

Clarifying the role of target similarity, task relevance and feature-based suppression during sustained inattentive blindness

Trafton Drew

Department of Psychology, University of Utah,
Salt Lake City, UT, USA



Cary Stothart

Department of Psychology, University of Notre Dame,
Notre Dame, IN, USA



How is feature-based attention distributed when engaged in a challenging attentional task? Thanks to formative electrophysiological and psychophysical work, we know a great deal about the spatial distribution of attention, but much less is known about how feature-based attention is allocated. In a large-scale online study, we investigated the distribution of attention to color space using a sustained inattentive blindness task. In order to query what parts of color space were being attended or inhibited, we varied the color of an unexpected stimulus on the final trial. Noticing rates for this stimulus indicate that when engaged in a difficult task that involves tracking items of one color and ignoring items of two different colors, observers attend the target color and inhibit the to-be ignored colors. Further, similarity to the target drives detection such that colors more similar to the target are more likely to be detected. Finally, our data suggest that when possible, observers inhibit regions of color space rather than individuating specific colors and adjusting the level of inhibition for a particular color accordingly. Together, our data support the notion of feature-based suppression for task relevant (to-be ignored) information, but we found no evidence of an inhibitory surround based on target color similarity.

thought to increase the neural response to particular features, attentional suppression is thought to push the neural response in the opposite direction. Whereas attentional facilitation is thought to help the observer find the target more efficiently, attentional suppression is proposed to decrease the likelihood of confusing a distractor with a target. Importantly, just as attentional facilitation is dependent upon having a clear target representation, recent evidence suggests that attentional suppression depends on a clear representation of what features may be mistaken for a target.

A vivid illustration of this effect comes from the work of Jens-Max Hopf and colleagues (Hopf, Boehler, & Luck, 2006; Hopf, Boehler, Schoenfeld, Heinze, & Tsotsos, 2010). Prior to this work, predominant theories suggested a simple monotonic distribution of spatial attention such that the amount of attention devoted to an area peaks at attended locations and decreases at locations further away (Eriksen & St James, 1986; Eriksen & Yeh, 1985; Heinze et al., 1994; Posner, Snyder, & Davidson, 1980). However, when asked to perform a difficult detection task at a consistent location, Hopf and colleagues found that the neural response was most consistent with an excitatory peak followed by an inhibitory surround. They argued that performance on this task was driven by at least two distinct processes: an excitatory effect that enhances information in the attended location, and an inhibitory effect tuned to the region just outside of the target region. This is consistent with Navalpakkam & Itti's optimal feature gain modulation theory, which seeks to optimize information from both targets and distractors to maximize target salience (Navalpakkam & Itti, 2007).

Until recently, it was unclear whether such an inhibitory surround was a unique characteristic of spatial attention or a more general attentional mech-

Introduction

One of the primary functions of attention is to facilitate processing of features thought to be indicative of important information, such as the sound of your ringtone while waiting for a phone call. Whereas there is a large amount of behavioral and neural data demonstrating the importance of this role of attention, there is increasing evidence that attentional suppression is equally important. Whereas attentional facilitation is

Citation: Drew, T., & Stothart, C. (2016). Clarifying the role of target similarity, task relevance, and feature-based suppression during sustained inattentive blindness. *Journal of Vision*, 16(15):13, 1–9, doi:10.1167/16.15.13.

doi: 10.1167/16.15.13

Received July 29, 2016; published December XX, 2016

ISSN 1534-7362



This work is licensed under a Creative Commons Attribution 4.0 International License.

anism. Based on recent evidence, it appears clear that suppression plays an important role in feature-based attention as well. Störmer and Alvarez (2014) recently found both behavioral and electrophysiological evidence that distractor items with features similar to the target were suppressed relative to items with features more distinct from target features during a visual search task. Similarly Moher and colleagues (Moher, Lakshmanan, Egeth, & Ewen, 2014) found that the neural response to distractor-colored probes was reliably suppressed relative to both neutral and target-colored probes. In both cases, feature-based suppression was found in the context of a task where observers were asked to detect a brief, subtle change in the attended objects.

While these findings clearly demonstrate that feature-based suppression exists, part of the reason that these results are so interesting is that they suggest that feature-based suppression is a fundamental part of our attentional system and therefore has important implications for how we interact with the world. Therefore, the purpose of the current paper was to determine whether these effects generalize to a situation that, while still quite artificial, moves a number of important steps closer to the real world. For example, there are many vivid illustrations that show that when we are looking for one thing, attention renders us effectively “blind” to otherwise salient stimuli such as bears riding unicycles (Hyman, Boss, & Wise, 2009), fights breaking out (Chabris, Weinberger, Fontaine, & Simons, 2011), and gorillas—even if they appear in medical images (Drew, Vo, & Wolfe, 2013; Simons & Chabris, 1999). These demonstrations of inattention blindness (IAB) are thought to be caused by the same basic attentional mechanisms outlined previously.

An important distinction between these inattention blindness paradigms and previously mentioned demonstrations of feature-based suppression is the time-scale of the task. The paradigm employed by Moher and colleagues asked observers to detect a brief (500 ms) luminance change during a long (15 s) trial. The Störmer and Alvarez (2014) paradigm asked observers to monitor a relatively brief (2.6 s) random dot kinematogram for a brief (230 ms) period of coherent motion. In contrast, the target in IAB paradigms is typically fully visible for multiple seconds. Whereas understanding the ability to quickly detect a brief stimulus provides an important window into the initial processing of visual stimuli, with a longer time-course the IAB paradigm provides a window into what may be the consequences of this initial processing on detection of stimuli that are fully visible and easy to detect in the absence of another task.

The purpose of the current study is to use the inattention blindness paradigm to uncover the rules that govern attentional deployment during a sustained

attention task. This is not the first study to use this approach. In a series of important experiments, Most and colleagues (Most et al., 2001) manipulated the relationship between the tracked (target) items, distractors, and the unexpected stimulus (US). These authors were the first to employ a simplified IAB experimental paradigm that asked observers to count the number of times target items of a particular color bounced off the edge of the movement area on a computer screen. This approach has a number of advantages over the more traditional method of displaying prerecorded videos (e.g., Simons & Chabris, 1999) or asking trained confederates to replay a specific scenario for each new observer (Chabris et al., 2011; Hyman et al., 2009). Whereas these sorts of realistic stimuli are useful for demonstrating that certain objects are or are not perceived in a given context, using computer-generated displays allowed the fine level of control necessary to address the interest of the current work: the distribution of feature-based attention during a sustained attention task.

Most and colleagues (2001) created a modified version of the IAB task to explore the role of selectivity and selective ignoring on unexpected stimulus detection. By manipulating the color of the US across observers, they found that the relationship between the target color and the US largely determined the likelihood that the US was detected: When the target was white and the US was white, almost all observers (94%) detected the US. Similarly, when Simons and Chabris asked their subjects to count the number of passes and dribbles from members of the “black team” rather than the “white team” (while detecting a black gorilla), the rate of IAB was much reduced relative to the archetypical example where observers are asked to track the “white team” (Simons & Chabris, 1999).

In a second experiment, Most and colleagues (2001) manipulated whether the US was the same or a different color than the irrelevant distractor letter items that shared the screen with a set of gray target letters. Observers were instructed to count the number of bounces of the gray target letters while ignoring the distractor letters. They found that the relationship between the distractor letters and the US had a dramatic effect on US detection: When the US was the same color as the distractor stimuli, detection was very low. When they were different, detection was high. These results suggest that IAB rates may be driven, at least in part, by an active suppression of information that is thought to be irrelevant to the task at hand, but there is at least one other explanation: Perhaps the observed effect was driven by the novelty of the US when it did not match the distractor items. For example, if the target items are gray and the distractor items are black, when the white US enters the screen, it is the first white stimulus to be seen by the observer,

Exclusion rule	Excluded	
	%	<i>N</i>
Failed an attention test	4.76	55
Reported not having normal or corrected-to-normal vision	4.68	54
Failed a colorblindness test	2.86	33
Already participated in the study	4.94	57
Reported being familiar with the inattentive blindness task	6.06	70
Reported not using either a laptop or desktop computer	1.56	18
Failed to get 100% on the color discrimination task	21.13	244
Reported the task had a glitch or provided an unexplained open response	1.47	17
Total excluded	38.7	447

Table 1. Exclusion rules. *Notes:* Observers could be excluded for multiple reasons.

which may thereby automatically attract attention and increase the noticing rate.

To address this concern, Most and colleagues (2001) ran an additional study where they measured the rate of IAB for a red stimulus while tracking black/white stimuli while simultaneously ignoring white/black stimuli. They found that noticing rates for the red item were actually lower than when the US was black and the distractor stimulus was white (or vice versa). Thus, they concluded that novelty alone did not drive the high detection rate of the US in an otherwise irrelevant color. One concern over this conclusion is that it was conducted across two experiments with different stimulus parameters, thereby complicating the interpretation of IAB rates under these different conditions. We will return to the issue in the Discussion.

In order to examine the role of feature-based suppression during IAB and address these difficulties in comparing performance across experiments, we modified Most et al.'s (2001) paradigm such that all items (including the US) were chosen from an equiluminant ring through color space. This allowed us to more directly address the concern that evidence for suppression of a specific color may be inflated by the novelty of a control color. Moreover, this design allowed us to examine a number of conditions that were difficult to address using the mostly gray-scale stimuli employed in previous work. Finally, by adding a second distractor color, we believe the current design allows for a cleaner method of determining whether there is indeed evidence for feature-based suppression during inattentive blindness. Interestingly, the data appear to support the notion of suppression based on target relevance, but not due to target similarity.

Methods

We preregistered the study on the Open Science Framework where we provide the materials, analyses,

and data (<https://osf.io/ywnjp>). The study was approved by the Institutional Review Board at Florida State University where the second author is a former graduate student and current research consultant. A waiver of the requirement for signed consent was approved by this board as participation in the study was anonymous.

Observers

Observers were from the United States and completed the study online using Amazon Mechanical Turk. We collected data on 1155 observers with the goal of obtaining useable data from at least 100 observers per condition. We considered a useable observer to be one that didn't meet any of the exclusion rules listed in Table 1. A total of 744 observers didn't meet any of these rules and were included in the analysis (32.39% male, 67.61% female; median age = 32, *IQR* = 14).

Inattentive blindness task

Figure 1 shows a screenshot of the task. Observers viewed a 666 pixels × 546 pixels gray (#777777) display and visually tracked a set of four moving objects (2 Ts and 2 Ls; 43 pixels × 43 pixels) while ignoring two other sets of four moving objects. Each set had a unique color, and the color of the sets differed in their position on a color wheel in HSV color space (see Figure 2; value = 100%, saturation = 100%). The color of the target set was at 0°, the color of one of the ignored sets was at 330° (near distractors), and the color of the other ignored set was at 195° (far distractors). The color wheel was randomly rotated across observers. The objects randomly moved around the display at speeds ranging from 30 pixels/s to 90 pixels/s, changing direction and/or speed every 1 to 4 s and bouncing off of the display's edges when they came into contact.

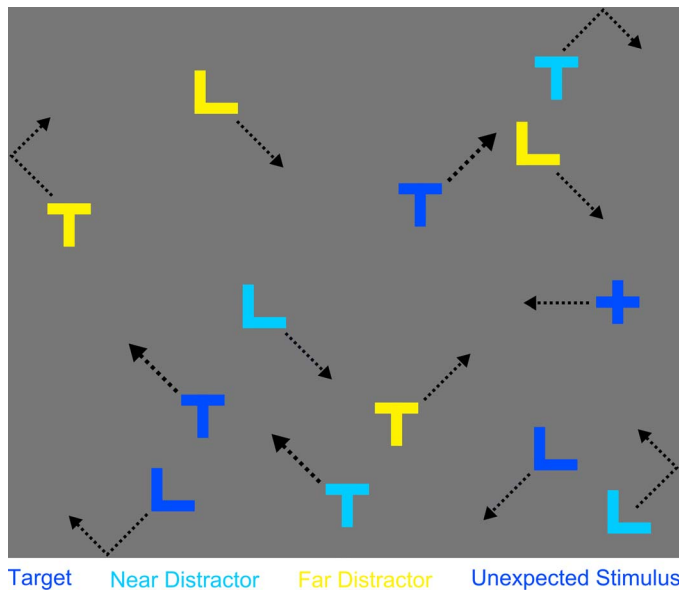


Figure 1. The inattention blindness task. The US in this example is in the Target Color. See text for additional details.

There were a total of six, 12-s trials. During each trial, observers counted the number of times a target object bounced off of one of the display's edges. There was an average of 8.2 bounces per trial (minimum: 0; maximum: 22). One second into the last trial, an unexpected cross (43 pixels \times 43 pixels) entered the display from the right, moved towards the left at 90 pixels/s, and disappeared once it reached the left side of the display. This object took 7.4 s to cross the display. And, its color was either at the 0°, 15°, 30°, 165°, 195°, 330°, or 345° mark on the color wheel.

Procedure

Observers first completed a color discrimination task to make sure that they could distinguish the target color from the near distractor color. In this task, they saw a 4 \times 3 array of squares where the color of half of the squares matched the color of the targets, and the color of the other half matched the color of the near distractors. We randomized the position of the squares across observers. Observers were instructed to click on the squares that had the same color as the instructions—the color of the targets. After completing this task, observers completed the inattention blindness task. After the final trial, they were asked: “On that last trial of the task, did you notice anything that was not there on previous trials?” They then answered questions about the features of the unexpected stimulus on a separate screen (color, shape, movement, and direction of movement). They were asked these questions irrespective of their answer on the first question. Observers then answered questions about their vision,

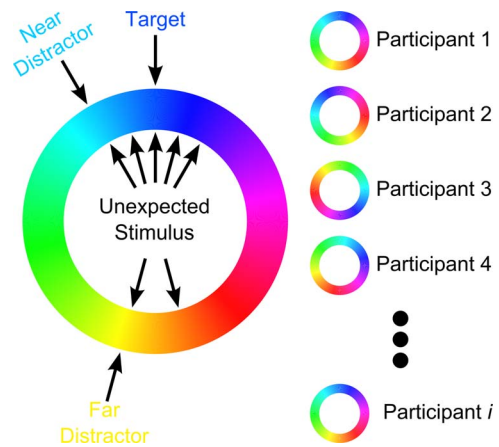


Figure 2. The positions of the task objects on the HSV color wheel. The color wheel was rotated randomly across observers.

computer, and demographics. Finally, they completed an attention test where they selected the middle number in a list of numbers and entered it on a separate screen. We programmed the entire experiment using JavaScript, PHP, and HTML/CSS.

All conditions were preregistered at Open Science Framework with the exception of the 15° “relevant” condition. After finding that noticing rate was unexpectedly high in our 15° condition, we realized that the US in this condition was always in the direction away from the relevant (to-be ignored) items. In order to evaluate whether this direction had a reliable effect on noticing rate, we ran an additional set of observers in the 15° “relevant” condition. Thus, we examined two 15° conditions: a “relevant” condition that was 15° from the target in the direction of the relevant distractors and an irrelevant condition that was 15° away from the relevant distractors.

Results

In order to ensure that the observers were attending to the primary task, observers whose performance on the counting task was far (>50%) from the correct answer on the final trial were excluded from further analysis. However, the pattern of results does not change if we remove this stipulation and examine all the data, or if we adopted a more rigid task performance filter. In order to ensure that the US in different conditions was not affecting performance on the counting task and that there were no differences across groups for ability on the counting task, we performed a one-way ANOVA across US condition for both overall counting performance and performance on the critical trial. There was no effect of condition in either case ($F_s < 1.5$, $p_s > 0.16$).

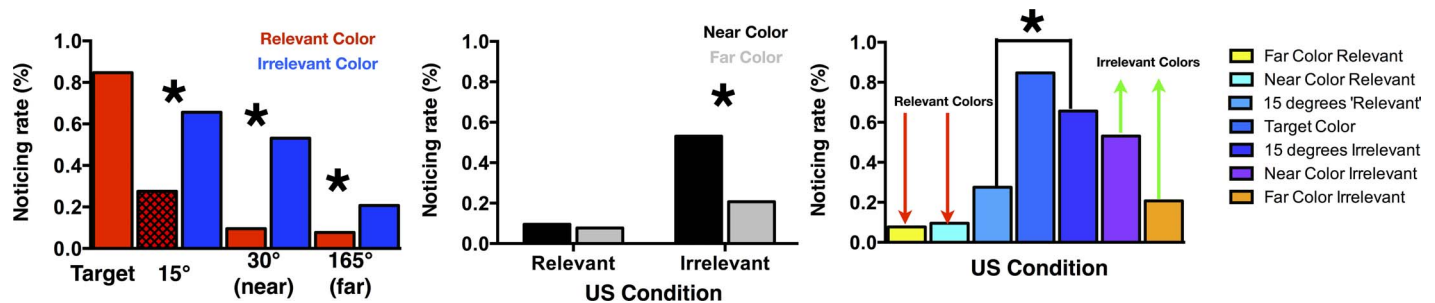


Figure 3. Noticing rates as a function of US color. Data has been replotted in three different ways to highlight comparisons made in the text. Note that in Figure 3a, both conditions at 15° are technically “irrelevant,” but the red hashed bar is on the same side of color space as the to-be ignored distractors. See text for additional details. Stars denote statistically significant differences ($p < 0.05$) highlighted in the text. Colors in Figure 3c denote approximate positions in color space for display purposes.

We compared noticing rates for the different conditions using chi-square tests with Yate’s continuity correction. The results revealed that the relationship between target and the US strongly drove noticing rates, consistent with previous results (Most et al., 2001; Figure 3). Observers were much more likely to notice the US when it was the same color as the target than when it was the same color as either the near, $X^2(1, N = 205) = 105, p < 0.001, \Phi = 0.75$, or far distractor, $X^2(1, N = 205) = 118, p < 0.001, \Phi = 0.77$; see Figure 3a). However, similarity alone did not drive noticing rates: While the irrelevant color and the similar distractor were equally close to the target color (30° away), noticing rates were much higher for the irrelevant color, $X^2(1, N = 190) = 39.7, p < 0.001, \Phi = 0.47$. We also found a similar, though reduced, effect when comparing the Far Color Relevant condition to the Far Color Irrelevant condition, $X^2(1, N = 215) = 6.37, p = 0.01, \Phi = 0.19$. As suggested previously, this result could be driven by the novelty of the “irrelevant” US color. If novelty alone drives noticing rates, there should be no effect of the similarity of the novel color to the target color. However, we found that there was a large effect of similarity to target color, even when the colors were novel: The “near” irrelevant color (55%) was noticed much more frequently than the “far” irrelevant color, 22%, $X^2(1, N = 207) = 22.1, p < 0.001, \Phi = 0.34$; see Figure 3B).

With these data, there are number of ways to evaluate whether there is evidence for feature-based suppression. The simplest method, similar to those used by Störmer and Alvarez is to compare noticing rates for the two types of irrelevant unexpected stimuli. As previously mentioned, we found a large effect of target similarity. However, the effect was in the opposite direction of Störmer and Alvarez: higher detection for near than far. Another way to evaluate the question of feature-based suppression due to target similarity is to examine noticing rates for the near and far relevant (to-be ignored) stimuli. A feature-suppression account would predict that noticing rates for the near relevant

item would be lower than far relevant item. However, we found no difference in noticing rate for the relevant items as a function of proximity to the target, $X^2(1, N = 198) = 0.05, p = 0.83, \Phi = 0.03$.

In order to further evaluate the possibility of feature-based suppression during IAB, we included two additional conditions that were both 15° from the target position. These conditions varied in the *direction* of the additional distractor items that the observers needed to ignore during the counting task. Both conditions were therefore irrelevant to the counting task, but differed in their proximity to other relevant colors. Therefore, if task relevance alone drives detection, performance on these conditions should be equivalent. However, if detection is driven by proximity to relevant colors, the item closer to relevant distractors should be inhibited such that detection for these items should be lower than detection for items that are farther from relevant distractors (but equally close to the target). We found strong evidence that proximity to relevant distractors drove detection: Noticing rates for the 15° “relevant” items was much lower (30%) than for the 15° “irrelevant” ones, 64%, $X^2(1, N = 205) = 28.4, p < 0.001, \Phi = 0.38$.

Discussion

In a large-scale online study, we used the IAB task developed by Most and colleagues (2001) to explore the distribution of feature-based attention during a difficult tracking task. Previous work in the IAB literature has suggested that both task relevance of the US and US similarity to the target drive US detection rates (Most et al., 2001). Coming from a different literature, Störmer and Alvarez (2014), found evidence of feature-based surround suppression such that colors similar to a target color were inhibited relative to a less similar target color by examining both behavioral and electrophysiological data. In order to determine

whether the sort of feature-based surround suppression observed by Störmer exist at the longer time-scale of the IAB paradigm, we varied Most et al.'s paradigm such that observers were asked to track a target color while ignoring two distractor colors. This design allowed us to carefully examine the roles that target relevance, proximity to target color, and novelty played in noticing rates of US of different colors. It also allowed us to more directly compare Störmer and Alvarez's strong evidence for feature-based surround suppression.

Our results largely conform to Most et al.'s (2001) previous work. They found that noticing rates generally decreased as the color of the US moved further from the target color. Specifically, when tracking white targets, Most and colleagues found an overall effect of target similarity as the color of the US went from light gray, to dark gray, and then black. However, whereas the trend is quite clear, the overall effect may be driven by the very low detection rate for black items and high detection for white items, which were task relevant in this paradigm: The task was to track white items amongst black distractors. Thus, this overall effect is at least partially driven by both target similarity and task relevance. In fact, if we reanalyze Most's data by focusing on the light and dark gray conditions, noticing rates do not vary reliably, $X^2(1, N = 32) = 0.55$, $p = 0.46$, $\Phi = 0.20$. This may be driven by the relatively small sample size ($N = 16$ per group). In the current study, given our larger sample size, we can more cleanly address the question by comparing detection rates between two conditions: our irrelevant near and irrelevant far conditions. The large effect of target similarity when comparing near and far irrelevant conditions, $X^2(1, N = 207) = 22.1$, $p < 0.001$, $\Phi = 0.34$ (see Figure 3B), supports Most et al.'s conclusions without any contribution from task relevance.

Most et al. (2001) also found that US relevance played an important role in US noticing rates. Here, when tracking gray items amongst black items, noticing rates for black items was much lower than for white items. However, in this scenario, the high noticing rate for white items may have been driven by novelty: It is the first white item to be observed in the study and may therefore elicit more attention. In order to address this concern, the authors conducted an additional study where the US was a red item while the observers tracked either white among black, or black amongst white items. Noticing rates for the red item were statistically equivalent to noticing rates for a novel stimulus in the previously outlined experiment. Interpreting these results is complicated by the fact that the stimuli and task difficulty varied across the two experiments. Given the evidence that noticing rates vary with difficulty of the IAB task (Simons & Jensen, 2009) and the fact that counting task performance was

not reported by Most et al., it is not currently clear the degree to which novelty contributed to Most's target-relevance result. In the current study, we have four different novel conditions that systematically vary in their similarity to the target color. Examining these conditions, it is again clear that Most's general conclusions hold true: Novelty alone does not drive high levels of noticing. The clearest illustration of this effect is that noticing rates for the far irrelevant color is substantially lower than those for the target and near irrelevant colors.

Most et al.'s (2001) study was not designed to evaluate the possibility of feature-based surround suppression, but based on the results of both Störmer and Moher; the current study evaluated this question for the first time in the IAB paradigm. Both studies provide electrophysiological evidence of suppression for a distractor of a particular color. Interestingly, in both cases, this evidence comes from responses to stimuli on the unattended side of the screen. Given the evidence that feature-based attention appears to be deployed globally (as opposed to space-based attention), examining responses to stimuli on the unattended side of the screen provides a cleaner view of the neural response to specific colors (Saenz, Buracas, & Boynton, 2002, 2003). Moher and colleagues randomized across three different colors throughout their experiments in order to focus on target relevance. They found evidence for suppression of a relevant distractor color that shared space with a target color relative to a task-irrelevant color (Moher et al., 2014). In Störmer's paradigm, the target color shared space with a distractor 180° from the target color. Electrophysiological responses for distractors 30° from the target on the unattended side of screen were lower than responses at 60° and equivalent to responses to the distractor 180° from the target. Thus, while Moher's study provides evidence of feature-based suppression based on task relevance, Störmer's study provides evidence based on target similarity. As outlined previously, our data support the notion of feature-based suppression based on task relevance. This can be seen in the lower noticing rates for relevant than for irrelevant colors at both the near and far color positions.

The evidence for suppression based on target color is more complex. Störmer's critical comparison is 30° to 60° for colors in the unattended field that were not task relevant. The most logical comparison within our own data is the near/far irrelevant color comparison. Noticing rates for these two conditions followed the opposite pattern as Störmer's, with noticing rates higher for the near condition (53%) than the far condition (21%). However, based on Störmer's behavioral data, which suggests that inhibition serves the function of attenuating the response to items most likely to interfere with target representation, perhaps a

more appropriate comparison is the relevant color conditions. The counting task is made much more difficult by these items, and one would imagine attention plays an important role in ensuring that distractor bounces are not inappropriately counted. From this perspective, we should expect that noticing rate for the attended items should follow a pattern similar to the one observed by Störmer and Alvarez: Near distractors suppressed relative to far distractors. However, we found no difference in the noticing rates for near and far relevant distractors. In sum, neither method of probing for evidence of a feature-based inhibitory surround based on target similarity yielded positive evidence that this effect exists in this paradigm.

Noticing rates for the relevant colors suggest that, rather than inhibiting specific colors such that more difficult to ignore items are more inhibited than easier to ignore items, attention may instead inhibit categories of colors as either relevant (to be inhibited), or not. If this is the case, colors that are near task relevant colors, but are not task relevant should be missed more often than equivalent colors that are farther from relevant colors. To test this idea, we asked two groups of observers to detect items that were 15° from the target, either in the direction of the relevant (to be ignored) stimuli or in the opposite direction in color space.

Thus far, noticing rate appears to be driven by target relevance, color similarity, and to a lesser extent, novelty. According to this model, noticing rate for two irrelevant colors that are both 15° from the target should be equivalent. However, if feature-based attention tends to suppress regions of color space (e.g., all the greenish colors rather than pea green and chartreuse) rather than on focusing on particular colors, one would predict lower noticing rates for the color closest to the to-be-ignored stimuli. This is exactly what we found: Noticing rates for the color 15° from the target in the direction of the “relevant” items was much lower (28%) than in the opposite direction (66%), which suggests that feature-based suppression in this task was based on regions of color rather than specific colors that must be ignored.

It would be interesting to explore this new effect in future work. We created a situation where observers were capable of suppressing a part of color space by asking them to ignore two colors on one side of color space. The idea that attention may alter our representation of color space is consistent with some recent work from Golomb (2015). The act of combining colors to a group might be much more difficult if the two colors were on opposite sides of color space. Perhaps in this situation, it would be necessary to differentiate between the unattended colors and we would not see this same pattern of results. Along similar lines, if the observers were asked to track two colors, we would predict high levels of inhibition of a

distractor color between the two target colors, but that pattern should dramatically shift if there were no distractor items in the space between the target colors.

Conclusions

The current study used a simple extension of Most et al.’s (2001) IAB paradigm to explore the distribution of feature-based attention during an attentionally demanding task. We found that both target similarity and task relevance strongly drive rates of noticing unexpected stimuli, replicating, and extending Most et al.’s work. Most notably, by using many different types of US within the same paradigm and many observers (~100) in each condition, we were able to show for the first time that target similarity drives noticing rates even when none of the comparison items are relevant to the counting task. We also found strong evidence for feature-based suppression based on target relevance very similar to the electrophysiological effects found by Moher and colleagues, but did not replicate Störmer and Alvarez’s finding of a band of feature-based suppression based on target similarity. The discrepancy between our result and Störmer and Alvarez is likely driven by the large differences in our tasks. Whereas Störmer’s observers were tasked with identifying brief (230 ms) changes in motion in random dot kinematograms, observers in the current study had 7.4 s to detect the US. It is possible that there is some early suppression of US stimuli that are similar to the target, but that this initial difference was over-ridden by later attentional processing that was not subject to this same sort of inhibitory surround.

Moreover, whereas the target color changed on each trial in Störmer’s study (2014), in the current study the target colors were held constant throughout the experiment. In this way, our study was more similar to Moher and colleagues’ previously discussed study (2014), where target color was consistent across the entire experiment. It is possible that placing the observers in a situation where there is more uncertainty about what colors are relevant on each trial makes it more likely for near-target colors to be suppressed, but future research will be necessary to determine whether this difference is an important factor in the observed differences between our results and those of Störmer and Alvarez (2014). Along similar lines, the modifications we have added to Most et al.’s (2001) paradigm provide an avenue for future researchers to address fine-grained questions about the distribution of feature-based attention. In addition to examining whether history associated with a given color influences detection of unexpected items of that color, in future research we hope to more closely evaluate the idea that

observers tend to orient attention to regions of color rather than to specific colors. The answer likely depends on the task. For instance, we predict that observers are more likely to attend specific colors when tracking targets or more than one color.

Based on our experience performing this task, we were surprised to find that there was no evidence of suppression based on target color similarity. The “near” distractor seems to be much more difficult to ignore than the “far” distractor, which led us to believe that this perceived difficulty would translate to a decreased noticing rate. This highlights one of the attributes of the IAB task that makes it so interesting, and sometimes challenging to study: Intuitions about what sort of stimuli will be noticed are frequently wrong. Of course, the attention literature is rife with examples of humans overestimating their attentional abilities (e.g., Sanbonmatsu, Strayer, Medeiros-Ward, & Watson, 2013), but IAB research presents the unique challenge that observing a particular paradigm prior to beginning a study provides almost no insight into the results of the study. From this vantage point, Mechanical Turk is an invaluable tool for advancing our understanding of what kinds of unexpected stimuli are often missed. More generally, the IAB task provides an important window into the distribution of attentional priority while engaged in a challenging task. As evidenced by the current study, it is clear that subtle differences in stimuli (i.e., whether an US is on the same side of color space as distractors or not) can have strong effects on whether something is noticed.

Keywords: inattention blindness, attention, feature-based attention, inhibition

Acknowledgments

TD wishes to thank DOD grant # W911NF-15-2-0046 for partially supporting this work.

Commercial relationships: none.

Corresponding author: Trafton Drew.

Email: trafton.drew@psych.utah.edu.

Address: Department of Psychology, University of Utah, Salt Lake City, UT, USA.

References

- Chabris, C. F., Weinberger, A., Fontaine, M., & Simons, D. J. (2011). You do not talk about Fight Club if you do not notice Fight Club: Inattention blindness for a simulated real-world assault. *I-Perception*, 2(2), 150–153, doi.org/10.1068/i0436.
- Drew, T., Vo, M. L. H., & Wolfe, J. M. (2013). The invisible gorilla strikes again: Sustained inattention blindness in expert observers. *Psychological Science*, 24(9), 1848–1853, doi.org/10.1177/0956797613479386.
- Eriksen, C., & St James, J. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception and Psychophysics*, 40, 225–240.
- Eriksen, C., & Yeh, Y. (1985). Allocation of attention in the visual field. *Journal of Experimental Psychology: Human Perception and Performance*, 11, 583–597.
- Golomb, J. D. (2015). Divided spatial attention and feature-mixing errors. *Attention, Perception & Psychophysics*, 77(8), 2562–2569, doi.org/10.3758/s13414-015-0951-0.
- Heinze, H., Mangun, G., Burchert, W., Hinrichs, H., Scholz, M., Münte, T., et al. (1994). Combined spatial and temporal imaging of brain activity during visual selective attention in humans. *Nature*, 372, 543–546.
- Hopf, J. M., Boehler, C. N., Luck, S. J., Tsotsos, J. K., Heinze, H. J., & Schoenfeld, M. A. (2006). Direct neurophysiological evidence for spatial suppression surrounding the focus of attention in vision. *Proceedings of the National Academy of Sciences of the United States of America*, 103(4), 1053–1058.
- Hopf, J.-M., Boehler, C. N., Schoenfeld, M. A., Heinze, H.-J., & Tsotsos, J. K. (2010). Vision Research. *Vision Research*, 50(14), 1312–1320, doi.org/10.1016/j.visres.2010.01.015.
- Hyman, I. E., Boss, S. M., Wise, B. M., McKenzie, K. E., & Caggiano, J. M. (2010). Did you see the unicycling clown? Inattention blindness while walking and talking on a cell phone. *Applied Cognitive Psychology*, 24(5), 597–607.
- Moher, J., Lakshmanan, B. M., Egeth, H. E., & Ewen, J. B. (2014). Inhibition Drives Early Feature-Based Attention. *Psychological Science*, 25(2), 315–324, doi.org/10.1177/0956797613511257.
- Most, S. B., Simons, D. J., Scholl, B. J., Jimenez, R., Clifford, E., & Chabris, C. F. (2001). How not to be seen: The contribution of similarity and selective ignoring to sustained inattention blindness. *Psychological Science*, 12(1), 9–17, doi.org/10.2307/40063560?ref=search-gateway:4a2aa33-f69e30e53c1533c1b74546cb7.
- Navalpakkam, V., & Itti, L. (2007). Search goal tunes visual features optimally. *Neuron*, 53(4), 605–617, doi.org/10.1016/j.neuron.2007.01.018.
- Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of*

- Experimental Psychology: General*, 109(2), 160–174, doi.org/10.1037/0096-3445.109.2.160.
- Saenz, M., Buracas, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience*, 5(7), 631–632, doi.org/10.1038/nn876.
- Saenz, M., Buracas, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision Research*, 43(6), 629–637, doi.org/10.1016/S0042-6989(02)00595-3.
- Sanbonmatsu, D. M., Strayer, D. L., Medeiros-Ward, N., & Watson, J. M. (2013). Who multi-tasks and why? Multi-tasking ability, perceived multi-tasking ability, impulsivity, and sensation seeking. *PLoS ONE*, 8(1), e54402–8, doi.org/10.1371/journal.pone.0054402.
- Simons, D. J., & Chabris, C. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28(9), 1059–1074, doi.org/10.1068/p281059.
- Simons, D. J., & Jensen, M. S. (2009). The effects of individual differences and task difficulty on inattention blindness. *Psychonomic Bulletin and Review*, 16(2), 398–403, doi.org/10.3758/PBR.16.2.398.
- Störmer, V. S., & Alvarez, G. A. (2014). Feature-based attention elicits surround suppression in feature space. *Current Biology*, 24(17), 1985–1988, doi.org/10.1016/j.cub.2014.07.030.