Evidence Against the Low-Salience Account of Attentional Suppression

Manuscript in press at *JEP:HPP*

Brad T. Stilwell
Texas A&M University

Howard E. Egeth
Johns Hopkins University

Nicholas Gaspelin
University of Missouri

Do salient distractors have the power to automatically capture attention? This question has led to a heated debate concerning the role of salience in attentional control. A potential resolution, called the signal suppression hypothesis, has proposed that salient items produce a bottom-up signal that vies for attention, but that salient stimuli can be suppressed via top-down control to prevent the capture of attention. This hypothesis, however, has been criticized on the grounds that the distractors used in initial studies of support were weakly salient. It has been difficult to know how seriously to take this low-salience criticism because assertions about high and low salience were made in the absence of a common (or any) measure of salience. The current study used a recently developed psychophysical technique to compare the salience of distractors from two previous studies at the center of this debate. Surprisingly, we found that the original stimuli criticized as having low salience were, if anything, more salient than stimuli from the later studies that purported to increase salience. Follow-up experiments determined exactly why the original stimuli were more salient and tested whether further improving salience could cause attentional capture as predicted by the low-salience account. Ultimately, these findings challenge purely stimulus-driven accounts of attentional control.

Keywords: visual attention, attentional capture, distraction, suppression, salience

Public Significance Statement

This study aimed to understand whether highly salient stimuli have the power to automatically attract attention. Using a new method to measure salience, we found evidence that is inconsistent with theoretical accounts claiming that stimuli of sufficient salience cannot be ignored

Our visual systems have the crucial task of determining which of the many pieces of information available to us should be attended and which should be ignored. Salient stimuli, such as brightly colored objects, have proven particularly problematic for theories of attentional control. Here, salience refers to the degree by which an object differs from other objects in low-level features. For decades, researchers have debated whether salient stimuli can automatically capture attention even if they are unrelated to our current goals (see review by Luck et al., 2021). Recent evidence has suggested a potential resolution called the *signal suppression hypothesis*. According to this hypothesis, salient stimuli automatically vie for attention in a bottom-up manner, but top-down mechanisms can be used to suppress these salient distractors to prevent attentional capture (Gaspelin & Luck, 2018d, 2019).

Although the signal suppression hypothesis has garnered much support, a point of contention remains regarding the role of salience. It has been claimed that the salient distractors used in the original studies of signal suppression were only weakly salient (Wang & Theeuwes, 2020; see also Theeuwes, 2004). If the stimuli were made more salient, according to this criticism, they could not be suppressed and would capture attention. If true, this would suggest that attentional control is entirely bottom-up, at least when stimuli are highly salient. This claim, however, has been difficult to test because the field currently has no agreed-upon methods to measure salience. The current study will therefore use a recently developed technique to compare the salience of distractors from two previous studies at the center of this debate (Gaspelin et al., 2015; Wang & Theeuwes, 2020). To preview the results, we find that the critiques of the signal suppression hypothesis regarding low salience were unfounded.

The Attentional Capture Debate

Historically, there were two opposing theoretical positions concerning whether salient stimuli automatically capture attention. *Stimulus-driven accounts* posit that salient objects will automatically capture attention, even when they conflict with the current goals of the observer (Theeuwes, 1991, 1992; Yantis & Jonides, 1984). Much support for stimulus-driven accounts has come from the *additional singleton paradigm* (Theeuwes, 1992). In this paradigm, participants searched for a unique shape target (e.g., a circle) amongst homogenously shaped distractors (e.g., diamonds) and made a speeded response indicating the tilt of a line inside the target. Critically, on some trials, a non-target shape was uniquely colored. Although this *color singleton* is never the target and should therefore be ignored, it interferes with visual search. When the singleton distractor is present, response times (RTs) to detect the target are *slower* than when the singleton is absent. This *singleton-presence cost* has been used to argue that the singleton distractor automatically captures attention, despite being completely task irrelevant.

According to *goal-driven accounts*, however, salient stimuli only capture attention when they match the attentional control settings of the observer (Folk et al., 1992; Folk & Remington, 2010). Several initial studies supported these accounts by showing that salient distractors only yielded capture effects when they matched the features of the target. In addition, other studies noted that in the additional singleton paradigm, the target is defined as the unique shape amongst homogenous shapes (i.e., a shape singleton). This might encourage participants to develop an attentional template for any unique item. This *singleton-detection mode* would allow any unique item to capture attention, which can explain why the color singleton captures attention despite being ostensibly irrelevant to the task as defined by the experimenter. Importantly, if the search task is modified to discourage singleton-detection mode (e.g., so that the target shape appears among heterogenous non-target shapes), the singleton-presence cost is eliminated, suggesting that the singleton distractor no longer captures attention (Bacon & Egeth, 1994; see also Leber & Egeth, 2006).

In sum, stimulus-driven and goal-driven accounts make competing predictions about whether salient distractors have the power to automatically capture attention. The debate between these two accounts has lasted several decades without a resolution.

The Signal Suppression Hypothesis

One proposed resolution to this debate has been the *signal suppression hypothesis* (Gaspelin & Luck, 2018d; Sawaki & Luck, 2010). According to this account, salient stimuli generate a bottom-up salience signal that automatically vies for attention, consistent with stimulus-driven accounts. However, salient distractors can be proactively suppressed to prevent attentional capture, consistent with goal-driven accounts. Thus, the signal suppression hypothesis is a hybrid model that predicts that bottom-up capture will occur; but only when participants are not prepared to exert top-down control to prevent attentional capture.

One line of support for the signal suppression hypothesis has come from the *capture-probe paradigm* (Gaspelin et al., 2015). In this task, participants searched for a target shape while attempting to ignore a color-singleton distractor. On some trials, probe letters were briefly presented superimposed on each search item and then quickly disappeared. Participants were asked to recall as many letters as they could. Probe report accuracy was used to estimate the relative probability that a given search item was attended or suppressed: If an item was attended, the probe letter at that location should be more likely to be reported. If an item was suppressed, the probe letter at that location should be less likely to be reported than other items. Importantly, the probe letters at singleton-distractor locations were reported below the baseline level of letters at nonsingleton-distractor locations. This *probe suppression effect* suggests that color singletons were suppressed in order to prevent attentional capture (see also Chang & Egeth, 2019, 2021; Ma & Abrams, 2023, Exp. 3; but see Oxner et al., 2022). Similar results were later obtained with paradigms that studied eye movements generated during visual search: First eye movements were directed to the singleton distractors below the baseline level of the nonsingleton distractors, suggesting they were suppressed (Gaspelin et al., 2017; see also Adams et al., 2022; Gaspelin et al., 2019; Gaspelin & Luck, 2018b; Hamblin-Frohman et al., 2022).

Other evidence for signal suppression has come from studies of event-related potentials, which have shown that color-singleton distractors elicit a P_D component and no subsequent N2pc component, which suggests that salient items can be suppressed to prevent attentional capture (see review by Gaspelin et al., 2023; see also Drisdelle & Eimer, 2021, 2023; Feldmann-Wüstefeld et al., 2020; Gaspar & McDonald, 2014; Gaspelin & Luck, 2018a; Sawaki & Luck, 2010; van Moorselaar & Slagter, 2019). There has also been evidence from single-unit recordings in non-human primates that singleton distractors are suppressed below baseline levels early in visual cortex and that this suppression corresponds to a primate homologue to the P_D component (Cosman et al., 2018).

In sum, an abundance of recent evidence has supported the signal suppression hypothesis by demonstrating that observers can suppress salient distractors to prevent attentional capture. Recent formulations of the signal suppression hypothesis have emphasized that suppression largely occurs via implicit learning of the expected features or locations of salient stimuli. Much evidence has now shown that repeating the features of a singleton reduces, and can even eliminate, attentional capture (e.g., Adam & Serences, 2021; Gaspelin et al., 2019; Ramgir & Lamy, 2023; Vatterott & Vecera, 2012; but see also Won et al., 2019; Ma & Abrams, 2022). If true, the signal suppression hypothesis could be a coherent resolution to the attentional capture debate.

Can High Salience Overpower Suppression?

The signal suppression hypothesis, however, has recently been challenged by stimulus-driven accounts on the grounds that the initial studies supporting signal suppression may have used stimuli that were not salient enough to capture attention (Wang & Theeuwes, 2020; see also Theeuwes in Luck et al., 2021, pp. 13-17). According to this *low-salience criticism*, previous studies of signal suppression used search displays that contained relatively few items (e.g., 4 to 6 items) and this may have resulted in a relatively weak color singleton. If salience could be boosted, according to this criticism, by increasing the overall display size to increase the number of homogenously colored items, the color singleton may overpower suppression and capture attention. Indeed, there are reasons to suspect that increasing display size could improve the salience of a feature singleton by increasing the number of objects that contrast with it (Duncan & Humphreys, 1989; Nothdurft, 1993). There are, for example, studies showing that the detection of a salient target in a homogeneous background is faster when display size increases (Bacon & Egeth, 1991; Bravo & Nakayama, 1992; Maljkovic & Nakayama, 1994).

As evidence for the low-salience criticism, Wang and Theeuwes (2020) modified the capture-probe paradigm of Gaspelin et al. (2015). To manipulate salience, they varied the display size of the search array using 4, 6, and 10 items. At a low display size of 4 items, they found evidence for a probe suppression effect, replicating the original results supporting the signal suppression hypothesis. At the higher display size of 10, they found a probe capture effect: the singleton distractors were reported at a slightly *higher* rate than the nonsalient items. These results were taken to suggest that improving the salience of a color singleton could lead the distractor to overpower suppression.

There are some reasons, however, to doubt the claim that highly salient distractors can overpower suppression. Namely, several studies have shown that color singletons can be suppressed even at large display sizes where they should be highly salient. For example, Stilwell and Gaspelin (2021) found that the probe task used by Wang and Theeuwes (2020) had a design flaw that may have caused a floor effect on probe report accuracy at high display sizes. When this flaw was corrected, probe suppression effects were observed even at exceptionally high display sizes (e.g., 30 items). Also, many ERP studies have shown that color singletons elicit a P_D component and no N2pc component even at large display sizes (Drisdelle & Eimer, 2023; Gaspar & McDonald, 2014; Sawaki & Luck, 2010; Stilwell et al., 2022), which also seems to indicate that color singletons can be ignored even when highly salient.

In sum, the low-salience criticism has challenged the signal suppression hypothesis. A key limitation, however, is that there has been no method to evaluate the claim that increasing display size will improve the salience of a color singleton. This assumption—which we have also assumed (e.g., Stilwell et al., 2022; Stilwell & Gaspelin, 2021)—has been difficult to test because there have not been established methods to compare salience across displays. Importantly, although the display-size manipulation of Wang and Theeuwes (2020) could well have improved salience compared to Gaspelin et al. (2015), there were also many unsystematic changes to the search displays that might have *decreased* the salience of the color singletons (see Figure 1). For example, as noted by Chang and colleagues (2021), Wang and Theeuwes used unfilled shapes on a gray background which would reduce the color contrast between neighboring objects, as would be needed for a salience computation. This issue will be the focus of the current study.

The Oddball Detection Task

We recently developed a psychophysical technique to compare the salience of objects (Stilwell et al., 2023). In the *oddball detection task*, participants attempted to detect the presence or absence of a color singleton in brief displays that were immediately post-masked. A staircasing procedure was used to manipulate the display duration and derive the minimum duration at which the singleton could be reliably detected (an *exposure threshold*). Salience was varied via a manipulation of color contrast between the color singleton and other objects in the display. For example, a high-contrast singleton might be a blue singleton amongst red items, whereas a low-contrast singleton might be a pink singleton amongst red items. Importantly, the singleton color was randomized, meaning it had to be detected based upon its salience alone. The basic logic was that, if a singleton was more salient, it should be easier to detect resulting in a shorter exposure threshold. Indeed, exposure thresholds were reliably shorter for the high-contrast singletons than low-contrast singletons, providing evidence that the manipulation of salience was successful (see the General Discussion for consideration of the exact cognitive mechanisms involved in the task). A follow-up study further supported this claim by using the same displays in an eye-tracking task where the singleton was the target (Zhang & Gaspelin, under revision). Importantly, the specific color of the singleton was randomized and it therefore had to be found based upon salience alone. High-contrast singletons were more easily found than low-contrast singletons, again suggesting they were indeed more salient.

In Stilwell et al. (2023), the same displays were used in an attentional capture task, in which participants attempted to *ignore* the color singletons as a distractor. Eye movements revealed that, if anything, suppression was stronger for high-salience singletons than low-salience singletons (replicated by Zhang & Gaspelin, under revision). At face value, this result would seem to contradict the low-salience criticism proposed by Wang and Theeuwes (2020). But, this study used a different manipulation of salience than that of Wang and

Theeuwes (i.e., color contrast rather than display size), which may have caused the discrepant result.

The Current Study

The present study will use the oddball detection task to evaluate the low-salience criticism of the signal suppression hypothesis. According to the low-salience criticism, the original evidence of signal suppression (Gaspelin et al., 2015) can be attributed to weak salience of the color singleton, which was supposedly improved via manipulation of display size in Wang and Theeuwes (2020). Until now, this claim has been difficult to evaluate because there was no technique to compare salience across studies. The oddball detection technique provides one method of trying to assess salience of stimuli independent of their ability to capture attention and could therefore be an important tool in resolving this dispute. To preview the results, we found a surprising outcome: the color singletons used by Wang and Theeuwes were, if anything, *less salient* than the singletons used by Gaspelin et al. Experiment 2 explores this outcome in an experiment to ascertain how another key factor other than display size, such as color fill, might have influenced salience. Experiments 3 and 4 then test whether increasing salience (using what was learned from Experiments 1 and 2) leads to capture by the color singleton, as predicted by the low-salience criticism.

EXPERIMENT 1

As shown in Figure 1, an oddball detection task was used to compare the salience of color singletons in the 10-item displays of Wang and Theeuwes (2020) and the 4-item displays of Gaspelin et al. (2015). Participants detected the presence or absence of a color singleton in displays that were briefly presented. A staircasing procedure was then used to adjust the duration of the next search display, and this was ultimately used to determine the participant's minimum exposure threshold needed to detect the color singleton. If a color singleton is more salient in one display versus another, it should be easier to detect, resulting in a shorter exposure threshold.

According to the low-salience criticism, the Gaspelin et al. (2015) displays produced a *less salient* singleton than the Wang and Theeuwes (2020) displays. If true, the exposure thresholds to detect the color singleton should be longer for the Gaspelin et al. displays than the Wang and Theeuwes displays. An alternative possibility, however, is that the Wang and Theeuwes may have not improved salience as claimed. There were many unsystematic differences between displays other than display size (e.g., fill of the shapes, background color) and these differences may have reduced the salience the color singleton. Thus, it is possible that the colors singletons in Wang and Theeuwes were not actually more salient than those of Gaspelin et al., despite the increased display size. According to this account, exposure thresholds to detect the color singletons might be equal or even shorter in the Gaspelin et al. displays than the Wang and Theeuwes displays.

Method

Participants

A sample size of 24 participants from State University of New York at Binghamton participated for course credit. This sample size was determined a priori based upon a previous study using the oddball detection task (Stilwell et al., 2023). Using the effect size of the difference between exposure thresholds (d_z = 1.71) for low- and high-salience singletons from Stilwell et al. (2023), 9 participants would be needed to obtain 99% power. Because the oddball detection technique is relatively new, we chose to err on the side of caution and collect a larger sample size than needed.

The participants were 18 women and 6 men (M_{age} = 19.1 years, *SD* = 1.1 years). For the sake of convenience, the current study had a target population of local undergraduates from the university. Demographic information beyond age and gender was not collected. It is therefore possible that the current results may not apply to all populations. All participants had normal color vision as indicated by an Ishihara test and had normal or corrected-to-normal visual acuity. This study was approved by a research ethics committee at the university.

Apparatus

Stimuli were presented using PsychToolbox for MATLAB on a Dell Precision 3660 with the Linux operating system (Kleiner et al., 2007). An Asus VG248QG monitor with a refresh rate of 100 Hz

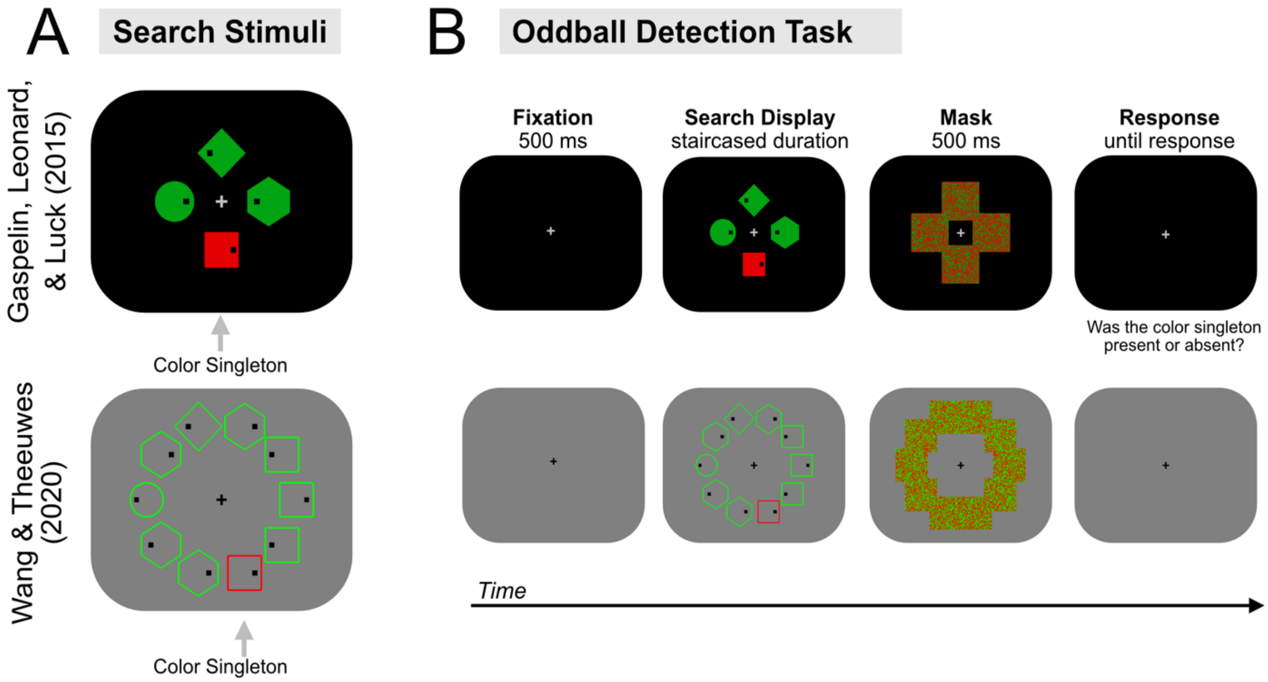

Figure 1. Stimuli and task for Experiment 1. (A) Stimuli were identical to either 4-item displays from Gaspelin et al. (2015) or 10-item displays from Wang & Theeuwes (2020). (B) In the oddball detection task, participants attempted to detect the presence of a color singleton that was briefly presented and then immediately masked. A staircase procedure was used to determine minimum exposure duration needed to detect the singleton.

presented stimuli at a viewing distance of 100 cm in a dimly lit room.

Stimuli

The search array stimuli were identical to the set-size 4 displays of Gaspelin, Leonard, and Luck (2015, Exp. 3) and the set-size 10 displays of Wang and Theeuwes (2020). Examples of the search displays are depicted in Figure 1A.

Gaspelin, Leonard, and Luck (2015). Four shapes were arranged in a notional circle with an eccentricity of 1.65° from the center of the screen. Each array contained one diamond (1.6° x 1.6°), one square (1.2° x 1.2°), one circle (1.4° diameter), and one hexagon (1.5° in width and height). Each shape contained a small black square (0.2° x 0.2°) that appeared 0.2° from the left or right edge of the shape. The position of the black square was selected at random for each shape. A gray fixation cross (30.0 cd/m^2; 0.4° x 0.4°) appeared at the center of the display. The colors of the shapes were either red (30.0 cd/m^2, x = .627, y = .330) or green (30.0 cd/m^2, x = .292, y = .631) and the shapes were filled with this color. Stimuli appeared on a black background.

Wang and Theeuwes (2020). Ten shapes were arranged in a notational circle with an eccentricity of 3.0° from the center of the screen. Each array contained one diamond (1.6° by 1.6°), four squares (1.6° by 1.6°), one circle (1.4° diameter), and four hexagons (1.6° by 1.6°). Each shape contained a small black square (0.2° x 0.2°) that appeared 0.2° from the left or right edge of the shape. A black fixation cross (0.4° x 0.4°) appeared at the center of the display. The colors of the shapes were either red (63.4 cd/m^2, x = .641, y = .331) or green (226 cd/m^2, x = .304, y = .639) and the shapes were outlined with this color. Stimuli appeared on a gray (58.3 cd/m^2, x = .302, y = .322) background.

Procedure

Participants completed the *oddball detection task* developed by Stilwell et al. (2023), in which they attempted to detect the presence of a color singleton in a search display that was briefly presented and then immediately masked (Figure 1B). A staircase procedure was used to titrate the duration of the search display on the current trial based upon performance on the previous trial. This allowed us to determine the minimum exposure duration (i.e., the *exposure threshold*) needed to detect the color singleton.[1] Exposure thresholds were then compared for the color singletons used by Gaspelin et al. (2015) and Wang and Theeuwes (2020).

The color singleton was present on half of trials and, when present, it was equally likely to appear at any of the search locations. An important aspect of the oddball detection task is that the color singleton must be detected solely based upon its color contrast with other objects. To prevent an attentional set for the specific colors, the color of the singleton and other items were randomly selected on each trial. Thus, the color singleton was equally likely to be the green item amongst red items or the red item amongst green items. In addition, the specific shape of the color singleton was also randomly selected on each trial. This prevents participants from using feature-based attention to detect an item of a specific color or a specific shape. This is crucial because it means that the exposure thresholds measure the speed to detect an object based upon its bottom-up salience alone (see also the General Discussion).

Each trial began with a fixation cross for 500 ms, followed by the search display which appeared for a variable duration (see next paragraph) and was then followed immediately by a pattern mask for 500 ms (see Figure 1B). The purpose of the pattern masks were to prevent any image of the search items from remaining in iconic memory (Loftus et al., 1985; Sperling, 1960). Each mask contained a random array of colored dots (0.03° by 0.03°) and an equal number of red and green pixels. The specific colors in the mask were matched to the colors of the search items. A pattern mask appeared at each search position and was 125% larger than the masked stimulus (i.e., an additional 0.4° in width and height). After the pattern mask, participants made an unspeeded response to the presence or absence of the color singleton via the right and left shoulder buttons on a gamepad, respectively.

The search-display duration was adjusted using a weighted up-down staircase procedure with a step-size ratio of 1/3 (Kaernbach, 1991; see also Greene & Oliva, 2009). At the beginning of each half of the experiment, the search display duration started at 50 ms. If the participant responded incorrectly (either failing to detect the singleton when present, a miss, or falsely reporting a singleton when it was absent, a false alarm), the next trial's exposure duration was 30 ms *slower*, with a maximum of 300 ms. If the participant responded correctly (either correctly detecting a singleton when it was present, a hit, or correctly rejecting the singleton when it was absent, correct rejection), the next trial's exposure duration was 10 ms *faster*, with a minimum duration of 10 ms. This staircase procedure ensures that accuracy asymptotes at approximately 75% correct.

The oddball-detection task consisted of two halves, one for each type of display. The order of display configuration was counterbalanced across participants. At the beginning of each half, an example search display depicting the upcoming stimuli configurations was presented to illustrate which type of singletons participants were to search for. In each half, participants first performed one practice block of 80 trials followed by four regular blocks of 80 trials. This yielded 640 total experimental trials. For each display type, there were 320 trials of which 160 were singleton-present trials and 160 were singleton-absent trials. Participants received block-by-block feedback on mean accuracy.

Transparency and Openness

All data and stimulus presentation programs are available on the Open Science Framework (OSF.io) at https://osf.io/bpjrf/. Experiments 1–3 were not preregistered, and the data was collected in 2023.

Results

Overall accuracy was 76.6%, which is to be expected with the weighted up-down staircase procedure (Kaernbach, 1991). The average false alarm rate was 34.5% and the average hit rate was 87.8%.

Exposure Thresholds

The primary dependent measure was the exposure threshold, which we define as the minimum exposure duration needed to detect the color singleton at 75% accuracy. To derive the exposure threshold, we first generated a plot of each participants accuracy at each potential exposure duration for each display type (Figure 2A; see also

[1] According to the signal detection theory, there are no discrete perceptual thresholds (Green & Swets, 1974). Our usage of "exposure threshold" is meant to simply denote the critical exposure duration whereby 75% accuracy is achieved.

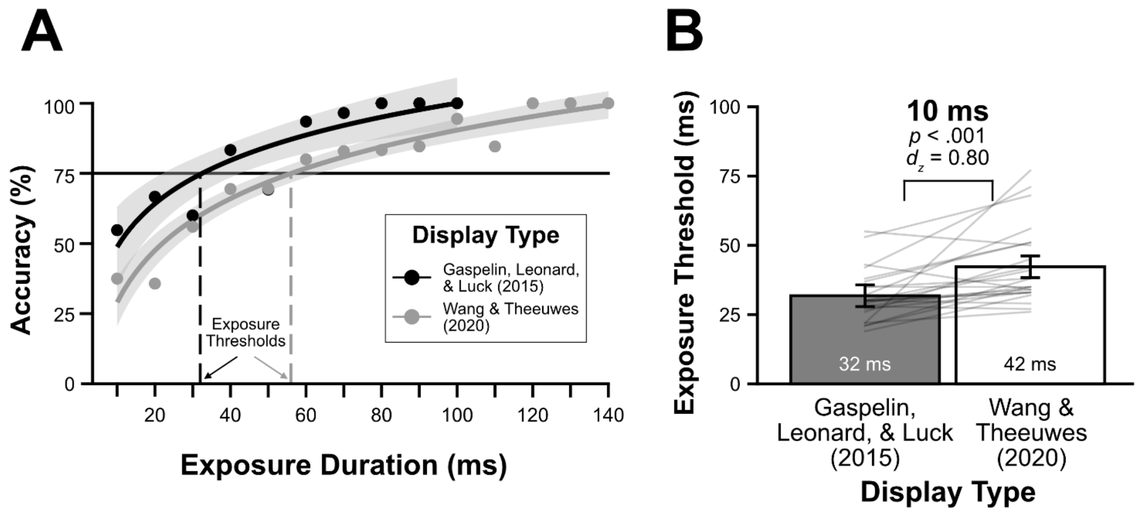

Figure 2. Results from the oddball detection task in Experiment 1. (A) An example of an individual participant's exposure threshold estimations. Shading depicts the standard error of the mean. (B) Grand averaged exposure thresholds for Gaspelin et al. (2015) and Wang and Theeuwes (2020) displays. Each participant's score is shown as a whisker plot. Error bars represent within-subject 95% confidence intervals in all figures of this paper (Cousineau, 2005; Morey, 2008).

Stilwell et al., 2023). A logarithmic function was then fit to these data, and we estimated the exposure threshold needed to reach 75% accuracy. If a singleton is more salient in one display than another, it should be detected more quickly resulting in a lower exposure threshold. Surprisingly, we found that exposure thresholds were significantly *lower*

for the color singletons appearing in the Gaspelin et al. (2015) displays (32 ms) than the Wang and Theeuwes (2020) displays (42 ms), $t(23) = 3.90$, $p < .001$, $d_z = 0.80$. These results suggest that the color singletons in Gaspelin et al. (2015) were, if anything, *more* salient than the color singletons in Wang and Theeuwes (2020), opposite to the claims of the low-salience criticism.

Sensitivity (d') by Exposure Duration

An exploratory analysis compared the sensitivity (*d'*) to detect the singleton at each exposure duration for the two types of displays. Sensitivity provides an unbiased estimate of the perceptual discriminability of stimulus, unlike overall accuracy.[2] If the color singleton was more salient in the displays of Gaspelin et al. (2015) than Wang and Theeuwes (2020), then there should be an enhancement in sensitivity (*d'*) to detect the singleton, especially at the shorter exposure durations

The number of hits, misses, false alarms, and correct rejections were calculated for each subject at each potential exposure duration for both display types. Many participants had few or no trials at longer exposure durations (i.e., greater than 80 ms) because they did not make enough errors for staircasing procedure to allow many trials with these exposure durations. We truncated our analysis from 10 to 70 ms, which prevented any participants from having missing values. The data were then corrected using the method proposed by Hautus (1995) which adds 0.5 to each cell (hits, false alarms, misses, and correct rejections) to prevent zeros that cause false alarm rates and hit rates to approach infinity. Sensitivity (*d'*) was calculated using the *z*-score of hit rates minus the *z*-score of false alarm rates.

As shown in Table 1, sensitivity (*d'*) to detect the singleton was higher for the Gaspelin et al. (2015) displays than the Wang and Theeuwes (2020) displays at several durations. Paired-sample

[2] As a cautionary note, it would be inappropriate to calculate an overall *d'* in the oddball detection task because a staircasing procedure adjusted the exposure durations separately between the two display types. Thus, the mean exposure duration for Wang and Theeuwes (2020) was, on average, longer than that of Gaspelin et al. (2015). We therefore compared sensitivity (*d'*) between the two display types at each potential exposure duration.

Table 1

Sensitivity (d') and Response Criterion (c) by Exposure Duration for Each Display Type

Measure	Display Type	Exposure Duration						
		10 ms	20 ms	30 ms	40 ms	50 ms	60 ms	70 ms
d'	Gaspelin et al. (2015)	0.28	0.88	1.57	2.49	2.81	2.84	2.48
	Wang & Theeuwes (2020)	0.41	0.43	1.13	1.80	2.33	2.51	2.55
	Significance	n.s.	*	*	*	*	*	n.s.
c	Gaspelin et al. (2015)	-0.14	-0.43	-0.75	-0.71	-0.48	-0.27	-0.22
	Wang & Theeuwes (2020)	0.04	-0.33	-0.76	-0.94	-0.75	-0.62	-0.36
	Significance	n.s.	n.s.	n.s.	n.s.	*	*	*

Note. Significance levels are from a paired-samples t-test comparing the two display types adjusted for multiple comparisons using false discovery rate (FDR).

t tests were used to compare sensitivity at each individual exposure duration using a false discovery rate correction for multiple comparisons (Benjamini & Hochberg, 1995; Benjamini & Yekutieli, 2001). As depicted, sensitivity was significantly improved for the Gaspelin et al. (2015) displays from 20 to 60 ms (p's < .05). There were no observed differences in sensitivity at 10 ms or 70 ms (p's > .10). This provides further evidence, using an unbiased measure of accuracy, that detection thresholds for the singletons were enhanced for the Gaspelin et al. (2015) displays.

We also analyze response bias (c) as a function of exposure duration. We had no a priori hypotheses about response bias, but did this analysis for the sake of completeness. The analysis showed an interesting pattern. There was no significant difference in response bias (c) between the two display types for the shorter exposure durations (10, 20, 30, and 40 ms), but there was a significantly larger liberal bias ($c < 0$) at the longer durations (50, 60, and 70 ms). It is not immediately clear why this should have been the case. However, the critical observation is that the measure of sensitivity, *d'*, which is independent of criterion, showed that the Gaspelin et al (2015) stimuli yielded significantly higher *d'* values than the Wang and Theeuwes (2020) stimuli at 20, 30, 40, 50, and 60 ms durations.

Discussion

The low-salience criticism claims that the original evidence of signal suppression can be attributed to weak salience of the color singleton. Although Wang and Theeuwes (2020) suggested they improved the salience of the color singleton compared to Gaspelin et al. (2015) by increasing display size, there were many uncontrolled differences between the search displays which could have also influenced salience. We therefore assessed the salience of the singleton using exposure thresholds derived from an oddball detection task (Stilwell et al., 2023). Both display types produced relatively short exposure thresholds (i.e., less than 45 ms), suggesting that both displays used color singletons that were highly salient. Importantly, the color singletons in the Gaspelin et al. displays were detected at shorter exposure thresholds than in the Wang and Theeuwes displays, suggesting they were more salient in the original signal suppression study. This result challenges the low-salience criticism, which predicts the opposite pattern of results.

EXPERIMENT 2

Experiment 1 produced a surprising outcome: The color singletons in Gaspelin et al. (2015) were, if anything, *more salient* than those in Wang and Theeuwes (2020), at least given our operational definition of salience. As previously noted, there were many uncontrolled differences between the two studies, such as color fill of the

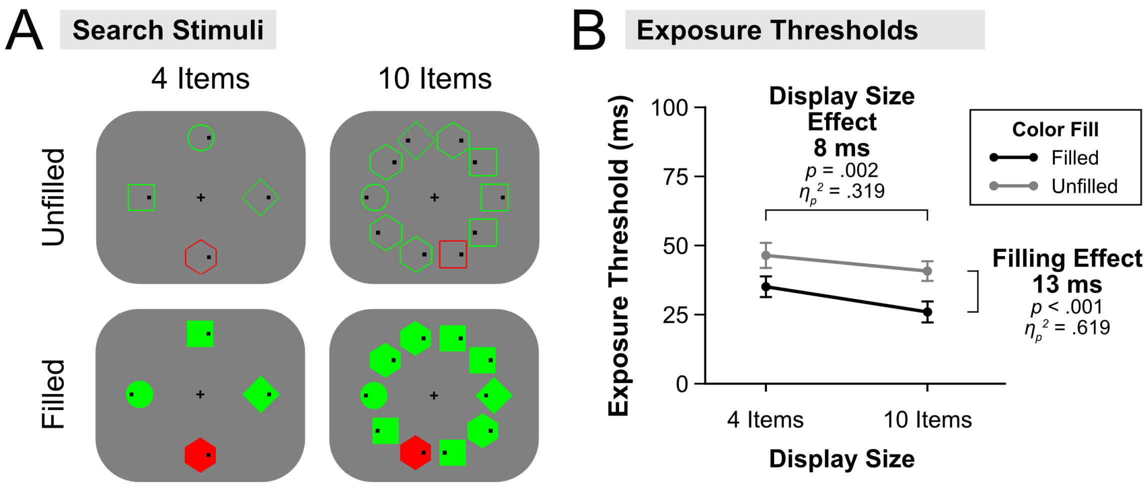

Figure 3. Stimuli and results from Experiment 2. (A) Display size (4 or 10 items) and color fill (filled or unfilled) were manipulated in a fully crossed design. (B) Grand averages of exposure thresholds as a function of display size and color fill.

shapes[3] and background color (see Figure 1). A straightforward explanation is therefore that one of these factors led to the unexpected difference in salience. Perhaps the most suspicious of these factors is color fill. Using unfilled shapes greatly reduces the amount of color information available to contrast with other objects and could thereby reduce the salience of the color singleton. This could explain why the color singletons in Wang and Theeuwes (2020) were less salient than those in Gaspelin et al. (2015).

Experiment 2 tested this explanation by manipulating (a) color fill and (b) display size to test the influence on salience in oddball detection task. The same displays, based upon Wang and Theeuwes (2020), were used in all conditions to ensure observed differences in salience would be attributable to these factors only. The key question is whether filled shapes will produce singletons of higher salience than unfilled shapes, which could help explain why the color singletons in the Gaspelin et al. (2015) displays were more salient despite the lower display size. If true, we should observe lower exposure thresholds for filled shapes than unfilled in the oddball detection task.

Method

A new sample of 24 students from State University of New York at Binghamton, participated for course credit (18 women and 6 men; M_{age} = 19.8 years, *SD* = 2.5 years).

Participants performed the same oddball detection task as in Experiment 1, except for the following changes. As depicted in Figure 3A, the eccentricity, shape dimensions, and colors matched those used by (Wang & Theeuwes, 2020). We then manipulated display size and shape fill in a fully crossed design. Display size was manipulated by generating displays which contained 4 items or 10 items. Displays that contained 4 items had one circle, one diamond, one square, and one hexagon, and displays that contained 10 items had one circle, one diamond, four squares, and four hexagons. Shape fill was manipulated by using shapes with were completely filled with color (filled) or by using outlines of shapes (unfilled). Thus, there were four

[3] It is worth clarifying that we use the term “color fill” to represent the difference between the Wang & Theeuwes (2020) and the Gaspelin et al. (2015) stimuli. However, a glance at Figure 1 makes it clear that the outline shapes and the filled-in shapes differ in more ways than just how their borders are defined. In particular, the number of color pixels is very different in the two cases. We make no effort here to tease apart the contributions of number of color pixels and border type to saliency.

conditions varying in display size (4 or 10 items) and color fill (unfilled or filled).

Each of the four display type sections began with one block of 80 practice trials. After practice, there were two blocks of 80 trials for each of the four display types, which produced 160 trials for each condition (i.e., 640 trials total). As in Experiment 1, the conditions were blocked, and the order of condition blocks was counterbalanced across participants.

Results

Overall accuracy was 77.0%, which is expected in a staircase procedure (Kaernbach, 1991). The average false alarm rate was 36.5% and the average hit rate was 90.5%.

As in Experiment 1, logarithmic functions were fit to each participant's data for each display type and were used to estimate the minimum exposure duration needed to obtain 75% accuracy. These exposure thresholds were analyzed using a two-way within-subject ANOVA with factors display size (4 vs. 10) and color fill (unfilled vs. filled). There was a main effect of display size, $F(1,23) = 12.24$, $p = .002$, *adj.* $\eta_p^2 = .319$, indicating that exposure thresholds were lower for 10-item than 4-item displays. This suggests that increasing the display size *did* increase the salience of color singleton. There was an even larger main effect of color fill, $F(1,23) = 40.05$, $p < .001$, *adj.* $\eta_p^2 = .619$, indicating that exposure thresholds were lower for filled (31 ms) than unfilled (44 ms) displays. This suggests that filling the shapes with color also increased the salience of the color singleton. The interaction between display size and color fill was nonsignificant, $F(1,23) = 1.45$, $p = .24$, *adj.* $\eta_p^2 = .018$. Altogether, these results suggest that both display size and color fill independently contribute to salience.

As an exploratory analysis, we compared two conditions that are directly relevant to Experiment 1: the unfilled set-size 10 condition (which resembles Wang & Theeuwes, 2020), and the filled set-size 4 condition (which resembles Gaspelin et al., 2015). We were interested in whether we would observe the same result from Experiment 1, in which color singletons in low set-size, filled display were more easily detected than a high set-size, unfilled display. Indeed, a preplanned *t* test indicated that the exposure threshold of the unfilled, set-size 10 condition (41 ms) was significantly higher than that of the filled, set-size 4 condition (35 ms), $t(23) = 2.10$, $p = .047$, $d_z = 0.43$. Interestingly, the exposure thresholds are remarkably similar to those observed in Experiment 1 (42 ms and 32 ms, respectively). This result demonstrates that a set-size 4 display can be at least as salient as a set-size 10 display, if other factors such as color fill are not properly controlled (e.g., as in Wang & Theeuwes, 2020).

Discussion

Experiment 2 assessed how color fill and display size contribute to the perceptual salience of color singletons using the oddball detection task. The results demonstrated that both factors contributed to the salience of a color singleton, at least as measured by the oddball detection task. This provides a potential explanation for the results of Experiment 1, which showed the color singletons at set-size 4 in Gaspelin et al. (2015) were more salient than those at set-size 10 in Wang and Theeuwes (2020). Although salience was improved by increasing display size by Wang and Theeuwes, salience was also reduced by using unfilled color shapes. This resulted in an overall salience level that was lower than the original Gaspelin et al. experiments that were being criticized for having a low salience.

EXPERIMENT 3

Together, the results of Experiments 1 and 2 demonstrated that both display size and color fill contribute to the salience of a singleton. However, those experiments leave it unclear whether highly salient singletons can overpower suppression and capture attention, as originally suggested by the low-salience criticism. Experiment 3 therefore tested whether the highly salient singletons, generated in Experiment 2, can be ignored when they were task irrelevant. Participants performed an additional singleton paradigm in which they searched for a target defined by shape and color (e.g., green circle) while attempting to ignore a color singleton distractor (see Figure 4). Importantly, the displays were identical to the set-size 10 of Experiment 2 and we compared capture effects for displays with unfilled shapes and filled shapes.

The low-salience criticism is derived from a stimulus-driven account of attentional capture. According to a purely stimulus-driven account, the singleton distractor should automatically capture attention, and capture should be more likely to occur as the distractor becomes more salient. This account therefore predicts that there should be a singleton-presence cost in both conditions because the singleton should be highly salient at set-size 10. Moreover, the singleton presence cost should be larger for the filled displays which produced a more salient color singleton in Experiment 2.

Method

Participants

A new sample of 48 participants (29 women, 18 men, 1 nonbinary individual; M_{age} = 19.1 years, *SD* = 1.3 years), who were students from State University of New York at Binghamton, participated for course credit. A power analysis using the effect size of the singleton-presence cost (d_z = 2.21) from a similar experiment (Wang & Theeuwes, 2020, display size 10 condition) suggested a small sample size (N = 7) should be sufficient to obtain 99% power. However, we collected a large sample size to err on the side of caution because prior studies have had challenges replicating the capture effect (Stilwell et al., 2022; Stilwell & Gaspelin, 2021).

Apparatus

The stimulus presentation system and viewing conditions were identical to Experiment 1 and 2, except that the refresh rate was 60 Hz. A photosensor was used to measure the timing delay of the video system at this refresh rate (12 ms) and this delay was subtracted from RTs in Experiment 3.

Stimuli and Procedure

The additional singleton paradigm was used (Theeuwes, 1992). Stimuli were identical to the unfilled and filled displays from set-size 10 in Experiment 2 (see Figure 4A). The main difference was that the singleton was now task irrelevant. Participants searched for a target that was defined by shape and color (e.g., green circle). They made a speeded button press regarding the location of black dot inside the target shape (left- or right-side) using a gamepad. On half of trials, all search items were the same color (*singleton-absent trials*). On the other half of trials, a randomly selected distractor was rendered in a unique color (*singleton-present trials*), creating a *singleton distractor*. Participants were told to ignore this singleton distractor as it would never be the target and could therefore be ignored. The color of the target (e.g., green) and singleton (e.g., red) were consistent for the entire experiment and were counterbalanced across participants. This was done because previous studies have shown that distractor suppression is feature-based, meaning that participants learn to ignore a specific color (e.g., Chang & Egeth, 2019; Gaspelin et al., 2019; Gaspelin & Luck, 2018b; Ramgir & Lamy, 2023; Savelson & Leber, under review; Stilwell et al., 2019; Vatterott & Vecera, 2012). Similarly, the shape of the target (circle or diamond) was also counterbalanced across participants. The filled and unfilled color conditions were separated by experiment half and the order of

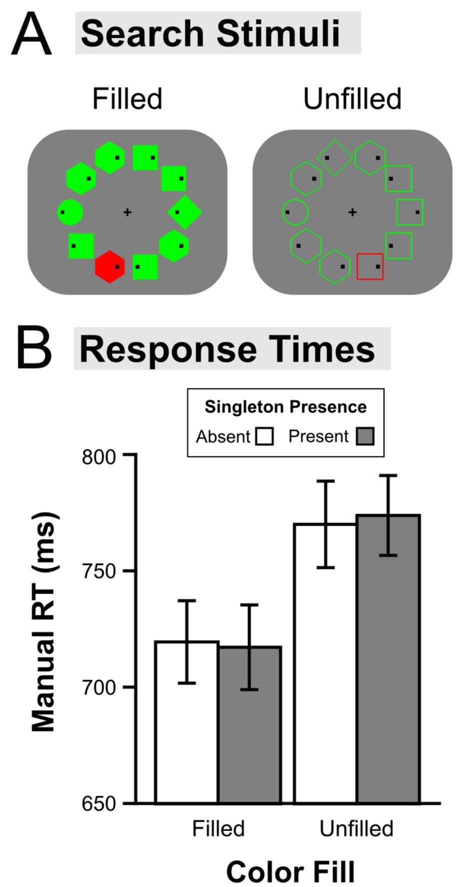

Figure 4. Stimuli and results from Experiment 3. (A) Participants searched for a shape-defined target (e.g., the green circle) and responded to the location of a dot inside (left vs. right). A color singleton distractor was present on half of trials. All search items were either filled or unfilled with color. (B) Mean response times (ms) for each color-fill condition.

these conditions were counterbalanced across subjects.

Each trial began with the fixation display for 500 ms. Next, the search array appeared until a response was made. If participants made an inaccurate response, a 200 Hz tone played for 300 ms accompanied by the words “Incorrect” for 1500 ms. If a participant took more than 2000 ms to respond, a 200 Hz tone played for 300 ms accompanied by “Too Slow” in the center of the screen for 1500 ms. Feedback was provided at the end of each block about mean accuracy and mean RT.

For each color-fill condition, participants performed two practice blocks of 40 trials (one for each experimental half) followed by 12 regular blocks of 40 trials each. This resulted in a total of 1040 trials (i.e., 520 trials for each color-fill condition).

Results

Trials with inaccurate responses (2.8%) and trials with RTs faster than 200 ms or slower than 2000 ms (0.8%) were removed from analysis, resulting in 3.7% total trials removed.

Response Time

As shown in Figure 4B, there were no singleton-presence costs for either filled or unfilled displays indicating that the salient distractors did not capture attention. To formally analyze this, we conducted a two-way within-subject ANOVA on mean RTs with factors color fill (filled vs. unfilled) and singleton presence (absent vs. present). There was a main effect of color fill, $F(1,47) = 12.71, p < .001$, *adj.* $\eta_p^2 = .20$, indicating that mean RTs were faster for filled displays than unfilled displays. This likely occurred because the filled displays made the shapes easier to see parafoveally, and therefore improved attentional guidance toward the target shape. There was no main effect of singleton presence, $F(1,47) = 0.10, p = .76$, *adj.* $\eta_p^2 = -.0.02$, suggesting no attentional capture by the color singleton distractor in either condition. There was also no interaction between singleton presence and color fill, $F(1,47) = 1.68, p = .20$, *adj.* $\eta_p^2 = .01$, suggesting that the singleton-presence costs were not modulated by color fill.

Singleton-presence costs were calculated as a difference score of mean RT on singleton-present trials minus mean RT on singleton-absent trials. Preplanned one-sample *t*-tests then evaluated the significance of singleton-presence costs for each display type. The singleton-presence cost was not significant for either filled displays (-2 ms), $t(47) = 0.62, p = .54, d = 0.09$, or unfilled displays (4 ms), $t(47) = 1.11, p = .27, d = 0.16$, further suggesting that the singleton did not capture attention. In addition, Bayesian *t*-tests with default priors of 0.707 (Morey et al., 2016) were computed for each singleton-presence cost. The analyses resulted in a $BF_{01} = 5.29$ for filled and a $BF_{01} = 3.45$ for unfilled, indicating strong evidence in favor of the null hypothesis.

Error Rates

Error rates were generally quite low (<4%). Error rates were submitted to the same two-way within-subject ANOVA as mean RTs with factors color fill (filled vs. unfilled) and singleton presence (absent vs. present). There were no differences in error rates between any of the conditions (all *p*’s > .14).

Discussion

Experiment 3 further tested the low-salience criticism by assessing whether the high-salience singletons could overpower suppression and capture attention. Participants performed an additional singleton paradigm in which they searched for a target and attempted to ignore a singleton distractor. The results refuted the low-salience criticism. First, in both filled and unfilled conditions, the singleton was successfully ignored despite their apparent high salience in Experiment 2. Furthermore, using filled displays instead of unfilled displays, which was shown to improve salience in Experiment 2, did not modulate capture by the color singleton. The results instead indicate that the singletons did not capture attention, regardless of their salience.

One might wonder why Experiment 3 did not find a singleton-presence benefit—faster mean RTs when the singleton is present than absent—a pattern associated with suppression (Gaspelin et al., 2015; see also Chang & Egeth, 2019; Ma & Abrams, 2023, 2022; Lien et al., 2022). A potential explanation of the singleton-presence benefit is that, if the singleton is suppressed, the display size of the search array will be reduced by one item. This will reduce the effective display size of items that need to be searched to locate the target shape. When the display size is low (e.g., 4 items), the effective display size of the display would be substantially reduced if the singleton is suppressed, leading to a larger proportion of items dismissed as candidate target items (i.e., eliminating 1 out of 4 items reduces the overall display size by 25%). However, as the display size increases (e.g., 10 items), the utility of ruling out the singleton decreases (i.e., eliminating 1 of 10 items only reduces the overall display size by 10%). Consistent with this interpretation, the original Gaspelin et al. (2015) study only observed a significant singleton-presence benefit in experiments that used lower display sizes. We will address the specific issue of whether the singleton was suppressed in Experiment 4.

EXPERIMENT 4

A shortcoming of Experiment 3 is that the experimental approach did not allow us to compare attentional processing at each location to evaluate whether the color singleton was suppressed below baseline levels of other objects, as predicted by the signal suppression hypothesis. Experiment 4 therefore used a *capture-probe paradigm* to evaluate whether the color singletons in Experiment 3 were suppressed. The capture-probe paradigm was chosen because it is the same paradigm used to evaluate suppression by both of the original studies in question (Gaspelin et al., 2015; Wang & Theeuwes, 2020). This experiment was also preregistered.

The capture-probe paradigm involves two types of trials (Figure 5A; Gaspelin et al., 2015). On most trials, participants search for a target stimulus and attempt to ignore a salient distractor. On a random subset of trials, letters are briefly superimposed over search objects before disappearing. Participants are then asked to recall as many letters as possible from the probe array. Probe report accuracy is used as a proxy measure of attentional allocation. If the color singleton is suppressed, the probe letter at the singleton distractor should be reported at a lower probability than the letter at the average nonsingleton distractor: a *probe suppression effect* (Gaspelin et al., 2015; Gaspelin & Luck, 2018c, 2018a; Wang & Theeuwes, 2020). In the context of the current experiment, the signal suppression hypothesis predicts that salient distractors are suppressed to prevent attentional capture. Therefore, probe suppression effects should be apparent for both filled and unfilled distractors.

Method

Participants

A new sample of 32 students from the University of Missouri participated for course credit. The sample size was determined a priori based on a power analysis of the probe suppression effect from Gaspelin et al. (2015, Exp. 4), which used similar methods and stimuli to the current experiment. Based upon the observed effect size (d_z = .94), 23 participants would be needed to obtain 99% power. A sample size divisible by 8 was needed to counterbalance different versions of the experiment and it was unknown how strongly stimulus salience would influence probe suppression effects. We therefore chose to err on the side of caution and collect 32 participants. This sample size was determined a priori in the preregistration. One participant was replaced for having a mean accuracy on the search task that was 3 standard deviations less than the group mean. The final sample of 32 participants consisted of 19 women, 12 men, and 1 nonbinary individual (M_{age} = 18.9 years, *SD* = 1.0 years).

Apparatus

The stimulus presentation system and viewing conditions were similar to Experiment 1–3, except that a refresh rate of 120 Hz was used. A photosensor was used to measure the timing delay of the video system at this refresh rate (7 ms) and this delay was subtracted from RTs in Experiment 4.

Stimuli and Procedure

Stimulus displays were identical to those used in Experiment 3, but the task was changed to a capture-probe paradigm (Figure 5A). On *search trials* (75% of trials), participants searched for a specific target shape (e.g., green circle) amongst heterogenous shapes and reported the location of a dot inside as quickly as possible (left vs. right). Responses were

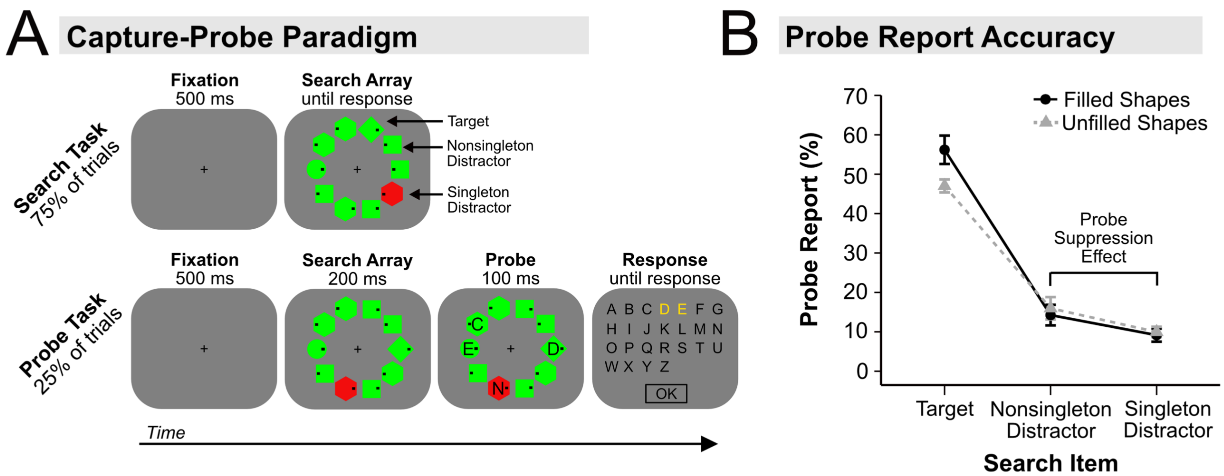

Figure 5. Stimuli and results from Experiment 4. (A) Participants completed a capture-probe paradigm to evaluate whether the color singletons were suppressed. There were separate halves of the experiment with both filled and unfilled shapes (not depicted here; see Figure 4A). (B) Probe report accuracy for displays with filled and unfilled shapes. Both display types yielded probe suppression effects. Error bars represent within-subject standard error on the mean.

made by clicking the left- or right- button on a mouse. The color (green or red) and shape (circle or diamond) of the target were constant throughout the experiment and counterbalanced across participants. On every trial, a color singleton was presented at a randomly selected nontarget location. The singleton was present on every trial to maximally encourage suppression (e.g., see Won et al., 2019). Participants were informed of this and were instructed to ignore the color singleton.

On *probe trials* (25% of trials), participants instead performed a letter-probe task (Gaspelin et al., 2015). After fixation, the search array appeared briefly and then letters were superimposed over the shapes before disappearing. To be consistent with Gaspelin et al. (2015; Experiments 1–3) and Wang and Theeuwes (2020), letter masks were not used. Probe letters were black and approximately 1° in width and height. The letters were selected at random from the English alphabet with the exception that no letter repeated. To prevent floor effects on probe report (Stilwell & Gaspelin, 2021), letters appeared at only four locations (target, singleton distractor, and two randomly selected nonsingleton distractors). After the probe display disappeared, a response display appeared that consisted of all letters in the English alphabet. Participants were asked to report which letters they saw by clicking on them, which made them turn yellow. Letters could be unselected by clicking on them again, to turn them black. A final response was submitted by clicking the OK button. Probe report was untimed, and participants received no feedback about accuracy of their responses. Probe trials occurred on a random subset of trials with the exception that probe trials could not occur for two trials in a row. This was meant to maximize the likelihood of observing probe suppression effects by ensuring participants had a strong attentional set from the search task on the previous trial. In other words, multiple subsequent probe trials could temporarily reduce the attentional set (Kim & Cave, 1995).

Search trials began with a fixation cross for 500 ms followed by the search display until a response was made. After a response, participants were given immediate feedback about the accuracy of their response via tone feedback. If no response was made within 3000 ms, the trial was considered a timeout and marked as inaccurate. Probe trials began the same as search trials. A fixation cross appeared for 500 ms, followed by a search display for 200 ms. Next, the letter probe displays appeared for 100 ms and then disappeared. A probe response screen appeared until a response was made. No feedback about the accuracy of the probe response was provided.

For one half of the experiment, the shapes were filled, and for the other half, shapes were unfilled. The order of halves was counterbalanced across participants. Each block consisted of 60 trials. Each half consisted of one practice block of only search trials, a practice block with both probe and search trials intermixed, and then four experimental blocks. This resulted in 12 blocks of 60 trials each (720 trials). In the final data set, there were 120 probe trials per subject, 60 for each shape fill condition (filled and unfilled).

Transparency and Openness

The methods of this experiment were preregistered at https://osf.io/4gtkn and data was collected in 2024.

Results

Search Trials

Trials with inaccurate responses (1.3%) and trials with RTs faster than 200 ms or slower than 2000 ms (1.1%) were removed from search trial analysis, resulting in 2.3% total trials removed. The singleton was present on every trial in this experiment to maximally incentivize suppression and we therefore could not assess singleton-presence costs or benefits (but see Experiment 3 for RTs in the same task). RTs were generally slower for unfilled shapes (847 ms) than filled shapes (812 ms), $t(31) = 2.70$, $p = .01$, $d_z = .48$. Error rates did not significantly differ between filled (1.1%) and unfilled (1.3%) displays, $t(31) = 1.02$, $p = .31$, $d_z = .18$.

Probe Trials

Participants reported an average of 1.3 letters, and this did not significantly differ between display fill conditions, $t(31) = .82$, $p = .42$.

The key question in Experiment 4 was whether the singleton could be suppressed. Figure 5B shows probe report accuracy as a function of item type and display fill. The nonsingleton distractor data have been divided by 2 to provide a per item estimate of probe report accuracy. As can be seen, probe report accuracy for the letter at the target was generally improved compared to letters at the average nonsingleton distractor (a target enhancement effect). Similarly, probe report accuracy for the letter at the singleton distractor was impaired compared to the average nonsingleton distractor (a probe suppression effect).

The primary question was whether the singleton would be suppressed. We first computed probe suppression effects as probe report accuracy at the average nonsingleton distractor minus the singleton distractor. A positive score would indicate suppression, whereas a negative score would indicate capture. One-sample *t* tests confirmed that probe suppression effects were significant for both filled displays (5.1%), $t(31) = 2.13$, $p = .041$, $d_z = .38$, and unfilled shapes (6.0%), $t(31) = 4.99$, $p < .001$, $d_z = .88$. There was not a significant difference in probe suppression effects between the display fill types, $t(31) = .355$, $p = .725$, $d_z = .06$, suggesting that singletons were successfully suppressed for both filled and unfilled displays.

For transparency, an outlier participant had a -60.5% suppression effect (i.e., a large capture effect) in the filled condition that was 4.9 standard deviations from the group mean. If excluded, the basic pattern of results remains the same, except that the probe suppression effect is larger and more robust for filled displays (7.2%), $t(30) = 6.33$, $p < .001$, $d_z = 1.14$. In short, excluding this outlier would not change the results and would only strengthen the observed probe suppression effect in the filled condition. In any case, the preregistration contained no a priori reason for excluding extreme outliers on probe report accuracy.

We also assessed whether the target shape was enhanced above baseline levels of the nonsingleton distractors. Target enhancement effects were computed as probe report accuracy at the target location minus the average nonsingleton distractor (Gaspelin et al., 2015). A positive score indicates enhancement of the target shape, whereas a negative score would indicate suppression of the target shape. One-sample *t* tests indicated that target enhancement effects were significant for both filled displays (41.9%), $t(31) = 9.41$, $p < .001$, $d_z = 1.66$, and unfilled displays (31.0%), $t(31) = 8.06$, $p < .001$, $d_z = 1.43$. Target enhancement effects were smaller for unfilled displays than filled displays, $t(31) = 4.22$, $p < .001$, $d_z = .75$. This, along with the finding that overall RTs were increased for unfilled shapes in Experiments 3 and 4, suggests that the unfilled shape condition may have been a more difficult search. This is likely because the unfilled shapes are more difficult to see parafoveally, and this may weaken top-down guidance toward the target shape.

Discussion

Experiment 4 tested whether the color singletons would elicit probe suppression effects, as predicted by the signal suppression hypothesis (Gaspelin et al., 2015). Consistent with this prediction, probe report accuracy for letters at the singleton distractor were lower than the baseline level of that for the nonsingleton distractor, regardless of display type. These results indicate that the color singletons were successfully suppressed to prevent capture, regardless of display type.

GENERAL DISCUSSION

The signal suppression hypothesis was developed to help resolve the attentional capture debate. Although it has garnered much empirical support (see reviews by Gaspelin & Luck, 2018d, 2019; Luck et al., 2021), it has been challenged on the grounds that the singletons used in studies supporting it might have been weakly salient (Wang & Theeuwes, 2020). Although the role of salience in capture is not a new concern (e.g., Moher et al., 2015; Nothdurft, 1993; Yantis & Egeth, 1999), this criticism has led to many recent investigations of whether highly salient stimuli can automatically overpower suppression and capture attention (e.g., Drisdelle & Eimer, 2023; Hauck et al., 2023; Lien et al., 2022; Ramgir & Lamy, 2023a; Stilwell et al., 2022, 2023; Stilwell & Gaspelin, 2021). A major shortcoming of these studies, including our own, has been the lack of an independent measure of salience. This has made it difficult to evaluate whether display size actually influences salience as was claimed by Wang and Theeuwes (2020). The current study therefore used a new psychophysical technique to measure salience to test the claims of the low-salience criticism.

Experiment 1 used the oddball detection task to compare the salience of stimuli in Gaspelin et al. (2015) and Wang and Theeuwes (2020). In this task, participants attempted to detect color singletons in briefly presented displays. Importantly, the singletons could only be detected based upon their color popout (i.e., their salience). According to the low-salience criticism, the displays in Wang and Theeuwes increased the salience of the singletons compared to Gaspelin et al. via a manipulation of display size. If true, the singletons should have been easier to detect, resulting in shorter exposure thresholds in the oddball detection task. Surprisingly, we found no evidence for this claim: the exposure thresholds to detect color singletons were actually shorter for the Gaspelin et al. displays than the Wang and Theeuwes displays, indicating that salience was actually stronger in the Gaspelin et al. displays.

Experiment 2 explored potential reasons why the Gaspelin et al. (2015) displays were more salient, despite the lower display size. One hypothesis is that the unfilled shapes in Wang and Theeuwes (2020) led to weak color singletons by reducing the amount of color information available to produce a color popout. To test this, we manipulated both display size (4 vs. 10) and color fill (unfilled vs. filled) in the displays that were otherwise identical. The oddball detection task revealed that both factors influenced the ability of a color singleton to be detected. Importantly, the color fill seemed to produce the largest effects suggesting that it can powerfully influence salience. In fact, filled singletons at display size 4 were detected more quickly than unfilled singletons at display size 10. These results suggest that Wang and Theeuwes (2020) may have failed to maximize salience due to a lack of filled stimuli.

Experiments 3 and 4 tested whether the high-salience displays generated in Experiment 2 would lead to capture. In Experiment 3, participants performed an additional singleton paradigm in which they searched for a target shape (e.g., green circle) and attempted to ignore the salient distractor. To maximize salience, we used high set-size displays and then compared whether improving salience by using filled versus unfilled shapes would modulate capture. We found no evidence of capture in either display, despite the apparent high salience of these displays in Experiment 2. In Experiment 4, a capture-probe paradigm was used to demonstrate that the singleton was suppressed below baseline levels for both filled and unfilled displays. This result provides further evidence against the low-salience criticism which directly proposes that singletons with sufficient salience should automatically capture attention.

Altogether, the results provide evidence against the low-salience criticism of signal suppression. First, the oddball detection task suggested that both Gaspelin et al. (2015) and Wang and Theeuwes (2020) used highly salient singletons, as the exposure threshold for both stimuli were relatively low. Similarly, computational models of salience have also suggested that the color singletons in the original signal suppression studies were highly salient (Chang et al., 2021). Second, additional evidence against the low-salience criticism comes from the finding that highly salient singletons produced no evidence of capture in Experiments 3 and 4. This has also been suggested by many other studies which have found that improving the salience of distractors does not automatically lead to capture (Drisdelle & Eimer, 2023; Gaspar & McDonald, 2014; Moher et al., 2015; Stilwell et al., 2022, 2023; Stilwell & Gaspelin, 2021). This latter finding challenges stimulus-driven accounts, more generally, which propose that salient stimuli should automatically capture attention.

The lack of a capture effect in the unfilled condition of Experiment 3 was somewhat surprising, given that Wang and Theeuwes (2020) reported a 43-ms capture effect using the same stimuli. The current study used more trials (1040 trials compared to 480 trials) and more participants (*N* = 48 compared to *N* = 24) than Wang and Theeuwes (2020), which should only reduce the likelihood of a Type I or II error. Other studies using these stimuli have also produced evidence that the capture effects were either more short-lived (Ramgir & Lamy, 2023) or weaker in magnitude (Stilwell & Gaspelin, 2021) than in the original study. It is also possible that other unknown factors might have also led to the discrepant results, such as differences in the ways participants were instructed or participants' level of experience with psychophysical tasks. In any case, the current results provide evidence that capture does not mandatorily occur at high display sizes, as was initially claimed by the low-salience account.

One question is whether the oddball detection task directly measures salience or instead measures salience indirectly via some other cognitive process. It is important to highlight that our task did vary the color singleton, meaning that it had to be detected based upon its contrast with other objects and could not be detected, for example, via its specific color. One possible explanation for why exposure thresholds were shorter in the Gaspelin et al. displays than in the Wang and Theeuwes displays in Experiment 1is that display objects were compared in a serial manner to determine whether one object mismatches the color of other objects. This could explain why exposure thresholds were faster for the Gaspelin et al. displays than Wang and Theeuwes displays in Experiment—there were fewer comparisons to make in the former than the latter. There are, however, problems with such an explanation. Most notably, Experiment 2 directly contradicts this explanation by showing that increasing display size *decreased* exposure thresholds. If there was a serial comparison process, exposure thresholds should increase with display size because there are, on average, more comparisons needed to determine whether or not a differently colored item was presented. An additional problem is that such an explanation would require multiple serial shifts of covert attention between objects to compare their colors. The exposure thresholds, which were as fast as 30 ms in some conditions, were simply too fast to permit multiple shifts of covert attention, which are estimated to take 35–100 ms to execute (Horowitz et al., 2009). That being said, we remain agnostic as to whether attention was distributed serially or in parallel in the oddball detection task. Although we contend that our measure is a reasonable proxy measure for salience, more research is needed to provide a definitive linkage.

There is still much to be understood about how salience is computed at a cognitive level. There has been some recent evidence from studies of monkey neurophysiology that bottom-up salience may be computed rapidly in V4 (e.g., within 50-60 ms after stimulus onset; Westerberg et al., 2023) and can be strongly modulated by prior selection history (see also Adam & Serences, 2021). Other recent evidence from fMRI decoding has suggested that different forms of feature singletons (e.g., motion vs. color) may have distinct representations in visual cortex that feed into a final priority map (Thayer & Sprague, 2023). Our new oddball detection technique could be used to help improve linkages between perceived differences in salience with neural measures of salience.

A few prior studies have used computational models of to verify the salience of color singletons (Chang et al., 2021; Stilwell et al., 2022; Stilwell & Gaspelin, 2021) and it is interesting to consider how computational models may relate to the current psychophysical approach. For example, Chang et al. (2021) applied two computational models of salience (Jeck et al., 2019; Russell et al., 2014) to compare the displays of Wang and Theeuwes (2021) and Gaspelin et al. (2015). The results suggested that the color singletons used in both studies were salient, which fits well with the conclusions of the current Experiment 1. However, there are two major limitations with using computational models of salience to resolve debates about attentional capture. First, computational models of salience often do not perform well with artificial displays such as those used in laboratory experiments (e.g., Jeck et al., 2019; Kotseruba et al., 2020). Second, there are many computational models to choose from and these models have several free parameters that are set by the user. Thus, conclusions drawn from these models will depend on both the veridicality of the specific model used and the parameters selected during the analysis. It is therefore difficult to imagine how computational models alone could definitively resolve debates in attentional capture about salience, which strongly highlights the need for psychophysical methods to measure salience.

A potential issue, pointed out by a reviewer, is that the low-salience account could always argue that a given manipulation of salience was not strong enough. In the current study, this would be an unusual position to take given we used the exact same stimuli that were previously purported to be "salient enough." (i.e., Wang & Theeuwes, 2020). Regardless, the current study provides clear evidence against the claim of Wang and Theeuwes (2020) that salience was improved compared to the initial Gaspelin et al. (2015) study. Additionally, we believe some consideration of falsifiability of the low-salience account is needed. Given that highly salient stimuli did not capture attention in the current study and many of our previous studies (Stilwell et al., 2022, 2023; Stilwell & Gaspelin, 2021; Zhang & Gaspelin, under revision), we see little evidence that increasing the salience of a color singleton can cause it to capture attention.

In conclusion, the current study refutes the low-salience criticism of the signal suppression hypothesis by showing that the purported manipulations of salience may have not been as successful as previously thought. Our findings suggest that display size may not be sufficient to improve salience if other factors are not controlled (e.g., filling in the stimuli with color). Furthermore, these findings highlight the need to empirically verify manipulations of salience when testing how salience influences attentional control.

Author Note

This project was supported by the National Science Foundation Grant BCS-2345898 to N.G. This project was conducted while N.G. and B.T.S. were at State University of New York at Binghamton.

Correspondence concerning this article should be directed to Nicholas Gaspelin, Department of Psychological Sciences, McAlester Hall, 320 S. Sixth Street, University of Missouri, Columbia, MO, 65211 email: ngaspelin@missouri.edu, phone: (573) 882-1445.

REFERENCES

Adam, K. C. S., & Serences, J. T. (2021). History modulates early sensory processing of salient distractors. *Journal of Neuroscience*.

Adams, O. J., Ruthruff, E., & Gaspelin, N. (2022). Oculomotor suppression of abrupt onsets versus color singletons. *Attention, Perception, & Psychophysics*. https://doi.org/10.3758/s13414-022-02524-0

Bacon, W. F., & Egeth, H. E. (1991). Local processes in preattentive feature detection. *Journal of Experimental Psychology: Human Perception and Performance*, *17*(1), 77.

Bacon, W. F., & Egeth, H. E. (1994). Overriding stimulus-driven attentional capture. *Perception & Psychophysics*, *55*(5), 485–496.

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, *57*(1), 289–300.

Benjamini, Y., & Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics*, 1165–1188.

Bravo, M. J., & Nakayama, K. (1992). The role of attention in different visual-search tasks. *Perception & Psychophysics*, *51*(5), 465–472.

Chang, S., & Egeth, H. E. (2019). Enhancement and Suppression Flexibly Guide Attention. *Psychological Science*. https://doi.org/10.1177/0956797619878813

Chang, S., & Egeth, H. E. (2021). Can salient stimuli really be suppressed? *Attention, Perception, & Psychophysics*, *83*(1), 260–269. https://doi.org/10.3758/s13414-020-02207-8

Chang, S., Niebur, E., & Egeth, H. E. (2021). *Standing out in a small crowd: The role of display size in attracting attention. Visual Cognition*.

Cosman, J. D., Lowe, K. A., Woodman, G. F., & Schall, J. D. (2018). Prefrontal Control of Visual Distraction.

Current Biology, *28*(3), 414-420.e3. https://doi.org/10.1016/j.cub.2017.12.023

Drisdelle, B. L., & Eimer, M. (2021). PD components and distractor inhibition in visual search: New evidence for the signal suppression hypothesis. *Psychophysiology*, *58*(9), e13878. https://doi.org/10.1111/psyp.13878

Drisdelle, B. L., & Eimer, M. (2023). Proactive suppression can be applied to multiple salient distractors in visual search. *Journal of Experimental Psychology. General*, *152*(9), 2504–2519. https://doi.org/10.1037/xge0001398

Duncan, J., & Humphreys, G. W. (1989). Visual Search and Stimulus Similarity. *Psychological Review*, *96*(3), 433–458. https://doi.org/10.1037/0033-295X.96.3.433

Feldmann-Wüstefeld, T., Busch, N. A., & Schubö, A. (2020). Failed suppression of salient stimuli precedes behavioral errors. *Journal of Cognitive Neuroscience*, *32*(2), 367–377.

Folk, C. L., & Remington, R. (2010). A critical evaluation of the disengagement hypothesis. *Acta Psychologica*, *135*(2), 103–105. https://doi.org/10.1016/j.actpsy.2010.04.012

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*(4), 1030–1044. https://doi.org/10.1037/0096-1523.18.4.1030

Gaspar, J. M., & McDonald, J. J. (2014). Suppression of Salient Objects Prevents Distraction in Visual Search. *Journal of Neuroscience*, *34*(16), 5658–5666. https://doi.org/10.1523/JNEUROSCI.4161-13.2014

Gaspelin, N., Gaspar, J. M., & Luck, S. J. (2019). Oculomotor inhibition of salient distractors: Voluntary inhibition cannot override selection history. *Visual Cognition*, *27*(3–4), 227–246. https://doi.org/10.1080/13506285.2019.1600090

Gaspelin, N., Lamy, D., Egeth, H. E., Liesefeld, H. R., Kerzel, D., Mandal, A., Müller, M. M., Schall, J. D., Schubö, A., Slagter, H. A., Stilwell, B. T., & van Moorselaar, D. (2023). The Distractor Positivity Component and the Inhibition of Distracting Stimuli. *Journal of Cognitive Neuroscience*, *35*(11), 1693–1715. https://doi.org/10.1162/jocn_a_02051

Gaspelin, N., Leonard, C. J., & Luck, S. J. (2015). Direct evidence for active suppression of salient-but-irrelevant sensory inputs. *Psychological Science*, *22*(11), 1740–1750.

Gaspelin, N., Leonard, C. J., & Luck, S. J. (2017). Suppression of overt attentional capture by salient-but-irrelevant color singletons. *Attention, Perception, and Psychophysics*, *79*(1), 45–62. https://doi.org/10.3758/s13414-016-1209-1

Gaspelin, N., & Luck, S. J. (2018a). Combined electrophysiological and behavioral evidence for the suppression of salient distractors. *Journal of Cognitive Neuroscience*, *30*(9), 1265–1280. https://doi.org/10.1162/jocn_a_01279

Gaspelin, N., & Luck, S. J. (2018b). Distinguishing among potential mechanisms of singleton suppression. *Journal of Experimental Psychology: Human Perception and Performance*. https://doi.org/10.1037/xhp0000484

Gaspelin, N., & Luck, S. J. (2018c). Distinguishing among potential mechanisms of singleton suppression. *Journal of Experimental Psychology: Human Perception and Performance*, *44*(4), 626–644.

Gaspelin, N., & Luck, S. J. (2018d). The Role of Inhibition in Avoiding Distraction by Salient Stimuli. *Trends in Cognitive Sciences*, *22*(1), 79–92. https://doi.org/10.1016/j.tics.2017.11.001

Gaspelin, N., & Luck, S. J. (2019). Inhibition as a potential resolution to the attentional capture debate. *Current Opinion in Psychology*, *29*, 12–18. https://doi.org/10.1016/j.copsyc.2018.10.013

Green, D. M., & Swets, J. A. (1974). *Signal detection theory and psychophysics.* Robert E. Krieger.

Greene, M. R., & Oliva, A. (2009). The briefest of glances: The time course of natural scene understanding. *Psychological Science*, *20*(4), 464–472.

Hamblin-Frohman, Z., Chang, S., Egeth, H., & Becker, S. I. (2022). Eye movements reveal the contributions of early and late processes of enhancement and suppression to the guidance of visual search. *Attention, Perception, & Psychophysics*, *84*(6), 1913–1924.

Hauck, C., Ruthruff, E., & Lien, M.-C. (2023). On preventing capture: Does greater salience cause greater suppression? *Attention, Perception, & Psychophysics*, *85*(8), 2553–2566. https://doi.org/10.3758/s13414-023-02694-5

Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values ofd′. *Behavior Research Methods, Instruments, & Computers*, *27*(1), 46–51. https://doi.org/10.3758/BF03203619

Horowitz, T. S., Wolfe, J. M., Alvarez, G. a, Cohen, M. A., & Kuzmova, Y. I. (2009). The speed of free will. *The Quarterly Journal of Experimental Psychology*, *62*(August 2012), 2262–2288. https://doi.org/10.1080/17470210902732155

Jeck, D. M., Qin, M., Egeth, H., & Niebur, E. (2019). Unique objects attract attention even when faint. *Vision Research*, *160*, 60–71.

Kaernbach, C. (1991). Simple adaptive testing with the weighted up-down method. *Perception & Psychophysics*, *49*(3), 227–229.

Kim, M.-S., & Cave, K. R. (1995). Spatial Attention in Visual Search for Features and Feature Conjunctions. *Psychological Science*, *6*(6), 376–380. https://doi.org/10.1111/j.1467-9280.1995.tb00529.x

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

Kotseruba, I., Wloka, C., Rasouli, A., & Tsotsos, J. K. (2020). Do Saliency Models Detect Odd-One-Out Targets? New Datasets and Evaluations. *arXiv Preprint arXiv:2005.06583*.

Leber, A. B., & Egeth, H. E. (2006). It's under control: Top-down search strategies can override attentional capture. *Psychonomic Bulletin & Review*, *13*(1), 132–138. https://doi.org/10.3758/BF03193824

Lien, M.-C., Ruthruff, E., & Hauck, C. (2022). On preventing attention capture: Is singleton suppression actually singleton suppression? *Psychological Research*, *86*(6), 1958–1971. https://doi.org/10.1007/s00426-021-01599-y

Loftus, G. R., Johnson, C. A., & Shimamura, A. P. (1985). How much is an icon worth? *Journal of Experimental Psychology: Human Perception and Performance*, *11*, 1–13. https://doi.org/10.1037/0096-1523.11.1.1

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

Ma, X., & Abrams, R. A. (2022). Ignoring the unknown: Attentional suppression of unpredictable visual distraction. *Journal of Experimental Psychology: Human Perception and Performance*.

Ma, X., & Abrams, R. A. (2023). Visual distraction's "silver lining": Distractor suppression boosts attention to competing stimuli. *Psychological Science*. https://doi.org/10.177/09567976231201853

Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition*, *22*(6), 657–672. https://doi.org/10.3758/BF03209251

Moher, J., Anderson, B. A., & Song, J.-H. (2015). Dissociable effects of salience on attention and goal-directed action. *Current Biology*, *25*(15), 2040–2046.

Morey, R. D., Romeijn, J.-W., & Rouder, J. N. (2016). The philosophy of Bayes factors and the quantification of statistical evidence. *Journal of Mathematical Psychology*, *72*, 6–18. https://doi.org/10.1016/j.jmp.2015.11.001

Nothdurft, H.-C. (1993). The role of features in preattentive vision: Comparison of orientation, motion and color cues. *Vision Research*, *33*(14), 1937–1958. https://doi.org/10.1016/0042-6989(93)90020-W

Oxner, M., Martinovic, J., Forschack, N., Lempe, R., Gundlach, C., & Müller, M. (2022). Global enhancement of target color-not proactive suppression-explains attentional deployment during visual search. *Journal of Experimental Psychology: General*.

Ramgir, A., & Lamy, D. (2023). Distractor's salience does not determine feature suppression: A commentary on Wang and Theeuwes (2020). *Journal of Experimental Psychology: Human Perception and Performance*, *49*(6), 852–861. https://doi.org/10.1037/xhp0001119

Russell, A. F., Mihalaş, S., von der Heydt, R., Niebur, E., & Etienne-Cummings, R. (2014). A model of proto-object based saliency. *Vision Research*, *94*, 1–15.

Savelson, I., & Leber, A. B. (under review). *How we learn to ignore singleton distractors: Suppressing saliency signals or specific features?*

Sawaki, R., & Luck, S. J. (2010). Capture versus suppression of attention by salient singletons: Electrophysiological evidence for an automatic attend-to-me signal. *Attention, Perception, and Psychophysics*, *72*(6), 1455–1470. https://doi.org/10.3758/APP.72.6.1455

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, *74*(11), 1–29.

Stilwell, B. T., Adams, O. J., Egeth, H. E., & Gaspelin, N. (2023). The role of salience in the suppression of distracting stimuli. *Psychonomic Bulletin & Review*. https://doi.org/10.3758/s13423-023-02302-5

Stilwell, B. T., Bahle, B., & Vecera, S. P. (2019). Feature-based statistical regularities of distractors modulate attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(3), 419.

Stilwell, B. T., Egeth, H., & Gaspelin, N. (2022). Electrophysiological Evidence for the Suppression of Highly Salient Distractors. *Journal of Cognitive Neuroscience*, *34*(5), 787–805. https://doi.org/10.1162/jocn_a_01827

Stilwell, B. T., & Gaspelin, N. (2021). Attentional suppression of highly salient color singletons. *Journal of Experimental Psychology: Human Perception and Performance*, *47*(10), 1313–1328. https://doi.org/10.1037/xhp0000948

Thayer, D. D., & Sprague, T. C. (2023). Feature-specific salience maps in human cortex. *Journal of Neuroscience*, *43*(50), 8785–8800.

Theeuwes, J. (1991). Exogenous and endogenous control of attention: The effect of visual onsets and offsets. *Perception & Psychophysics*, *49*(1), 83–90. https://doi.org/10.3758/BF03211619

Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics*, *51*(6), 599–606. https://doi.org/10.3758/BF03211656

Theeuwes, J. (2004). Top-down search strategies cannot override attentional capture. *Psychonomic Bulletin and Review*, *11*(1), 65–70. https://doi.org/10.3758/BF03206462

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.

Vatterott, D. B., & Vecera, S. P. (2012). Experience-dependent attentional tuning of distractor rejection. *Psychonomic Bulletin & Review*, *19*(5), 871–878. https://doi.org/10.3758/s13423-012-0280-4

Wang, B., & Theeuwes, J. (2020). Salience determines attentional orienting in visual selection. *Journal of Experimental Psychology: Human Perception and Performance*, *46*(10), 1051–1057. https://doi.org/10.1037/xhp0000796

Westerberg, J. A., Schall, J. D., Woodman, G. F., & Maier, A. (2023). Feedforward attentional selection in sensory cortex. *Nature Communications*, *14*(1), 5993.

Won, B.-Y., Kosoyan, M., & Geng, J. J. (2019). Evidence for second-order singleton suppression based on probabilistic expectations. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(1), 125–138. psyh. https://doi.org/10.1037/xhp0000594

Yantis, S., & Egeth, H. E. (1999). On the distinction between visual salience and stimulus-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *25*(3), 661–676. https://doi.org/10.1037/0096-1523.25.3.661

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, Y., & Gaspelin, N. (n.d.). *Salience effects are enabled by task relevance*.