theeuwes, J. (2020). Independent effects of statistical learning and top-down attention. *Attention, Perception, & Psychophysics*, *82*(8), 3895–3906.

Gao, Y., & Theeuwes, J. (2022). Learning to suppress a location does not depend on knowing which location. *Attention, Perception, & Psychophysics*, *84*(4), 1087–1097.

Goschy, H., Bakos, S., Müller, H. J., & Zehetleitner, M. (2014). Probability cueing of distractor locations: Both intertrial facilitation and statistical learning mediate interference reduction. *Frontiers in Psychology*, *5*, 1195.

Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., Parkkonen, L., & Hämäläinen, M. S. (2014). MNE software for processing MEG and EEG data. *Neuroimage*, *86*, 446–460.

Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2017). Decoding dynamic brain patterns from evoked responses: A tutorial on multivariate pattern analysis applied to time series neuroimaging data. *Journal of Cognitive Neuroscience*, *29*(4), 677–697.

Hand, D. J., & Till, R. J. (2001). A Simple Generalisation of the Area Under the ROC Curve for Multiple Class Classification Problems. *Machine Learning*, *45*(2), 171–186. https://doi.org/10.1023/A:1010920819831

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

Huang, C., Donk, M., & Theeuwes, J. (2022). Proactive enhancement and suppression elicited by statistical regularities in visual search. *Journal of Experimental Psychology: Human Perception and Performance*.

Huang, C., Vilotijević, A., Theeuwes, J., & Donk, M. (2021). Proactive distractor suppression elicited by statistical regularities in visual search. *Psychonomic Bulletin & Review*, *28*(3), 918–927.

Huang, Y., & Rao, R. P. (2011). Predictive coding. *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(5), 580–593.

Huang, Z., Zhang, J., Longtin, A., Dumont, G., Duncan, N. W., Pokorny, J., Qin, P., Dai, R., Ferri, F., & Weng, X. (2015). Is there a nonadditive interaction between spontaneous and evoked activity? Phase-dependence and its relation to the temporal structure of scale-free brain activity. *Cerebral Cortex*, bhv288.

Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience*, *2*(3), 194–203. https://doi.org/10.1038/35058500

Jiang, Y. V., Swallow, K. M., Rosenbaum, G. M., & Herzig, C. (2013). Rapid acquisition but slow extinction of an attentional bias in space. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(1), 87.

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.

Liesefeld, H. R., & Müller, H. J. (2019). Distractor handling via dimension weighting. *Current Opinion in Psychology*, *29*, 160–167.

Linkenkaer-Hansen, K., Nikouline, V. V., Palva, J. M., & Ilmoniemi, R. J. (2001). Long-range temporal correlations and scaling behavior in human brain oscillations. *Journal of Neuroscience*, *21*(4), 1370–1377.

Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition*, *22*(6), 657–672.

Mathôt, S., Schreij, D., & Theeuwes, J. (2012). OpenSesame: An open-source, graphical experiment builder for the social sciences. *Behavior Research Methods*, *44*(2), 314–324.

Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic theory of working memory. *Science*, *319*(5869), 1543–1546.

Moorselaar, D. van, & Slagter, H. A. (2019). Learning What Is Irrelevant or Relevant: Expectations Facilitate Distractor Inhibition and Target Facilitation through Distinct Neural Mechanisms. *Journal of Neuroscience*, *39*(35), 6953–6967. https://doi.org/10.1523/JNEUROSCI.0593-19.2019

Mostert, P., Albers, A. M., Brinkman, L., Todorova, L., Kok, P., & De Lange, F. P. (2018). Eye movement-related confounds in neural decoding of visual working memory representations. *Eneuro*, *5*(4).

Naselaris, T., & Kay, K. N. (2015). Resolving ambiguities of MVPA using explicit models of representation. *Trends in Cognitive Sciences*, *19*(10), 551–554.

Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J.-M. (2011). FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Intelligence and Neuroscience*, *2011*.

Palazzo, S., Spampinato, C., Schmidt, J., Kavasidis, I., Giordano, D., & Shah, M. (2020). Correct block-design experiments mitigate temporal correlation bias in EEG classification. *ArXiv Preprint ArXiv:2012.03849*.

Pearson, D., Watson, P., Albertella, L., & Le Pelley, M. E. (2022). Attentional economics links value-modulated attentional capture and decision-making. *Nature Reviews Psychology*, 1–14.

Peirce, J. W. (2007). PsychoPy—Psychophysics software in Python. *Journal of Neuroscience Methods*, *162*(1–2), 8–13.

Perrin, F., Pernier, J., Bertrand, O., & Echallier, J. F. (1989). Spherical splines for scalp potential and current density mapping. *Electroencephalography and Clinical Neurophysiology*, *72*(2), 184–187. https://doi.org/10.1016/0013-4694(89)90180-6

Ptak, R. (2012). The frontoparietal attention network of the human brain: Action, saliency, and a priority map of the environment. *The Neuroscientist*, *18*(5), 502–515.

Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., & Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature*, *435*(7045), 1102–1107.

Rademaker, R. L., & Serences, J. T. (2017). Pinging the brain to reveal hidden memories. *Nature Neuroscience*, *20*(6), 767–769.

Rensink, R. A., O'regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, *8*(5), 368–373.

Rose, N. S., LaRocque, J. J., Riggall, A. C., Gosseries, O., Starrett, M. J., Meyering, E. E., & Postle, B. R. (2016). Reactivation of latent working memories with transcranial magnetic stimulation. *Science*, *354*(6316), 1136–1139.

Samaha, J., Sprague, T. C., & Postle, B. R. (2016). Decoding and reconstructing the focus of spatial attention from the topography of alpha-band oscillations. *Journal of Cognitive Neuroscience*, *28*(8), 1090–1097.

Schneegans, S., & Bays, P. M. (2017). Restoration of fMRI decodability does not imply latent working memory states. *Journal of Cognitive Neuroscience*, *29*(12), 1977–1994.

Serences, J. T., & Yantis, S. (2006). Selective visual attention and perceptual coherence. *Trends in Cognitive Sciences*, *10*(1), 38–45.

Silver, M. A., Ress, D., & Heeger, D. J. (2005). Topographic Maps of Visual Spatial Attention in Human Parietal Cortex. *Journal of Neurophysiology*, *94*(2), 1358–1371. https://doi.org/10.1152/jn.01316.2004

Sprague, T. C., Itthipuripat, S., Vo, V. A., & Serences, J. T. (2018). Dissociable signatures of visual salience and behavioral relevance across attentional priority maps in human cortex. *Journal of Neurophysiology*, *119*(6), 2153–2165. https://doi.org/10.1152/jn.00059.2018

Sprague, T. C., & Serences, J. T. (2013). Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. *Nature Neuroscience*, *16*(12), 1879–1887. https://doi.org/10.1038/nn.3574

Srinivasan, M. V., Laughlin, S. B., & Dubs, A. (1982). Predictive coding: A fresh view of inhibition in the retina. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, *216*(1205), 427–459.

Stokes, M. G. (2015). 'Activity-silent'working memory in prefrontal cortex: A dynamic coding framework. *Trends in Cognitive Sciences*, *19*(7), 394–405.

Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., & Duncan, J. (2013). Dynamic coding for cognitive control in prefrontal cortex. *Neuron*, *78*(2), 364–375.

Sugase-Miyamoto, Y., Liu, Z., Wiener, M. C., Optican, L. M., & Richmond, B. J. (2008). Short-term memory trace in rapidly adapting synapses of inferior temporal cortex. *PLoS Computational Biology*, *4*(5), e1000073.

Thaler, L., Schütz, A. C., Goodale, M. A., & Gegenfurtner, K. R. (2013). What is the best fixation target? The effect of target shape on stability of fixational eye movements. *Vision Research*, *76*, 31–42.

Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics*, *51*(6), 599–606.

Theeuwes, J. (2010). Top–down and bottom–up control of visual selection. *Acta Psychologica*, *135*(2), 77–99.

Theeuwes, J. (2018). Visual selection: Usually fast and automatic; seldom slow and volitional. *Journal of Cognition*.

Theeuwes, J., Bogaerts, L., & van Moorselaar, D. (2022). What to expect where and when: How statistical learning drives visual selection. *Trends in Cognitive Sciences*.

Theeuwes, J., & Failing, M. (2020). *Attentional selection: Top-down, bottom-up and history-based biases*. Cambridge University Press.

Theeuwes, J., & Van der Burg, E. (2011). On the limits of top-down control of visual selection. *Attention, Perception, & Psychophysics*, *73*(7), 2092–2103.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136.

Tse, P. U., Baumgartner, F. J., & Greenlee, M. W. (2010). Event-related functional MRI of cortical activity evoked by microsaccades, small visually-guided saccades, and eyeblinks in human visual cortex. *Neuroimage*, *49*(1), 805–816.

Turatto, M., Bonetti, F., & Pascucci, D. (2018). Filtering visual onsets via habituation: A context-specific long-term memory of irrelevant stimuli. *Psychonomic Bulletin & Review*, *25*(3), 1028–1034. https://doi.org/10.3758/s13423-017-1320-x

Vadillo, M. A., Linssen, D., Orgaz, C., Parsons, S., & Shanks, D. R. (2020). Unconscious or underpowered? Probabilistic cuing of visual attention. *Journal of Experimental Psychology: General*, *149*(1), 160.

Valsecchi, M., & Turatto, M. (2021). Distractor filtering is affected by local and global distractor probability, emerges very rapidly but is resistant to extinction. *Attention, Perception, & Psychophysics*, *83*(6), 2458–2472.

Van Essen, D. C., Newsome, W. T., & Maunsell, J. H. (1984). The visual field representation in striate cortex of the macaque monkey: Asymmetries, anisotropies, and individual variability. *Vision Research*, *24*(5), 429–448.

van Moorselaar, D., Daneshtalab, N., & Slagter, H. A. (2021). Neural mechanisms underlying distractor inhibition on the basis of feature and/or spatial expectations. *Cortex*, *137*, 232–250. https://doi.org/10.1016/j.cortex.2021.01.010

van Moorselaar, D., Lampers, E., Cordesius, E., & Slagter, H. A. (2020). Neural mechanisms underlying expectation-dependent inhibition of distracting information. *Elife*, *9*, e61048.

van Moorselaar, D., & Slagter, H. A. (2020). Inhibition in selective attention. *Annals of the New York Academy of Sciences*, *1464*(1), 204–221.

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

Vicente-Conesa, F., Fernández, T. G., Luque, D., & Vadillo, M. A. (2021). *Learning to suppress a distractor may not be unconscious*.

Vries, I. E. J. de, Driel, J. van, & Olivers, C. N. L. (2017). Posterior α EEG Dynamics Dissociate Current from Future Goals in Working Memory-Guided Visual Search. *Journal of Neuroscience*, *37*(6), 1591–1603. https://doi.org/10.1523/JNEUROSCI.2945-16.2016

Wang, B., & Theeuwes, J. (2018). Statistical regularities modulate attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *44*(1), 13.

Wang, B., van Driel, J., Ort, E., & Theeuwes, J. (2019). Anticipatory Distractor Suppression Elicited by Statistical Regularities in Visual Search. *Journal of Cognitive Neuroscience*, *31*(10), 1535–1548. https://doi.org/10.1162/jocn_a_01433

Wolfe, J. M. (1994). Guided search 2.0 a revised model of visual search. *Psychonomic Bulletin & Review*, *1*(2), 202–238.

Wolff, M. J., Akyurek, E., & Stokes, M. G. (2021). *What is the functional role of delay-related alpha oscillations during working memory?*

Wolff, M. J., Ding, J., Myers, N. E., & Stokes, M. G. (2015). Revealing hidden states in visual working memory using electroencephalography. *Frontiers in Systems Neuroscience*, *9*, 123.

Wolff, M. J., Jochim, J., Akyürek, E. G., Buschman, T. J., & Stokes, M. G. (2020). Drifting codes within a stable coding scheme for working memory. *PLoS Biology*, *18*(3), e3000625.

Wolff, M. J., Jochim, J., Akyürek, E. G., & Stokes, M. G. (2017). Dynamic hidden states underlying working-memory-guided behavior. *Nature Neuroscience*, *20*(6), 864–871.

Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific α-bank electroencephalography increases over occipital cortex. *Journal of Neuroscience*, *20*(6), RC63–RC63.

Zelinsky, G. J., & Bisley, J. W. (2015). The what, where, and why of priority maps and their interactions with visual working memory. *Annals of the New York Academy of Sciences*, *1339*(1), 154–164.

Zhang, B., Weidner, R., Allenmark, F., Bertleff, S., Fink, G. R., Shi, Z., & Müller, H. J. (2022). Statistical learning of frequent distractor locations in visual search involves regional signal suppression in early visual cortex. *Cerebral Cortex*, *32*(13), 2729–2744.