Title: Memory precision for salient distractors decreases with learned suppression

Bo-Yeong Won^{1*}, Aditi Venkatesh¹, Phillip P. Witkowski^{1,2}, Timothy Banh³, Joy J. Geng^{1,2*}

1. Center for Mind and Brain, University of California Davis
2. Department of Psychology, University of California Davis
3. Department of Radiology, University of California San Francisco

*address correspondence to:

Bo-Yeong Won, bywon@ucdavis.edu

Joy J. Geng, jgeng@ucdavis.edu

Acknowledgements: This work was supported by NIH R01 MH113855-01 to JJG.

Abstract

Attention operates as a cognitive gate that selects sensory information for entry into memory and awareness (Driver, 2001). Under many circumstances, the selected information is task-relevant and important to remember, but it is also well known that perceptually salient non-target objects will capture attention and induce awareness even when they are irrelevant (Adams and Gaspelin, 2020). Recent studies have shown that task interference by salient distractors diminishes after repeated exposures, but the relationship between suppression and memory remain unclear. If attentional suppression (indicated by a reduced capture costs) occurs at the sensory level and prevents readout to other cognitive processes, one might expect memory and awareness to diminish commensurate with improved suppression. Here we test this hypothesis by measuring memory precision, awareness, and confidence of salient non-targets over repetitions as capture costs are reduced. Our results show that better attentional suppression is accompanied by reductions in memory precision and awareness, suggesting that reductions in capture costs reflect a reduction in the likelihood of the distractor entering memory and awareness.

Introduction

In order to accomplish goal-oriented behaviors, it is necessary to suppress distractions. Failures to do so will delay task completion or may even derail it completely. While it is clear that distractor suppression is integral to goal-oriented behaviors, how this is accomplished and how this affects downstream cognitive processes is still poorly understood. All models of attention posit that attention operates as a gating mechanism in which selected objects are processed more deeply and have a greater likelihood of entry into memory and awareness (Driver, 2001). In these studies, we test the attentional gating hypothesis for distractor suppression and ask if memory and awareness for salient distractors are related to the initial strength of attentional capture (or strength of suppression) over time.

The bulk of research on attentional mechanisms has focused on target selection, and it has been debated whether suppression occurs outside of target selection. It is becoming increasingly clear, however, that suppression of task-irrelevant information can occur independently from target selection (Chang & Egeth, 2019) and is likely supported by multiple cognitive mechanisms (Chelazzi, Marini, Pascucci, & Turatto, 2019; Fang, Becker, & Liu, 2019; Gaspelin & Luck, 2018b; Geng, Won, & Carlisle, 2019; Moher & Egeth, 2012; Noonan, Crittenden, Jensen, & Stokes, 2018). A number of these studies have found that suppression of both spatial and non-spatial distractor features is particularly effective when properties of task-irrelevant stimuli are predictable (Chetverikov, Campana, & Kristjansson, 2017; Ferrante et al., 2018; Gaspelin & Luck, 2018a; Geng & Diquattro, 2010; Geyer, Muller, & Krummenacher, 2006; Stilwell & Vecera, 2019; van Moorselaar & Slagter, 2019; Vatterott, Mozer, & Vecera, 2018; Vatterott & Vecera, 2012; Wang & Theeuwes, 2018). For example, Vatterot and Vecera (2012) found that distractor interference decreased when a salient color singleton distractor

appeared repeatedly, but attentional capture returned when the color changed. This rebound did not occur, however, when the color of the singleton distractor was expected to vary (Vatterott et al., 2018; Won, Kosoyan, & Geng, 2019). These results suggest remarkable flexibility in using feature-specific properties or abstract rules to actively suppress distractors.

Although task-irrelevant, but salient, stimuli are more likely to capture attention, the relationship between involuntary attentional capture and subsequent attentional engagement, memory, and awareness is still debated (Zivony & Lamy, 2016, 2018). Two recent studies, however, have shown a relationship between perceptual salience, attention, and awareness (Adams & Gaspelin, 2020; Constant & Liesefeld, 2020). Constant and Liesefeld (2020) used a parametric manipulation of salience and found a monotonic relationship between bottom-up saliency and the probability of memory for the item, although interestingly, the precision of the memory was not affected by saliency. Relatedly, Adams and Gaspelin (2019) found that participants were more likely to report awareness of a color singleton distractor from visual search trials with larger RT capture effects (Beloposky et al., 2008). These results lead to the prediction that the opposite should also be true: memory precision and awareness should decline as suppression improves. Such a result would be consistent with recent findings that reductions in distractor interference are related to mechanisms within visual cortex that prevent the “readout” of attentional priority signals to later stages of processing (Adam & Serences, 2020; Birman & Gardner, 2019; Won, Forloines, Zhou, & Geng, 2020).

The purpose of the current studies was to determine if learned suppression of salient distractors during visual search produces a concomitant reduction in subsequent memory probability and precision for the salient distractor. Measuring the outcome of distractor processing can be difficult because it requires participants to report on a feature that they are

supposed to ignore. Previous studies circumvented this issue by measuring residual spatial attention as a proxy for which objects were attended (Gaspelin, Leonard, & Luck, 2015; Kim & Cave, 1995; Won et al., 2019). While this is an elegant method, it does not provide a direct measurement of attention to specific stimulus features. In the present study, we directly measure the consequence of attentional suppression of color singleton distractors as a function of repeated exposure. We do so by harnessing large numbers of participants engaging in a one-trial memory probe following a varying number of trials of color singleton distractors (Mack & Rock, 1998; Won et al., 2019). Our results provide direct evidence for a relationship between learned distractor suppression and memory and awareness of the distractor.

Experiment 1A

Method

Participants. Four hundred twenty-eight subjects recruited from UC Davis participated in an online experiment (testable.org) for course credit through SONA. One hundred eight subjects (25.2% of subjects) were excluded from analyses due to performance lower than 80% accuracy on the visual search trials. The high number of individuals with poor performance is likely due to the fact that this was an online experiment administered through SONA and course credit was not tied to performance. This resulted in a total of 320 participants (mean age = 20.6, SD = 2.4, female = 251, left-handed = 30). Forty participants were randomly assigned to each of eight groups. Each group consisted of two sub-groups to counterbalance the two *critical singleton* colors, but no differences in color were expected, and the data were collapsed to increase statistical power. The sample size of N=20 for each subgroup was determined based on a previous study from which we adopted the experimental design (Won et al., 2019). A power

analyses was conducted based on differences between group means in RT for the high frequency singleton and the low frequency singleton conditions in Experiment 1 of Won et al., 2019 which had an effect size of $d=1.2$, and an alpha of .05. Results showed that a sample of 16 participants was required to achieve a power of .80. Our final sample size of 20 participants per subgroup (40 per condition of interest after collapsing between singleton colors, which was expected to be a manipulation of non-interest) was used to buffer against additional noise expected from online experiments and the single-trial approach. All participants had normal or corrected-to-normal vision and provided informed consent in accordance with NIH guidelines provided through the UCD Institutional Review Board.

Stimuli and apparatus

Search displays: Search displays were generated from Matlab 2019a (MathWorks). Search displays contained six shapes drawn in gray, five diamonds, ($1.7^\circ \times 1.7^\circ$) and one circle (1.5° in diameter) or one diamond and five circles on a black background. The eccentricity (center of each item to the center of the screen) was 4° based on an estimated viewing distance of 60 cm. Each shape contained a black line subtending $0.1^\circ \times 0.5^\circ$ that was randomly tilted 45° to the left or right (Figure 1A). Note that the stimulus size and color might vary depending on the participant's environment (e.g., monitor, video card, screen specifications, and lighting conditions). The target was defined by its shape and appeared equally often in six positions and was always gray. The other five positions were occupied by a non-target. On 6 out of the total 30 trials, all non-targets were an identical gray color to the target. On 24 out of 30 trials, one distractor was a unique color. We refer to this salient distractor as the color "singleton." The singleton color was different on each trial it was present; the colors were equidistant and selected

from CIE Lab color wheel (radius: 39, luminance: 70, a=0, b=0). The critical singleton was always one of two colors, counterbalanced across participants: RGB for the pink color = [218;148;208] and RGB for the green color = [110;187;134]. Two colors were used to ensure that our results were not due to spurious effects associated with one color and to minimize color variation due to variability in monitor settings and differences in color perception.

Probe displays: The set of probe displays consisted of four questions on three consecutive displays. In the first display, two questions were shown, “Was there a colored item in the previous display (Yes or No)?” and “How confident are you in your answer (10 is very confident and 1 is not confident at all)?” Regardless of their answers, a second display containing the question, “What color was the colored item?” and a color wheel that consisted of 24 colors (15 degrees apart) was shown. Finally, a location probe appeared with the question, “Where was the colored item?” with six alternative location choices (Figure 1B). The color probes were of primary interest given the singleton was defined by color saliency. The location probes were included for exploratory analyses. The location probe display always came last and were temporally distant from the actual search display. We therefore expected the results to be noisier than expected for a targeted working memory assessment of location. For these reasons, the data are only reported in the Supplemental Materials.

Figure 1. Search and probe displays. **A.** An example of a singleton-present trial. Participants were asked to find the odd shape and indicate the orientation of the bar inside with a manual

response. In this example, the target is the circle, and the bar is tilted to the right. **B.** After a fixed number of trials (determined by the *critical singleton group*), participants encountered a set of surprise probe displays that assessed memory and awareness of the preceding color singleton.

Design and procedure. We collected data through Testable (testable.org), an online experiment platform. The experiment began with instructions and four practice trials followed by the main experiment. The main experiment consisted of 30 search trials (24 singleton-present trials and six singleton-absent trials presented in random order) and two sets of probe displays. The probe displays were introduced without any instruction. The first set of probe displays was inserted after a specific number of singleton distractor trials that differed between groups. The *critical singleton trial* refers to the search trial that immediately preceded the probe displays. There were eight *critical singleton* groups: the *trial-1* group saw the probe trials after the first trial with a singleton; the *trial-3* group saw probe trials after the third singleton trial. The *critical singleton trial* in each successive group occurred in increasing intervals of three trials. The last *critical singleton* group, *trial-21*, saw the memory probe on the 21st singleton trial (Figure 2). These probe displays were presented without prior instruction and were, therefore, a “surprise” to the participant. All groups also experienced a second memory probe on the 24th trial as a control memory probe. However, there were no systematic differences of interest between groups, therefore the data were not included for report. Each search display remained visible until participants made a response and was followed by 200-ms of visual feedback (“correct” or “incorrect”). The trial ended with a 500-ms blank screen intertrial interval. We excluded critical search trials longer than 10-sec RT from analyses. This liberal criterion was applied to only eliminate participants that were clearly disengaged from the task without unfairly excluding data from slower, more deliberate, participants. This criterion resulted in elimination of two participants.

Analyses. Data from each of the dependent variables were analyzed across critical singleton groups, defined by the number of singleton-present trials experienced before the *critical singleton trial* and the set of probe displays. In each case, we expected the data to show a monotonic change over groups until an asymptote was reached. For example, the RT data were expected to decrease with the number of experienced singleton distractors, indicating improved suppression, until a plateau was reached. In order to capture the expected linear change followed by a plateau, the data from each of the dependent measures were fitted using a two-part spline model. Spline models estimate multiple parameters of non-linear data and include a “knot” point, which indicates the point at which monotonically decreasing data reaches a statistical asymptote (Cudeck and Klebe, 2002). A spline model fit with no knot point indicates that the data are fit by a single continuous line (see also Geng, DiQuattro, and Helm, 2017). To determine the best inflection point, we fit spline models with no knot point and knot points at different trial positions. The best model was determined by the one with the lowest BIC value, and model comparison was calculated with Bayes Factors (BF) (Wagenmakers, 2007). The BF we calculated reflects *evidence in favor of the best fitting model* (Table 1). As a heuristic, BF values less than 3 are considered weak evidence, between 3-10 as moderate evidence, up to 30 as strong evidence, up to 100 as very strong evidence, and above 100 is considered extreme evidence (Stefan, Gronau, Schonbrodt, & Wagenmakers, 2019).

Experiment 1A

8 groups (N=40 each)

Figure 2. Trial procedure in Experiment 1A. Each group consisted of 40 participants. Participants in all eight groups saw the same visual search and probe displays. The only difference between them was the trial position of the surprise set of memory and awareness probe displays. The singleton-present trial immediately preceding the probe displays is referred to as the *critical singleton trial*.

Experiment 1B

The purpose of Experiment 1A was to test if repeated singleton distractors lead to better attentional suppression and decrease in memory and awareness. However, since *critical singleton trials* occurring after more exposures always occurred later in the experiment, it is possible that observed effects are due to trial order (e.g., practice effects or fatigue) and not singleton exposure. To control for this possibility, Experiment 1B was identical to 1A except that all singleton-present trials, save the *critical singleton* trial and the last trial, were replaced with singleton-absent trials. Thus, the *critical singleton trial* was now the first singleton distractor trial in all groups, but the trial number of the *critical singleton* still differed across groups as before. If exposure to singleton distractors is necessary for learned suppression, then there should be no evidence of suppression as a function of trial sequence in this experiment. However, if the

previous results were due to order effects, the results should be identical to those of Experiment 1A despite the elimination of recurring singleton distractor trials.

Participants. An independent sample of 381 undergraduates recruited from UC Davis participated in an online experiment (testable.org) for course credit through SONA, but 61 subjects were excluded (16.0% of subjects) due to the lower than 80% of accuracy, which led to the total of 320 undergraduates (mean age = 19.6, SD = 1.9, female = 272, left-handed = 23). Forty participants were randomly assigned to each of eight groups. All participants had normal or corrected-to-normal vision and provided informed consent in accordance with NIH guidelines provided through the UCD Institutional Review Board.

Stimuli, apparatus, design, and procedure. All aspects of Experiment 1B were identical to Experiment 1A except that the 30 trials consisted of 28 singleton-absent trials and there was only one *critical singleton trial*. The positions of the *critical singleton trials* for eight groups were identical with those in Experiment 1A (Figure 3).

Experiment 1B 8 groups (N=40 each)

Figure 3. Trial procedure in Experiment 1B. All trials were singleton-absent trials except for the *critical singleton trial* and the last trial (i.e., the 2nd singleton). The trial positions of *critical singletons* and probe displays were identical with those in Experiment 1A.

Results: Experiments 1A and 1B

Search RTs on Critical Singleton Trials

To estimate the time course of learned suppression, we analyzed search RT as a function of the *critical singleton* group (Figure 4) using a two-part spline model (see Methods). The spline model with the lowest BIC for Experiment 1A was Model 4, which had a knot point set at the *trial-9* group. The model showed that attentional capture decreased linearly until about the 9th singleton repetition, after which RT reached asymptote, indicating the reduction in singleton interference reached its maximum (Figure 4A, Table 1). In contrast, the linear regression for data from Experiment 1B was the best fit by a model without a knot point, suggesting no consistent inflection point in performance as a function of when the singleton trial first appeared (Figure 4B, Table 1).

Table 1: Spline model fits for data in Experiments 1A and 1B. Values are model BIC and in parentheses are BF values comparing the best fitting model to each other model. BF values indicate evidence in favor of the best fitting model. * = best fitting model						
	Exp 1A Search RT	Exp 1B Search RT	Exp 1A Color memory	Exp 1B Color memory	Exp 1A Awareness confidence	Exp 1B Awareness confidence
Model 1: linear regression	5248.93 (BF ₄₁ =7.57)	5222.52*	1649.24 (BF ₂₁ =244.69)	1465.56*	1404.24 (BF ₅₁ =1.34)	1272.74*
Model 2: knot at trial 3	5248.97 (BF ₄₂ =7.73)	5228.27 (BF ₁₂ =17.73)	1638.24*	1470.54 (BF ₁₂ =12.06)	1407.664 (BF ₅₂ =7.40)	1276.34 (BF ₁₂ =6.05)
Model 3: knot at trial 6	5247.88 (BF ₄₃ =4.48)	5228.06 (BF ₁₃ =15.96)	1640.93 (BF ₂₃ =3.84)	1470.71 (BF ₁₃ =13.13)	1405.21 (BF ₅₃ =2.18)	1277.48 (BF ₁₃ =10.70)
Model 4: knot at trial 9	5244.88*	5227.58 (BF ₁₄ =12.55)	1644.43 (BF ₂₄ =22.08)	1471.32 (BF ₁₄ =17.81)	1403.85 (BF ₅₄ =1.10)	1276.38 (BF ₁₄ =6.17)
Model 5: knot at trial 12	5246.42 (BF ₄₅ =2.16)	5227.97 (BF ₁₅ =15.25)	1646.72 (BF ₂₅ =69.41)	1471.29 (BF ₁₅ =17.55)	1403.66*	1274.57 (BF ₁₅ =2.50)
Model 6: knot at trial 15	5247.73 (BF ₄₆ =4.16)	5228.26 (BF ₁₆ =17.63)	1649.55 (BF ₂₆ =285.72)	1471.33 (BF ₁₆ =17.90)	1405.23 (BF ₅₆ =2.19)	1274.9 (BF ₁₆ =2.94)
Model 7: knot at trial 18	5249.44 (BF ₄₇ =9.77)	5228.19 (BF ₁₇ =17.03)	1648.15 (BF ₂₇ =141.88)	1471.20 (BF ₁₇ =16.78)	1406.25 (BF ₅₇ =3.64)	1276.4 (BF ₁₇ =6.23)

Figure 4. Mean search RT for the *critical singleton trial* across eight groups. A. Experiment 1A. The “*” indicates the knot point of the best fitting model. B. Experiment 1B. Error bars indicate ± 1 standard error of the mean.

Color memory probe for the critical singleton

Having confirmed that singleton distractors were increasingly suppressed over repetitions, we turn next to our main research question regarding the consequence of attentional

suppression on memory. We quantified memory for the color of the singleton distractor by calculating the absolute circular distance (error) between the actual singleton color and the color wheel responses on the surprise memory probe trial. Larger values indicate greater deviations from the true singleton color and, therefore, less precise memories.

Similar to the search RT results, we fitted spline models with varying knot points to memory performance for the color memory probe trial from Experiments 1A and 1B (Figure 5). The results show that in Experiment 1A, Model 2 had the lowest BIC with a knot point at the *critical singleton trial-3* group: memory error increased steeply between the first and the third singleton trials and then plateaued (Figure 5A, Table 1). This suggests that by the time the third singleton distractor was experienced, its color information was already filtered from memory. The fact that memory precision for the singleton distractor reached asymptote more rapidly than the search RT (i.e., by the 9th singleton) suggests partial suppression is sufficient to reduce memory for the color of singleton. In contrast, the model with no knot point was the best fit for the color memory data in Experiment 1B (Figure 5B, Table 1), showing that memory remained similarly precise irrespective of when the first singleton was seen within the sequence of trials.

Figure 5. Memory for the color of the singleton distractor. A. Experiment 1A data from color memory probes of the critical singleton distractor across eight groups. The “*” indicates the knot point of the best fitting model. B. Experiment 1B data of color memory probes. The *critical*

singleton was the first singleton encountered in all eight groups even though they appeared in a different trial position within the stream of search trials. Error bars indicate ± 1 standard error of the mean.

Awareness and confidence ratings

Next, we assessed awareness and confidence of the color singleton's presence.

Confidence ratings were only analyzed for participants who answered "Yes" to whether or not they had seen a color singleton. First, visual inspection of the data suggests that awareness of the singleton was uniformly high in both experiments, 91.9% (N=294) in Experiment 1A and 97.8% (N=313) in Experiment 1B. This suggests that subjects were overall aware of the presence of singleton distractors when they occurred, although it is worth noting that responses may be biased in Experiment 1 because of the high frequency of singleton distractors – i.e., participants may have tended to guess that a singleton was present if they were not sure.

Turning to the confidence data, the pattern was similar to RT (compare Figures 4 and 6), showing a clear decrease in confidence that a singleton was seen on later critical singleton trials in Experiment 1A. No such pattern was seen in Experiment 1B. This observation was verified by the spline models. The best fitting model was Model 5 with a knot point at *trial-12*, but there was only very weak evidence for this model compared to the model with a knot point at critical singleton *trial-9* (Table 1). This suggests that at most, by the time a color singleton was seen for the 12th time, confidence in having seen the singleton had depreciated to an asymptotic level. This pattern mirrors that of RT and suggests that as attentional suppression increased, there was a concomitant decrease in participant confidence in having seen the preceding color singleton distractor.

Figure 6. Awareness of color singleton and confidence rating. The proportion of people who said “Yes” to the question, “Was there a colored item in the previous display?” and confidence rating across eight groups. AC. Experiment 1A. The “*” indicates the knot point of the best fitting model. BD. Experiment 1B.

Experiment 2

In Experiment 1, we found that learned suppression of salient distractors during visual search produced reduced memory precision for the color of singleton and confidence that a color singleton distractor had appeared on the previous trial. The results suggested that memory precision for the singleton declined as attentional suppression improved, however, our estimates of precision were limited by the relatively coarse color wheel and small sample size given the necessity of the one-trial surprise memory probe. In Experiment 2, we increased the sample size to 200 per group and used a more fine-grained continuous measurement of color to more precisely measure the distribution of responses for the singleton memory probe. Moreover, we

chose colors that were within a single category (blue, orange) but shifted off the focal color in order to test if working memory biases towards the category center decay with learned suppression {Hardman, 2017 #2691}{Bae, 2015 #204}.

Method

Participants. Five hundred seventy-seven undergraduates from UC Davis participated in Experiment 2 for course credit. We excluded 177 subjects (30.7%) due to lower than 80% accuracy or if the mouse click response was in an area outside the color and location wheels entirely. In total we had 400 participants (mean age = 20.16, SD = 2.73, female= 281, left-handed = 30). The experiment was run online on the Testable platform (testable.org). Two hundred participants were randomly assigned to each of two *critical singleton* groups. All participants had normal or corrected-to-normal vision and provided informed consent in accordance with NIH guidelines provided through the UCD Institutional Review Board.

Apparatus and Stimuli. All apparatus and stimuli were identical with Experiment 1 except that a color wheel on color memory probes now consisted of 72 colors; location memory was measured using a gray location wheel (Figure 7). Also, 24 equidistant singleton colors were chosen from 72 CIE Lab color wheel and we chose two critical singleton colors (RGB for the orange critical singleton: [225;153; 118]; RGB for the blue critical singleton color: [47;190;194]). These colors were approximately 10 and 35 degrees from their category focal points, respectively. The distance of the blue singleton was farther because the blue category is wider (Bae, Olkkonen, Allred, Flombaum, 2015). Both critical singleton colors were boundary colors located in the leftward (counterclockwise) from the prototypical category colors. If a bias occurs in working

memory report of the color towards the category center, the memory for the singleton color should be biased rightward (clockwise) in our measurement.

Figure 7. Search and probe display in Experiment 2. A. Visual search and B. probe displays were identical with those in Experiment 1 except that a color wheel that consisted of 72 colors and a gray location wheel were used for memory probes. Participants were asked to click the singleton color and singleton location on a color wheel and a location wheel, respectively.

Design and Procedure. Design and procedure were identical with Experiment 1 except that there were only two *critical singleton* groups (*trial-1* and *trial-12*). The *trial-1* group saw the set of probe displays after the 1st singleton trial and the other group after the 12th singleton trial. The 1st and 12th trials were chosen because they fall on the two sides of the knot-points found in Experiment 1 such that we expected maximum capture by the 1st singleton and asymptotic suppression by the 12th singleton.

Model fitting analysis. We tested whether the precision of memory for the critical singleton differed between *critical singleton* groups by fitting response data from the memory probe trials to a “mixture model” of working memory (Zhang and Luck, 2008). The model assumes that there are two sources of error in the representation of the item in memory, and that these errors can be modeled by two separate underlying distributions: the first is a von Mises distribution, with a concentration parameter (κ) that reflects the precision of information held in working

memory; the second component is a uniform distribution which captures “guess” responses thought to result from information lost from working memory entirely. Thus, the mixture model dissociates response errors from imperfect memories and those from “guessing”.

We fit this model to the color memory probe data using custom python code with the pymc3 package (Liew et al., 2019). We supplied priors for two free parameters in the model, the center (μ) and concentration (κ) of the von Mises distribution. We used a prior distribution for μ that was itself a von Mises distribution with center on 0 and precision of .05 ($\mu=0$, $\kappa=.05$). For κ , we assumed a uniform distribution between 0 and 1000. κ is bounded on the lower end by 0 representing complete imprecision. The uniform distribution was set between negative π and π , representing an equal likelihood of all responses during guessing.

Results

Critical singleton search RTs

We excluded trials with RTs over 10-sec from analyses, which was .5% of data. We conducted an independent t-test to compare RT between *critical singleton* groups. RT in the *trial-12* group was shorter than the *trial-1* group, $t(294.29) = 10.35$, $p < .001$, Cohen’s $d = 1.04$, $BF > 100$, which is a replication of the suppression effect in Experiment 1 (Figure 9).

Color memory distributions for critical singleton color

To test whether the error in color response distribution differed between groups, we looked at the posterior distribution of the concentration parameter (κ) estimated from fitting each model. Figure 8 shows the distribution and marks the Highest Density Interval (HDI) for each group, which describes the 95% most credible values for the parameter (Kruschke, 2018).

The fitted models estimated the value of kappa to be 5.96 (HDI = 4.58,7.46) for the *trial-1* group and 1.99 (HDI = 0.59, 4.28) for the *trial-12* group. The fact that these intervals are non-overlapping suggests that the distribution for the *trial-1* group was more precise than that of the *trial-12* group. This indicates that while memory representations were relatively precise for the first color singleton distractor, representations of the 12th singleton distractor were poorer.

Similarly, we tested whether the guessing rate differed between groups. The results showed that the guessing rate estimated by the model was higher in the *trial-12* (Mean=0.43, HDI = 0.31,0.56) group compared to the *trial-1* group (Mean=0.2, HDI = 0.15,0.26), indicating that the *trial-12* group was more likely to randomly guess the color of the singleton distractor.

Figure 8: A. Best fit working memory mixture model for the *trial-1* group (blue), overlaid on the distribution of actual responses. Distance from the true singleton color value is in color degrees. The Center parameter reflects the central tendency of the distribution. The rightward

shift is towards the color category center, suggesting a center-bias in memory. Kappa reflects the precision of the memory where larger numbers indicate greater precision. Guess indicates the guessing rate where larger values indicate more guessing about the singleton color. B. Same as A but for the *trial-12* group (orange). C. Distribution of posterior estimates of the concentration parameter (kappa) for each group. Black bars at the bottom indicate the 95% HDI for each group. Precision was significantly greater for the *trial-1* group compared to the *trial-12* group.

Finally, it is worth noting that the center values of the data were both right-shifted although clearly more so for the *trial-1* group (stats?). Recall that colors were between the color category center (i.e., defined by Bae et al., 2015) and a color boundary. The orange and blue singleton colors were approximately 10 and 35 degrees from the category center, respectively. The shift therefore in the *trial-1* group represents a bias to report colors closer to the category center than the true color and replicates previous research {Hardman, 2017 #2691}{Bae, 2015 #204}, but this effect dissipated in the *trial-12* group.

Awareness assessment and confidence ratings

Overall awareness 95.5% for the *trial-1* group and 83.0% for the *trial-12* group (Figure 9) again indicating relatively high reports of awareness in both groups. The *trial-1* group produced higher confidence in their awareness response than the *trial-12* group, $t(307.36) = 7.298$, $p < .001$, Cohen's $d = .79$, $BF > 100$. These data replicate those from Experiment 1A, confirming that while awareness was relatively high in both Groups, there was an erosion of confidence in awareness when attentional suppression improved after 12 repetitions of singleton distractors.

Figure 9. A. Mean search RT for the *critical singleton trial* between the *trial-1* group (blue circle) and the *trial-12* group (orange circle). B. Awareness proportion. C. Confidence ratings. Error bars indicate ± 1 standard error of the mean.

General Discussion

There has been substantial evidence over the last decade that salient distractors can be suppressed during visual search, particularly when the distractor features recur and are predictable (Chelazzi et al., 2019; Gaspelin & Luck, 2018b; Geng et al., 2019; Noonan et al., 2018). Recent work has provided further evidence that suppression is related to participant awareness of the salient distractor (Adams and Gaspelin, 2020; Adams and Gaspelin, 2021; Constant and Lisefield, 2020)(Won et al., 2019). However, the relationship between attentional suppression of salient distractors and memory for their features and confidence in awareness has not yet been assessed. In other words, while previous studies such as Adams and Gaspelin (2020, 2021) assessed awareness of capture itself, our current study assess the downstream cognitive consequences of capture (for early probe groups) and learned suppression (late probe groups). In

two experiments, we used a “one-shot” memory probe paradigm to assess the quality of memory and awareness for a color singleton distractor on an immediately preceding visual search trial.

The visual search paradigm we used was based on Won et al., (2019) and the RT results replicated the previous pattern: interference from a color singleton distractor declined over repetitions even though the exact color and location of the singleton could not be predicted. Furthermore, in this sample, we found that RT reached asymptote by the time the color singleton appeared for the ninth time. While we do not expect that asymptote will always be reached following a fixed number of exposures, the results suggest that learned suppression operates rapidly when the expected frequency of color singletons is high. These results are consistent with a number of studies in the literature showing that distractor suppression improves with exposure (Awh, Matsukura, & Serences, 2003; Noonan et al., 2016; Turatto, Bonetti, Pascucci, & Chelazzi, 2018; Vatterott et al., 2018) and learning occurs rapidly (Vatterott et al., 2018; Vatterott & Vecera, 2012). While the specific mechanisms are still debated, it appears that learned suppression for distractor features operates in visual cortex, directly attenuating the attend-to-me signal and preventing attentional capture (Adam & Serences, 2020; van Moorselaar, Lampers, Cordesius, & Slagter, 2020; Won et al., 2020). Such results suggest that learned suppression attenuates the readout of visual information about the distractor. If this is true, then we would expect memory representations of the distractor color and awareness of the singleton to decline with increased attentional suppression over exposures.

The current experiments tested this directly by introducing a “one-shot” memory probe following visual search trials with different numbers of preceding singleton trials. Because it is impossible to ask subjects to report on distractors more than once without the distractors becoming “task-relevant”, we chose to recruit many participants and acquire responses from only

one critical trial per person. Consistent with the notion that learned suppression operates during visual processing, we found that participants who saw the memory probe after a greater number of exposures had poorer memory representations for the color of the distractor singleton and they had lower confidence in their awareness report of the singleton being present. This finding from Experiment 1 was bolstered in Experiment 2 using the mixture model to formally estimate memory precision and guess rates: memories after the 12th singleton was significantly less precise than after the first singleton, and guess rates were significantly higher. Notably, the expected response bias towards the category center {Hardman, 2017 #2691}{Bae, 2015 #204}, was more pronounced after *trial-1* than *trial-12*, suggesting that the category bias is stronger when memories are relatively precise.

Although memory degraded with increased learned suppression, we found that the patterns of memory degradation and RT were not identical in Experiment 1. The maximum degradation in memory from Experiment 1 occurred rapidly, reaching asymptote after just three experiences with the singleton distractor. The fact that memory performance reached asymptote so early, compared to RT declines, suggests that as soon as a visually salient object is tagged as being “task-irrelevant”, attentive processing of that stimulus is attenuated, reducing the likelihood of its feature information being stored in working memory and entering awareness.

In contrast to the decay in memory precision, which occurred early and was sustained, confidence in awareness decayed more slowly over time. The confidence results were more similar to RT, suggesting there might be a closer link between the strength of suppression and confidence in having seen a distractor at all. In these experiments, confidence appeared to track RT costs better than awareness, but the results were based on a single trial per participant and therefore may not have had elicited the range of responses possible from repeated measures

(Adams & Gaspelin, 2020). Furthermore, the binary awareness question may not have been a good assessment of actual awareness from the previous trial since most trials had a color singleton and participants may have used that to “guess” that a singleton was present.

In sum, our results provide evidence that the attentional suppression operates as a gating mechanism that reduces the likelihood of task-irrelevant information being readout into memory and awareness. This “gating” function that prevents further processing of distractors appears analogous and opposite to effects of target selection in which processing of sensory features that match the target template is facilitated and pass through into memory and awareness. Future work is needed to map the exact relationship between learned suppression, memory, and awareness, but these data contribute to a better understanding of how learned suppression prevents sensory readout and affects information representations in memory and awareness.

Open Practices Statement

The data for all experiments are available for download at OSF (osf.io/; the full link will be made available upon acceptance for publication). None of the experiments were preregistered.

References

- Adam, K. C. S., & Serences, J. T. (2020). History-driven modulations of population codes in early visual cortex during visual search. *bioarxiv*.
- Adams, O. J., & Gaspelin, N. (2020). Assessing introspective awareness of attention capture. *Atten Percept Psychophys*, 82(4), 1586-1598. doi:10.3758/s13414-019-01936-9
- Awh, E., Matsukura, M., & Serences, J. T. (2003). Top-down control over biased competition during covert spatial orienting. *Journal of Experimental Psychology: Human Perception and Performance*, 29(1), 52-63. doi:10.1037/0096-1523.29.1.52

- Birman, D., & Gardner, J. L. (2019). A flexible readout mechanism of human sensory representations. *Nat Commun*, 10(1), 3500. doi:10.1038/s41467-019-11448-7
- Chang, S., & Egeth, H. E. (2019). Enhancement and suppression flexibly guide attention. *Psychological Science*.
- Chelazzi, L., Marini, F., Pascucci, D., & Turatto, M. (2019). Getting Rid of Visual Distractors: The Why, When, How and Where. *Current Opinion in Psychology*. doi:10.1016/j.copsyc.2019.02.004
- Chetverikov, A., Campana, G., & Kristjansson, A. (2017). Representing color ensembles. *Psychological Science*.
- Constant, M., & Liesefeld, H. R. (2020). Massive effects of saliency on information processing in visual working memory. *Psychological Science*.
- Driver, J. (2001). A selective review of selective attention research from the past century. *British Journal of Psychology*, 92, 53-78.
- Fang, M. W. H., Becker, M. W., & Liu, T. (2019). Attention to colors induces surround suppression at category boundaries. *Sci Rep*, 9(1), 1443. doi:10.1038/s41598-018-37610-7
- 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. doi:10.1016/j.cortex.2017.09.027
- Gaspelin, N., Leonard, C. J., & Luck, S. J. (2015). Direct Evidence for Active Suppression of Salient-but-Irrelevant Sensory Inputs. *Psychol Sci*, 26(11), 1740-1750. doi:10.1177/0956797615597913
- Gaspelin, N., & Luck, S. J. (2018a). Inhibition as a potential resolution to the attentional capture debate. *Curr Opin Psychol*, 29, 12-18. doi:10.1016/j.copsyc.2018.10.013
- Gaspelin, N., & Luck, S. J. (2018b). The Role of Inhibition in Avoiding Distraction by Salient Stimuli. *Trends Cogn Sci*, 22(1), 79-92. doi:10.1016/j.tics.2017.11.001
- Geng, J. J., & Diquattro, N. E. (2010). Attentional capture by a perceptually salient non-target facilitates target processing through inhibition and rapid rejection. *J Vis*, 10(6), 5. doi:10.1167/10.6.5
- Geng, J. J., Won, B.-Y., & Carlisle, N. B. (2019). Distractor ignoring: Strategies, learning, and passive filtering. *Current Directions in Psychological Science*.
- Geyer, T., Muller, 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.
- Kim, M.-S., & Cave, K. R. (1995). Spatial attention in visual search for features and feature conjunctions. *Psychological Science*(6), 376-380.
- Mack, A., & Rock, I. (1998). *Inattentional Blindness*. Cambridge, MA: MIT Press.
- Moher, J., & Egeth, H. E. (2012). The ignoring paradox: cueing distractor features leads first to selection, then to inhibition of to-be-ignored items. *Atten Percept Psychophys*, 74(8), 1590-1605. doi:10.3758/s13414-012-0358-0
- Noonan, M. P., Adamian, N., Pike, A., Printzlau, F., Crittenden, B. M., & Stokes, M. G. (2016). Distinct Mechanisms for Distractor Suppression and Target Facilitation. *J Neurosci*, 36(6), 1797-1807. doi:10.1523/JNEUROSCI.2133-15.2016
- Noonan, M. P., Crittenden, B. M., Jensen, O., & Stokes, M. G. (2018). Selective inhibition of distracting input. *Behav Brain Res*, 355, 36-47. doi:10.1016/j.bbr.2017.10.010

- Stefan, A. M., Gronau, Q. F., Schonbrodt, F. D., & Wagenmakers, E. J. (2019). A tutorial on Bayes Factor Design Analysis using an informed prior. *Behav Res Methods*, 51(3), 1042-1058. doi:10.3758/s13428-018-01189-8
- Stilwell, B. T., & Vecera, S. P. (2019). Learned and cued distractor rejection for multiple features in visual search. *Atten Percept Psychophys*, 81(2), 359-376. doi:10.3758/s13414-018-1622-8
- Turatto, M., Bonetti, F., Pascucci, D., & Chelazzi, L. (2018). Desensitizing the attention system to distraction while idling: A new latent learning phenomenon in the visual attention domain. *Journal of Experimental Psychology: General*, 147(12), 1827-1850. doi:10.1037/xge0000503
- van Moorselaar, D., Lampers, E., Cordesius, E., & Slagter, H. A. (2020). Neural mechanisms underlying expectation-dependent inhibition of distracting information. *Elife*, 9. doi: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. *J Neurosci*, 39(35), 6953-6967. doi:10.1523/JNEUROSCI.0593-19.2019
- Vatterott, D. B., Mozer, M. C., & Vecera, S. P. (2018). Rejecting salient distractors: Generalization from experience. *Atten Percept Psychophys*, 80(2), 485-499. doi:10.3758/s13414-017-1465-8
- Vatterott, D. B., & Vecera, S. P. (2012). Experience-dependent attentional tuning of distractor rejection. *Psychon Bull Rev*, 19(5), 871-878. doi:10.3758/s13423-012-0280-4
- Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, 14(5), 779-804.
- Wang, B., & Theeuwes, J. (2018). Statistical regularities modulate attentional capture independent of search strategy. *Attention, Perception, & Psychophysics*, 80(7), 1763-1774. doi:10.3758/s13414-018-1562-3
- Won, B.-Y., Forloines, M., Zhou, Z., & Geng, J. J. (2020). Changes in visual cortical processing attenuate singleton distraction during visual search. *Cortex*, 132, 309-321. doi:10.1016/j.cortex.2020.08.025
- Won, B.-Y., Kosoyan, M., & Geng, J. J. (2019). Evidence for second-order singleton suppression based on probabilistic expectations. *J Exp Psychol Hum Percept Perform*, 45(1), 125-138. doi:10.1037/xhp0000594
- Zivony, A., & Lamy, D. (2016). Attentional capture and engagement during the attentional blink: A "camera" metaphor of attention. *J Exp Psychol Hum Percept Perform*, 42(11), 1886-1902. doi:10.1037/xhp0000286
- Zivony, A., & Lamy, D. (2018). Contingent attentional engagement: stimulus- and goal-driven capture have qualitatively different consequences. *Psychological Science*, 29(12), 1930-1941.

Supplementary Information

S1. Participant demographic information

Experiment 1A

trial-1 group: N=40, mean age=20.7 years, SD=3.0, 29 females, 4 left-handed; *trial-3 group*: N=40, mean age=20.0 years, SD=2.6, 30 females, 4 left-handed; *trial-6 group*: N=40, mean age=20.4 years, SD=1.6, 30 females, 4 left-handed; *trial-9 group*: N=40, mean age=20.5 years, SD=1.6, 28 females, 2 left-handed; *trial-12 group*: N=40, mean age=21.0 years, SD=2.9, 21 females, 2 left-handed; *trial-15 group*: N=40, mean age=20.2 years, SD=2.4, 24 females, 5 left-handed; *trial-18 group*: N=40, mean age=20.2 years, SD=1.7, 29 females, 2 left-handed; *trial-21 group*: N=40, mean age=21.1 years, SD=3.0, 28 females, 4 left-handed.

Experiment 1B

trial-1 group: N=40, mean age=19.9 years, SD=1.8, 36 females, 2 left-handed; *trial-3 group*: N=40, mean age=19.4 years, SD=1.8, 34 females, 6 left-handed; *trial-6 group*: N=40, mean age=20.1 years, SD=3.4, 34 females, 3 left-handed; *trial-9 group*: N=40, mean age=19.6 years, SD=1.0, 35 females, 4 left-handed; *trial-12 group*: N=40, mean age=19.5 years, SD=1.8, 32 females, 1 left-handed; *trial-15 group*: N=40, mean age=19.4 years, SD=1.3, 37 females, 6 left-handed; *trial-18 group*: N=40, mean age=19.6 years, SD=1.7, 31 females, 0 left-handed; *trial-21 group*: N=40, mean age=19.4 years, SD=1.1, 32 females, 1 left-handed.

Experiment 2

trial-1 group with orange critical singleton: N=100, mean age=20.33 years, SD=3.82, 78 females, 7 left-handed; *trial-12 group with orange critical singleton*: N=100, mean age=19.90

years, SD=1.97, 66 females, 11 left-handed; *trial-1 group with blue critical singleton*: N=100, mean age=20.57 years, SD=2.89, 68 females, 10 left-handed; *trial-12 group with blue critical singleton*: N=100, mean age=19.82 years, SD=1.66, 69 females, 2 left-handed.

S2. Overall search accuracy and RTs across groups

Table S1. Average search accuracy of eight groups in Experiment 1A and Experiment 1B, and two groups in Experiment 2 (%). The value in the parentheses is the standard deviation.

	<i>trial-1</i>	<i>trial-3</i>	<i>trial-6</i>	<i>trial-9</i>	<i>trial-12</i>	<i>trial-15</i>	<i>trial-18</i>	<i>trial-21</i>
Exp1A	96.3 (4.1)	95.3 (6.5)	96.0 (3.9)	95.1 (4.8)	96.3 (3.9)	96.5 (4.5)	96.5 (5.3)	95.4 (5.1)
Exp1B	95.3 (6.6)	94.3 (6.5)	94.6 (5.6)	95.3 (6.2)	97.1 (4.1)	96.0 (6.1)	95.9 (4.5)	95.7 (4.7)
Exp2	95.9 (5.07)				96.2 (4.93)			

(*Sub-group accuracy in Experiment 2: *trial-1 group with orange critical singleton* = 95.6%, SD = 4.8; *trial-12 group with orange critical singleton* = 96.5%, SD = 3.9; *trial-1 group with blue critical singleton* = 96.2%, SD = 5.3; *trial-12 group with blue critical singleton* = 95.9%, SD = 4.8)

Table S2. Average search RT of eight groups in Experiment 1A and Experiment 1B, and two groups in Experiment 2 (ms). The value in the parentheses is the standard deviation.

	<i>trial-1</i>	<i>trial-3</i>	<i>trial-6</i>	<i>trial-9</i>	<i>trial-12</i>	<i>trial-15</i>	<i>trial-18</i>	<i>trial-21</i>
--	----------------	----------------	----------------	----------------	-----------------	-----------------	-----------------	-----------------

Exp1A	1798.5 (573.7)	1813.2 (498.4)	1920.4 (573.6)	1649.2 (422.3)	1807.6 (624.5)	1613.6 (362.4)	1785.7 (492.6)	1704.0 (399.4)
Exp1B	1822.9 (878.0)	1659.4 (635.7)	1720.4 (656.3)	1625.3 (606.6)	1511.0 (346.2)	1504.7 (424.2)	1547.3 (383.7)	1343.9 (316.6)
Exp2	1840.5 (1173.7)	1777.0 (983.9)						

(*Sub-group RTs in Experiment 2: *trial-1* group with orange critical singleton = 1841.3, SD = 488.5; *trial-12* group with orange critical singleton = 1821.6; SD = 590.5; *trial-1* group with blue critical singleton = 1839.7, SD = 641.6; *trial-12* group with blue critical singleton = 1684.9, SD = 443.7)

S3. Location memory probe for the critical singleton

Figure S1. Memory for the location of the singleton distractor. Error bars indicate ± 1 standard error of the mean.

S4. Location memory distributions for critical singleton

We followed the same procedure for fitting response about the location of the critical singleton that we did with color responses. We used the same priors for all parameters. We found that the location responses for *trial-1* group (Mean = 43.45, HDI= 26.11, 63.35) were more precise compared to that of *trial-12* group (Mean = .38, HDI=.04, .7). This shows that subjects in *trial-1* group were better able to remember the location of the singleton distractor.

Figure S2. A. Best fit working memory mixture model for the *trial-1* group (blue), overlaid on the distribution of actual responses. B. Same as A but for the *trial-12* group (orange). C. Distribution of posterior estimates of the concentration parameter (kappa) for each group. Black bars at the bottom indicate the 95% HDI for each group.