

Pinging the Hidden Attentional Priority Map: Suppression Needs Attention
Changrun Huang^{1,2*}, Dirk van Moorselaar^{1,2}, Joshua J. Foster⁴, Mieke Donk^{1,2}, and Jan

Theeuwes^{1,2,3}

¹Department of Experimental and Applied Psychology, Vrije Universiteit Amsterdam,
Amsterdam, the Netherlands


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


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


⁴Department of Psychological and Brain Sciences, Boston University


Author notes

Changrun Huang  <https://orcid.org/0000-0002-1627-0887>

Dirk van Moorselaar  <https://orcid.org/0000-0002-0491-1317>

Joshua J. Foster  <https://orcid.org/0000-0001-7034-9636>

Mieke Donk  <https://orcid.org/0000-0001-9310-8210>

Jan Theeuwes  <https://orcid.org/0000-0002-5849-7721>

Jan Theeuwes was supported by a European Research Council (ERC) advanced grant 833029 – [LEARNATTEND] and Changrun Huang was supported by a China Scholarship Council (CSC) scholarship [201908440284]. The data and analysis materials are available upon request.

Correspondence concerning this article should be addressed to Changrun Huang, Department of Experimental and Applied Psychology, Vrije Universiteit Amsterdam, Van der Boechorststraat 7-9, 1081 BT Amsterdam, The Netherlands, Email: changrunhuang@gmail.com

Abstract

Attentional capture by an irrelevant salient distractor is attenuated when the distractor is presented more frequently in one location compared to other locations, suggesting that people learn to suppress an irrelevant salient location. However, to date it is unclear whether this suppression is proactive, applied before attention has been directed to the distractor location, or reactive, occurring after attention has been directed to that specific location. The aim of the present study is to investigate how suppression is accomplished by using the pinging technique which allows one to probe how attention is distributed across the visual field prior to the presentation of the search display. In an EEG experiment, participants performed a visual search task wherein they were tasked with identifying a shape singleton in the presence of an irrelevant color singleton. Compared to all other locations, this color singleton appeared more frequently at a specific location, which was termed the high-probability location. Prior to the search task, we introduced a continuous recall spatial memory task to reveal the hidden attentional priority map. Participants had to memorize the location of a memory cue continuously and report this location after the visual search task. Critically, after the presentation of the memory cue but before the onset of the search display, a neutral placeholder display was presented to probe how hidden priority map is reconfigured by the learned distractor suppression. Behaviorally, there was clear evidence that the high-probability location was suppressed, as search was more efficient when the distractor appeared at this location. To examine the priority map prior to search, we adopted an inverted encoding approach to reconstruct the tuning profile of the memorized position in the spatial memory task. Inverted modeling resulted in reliable tuning profiles during memory maintenance that gradually decayed and that were revived again by the onset of a neutral placeholder display preceding search. After the onset of the placeholders, the tuning profile observed was characterized by a spatial gradient centered over the high-probability location, with tuning being most pronounced at the-to-be suppressed location. This finding suggests that while learned suppression is initiated prior to search display onset, it is preceded by an initial phase of spatial selection, which is in line with a reactive suppression account. Together these results further our understanding of the mechanism of spatial distractor suppression.

Keywords: distractor suppression, priority map, statistical learning, visual selection

Pinging the Hidden Attentional Priority Map: Suppression Needs Attention

Even though large amounts of information constantly bombard our senses, we can effortlessly direct attention to relevant information and ignore information that may distract us. Recent studies have demonstrated that attentional selection can be so efficient because visual input is highly repetitive and structured, which makes it possible to predict what information will appear next based on the current sensory input (Friston, 2009; Kok et al., 2017). Extracting these regularities from the environment, often called statistical learning (Chun & Jiang, 1999; Frost et al., 2015), occurs effortlessly across trials, operates largely outside the realm of conscious awareness, and is not contingent on explicit knowledge of the regularity (Gao & Theeuwes, 2022; Goujon et al., 2015; Turk-Browne et al., 2005; but see Vicente-Conesa et al., 2023). In addition to top-down and bottom-up control processes (Corbetta & Shulman, 2002; Desimone & Duncan, 1995; Theeuwes, 2010), statistical learning plays a crucial role in attentional selection (Awh et al., 2012; Failing & Theeuwes, 2018; Theeuwes et al., 2022). According to the tripartite model of attention (Awh et al., 2012; Theeuwes, 2019; Theeuwes et al., 2022) the interaction between top-down, bottom-up, and selection history, a category which encompasses statistical learning as well as other history-based effects, jointly determine the weights in a spatial priority map, which determines, in a winner take all fashion, which object is selected (Chelazzi et al., 2019; Zelinsky & Bisley, 2015).

Previous studies have demonstrated that observers can learn which location is most likely to contain the target. Geng and Behrmann (2005) showed that targets presented in high-probability locations are detected faster than those in low-probability locations (see also Ferrante et al., 2018; Jiang et al., 2013). In a related vein, Huang et al. (2022) used the additional singleton visual search task in which the target was presented more often in one location than in all other locations. This task was interleaved with probe trials, enabling an exploration of the distribution of attention across the display in the period preceding search display onset. Crucially, the probe task showed spatial enhancement for the location with the highest likelihood of containing the target. Based on these findings, it was argued that the amplification of weights within the spatial priority map, driven by statistical learning, takes place preemptively before the actual display presentation, implying that priority was changed *proactively*, prior to the allocation of attention. Consistent with this, an innovative EEG study by Duncan et al. (2023) employed a "ping" technique, commonly used in the realm of working memory (Wolff et al., 2015, 2017), to unveil

the weights of the concealed attentional priority. This technique pushes a wave of activity, often via presentation of a high-contrast stimulus, through the visual system to reveal hidden neural representations within networks of altered synaptic weights. Duncan et al. (2023) demonstrated that, following the acquisition of a learned priority for a specific spatial location, there was reliable decoding of the high-probability target location in the evoked EEG signal when a visual ping occurred during the interval preceding the presentation of the search display. The above change decoding is assumed to reveal the prioritized (enhanced) location within the “activity-silent” priority map.

These previous studies indicate that people easily pick up and learn the statistical regularities associated with the location of the target. Recently however, a large number of studies demonstrated that not only target but also distractor-based regularities affect the attention deployment (e.g., Failing & Theeuwes, 2020; Ferrante et al., 2018; Goschy et al., 2014; Huang, Theeuwes, et al., 2021; van Moorselaar & Theeuwes, 2021, 2022; Wang & Theeuwes, 2018b, 2018a, 2018c; feature-based: Failing et al., 2019; Vatterott & Vecera, 2012). Employing a modified version of the additional singleton paradigm, where the color singleton distractor appeared with a higher probability at a specific location relative to other locations, Wang and Theeuwes demonstrated that distractor interference was reduced when distractors appeared at this high probability distractor location than at other locations. Furthermore, participants exhibited slower responses when the target was presented at this high-probability distractor location (Wang & Theeuwes, 2018b, 2018a, 2018c; also see Ferrante et al., 2018; Goschy et al., 2014), suggesting that the high probability location was suppressed. Similar to the method used for target learning, Huang and colleagues (2022, 2023) randomly interleaved probe trials to explore whether learned distractor suppression was implemented before search display onset. Across multiple studies it was found that responses were slowed when the probe appeared at the high probability distractor location (Huang et al., 2022; Huang, Vilotijević, et al., 2021; Kong et al., 2020), even when the probe display was presented prior to the expected onset of the search display (Huang, Donk and Theeuwes, 2023) leading to the conclusion that suppression operates prior to search display onset. However, others have argued that suppression can only be instantiated *reactively*, after attention has been directed to that location (Chang et al., 2023; Moher & Egeth, 2012). In this respect, it is noteworthy that in the capture probes studies (Huang et al., 2022, 2023) probes that revealed suppression were presented after presenting a neutral

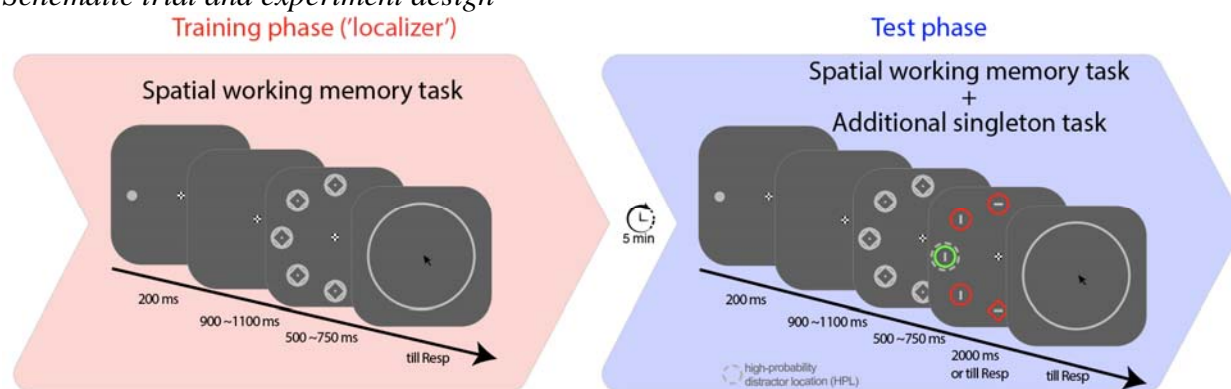
placeholder display. Therefore, it is possible that while suppression may have been in place at the moment in time the additional singleton search display was presented, the earlier presented placeholder display may have triggered an initial shift of attention to the high probability distractor location which was immediately followed by rapid attentional disengagement and suppression (Theeuwes et al., 2000; Theeuwes & Chen, 2005). Taking this concern into account, in a variant of the capture-probe paradigm, where participants discriminated the orientation of a tilted bar presented at one of the search locations, (Chang et al., 2023) revealed initial enhancement at the high probability distractor location before suppression. This finding implies that counter to the proactive suppression account, attention may be a prerequisite for suppression to occur.

The current study was designed to establish how suppression is implemented within the attentional priority map. For this purpose, analogous to previously used pinging technique (Duncan et al., 2023), search display onset was preceded by a task-irrelevant neutral placeholder display that served as a visual ping (see Figure 1). The attentional profile elicited by this visual ping was reconstructed in a time-resolved manner with a spatial inverted encoding model (Brouwer & Heeger, 2009, 2011; Foster et al., 2016, 2017; Sprague et al., 2016), generating location-selective channel tuning functions (CTFs) over time. In order to model the spatial response profile associated with learned suppression, a prerequisite is that the high-probability distractor location periodically shifts across space such that the topographic distribution of different spatial channels (i.e., neural populations) can be considered. Periodically shifting the high-probability distractor location however inevitably introduces temporal confounds (Duncan et al., 2023), such as lingering and enduring effects from initial learning experiences on subsequent attentional biases (Wang & Theeuwes, 2020). In light of these challenges, we opted to not center the tuning profiles around the high-probability distractor location. Instead, the high probability location remained static throughout the experiment. We then measured its influence on the priority landscape via a spatially specific modulation of a top-down attention signal, which can be flexibly adjusted on a trial-by-trial basis (Posner et al., 1980; Theeuwes, 2019).

To this end, we embedded the additional singleton task, which was always preceded by a placeholder display, within the maintenance period of a spatial working memory task (see Figure 1), with the spatial memory serving as a proxy for top-down attention (Awh & Jonides, 2001). Participants performed a visual search task wherein they were tasked with identifying a shape

singleton in the presence of an irrelevant color singleton. Compared to all other locations, this color singleton appeared more frequently at a specific location, which was termed the high-probability location. Prior to the search task, we introduced a continuous recall spatial memory task to reveal the hidden attentional priority map. Participants had to memorize the location of a memory cue continuously and report this location after the visual search task. Critically, after the presentation of the memory cue but before the onset of the search display, a neutral placeholder display was presented to probe how the hidden priority map is reconfigured by the learned distractor suppression. This design not only enabled us to investigate whether a ping could also unveil learned attentional biases associated with suppression but, crucially, also facilitated the examination of the tuning profile of this suppression in a time-resolved manner. This approach allows for a dissociation between proactive and reactive mechanisms. Specifically, within a proactive account memory specific tuning should be attenuated at the high probability location immediately following placeholder onset (Huang et al., 2022; Huang, Theeuwes, et al., 2021; Kong et al., 2020), whereas a reactive suppression account predicts that the to-be suppressed location is initially attended resulting in temporarily enhanced tuning at that location (Chang et al., 2023; Moher & Egeth, 2012).

Figure 1
Schematic trial and experiment design



Note. The experiment commenced with a training phase, during which participants were tasked with retaining the spatial cue's location in their memory for subsequent testing following a delay. After a 5 minutes break, the training phase transitioned to the test phase, where a search task was embedded during memory maintenance. Upon search display onset, participants were instructed to search for a unique shape singleton while ignoring the color singleton. Unbeknownst to the participants, the color singleton was presented more frequently at a location referred to as the high-probability location (HPL).

Results

Behavior: Spatial distractor learning and precise spatial memory maintenance

We conducted a pairwise t-test comparing memory recall deviations between training and test phases to examine how memory performance varied with or without a concurrent secondary search task. Our findings revealed a significantly larger recall deviation at the test phase compared to the training phase ($6.27^\circ \pm 0.95$ vs. $11.94^\circ \pm 3.55$, $t(23) = 5.67$, $p < .001$, $d = 2.52$), indicating that the process of tracking the spatial memory cue was substantially disrupted by the concurrent secondary task in the test phase. Note, that despite being disrupted by the search task, overall memory recall performance in the test phase was nevertheless high, indicating that observers were able to maintain a relatively precise memory representation outside the current focus of attention.

Next, we examined whether participants learned to inhibit the high-probability location in the visual search task while simultaneously maintaining an online representation of a spatial location in WM. A paired t-test was performed on mean RTs and accuracy, comparing trials where the distractor appeared at the high-probability location with those in which it appeared at one of the low-probability locations¹ (see Figure 2A and 2B). The planned paired t-test revealed that participants exhibited faster ($t(23) = 10.14$, $p < .001$, $d = 0.43$) and more accurate ($t(23) = 5.19$, $p < .001$, $d = 0.71$) responses when the distractor appeared at the high-probability location compared to the low-probability locations, showing a clear effect of statistically learned distractor suppression and replicating previous findings (Ferrante et al., 2018; Wang & Theeuwes, 2018b). If the statistically learned suppression was spatial-based and feature-blind, one would also expect impaired target processing at the high-probability location. Consistent with this expectation, a planned paired t-test on RTs for target location (collapsed across distractor-absent and distractor-present trials) demonstrated a substantial delay in participants' responses when the target was presented at the high-probability location ($M = 1018$ ms), as opposed to the low-probability locations ($M = 958$ ms, $t(23) = 5.45$, $p < .001$, $d = 0.48$). The

¹ When contrasting RTs between trials featuring distractors at the high-probability location and those at the low-probability location, potential RT slowdown in the latter scenario could arise from the target occupying the high-probability location, a condition known to yield impaired target processing. To disentangle the effects of distractor suppression and target impairment and mitigate confounding, we categorized trials as low-probability only when both the target and distractor were not presented at the high-probability location.

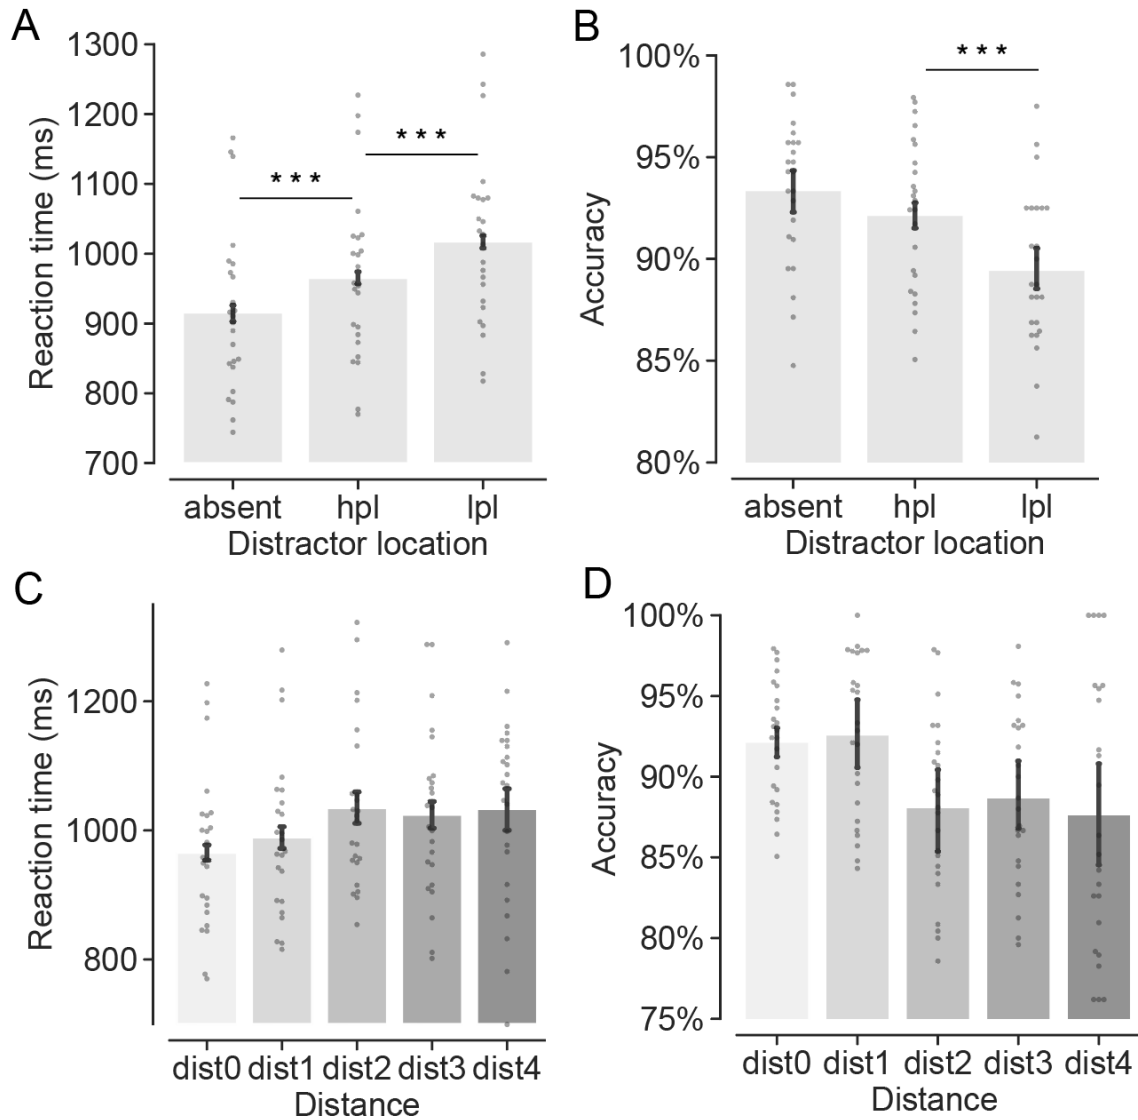
same paired t-test on the mean accuracy (M_{HPL} vs. M_{LPL} : 90.1% vs. 91.7%) did not reveal any difference ($t(23) = 1.68$, $p = .106$, $d = 0.36$, $\text{BF}_{01} = 1.37$).

To further characterize whether the learned suppression scaled with the relative distance to the high-probability location, we assigned the low-probability trials into one of four distance groups depending on the relative distance between the distractor and the high-probability location. The mean RTs and mean accuracy were submitted to a repeated-measures ANOVA with a within-subject factor Distance (dist-1, dist-2, dist-3, dist-4). A significant main effect was found in the analysis of RTs ($F(2.22, 51.06) = 3.87$, $p = .023$, $\eta_p^2 = .14$). As shown in Figure 2C, responses in the dist-1 condition were faster than those in the dist-2 condition ($t(23) = 3.96$, $p = .003$, $d = 0.36$), and marginally faster than those in the dist-3 condition ($t(23) = 2.83$, $p = .058$, $d = 0.28$) and dist-4 condition ($t(23) = 2.75$, $p = .069$, $d = 0.33$). All other comparisons did not reach significance (all $ps = 1$). Mean accuracy analyses mimicked these findings (see Figure 2D), showing a main effect of Distance ($F(2.13, 49) = 4.20$, $p = .019$, $\eta_p^2 = .15$). Performance in the dist-1 condition was more accurate than in the dist-2 condition ($t(23) = 4.24$, $p = .002$, $d = 0.86$), and in the dist-3 condition ($t(23) = 2.97$, $p = .041$, $d = 0.75$), and marginally more accurate than in the dist-4 condition ($t(23) = 2.87$, $p = .052$, $d = 0.73$).

In summary, observers remained sensitive to the spatial distractor regularity when the visual search task was embedded within the maintenance period of a spatial memory task. Consistent with previous work, distractor learning was characterized by a spatial gradient centered at the high probability distractor location, where distractor interference was attenuated, and target processing was impaired. After having validated that processing at the high probability distractor location was suppressed, we set out next to establish whether this suppression could be revealed in anticipation of search display onset, and critically, whether it did attenuate or enhance memory-specific tuning at the high probability distractor location relative to the other display locations.

Figure 2

Behavioral findings in the search task



Note. The top panels show (A) the mean RTs and (B) the mean accuracy under conditions where the distractor was absent, presented at either the high-probability location (hpl) or at the low-probability locations (lpl). The bottom panels show (C) the mean RTs and (D) the mean accuracy in relation to the relative distance between the distractor and the high-probability location. Specifically, dist0, dist1, dist2, dist3, and dist4 signify instances where the distractor was at the hpl, one position, two positions, three positions, and four positions away from the hpl. Small grey dots show data for individual participants. The presence of '***' denotes statistical significance at the level of $p < .001$. Error bars indicate condition-specific, within-subject 95% confidence intervals (Morey, 2008).

Localizer data: alpha-band tuned ping of spatial working memory representations

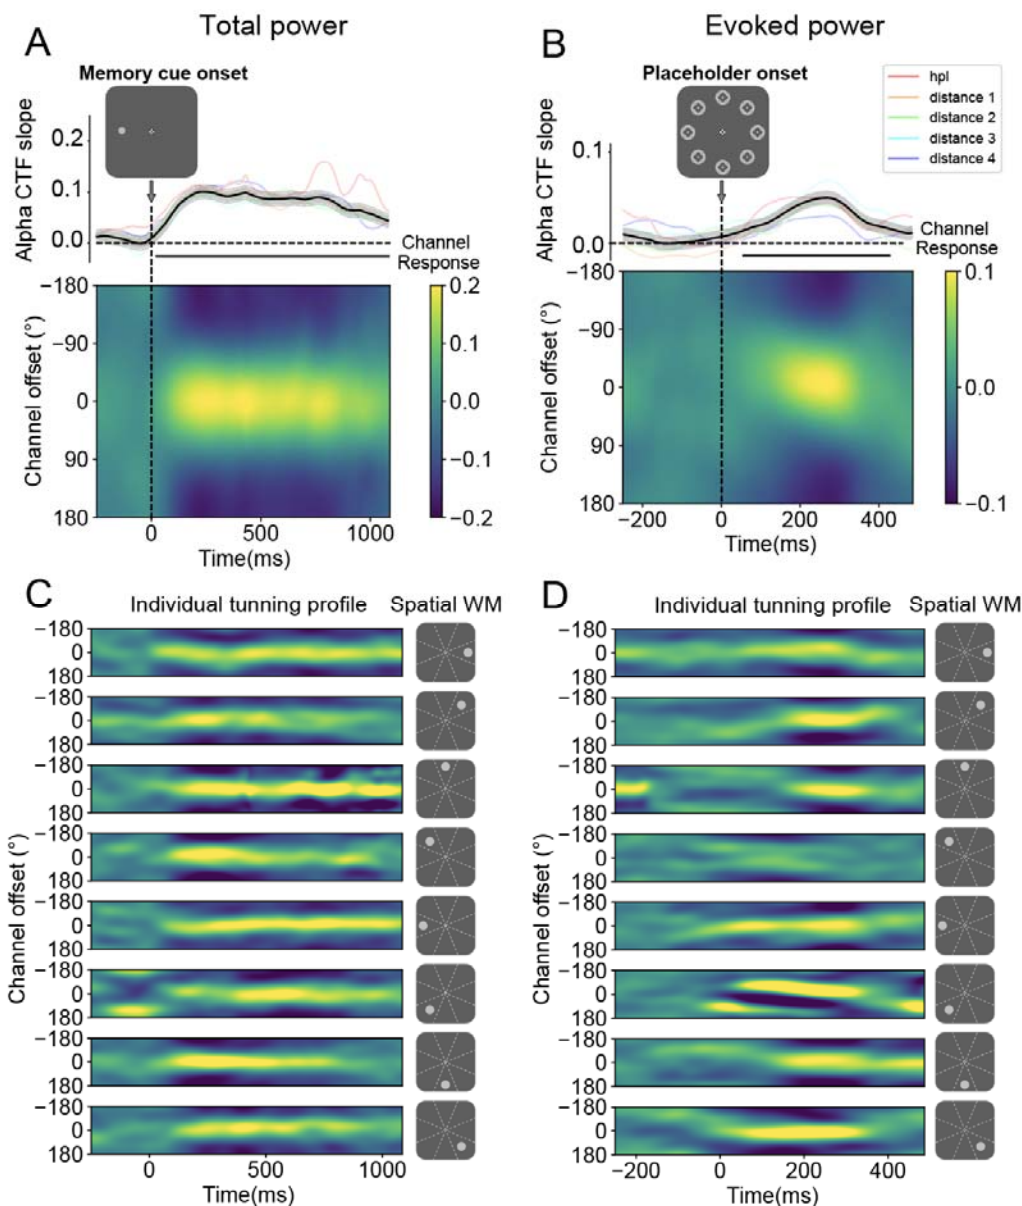
Before investigating whether attentional tuning profiles elicited by the placeholder displays would be modulated by distractor-based learning, we first characterized the spatial memory code in the training phase (i.e., in the absence of selection history effects). Given the robust association with spatial attention and working memory maintenance, these analyses focused on oscillatory power within the alpha-band (8 - 12 Hz). Consistent with previous work (Foster et al.,

2016, 2017; van Moorselaar et al., 2018), both evoked and total alpha power generated reliable CTFs in response to memory display onset ($p < .05$), with evoked power returning to baseline relatively early in the delayed period (see Supplementary Figure S1A), whereas total alpha power enabled reliable CTF reconstruction throughout the entire delay period (Figure 3A). Notably, as depicted in Figure 3B, the evoked CTFs were distinctly revived upon the presentation of the placeholder display ($p < .05$).

Given that the primary objective of this study was to monitor location-specific modulations of the attention set as a function of the distance from the high probability location, we also ensured that the observed reconstructions were not driven by a subset of locations, but instead were homogeneous across all possible locations. As visualized in Figure 3C/D, the generated CTFs were not the result of imbalanced spatial selectivity of specific locations, as similar tuning was observed across all eight individual locations. Indeed, an artificially created distance-based gradient centered around the location that would become the high probability location in the subsequent test phase did not exhibit any unbalanced spatial selectivity (see colored coded lines in Figure 3A and 3B). These analyses confirm that the data from the training set did not exhibit a spatial bias from the outset and can thus serve as a neutral independent data set to investigate how distractor learning would, if at all, modulate attentional tuning across all display locations.

Figure 3

The performance of the Localizer in the training phase



Note. (A) The average total alpha-band channel tuning function (CTF) profiles time-locked to memory onset (B) The average evoked alpha CTF profiles time-locked to placeholder onset. The lower image depicts the responses across channels, while the plot above shows CTF slopes, with amplitude signifying spatial selectivity. Shaded areas reflect bootstrapped SEM. Time points exhibiting significant differences in CTF slopes, identified through a cluster-based permutation test ($p < .05$), are marked with horizontal black insets. The light color lines in the background indicate the CTF slopes tuned to different memory cue locations, grouped by their distance to the artificial high-probability location. (C) Individual total CTF profiles, synchronized with memory cue onset, finely tuned to each of the eight spatial cue locations, respectively. (D) Individual evoked CTF profiles, synchronized with placeholder onset, finely tuned to each of the eight spatial cue locations, respectively.

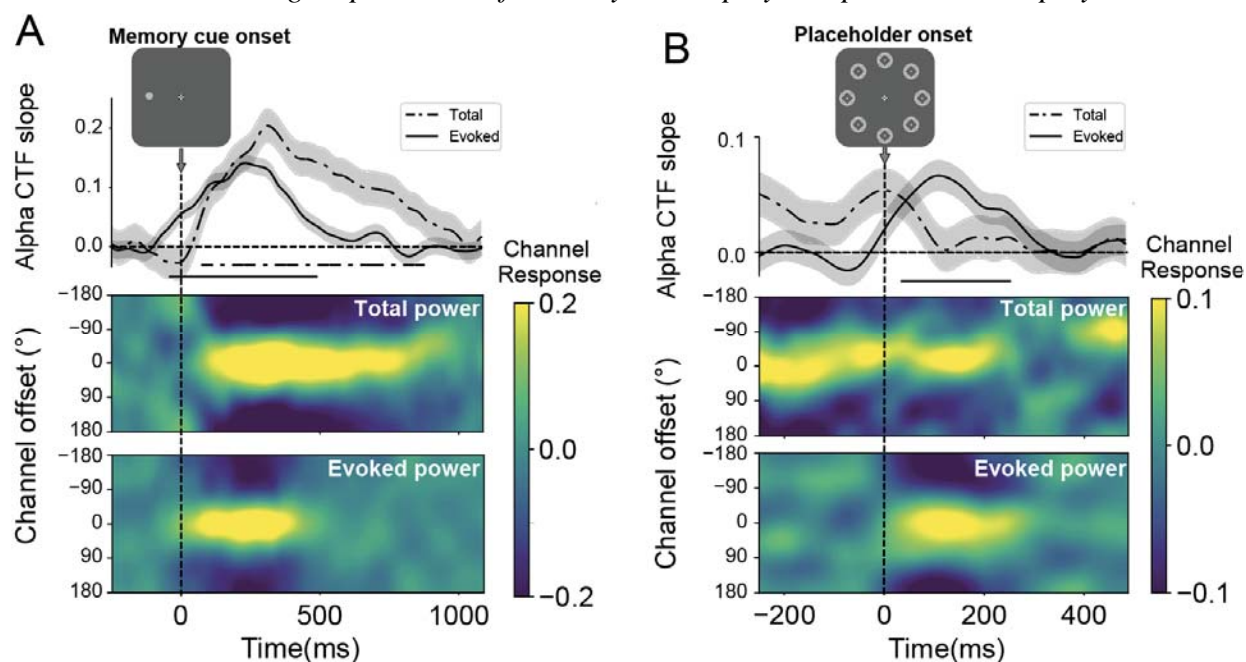
Cross-session encoding: No evidence for proactive suppression of top-down attentional selection biases

After having established that weights obtained from the training set reliably captured working memory maintenance across all eight locations without a spatial bias, we next examined how this

model generalized to the test set. In doing so, we ensured that data points in the training and test set were aligned such that training and testing were done on matching time samples. As depicted in Figure 4, and in line with van Moorselaar et al. (2018) the averaged tuning profiles resulting from this cross-session encoding procedure mimicked the pattern of the localizer, with the main difference that total power no longer tracked the memorized position throughout the entire delay interval, but instead already returned to baseline prior to placeholder onset (Figure 4A). Critically, however, this information was revived again by the placeholder, an effect that was selective to evoked power (Figure 4B). Together these findings are consistent with a model wherein the prospect of a secondary task during the memory maintenance interval shifted the memory from an activity-driven into an activity-silent representation, which was then revived again by placeholder onset (analogous to the pinging procedure). Despite a numerical trend, this pattern did not manifest in the total alpha power CTF (see Figure 4B), indicating that the revived spatial selectivity may be attributed to the phase-locked information carried by the placeholder display.

Figure 4

Cross-session encoding: Alpha CTFs of Memory cue display and placeholder display

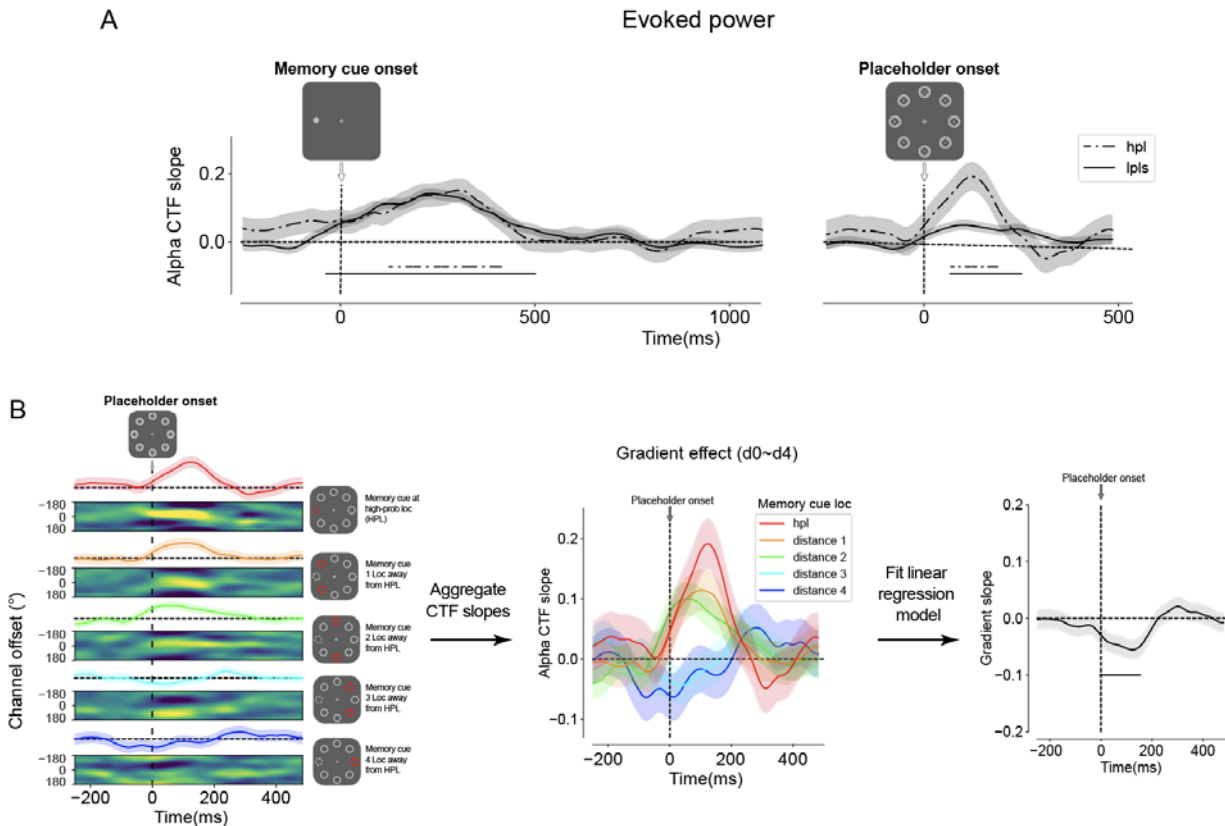


Note. The location CTFs obtained by generalizing the trained model to the test set. CTFs were reconstructed from evoked and total alpha power following the onset of (A) the memory cue display and (B) the placeholder display. The lower image depicts the responses across channels, while the plot above shows CTF slopes, with amplitude signifying spatial selectivity. Shaded areas reflect bootstrapped SEM. Time points exhibiting significant differences in CTF slopes, identified through a cluster-based permutation test ($p < .05$), are marked with horizontal black insets.

To further investigate how this revival of the memory representation was modulated by the distractor regularity, we analyzed the spatial selectivity indexed by the reconstructed CTF profiles as a function of the memory cue location (HPL-matched, LPLs-matched). As visualized in Figure 5A, although reliable CTF reconstruction was observed in both conditions ($p < 0.05$), the HPL-matched CTFs were numerically more pronounced than the LPLs-matched CTFs. Although statistically this conditional difference was not reliable, the observed pattern seems inconsistent with the notion that distractor learning leads to the proactive suppression of the high probability distractor location, which should arguably have resulted in an attenuation of the memory-specific revival at that location. To further explore this finding, we leveraged the observation that, as was also the case here, suppression often exhibits a spatial gradient. Specifically, we organized the evoked CTFs based on the distance from the memory location to the high-probability distractor location. As shown in Figure 5B, in line with a spatial gradient, the resulting CTF profiles varied depending on their distances to the high-probability distractor location. To evaluate the reliability of this observed gradient pattern in spatial selectivity, we employed a linear regression model to fit the tuning profiles. A cluster-based permutation test conducted on the slopes of the regression model identified a reliable gradient ($p < .05$) centered around the high probability location following placeholder onset. Critically, this gradient reflected an increase of spatial selectivity centered at the high probability distractor location rather than an attenuation, suggesting that immediately following placeholder onset the to-be-suppressed location was at least transiently attended before being suppressed during visual search in the subsequent display.

Figure 5

The gradient pattern of Alpha CTFs following the onset of the placeholder display



Note. (A) Numerically larger CTFs were observed post-placeholder onset when the spatial memory cue matched the high-probability location compared to the low-probability locations. (B) The location-specific CTFs were categorized based on the relative position distance between the spatial memory cue and the high-probability location. To statistically evaluate the observed gradient pattern, a linear regression model was applied to the data points, followed by a permutation test on the slope of the regression model. Shaded areas depict bootstrapped SEM, and time points with significant differences in regression model slopes are indicated with horizontal black insets.

Discussion

The aim of the current study was to examine whether learned distractor suppression is in place before search display onset (i.e., proactive suppression) or instead whether suppression is instantiated following attentional selection (i.e., reactive suppression). To this end, we introduced a variant of the additional singleton task with a high probability distractor location in the maintenance period of a spatial working memory task. At the behavioral level, responses were faster when the distractor occurred at the high-probability location, but slower when the target occurred at this location, indicative of generic spatial suppression, an effect that was characterized by a discernible spatial gradient centered at the high-probability distractor location. Yet, critically, at the neural level, there was no evidence that this spatial suppression attenuated processing at the high probability distractor location in anticipation of a new search episode. The neutral placeholder display, serving as a ping, presented prior to search display onset revived the

reconstruction of the CTF profile of the current memory representation. Instead of a model that assumes that the high probability location was proactively suppressed, the revival of the CTF profile within the alpha-band showed the largest tuning at the learned suppressed location. Moreover, the modulation of the memory-tuned CTF profile was characterized by a spatial gradient that was centered over the high-probability location. This finding offers evidence that suppression comes into play only after the location is initially selected. Suppression occurs only following attentional enhancement, indicating a reactive suppression mechanism rather than proactive suppression.

Within the realm of learned distractor suppression, an ongoing debate centers around the question of whether, and precisely when, visual distractors can be proactively suppressed. As noted, the idea that learned spatial distractor suppression is applied proactively is largely based on the finding that the behavioral benefit observed when distractors appear with a higher probability at a given location is accompanied by a probe detection cost (measured via dot offset detection) at the high probability distractor location (Huang et al., 2022, 2023; Huang, Vilotijević, et al., 2021). As also pointed out by Chang et al. (2023) however, this assumption is largely grounded in findings that rely on behavioral measures, leaving open the possibility that the suppression of a distractor location only occurs after it was initially attended. To bypass the ambiguity in behavioral measures, other studies have turned to measures of the brain in the window prior to search display onset. Studies investigating the preparatory bias in response to distractor regularities largely examined topographical modulations of alpha-power, given its functional link to reduced cortical excitability. These studies however, mainly failed to find evidence in support of active preparatory inhibition (van Moorselaar et al., 2020, 2021; van Moorselaar & Slagter, 2019), with only one study observing increased preparatory alpha contralateral to the high probability distractor location (Wang et al., 2019). This absence of anticipatory tuning towards the to-be suppressed location has given rise to the idea that distractor learning operates by changing synaptic efficiency within these regions (van Moorselaar & Slagter, 2020). Consistent with this, a recent study using rapid invisible frequency tagging demonstrated reduced neural excitability at the high-probability distractor location in the absence of any alpha-band modulations (Ferrante et al., 2023). However, the correlational approach in that study, prevented a time-resolved analysis, leaving uncertainties about whether suppression was genuinely proactive or triggered by placeholder onset. Thus, while it is generally assumed

that learned suppression is proactive in nature, the evidence in support of that notion remains equivocal.

The observed modulation of the revival of the CTF profile casts further doubt on the assumption that learned distractor suppression is implemented proactively. While we hypothesized that top-down tuning (induced by holding a location in memory) would be least pronounced at the high-probability distractor location, we observed the opposite pattern. Rather than being suppressed the revived CTFs were characterized by a spatial gradient centered over the high-probability location, with tuning being most pronounced at the-to-be suppressed location. In contrast, in the training phase, tuning was homogeneous across all eight search locations. This finding clearly indicates that the putative priority map, initially tuned by maintaining a spatial location in WM, was reconfigured by statistical regularities across search displays to align with the imminent search. Thus, while learning-dependent suppression may have been initiated prior to the onset of the search display, this suppression appears to be preceded by an initial phase of spatial selection. Although this result is based on an exploratory analysis and should therefore be interpreted with caution, it is noteworthy that a recent behavioral study observed a similar pattern of results (Chang et al., 2023). In a modified version of the capture-probe paradigm, where the probe display onset was not preceded by a placeholder display, participants discriminated the orientation of a tilted bar presented at one of the search locations, revealing initial enhancement at the high probability distractor location before suppression. Although the exploratory nature of our findings should again be stressed, they may call for a reinterpretation of how learned suppression might take place. Instead of proactively suppressing specific locations, individuals may first direct attention to the location that they have implicitly learned to expect a distractor, enabling suppression of that location following rapid attentional disengagement.

A novel aspect of the current study is that visual pings, here in the shape of a neutral placeholder display, can effectively unveil spatial memories hidden from the ongoing EEG signal. Recently, we demonstrated that the otherwise, invisible attentional priority map induced by a spatial imbalance of target probability across trials could also be revealed by the ‘pinging’ technique in conjunction with multivariate pattern analysis and EEG (Duncan et al., 2023). However, while it is well established that early ERP components are enhanced in response to visual stimulation at memorized locations (Awh et al., 2000), to date the pinging technique has

exclusively been used to reveal feature-specific information within working memory (Wolff et al., 2015, 2017). Despite initial enthusiasm within this field, there is an ongoing debate surrounding the precise mechanisms underlying the ping effect, particularly whether it reactivates latent networks or merely amplifies existing, yet below-threshold representations within ongoing neural activity (Barbosa et al., 2021). Additionally, skepticism arises from demonstrations that measurable neural activity often underlies working memory maintenance (Schneegans & Bays, 2017). These two, not necessarily exclusive scenarios, were also evident in the present study. During the training phase, where the memory task was the only task at hand, CTFs continuously tracked the position of the spatial memory cue. However, the same reconstruction returned to baseline when the model was applied to the test session that incorporated a search task during the maintenance interval. This dissociation aligns intriguingly with a study by van Moorselaar et al. (2018), where continuous reconstruction of spatial memory was disrupted during a secondary task introduced in the maintenance interval. This suggests the possibility that under dual-task conditions memories might be strategically offloaded to a mechanism relying on weak neural activity or even one not dependent on sustained neural firing at all, such as an activity-silent representation facilitated by synaptic plasticity or long-term memory. Regardless of the precise underlying mechanism, our results demonstrate in a compelling way that the content of a hidden spatial representation can be revived by flooding the visual system with sensory input. This demonstrates that the pinging technique can not only be used to investigate feature-based working memory but also the dynamics of spatial memories.

In summary, the present study is the first to show that a spatial memory representation can be reconstructed based on a ping, a neutral placeholder display, and used to infer how distractor suppression affects the priority map prior to search. Prior to search and in response to the ping, the tuning profile of the memorized location was most pronounced at the high probability-distractor location and exhibited a spatial gradient centered over that location. These results are not in line with a proactive suppression account but support the idea that learned suppression follows an initial phase of spatial selection.

Methods

Participants

In line with previous EEG experiments (Duncan et al., 2023; Noonan et al., 2016; van Moorselaar et al., 2020, 2021, 2023), we used a predetermined sample size of 24 participants. Participants ($N=24$, 22 females, 2 males, $M_{age} = 20.8$, $SD_{age} = 2.9$) were recruited from the research pool of the Vrije Universiteit Amsterdam. Participants received either course credits or a monetary reward (€35) for 3.5 hours of participation. To ensure data quality and maintain the predefined sample size, three participants were replaced. One participant was substituted due to a high number of trial removals during preprocessing ($>30\%$). Another participant, whose accuracy fell below 2.5 standard deviations from the overall mean for the search task, was also replaced. The third participant was substituted because of poor memory recall, deviating more than 2.5 standard deviations from the overall mean. Written informed consent was obtained from all participants before the experiment. The study was approved by the Ethical Review Committee of the Faculty of Behavioral and Movement Sciences of Vrije Universiteit Amsterdam and was conducted following the guidelines of the Helsinki Declaration.

Apparatus, Task and Stimuli

The experiment was created on a Dell Precision 3640 Windows 10 computer equipped with an NVIDIA Quadro P620 graphics card in OpenSesame (Mathôt et al., 2012) using PsychoPy functionality (Peirce, 2007). Stimuli were presented on a 23.8-in. ASUS XG248Q monitor with a 240-Hz refresh rate. Participants were seated in a dimly lit room, positioned at a distance of 70 cm from the monitor, with the aid of a chinrest to ensure stable head position. Throughout the experiment, participants' right eye movements were tracked by Eyelink 1000 (SR Research) eye tracker at a sample rate of 1000 Hz. Participants were given specific instructions to maintain their fixation on the central point during the experiment. Auditory feedback was provided whenever the gaze position deviated by more than 1.5° from the central point. All visual stimuli were presented against a light grey background with RGB values of 224/224/224. Henceforth, the colors of the stimuli were consistently delineated in RGB values.

The spatial working memory task was adapted from Foster et al. (2016) and required participants to memorize the angular location of a circle stimulus (0.9° in radius, 196/196/196) positioned 4.2° away from the central fixation marker (0.35° in radius). This fixation marker was designed as a combination of a bull's eye and crosshair, a feature known to improve stable fixation (Thaler et al., 2013). The angular location was sampled with equal probability from one

of eight locations bins spanning 0° - 315° , with jitter added (-22.5° - 22.5°) to cover all possible locations to avoid categorical coding. At test, participants were required to report the memorized location via a mouse by clicking on a ring (4.2° in radius).

The intermediate visual search task was modeled after the additional singleton paradigm (Theeuwes, 1992). Participants were instructed to search for a unique shape (i.e. target) within a further homogeneous display and report the orientation of the line within the target shape. The search array consisted of eight evenly spaced items in a circular configuration (4.2° in radius) around central fixation. Each item within the search array contained either a vertical or horizontal white bar ($0.1^\circ \times 1^\circ$, 255/255/255). The target could either be a diamond ($2.3^\circ \times 2.3^\circ$) among seven circles (2° in diameter) or vice versa. In 74% of the trials, a distractor was present: one of the non-target items was a color singleton (either red (255/0/0) or green (0/131/0)). In these trials, the distractor was more likely (64%) to be placed at one location (64%), called the *high-probability location* (HPL), than at the other locations (5% at each of these seven locations), to induce statistically learned suppression of the HPL. The HPL remained constant throughout the experiment and was counterbalanced across participants. The target color, shape, and line orientation within the target were randomized on each trial, and the target location was equally likely across all locations.

The placeholder display, which served as visual pings, consisted of eight grey shapes (128/128/128), each created by superimposing a diamond shape onto a circle shape. Importantly, the spatial locations that were occupied by the placeholder and the search array were spatially matched with the eight position bins in the spatial working memory task.

Design and procedure

The current study followed a structured protocol comprising two phases, a training phase, which served as an independent localizer to train the encoding model (details below), and a subsequent test phase combining the spatial working memory task with the visual search task. During the training phase, each trial began with a blank display for a randomly jittered duration of 200 – 400 ms, followed by a fixation dot for 250 ms. Participants were explicitly instructed to maintain their gaze on the fixation dot as long as it was visible. Subsequently, a memory cue was presented on the screen, and participants were instructed to remember its location until the end of the trial. The memory cue disappeared after 200 ms, leaving only the fixation dot visible for a

duration that ranged randomly between 900 to 1100 ms. Next, a placeholder display appeared for a randomly varying duration between 500 and 750 ms. Finally, a memory test display was presented until a response was made by participants.

In the test phase, the trial structure closely resembled that of the training phase, with two important distinctions: First, following the placeholder display, a search display was presented for a maximum duration of 2000 ms or until a response was made. Second, on a small subset of trials (13%), the spatial memory cue display was replaced by a display containing a yellow fixation, signaling participants that they only needed to perform the visual search task².

The training phase consisted of 10 blocks, each containing 80 trials, while the test phase comprised 10 blocks of 92 trials. Between blocks, participants were given the opportunity of a short break, during which feedback on the mean memory recall deviation and search performance (in the test phase) for both the latest block and all finished blocks was provided. After the completion of the training phase, participants were given a 5-minute break before proceeding to the test phase. Before entering the training phase, participants performed 16 practice trials randomly drawn from a training block. Likewise, preceding the test phase, participants underwent 48 practice trials randomly drawn from a testing block. This process continued until participants demonstrated a thorough understanding of the assigned tasks. It is important to emphasize that during the test phase, the location of the memory cue was not contingent on the location of the singleton in the search task (i.e., these two tasks were independent of each other).

Behavioral analysis

All data were preprocessed using Python, and statistical analysis was done using R (R Core Team, 2020). Conditional mean reaction times (RTs) and accuracy of the visual search task were analyzed with repeated measures ANOVA, followed by planned comparisons with paired t-tests. For RT analyses, we excluded incorrect responses and RTs < 200 ms. For each participant, we also excluded RTs that exceeded the ± 2.5 standard deviation from the overall mean RT (collapsed across conditions). The exclusion of incorrect responses and data trimming resulted in an average loss of 10.5% of trials. Participants whose accuracy or memory deviation fell above

² The goal of these no-memory trials was to track learned attentional priority in the absence of a top-down attentional set. The resulting model, however, failed to capture any reliable spatial tuning relative to the high-probability location, leading us to omit reporting this particular outcome.

or below 2.5 SD from the overall mean were replaced during the data collection phase (see Participants). P-values were Greenhouse-Geiser corrected in case where the assumption of sphericity was violated, and were corrected with the Holm-Bonferroni method for multiple-level comparisons. In cases of non-significant findings, we also provided the Bayes factor (BF_{01}) to support the null model.

EEG recording and preprocessing

EEG data were collected using a 64-electrode cap with electrodes placed according to the 10-10 system (Biosemi ActiveTwo system; www.biosemi.com). To monitor eye movements, in case of missing eye tracker data, vertical and horizontal electrooculogram (VEOG/HEOG) signals were recorded via external electrodes placed ~2 cm above and below the right eye, and ~1 cm lateral to the external canthi. Two additional electrodes were placed on the left and right earlobe for offline reference. Electrode impedances were kept below 20 k Ω . Signals were amplified (100 Hz low-pass filter, 0.16 Hz high-pass filter; ActiveTwo AD-box, ActiveView;) and sampled at 512 Hz.

EEG data were preprocessed using a customized Python script and the MNE package (Gramfort et al., 2013). During preprocessing, the data were re-referenced to the average of the left and the right earlobe and high-pass filtered using a zero-phase ‘firwin’ filter at 0.01 Hz to remove slow drifts. Malfunctioning electrodes detected during recording were temporarily removed in offline analysis such that subsequent preprocessing steps were not influenced by these electrodes. The continuous data were then epoched from 750 ms to 1600 ms relative to the memory cue onset and from 750 ms to 1000 ms relative to the placeholder display onset, with the windows of interest being -250 ms to 1100 ms and -250 ms to 500 ms respectively (centered around memory cue and placeholder display onset). Eye-blink components were removed after performing an independent components analysis as implemented in MNE (method = “picard”) on 1 Hz filtered epochs. We used an automatic trial-rejection procedure on the EEG signal to remove noise-contaminated epochs. Specifically, the EEG signal was further processed by applying a 110 to 140 Hz band-pass filter to capture muscle activity and transform it into z scores. A subject-specific z-score threshold was determined based on the within-subject variance of z-scores within the windows of interest (de Vries et al., 2017; Duncan et al., 2023). To minimize false alarms, instead of immediately removing epochs that exceeded the z-score

threshold, an iterative procedure was employed. For each marked epoch, the five electrodes contributing most to the accumulated z score within the time period containing the marked artifact were identified. These electrodes were then interpolated one by one using spherical splines (Perrin et al., 1989). After each interpolation, the epoch was checked to see if it still exceeded the z-score threshold. Epochs that still exceeded the threshold after the iterative interpolation were dropped which led to an average loss of 9.1% of all trials (range 0.1% - 15.3%) for epochs time-locked to memory cue onset and an average loss of 6.1% of all trials (range 0% - 12.3%) for epochs time-locked to placeholder onset. Lastly, malfunctioning electrodes were interpolated using spherical splines (Perrin et al., 1989).

Samples of the eye positions were aligned with the EEG data during offline analysis and converted to visual degrees that deviated from the central fixation point. To prevent potential confounds in interpreting the results, epochs were excluded if the gaze position exceeded 1.2° of the central fixation point anytime during the time range of -100 to 500 ms relative to the memory cue onset or during the time range of -200 to 300 ms relative to the placeholder onset. In case of missing Eyelink data, epochs were removed when detecting a sudden increase in HEOG amplitude via an algorithm with a window size of 200 ms, a step of 10 ms, and a threshold of 15 μV . In total, 6.8% of the cleaned data for memory epochs (range 0.4% - 18.7%) and 3.2% of the cleaned data for placeholder epochs (range 0% - 8.2%) were excluded due to the detection of eye movements.

Time-frequency analysis

After performing data cleaning procedures, the artifact-free EEG signals were filtered with a fifth-order Butterworth bandpass filter (8-13 Hz) implemented within MNE. The subsequent step involved computing the alpha-band power via a Hilbert transform. To distinguish between evoked and total power, distinct methodologies were applied. Evoked power was derived by averaging the complex analytic signals obtained from the Hilbert transform across trials, followed by squaring the resulting complex magnitude. Conversely, total power was computed by squaring the complex magnitude for each trial individually and then averaging these values. This distinction in procedures results in evoked power capturing phase-locked information aligned with stimulus onset, whereas total power reflects non-phase-locked information.

Following time-frequency decomposition, the resulting time courses were smoothed with a sliding window approach (window size = 8, step =1). Subsequently, a principal component analysis (n_components = 16) was performed on the smoothed data of all electrodes to enhance the signal-to-noise ratio.

Inverted Encoding Model

We used an inverted encoding model (IEM) to reconstruct location-selective channel tuning functions (CTFs) from alpha-band power (Foster et al., 2016). The CTFs were calculated using the alpha-band topographical power for its functional role of tracking spatial attention (Foster et al., 2016, 2017; Rihs et al., 2007; Samaha et al., 2016; van Moorselaar et al., 2018).

The topographical power measured at each electrode was assumed to reflect the weighted sum over 8 spatial channels (i.e., neuronal populations), each tuned to a different angular location (Foster et al., 2016). The response profile of each spatial channel across angular locations was modeled as the below function:

$$R = \sin(0.5\theta)^7$$

, where θ represents angular locations ranging from 0° to 359° and R is the response of the spatial channel. This response profile was shifted circularly for each channel such that the peak response of each spatial channel revolved around one of the polar angles (0° , 45° , 90° , etc.).

An IEM routine was applied to each data point in the alpha-band power in two stages. In the first stage, the response profile of the spatial channel and the data acquired during the experimental training phase were used to calculate a weight matrix by fitting them into a general linear model represented by the formula below:

$$B_1 = WC_1$$

, where B_1 (m electrodes \times n trials) is the observed power (evoked or total) at each electrode for each trial in the training set, C_1 (k channels \times n trials) is a matrix of predicted responses for each spatial channel on each trial, and W is a weight matrix that characterizes the mapping from “channel space” to “electrode” space. The weight matrix W is solved using the least square solution with the Python function `np.linalg.lstsq(C1, B1)`. In the second stage, We derived a set of estimated channel responses from inverting the model with the input of the obtained weight matrix W and the independent observed data B_2 (m electrodes \times n_2 trials) using the Python

function `np.linalg.lstsq(W.T, B2.T)`. The estimated channel responses (8 channels \times 8 location bins) were shifted to a common center (0°).

The IEM procedures underwent 100 iterations to mitigate the impact of noisy or erroneous data points, independently for both the localizer and cross-phase encoding. In each iteration for the localizer, the training phase data was randomly partitioned into three sets, ensuring an equitable distribution of observations across location bins (i.e., memory locations). Two sets were designated as the observed data B_1 to derive the weight matrix, while the remaining set served as the observed data B_2 . For cross-phase encoding, the training phase data was randomly split into two sets, each serving as observed data B_1 to compute the weight matrix. The test phase data was utilized as observed data B_2 . The random equating of observations across location bins occurred independently for both the training and test sets. The resulting estimated channel responses were averaged across all iterations to derive the final CTFs. This procedure ensured the independence of training and test data throughout the analysis.

Statistical Analysis

To quantify the spatial selectivity of the CTFs, we employed a linear polynomial fitting approach to estimate the slopes of the CTFs. Positive slopes would signify spatial selectivity, while negative slopes would indicate deviation from spatial selection (i.e., suppression). These slope values were then subjected to statistical analysis to assess the validity of the reconstructed CTFs and conditional differences, using a cluster-based one-sample pair t-test combined with Monte Carlo randomization. This non-parametric statistical testing considers the data's temporal correlation structure and controls for multiple comparison problems (Maris & Oostenveld, 2007). A new dataset was randomly resampled from the observed data in each iteration. The sign of the resample data was randomly flipped, and clusters were identified if the t-values of the adjacent data points exceeded the threshold. The cluster with the largest sum of t-values was retained. This process was repeated 1024 times, yielding a permutation distribution from the retained clusters. The clusters obtained from the veridical data were then compared with the form distribution. The slope or the conditional difference was deemed statistically reliable if the cluster exceeded the 95th percentile of the distribution.

Acknowledgments

SUPPRESSION NEEDS ATTENTION

25

We would like to thank Zhenzhen Xu and Ya Gao for their invaluable assistance in data collection. We also appreciate Johannes Fahrenfort for his insightful discussion on data analysis.

References

- Awh, E., Anllo-Vento, L., & Hillyard, S. A. (2000). The Role of Spatial Selective Attention in Working Memory for Locations: Evidence from Event-Related Potentials. *Journal of Cognitive Neuroscience*, *12*(5), 840–847. <https://doi.org/10.1162/089892900562444>
- 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. <https://doi.org/10.1016/j.tics.2012.06.010>
- Awh, E., & Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends in Cognitive Sciences*, *5*(3), 119–126. [https://doi.org/10.1016/S1364-6613\(00\)01593-X](https://doi.org/10.1016/S1364-6613(00)01593-X)
- 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. <https://doi.org/10.1371/journal.pbio.3001436>
- Brouwer, G. J., & Heeger, D. J. (2009). Decoding and Reconstructing Color from Responses in Human Visual Cortex. *The Journal of Neuroscience*, *29*(44), 13992–14003. <https://doi.org/10.1523/JNEUROSCI.3577-09.2009>
- Brouwer, G. J., & Heeger, D. J. (2011). Cross-orientation suppression in human visual cortex. *Journal of Neurophysiology*, *106*(5), 2108–2119. <https://doi.org/10.1152/jn.00540.2011>
- Chang, S., Dube, B., Golomb, J. D., & Leber, A. B. (2023). Learned spatial suppression is not always proactive. *Journal of Experimental Psychology: Human Perception and Performance*, *49*(7), 1031–1041. <https://doi.org/10.1037/xhp0001133>
- 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. <https://doi.org/10.1016/j.copsyc.2019.02.004>
- Chun, M. M., & Jiang, Y. (1999). Top-down attentional guidance based on implicit learning of visual covariation. *Psychological Science*, *10*(4), 360–365. <https://doi.org/10.1111/1467-9280.00168>
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*(3), 201–215. <https://doi.org/10.1038/nrn755>
- de Vries, I. E. J., van Driel, J., & 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>
- Desimone, R., & Duncan, J. (1995). Neural Mechanisms of Selective Visual Attention. *Annual Review of Neuroscience*, *18*(1), 193–222. <https://doi.org/10.1146/annurev.ne.18.030195.001205>
- Duncan, D. H., van Moorselaar, D., & Theeuwes, J. (2023). Pinging the brain to reveal the hidden attentional priority map using encephalography. *Nature Communications*, *14*(1), 4749. <https://doi.org/10.1038/s41467-023-40405-8>
- Failing, M., Feldmann-Wüstefeld, T., Wang, B., Olivers, C., & Theeuwes, J. (2019). Statistical regularities induce spatial as well as feature-specific suppression. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(10), 1291–1303. <https://doi.org/10.1037/xhp0000660>
- Failing, M., & Theeuwes, J. (2018). Selection history: How reward modulates selectivity of visual attention. *Psychonomic Bulletin & Review*, *25*(2), 514–538. <https://doi.org/10.3758/s13423-017-1380-y>

- Failing, M., & Theeuwes, J. (2020). More capture, more suppression: Distractor suppression due to statistical regularities is determined by the magnitude of attentional capture. *Psychonomic Bulletin & Review*, 27(1), 86–95. <https://doi.org/10.3758/s13423-019-01672-z>
- 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. <https://doi.org/10.1016/j.cortex.2017.09.027>
- Ferrante, O., Zhigalov, A., Hickey, C., & Jensen, O. (2023). Statistical Learning of Distractor Suppression Downregulates Prestimulus Neural Excitability in Early Visual Cortex. *The Journal of Neuroscience*, 43(12), 2190–2198. <https://doi.org/10.1523/JNEUROSCI.1703-22.2022>
- Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K., & Awh, E. (2016). The topography of alpha-band activity tracks the content of spatial working memory. *Journal of Neurophysiology*, 115(1), 168–177. <https://doi.org/10.1152/jn.00860.2015>
- 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. <https://doi.org/10.1177/0956797617699167>
- Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, 13(7), 293–301. <https://doi.org/10.1016/j.tics.2009.04.005>
- Frost, R., Armstrong, B. C., Siegelman, N., & Christiansen, M. H. (2015). Domain generality versus modality specificity: the paradox of statistical learning. *Trends in Cognitive Sciences*, 19(3), 117–125. <https://doi.org/10.1016/j.tics.2014.12.010>
- Gao, Y., & Theeuwes, J. (2022). Learning to suppress a location does not depend on knowing which location. *Attention, Perception, & Psychophysics*, 84(4), 1087–1097. <https://doi.org/10.3758/s13414-021-02404-z>
- Geng, J. J., & Behrmann, M. (2005). Spatial probability as an attentional cue in visual search. *Perception & Psychophysics*, 67(7), 1252–1268. <https://doi.org/10.3758/BF03193557>
- 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. <https://doi.org/10.3389/fpsyg.2014.01195>
- Goujon, A., Didierjean, A., & Thorpe, S. (2015). Investigating implicit statistical learning mechanisms through contextual cueing. *Trends in Cognitive Sciences*, 19(9), 524–533. <https://doi.org/10.1016/j.tics.2015.07.009>
- Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., Goj, R., Jas, M., Brooks, T., Parkkonen, L., & Hämäläinen, M. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, 7(7 DEC), 267. <https://doi.org/10.3389/fnins.2013.00267>
- 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*. <https://doi.org/10.1037/xhp0001002>
- Huang, C., Donk, M., & Theeuwes, J. (2023). Attentional suppression is in place before display onset. *Attention, Perception, & Psychophysics*, 85(4), 1012–1020. <https://doi.org/10.3758/s13414-023-02704-6>
- Huang, C., Theeuwes, J., & Donk, M. (2021). Statistical learning affects the time courses of salience-driven and goal-driven selection. *Journal of Experimental Psychology: Human Perception and Performance*, 47(1), 121–133. <https://doi.org/10.1037/xhp0000781>

- Huang, C., Vilotijević, A., Theeuwes, J., & Donk, M. (2021). Proactive distractor suppression elicited by statistical regularities in visual search. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-021-01891-3>
- 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), 87–99. <https://doi.org/10.1037/a0027611>
- Kok, P., Mostert, P., & de Lange, F. P. (2017). Prior expectations induce prestimulus sensory templates. *Proceedings of the National Academy of Sciences*, 114(39), 10473–10478. <https://doi.org/10.1073/pnas.1705652114>
- Kong, S., Li, X., Wang, B., & Theeuwes, J. (2020). Proactively location-based suppression elicited by statistical learning. *PLOS ONE*, 15(6), e0233544. <https://doi.org/10.1371/journal.pone.0233544>
- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 164(1), 177–190. <https://doi.org/10.1016/j.jneumeth.2007.03.024>
- 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. <https://doi.org/10.3758/s13428-011-0168-7>
- Moher, J., & Egeth, H. E. (2012). The ignoring paradox: Cueing distractor features leads first to selection, then to inhibition of to-be-ignored items. *Attention, Perception, & Psychophysics*, 74(8), 1590–1605. <https://doi.org/10.3758/s13414-012-0358-0>
- Morey, R. D. (2008). Confidence Intervals from Normalized Data: A correction to Cousineau (2005). *Tutorials in Quantitative Methods for Psychology*, 4(2), 61–64. <https://doi.org/10.20982/tqmp.04.2.p061>
- Noonan, M. A. P., Adamian, N., Pike, A., Printzlau, F., Crittenden, B. M., & Stokes, M. G. (2016). Distinct mechanisms for distractor suppression and target facilitation. *Journal of Neuroscience*, 36(6), 1797–1807. <https://doi.org/10.1523/JNEUROSCI.2133-15.2016>
- Peirce, J. W. (2007). PsychoPy—Psychophysics software in Python. *Journal of Neuroscience Methods*, 162(1–2), 8–13. <https://doi.org/10.1016/j.jneumeth.2006.11.017>
- 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](https://doi.org/10.1016/0013-4694(89)90180-6)
- Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, 109(2), 160–174. <https://doi.org/10.1037/0096-3445.109.2.160>
- R Core Team. (2020). *A language and environment of statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.r-project.org/>
- Rihs, T. A., Michel, C. M., & Thut, G. (2007). Mechanisms of selective inhibition in visual spatial attention are indexed by α -band EEG synchronization. *European Journal of Neuroscience*, 25(2), 603–610. <https://doi.org/10.1111/j.1460-9568.2007.05278.x>
- 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. https://doi.org/10.1162/jocn_a_00955
- 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. https://doi.org/10.1162/jocn_a_01180

- Sprague, T. C., Ester, E. F., & Serences, J. T. (2016). Restoring Latent Visual Working Memory Representations in Human Cortex. *Neuron*, *91*(3), 694–707. <https://doi.org/10.1016/j.neuron.2016.07.006>
- 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. <https://doi.org/10.1016/j.visres.2012.10.012>
- Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics*, *51*(6), 599–606. <https://doi.org/10.3758/BF03211656>
- Theeuwes, J. (2010). Top-down and bottom-up control of visual selection. *Acta Psychologica*, *135*(2), 77–99. <https://doi.org/10.1016/j.actpsy.2010.02.006>
- Theeuwes, J. (2019). Goal-driven, stimulus-driven, and history-driven selection. *Current Opinion in Psychology*, *29*, 97–101. <https://doi.org/10.1016/j.copsyc.2018.12.024>
- Theeuwes, J., Atchley, P., & Kramer, A. F. (2000). On the time course of top-down and bottom-up control of visual attention. *Control of Cognitive Processes: Attention and Performance XVIII*, 105–124.
- Theeuwes, J., Bogaerts, L., & van Moorselaar, D. (2022). What to expect where and when: how statistical learning drives visual selection. *Trends in Cognitive Sciences*, 1–13. <https://doi.org/10.1016/j.tics.2022.06.001>
- Theeuwes, J., & Chen, C. Y. D. (2005). Attentional capture and inhibition (of return): The effect on perceptual sensitivity. *Perception & Psychophysics*, *67*(8), 1305–1312. <https://doi.org/10.3758/BF03193636>
- Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology: General*, *134*(4), 552–564. <https://doi.org/10.1037/0096-3445.134.4.552>
- van Moorselaar, D., Daneshmand, 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., Foster, J. J., Sutterer, D. W., Theeuwes, J., Olivers, C. N. L., & Awh, E. (2018). Spatially Selective Alpha Oscillations Reveal Moment-by-Moment Trade-offs between Working Memory and Attention. *Journal of Cognitive Neuroscience*, *30*(2), 256–266. https://doi.org/10.1162/jocn_a_01198
- van Moorselaar, D., Huang, C., & Theeuwes, J. (2023). Electrophysiological Indices of Distractor Processing in Visual Search Are Shaped by Target Expectations. *Journal of Cognitive Neuroscience*, *35*(6), 1032–1044. https://doi.org/10.1162/jocn_a_01986
- van Moorselaar, D., Lampers, E., Cordesius, E., & Slagter, H. A. (2020). Neural mechanisms underlying expectation-dependent inhibition of distracting information. *eLife*, *9*, 1–26. <https://doi.org/10.7554/eLife.61048>
- van Moorselaar, D., & Slagter, H. A. (2019). Learning What Is Irrelevant or Relevant: Expectations Facilitate Distractor Inhibition and Target Facilitation through Distinct Neural Mechanisms. *The Journal of Neuroscience*, *39*(35), 6953–6967. <https://doi.org/10.1523/JNEUROSCI.0593-19.2019>
- van Moorselaar, D., & Slagter, H. A. (2020). Inhibition in selective attention. *Annals of the New York Academy of Sciences*, *1464*(1), 204–221. <https://doi.org/10.1111/nyas.14304>
- van Moorselaar, D., & Theeuwes, J. (2021). Statistical distractor learning modulates perceptual sensitivity. *Journal of Vision*, *21*(12), 3. <https://doi.org/10.1167/jov.21.12.3>

- van Moorselaar, D., & Theeuwes, J. (2022). Spatial suppression due to statistical regularities in a visual detection task. *Attention, Perception, & Psychophysics*, *84*(2), 450–458. <https://doi.org/10.3758/s13414-021-02330-0>
- Vatterott, D. B., & Vecera, S. P. (2012). Experience-dependent attentional tuning of distractor rejection. *Psychonomic Bulletin & Review*, *19*(5), 871–878. <https://doi.org/10.3758/s13423-012-0280-4>
- Vicente-Conesa, F., Giménez-Fernández, T., Luque, D., & Vadillo, M. A. (2023). Learning to suppress a distractor may not be unconscious. *Attention, Perception, & Psychophysics*, *85*(3), 796–813. <https://doi.org/10.3758/s13414-022-02608-x>
- Wang, B., & Theeuwes, J. (2018a). How to inhibit a distractor location? Statistical learning versus active, top-down suppression. *Attention, Perception, & Psychophysics*, *80*(4), 860–870. <https://doi.org/10.3758/s13414-018-1493-z>
- Wang, B., & Theeuwes, J. (2018b). Statistical regularities modulate attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *44*(1), 13–17. <https://doi.org/10.1037/xhp0000472>
- Wang, B., & Theeuwes, J. (2018c). Statistical regularities modulate attentional capture independent of search strategy. *Attention, Perception, & Psychophysics*, *80*(7), 1763–1774. <https://doi.org/10.3758/s13414-018-1562-3>
- Wang, B., & Theeuwes, J. (2020). Implicit attentional biases in a changing environment. *Acta Psychologica*, *206*, 103064. <https://doi.org/10.1016/j.actpsy.2020.103064>
- 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
- 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. <https://doi.org/10.3389/fnsys.2015.00123>
- 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. <https://doi.org/10.1038/nn.4546>
- 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. <https://doi.org/10.1111/nyas.12606>