Learned relevance of a distracting cue can influence feature interference errors

John A. McNally, William Narhi-Martinez, Andrew B. Leber, Julie D. Golomb

Department of Psychology, The Ohio State University

*Correspondence should be addressed to Julie Golomb, Department of Psychology, The Ohio State University, Columbus, OH, 43210. Email: golomb.9@osu.edu

Keywords: attentional capture, salient distractor, swap errors, feature interference, learned regularities

Abstract

Previous research has shown that attention can be guided via past experiences and learned regularities of target and/or distractor location. It has also been suggested that expectations surrounding salient distractors can potentially make them easier to suppress, thereby improving performance. Here, we ask: could the learned *relevance* of a distracting cue affect our ability to suppress it and protect against feature interference errors? Participants performed a delayed estimation task reporting the color of a target item, with a salient distractor cue appearing at a nontarget item's location on some trials. Critically, the experiment was split into two contexts presented in separate halves of the experiment, which differed based on whether the target and distracting cue could appear at the same location. In the "0 percent match" context, none of the trials had the salient cue appear at the target location; i.e., no 'valid' trials; whereas in the "50 percent match" context, 50 percent of distractor-present trials had the salient cue appear at the target location, increasing its potential relevance. We used a probabilistic mixture model to estimate generic performance measures (guess rate and standard deviation) as well as feature interference measures (swap rate and mean shift) for each condition. We found a significant difference in swap rates (misreporting the color of the item at the salient distractor location instead of the target color) between the two contexts depended on the order they were experienced. These results suggest that the learned relevance of a distractor cue can affect how likely participants were to be captured by a salient distractor and its resulting impact on target feature perception, and that statistical regularities relating to the relevance of salient items can affect the perception and encoding of stimulus features.

Introduction

The ability to intentionally direct our attention is crucial for day-to-day tasks and enables us to focus on what is most important at a given time, but that does not always mean we are veridically encoding everything that we see. For example, a daily task such as driving to work requires attention to be focused on the road ahead while also shifting attention to passing cars and ignoring distractors such as getting a phone call. Attention can be conceptually described as a multi-level system of weights and balances that leads to us to prioritize certain stimuli more than others (Narhi-Martinez, Dube, & Golomb, 2023), which will result in some stimuli either being processed poorly or not at all. Thus, the distracting phone call – or a car weaving into our lane – may alternatively capture our attention or not even be noticed at all depending on our attentional state in that moment.

Attention is important in not only determining how likely we are to perceive something, but it also has a role in feature-binding (Reynolds & Desimone, 1999; Treisman & Gelade, 1980) and has been theorized to be the 'glue' that binds features together to make cohesive objects. Recent studies have found that when spatial attention is unstable there is a higher chance of making feature-binding errors (Dowd & Golomb, 2019; Golomb, 2015; Golomb et al., 2014). One example of a feature-binding error is color swaps: when asked to report the color of an item at a target location, subjects sometimes mistakenly report the color of an item at a distractor location instead of the target's color.

Swap errors have been found to be particularly prevalent in the context of attentional capture by salient distractors (Chen et al., 2019). Chen and colleagues (2019) revealed the impacts of attentional capture on feature perception and recall by presenting participants with a brief (50ms) array of four colored items, one of which was highlighted with a thick border as the target, and comparing three conditions: neutral (no salient distractor cue), valid (salient cue appears surrounding the same stimulus as the target probe), and invalid (salient cue appears surrounding a nontarget stimulus). Probabilistic mixture modeling (Bays et al., 2009; Wilken & Ma, 2004; Zhang & Luck, 2008) of the participants' responses along a continuous color wheel revealed a significant amount of swapping errors in the invalid condition due to the salient distractor capturing the attention on invalid trials, thus disrupting the feature encoding of that target color. Interestingly, participants made these large swapping errors with high confidence, seemingly unaware they were reporting a distractor color instead of the target color (Chen et al., 2019).

A follow-up study by Narhi-Martinez et al. (2023) investigated how these swap errors might be impacted by learned spatial suppression elicited via a predictable salient distractor location. Previous studies have shown that we exploit learned target and/or distractor regularities to use attention more efficiently (Chun, 2005; Chun & Jiang, 1998; Ferrante et al., 2018; Geng & Behrmann, 2005; Jiang, 2018). Additionally, it has been found that previous experience with distractors plays a large role in our ability to successfully suppress them (Cunningham & Egeth, 2016; Goschy et al., 2014; Leber et al., 2016; Noonan et al., 2016; Wang & Theeuwes, 2018). Narhi-Martinez et al. (2023) showed that participants were significantly less likely to make swap errors when the salient distractor was in a highly likely location compared to a less likely location.

Interestingly, however, Narhi-Martinez et al. (2023) also found a lower baseline swap rate; even when the salient distractor appeared in a less likely location, the proportion of swap errors was, on average, half of what had been measured in Chen et al. (2019). The two studies shared largely similar designs, though one difference was that the "valid" trial condition was eliminated in the Narhi-Martinez et al. (2023) design to increase statistical power for the critical distractor location likelihood manipulation. Is it possible that the participants' overall distractor processing could have differed during the experiment relative to Chen et al. (2019) because the distractor never appeared around the same location as the target? In other words, the lack of the "valid" condition could have made the salient cue more irrelevant compared to the one used in Chen et al. (2019), and therefore potentially easier to suppress overall. It has been found that when a repeated distractor becomes expected or irrelevant, the distractor is potentially no longer classified as a distractor during processing (Sawaki & Katayama, 2006; van Moorselaar et al., 2021; van Moorselaar & Slagter, 2019). Are we able to implicitly learn the current relevance of a distractor, ultimately changing our processing of that distractor and reducing feature-binding errors? Although the prior studies have examined how learned distractor expectations affect processing, there has not been a study that directly compared multiple contexts in which the distractor relevancy differs within a single session.

In the current, preregistered study, we manipulate the relevance of the salient distractor cue in separate halves of the experiment, in which the salient cue did or did not appear at the same location as the target probe (valid trials), respectively. By comparing these different contexts in which the relevance of the salient cue differed, we aimed to test how the learned relevancy of the distractor might influence feature binding errors, and if participants could update that relevancy when a new context is introduced.

We reasoned that in contexts when valid trials are present, the relevance of the salient cue could be higher, since the salient cue may be helpful in some of the trials due to its ability to assist in directing the participants' attention to the target location. We hypothesized that participants may consequently have a harder time suppressing the salient cue when it appears around a distractor, thus leading to a higher amount of swapping errors compared to contexts when there are no valid trials and the salient cue may be more easily suppressed because it never guides attention to the target location. We investigated whether feature-binding errors differ in these contexts, and further, whether the order of contexts experienced matters, such that participants might either dynamically update learned relevance when the context changes, or persist based on the original learned context.

Methods

This experiment was preregistered on the Open Science Framework (https://osf.io/rgc3a/) before the start of data collection. The preregistration explains the motivation for the experiment, sample size, exclusion criteria, methods. We deviated from the preregistration in one way: while we initially planned 122 trials per condition as stated in the preregistration, we increased this to 144 trials per condition to ensure sufficient data to model. Analyses that were not listed in the preregistration but included in this paper are declared as exploratory.

Sample Size:

We included 56 participants in this experiment. This sample size was chosen to have sufficient power to detect a within-participant difference in swap rates between the invalid conditions in each of the two contexts, if present. Power analyses were conducted on the datasets from Chen et al. (2019) and Narhi-Martinez et al. (2023). The "swap rate" measure is defined as the difference in the probability of misreporting the color of the salient distractor compared to a misreport of a control nontarget's color. In Chen et al. (2019), the average effect size of the swap rate between the invalid conditions in both experiments that were conducted was d = 0.643. A priori power analyses using G*Power (Faul et al., 2007) on this averaged effect size (utilizing two-tailed, paired-samples t-tests, an alpha of .05, and a power of 80%) resulted in an estimation of 17 participants needed to detect significant swap-errors when valid trials are present. In Narhi-Martinez et al. (2023), they compared swap rates in two “invalid” conditions: when the distractor appeared in a highly likely location and when the distractor appeared in a less likely location. The effect size for swap errors in the latter condition (our best estimate for the no-valid context) was d = .510; resulting in a power analysis estimation of 33 participants needed to detect swap errors in this condition. Because we are interested in a difference between invalid conditions, we also conducted another power analysis on the Narhi-Martinez data for the *difference in* swap rates across those two conditions (effect size of d = 0.403). A priori analysis on this effect size resulted in an estimation of at least 51 participants. To be conservative and facilitate counterbalancing of block order, we pre-registered our sample size at 56 participants.

Our sample included 25 male and 31 female participants (mean age: 19.52), and every participant reported having normal or corrected-to-normal color vision and visual acuity. As stated in our preregistration, participants were excluded and replaced if (1) they did not have at least 80 trials per condition in which they successfully maintained fixation to include in the mixture models; (2) their performance on the task was very poor even on Neutral trials (standard deviation parameter greater than 80 or pT less than 0.5; see **Analyses** below); or (3) the probability of reporting a nontarget (β_1 or β_2) on Invalid trials was greater than 0.5, suggesting they may have misunderstood the task. Fifteen additional participants completed the experiment and were excluded and replaced because they did not meet criteria #1, and 1 additional participant was excluded and replaced for criterial #2.

Setup:

Each participant was seated and placed their head against a chin and forehead rest 60cm away from the monitor. The 62cm LCD monitor's resolution was adjusted to display a 4x3 presentation window (resolution: 1280x960, refresh rate: 200Hz) and was color calibrated with a Minolta CS-100 colorimeter. Stimuli were generated using MATLAB (Mathworks) and the Psychophysics Toolbox (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997) on a Windows computer. Eye position was recorded using an Eyelink 1000 eye-tracker (SR Research).

Procedure:

The design for this experiment (Figure 1) was similar to Chen et al. (2019). Each trial began with the participant fixating on the center of the screen at a black cross on a grey background (RGB [127.5, 127.5, 127.5]). Once the participant fixated for 700ms within 2° (visual degrees) of the cross, four thin white framed squares appeared for 300ms, and the cross turned into a black dot. If participants did not maintain fixation on the cross and dot for the entire 1000ms, the cross would reappear on the screen, and they would keep repeating the process until they could maintain fixation for the entire 1000ms.

Four colored squares then appeared for 50ms, each sized 2° x 2°, equally spaced at an eccentricity of 4° from the center of the screen. The colors presented varied on each trial. The color of the upper left square was chosen randomly from 180 different color values, spaced along a color wheel in CIE L*a*b* color space with L* = 60, a* = 22, b* = -1, and radius = 50. The colored squares in the upper right and lower left were then randomly assigned on each trial to be +90° and -90° along the color wheel from the color in the upper left square. The lower right square was always colored 180° away in color space from the color in the upper left square. In each trial, one colored square had a thick, white border around it, indicating it as the target. In some of the trials, four white dots ('salient cue') also appeared around one of the squares. This stimulus was chosen because previous studies have shown that when the distractor is similar to the target (white border around one square), attentional capture is the greater compared to when the distractor is not similar to the target (Folk et al., 1992, 2002; Folk & Remington, 1998). Trials could be either valid (target and salient cue highlight same location), invalid (salient cue highlights different location than target), or neutral (no salient cue present). After the items were removed from the screen, there was a 100ms delay, followed by a 200ms mask, before the color wheel (diameter = 6.5°, width = 1°) appeared on the screen along with a post cue to remind participants of the target location. The participant then used the mouse to click the color on the color wheel that they believed was in the target location.

Two different contexts differing in the relevance of the salient cue were tested in separate halves of the experiment: a "0 percent match" context (0 percent of trials had the salient cue appear at the target location; i.e., 0 valid trials) and a "50 percent match" context (50 percent of distractor-present trials had the salient cue appear at the target location; i.e. 1/3 of total trials were valid). For the 0 percent match context, participants completed 144 trials of the neutral condition and 144 trials of the invalid condition, randomly intermixed. For the 50 percent match context, participants completed 144 trials of the neutral condition, 144 trials of the invalid condition, and 144 trials of the valid condition, randomly intermixed. There were 12 blocks total in the experiment (6 for each context), with 48 trials per block for the 0 percent match and 72 trials per block in the 50 percent match context. Participants were assigned to one of two order groups: Half of the participants performed the 0 percent match blocks in the first half of the experiment and the 50 percent match blocks in the second half (0% → 50% group), and the other participants performed the task in the opposite order (50% → 0% group).

Participants were given instructions (fixate at the center of the screen and report the color of the square that had the thick white border around it) and 10 practice trials (excluded from analyses) at the start of the experiment to make sure they understood the task and were able to successfully fixate at the center of the screen. The salient cue was not mentioned in the instructions nor was the concept of the differing contexts in each half of the experiment. This information was omitted from instructions with the intention of creating an environment where the participants would incidentally learn the relationship of the distractor to the target and minimizing explicit consideration of the salient cue in a top-down manner. At the end of the experiment, there were exit questions that participants filled out which asked questions regarding how they thought the task was and their overall performance.

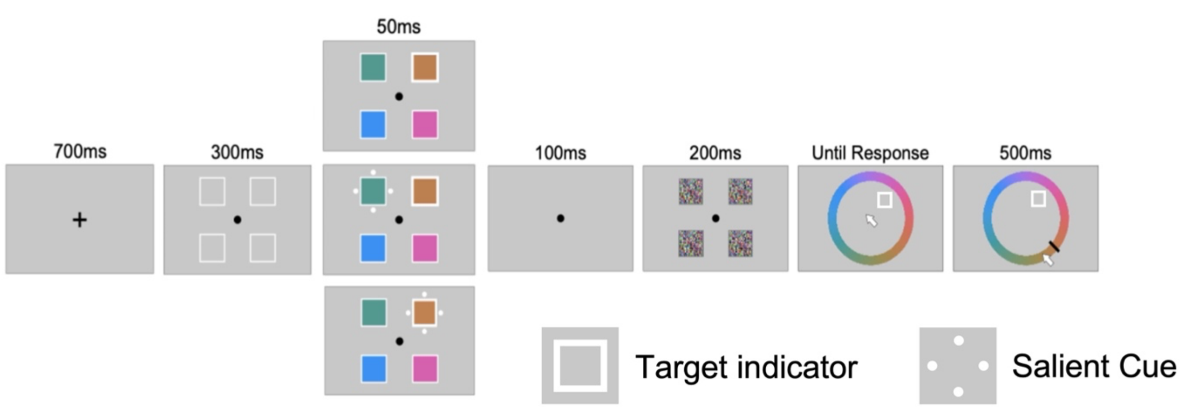

Figure 1. Experimental procedure (not drawn to scale). For every trial, participants were shown 4 colored squares and asked to report the color of the square that had a thick white frame around it (target location). The "0 percent match" context included neutral (top) and invalid (middle) trials, while the "50 percent match" context included neutral, invalid, and valid (bottom) trials.

Analyses

Any trials during which the participant broke fixation (gaze deviated more than 2 visual degrees from the central fixation cue) were discarded and not analyzed. On invalid trials, we only analyzed trials in which the salient distractor was adjacent to the target (2/3 of invalid trials), as in Chen et al (2019) and Narhi-Martinez et al. (2023). Color report errors were calculated for each trial, aligned so that the correct target color was 0° error, the salient distractor item's color was aligned to +90°, and the control (non-salient) nontarget location was -90°. Thus, positively signed errors mean color errors shifted 'towards' the distractor and negatively signed errors 'away' from the distractor. On neutral and valid trials, there was no salient distractor, so the two control nontargets adjacent to the target were assigned +90° or -90° at random to avoid any directional biases.

A probabilistic mixture model (see formula below) was used to fit the distribution of response errors for each of the 5 conditions (Neutral, Invalid, and Valid for 50% match context and Neutral and Invalid for 0% match context) for each participant. A Markov chain Monte Carlo method using MemToolbox (Suchow et al., 2013) was utilized to calculate the best-fitting parameter estimates.

$$p(\theta) = (1 - \beta_1 - \beta_2 - \gamma)\phi_{\mu,\kappa} + \beta_1\phi_{90°,\kappa} + \beta_2\phi_{-90°,\kappa} + \gamma(\frac{1}{2\pi})$$

The error between the target and reported colors was represented by θ which is the difference in degrees between target color and reported color. Probability of reporting the target (pT) was represented by 1 - β_1 - $\beta_2 - \gamma$. ϕ is the von Mises distribution with mean μ and concentration κ parameters. The variables κ and γ estimated the precision and guess rate, respectively, and were used to assess the overall performance of a participant for each condition. For the invalid condition, the β_1 - β_2 difference was utilized to measure the rate of swapping errors, where β_1 was the probability of misreporting the salient distractor item's color and β_2 was the probability of misreporting the control nontarget color. The variable μ was used to measure mean shift of the target distribution, if shifted from 0°. These measures were calculated for each condition and individual participant. The parameter estimates were analyzed in JASP software (Version 0.16.03) and MATLAB (MathWorks) using repeated measures ANOVA, paired samples t-tests, and one sample t-tests, frequentist and Bayesian.

To ensure that the premise of attentional capture by an unpredictable salient distractor was achieved, we first compare the above measures between our neutral and invalid conditions within each context. The main comparison of interest is a within-subject analysis comparing systematic feature binding errors, particularly swap rates (β_1 - β_2) and mean shifts induced by the salient distractor in the invalid condition, between the two contexts. We then perform between-subjects analyses examining potential effects of context order and learning.

Results

Generic Performance Indicators:

Our generic performance analyses included measuring the guess rate and standard deviation (SD = $\sqrt{1/k}$) for each of the conditions in both contexts (Figure 2). Paired-sample t-tests showed that the invalid condition guess rate was significantly higher than the neutral condition guess rate within 0 percent context, $t(55) = -4.232$, $p < .001$, $d = 0.566$, $BF_{10} = 248$ as well as the in the 50 percent match context, $t(55) = -4.001$, $p < .001$, $d = 0.535$, $BF_{10} = 1.23$ X 10^2. There was also a significantly higher guess rate in the invalid condition compared to the valid condition in the 50 percent match context, $t(55) = -5.094$, $p < .001$, $d = 0.681$, $BF_{10} = 3.98$ x 10^3. For standard deviation, the invalid condition had a marginally higher standard deviation compared to the neutral condition in the 0 percent match context, $t(55) = -1.938$, $p = 0.058$, $d = 0.259$, $BF_{10} = 0.829$. In the 50 percent match context, the invalid condition standard deviation was significantly higher than the neutral condition, $t(55) = -3.125$, $p = 0.003$, $d = 0.418$, $BF_{10} = 10.8$, as well as valid condition, $t(55) = 4.699$, $p < .001$, $d = 0.628$, $BF_{10} = 1.08 \times 10^3$. The standard deviation for the neutral condition was marginally higher than the valid condition in the 50 percent match context, $t(55) = 1.681$, $p = 0.098$, $d = 0.225$, $BF_{10} = 0.545$. These results parallel those reported in Chen et al (2019) and Narhi-Martinez et al. (2023).

When comparing the guess rate and standard deviation for the invalid trials across contexts, there were no significant differences for guess rate: $t(55) = 1.011$, $p = 0.317$, $d = 0.135$, $BF_{10} = 0.237$, nor standard deviation: $t(55) = -0.152$, $p = 0.879$, $d = 0.020$, $BF_{10} = 0.148$, suggesting that the difference in relevance of the salient cue between the two contexts did not affect general performance.

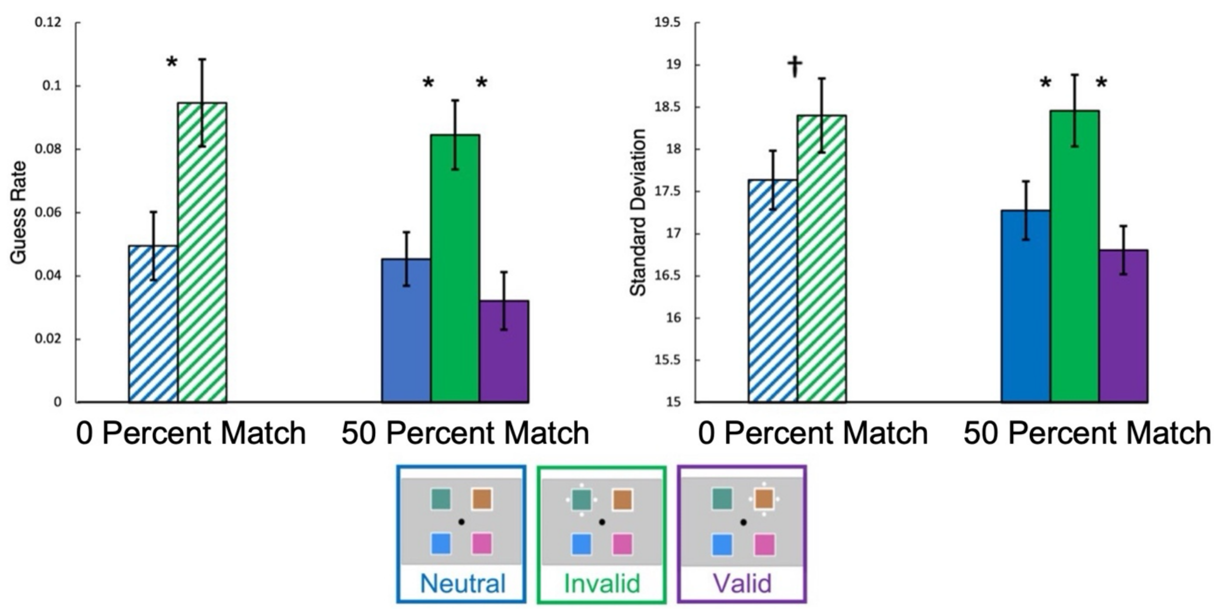

Figure 2. General performance indicators. Probabilistic mixture model results for guess rate (left) and standard deviation (right). The error bars represent SEM, *: $p < 0.05$, †: $p < 0.10$, N = 56

Systematic Feature Binding Errors:

Our main question of interest concerned potential differences in systematic feature binding errors induced by the salient distractor cue under the different contexts. We first examined swap errors (Figure 3A), assessed by the probability of misreporting the color of the salient distractor item (β_1) relative to the probability of misreporting the color of the control nontarget item (β_2). A paired-sample t-test showed that there was a significant difference between β_1 and β_2 for invalid trials in the 0 percent match context, $t(55) = 4.065$, $p < .001$, $d = 0.543$, $BF_{10} = 1.49 \times 10^2$, and in the 50 percent match context, $t(55) = 5.453$, $p < .001$, $d = 0.729$, $BF_{10} = 1.34 \times 10^4$. In other words, participants made significant swap errors in both contexts.

In order to compare the amount of swap errors across contexts, the swap rate (β_1 - β_2) was calculated for each participant in the two contexts (Figure 3B). We found a marginally higher swap rate in the 50 match context compared to the 0 match context, $t(55) = -1.808$, $p = 0.076$, $d = 0.242$, $BF_{10} = 0.667$. This difference was in the predicted direction – such that the salient distractor would be harder to ignore and produce more feature interference when it sometimes overlaps with the target location compared to when it is purely distracting – though the difference between contexts was not statistically significant, at least when ignoring potential order effects (see next section).

The mean shift parameter was also analyzed between the invalid conditions of each context to examine potential repulsion or attraction errors induced by the salient distractor (Figure 3C). The mean shift in the invalid condition was not significantly different from zero in either condition: 0 percent match context: $t(55) = -0.683$, $p = 0.498$, $d = -0.079$, $BF_{10} = 0.182$; 50

percent match context: $t(55) = -0.594$, $p = 0.555$, $d = -0.079$, $BF_{10} = 0.173$. A paired-sample t-test found no significant difference between the mean shifts in the invalid conditions across contexts, $t(55) = -0.073$, $p = 0.942$, $d = 0.010$, $BF_{10} = 0.146$.

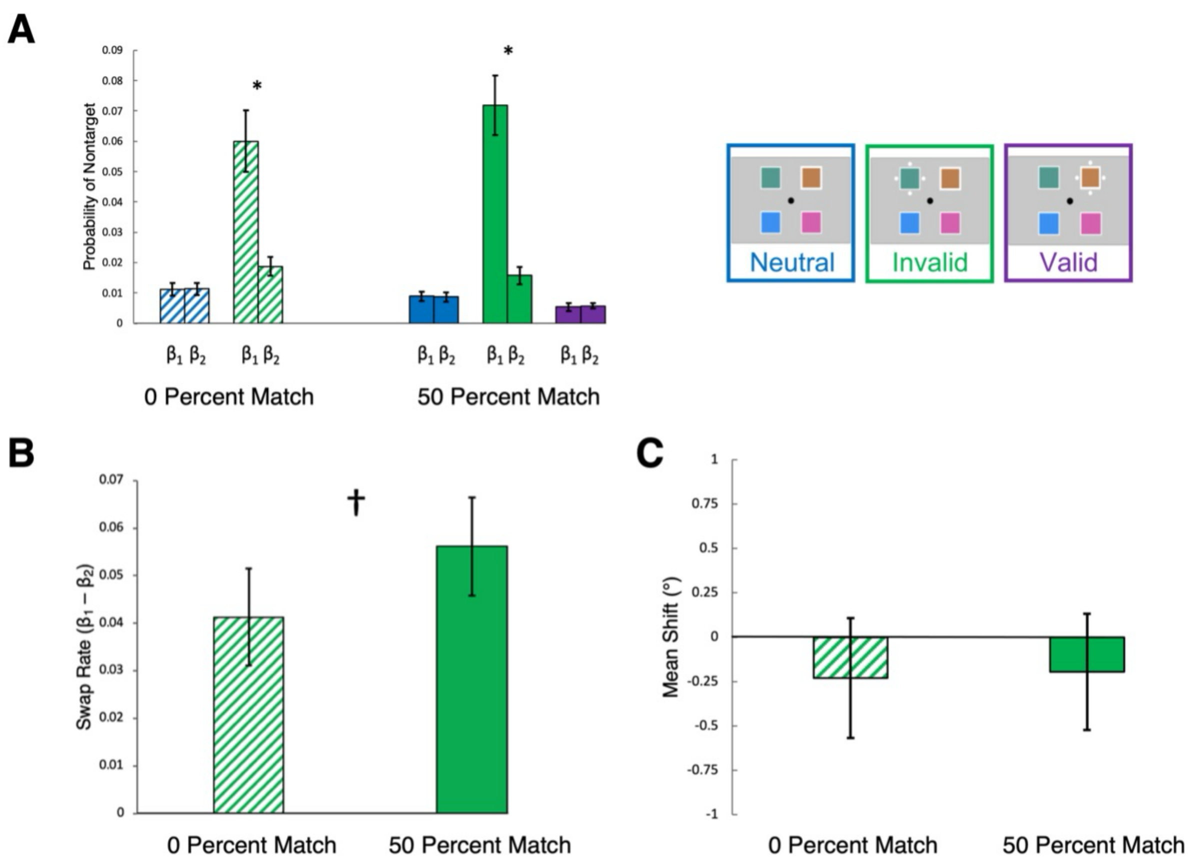

Figure 3. Systematic feature binding errors. **(A)** Probabilistic mixture model results for swap errors, plotting the probability of misreporting the color of the item at the salient distractor location (β_1) and the probability of misreporting the color of the item at the control nontarget location (β_2). **(B)** Distractor-induced swap rates (β_1 - β_2) for the invalid conditions. **(C)** Probabilistic mixture model results for mean shift parameter (μ). Positively-signed shifts indicate attraction towards the color of the item at the salient distractor location and negatively-signed shifts indicate repulsion away from the color of the item at the salient distractor location. The error bars represent SEM, *: $p < 0.05$, † = $p < 0.10$, N = 56

Context Order Effects:

We next investigated whether the order of which context was presented first affected participants' ability to suppress the salient distractor and protect the target from feature interference. Because the mean shift results were negligible, we focus the subsequent analyses on the swap rate measure. We first report the results of a pre-registered analysis between the two groups of participants (0% → 50% group vs. 50% → 0% group) comparing their first and second halves of the experiment. We predicted that there may be a difference in the amount of swap errors depending on the order the contexts were experienced. In other words, if participants persisted in the behavior associated with their initial context exposure, results from their second half of the experiment might mimic the first half, ignoring the context change. Such a pattern would weaken the overall context effects reported above. On the other hand, if participants rapidly adapt to the new contingencies in the second half, their performance should only vary with context, with no order or group effects, such that their swap rate for the invalid trials might always be slightly higher in the 50 percent match context compared to the 0 percent match context. It is also possible the results could reflect additional possibilities, for example, that participants from both groups get better over time due to practice and exposure, and swap rate for the first half of the experiment would always be higher compared to the swap rate for the second half of the experiment (such that context differences may be most evident between groups in the first half of the experiment), or even that the *change* from one context to another triggers a more salient effect (such that context differences may be most evident between groups in the second half of the experiment).

Figure 4 shows the results of this context order analysis, with invalid condition swap rates broken down by the order of which context a participant saw first. A repeated measures ANOVA comparing swap rates by context (0% match vs. 50% match) and group (between-subjects factor; 0% → 50% group vs. 50% → 0% group) revealed no significant main effect for group, $F(1,54) = 0.894$, $p = 0.349$, $\eta^2 = 0.014$, group model $BF_{10} = 0.518$, but a significant context x group interaction, $F(1,54) = 6.115$, $p = 0.017$, $\eta^2 = 0.016$, with strongest evidence for the context + group + context x group model $BF_{10} = 1.361$ compared to the null model. Consistent with the overall comparison above, we also found a marginal main effect of context, $F(1,54) = 3.573$, $p = 0.064$, $\eta^2 = 0.009$, context model $BF_{10} = 0.814$. These results suggest that the order in which the contexts were experienced did impact the amount of swapping for each context. As can be seen in Figure 4, participants who started with the 50 percent match context made significantly fewer swap errors in the subsequent half of the experiment where they experienced the 0 percent match context, $t(27) = 2.672$, $p = 0.013$, $d = 0.505$, $BF_{10} = 3.790$. However, when the 0 percent match context was experienced first, the swap rates in the second half of the experiment (50 percent match context) were not significantly different relative to the first half, $t(27) = 0.504$, $p = 0.618$, $d = 0.095$, $BF_{10} = 0.225$. When comparing the second half of the experiment across groups, there was a marginal difference in swap rates, $t(27) = -1.856$, $p = 0.074$, $d = 0.351$, $BF_{10} = 0.900$. We further explore this finding in the exploratory block-by-block timecourse section below.

We also conducted similar ANOVAs on the mean shift, guess rate, and standard deviation parameters, as pre-registered, to determine if the context order influenced any of these measures. We found no significant main effects or interactions for any of these measures (Table 1).

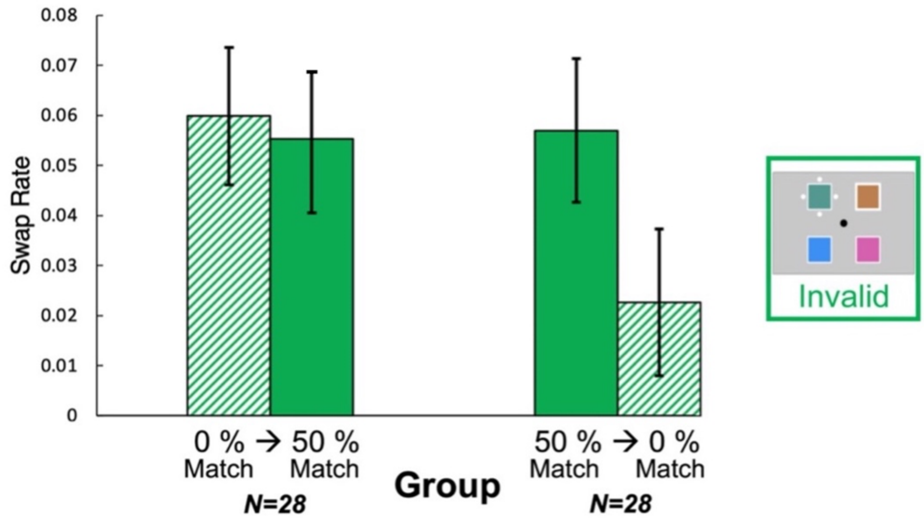

Figure 4. Context order effects. Swap rates in the invalid conditions broken down by participant group, based on the order that participants experienced each context. The error bars represent SEM, N = 56 (N=28 per group).

Source	Cases	df	F	p	η^2
Swap Rate Invalid Conditions by Group					
Within Subjects	Context	1	3.573	0.064	0.009
	Context * Group First	1	6.115	0.017	0.016
	Residuals	54			
Between Subjects	Group First	1	0.894	0.349	0.014
	Residuals	54			
Mean Shift Invalid Conditions by Group					
Within Subjects	Context	1	0.005	0.942	5.173e-5
	Context * Group First	1	3.021e-5	0.996	2.974e-7
	Residuals	54			
Between Subjects	Group First	1	0.039	0.844	3.404e-4
	Residuals	54			
Guess Rate Invalid Conditions by Group					
Within Subjects	Context	1	1.049	0.310	0.003
	Context * Group First	1	2.476	0.121	0.007
	Residuals	54			
Between Subjects	Group First	1	0.095	0.760	0.001
	Residuals	54			
Standard Deviation Invalid Conditions by Group					
Within Subjects	Context	1	0.024	0.878	8.065e-5
	Context * Group First	1	2.314	0.134	0.008
	Residuals	54			
Between Subjects	Group First	1	0.018	0.893	2.742e-4
	Residuals	54			

Table 1. ANOVA results examining context order effects for each of the different measures on invalid trials.

Exploratory Block-by-Block Timecourse Analysis

In order to better understand how the pattern of swapping rates may have changed during and across each half of the experiment, we conducted an exploratory block-by-block analysis of swap rates for the 12 total blocks (6 in each half of the experiment). Each block did not have a sufficient number of trials to run a mixture model on individual participants, so for this exploratory analysis we aggregated the data from all participants in each group into two 'super subjects' and modeled the errors by block. We calculated the swap rate in each of the blocks for each of the two groups of participants. Figure 5A shows the block-by-block timecourse of swap rates for the group of participants that experienced the 0% match context in the first six blocks followed by the 50% context in the second six blocks. Figure 5B shows the block-by-block timecourse of swap rates for the group that experienced the contexts in the opposite order.

A few interesting patterns emerge. First, the swap rate decreased rapidly over the first six blocks in both groups. There was a significant negative correlation between block number and swap rate for both the group that experienced the 0% match context first and the group that experienced the 50% match context first, $r(4) = -0.962$, $p = 0.002$, and $r(4) = -0.901$, $p = 0.014$, respectively. These results suggest that swapping errors reduced over time within both contexts during the first half of the experiment, presumably due to practice effects and/or accumulated exposures to the salient distractor.

Strikingly, the swap errors show different temporal patterns in the second half of the experiment for the two groups. For the group that experienced the 0 percent match context first (no valid trials), when they switch to the new 50 percent match context where the salient cue is sometimes overlapping with the target (valid), swap rates on invalid trials increase over blocks 7-8 and then remained consistently higher throughout this latter half of the experiment (no significant correlation between block number and swap rate: $r(4) = -0.216$, $p = 0.680$). On the other hand, for the group that experienced the 50 percent match context first, when they switch to the 0 percent match context containing no valid trials, there was a brief initial boost in swap errors in block 7, followed by a rapid decrease and sustained lower swap rate over the course of the second half, $r(4) = -0.429$, $p = 0.396$. Although this analysis is exploratory, it suggests that participants may have been sensitive to the change in contexts, such that overall, there were fewer swapping errors when the 0 percent match came second, compared to when the 50 percent match came second.

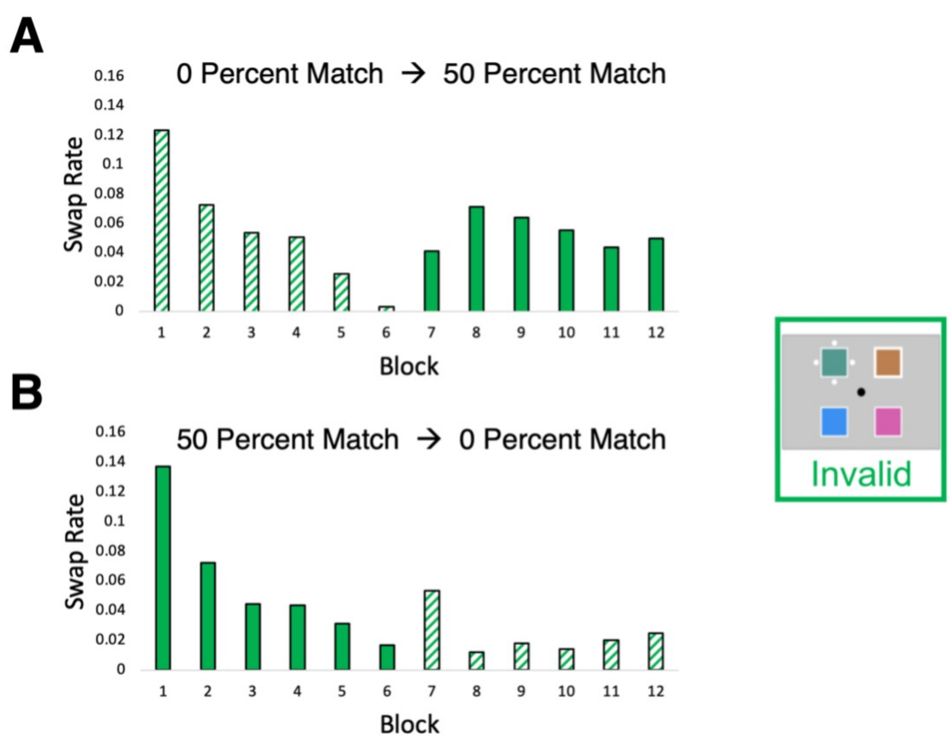

Figure 5. Exploratory block-by-block timecourse analysis on swap rates for invalid trials, broken down by group. **(A)** Block-by-block swap rate for the group of participants who experienced the 0 percent match context first. **(B)** Block-by-block swap rate for the group of participants who experienced the 50 percent match context first. Mixture model parameters for this analysis were fit on aggregate subject data (N=28 for each group).

Exit Question Results:

We also had participants answer exit questions after they were done with the experiment. The first question asked, “Did you notice white dots sometimes flash around one of the colored squares?” 18/56 participants reported that they did not notice, meaning that they were not consciously aware that a distractor was present in some trials. A context x awareness ANOVA conducted on swap rates (β_1 - β_2) in the invalid conditions found no significant main effect for context, $F(1,54) = 1.888$, $p = 0.175$, $\eta^2 = 0.006$, context model $BF_{10} = 0.809$, nor for awareness, $F(1,54) = 0.156$, $p = 0.694$, $\eta^2 = 0.002$, awareness model $BF_{10} = 0.411$. There was no significant interaction between the two factors, $F(1,54) = 0.759$, $p = 0.387$, $\eta^2 = 0.002$, context + awareness + context x awareness model $BF_{10} = 0.142$. These results suggest that the swap rates in the invalid conditions were not affected by whether the participant was consciously aware of the salient distractor.

The second question asked, “Did you find that the first or second half of the experiment easier?”. After taking into account which context-order group they were assigned to, we found that 33/56 participants felt that the 0 percent match context was easier, consistent with the overall trend for fewer swap errors in this context. However, the pattern of swap errors did not significantly differ across groups. A context x response group ANOVA conducted on swap rates (β_1 - β_2) in the invalid conditions found no significant main effect of response group, $F(1,54) = 0.970$, $p = 0.329$, $\eta^2 = 0.015$, response group model $BF_{10} = 0.534$. Additionally, no significant interaction between the two factors, $F(1,54) = 0.019$, $p = 0.891$, $\eta^2 = 5.593$ e-5, context + response group + context x response group model $BF_{10} = 0.120$.

Discussion

Many previous studies have examined how visual distractors affect attention, commonly focusing on different aspects of attentional capture, including the ubiquitous finding that when a salient distractor is present, reaction times slow down and higher error rates occur compared to when there was no distractor present (Folk et al., 1992, 2002; Luck et al., 2021; Theeuwes, 1996; Theeuwes & Burger, 1998; Yantis & Jonides, 1984). More recently, it has been demonstrated that attentional capture can also lead to perceptual errors such as feature swapping and repulsion (Chen et al., 2019).

An important focus of the attentional capture literature has been understanding how and when distractors can be suppressed (Geng et al., 2019; Geng, 2014; van Moorselaar & Slagter, 2019; Wang & Theeuwes, 2018). In particular, learned statistical regularities – such as learning a certain location is more likely to contain a distractor – can aid in the ability of a person to suppress a salient object (Failing & Theeuwes, 2020; Huang et al., 2021, 2022; Kong et al., 2020; Leber et al., 2016; Wang et al., 2019). Indeed, learned spatial suppression can also protect against feature interference, with fewer swap errors evoked by a salient distractor in a high-probability location than a low-probability location (Narhi-Martinez et al., 2023).

One proposed idea as to why participants are able to suppress distractors is due to their ability to learn and expect the salient distractor over time, thus potentially changing the way the distractor is processed (van Moorselaar & Slagter, 2019). More frequent distractors – regardless of spatial regularities – also become easier to ignore compared to contexts where the distractors are more rare (Geng et al., 2019; Geyer et al., 2006; Won & Geng, 2020). In the current study, we focused on another type of learned regularity affecting the relevance of a distracting cue: whether the salient distractor ever overlaps with the target location.

We examined how the learned reliability/relevance of a salient cue would affect the amount of feature errors made within different relevancy contexts. In the current study the two contexts (0 percent match and 50 percent match) allowed participants to learn two different relevancies of the salient cue, asking whether this learned relevance would impact the extent to which the distractor induced feature binding errors. The 0 percent match context represented the low relevance salient cue due to the absence of any target-distractor overlap, while the 50 percent match context represented the higher relevance salient cue, because there was the possibility of overlap between the distractor and target, such that in some trials the salient distractor appeared in the same location as the target, drawing attention towards the participant's goal in a potentially helpful way. We predicted that this greater relevance in the 50 percent match context might lead to a higher swap rate compared to the 0 percent match context due to participants being less likely to suppress the salient cue in the former.

Our results showed that in both contexts, the salient distractor captured attention, resulting in worse overall performance on the invalid (distractor present) trials compared to neutral (distractor absent) trials, measured by the general performance indicators of precision and guess rate. On invalid trials in both contexts, participants also exhibited swap errors: a significantly higher probability of misreporting the non-target color that appeared at the location where the salient cue was compared to a control nontarget, replicating the overall trend as seen in Chen et al. (2019) and Narhi-Martinez et al. (2023).

Importantly, the two contexts differed only in the presence or absence of an additional "valid" condition. This contextual manipulation did not seem to influence the general performance measures, with no significant difference in precision and guess rate on the invalid trials across the two contexts, suggesting that general performance was not affected when the relevance of a salient cue was altered. When looking at the systematic feature errors for the invalid conditions in each of the contexts, we found a slightly higher swap rate in the 50 percent match context compared to the 0 percent match context, though this difference was not statistically significant.

Our pre-registered follow-up analyses did, however, find that the order in which the contexts were experienced had a significant effect on swap rates. We designed this study such that the relevance of the salient cue changed halfway through the experiment, meaning that participants had the opportunity to update their perceived relevance of the salient cue. Previous studies have shown that the processing of a repeated distractor/salient cue can be altered when it becomes expected or irrelevant (Sawaki & Katayama, 2006; van Moorselaar et al., 2021; van Moorselaar & Slagter, 2019). Our results showed that when participants experienced the 50 percent match context first (higher relevance distractor), swap rates were significantly reduced in the second half of the experiment during the 0 percent match condition. However, when participants experienced the 0 percent match context first, the rate of swap errors did not change in the second half of the experiment during the 50 percent match condition. In pre-registering this analysis, we had aimed to distinguish between two main hypothetical patterns: that participants would be sensitive to the current context, updating with the change in context, or that participants would persist in the behavior associated with the original context, failing to update to the change in context. Instead we found a more complicated pattern, where there seems to be some interaction between context sensitivity and overall time in the task. It is not surprising that performance would improve over time and swap errors would reduce with distractor exposure regardless of context, and indeed the exploratory block-by-block analysis shows an expected effect of time in the first half of the experiment. What was less expected was that instead of an initial effect of relevance context in the first half of the experiment, the differences across relevance contexts only emerged in the second half of the experiment, after the switch.

There could be a few possible accounts for this pattern of results. One explanation is that the difference in relevance context only became salient to the participants after the switch from one context to another. In other words, salient distractors caused similar initial rates of swap errors in the invalid condition in both contexts, but when the salient distractor becomes *relatively more* relevant (0% → 50%), the distractor interferes more after that change in context compared to when the salient distractor becomes *relatively less* relevant (50% → 0%), Another explanation could be that context played a higher role in the second half of the experiment because distractor relevance simply takes more time to learn. Motivated by the large difference in baseline swap rates between Chen et. al (2019) and Narhi-Martinez et al. (2023), we designed the current study under the assumption that 6 blocks would be sufficient to detect differences in learned relevance context. However, if learned distractor relevance did not fully account for that difference across experiments, it is possible that more exposure time is needed to fully detect these context effects. Future studies may be better able to disentangle these possibilities by collecting datasets with additional exposure conditions, such as no-update groups (0% → 0% or 50% → 50%), or manipulations varying the relevance more finely (e.g. 0%, 25%, 50% match contexts) to compare effects of relative vs absolute differences in relevance context. Regardless, our findings highlight the idea that both relevance context and order can play a factor in the magnitude of systematic feature errors participants make.

Interestingly, although our results suggest an intriguing potential effect of relevance and context order, we only found a marginal difference between the two contexts overall, and the rate of swapping errors in the 50 percent match context was substantially less than in Chen et al. (2019). We also did not find evidence of feature repulsion in the 50 percent match context, though repulsion errors were reliable across both experiments in Chen et al. (2019). This suggests that the relevance of the distractor can influence feature interference, but it is not the only factor that does so, and there are likely other reasons for the differences across studies. Indeed, it can be difficult to isolate reasons for differences found across studies, especially when studies are run on different participants at different points in time. Some additional differences between Chen et al. (2019) and the current study's 50 percent match context include different groups of participants (e.g., pre- vs. post-pandemic), the type of monitor used (CRT vs. LCD), the amount of trials in each condition, etc. The feature interference errors found in both contexts of the current study more closely mimic Narhi-Martinez et al. (2023), which collected data within a more similar time period as the current study and used an identical testing environment. Thus, the differences in magnitude of feature interference reported in Narhi-Martinez et al. (2023) and Chen et al (2019) are unlikely to have been solely driven by the different distractor relevance contexts, although this could have been one of the contributing factors. We conclude that the degree of feature interference induced by a salient distractor may be influenced by multiple factors, some still unknown, and as such it is particularly important to conduct within-subject comparisons, or at least between-subject comparisons where the groups are drawn from the same subject pool at the same point in time on the same computers, as in the current study. That said, it is also important to emphasize that the existence of these feature interference errors – particularly swap errors – is consistently replicable across studies, even if the magnitude varies.

In conclusion, the current study investigated how the relevance of a distracting cue impacts the magnitude of systematic feature errors that we can make. Within the real world, there are copious salient cues with differing levels of relevance to our goals that, if repeated enough times, can be learned about and potentially used to our advantage (consciously or otherwise). Our results indicate that when a salient distractor is present to capture attention away from our goal, the learned relevance of that salient cue may be an important factor in how much we can suppress it. In other words, that learned relevancy of salient items could potentially change the way they are processed and impact the suppression of later salient distractors. Our results indicate that the ordering of salient cue context plays a key role in determining the swap rate, potentially via relative changes in context, indicating that the previous and current learned relevancies of a salient cue can interact to contribute to our processing of targets and distractors.

Acknowledgments

The authors gratefully acknowledge the assistance of Isabel Jaen and Jake Ferreira in the recruitment of participants and data collection. This research was supported by NIH Grant R01-EY025648 (JG); and NSF Grant BCS-1848939 (JG, AL).

Declaration of Interest Statement

The authors report there are no competing interests to declare.

References

Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), 7. https://doi.org/10.1167/9.10.7

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436. https://doi.org/10.1163/156856897X00357

Chen, J., Leber, A. B., & Golomb, J. D. (2019). Attentional capture alters feature perception. *Journal of Experimental Psychology: Human Perception and Performance*, *45*, 1443–1454. https://doi.org/10.1037/xhp0000681

Chun, M. M. (2005). CHAPTER 40—Contextual Guidance of Visual Attention. In L. Itti, G. Rees, & J. K. Tsotsos (Eds.), *Neurobiology of Attention* (pp. 246–250). Academic Press. https://doi.org/10.1016/B978-012375731-9/50044-6

Chun, M. M., & Jiang, Y. (1998). Contextual Cueing: Implicit Learning and Memory of Visual Context Guides Spatial Attention. *Cognitive Psychology*, *36*(1), 28–71. https://doi.org/10.1006/cogp.1998.0681

Cunningham, C. A., & Egeth, H. E. (2016). Taming the White Bear: Initial Costs and Eventual Benefits of Distractor Inhibition. *Psychological Science*, *27*(4), 476–485. https://doi.org/10.1177/0956797615626564

Dowd, E. W., & Golomb, J. D. (2019). Object-Feature Binding Survives Dynamic Shifts of Spatial Attention. *Psychological Science*, *30*(3), 343–361. https://doi.org/10.1177/0956797618818481

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

Folk, C. L., Leber, A. B., & Egeth, H. E. (2002). Made you blink! Contingent attentional capture produces a spatial blink. *Perception & Psychophysics*, *64*(5), 741–753. https://doi.org/10.3758/BF03194741

Folk, C. L., & Remington, R. (1998). Selectivity in distraction by irrelevant featural singletons: Evidence for two forms of attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 847–858. https://doi.org/10.1037/0096-1523.24.3.847

Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 1030–1044. https://doi.org/10.1037/0096-1523.18.4.1030

Geng, J. J. (2014). Attentional mechanisms of distractor suppression. *Current Directions in Psychological Science*, *23*, 147–153. https://doi.org/10.1177/0963721414525780

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

Geng, J., Won, B.-Y., & Carlisle, N. (2019). Distractor ignoring: Strategies, learning, and passive filtering. *Current Directions in Psychological Science*, *28*(6), 600–606. https://doi.org/10.1177/0963721419867099

Geyer, T., Müller, H. J., & Krummenacher, J. (2006). Cross-trial priming in visual search for singleton conjunction targets: Role of repeated target and distractor features. *Perception & Psychophysics*, *68*(5), 736–749. https://doi.org/10.3758/bf03193697

Golomb, J. D. (2015). Divided spatial attention and feature-mixing errors. *Attention, Perception, & Psychophysics*, *77*(8), 2562–2569. https://doi.org/10.3758/s13414-015-0951-0

Golomb, J. D., L'Heureux, Z. E., & Kanwisher, N. (2014). Feature-Binding Errors After Eye Movements and Shifts of Attention. *Psychological Science*, *25*(5), 1067–1078. https://doi.org/10.1177/0956797614522068

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

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*, *48*, 443–457. https://doi.org/10.1037/xhp0001002

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. https://doi.org/10.3758/s13423-021-01891-3

Jiang, Y. V. (2018). Habitual versus goal-driven attention. *Cortex*, *102*, 107–120. https://doi.org/10.1016/j.cortex.2017.06.018

Kleiner, M., Brainard, D. H., & Pelli, D. (2007). *What's new in Psychtoolbox-3?* 89.

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

Leber, A. B., Gwinn, R. E., Hong, Y., & O'Toole, R. J. (2016). Implicitly learned suppression of irrelevant spatial locations. *Psychonomic Bulletin & Review*, *23*(6), 1873–1881. https://doi.org/10.3758/s13423-016-1065-y

Luck, S. J., Gaspelin, N., Folk, C. L., Remington, R. W., & Theeuwes, J. (2021). Progress Toward Resolving the Attentional Capture Debate. *Visual Cognition*, *29*(1), 1–21. https://doi.org/10.1080/13506285.2020.1848949

Narhi-Martinez, W., Dube, B., Chen, J., Leber, A. B., & Golomb, J. (2023). *Suppression of a salient distractor protects the processing of target features*. PsyArXiv. https://doi.org/10.31234/osf.io/7c8gt

Narhi-Martinez, W., Dube, B., & Golomb, J. D. (2023). Attention as a multi-level system of weights and balances. *WIREs Cognitive Science*, *14*(1), e1633. https://doi.org/10.1002/wcs.1633

Noonan, M. 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

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.

Reynolds, J. H., & Desimone, R. (1999). The Role of Neural Mechanisms of Attention in Solving the Binding Problem. *Neuron*, *24*(1), 19–29. https://doi.org/10.1016/S0896-6273(00)80819-3

Sawaki, R., & Katayama, J. (2006). Stimulus context determines whether non-target stimuli are processed as task-relevant or distractor information. *Clinical Neurophysiology*, *117*(11), 2532–2539. https://doi.org/10.1016/j.clinph.2006.06.755

Theeuwes, J. (1996). Parallel search for a conjunction of color and orientation: The effect of spatial proximity. *Acta Psychologica*, *94*(3), 291–307. https://doi.org/10.1016/S0001-6918(96)00003-0

Theeuwes, J., & Burger, R. (1998). Attentional control during visual search: The effect of irrelevant singletons. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1342–1353. https://doi.org/10.1037/0096-1523.24.5.1342

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136. https://doi.org/10.1016/0010-0285(80)90005-5

van Moorselaar, D., Daneshtalab, N., & Slagter, H. A. (2021). *Neural mechanisms underlying distractor inhibition on the basis of feature and/or spatial expectations* (p. 2020.04.05.026070). bioRxiv. https://doi.org/10.1101/2020.04.05.026070

van Moorselaar, D., & 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

Wang, B., & Theeuwes, J. (2018). 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., 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

Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, *4*(12), 11. https://doi.org/10.1167/4.12.11

Won, B.-Y., & Geng, J. J. (2020). Passive exposure attenuates distraction during visual search. *Journal of Experimental Psychology. General*, *149*(10), 1987–1995. https://doi.org/10.1037/xge0000760

Yantis, S., & Jonides, J. (1984). Abrupt visual onsets and selective attention: Evidence from visual search. *Journal of Experimental Psychology. Human Perception and Performance*, *10*(5), 601–621. https://doi.org/10.1037//0096-1523.10.5.601

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*(7192), Article 7192. https://doi.org/10.1038/nature06860