IMPLICIT FEATURE-BASED SUPPRESSION IS EFFECTIVE AT GUIDING ATTENTION, EVEN WITH STRONG TARGET GUIDANCE, WHILE EXPLICIT FEATURE-BASED SUPPRESSION IS INEFFECTIVE

By

Andrew Rodriguez

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Psychology – Doctor of Philosophy

2025

ABSTRACT

Evidence for feature-based suppression (FBS) in visual search literature has been controversial and mixed. Much of the mixed findings are due to differences between implicit FBS, the unconscious ignoring of a feature that appears as a distractor with high regularity, and explicit FBS, the active biasing of attention away from a feature. Given this discrepancy, the goal of this dissertation is to examine the benefits of implicit and explicit FBS in a visual search task. Across four experiments, participants searched for a Landolt C with a horizontal gap among eight distractors with vertical gaps, with three items appearing in each of three colors. In Experiment 1, one color (ignored color) appeared on a higher percentage of trials (80%) and always appear as a distractor. Participants were not explicitly aware of the suppression contingency with the goal that they would implicitly learn to ignore this color. In Experiment 2, we included an additional implicit feature-based gain (FBG) contingency - a color (attend color) appeared on a large percentage of trials (40%) and always as the target - to see if suppression of the ignore color persists if there is an implicit feature-based gain (FBG) component present. Experiment 3 replicated Experiment 2 but made the FBG component explicit by cuing participants about the color that was likely to be the target. Experiment 4 replicated Experiment 1 but made the FBS component explicit by cuing what color should be ignored. Experiment 1 found a strong implicit suppression benefit that required minimal training of the to-be-ignored feature. However, in Experiment 2 where there were implicit FBS and FBG contingencies, only the FBG impacted search – suggesting that implicit FBG interfered with the development or implantation of an implicit FBS mechanism. However, when the FBG component was made explicit (Exp. 3), the implicit FBS effect reappeared – suggesting the failure to find FBS in Experiment 2 was likely due to a competition for implicit resources between the FBG and FBS components. Finally,

when FBS appeared alone but was made explicit (Experiment 4) there was week evidence for suppression. These findings suggest that implicit FBS can effectively allocate attention, but the circumstance under which they do so are limited – only when the contingencies are learned implicitly and in the absence of implicit feature-based gain contingencies. These restrictions may limit FBS's utility as a real-world mechanism for guiding attention.

ACKNOWLEDGMENTS

I want to start off by thanking my advisor and mentor, Mark W. Becker, who was instrumental in my development and growth during my graduate school training. Mark was always considerate with his time, being incredibly flexible with scheduling meetings and providing quick turnarounds on any material sent in his direction. While many would sacrifice quality for speed, his constructive criticism and feedback was often detailed, elaborating on elements that might be missing or lacking while providing directions to help fill in the missing gaps. He also provided a sense of freedom in my training, leading me to explore *PsychoPy* and learning how to program in Python and Java (too much lesser and frustrating degree). I am appreciative of all the work that he has done throughout my graduate training.

I am also grateful to my dissertation committee members, Jan Brascamp, Susan Ravizza, and Taosheng Liu, who have assisted me throughout my time here at MSU, whether that was in classroom as professors teaching critical skills and knowledge, in the laboratory as co-advisors contributing ideas and code to improve my experiments, or in the hallways as friends making small talk and creating a sense of warmth within our program.

I also want to extend my sincere gratitude to my undergraduate research assistants that collected data for all experiments in my dissertation: Alison Hanes, Ian Render Flores, Lauren Broersma, Maxwell Mccort, Paige Abraham, and Tristan Janisse. Without your hard work and dedication to running participants at an expeditious pace, I would not have completed my dissertation.

I also want to give a special thank you to the lab managers of the Attention & Perception Lab who were involved in my dissertation projects: Derrek Montalvo and Morgan Dodd. I sincerely appreciate the little things they did to make my life easier, from training the RAs to run my experiments to backing up the data on lab drive. I also want to recognize the previous lab managers of the Attention & Perception Lab: Dana Pontious, Katie Anderson, and Matt Kebede. While they were not actively involved in my dissertation work, their assistance through my years as a graduate student at Michigan State University facilitated both my productivity and happiness.

I am deeply grateful for the financial support that I received during my graduate training. Much of my funding came from the Agency for Healthcare Research and Quality grant assisting on collaborative project between Laura Bix at the Packaging Department and Mark Becker, with the College of Social Science funding my dissertation work through the Dissertation Completion Fellowship award.

Special mentions go to my friends and family who have supported me throughout my years. To my mother, Leticia, and my three younger siblings, Alex, Anthony, and Genevie, who have no idea what I am researching no matter how many times I have explained it to them but still support my dreams and career aspirations from afar. To my wife, Beth, who has supported my journey at MSU from my first year here in East Lansing through the COVID years all the way up to my sixth year and will (hopefully) continue to support me for many years to come. I couldn’t do this without your love and encouragement. To Izzy and Alexa, who are my only family in Michigan, were always willing to help me out regardless of the circumstances, and literally never gave “no" as an answer. To my former and only graduate school lab mates, Matthew and Stephanie, who were instrumental to my early success and provided my first home away from home during all my ups and downs at the University of Arkansas. To Clara and Chadwick, who selflessly offered me a place to stay out while I attempted to finish my dissertation. I love and appreciate you all.

Lastly, to all my other friends, family, and colleagues who have assisted me throughout the years. Thank you all for being a part of my journey and for being there when I needed you.

TABLE OF CONTENTS

LIST OF ABBREVIATIONS

BF	Bayes Factor
CDA	contralateral delay activity
EEG	electroencephalogram
ERP	event-related potential
Exp	experiment
FBA	feature-based attention
FBG	feature-based gain
FBS	feature-based suppression
ms	milliseconds
msec	milliseconds
RDP	random dot patterns
RDK	random dot kinematogram
RTs	reaction times
SD	standard deviation
SE	standard error
VSS	Vision Sciences Society

CHAPTER 1:

A REVIEW OF THE EFFECTIVENSS OF FEATURE-BASED SUPPRESSION IN ALLOCATING VISUAL ATTENTION

1.1 Abstract

There is strong evidence suggesting that feature-based attention facilitates the selection of relevant items through feature-based gain. More recently, several prominent papers have claimed that the suppression of irrelevant features also facilitates the selection of relevant items, and in some extreme cases, that feature-based attention is entirely driven by suppression. However, the evidence for suppression is less comprehensive and more controversial, raising questions about how central a role feature-based suppression (FBS) plays in the allocation of attention. To assess its role, here, I review previous literature on FBS and describe some of my previous attempts to find suppression effects with the goal of evaluating the quality of FBS evidence and identifying critical factors that may impact the effectiveness of FBS. To foreshadow, my review suggests that FBS may not generalize across tasks, and questions remain about whether making the knowledge of the to-be-ignored feature explicit interferes or reduces FBS relative to pure implicit learning and whether FBS occurs when there is the ability to engage feature-based gain mechanisms. I then propose a series of studies that attempt to address these questions to help clarify the importance of FBS in directing attention.

1.2 Introduction

Feature-based attention is one of many factors involved in the efficient allocation of attention toward relevant stimuli. Feature-based attention (FBA) is the ability to selectively attend to a specific feature (e.g., color, shape, size, etc.) while ignoring features that are irrelevant and distracting to the current goal (Maunsell & Treue, 2006). There is strong evidence that feature-based attention can facilitate the selection of relevant items by feature-based gain (Egeth, Virzi, & Garbart, 1984). Feature-based gain is the upweighting of the neural response to a target's feature, enhancing the perceptual processing of that feature and enabling more efficient attentional selection of items with that feature in behavioral tasks (White & Carrasco, 2011; Yeshurun & Carrasco, 1998).

For several decades, research on feature-based attention focused on this gain mechanism. This research has found that prior knowledge of a target's feature (e.g., motion) results in enhancement of the neural response of units coding that feature (Treue & Martinez-Trujillo, 2007). For example, Martinez-Trujillo and Treue (2004) performed single cell recordings in macaques who viewed two sets of random dot patterns (RDPs) either moving in the same or opposite directions. One RDP was located inside the receptive field of a recorded neuron but was not attended to by the monkey (a distractor RDP), while the other RDP was located outside the neuron's receptive field and was the attended RDP. They found significant neuronal enhancement as a function of how similar the distractor RDP was to the attended RDP, leading to their feature-similarity gain model of attention (Trueu & Trujillo, 1999). According to this model, the magnitude of enhancement is centered on the attended feature and decreases monotonically as the distance in feature space moves away from the attended feature. One notable aspect of this monotonic decrease is that it can result in suppression; the neural response to features that are distal in feature space from the attended feature are suppressed below baseline. However, this is a by-product of attending to a specific feature rather than an independent suppressive mechanism centered on a to-be-ignored feature

More recent research (Hopf et al. 2006; Stormer & Alvarez, 2014) suggests that feature-based gain has a Mexican Hat attentional profile with regions of suppression that are at intermediate distances in feature space when compared to the target feature. Others (Fang, Becker, & Liu, 2019) argue that both attentional profiles occur for feature-based gain: a broad feature-similarity gain profile with a local Mexican Hat profile centered on the target feature. Nevertheless, this upweighting gives items with that feature an advantage in the competition for attention, allowing attention to be effectively guided toward items with this feature (Williams, 1966).

Consistent with the idea that gain may be important to guiding attention toward relevant objects in the scene, there is evidence that feature-based gain occurs globally across the visual field (Liu & Hou, 2011; Saenz, Buracâs, & Boynton, 2002, 2003). An example of this global enhancement was observed in Zhang & Luck (2009) using electroencephalogram (EEG). They had observers view a random dot kinematogram (RDK) composed of two sets of colored dots that were superimposed on each other and were presented in one hemifield. Observers were to attend to the dots of one color while ignoring the dots in the other color and were to respond when the attended dots changed brightness. Occasionally, an irrelevant field of dots appeared briefly in the contralateral visual field and the P1 event-related potential (ERP) response to the onset of this irrelevant field of dots was monitored. The P1 response was larger when the irrelevant dots matched the attended color than when they matched the ignored color. The authors interpreted this as evidence for an early perceptual global gain for the attended color. Others have found similar feature-based gain advantages using a variety of methods including psychophysics (Yantis, 1993; Treisman & Gelade, 1980), eye tracking (Beck, Hollingworth, & Luck, 2012; Navalpakkam & Itti, 2006), and fMRI (Liu & Hou, 2011; Liu, Larsson, & Carrasco, 2007; Serences & Boynton, 2007). In short, this feature-based gain mechanism has been well-studied and characterized.

However, more recently several prominent papers have claimed that the suppression of irrelevant features also facilitates the selection of relevant items – this is known as feature-based suppression. Feature-based suppression (or feature-based inhibition) involves the downweighing of irrelevant stimuli features. Critically, this feature-based suppression differs from the suppression observed as a result of the feature-based gain attentional profile. It is not a byproduct of a gain mechanism but is instead a process that is directly applied to the features of a known distractor. Many of these papers argue that feature-based suppression of distractor features also leads to more efficient target detection on behavioral tasks (Chang & Egeth, 2019, 2021; Gaspelin, Leonard, & Luck, 2015).

In one extreme interpretation, Moher, Lakshmanan, Egeth, & Ewen (2014) argued that feature-based attention is driven almost entirely by suppressive mechanisms, suggesting that suppressive mechanisms play a major role in guiding attention toward relevant stimuli. Their experiment was an extension of Zhang and Luck's (2009) but proposed that the observed difference in P1 magnitude could be due to either upweighting of the attended color or suppression of the ignored color. To differentiate between these possibilities, Moher and colleagues replicated the experiment, adding a third set of irrelevant dots that matched neither the attended nor ignored color in the primary task (neutral color). While they replicated the P1 magnitude difference between the attended and ignored colors, they found that the response to the neutral color almost perfectly matched the response to the attended color. This result suggests that the difference between the response to the attended and ignored colors was due to suppression of the ignored color, rather than upweighting of the attended color. Their results also suggests that feature-based suppression operates globally across the visual field, downweighting items that appear in the to-be-ignored color.

In a behavioral follow-up, Moher and colleagues had observers perform the same RDK task for an extended period of time (~20 minutes) immediately followed by a visual search task in which observers looked for the presence of a Z or N among a large set of colored letters. Although color was irrelevant to the search task, some letters appeared in the previously attended or ignored color, while others appeared in colors that had not appeared in the RDK task. RTs were longer when the target appeared in the ignored color than when it appeared in the target color or a color that had not appeared in the RDK task, with no RT difference in the latter two. Again, they took this as evidence that selection was being driven almost exclusively by suppression. They also found that the carry-over effect from the RDK to the visual search task was short-lived; the RT differences between conditions disappeared by the second block of search trials. Even so, these findings suggest that suppression may be an important mechanism in the allocation of attention. Not only did they find evidence that suppression, rather than gain, was responsible for attentional selection, but they also found that the attentional weights set up during the RDK task generalized to a very different visual search task. Thus, these results present an extreme view that FBS is a stronger mechanism than feature-based gain, that it is global in nature, and that it is applied flexibly across contexts.

However, there are reasons to doubt this extreme view. First, it is worth noting that these results have not been replicated in the past decade. In particular, there is some debate about whether feature-based suppression occurs on a global level (Oxner et al., 2023). Forschack, Andersen, & Müller (2017) had subjects view two centrally superimposed RDKs (one attended and the other ignored) while task-irrelevant single-colored RDKs were presented to the right and left periphery. Each set of colored dots flickering at different rates allowing the use of steady state visually evoked potentials (SSVEPs) to isolate the response to each RDK. They found evidence that feature-based gain was set globally but suppression that only occurred in the center, leaving it unclear whether feature-based suppression occurs globally as suggested by Moher et al. (2014).

Second, there are many results (Beck, Luck, & Hollingworth, 2018; Olivers, Meijer, & Theeuwes, 2006; van Moorselaar, Theeuwes, & Olivers, 2014) which suggest that positive cues (for example, cues that indicate the color of the target) are more effective and can be used more flexibly than negative cues (cues that indicate the color that the target will not appear in). Finally, there are reports that find benefits of a positive cue in situations where there are heterogenous distractors (Beck & Hollingworth, 2015; Becker, Hemsteger, & Peltier, 2015), situations where suppression would be an ineffective mechanism. In short, while the view of feature-based suppression as a dominant mechanism is provocative, there is little corroborating evidence to support such an extreme position.

Another extreme claim comes from Arita, Carlisle, & Woodman (2012). Observers searched for a Landolt C with a break at the top or bottom, among an array with breaks on the right or left. In separate blocks of trials, search arrays were preceded with positive, negative, or neutral cues (a color that did not appear in the search array). The color of the cue varied on a trial-by-trial basis. Relative to the neutral block, they found much faster RTs in the positive cue block, demonstrating a strong role for feature-based gain. They also found faster RTs in the negative cue block than neutral cue block, providing evidence for a “template for rejection” or suppression of the cued color. However, it is worth noting the RTs were slower for negative than positive cues, suggesting that gain is a more potent mechanism than suppression. While the finding that suppression is not as strong as gain makes this claim less extreme than the Moher claim above, the fact that the suppressive mechanism can be set volitionally on a trial-by-trial basis makes this claim extreme.

However, multiple authors have criticized the Carlisle method. The main criticism focuses on the fact that the search arrays always contained only two colors, and the colors were segregated by hemifield. Beck & Hollingworth (2015) and Becker, Hemsteger, & Peltier (2015) failed to replicate the negative cue effect when they intermixed the display so that the two colors were no longer segregated by hemifield. In the Becker paper, they also included a condition where half the items appeared in the negative cue's color, while the other items appeared in heterogenous colors. If the cued color was suppressed, the effective set size should have been cut in half, thereby reducing RTs. However, no negative cue benefit was found in those conditions. Both sets of authors suggested that the benefit observed in the Carlisle method resulted from using the negative cue to quickly remap a positive cue for either the target hemifield or the target color. When these remapping strategies were made ineffective, the negative cue benefit disappeared. Some researchers argue that the failure to observe templates for rejection in the above studies may be due to the tasks being insufficiently difficult (Kerzel & Huynh Cong, 2022). For instance, Conci, Deichsel, Müller, & Töllner (2019) hypothesized that the use of negative templates may be effortful, and thus, is only utilized when the search task is difficult enough to encourage participants to expend this effort. They used an intermixed visual search task with positive and negative cues that change on a trial-to-trial basis with two different search difficulties (easy & difficult search tasks). They found a large benefit from the positive cue in the easy task, but no benefit for the negative cue. However, in the difficult search task, they found a benefit for the negative cue, though it was much smaller than that for the positive cue. They concluded that observers only utilize negative templates when the search task is sufficiently difficult.

While this finding might suggest that people will only utilize negative cues when there is a notable benefit to doing so, it seems difficult to dismiss the results of Beck & Hollingworth (2015) and Becker, Hemsteger, & Peltier (2015) by arguing that those tasks were not difficult enough to encourage the use of suppression. These studies used intermixed and/or heterogenous search arrays, which should make the tasks more difficult, rather than less difficult, than the Carlisle method. Furthermore, task difficulties can be subjective and it is not necessarily clear what constitutes an “easy” task and what constitutes a “difficult” task, making this claim a weak defense for templates for rejection.

While both Moher et al. (2014) and Arita, Carlisle, & Woodman (2012) make extreme claims about the role of feature suppression in the allocation of attention, there is other evidence for suppression that is less controversial. Gaspelin, Leonard, and Luck (2017) used the additional singleton search paradigm and found faster reaction times (RTs) on trials with a salient singleton than without. They also found less oculomotor capture for the singleton than other non-singleton distractors. Both findings suggest that the irrelevant feature is suppressed, leading to less attentional capture by that feature. While the results of this paper and others (Cunningham & Egeth, 2016; Sawaki & Luck, 2010; Won & Geng, 2020) are compelling, there is debate about the conditions under which feature-based suppression will occur. The resolution of this debate has implications for how prominent a role feature-based suppression should assume in theories about the allocation of visual attention. If suppression is used flexibly across a variety of contexts and tasks, it would suggest that it is a prominent mechanism in the guidance of attention. To the extent, however, that suppression is only effective in very specific circumstances that are likely to be rare in everyday life, it would suggest that suppression is likely not a main mechanism influencing the allocation of attention.

For example, some researchers suggest that suppression, unlike gain, cannot be applied on a trial-by-trial basis but requires repeated exposure to the same to-be-ignored features. Cunningham and Egeth (2016) had observers perform a visual search task in which they searched for a “B” or “F” among several distractor letters. In some trials, participants were given a neutral cue or a negative cue, with the negative cue staying the same color throughout the experiment. Initially, RTs were longer in the negative cue trials compared to the neutral cue trials, suggesting that attention was captured by stimuli that appeared in the negative cue color, rather than those items being suppressed. However, with enough practice and learning, participants eventually learned to suppress the negatively cued color, resulting in faster RTs in negative cue than neutral trials. In addition, if the color of the negative cue changed on a trial-by-trial basis, suppression never develops and there is not a benefit from the negative cue. Thus, they found that successful explicit FBS could only occur if the ignored feature remains the same through repeat exposure and learning.

Similarly, Vatterott & Vecera (2012) had observers perform an additional singleton search task, with the singleton distractor changing colors from block to block (e.g., red in block 1, yellow in block 2, purple in block 3, etc.). This manipulation allowed them to investigate whether observers learn to ignore the additional singleton distractor (regardless of color), or if they learn to ignore the specific color that is relevant to that block. They found the same pattern of results across every block: attention was captured by the singleton at the beginning of each block, as indicated by slower RTs with the singleton than without, but in the second half of the block, RTs were faster with the singleton than without. This pattern of reversal suggests that successful suppression of the singleton is color specific and must be learned through repeated exposure. This is in stark contrast with Arita, Carlisle, & Woodman (2012) where the to-be-ignored feature changes on a trial-to-trial basis, which eliminates the ability to learn to ignore the same feature.

Even when the to-be-ignored feature is held constant, there appear to be conditions where suppression is more robust. Notably, suppression seems more effective when the distractor feature is learned implicitly through repetitions and statistical regularities than when observers are explicit cued to ignore the same feature repeatedly (Stilwell & Vecera, 2019a; 2019b). Addleman & Stormer (2023) compared the magnitude of suppression effects when observers were explicitly cued about the irrelevant feature to cases where they implicitly learned about the irrelevant feature via repetition. Observers performed a Landolt C search task with half of the Cs in one color and the other half in a second color, intermixed in the search array rather than separated by color. In Experiment 1, the distractors appeared in a given color on 75% of trials, but the participants were not made aware of this contingency. RTs were faster when the distractor appeared in the more frequent color compared to less frequent colors, suggesting that participants learned these contingencies and used them to suppress the color that was most often associated with distractors. In Experiment 2, observers were explicitly cued about the distractor color, which was validly cued on 75% of trials. While RTs were faster when the distractor appeared in the valid cue condition compared to the invalid cue condition, the benefit for the explicit negative cue was not as large as the benefit of the implicitly learned feature in Experiment 1. Thus, it seems that feature-based suppression is more effective when implicitly learned than explicitly cued.

Additionally, much of the feature-based suppression literature examines the ability to downweight one specific feature dimension, mainly color. However, if FBS is a main mechanism in guiding attention, this ability to downweight should be observed across various feature dimensions. For example, there is evidence that feature-based gain not only enhances color features (i.e., Zhang & Luck, 2009) but also enhances specific orientations (Liu & Hou, 2011), motions (Martinez-Trujillo & Treue, 2004; Saenz, Buracas, & Boynton, 2002), and even in complex objects (Chelazzi, Duncan, Miller, & Desimone, 1998), indicating the robustness and applicability of the gain mechanism in everyday life. For FBS to be a critical mechanism in allocating attention, it should also be versatile enough to downweight more than just color features.

Adams, Ruthruff, & Gaspelin (2023) investigated whether other features besides color could be suppressed. They used the additional singleton paradigm with three conditions: distractor absent (no additional singleton was present), singleton distractor (a color singleton appeared), and onset distractor (four white circles flashed around a distractor, creating a non-colored singleton). While participants could actively suppress the colored singleton, leading to faster RTs and little oculomotor capture for that condition compared to the distractor absent condition, they struggled to suppress the onset singleton. In fact, RTs were the slowest and oculomotor capture was highest in the onset-singleton condition. Even when the authors controlled for luminance (changing the white dots to grey) and for color (changing the dots to red or green), the suppression of the onset singleton did not improve. The authors concluded that abrupt onsets could not be suppressed like color. This inability to suppress other feature types besides color indicates further limitations of FBS, weakening the argument that it is an important attentional mechanism.

In sum, there are reasons to doubt both the extreme views presented above. The evidence seems to suggest that feature-based gain is a more powerful and flexible mechanism than FBS. While feature-based gain upweights the to-be-attended feature on global level, FBS appears to only the downweight locally when there is spatial competition from a rival feature, limiting its benefit in larger scale visual searches. Further, it appears that FBS can occur under conditions where a given feature is consistently irrelevant, but does not seem to be effective if the irrelevant feature changes on a trial-by-trial basis. This suggests that FBS attention requires learning and evidence suggests that this learning is best done implicitly.

Given these limitations, while FBS may play a role in the allocation of attention, it seems like a less powerful mechanism than feature-based gain. Positive cues that allow for feature-based gain have been shown to produce large and robust effects (Beck, Hollingworth, & Luck, 2012; Vickery, King, & Jiang, 2005), remain effective even when features change on a trial-by-trial basis (Töllner, Conci, & Müller, 2015), and apply across various feature dimensions (Chelazzi, Duncan, Miller, & Desimone, 1998; Liu & Hou, 2011; Martinez-Trujillo & Treue, 2004; Zhang & Luck, 2009). These properties make gain mechanisms fairly flexible and suggest that they can efficiently guide attention to relevant stimuli in the world. By contrast, FBS is less flexible, only occurring when a feature (often color) is consistently irrelevant and is learned through statistical regularity. Thus, FBS is likely a less central mechanism in the allocation of visual attention.

This apparent difference in the effectiveness of suppressive versus gain mechanisms raises the question of whether feature-based gain interferes with feature-based suppression. It could be that gain masks the benefits of suppression or encourages observers to prioritize gain over suppression. While research on this subject is limited and relatively new, it is also somewhat mixed. Stilwell and Vecera (2020) provide evidence that feature-based gain and FBS mechanisms can co-occur and are additive. In their experiment, observers searched a circular display of eight isoluminate colored squares containing gaps (like a Landolt C). On a given trial, two colors (out of a total of six colors) appeared, with half of the items in one color and the other half in a second color. One of the six colors never contained the target (the consistent distractor color) and observers were not made aware of this contingency, providing a strong implicit feature-based suppression component. Prior to the search array appearing, there was also an explicit positive cue, providing a strong feature-based gain component. RTs were faster when the positive cue matched the target color than when the cue was neutral (when the cue matched neither color in the display), providing evidence for a strong feature-based gain mechanism. RTs were also faster when the four distractor squares appeared in the consistent distractor color compared to one of other colors, providing evidence for a FBS mechanism. Furthermore, the fastest RTs occurred for the target cue/consistent distractor color pairing, suggesting that FBS and feature-based gain were independent processes that could co-occur and were additive with one another.

However, other authors (Kawashima & Amano, 2022; Rajsic, Carlisle, & Woodman, 2020) have found that the ability to use feature-based gain interferes with the use of suppression. For example, Rajsic, Carlisle, & Woodman (2020) replicated the Landolt C search task from Arita, Carlisle, & Woodman (2012), but provided participants with four different levels of cuing information: a neutral cue, a positive cue, a negative cue, and a positive+negative cue that provided the color of both the target and distractors. There was a large benefit in RTs for the positive cue, suggesting a strong feature-based gain mechanism. However, the positive+negative cue was no faster than the positive cue alone, indicating that there was not an independent and additive effect of FBS. Using ERPs, they also measured the contralateral delay activity (CDA) elicited by each cue condition. The CDA is a lateralized component that tracks the number of items held in visual working memory. For the positive+negative cue condition, the negative and positive cues were presented to opposite hemifields, allowing the lateralization on the CDA to provide evidence about which of the two cues was used. Their results show that when participants only used the positive cue - the CDA magnitude indicated that only one cue was used and the lateralization tracked the hemifield of the positive cue. The authors concluded that when providing observers with the option of using both feature-based gain and suppression, observers choose to use gain in lieu of suppression, contradicting what Stilwell & Vecera (2020) found. However, it's worth noting that Stilwell & Vecera (2020) had their participants implicitly learn to ignore a specific color while Rajsic, Carlisle, & Woodman (2020) explicitly cued the to-be-ignored feature that changed on a trial-to-trial basis. This methodological difference could have led to differences in results. Nevertheless, it is still unclear if feature-based suppression could be effective in the presence of gain.

Finally, there are questions about how well FBS effects generalize across tasks. As pointed out above, Moher, Lakshmanan, Egeth, & Ewen (2014) found that attentional weights set up during a RDK task generalized to a subsequent visual search task. Such generalization would suggest that learning that a feature was likely irrelevant leads to a general turning down of the response to that feature, a process that could have important impacts on the allocation of attention that may generalize across tasks.

However, suppression effects do not always generalize across different tasks (De Waard, Van Moorselaar, Bogaerts, & Theeuwes, 2023; van Moorselaar, & Theeuwes, 2024). Britton & Anderson (2020) examined if suppression implicitly learned in one context generalized to a different context. Observers performed an additional singleton search task against various realistic image backgrounds (for example, pictures of forest or urban landscapes). Each image was associated with a highly probable distractor location (for example, the urban landscape might have the top right corner be the highly probable distractor location while the forest might have the bottom left corner be the highly probable distractor location). Observers learned the high-probability distractor locations for each background, but did not restrict them to the context that they were learned. RTs were faster when the additional singleton distractor appeared in either high-probability location regardless of if it matched the appropriate background, indicating that the learned suppression generalized across the different contexts. However, when the researchers decided to see if the high-probability distractor locations learned in the additional singleton task would generalize to a completely unrelated and different task (a spatial cuing task) using the same urban and forest backgrounds from the additional singleton task, they did not. While they found a benefit of cuing, as RTs were faster when the target appeared in the cued location, search times did not differ across the low- and high-probability distractor locations were similar. More specifically, RTs were not slower when the target appeared in the high-probability distractor location, even when the backgrounds matched. Thus, they found that suppression can generalize across contexts and backgrounds, but not across different task paradigms using unrelated stimuli, putting it at odds with what Moher, Lakshmanan, Egeth, & Ewen (2014) found.

In sum, while there has been a recent explosion in the number of papers investigating feature-based suppression mechanism, there is still considerable uncertainty about how central a role FBS mechanisms play in the allocation of attention. Determining how central a role FBS plays requires a better understanding of the cases when it will and will not occur. Specifically, there are open questions of whether it will occur when there is the ability to engage feature-based gain mechanisms, whether making the knowledge of the to-be-ignored feature explicit interferes or reduces FBS relative to pure implicit learning, and the extent to which the effects can generalize across tasks suggesting a general down-weighting of the to-be-ignored feature. If FBS effects generalize across tasks, maintain even when feature-based gain is possible, and maintain even in the face of explicit knowledge, then one can be confident that it plays a central role in the allocation of attention. If, however, these factors eliminate or diminish the effectiveness of FBS it would diminish how central a role FBS should play in models of the attentional allocation process. My previous work sought to investigate the first two of these issues within a generalization framework.

1.3 Prior Work

I have made several attempts to find evidence of feature-based suppression, with none of those experiments providing any compelling or robust evidence for FBS. All three experiments consisted of two tasks: an RDK task and a letter search task. The RDK task was similar to that used in Zhang & Luck (2009) and Moher et al. (2014) where two clouds of colored dots appeared in an invisible circular aperture and moved randomly (0 motion coherence) while occasionally changing luminance (see figure 1). This task was used to create attentional weights for the colors.

The letter search task that was used to assess suppression and gain was identical to the one in Cunningham & Egeth (2016). Participants searched for a “B” or “F” target in an array consisting of clock face containing 12 letters (see figure 2). The letter search task was used to evaluate how well the attentional weights from the RDK task generalized to this search task. To do so, participants performed a hybrid test block consisting of both the RDK task and letter

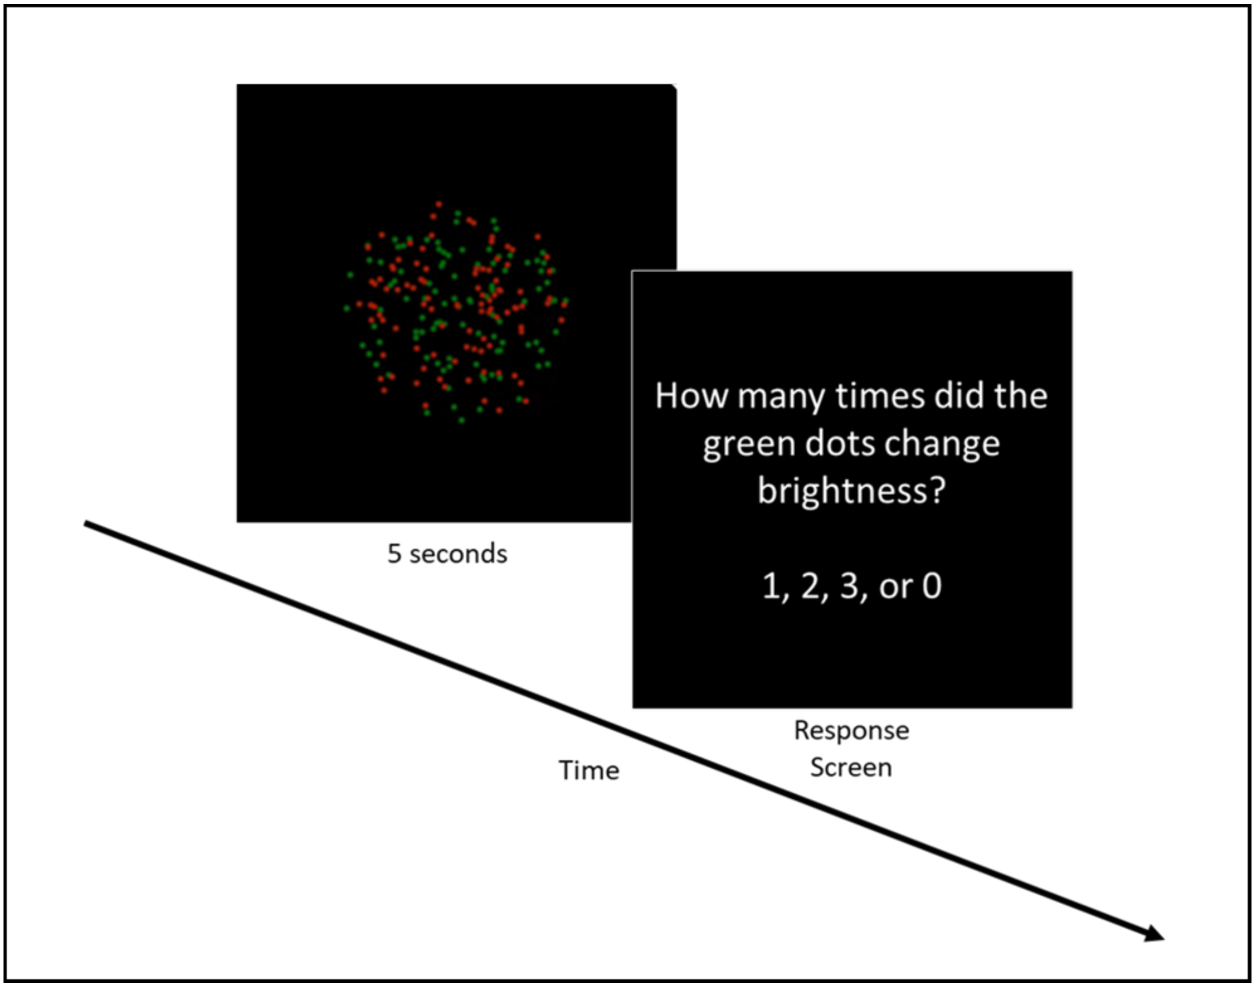

Figure 1. An example of an RDK task trial from Experiment 1 of the previous studies. An invisible circular aperture consisting of 100 red and 100 green dots was presented centrally against a black background. The dots will appear for 5 seconds, and during that period, they moved randomly (0 motion coherence) within the aperture. Both sets of dots would randomly change in luminance (either increasing or decreasing in brightness) but never changed brightness at the same time. Once the dot filtering was completed, a response screen appeared asking to indicate the number of times the target dots changed in brightness. This response screen stayed on until a button press was made.

search task (see figure 3). Each trial began with the RDK task. 75% of these trials were normal RDK trials as described below. The remaining 25% of trials began with the RDK display but after 2500 msec it was interrupted and changed to a letter search trial. The reasoning for interrupting the RDK task was to maximize the attentional weights for the attend and ignore colors before changing to the letter search task.

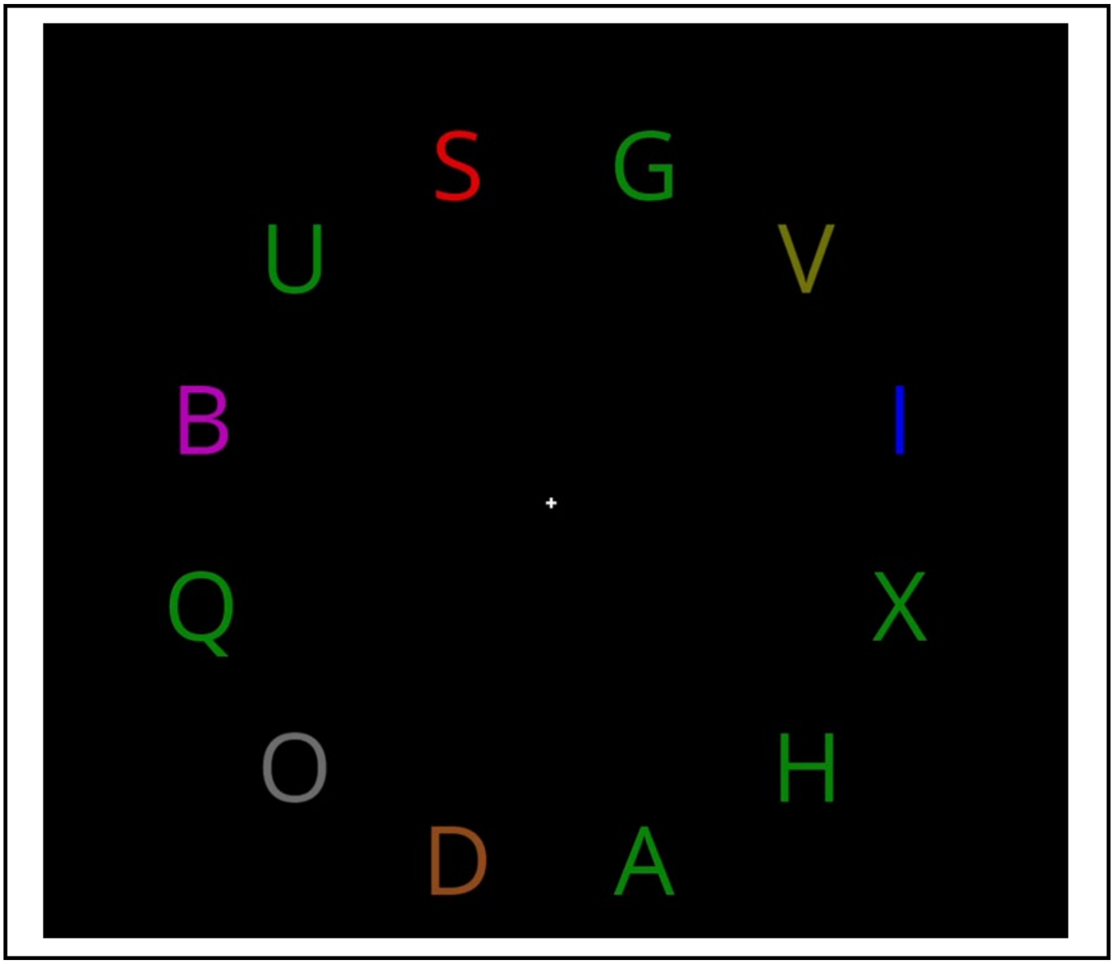

Figure 2. An example of the letter search display used in all three previous experiments. Participants searched a circular array containing 12 letters to find either a “B” or “F” target and indicate which letter appeared. Six letters will appear in the same color (green in the above example) and never contain the target. The other six letters will each appear in a unique color (either red, green, blue, yellow, purple, brown, or gray) and one of those letters will be the target (purple in the above example).

In experiment 1, for the RDK task, participants were assigned a target color (either red or green) – referred to as the attended color. They were instructed to monitor that color and ignore the other color. The task required them to count the number of luminance changes (0, 1, 2, or 3 changes) in the attended dots while ignoring changes in the ignored color. Thus, the task involved both explicit feature-based gain and suppression components. Luminance changes could occur in the to-be-ignored color, so accurate counts required selective attention to the

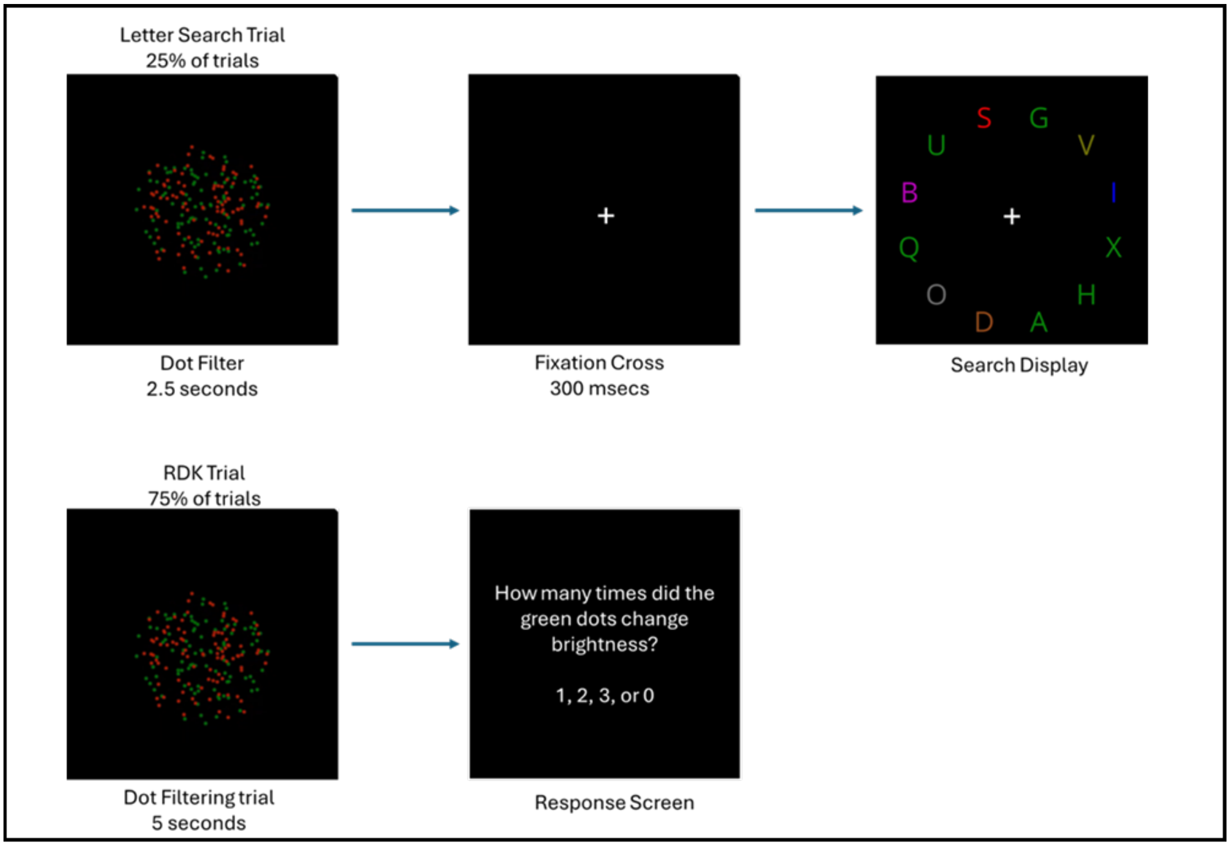

Figure 3. An example of a possible trial in the hybrid test block. The upper panel shows an example of an interrupting letter search trial. These trials occur 25% of the time and always begin with the RDK task. After 2.5 seconds, the task is interrupted with a fixation across appearing for 300 msecs before the letter search array appears. The bottom panel shows an example of a RDK trial.

attended color. For the letter search task, six of the 12 letters appeared in the same color and the target never appeared in this set of six. Across trials, this set appeared in one of three colors (the attend color, the ignore color, or a neutral color that did not appear in RDK task). Each of the other six letters appear in a unique color (either red, green, blue, yellow, purple, brown, or gray). The target letter appeared in three different colors (the attend color, the ignore, or one of the neutral colors).

An analysis of the letter search RTs as a function of target color (attend, ignore, and neutral color) found a significant difference, $F(2, 126) = 3.56$, $p = .031$, with faster RTs when the

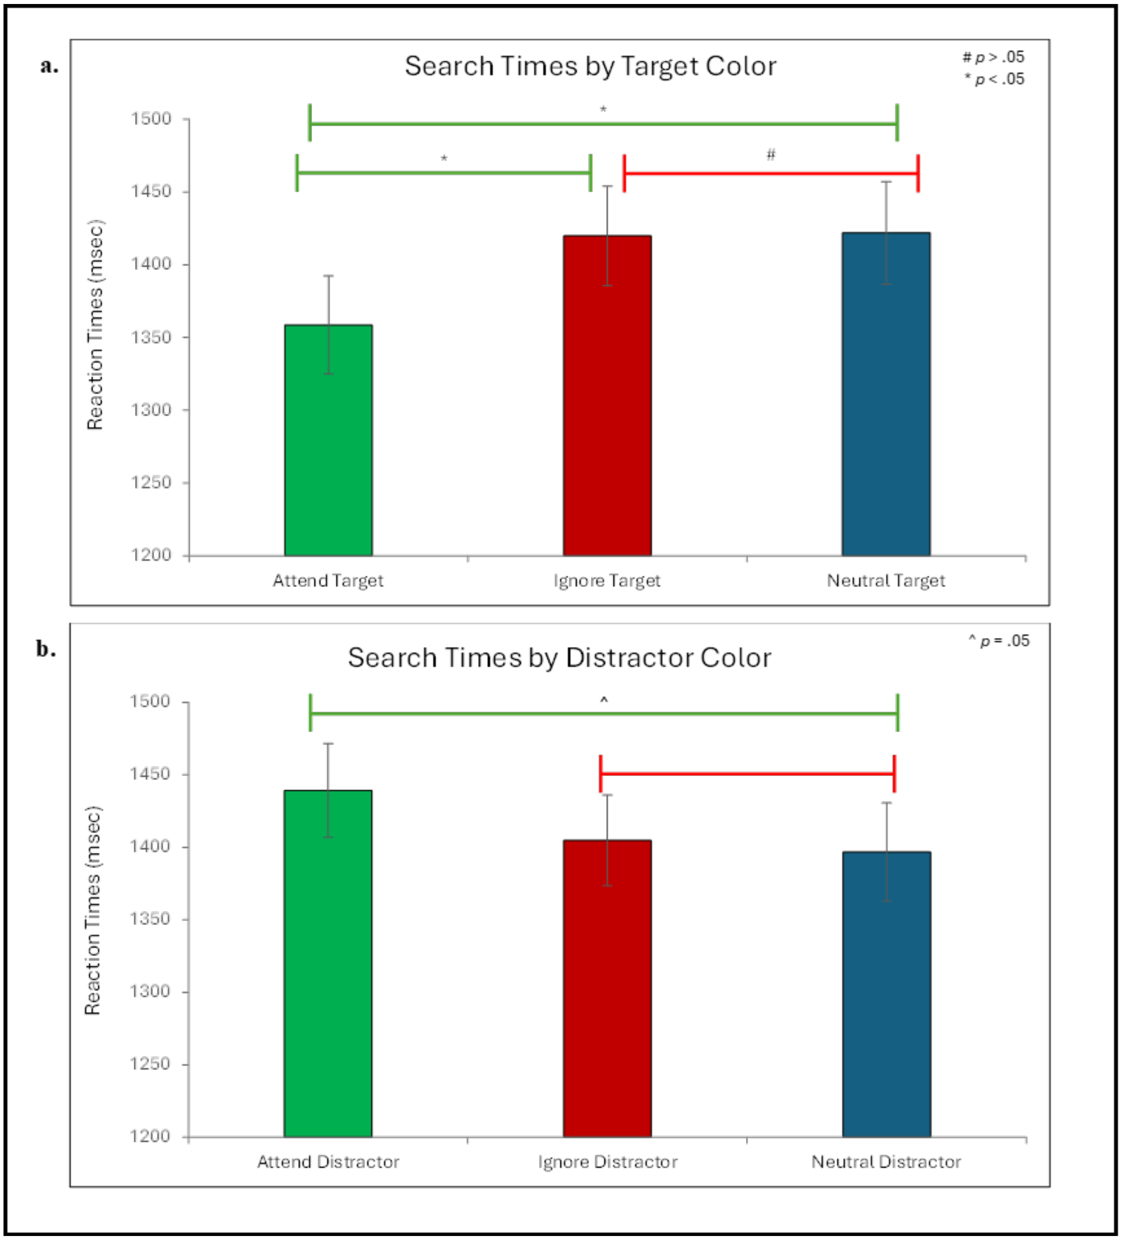

Figure 4. Search times for Experiment 1 by target color (figure 4a, top graph) and distractor color (figure 4b, bottom graph).

target was in the attend (M = 1358.62 ms, SE = 33.63 ms) than either the neutral color (M = 1419.82 ms, SE = 34.05 ms), $t(63) = 2.31$, $p = .024$, or the ignore colors (M = 1421.89 ms, SE = 35.18 ms), $t(63) = 2.08$, $p = .042$ (figure 4a). The ignore color did not differ significantly from the neutral color, $t(63) = .09$, $p = .931$. The faster RTs for the attend targets provide evidence for feature-based gain, but the fact that RTs were no slower when the target was in the ignored color than the neutral color provides no evidence of FBS.

An analysis of RTs as a function of the frequent distractor color (attend, ignore, and neutral color) found no significant differences in RTs, $F(2, 126) = 2.04$, $p = .135$ (figure 4b). Had there been feature-based suppression, we would have expected faster RTs when six of the distractors appeared in the ignore color. If there were feature-based gain, we might have expected slower RTs when the six distractors appeared in the attend color, but we did not find that effect here. In sum, the results from this study found some evidence for feature-based gain (from the target color analysis), but no evidence of FBS.

Given that some studies suggest that FBS is not implemented when it is possible to implement a feature-based gain mechanism (Kawashima & Amano, 2022; Rajsic, Carlisle, & Woodman, 2020), we conducted a follow-up experiment where we attempted to eliminate the gain component. To do this, we modified the RDK task, instructing participants to always ignore the dots of an assigned color (either red or green) and to report the luminance changes of the dots that appeared in the other color (the target color). The target color changed on a trial-by-trial basis and could be any of five other colors (green or red depending on the assigned color, blue, gray, purple, or cyan). Given this approach, there was an explicit instruction for the to-be-ignored feature, but there was no explicit cue for gain and the gain feature was not consistent across trials. We hoped this would minimize the use of feature-based gain. Furthermore, in the letter search task, the color of the attend dots from the previous RDK task never appeared. For example, if the attend color from the RDK task was red for a given trial, none of the letters in the subsequent search trial would appear in red. This was done to further limit the impact of the feature-based gain. For our analysis, we had three comparisons: when the target letter appeared

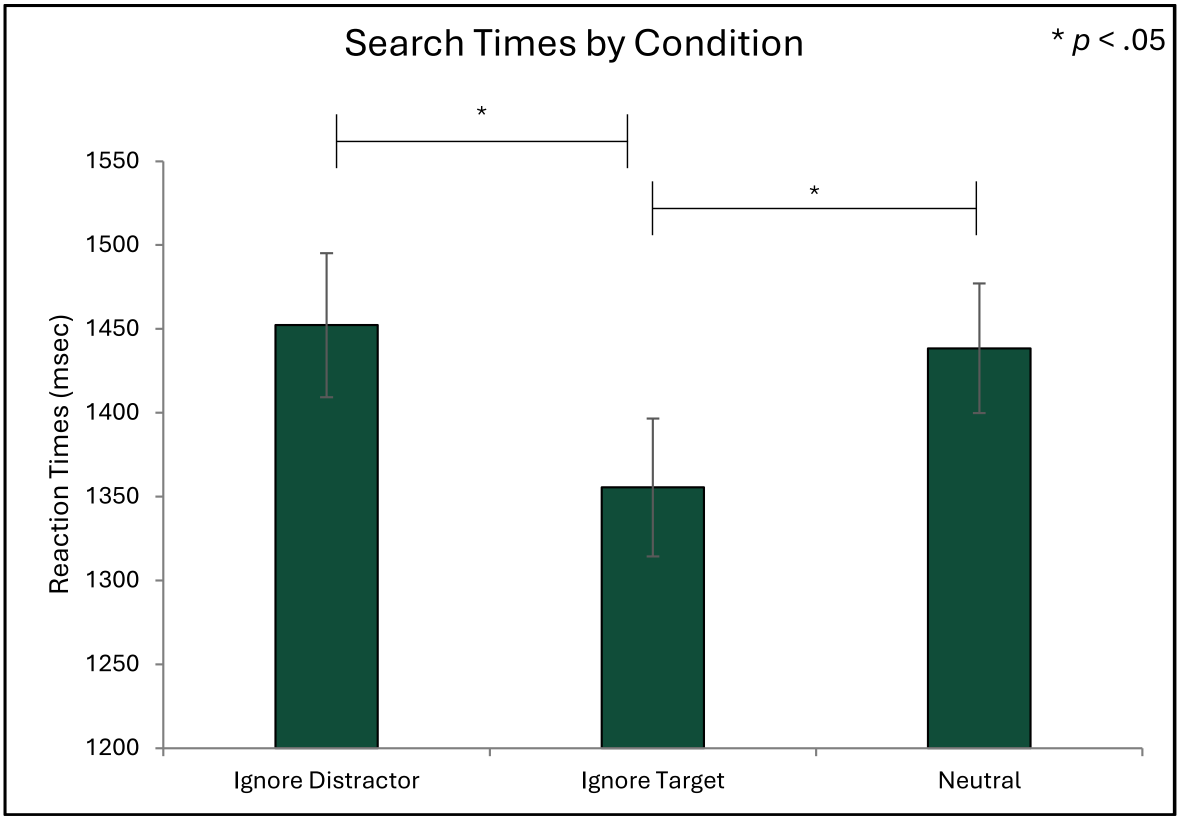

Figure 5. Search times for Experiment 2 by the three conditions.

in the ignore color (ignore target condition), when the six distractor letters appeared in the ignore color (ignore distractor condition), and when the ignore color did not appear as either target or six distractor letters (neutral condition).

Like Experiment 1, Experiment 2 also found no evidence for FBS. In fact, in Experiment 2, rather than finding evidence for FBS, we find evidence of attentional capture by the ignore color. A repeated measures ANOVA on the three conditions (ignore target, ignore distractor, and neutral conditions) revealed a difference in search times, $F(2, 118) = 12.83$, $p < .001$, with the ignore target RTs ($M = 1221.01$ ms, $SE = 35.13$ ms) being significantly faster than the ignore distractor ($M = 1320.10$ ms, $SE = 39.17$ ms), $t(59) = 4.52$, $p < .001$, and neutral conditions ($M = 1338.83$ ms, $SE = 40.50$ ms), $t(59) = 4.23$, $p < .001$ (figure 5). However, this result could have been a byproduct of explicitly instructing the participants to always ignore the same feature,

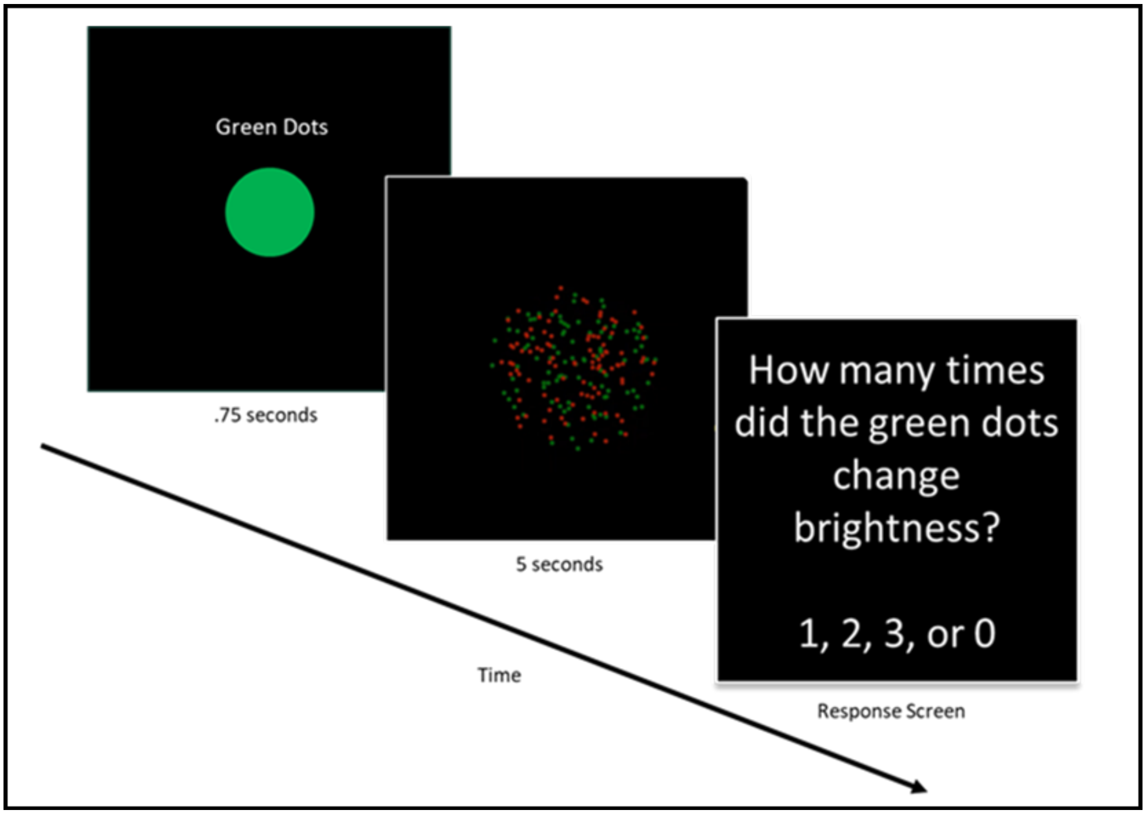

Figure 6. An example of an RDK task trial from Experiment 3. Participants are shown a 100% valid cue for 750 msecs, indicating which set of colored dots to attend to. After the cue is shown, the dot filtering occurs for 5 seconds followed by the response screen.

possibly biasing them to engage in a "search & destroy" search mode resulting in faster search times when the target is in the ignore color (Moher & Egeth, 2012). Given this possibility, we conducted another follow-up experiment.

Given there is evidence suggesting that implicitly learned feature-based suppression is more beneficial than explicitly learned/cued feature-based suppression (Addleman & Stormer, 2023; Stilwell & Vecera, 2019a; 2019b) and that this method could prevent participants from actively attending to the ignore color, Experiment 3 tested implicit feature-based suppression. For the RDK task, participants were given a 100% valid cue that told them which colored dots (either green, red, blue, or gray) to attend to as they changed on a trial-by-trial basis (see figure 6). The ignore dots were a given color (either red or green counterbalanced across participants) in 75% of all RDK trials. For the remaining 25% of trials, the color of the ignore dots were

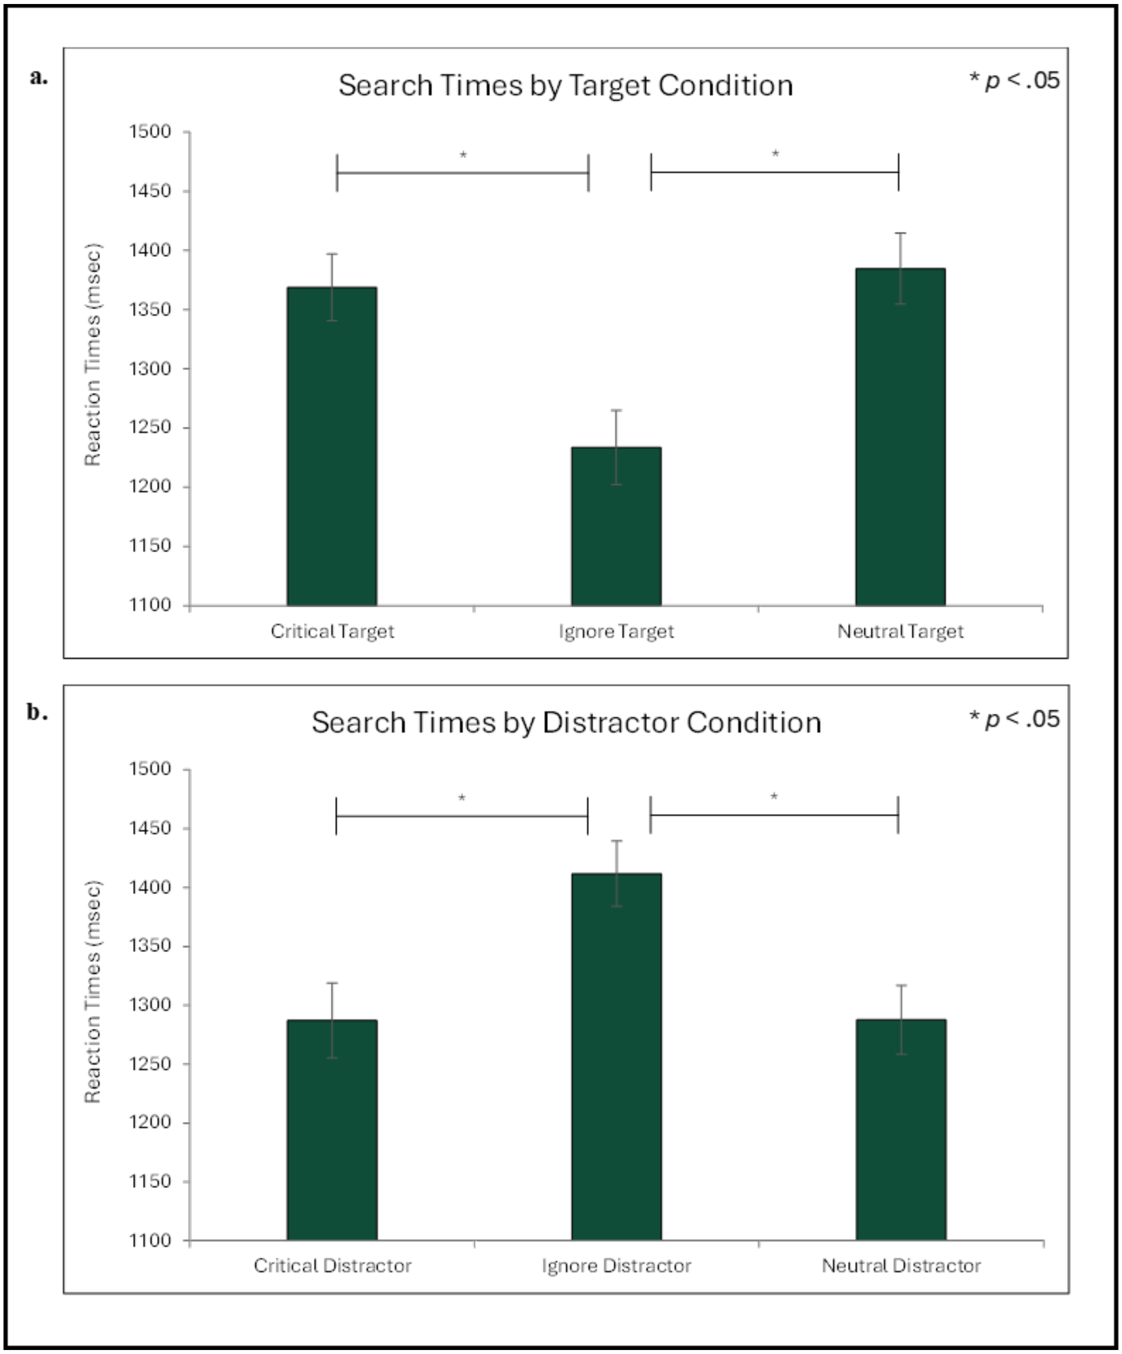

Figure 7. Search times for Experiment 3 by target color (figure 7a, top graph) and distractor color (figure 7b, bottom graph).

randomly assigned as either purple or cyan with the other color never appearing (this color is known as the "critical color" to differentiate them from the ignore and neutral colors). This approach included an explicit gain cue (that changed trial by trial) and implicit learning to suppress a given color. For the letter search task, the target letter appeared in either the ignore color, neutral color, or critical color. Similarly, the set of six similarly colored distractors appeared in either the ignore color, neutral color, or critical color.

Like Experiment 2, we found that the to-be-ignore feature produced attentional capture rather than suppression. An analysis of letter search RTs by target color (ignore, critical, and neutral color) found a significant difference, $F(2, 130) = 34.38$, $p < .001$, with faster RTs when the target was in the ignore color ($M = 1233.73$ ms, $SE = 31.28$ ms) than either the neutral ($M = 1384.74$ ms, $SE = 29.90$ ms), $t(65) = 7.05$, $p < .001$, or critical colors ($M = 1368.89$ ms, $SE = 28.24$ ms), $t(65) = 7.99$, $p < .001$ (figure 7a). An analysis of RTs as a function of the color of the set of six distractors that shared a color found a significant main effect, $F(2, 130) = 28.79$, $p < .001$, with slower RTs when the set of distractors appeared in the ignore color ($M = 1411.75$ ms, $SE = 27.63$ ms) than either the neutral ($M = 1287.61$ ms, $SE = 29.42$ ms), $t(65) = 6.31$, $p < .001$, or critical colors ($M = 1287.24$ ms, $SE = 29.42$ ms), $t(65) = 6.51$, $p < .001$ (figure 7b). This result was especially surprising given that most participants (75.71%) had no awareness of what the commonly ignored color was, making the explanation that they were actively attending to the to-be-ignored feature unlikely.

Altogether, we either found no evidence of feature-based suppression (Experiment 1) or evidence of attentional capture of the to-be-ignored feature (Experiments 2 & 3). The results from these studies raise serious concerns about the role that FBS plays in the allocation of attention and are broadly consistent with some recent criticisms of FBS. For instance, Kerzel and colleagues (Kerzel & Burra, 2020; Kerzel, Huynh Cong, & Burra, 2021) have recently criticized the FBS literature and suggested that FBS probably plays little or no role in the allocation of attention. However, before accepting that my prior results support the view that FBS plays no role in the allocation, I note that there are several limitations in the three experiments we ran.

First, it could be that suppression effects are built up and present in the RDK task, but do not transfer and generalize to the letter search task. While we used the same tasks and similar design as Moher, Lakshmanan, Egeth, & Ewen (2014) and they found that suppression effects transferred from the RDK task to the letter search task, that study may be an outlier in finding suppression effects that generalize across tasks. In fact, our letter search tasks began with the RDK task, so to the extent that filtering during the RDK task involved suppression of the ignored color, we would have expected the effect to have some impact on the search task. The finding that the suppressive effects do not generalize across tasks, even in this method, suggests that suppression may be very limited to within a particular task, raising questions about how central a role it plays in the typical allocation of attention. Even so, it may be ideal to have a task that remains the same throughout to study the possible benefits of feature-based suppression.

Furthermore, having the to-be-ignored feature occasionally be the target in the search task could be detrimental to development of feature-based suppression, in general, or may discourage participants from creating a suppressive component that generalizes across tasks. In our experiments, while the ignore color was never the target in the RDK task, it could be the target in the letter search task, potentially encouraging participants to ignore that color only in the RDK task. In addition, many studies find robust effects of suppression either never or rarely associate the to-be-ignored feature with the target (Addleman & Stormer, 2023; Stilwell & Vecera, 2020). It might be important for the development of feature-based suppression that the to-be-ignored feature is both constant and never a target. Again, while this approach might be ideal for demonstrating FBS effect, the need to do so to obtain FBS effects would limit how central FBS is to typical attentional allocation.

Lastly, it could be that the letter search task is not the most appropriate task to measure suppressive effects. Stilwell (personal communication, VSS 2024) noted that they had attempted to use the letter search task to evaluate suppressive effects but failed to find them. He hypothesized two potential reasons for this problem with letter stimuli. The first was that the color is not salient enough in a letter display as there is not much color shown (the only way to increase this is to either make the font boldface and/or larger). The second potential reason is that letters (and by extension, words) are processed in a way that observers read them automatically without much influence from the colored font. This led him and Vecera to use the colored squares with gaps (Stilwell & Vecera, 2020) to increase the amount of color and the saliency of the color presented in the search displays. With those stimuli, they did find suppressive effects. Thus, utilizing stimuli that increases the saliency and physical amount of the color instead of letters can increase the likelihood of finding suppression effects.

In summary, it is possible that FBS plays a role in the allocation of attention, but that it fails to generalize across tasks, may require circumstances where the to-be-ignored feature is never or almost never associated with the target, and where the to-be-ignored feature is salient. Considering these limitations from my previous experiments and the limitations from previous literature, that are number of factors and questions that need to be addressed to determine how prominent a role FBS plays in guiding attention.

Given that my pilot studies failed to find evidence of suppression, for my dissertation experiments, I will begin by attempting to find evidence of suppression in an “idealized” experimental design. Based on the literature reviewed above, this design would ideally avoid the need for cross-task generalization, use implicit learning of suppression, and employ stimuli where the color feature should be salient. Experiment 1 will implement this approach, and I expect to find evidence for suppression in this ideal method.

Assuming that this "ideal" method finds evidence for FBS, there remains the question of how many of these idealized factors can be altered before the suppression benefit disappears. Specifically, my further studies will investigate whether suppressive effects remain when the to-be-ignored color is made explicit, and when there is an ability to engage an attentional gain mechanism (that is learned implicitly or is made explicit).

If all of these studies find evidence for suppression it would suggest that suppression should play a prominent role in models of the allocation of attention, with the caveat that the effects may not generalize across tasks. If, however, the suppressive effects are disrupted by making the to-be-ignored feature explicit and/or allowing the implementation of a gain mechanism, the results would suggest that FBS is not a very robust mechanism and should not play a prominent role in models of the allocation of attention. Instead, it might be a mechanism that only is implemented in extremely limited circumstances, like the irrelevant singleton task, which are unlikely to occur in real-life.

CHAPTER 2:

THE BENEFIT OF IMPLICIT FEATURE-BASED SUPPRESSION IN AN IDEALIZED EXPERIMENT

2.1 Abstract

In the visual search literature, evidence for feature-based suppression, which is the suppression of irrelevant features that facilitates the selection of relevant items, has been mixed and controversial. Our previous attempts to find a benefit of feature-based suppression have either found no effect or the opposite, attentional capture of the to-be-ignore feature. However, those attempts utilized methods that may not have been conducive to finding a suppression benefit. Here, we designed an "idealized" task in attempt to find a baseline evidence of suppression benefit. Participants performed a Landolt C search task, where they searched for C with a horizontal gap among several distractor Cs with vertical gaps. This search task has two phases. The first phase (training task) had one color (ignore color) always appears as a distractor on 80% of trials to induce implicit learning of the to-be-ignored feature. The second phase (test block) removes the suppression contingency, with the previously ignored color now appearing as the target. Our results showed a suppression benefit in the training task, with faster searches when the ignore color was present compared to when it was absent. This suppression effect carried over to the test block, with the fastest searches occurring when the ignore color appears as a distractor and slowest searches occurring when the ignore color appears as the target, but this effect quickly dissipates. Thus, feature-based suppression can be beneficial in visual search, but potentially only under select circumstances.

2.2 Introduction

The goal of this experiment was to create an “idealized” task to find evidence that feature-based suppression can benefit visual search. Based on previous literature that found a suppression benefit, this task should require no generalization of suppression from one task to another (Britton & Anderson, 2020), have implicit learning of the to-be-ignored feature (Addleman & Stormer, 2023; Stilwell & Vecera, 2019a), have no feature-based gain component (Kawashima & Amano, 2022; Rajsic, Carlisle, & Woodman, 2020), and use stimuli where the color feature should be salient (Stilwell, personal communication, VSS 2024). In addition, an ideal task would include a post-training test block that removes the training contingencies, to eliminate the possibility that the results were influenced by differential intertrial priming effects across conditions (Addleman & Stormer, 2023; Britton & Anderson, 2020)

To meet these requirements, we designed a Landolt C search task with two phases. The first phase is a training task, where one color (which I will call the “ignore” color) appears on 80% of the trials and always appears as a distractor. Participants are not explicitly informed of this contingency with the expectation that they will implicitly learn to suppress this color. The second phase involved a test block, in which the contingency is removed so the ignore color appears at the same rate as other colors and is equally likely to be the target. Removing these contingencies equates the effect of inter-trial priming across all conditions, thereby allowing an evaluation of whether the suppressive effects are due to attentional biases or intertrial priming.

Given that we are using a single task design with implicit learning of to-be-ignored feature, salient colored stimuli, and no feature-based gain components, the expectation is that we will find an effect of feature-based suppression. To be more specific, we anticipate that there will be a suppression benefit in the training tasks, with faster reaction times (RTs) in trials where the ignore color appears compared to the trials where the ignore color is absent. Furthermore, we anticipate that this suppression benefit will increase in magnitude as learning increases, as previously found by Cunningham & Egeth (2016) and Vatterott & Vecera (2012), with the suppression benefit being weak or non-existent in the first quarter of the training task and the suppression benefit being the strongest in the last quarter of the training task. We also predict that suppression benefit will carryover into the test block, once the contingency is removed. However, this suppression effect should be weaker and rapidly disappear as participants learn that the ignore color can now appear as the target.

2.3 Methods

2.3.1 Participants

Sample size was estimated a priori based on Cohen's recommended small effect size of d = 0.2 (Cohen, 1992). Using G*Power version 3.1.9.7 with power = .95 and $\alpha = .05$ for repeated measures ANOVA with between-subjects interaction for the cross-experiment analyses, this gave a minimum sample size of 56.

As such, 59 undergraduate students from Michigan State University were recruited to participate in the experiment. Data from one participant was excluded from further analysis for having error rates greater than 3 standard deviations (SDs) from the overall mean, leaving 58 participants whose data was analyzed. All the participants had normal or corrected-to-normal acuity. Participants gave written informed consent under the study protocol approved by the Institutional Review Board and were compensated with research credits.

2.3.2 Stimulus and apparatus

The experiment was programmed using PsychoPy v. 2021.2.3 (Peirce, 2007; Peirce et al., 2019) set at 1920x1080-pixel resolution with a 120-Hz refresh rate. Participants viewed the

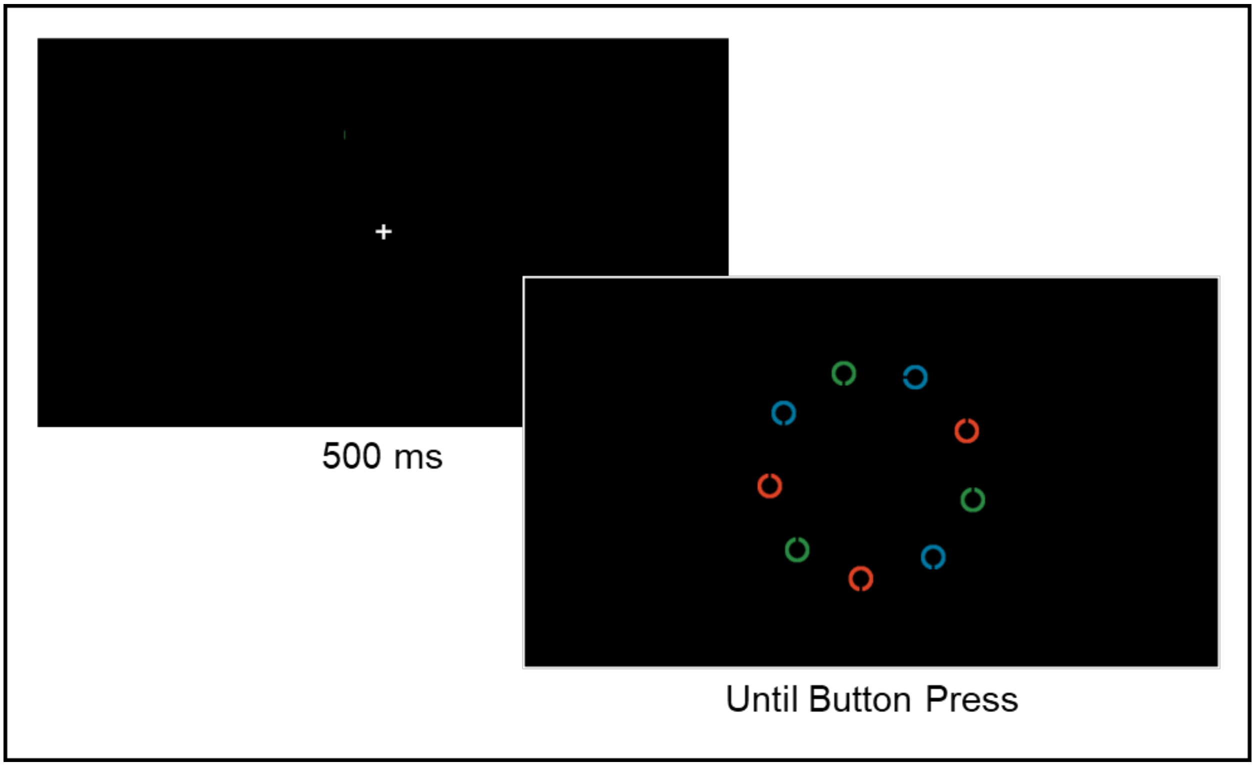

Figure 8. An example of a possible trial in the training task. Each trial begins with a fixation cross appearing for 500 ms. When the fixation cross disappears, the search display appears, consisting of nine Landolt Cs appearing in a circular array 40° away from each other, and stays on screen until a button press is made. Eight Cs are distractors with vertical gaps (either upwards or downward facing) and the target is the C with the horizontal gap (either rightward or leftward facing). Participants indicate the location of horizontal gap via button press. On a given trial, three colors appear, with three Cs in each color. The ignore color (for example, the red Cs) appear on 80% and are never the target.

screen from approximately 55 cm away in a dark, sound attenuated booth, and responded via keyboard button presses and mouse presses.

The search array consisted of nine stimuli, one target among eight distractors, in a clock face ring (radius ~ 8.32°). The stimuli were Landolt Cs (1.56°) that were equally spaced around an imaginary clock face. The line width of the Cs was .31° and the gaps were .21°. The Cs were presented on a black background and could appear in six distinct colors (red, green, yellow, blue, pink, or brown). Distractors had vertical gaps (either upward or downward) while the target had a horizontal gap (either leftward or rightward). Participants had to indicate whether the target's gap was leftward or rightward facing via button press.

2.3.3 Procedure

Each trial begins with a fixation cross (.52°) appearing for 500 ms, after which the fixation cross disappeared and the search array appeared. Participants searched for the target and reported its gap location using the left and right arrow keys on the keyboard. We instructed each participant to respond as quickly and as accurately as possible.

Participants began the experiment with a practice block, consisting of 24 trials. Here every C appeared gray and accuracy feedback was given after each trial. After successful completion of the practice block, participants moved on to the training blocks.

For the training task, each trial had three colors with three items appearing in each color (figure 8). Critically, the colors were spatially intermixed to prevent three Cs of the same color from appearing together. However, it was possible for two neighboring Cs to be of the same color. For a given subject, the ignore color was set to be either red or green. This ignore color appeared as a distractor color in 80% of trials and the target never appeared in this color. Participants were not explicitly told of this contingency, with the expectation that they would implicitly learn to suppress this color over time. Suppressing attention to the ignore color would reduce the effective set size by 33.3%, which should encourage the use of feature-based suppression. To the extent that this suppression occurred, RTs should be faster in trials in which the ignore color was present (ignore trials) compared to the trials where the ignore color was absent (neutral trials).

The target appeared equally often in each of the remaining five colors. Each of these five colors also appeared as a distractor color an equal number of times. Given that the ignore trial

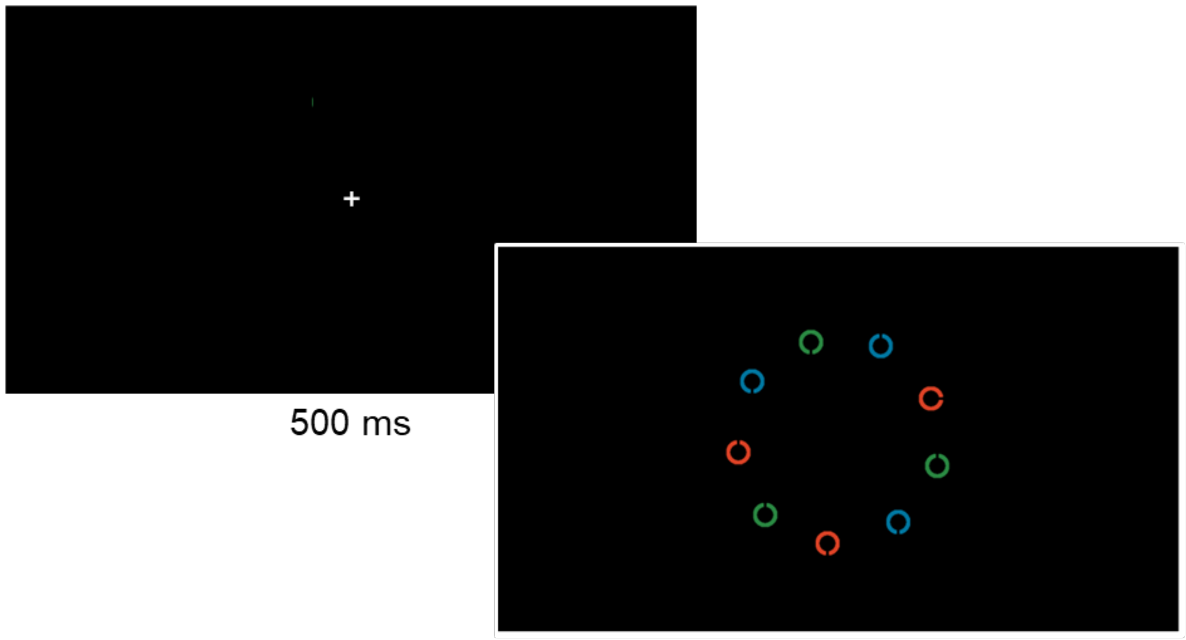

Figure 9. An example of a possible trial in the test block. The target (C with horizontal gap) can now appear in the ignore color (red in this example) when it previously could not in training task.

appeared in 80% of trials, there were 20% of trials in which both distractor colors were selected from the remaining five colors. As a result, across trials each of these five colors was the target color in 20% of trials and each of these colors appeared as a distractor on 24% of trials. Participants performed 240 trials of the training task split into two blocks of 120 trials, with a rest in between.

After the training task, participants completed a final test block consisting of 60 trials. In this block, the suppression contingencies from the training blocks were removed, with each of the six colors appearing as the target on 10 trials and each color appearing as one of the two distractor colors on 20 trials. Thus, the ignore color now appeared as the target on some trials (the ignore target condition) and as one of the two distractors on other trials (ignore distractor

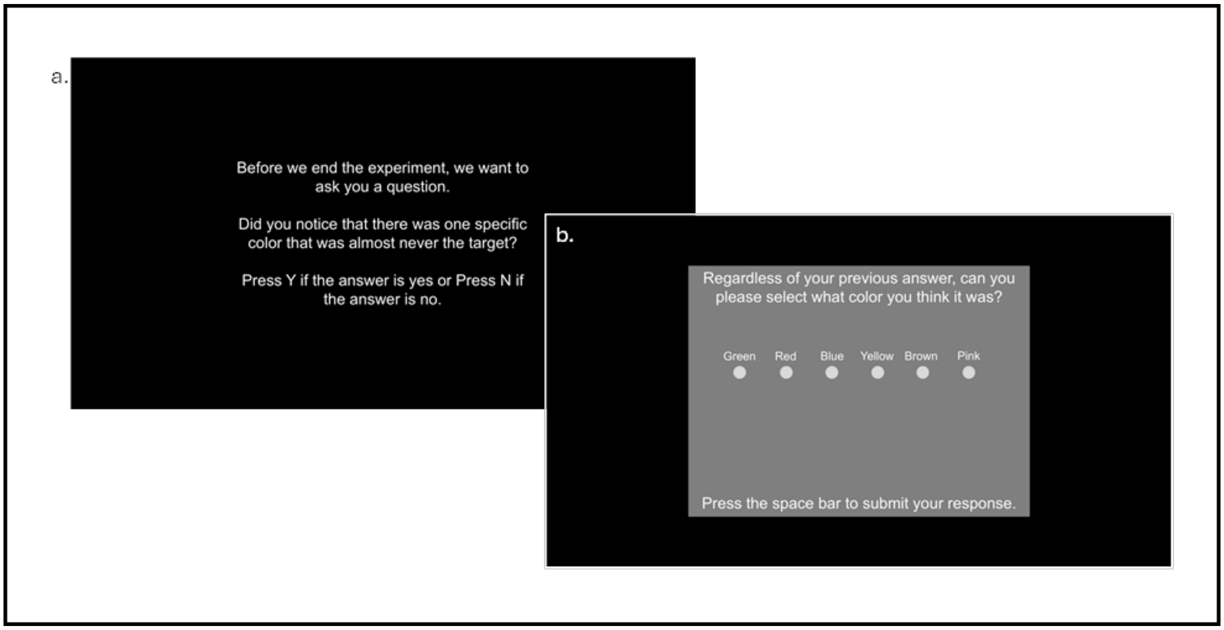

Figure 10. The end of experiment suppression contingency checks. (a.) Participants are first asked if they had any awareness of one color that almost never was the target and answer via button press. (b.) Following the first question, they then had to choose which color was the ignore color via mouse selection.

condition). There were also trials where the ignore color did not appear (neutral condition) (figure 9). If the ignore color was suppressed, we would expect the fastest RTs in the ignore distractor condition, followed by the neutral condition, with the ignore target condition having the slowest RTs.

After the experiment was completed, participants were assessed to see if they had any explicit awareness that the ignored color rarely was the target. They were first asked if they had any awareness (yes or no question) of one specific color rarely containing the target (figure 10a). Regardless of the answer to the first question, they were given a six alternative forced choice task in which they were asked to indicate which color was seldom the target (figure 10b).

2.4 Results

2.4.1 Accuracy

Accuracy was generally high across the training blocks (M = 99.20%; SE = .11%) and the

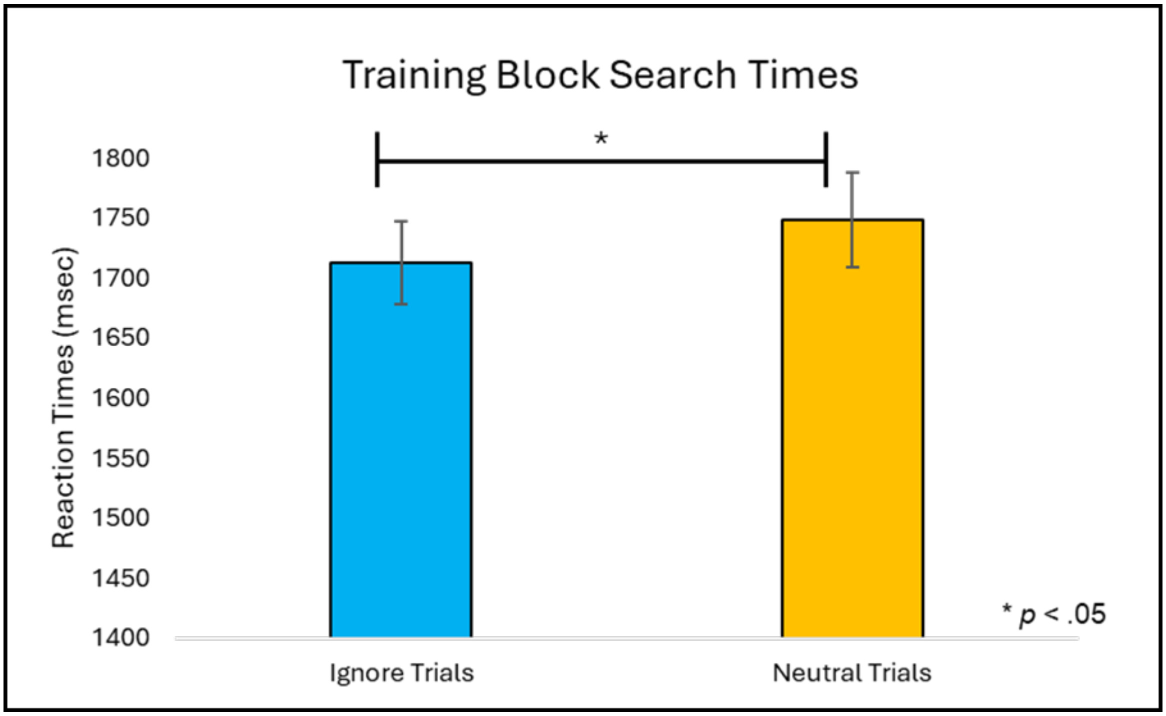

Figure 11. Search times for the training task.

test block (M = 99.22%; SE = .24%), likely indicating that participants were performing near ceiling, and thus it is unlikely that accuracy rates would differ between conditions in training or test blocks. Nevertheless for completeness, comparisons on accuracy were performed. A paired samples t-test on accuracy for the training block showed that accuracy did not differ between the ignored and neutral trials, $t(57) = .56$, $p = .576$. A repeated measures ANOVA on accuracy in the test block on the three conditions (ignore target, ignore distractor, and neutral conditions) also showed no difference in accuracy, $F(2, 114) = 1.38$, $p = .255$, partial eta squared = .024.

2.4.2 Training Block

Trials with reaction times (RTs) more than 3 SDs away from the overall mean and with incorrect responses were excluded from the reaction time analyses, resulting in removal of .80% of total trials. To examine if there is a benefit of implicit feature-based suppression, we

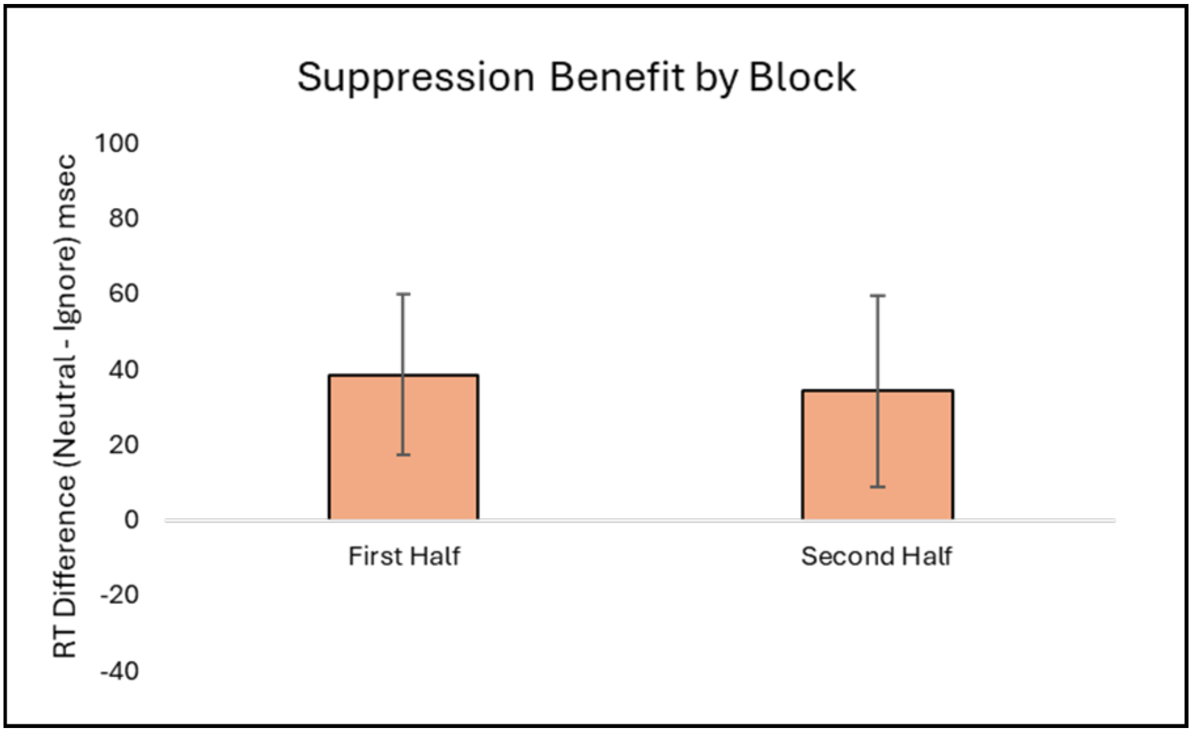

Figure 12. The suppression benefit comparing the first half of the training task to the second half of the training task.

conducted a paired samples t-test on RTs between the ignored and neutral trials. This analysis revealed that there was indeed a benefit of feature suppression, with faster RTs for the ignored trials (*M* = 1712.35 ms, *SE* = 34.49 ms) compared to the neutral trials (*M* = 1748.28 ms, *SE* = 39.66 ms), $t(57) = 2.05$, $p = .045$ (figure 11).

Given that previous research has found that suppression effects build over time and are the strongest after prolonged periods of training (Cunningham & Egeth, 2016; Vatterott & Vecera, 2012), we conducted follow-up block analysis splitting the training block into halves (first half vs. second half) to examine if there is a difference in the suppression benefit (difference in RTs by subtracting the ignore trials from the neutral trials). A paired samples t-test revealed no difference in suppression benefit between the first half (*M* = 38.58 ms, *SE* = 21.22 ms) and second half (*M* = 34.18 ms, *SE* = 25.23 ms) of the training block, $t(57) = .15$, $p = .882$,

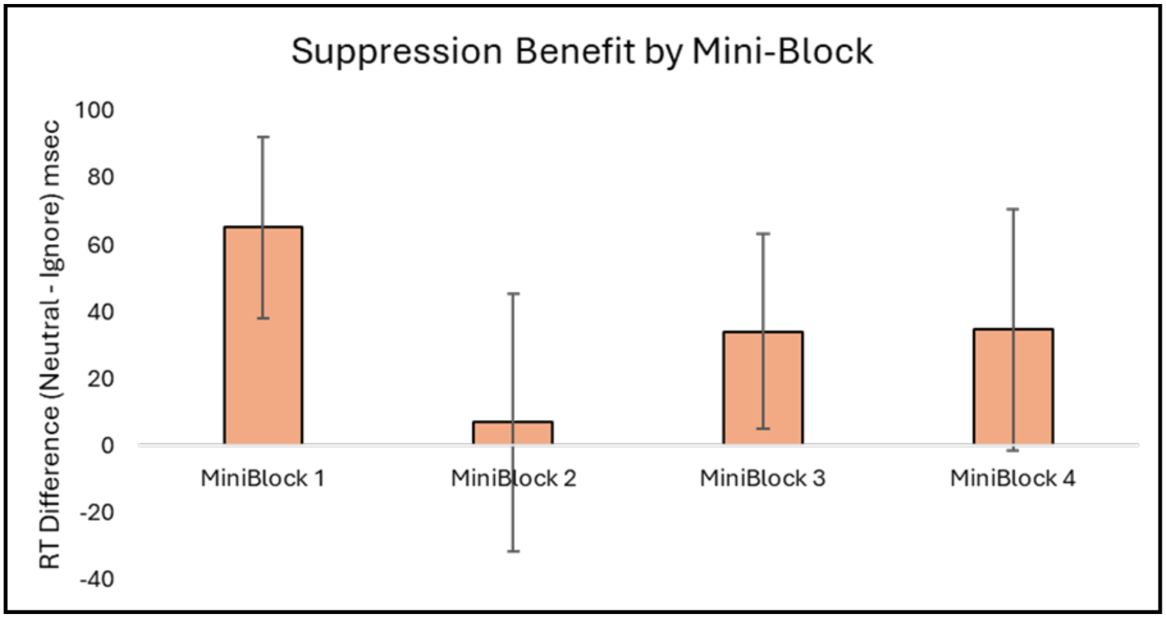

Figure 13. The suppression benefit splitting the training task into quarters.

indicating there is not a greater benefit of suppression from prolonged periods of learning (figure 12).

We also decided to investigate the suppression benefit by quarters as it may reveal more nuanced learning that is lost by examining the suppression benefit by halves. However, this analysis still found no differences in suppression benefit across the first quarter (*M* = 65.18 ms, *SE* = 27.00 ms), second quarter (*M* = 6.91 ms, *SE* = 38.63 ms), third quarter (*M* = 34.07 ms, *SE* = 29.13 ms), and fourth quarter (*M* = 34.62 ms, *SE* = 35.95 ms) of trials (figure 13), $F(3, 171) = .51$, $p = .680$, partial eta squared = .009. While these results are surprising given that previous studies found that feature-based suppression is initially costly towards search performance and takes time for it to become beneficial (Cunningham & Egeth, 2016; Vatterott & Vecera, 2012), those studies investigated explicit feature-based suppression rather than the implicit feature-based suppression. Considering that implicit feature-based suppression tends to be more robust than explicit feature-based suppression (Addleman & Stormer, 2023; Stilwell & Vecera, 2019), a crucial difference between the two could be that implicit suppression is beneficial after a few

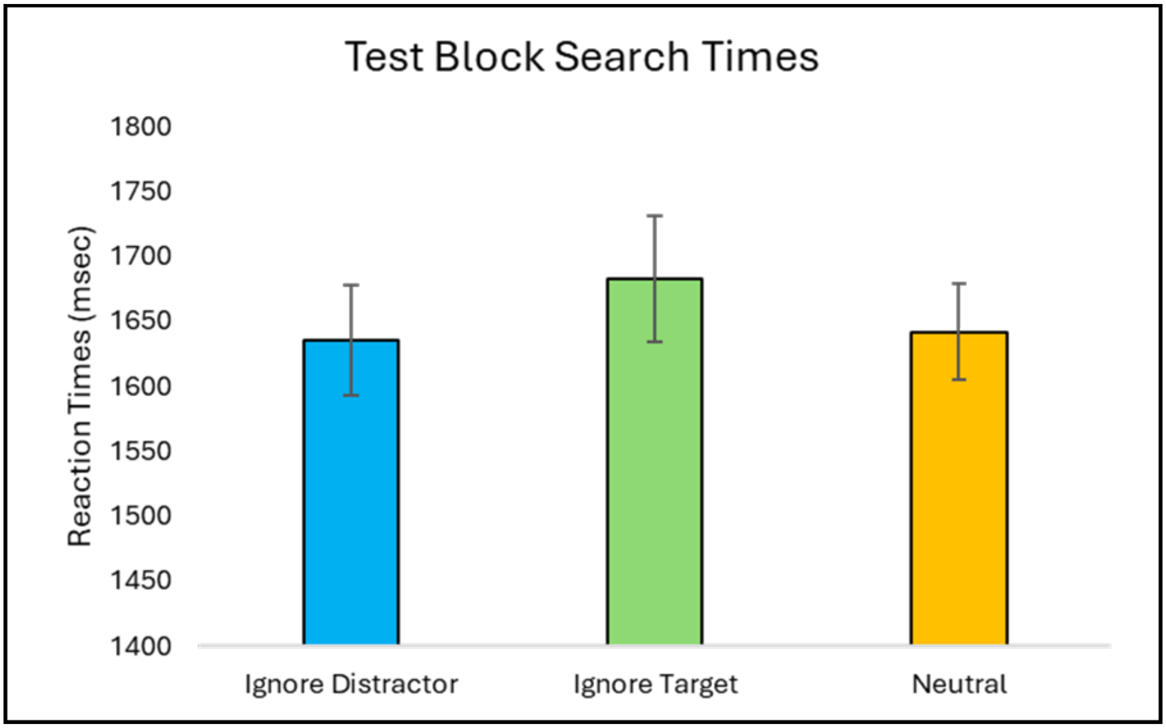

Figure 14. Search times in the test block.

trials (Golan & Lamy, 2022) while explicit suppression takes time to overcome the initial attentional capture of the to-be-ignored feature.

2.4.3 Test Block

To determine if the suppression effect persists in the test block without the inter-trial priming from suppression contingency in the training block, we conducted a repeated measures ANOVA on RTs for the three conditions (ignore target, ignore distractor, and neutral conditions). This analysis revealed that there were no differences in RTs between the ignore target (M = 1682.58 ms, SE = 48.81 ms), ignore distractor (M = 1635.18 ms, SE = 42.97 ms), and neutral (M = 1641.57 ms, SE = 36.82 ms) conditions (figure 14), $F(2, 114) = 1.10$, $p = .335$, partial eta squared = .019.

However, since there is a possibility that the suppression effects may quickly disappear, we conducted a similar analysis but just on the first half of the test block. Here we find a

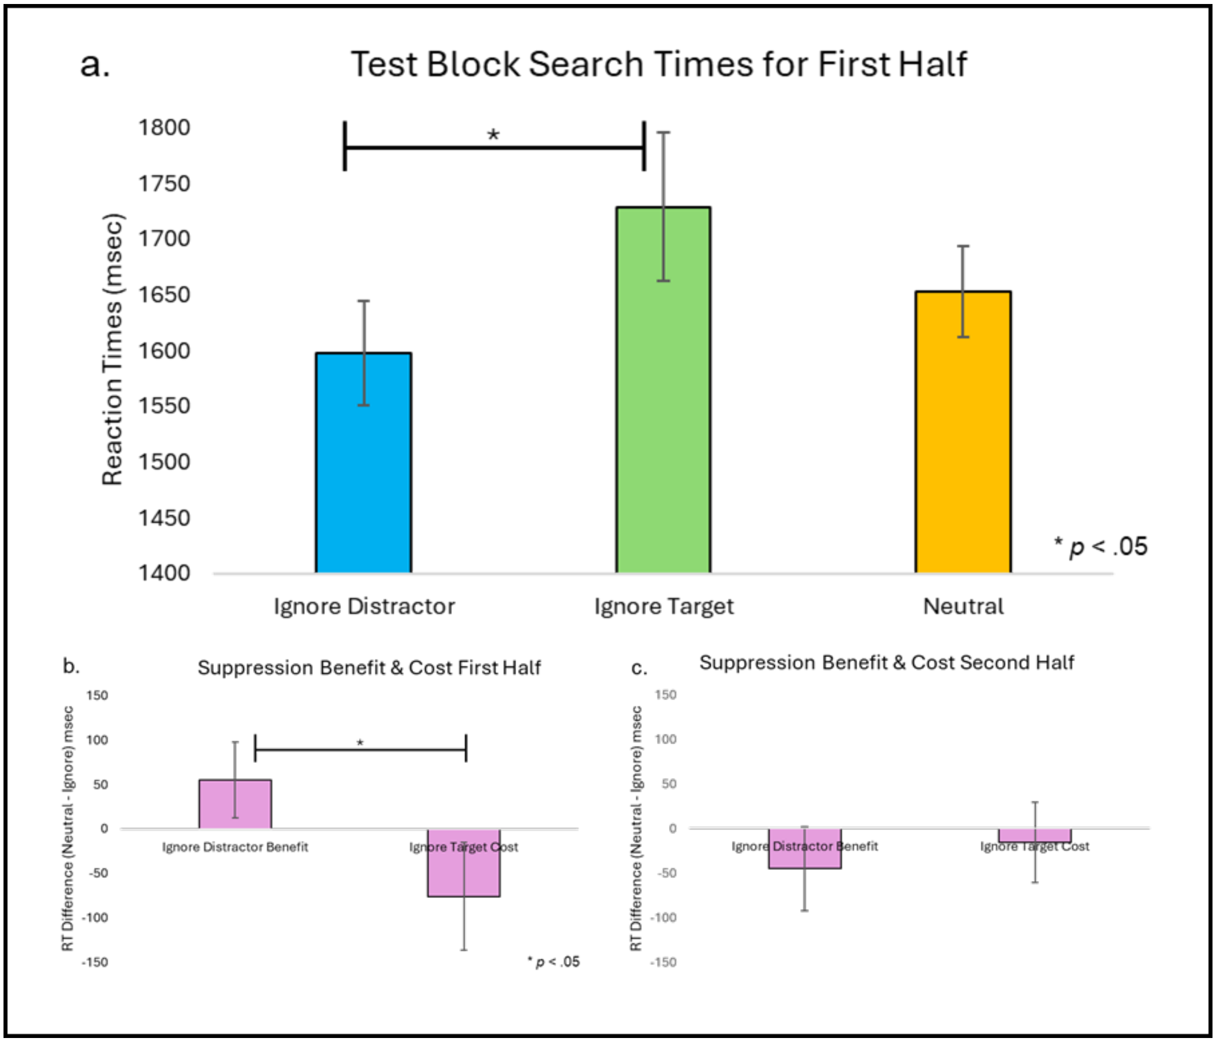

Figure 15. (a.) Search times for the first half of the test block. (b.) Comparing the suppression benefit (subtracting the RTs of ignore distractor condition from the neutral condition) to the suppression cost (subtracting the RTs of ignore target condition from the neutral condition) in the first half of the test block. (c.) Comparing the suppression benefit and suppression cost in the second half of the test block.

marginal difference in RTs (figure 15a), $F(2, 114) = 3.00$, $p = .054$, partial eta squared = .050, with the ignore distractor condition having the fastest search times ($M = 1598.17$ ms, $SE = 46.70$ ms) followed by the neutral condition ($M = 1653.11$ ms, $SE = 40.71$ ms) with the slowest search times occurring in the ignore target condition ($M = 1729.08$ ms, $SE = 66.38$ ms). Follow-up pairwise comparisons revealed that the source of this difference is that the ignore distractor and ignore target conditions are significantly different from each other, $t(57) = 2.30$, $p = .025$. This result indicates that there is an effect of suppression early in the test block, which benefits search when the distractor appears in the ignore color and is costly when the target appears in the ignore color (figure 15b), but rapidly disappear as participants learn that the ignore color can now be the target (figure 15c).

2.5 Discussion

The goal of this study was to find a benefit of feature-based suppression in a visual search task designed to be an ideal case for finding the effects of suppression. This idealized method produced evidence for feature-based suppression. There was a suppression benefit in the training task, with faster search times when the ignore color appeared as one of the distractors compared to the neutral trials. This suppression benefit also maintained for the first half of the test block, with the fastest search times occurring when the ignore color appeared as a distractor and slowest search times occurring when the target appeared in the ignore color. This maintenance rules out an explanation based solely on intertrial priming. While the literature on feature-based suppression has been mixed, these results show that under the right circumstances, feature-based suppression can guide attention effectively.

Given that we designed the method of this experiment to be the idealized circumstance to find evidence for suppression in a visual search task, if this method had not produced evidence of a suppression benefit, it would have provided a serious challenge to the view that feature-based suppression is an important mechanism for allocating attention. The evidence for suppression in this idealized method, suggests that suppression can play a role in the allocation of attention, but questions remain about how prominent a role it plays is still debatable – if it only plays a role in the idealized circumstance, it likely is not a mechanism that is often influencing attention.

Even so, our method may provide more evidence for the use of suppression in everyday activities relative to other paradigms that frequently find evidence of a suppression benefit. For example, the additional singleton task is a prevalent paradigm in studies supporting the role of feature-based suppression in guiding attention (Adam & Serences, 2021; Gaspelin, Leonard, & Luck, 2015; 2017). In this research, the additional singleton distractor initially captures attention and suppression effect build up over time after prolonged exposure to the same distractor (Vatterott & Vecera, 2012). However, this paradigm may have limited real-world applicability. In this method, the additional singleton is a distractor that is salient because it is the one item with a unique feature (color) in an otherwise homogenous field. It is very unlikely that real-world searches would involve repeated searches through homogeneous items that have a single distractor that violates the homogeneity and appears repeatedly. Thus, while this paradigm may be an effective method for finding feature-based suppression, it may be impractical for studying real-world feature-based suppression. By contrast, our search paradigm utilized heterogeneous distractors and targets, providing a more realistic search display, while also finding a suppression benefit.

Similarly, the Landolt C search paradigm in Arita, Carlisle, & Woodman (2010) may be problematic because it used a search array that was always comprised of two colors that were segregated by hemifield. While they found a suppression benefit from a negative cue (i.e., a cue that indicates the color that the target will not appear in), there have been many criticisms of their task design. Beck & Hollingworth (2015) and Becker, Hemsteger, & Peltier (2015) both had alternative hypotheses to explain Arita et al.'s supposed suppression benefit. Beck & Hollingworth (2015) argued that participants were merely spatially recoding the negative cue and found that the negative cue was ineffective in guiding attention when the search array was intermixed, leading to the *spatial cue recoding hypothesis*. Becker, Hemsteger, & Peltier (2015) had a similar argument, claiming that participants were using the negative cue to create a positive cue for the other color in the display, and similarly found that the negative cue was ineffective when the non cued items in the array were heterogenous colors (leading to the *color recoding hypothesis*). By contrast, our task spatially distributed Cs of the same color, ruling out that any suppression benefit is a product of spatial recoding. Furthermore, when the ignore color appeared in our study, there were two other colors in the display and these colors changed on a trial-by-trial basis. We believe this approach reduced the likelihood that participants were setting up positive cues for the colors that were not the ignore color. In short, we believe our results show the effect of suppression that cannot be explained by spatial or color recoding to create a positive cue.

It is worth noting that when we examined the suppression benefit in the training task by mini-block, the magnitude of the suppression benefit was similar across the four quarters. This was a surprising result given that we expected the suppression benefit to increase as a product of learning and repetition of the to-be-ignored color. However, it appears that our participants quickly learned the to-be-ignored color and suppressed it almost immediately (Golan & Lamy, 2022). This finding is in stark contrast to previous studies (Cunningham & Egeth, 2016; Vatterott & Vecera, 2012) that found an initial cost of the to-be-ignore feature which only switches to a benefit after significant training and repetition of the to-be-ignore feature.

A key difference between our study and those that find that suppression effects take time to build is that our study utilizes implicit FBS while those previous studies used explicit FBS. Some studies (Stilwell & Vecera, 2019) have reported that implicit feature-based suppression is more robust and effective than explicit feature-based suppression, with some concluding that implicit FBS is as effective as feature-based gain (Addleman & Stormer, 2019). With explicit FBS, the initial bias to attend to the to-be-ignored feature may be a product of holding that feature in active working memory. Only through repeated exposure with the to-be-ignored feature can one learn to suppress this initial bias towards it and successfully ignore it. In contrast with implicit FBS, since observers are not consciously aware of the to-be-ignored feature, perhaps that feature is not actively stored in working memory, and thus, there is not an initial bias towards that feature. Given that 1.72% of our participants had explicit knowledge of the to-be-ignored feature, it gives some plausibility that there was no top-down bias towards that feature, leading to faster learning of the to-be-ignored feature.

While we find evidence for suppression in a paradigm that may have more real-world validity than some of the prior methods, if suppression is a major factor in the guidance of attention in real-world search tasks, it should maintain even if we make the methods less ideal. It is important to determine if feature-based suppression mechanism is still utilized if a feature-based gain component is also present, as previous research has found that that observers ignore or do not engage with FBS if feature-based gain is an available mechanism (Kawashima & Amano, 2022; Rajsic, Carlisle, & Woodman, 2020). However, other research (Stilwell & Vecera, 2020) found that feature-based suppression and gain are independent processes that can both be utilized to improve search performance, leading to the most efficient searches when both mechanisms are used. Thus, it is unclear whether feature-based suppression can still be beneficial if a feature-based gain mechanism can be used.

As such, we decided to replicate the current Landolt C search task but also introduce an implicit feature-based gain component. To do so, there is one color (attend color) that appears on a large portion of trials and always appears as the target. If we find a suppression effect even with an implicit gain component, it provides evidence that feature-based suppression has real-world validity and can benefit everyday searches. However, if the suppression benefit is nullified by target guidance, then it limits the real-world applicability of feature-based suppression as it can only be implemented and effective in an idealized and controlled search paradigm.

CHAPTER 3:

NO BENEFIT OF IMPLICIT FEATURE-BASED SUPPRESSION WITH AN IMPLICIT FEATURE-BASED GAIN COMPONENT

3.1 Abstract

There is some evidence suggesting that the suppression of irrelevant features, also known as feature-based suppression, facilitates the selection of relevant items and is an important mechanism for allocating visual searches. However, in many of those studies, there was the absence of feature-based gain, a mechanism that upweights features associated with the target. As such, it is inconclusive if feature-based suppression is still beneficial in visual search if observers can also use feature-based gain. To address this, we replicated the Landolt C task from Exp. 1 but included an implicit feature-based gain component in addition to an implicit suppression component. In the training task, we had one color (attend color) always appears as the target on 40% of trials while another color (ignore color) always appears as the distractor on 80% of trials. The follow-up test block removed both gain and suppression contingencies to see if both persist without inter-trial priming. Our results showed a gain benefit in the training task, with fastest and most accurate searches occurring when the attend color was present. This gain benefit carries over to the test block. However, there was no suppression effect in either the training task or in the test block. Our findings suggest that feature-based suppression may only be beneficial if it is the only mechanism available, potentially limiting the real-world applicability of it and possibly suggesting that it is not an important mechanism for allocating visual attention.

3.2 Introduction

Given that we found that implicit feature-based suppression (FBS) benefitted visual search in Exp. 1 without a feature-based gain (FBG) component, there is still the question of whether FBS can still be beneficial and applied in visual search if there is also a FBG mechanism that can be utilized. For FBS to be considered an influential mechanism in allocating attention, it should also be efficacious in guiding attention when there is also the possibility to engage a FBG mechanism.

Previous literature has been mixed on whether FBS is beneficial if there is also a FBG component. Stilwell and Vecera (2020) found that implicit FBS benefited search even when providing an explicit target color cue that was 100% valid, and that the effects of FBG and FBS were additive when they were both present on a search trial. However, others (Kawashima & Amano, 2022; Kawashima & Matsumoto, 2018; Kugler et al., 2015) have found that observers prefer to utilize FBG mechanism in lieu of FBS, when both are possible. For example, Rajsic, Carlisle, & Woodman (2020) found that when participants are provided with both positive and negative cues, that participants ignore the negative cue in favor of the positive cue and encode the positive cue to working memory while the negative cue is not encoded. Given the uncertainty of FBS's role in allocating attention in the presence of FBG mechanism, the goal of this experiment is to find evidence that FBS is still beneficial in visual search even when there is a FBG component.

To test this, we used the same Landolt C search task from Exp. 1 but added an implicit FBG component. In the training task, there was one color (which I will call the "attend" color) that appeared in 40% of the trials and always appeared as the target. This in contrast to the "ignore" color, which appeared in 80% of trials and always appears as a distractor. While the ignore color appears twice as often as the attend color, these values were chosen to equalize the amount of information provided by each mechanism as the attend color reduces the search array by six items (from nine Landolt Cs down to three) as only three Landolt Cs appear in the attend color, while the ignore color reduces the search array by three items (from nine Landolt Cs down to six). Thus, we attempted to make both mechanisms equally beneficial to the visual search task by doubling the frequency of the suppression cue to make up for the fact that it was less spatial informative within a given trial. We hoped this would encourage participants to utilize both mechanisms. In the test block, the suppression and gain contingencies were removed so that both the ignore and attend colors appear at the same rate as the other colors and both were equally likely to appear as either the target or a distractor. Removing these contingencies equated the effect of inter-trial priming across all conditions, thereby allowing an evaluation of whether the suppressive and gain effects are due to attentional biases or intertrial priming.

There are three potential patterns of results. One, the inclusion of a FBG component could reduce or eliminate any suppression effects (Kawashima & Amano, 2022; Kawashima & Matsumoto, 2018; Kugler et al., 2015; Rajsic, Carlisle, & Woodman, 2020), resulting in faster RTs and potentially higher accuracy in trials where the attend color appears compared to the trials where the attend color is absent, but no improvements in either RTs or accuracy in trials where the ignore color appears. Two, participants might be able to utilize both the FBG and FBS mechanisms, resulting in faster RTs and potentially higher accuracy in trials where either the attend or ignore colors appear, but that these mechanisms may not be additive, meaning in trials where both the attend and ignore colors appear together are not faster or more accurate than trials where the attend or ignore colors appear alone. Third, participants can utilize both the FBG and FBS components and that these two mechanisms are additive (Stilwell & Vecera, 2020), resulting in fastest and most accurate searches when both the attend and ignore colors appear together. Since we found a robust suppression benefit in Exp. 1, my hypothesis is that we will find evidence of a suppression benefit even with the inclusion of FBG component. Additionally, there will also be gain benefit when the attend color appears as the target. However, given that Stilwell & Vecera (2020) is the only study to have reported an additive effect of FBS and FBG, I do not believe there will be an additive effect when both the attend and ignore colors appear together.

3.3 Methods

3.3.1 Participants

Sample size was estimated a priori based on Cohen's recommended small effect size of d = 0.2 (Cohen, 1992). Using G*Power version 3.1.9.7 with power = .95 and $\alpha = .05$ for repeated measures ANOVA with between-subjects interaction for the cross-experiment analyses, this gave a minimum sample size of 56.

As such, 65 undergraduate students from Michigan State University were recruited to participate in the experiment. Data from one participant was excluded from further analysis for having error rates greater than 3 SD from the overall mean, leaving 64 participants whose data was analyzed for the behavioral data. All the participants had normal or corrected-to-normal acuity. Participants gave written informed consent under the study protocol approved by the Institutional Review Board and were compensated with research credits.

3.3.2 Stimulus and apparatus

The experiment was programmed using PsychoPy v. 2021.2.3 (Peirce, 2007; Peirce et al., 2019) set at 1920x1080-pixel resolution with a 120-Hz refresh rate. Participants viewed the

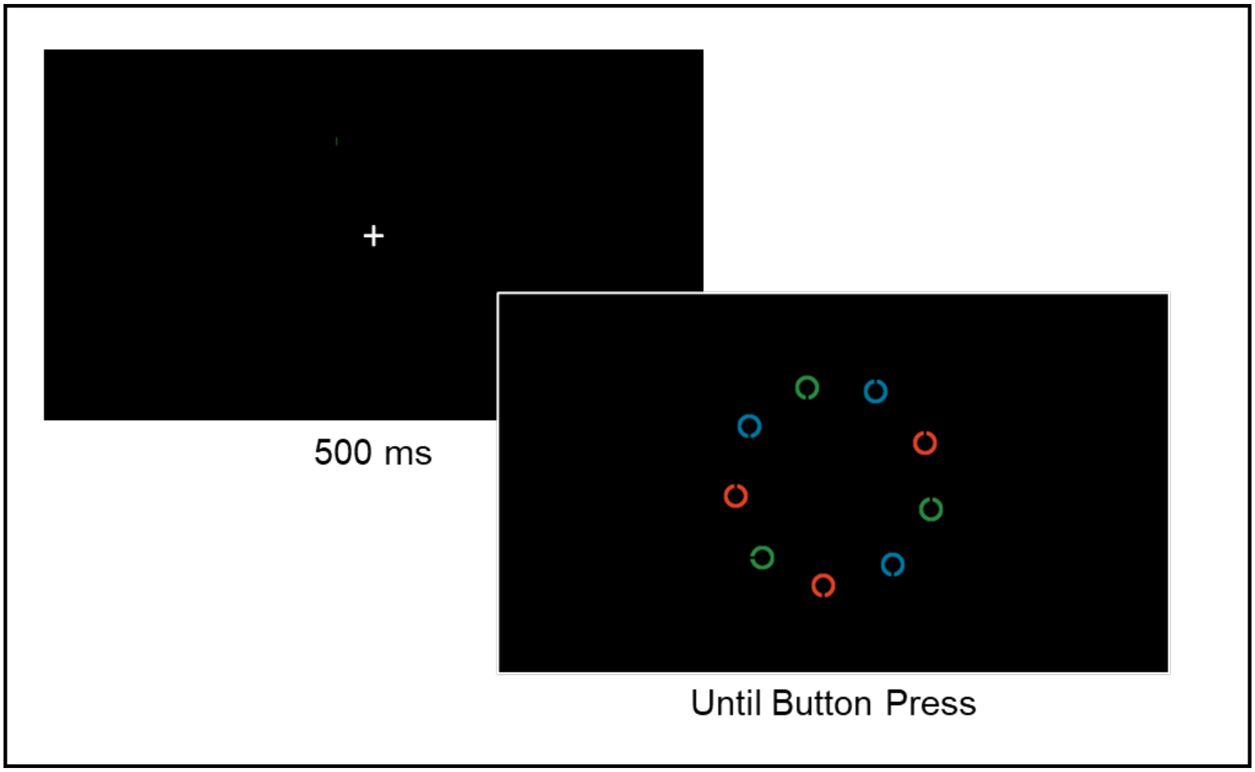

Figure 16. An example of a possible trial in the training task. Each trial begins with a fixation cross appearing for 500 ms. When the fixation cross disappears, the search display appears, consisting of nine Landolt Cs, and stays on screen until a button press is made. Eight Cs are distractors with vertical gaps (either upwards or downward facing) and the target is the C with the horizontal gap (either rightward or leftward facing). Participants indicate the location of horizontal gap via button press. On a given trial, three colors appear, with three Cs in each color. The ignore color (for example, the red Cs) appeared on 80% of trials and was never the target while the attend color (for example, the green Cs) appeared on 40% of trials and always as the target.

screen from approximately 55 cm away in a dark, sound attenuated booth, and responded via keyboard button presses and mouse presses.

The search array consisted of nine stimuli, one target among eight distractors, in a clock face ring (radius ~ 8.32°). The stimuli were Landolt Cs (1.56°) that were equally spaced around an imaginary clock face. The line width of the Cs was .31° and the gaps were .21°. The Cs were presented on a black background and could appear in six distinct colors (red, green, yellow, blue, pink, or brown). Distractors had vertical gaps (either upward or downward) while the target had a horizontal gap (either leftward or rightward). Participants had to indicate whether the target's gap was leftward or rightward facing via button press.

3.3.3 Procedure

Each trial begins with a fixation cross (.52°) appearing for 500 ms, after which the fixation cross disappeared, and the search array appeared. Participants searched for the target and reported its gap location using the left and right arrow keys on the keyboard. We instructed each participant to respond as quickly and as accurately as possible.

Participants began the experiment with a practice block, consisting of 24 trials. Here every C appeared gray and accuracy feedback was given after each trial. After successful completion of the practice block, participants completed the training blocks.

For the training task, each trial had three colors with three items appearing in each color (figure 16). Critically, the colors were spatially intermixed to prevent three Cs of the same color from appearing together. However, it was possible for two neighboring Cs to be of the same color. For a given subject, the ignore color was set to be either red or green. This ignore color appeared as a distractor color in 80% of trials – the target never appeared in this color. The remaining color of those two was the attend color. This attend color appeared on 40% of trials as the target. Importantly, there were some trials were the ignore color appeared without the attend color (ignore trials, figure 17a), trials were the attend color appeared without the ignore color (attend trials, figure 17b), trials were both colors appeared together (attend+ignore trials, figure 17c), and trials were neither color appeared (neutral trials, figure 17d). Participants were not explicitly told of either contingency, with the expectation that they should implicitly learn to bias their attention away from the ignore color and implicitly learn to bias their attention towards the attend color. The remaining four neutral colors were equally likely to be the target on a given

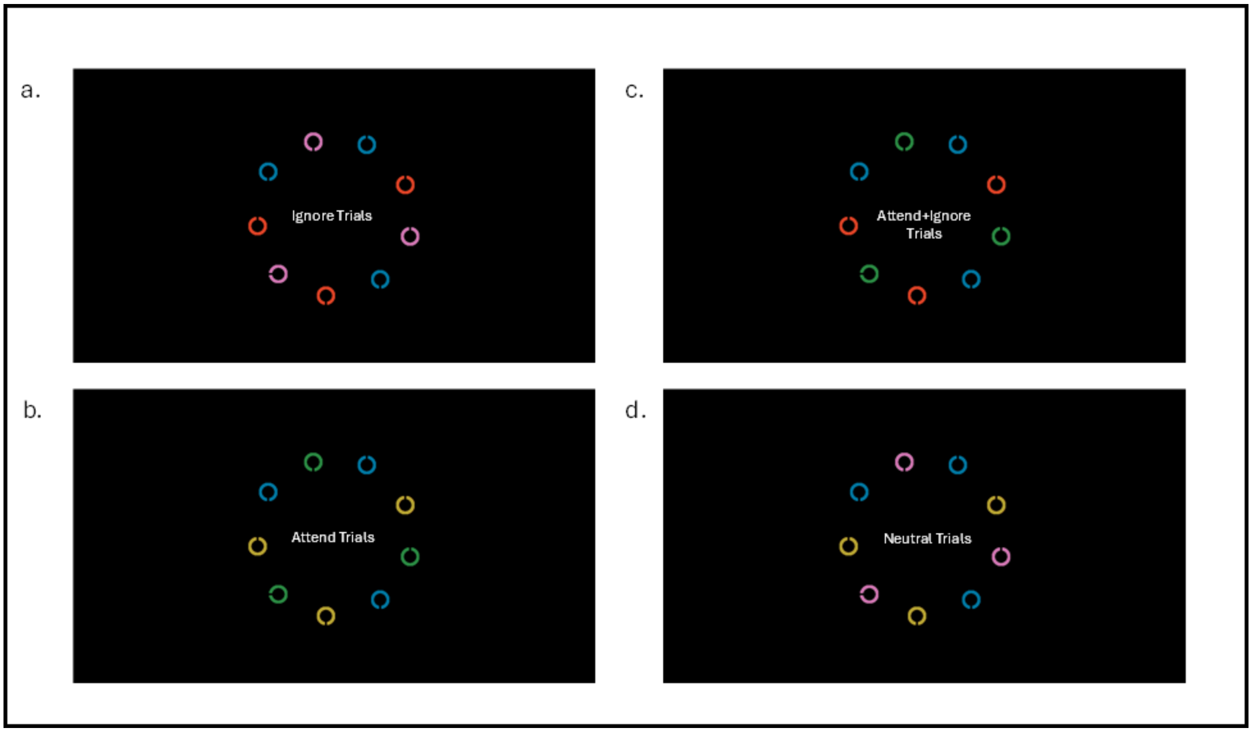

Figure 17. The four possible trial types for the training task. (a.) Example of an ignore trial, where the ignore color appears as a distractor without the attend color. (b.) Example of an attend trial, where the attend color appears as the target without the ignore color. (c.) Example of an attend+ignore trial, where both the attend and ignore colors appear together. (d.) Example of a neutral trial, where neither the attend nor ignore colors appear.

trial, each color appearing as the target on 15% of trials and each color appearing as a distractor on 30% of trials. Participants performed 250 trials of training split into two blocks of 125 trials, with a rest in between.

After the training task, participants completed a final test block consisting of 60 trials. In this block, the contingencies from the training task were removed, with each color (including the ignore and attend colors) appearing as the target for 10 trials and each color appearing as one of the two distractor colors for 20 trials. Thus, the ignore color now appeared as the target on some trials and the attend color now appeared as one of the two distractors on other trials. There were also trials where neither the ignore or attend colors appeared (neutral condition).

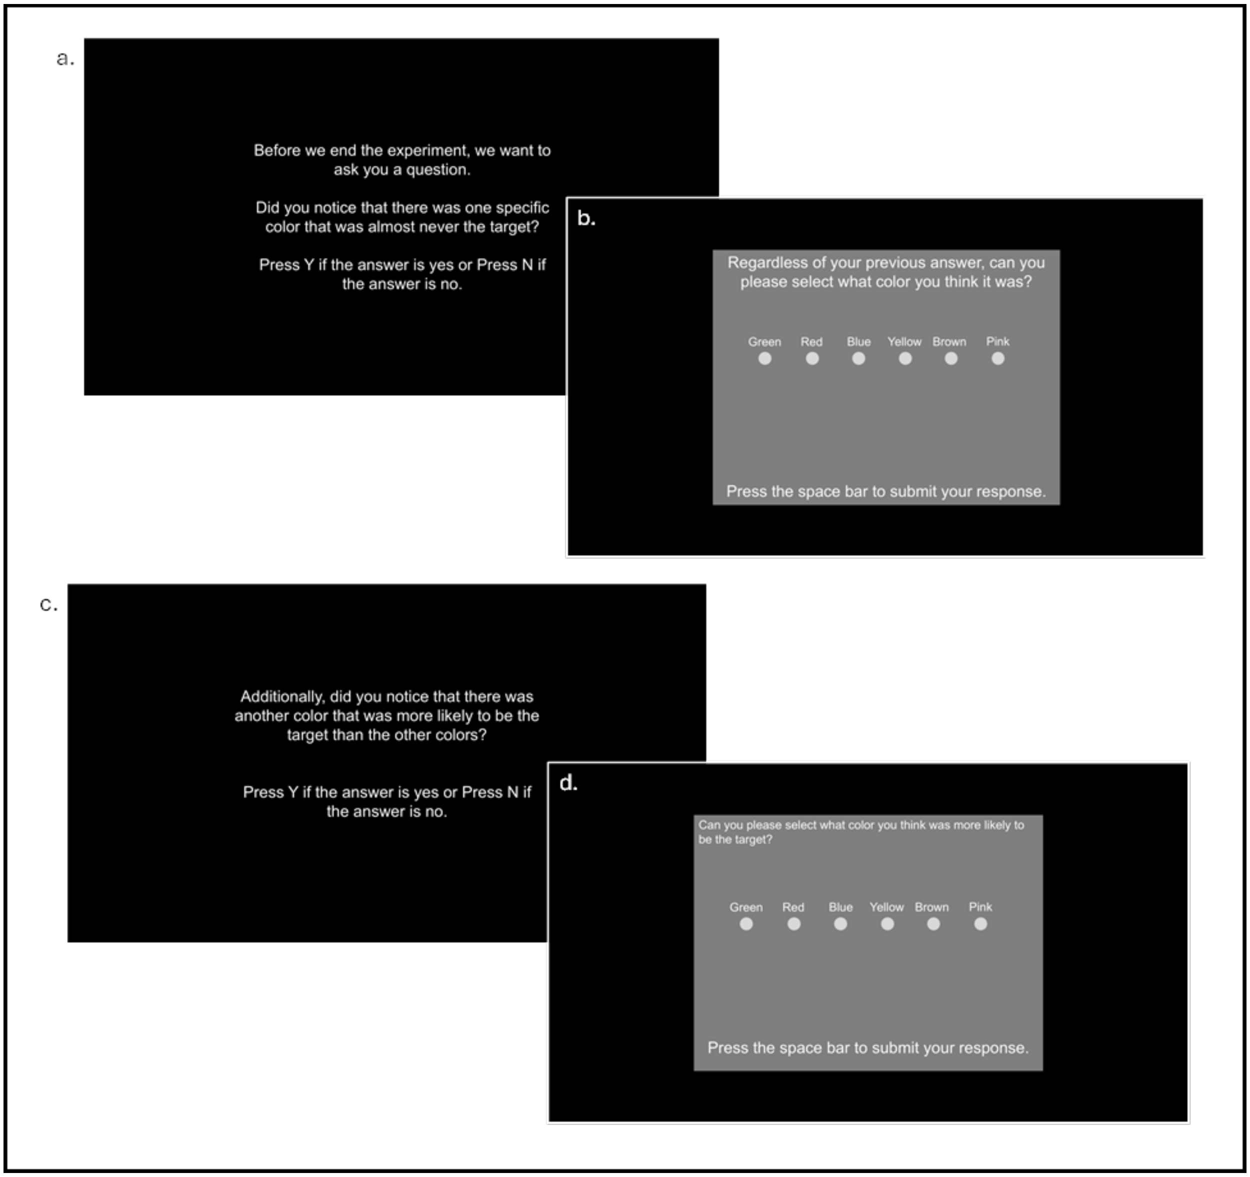

Figure 18. The end of experiment contingency checks. (a.) Participants are first asked if they had any awareness of one color that almost never was the target and answer via button press. (b.) Following the first question, they then had to choose which color was the ignore color via mouse selection. (c.) Participants are first asked if they had any awareness of one color that almost always the target and answer via button press. (b.) Following the third question, they then had to choose which color was the attend color via mouse selection.

After the experiment was completed, participants were assessed to see if they had any explicit awareness that the ignored color rarely was the target. They were first asked if they had any awareness (yes or no question) that one specific color rarely contained the target (figure 18a). Regardless of the answer to the first question, they were given a six alternative forced choice question in which they were asked to indicate which color was seldom the target (figure 18b). Furthermore, participants also were assessed to see if they had any explicit awareness of the attend color using the same two question procedure as described for the ignored color (figures 18c and 18d).

3.4 Results

3.4.1 Training Task

Trials with reaction times (RTs) more than 3 standard deviations away from the overall mean and with incorrect responses were excluded from the analyses, resulting in removal of .96% of total trials. A repeated measures ANOVA on RTs for the four conditions (attend, ignore, attend+ignore, and neutral conditions) revealed a significant difference (figure 19), $F(3, 189) = 7.99$, $p < .001$, partial eta squared = .112. Follow-up pairwise comparisons were performed to better understand the source of this difference.

The attend condition ($M = 1671.36$ ms, $SE = 39.09$ ms) had significantly faster RTs than the ignore condition ($M = 1744.27$ ms, $SE = 33.48$ ms), $t(63) = 2.93$, $p = .005$, and neutral condition ($M = 1778.69$ ms, $SE = 42.36$ ms), $t(63) = 3.39$, $p = .001$. The attend+ignore condition ($M = 1695.74$ ms, $SE = 37.62$ ms) also had significantly faster RTs than the ignore, $t(63) = 3.07$, $p = .003$, and neutral conditions, $t(63) = 3.56$, $p < .001$, but was not statistically different from the attend condition, $t(63) = 1.03$, $p = .309$. Thus, there was a benefit of feature-based gain, resulting in faster RTs in both the attend and attend+ignore conditions. However, we found no evidence of feature-based suppression – the attend+ignore condition was no faster than the attend alone condition and the ignore condition was no different than the neutral condition, $t(63) = 1.53$, $p = .131$. For completeness, a Bayesian paired samples t-test implemented in JASP (Version

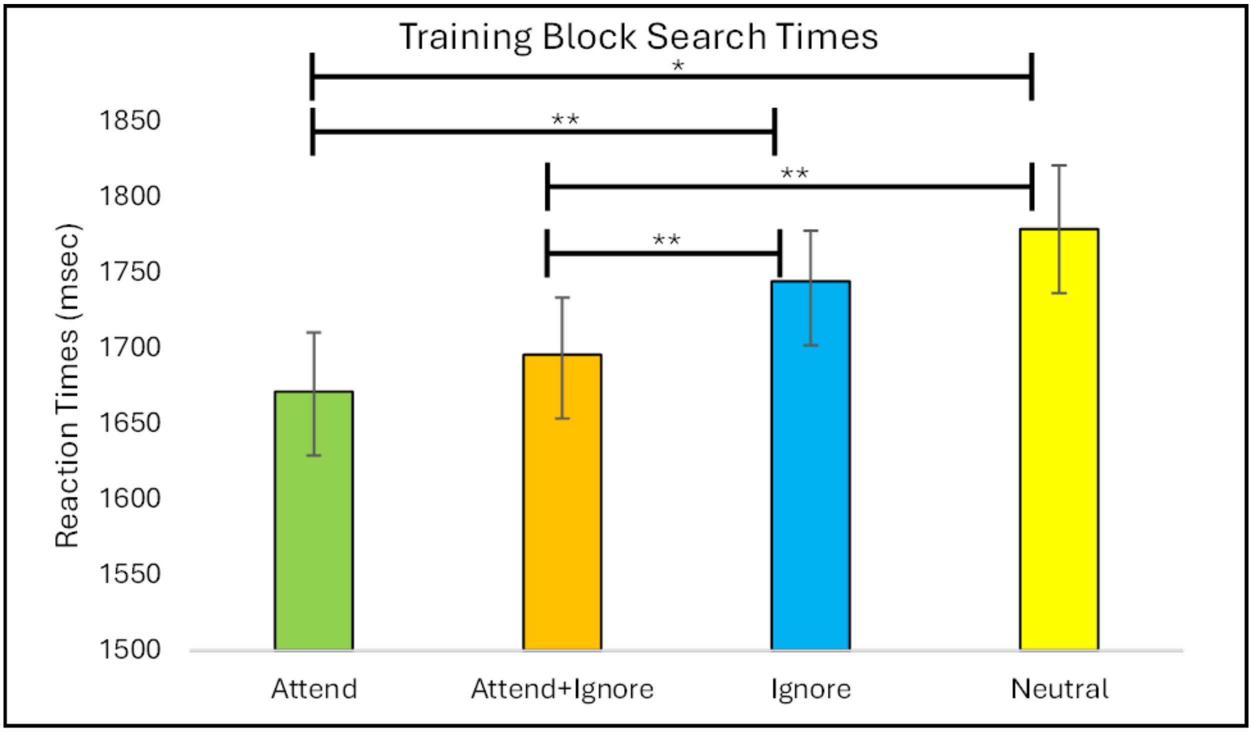

Figure 19. Search times for the training task by condition (* $p < .05$; ** $p < .005$).

0.19.1) compared the ignore and neutral conditions on RTs. This analysis provided moderate evidence for the null hypothesis with a Bayes Factor (BF_{10}) = .41.

Accuracy was generally high in the training block (M = 99.12%; SE = .15%), likely indicating that participants were performing at ceiling. Given that we found no differences in accuracy for Exp. 1, we wouldn't expect differences in accuracy between the different conditions. Nevertheless for completeness, comparisons on accuracy were performed. A repeated measures ANOVA on accuracy for the four conditions (attend, ignore, attend+ignore, and neutral conditions) revealed a significant difference (figure 20), $F(3, 189) = 3.56$, $p = .015$, partial eta squared = .083. Follow-up pairwise comparisons were performed to better understand the source of this difference.

Attend condition (*M* = 99.53%, *SE* = .21%) had significantly higher accuracy than the ignore (*M* = 98.96%, *SE* = .20%), $t(63) = 2.32$, $p = .012$, and neutral conditions (*M* = 98.80%, *SE*

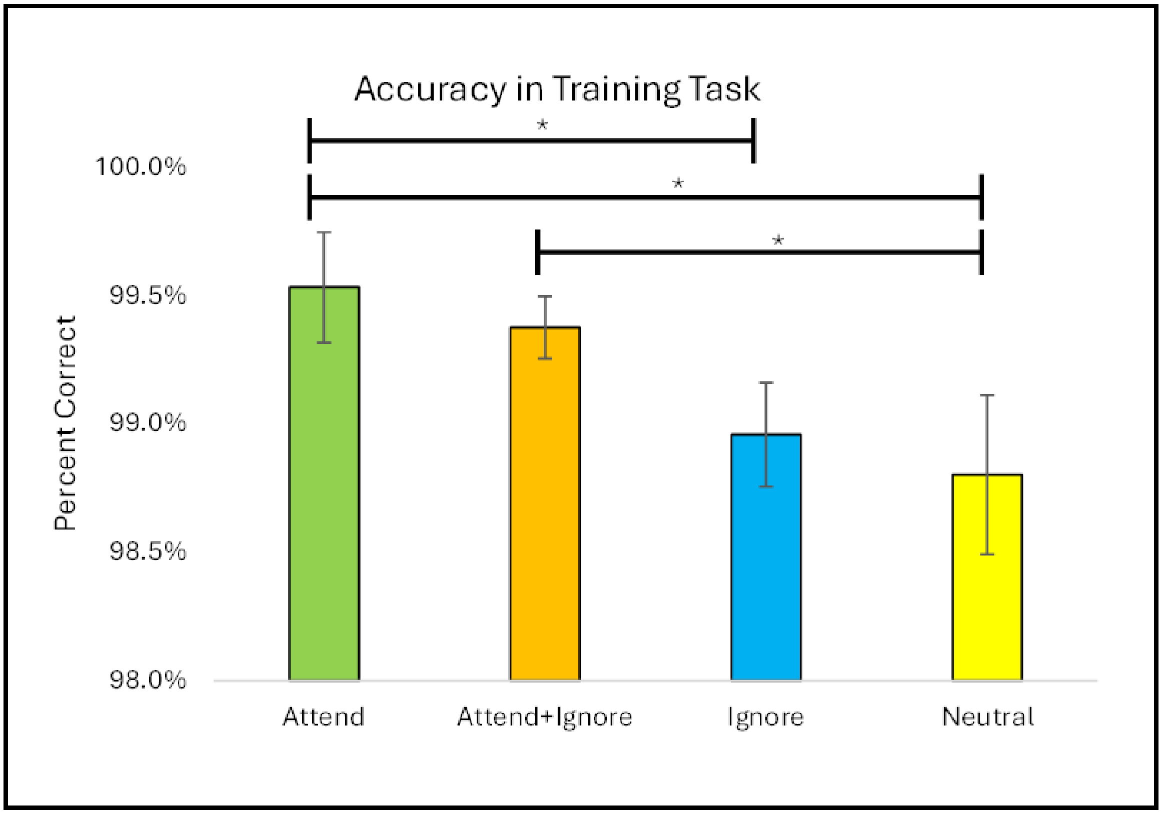

Figure 20. Accuracy for the training task by condition (* $p < .05$).

= .31%), $t(63) = 2.38$, $p = .010$. The attend+ignore condition ($M = 99.38\%$, $SE = .12\%$) also had significantly higher accuracy than the ignore, $t(63) = 2.34$, $p = .011$, and neutral conditions, $t(63) = 1.99$, $p = .025$, but not statistically different from the attend condition, $t(63) = .67$, $p = .253$. These results align with the RT data above, showing a benefit of feature-based gain, resulting in higher accuracy in both the attend and attend+ignore conditions, but no benefit of feature-based suppression as the ignore condition was no different than the neutral condition, $t(63) = .58$, $p = .282$. A Bayesian paired samples t-test was performed on the accuracy comparing the ignore and neutral conditions. This analysis provided moderate evidence for the null hypothesis with a $BF_{10} = .16$.

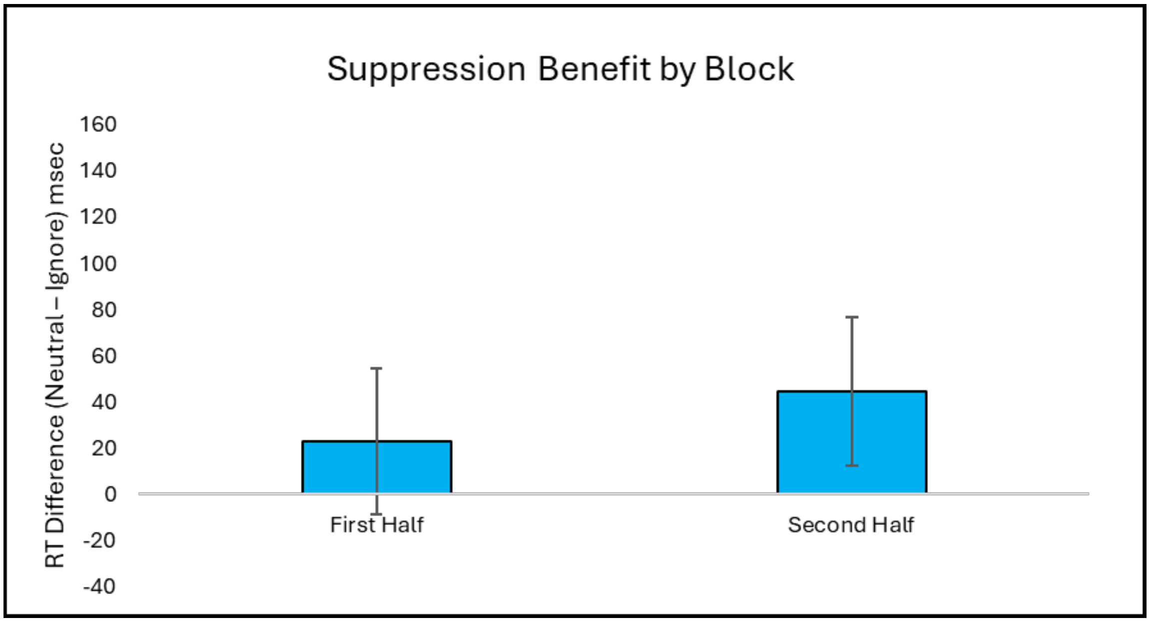

Figure 21. The suppression effect (subtracting the RTs in neutral trials from the RTs in ignore trials) in the training task split by half.

3.4.2 Time Course Analysis for Training Block

We conducted follow-up block analysis splitting the training block into halves to examine if there is a suppression benefit (difference in RTs by subtracting the ignore trials from the neutral trials) that develops in the second half of the training block after a prolong period of learning. While there was a numeric increase in suppression effect from the first half (M = 23.11 ms, SE = 31.64 ms) to the second half of the training block (M = 44.41 ms, SE = 32.07 ms) of the training block, this difference was not statistically significant, $t(63) = .48$, $p = .632$, indicating that the suppression benefit did not improve from training (figure 21). A Bayesian paired samples t-test confirmed that there is no improvement in suppression from the first half to the second half of the training task with a $BF_{10} = .15$. Furthermore, one-samples t-tests corroborate the lack of a suppression benefit, finding that the suppression effect in the first half, $t(63) = .73$, $p = .468$, and in the second half, $t(63) = 1.39$, $p = .171$, were not different than zero. Bayesian one-sample t-

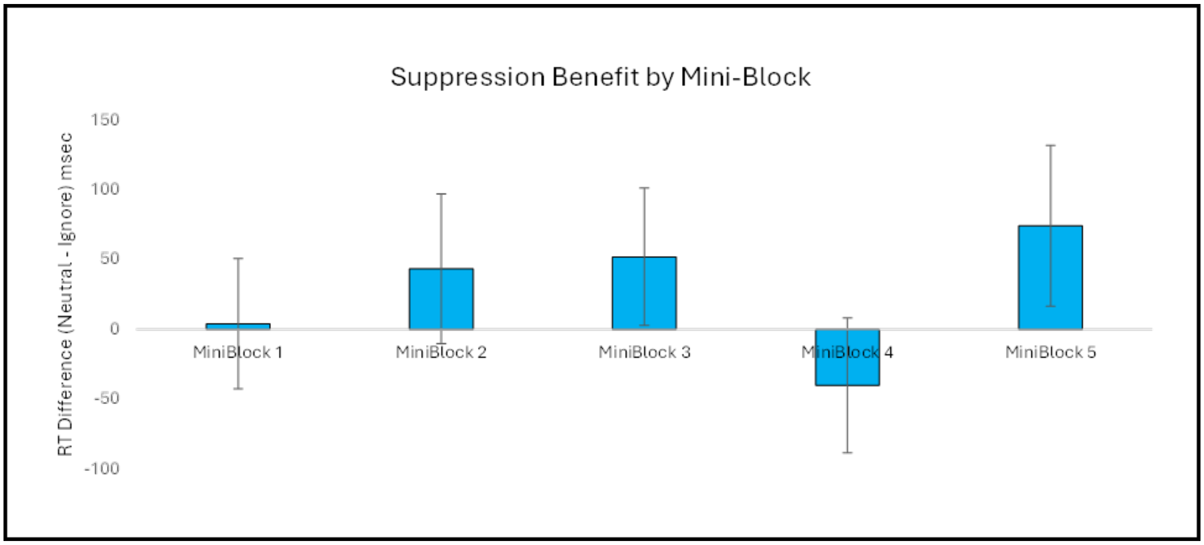

Figure 22. The suppression effect (subtracting the RTs in neutral trials from the RTs in ignore trials) in the training task split by mini-blocks.

tests also confirmed the lack of a suppression benefit, with BF_{10} = .17 and BF_{10} = .34 for the first and second halves respectively.

We also decided to investigate the suppression benefit by mini-blocks (each mini-block consisting of 50 trials) as it may reveal more nuanced learning that is lost by examining the suppression benefit by halves. However, this analysis still found no differences in suppression effect across the first (*M* = 4.03 ms, *SE* = 46.51 ms), second (*M* = 43.40 ms, *SE* = 53.62 ms), third (*M* = 51.91 ms, *SE* = 49.16 ms), fourth (*M* = -40.13 ms, *SE* = 48.26 ms), and fifth (*M* = 74.11 ms, *SE* = 57.64 ms) mini-blocks of the training block (figure 22), $F(4, 252) = .80$, $p = .530$, partial eta squared = .012. Thus, unlike Cunningham & Egeth (2016) which found an increase in suppression benefit over several mini-blocks, we found that the suppression benefit did not increase as the amount of training increased. A Bayesian repeated measures ANOVA confirmed that there are no differences in the suppression benefit among the different mini-blocks in the training task, finding strong evidence for the null hypothesis with a BF_M = .03.

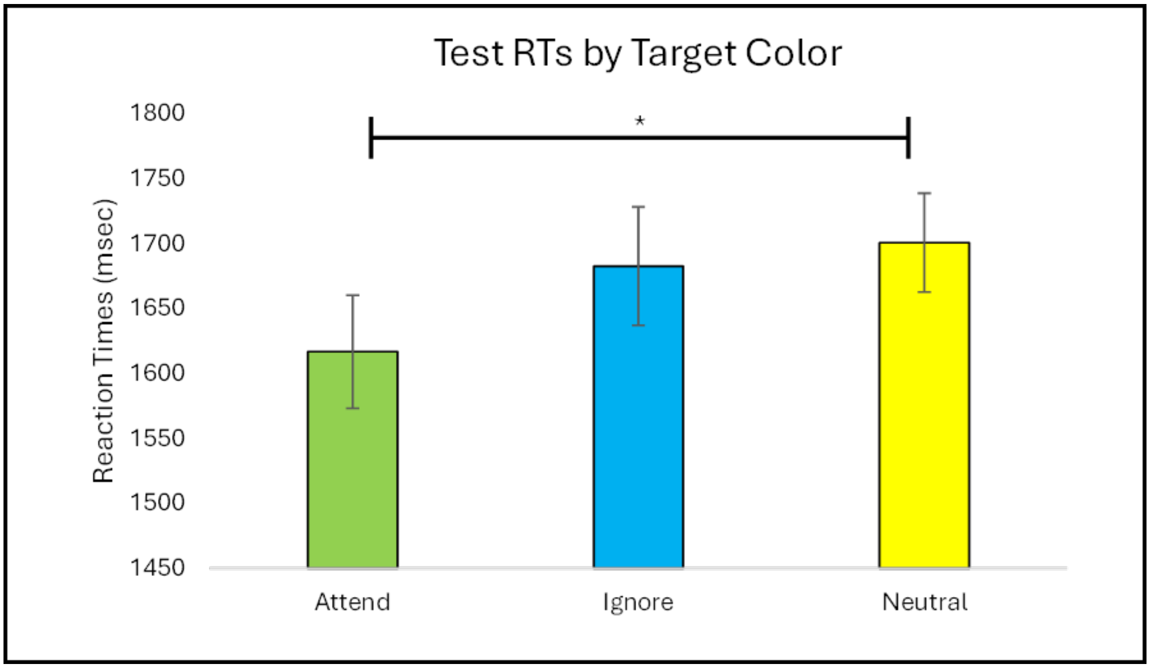

Figure 23. Search times for the test block by target color (* $p < .05$).

One-samples t-tests confirm the lack of a suppression benefit, finding that the suppression effect in the first, $t(63) = .09$, $p = .931$, second, $t(63) = .80$, $p = .421$, third, $t(63) = 1.06$, $p = .295$, fourth, $t(63) = -.83$, $p = .409$, and fifth, $t(63) = 1.29$, $p = .203$, mini-blocks were not different than zero. Bayesian one-sample t-tests reaffirmed the lack of a suppression benefit, with $BF_{10} = .14$, $BF_{10} = .19$, $BF_{10} = .23$, $BF_{10} = .19$, and $BF_{10} = .30$ for the first, second, third, fourth, and fifth mini-blocks respectively.

3.4.3 Test Block

To determine if the gain or suppression effects persist in the test block without the inter-trial priming from the training task, we conducted a repeated measures ANOVA on RTs for the three target conditions (attend, ignore, and neutral target colors). This analysis revealed that there is a marginal difference in RTs between the attend (*M* = 1616.88 ms, *SE* = 43.50 ms), ignore (*M* = 1682.73 ms, *SE* = 45.50 ms), and neutral (*M* = 1700.90 ms, *SE* = 38.03 ms) target trials (figure

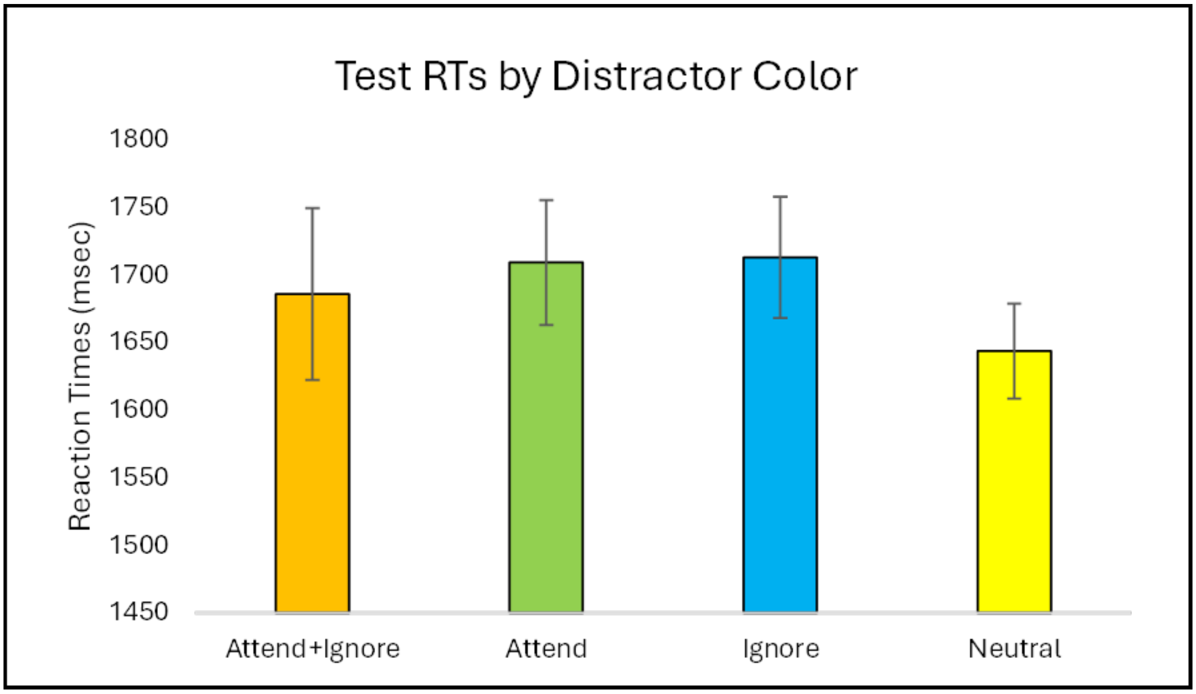

Figure 24. Search times for the test block by distractor color.

23), $F(2, 126) = 2.72$, $p = .070$, partial eta squared = .041. Follow-up pairwise comparisons were performed to investigate the source of this difference.

Attend target RTs were significantly faster than the neutral target, $t(63) = 2.42$, $p = .018$, but they were not statistically faster than the ignore target RTs, $t(63) = 1.57$, $p = .120$. Furthermore, the RTs for the ignore target condition did not statistically differ from the neutral target condition, $t(63) = .49$, $p = .623$, and a Bayesian paired samples t-test of this comparison found moderate evidence for the null with $BF_{10} = .15$. These results indicate the gain benefit from the training task carryover into the test block, while the suppression benefit still does not materialize in the test block.

Additionally, we conducted a repeated measures ANOVA on RTs for the four distractor color conditions (attend was a distractor, attend+ignore were distractors, ignore was a distractor, and neutral was a distractor). This analysis revealed no differences in RTs between the attend (*M*

= 1709.25 ms, *SE* = 46.23 ms), attend+ignore (*M* = 1685.92 ms, *SE* = 46.23 ms), ignore (*M* = 1713.05 ms, *SE* = 45.02 ms), and neutral (*M* = 1643.70 ms, *SE* = 35.30 ms) distractor conditions (figure 24), *F*(3, 189) = .83, *p* = .481, partial eta squared = .013.

Like the training task, accuracy was generally high in the test block (*M* = 98.70%; *SE* = .36%). Since we found evidence of gain benefit for accuracy in the training task, we also examined accuracy for the test block. A repeated measures ANOVA on accuracy for the target conditions showed no difference in accuracy between the attend (*M* = 98.91%, *SE* = .67%), ignore (*M* = 98.91%, *SE* = .39%), and neutral (*M* = 98.59%, *SE* = .36%) colors, *F*(2, 126) = .24, *p* = .784, partial eta squared = .004. A repeated measures ANOVA on accuracy for the distractor conditions also showed no difference in accuracy between the attend (*M* = 99.32%, *SE* = .34%), attend+ignore (*M* = 98.44%, *SE* = .76%), ignore (*M* = 98.24%, *SE* = .58%), and neutral (*M* = 98.63%, *SE* = .43%) colors, *F*(3, 189) = .88, *p* = .452, partial eta squared = .014. Overall, we found evidence of a gain benefit in the test block, but no evidence for a suppression effect.

3.4.4 Time Course Analysis for Test Block

While unlikely, there is a possibility that suppression benefit manifests in the test block but quickly disappears (as in Exp. 1). As such, we conducted the above analyses but only in the first half of the test block. A repeated measures ANOVA on RTs for the three target conditions (attend, ignore, and neutral target colors) revealed no significant differences (figure 25a), *F*(2, 126) = .33, *p* = .719, partial eta squared = .005. A repeated measures ANOVA on RTs for the four distractor color conditions (attend was distractor, attend+ignore were distractors, ignore was distractor, and neutral was distractor) also revealed no significant differences (figure 25b), *F*(3, 177) = 2.09, *p* = .103, partial eta squared = .034. Similar analyses were performed on accuracy. Both repeated measures ANOVAs for target, *F*(2, 126) = .16, *p* = .850, partial eta squared = .003,

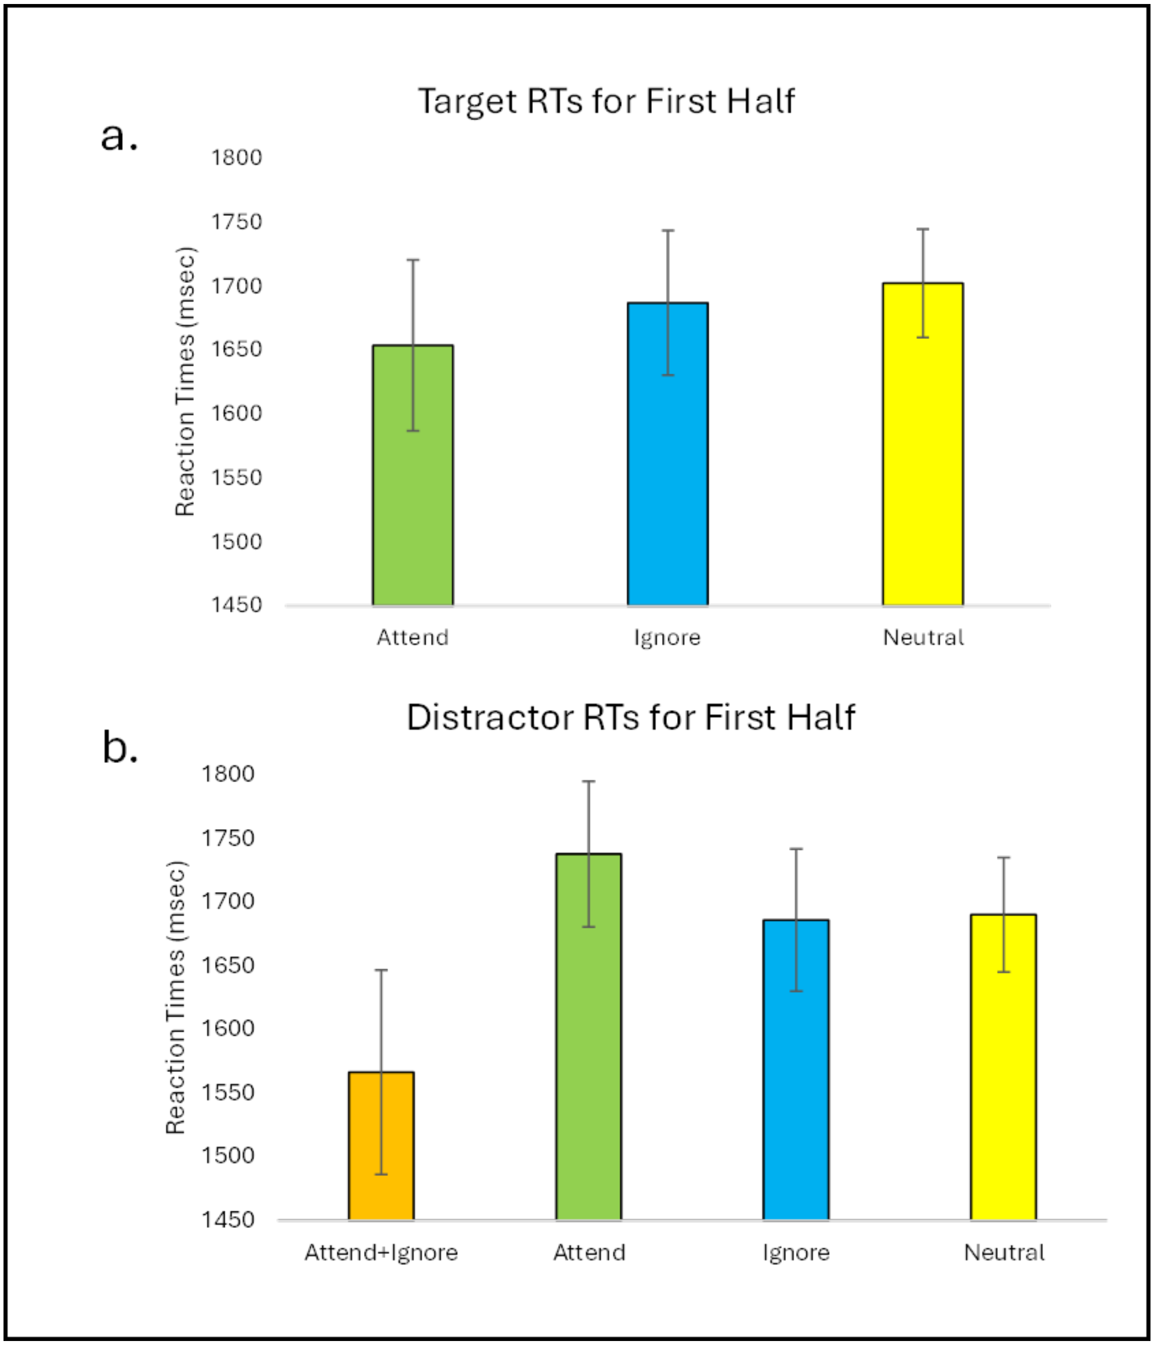

Figure 25. Search times for the first half of the test block. (a.) Search times by target color. (b.) Search times by distractor color.

and distractor conditions, $F(3, 180) = .19$, $p = .904$, partial eta squared = .003, found no differences in accuracy. Thus, there was no evidence of a suppression benefit in the test block.

3.4.5 Comparing the Suppression Effect between Exp.1 and 2

While it appears that the previous suppression effect observed both in the training and test blocks for Exp. 1 is nullified in Exp. 2 by including a feature-based gain component, we compared the magnitude of suppression effect between the two experiments to confirm this difference. To begin, we performed an independent samples t-test comparing the suppression effect (subtracting the RTs in neutral trials from the RTs in ignore trials) in training task for Exp. 1 and 2. Surprisingly, this analysis found no difference in suppression effect between Exp. 1 (M = 35.92 ms, SE = 17.50 ms) and Exp. 2 (M = 34.42 ms, SE = 22.48 ms), $t(120) = .05$, $p =.479$.

We also compared the difference in suppression effect (subtracting the RTs in ignore target trials from the RTs in ignore distractor trials) in test block for Exp. 1 and 2. Here we found a marginal difference in suppression effect, $t(120) = 1.35$, $p =.090$, showing a larger suppression benefit in Exp. 1 (M = 47.40 ms, SE = 38.08 ms) compared to Exp. 2 (M = -30.32 ms, SE = 42.75 ms). These results show some evidence that there is a stronger suppression effect in Exp. 1 that disappears in Exp. 2 with the inclusion of feature-based gain component.

3.5 Discussion

In an effort to create a task paradigm with more real-world validity, we used the same Landolt C search task from Exp 1. but added a feature-based gain component, as most real-world searches have well defined target that is the goal of the search. While we found a suppression benefit in Exp. 1, adding a FBG component to the same task paradigm not only created a gain benefit, with improvements in search times across the training task and the test block when the attend color appeared, but more importantly, resulted in no suppression benefit developing, as there was no improvement in search performance when the ignore color appeared as a distractor in either the training task or the test block. Furthermore, there was a lack of an additive effect when both the to-be-ignored and to-be-attended features appeared on the same trial. In fact, RTs and accuracy were numerically worse in the attend+ignore condition compared to the attend condition, indicating that, if anything, the presence of FBS component was actually making search performance more inefficient. While the above findings are somewhat surprising as we predicted that there would be a benefit of both gain and suppression, they aligned with previous studies (Kawashima & Amano, 2022; Kawashima & Matsumoto, 2018; Kugler et al., 2015; Rajsic, Carlisle, & Woodman, 2020) that found observers do not engage in FBS mechanisms when FBG mechanisms are available, potentially indicating the FBS is not an important mechanism for allocating visual attention.

However, those previous studies provided explicit gain and suppression mechanisms in the form of positive and negative cues while the current study surreptitiously included the gain and suppression components. This was evident as only 7.81% of our participants had explicit awareness of the to-be-ignored feature and 25.00% had explicit awareness of the to-be-attended feature. While the participants from the previous studies could consciously ignore the negative cue and choose to only use positive cue, resulting in a robust gain benefit and a lack of a suppression benefit, that explanation does not fit for the current study. In fact, the current study is more similar to Stilwell and Vecera (2020) as they implemented an implicit feature-based suppression and found a robust suppression benefit that was additive with the feature-based gain, making the lack of both a suppression benefit in our study even more surprising.

It is important to point out that Stilwell and Vecera (2020) provided participants with a 100% valid positive cue, making the FBG component in their study explicit rather than implicit as in our study. This difference in methodology could explain why the previously beneficial implicit FBS mechanism in Exp. 1 is no longer beneficial here in Exp. 2. It is possible that the resources and representation for implicit visual mechanisms are limited, and only one implicit mechanism can be encoded at a time. Given that FBG is a more robust mechanism than FBS (Arita, Carlisle, & Woodman, 2012; Beck & Hollingworth, 2012; Becker, Hemsteger, & Peltier, 2015), there is a possibility that participants were only encoding the implicit FBG mechanism and not the implicit FBS mechanism, explaining why there is gain benefit and a lack of a suppression effect.

As such, we decided to replicate the current task paradigm, but made the FBG component explicit via a positive cue that indicated the attend color. If this one change results in a suppression benefit, it would help elucidate the sensitive nature of implicit mechanisms, revealing that there are extremely limited resources for encoding them. Further, it would provide additional evidence that implicit FBS is an important mechanism for guiding attention as implicit FBS could be a complementary mechanism to explicit FBG, with the two co-operating to help allocate attention in the most efficient manner. However, if we do not find a suppression benefit after making the FBG component explicit, it weakens the argument that FBS is a main mechanism for guiding attention in everyday life.

CHAPTER 4:

STRONG BENEFIT OF IMPLICIT FEATURE-BASED SUPPRESSION IF THE FEATURE-BASED GAIN COMPONENT IS EXPLICIT

4.1 Abstract

Stilwell & Vecera (2020) suggested that observers can use both feature-based gain and feature-based suppression in conjunction to facilitate visual search. When we provided participants with both FBG and FBS mechanisms in Exp. 2, we found evidence that the gain mechanism improved search performance, but there was no evidence that suppression benefited search. However, differences in methodology could explain the incongruency of our findings, as Stilwell & Vecera (2020) used explicit FBG while we used implicit FBG in Exp. 2. Given that resources and representation for implicit visual mechanisms could be limited, it is possible that the lack of suppression benefit in our experiment was a product of employing two implicit visual mechanisms. As such, we replicated the Landolt C task from Exp. 2 but made the FBG component explicit. To do so, we explicitly cued participants during the training task on what color was associated with the target. The follow-up test block removed both gain and suppression contingencies to see if both persist without inter-trial priming. Our results showed a gain benefit in both the training task and test block. More importantly, we found a suppression benefit in training task, with fast searches occurring when the ignore color appeared as a distractor. This suppression effect carried over to the test block, with fast searches when the ignore color appeared as a distractor and slow searches when the ignore color appeared as the target. Our findings suggest that implicit FBS is beneficial to visual search even with the inclusion of a gain mechanism, showing that FBS may have real-world applicability and is potentially an important mechanism for allocating attention. Additionally, our results suggest that there is limited

representation for implicit mechanisms, and that only one mechanism can be encoded and used at a given time.

4.2 Introduction

The lack of suppression benefit in Exp. 2, when participants could engage in feature-based gain mechanisms, weakens the argument that feature-based suppression is an important mechanism for guiding attention. However, this result contradicts Stilwell & Vecera's (2020) finding that implicit FBS can benefit search even when provided with strong target guidance. A critical difference in methodology could explain why we did not replicate Stilwell & Vecera's suppression benefit, namely they explicitly cued the upcoming target's feature while Experiment 2 relied on implicit learning of the to-be-attended feature. Before concluding that FBS is not as beneficial as some studies (Arita, Carlisle, & Woodman, 2012; Chang & Egeth, 2019, 2021; Gaspelin, Leonard, & Luck, 2015; Moher, Lakshmanan, Egeth, & Ewen, 2014) suggest and is not an influential mechanism in allocating attention, we attempted to replicate Exp. 2 but with an explicit FBG cue. If we replicate the Experiment 2 finding of no evidence for FBS when gain is also possible, it would significantly diminish the importance of suppressive mechanism in the allocation of attention – since most times people have a sense of the features associated with the target they are looking for. By contrast, if in this experiment the suppression benefit reappears, the findings would suggest that implicit learning of suppression can occur in the face of top-down implementation of a FBG, and that the failure to find evidence for suppression in Experiment 2 might be due to the inability to implicitly learn to implement both a FBG and FBS mechanisms simultaneously.

In this experiment, participants were made explicitly aware that one color (attend color) would appear more frequently as the target. In the training task, a positive cue was shown on every trial, instructing participants to direct their attention to the attend color. However, participants were not aware that there was another color (ignore color) that appeared frequently and never as the target. This approach allowed us to determine whether participants learn to implicitly suppress the ignore color while implementing a top-down attentional bias towards the attend color. In the test block, the suppression and gain contingencies were removed so that both the ignore and attend colors appear at the same rate as the other colors and both were equally likely to appear as either the target or a distractor. Furthermore, the positive cue was also removed from the test block trials.

It is possible that changing the FBG component from an implicit mechanism (as in Experiment 2) to an explicit mechanism may reduce the demands on an implicit learning mechanism thereby allowing adequate resources for the implicit FBS to be learned and utilized. If this is true, we predict that in addition to a gain benefit improving search when the attend color appears as the target, that there will be a suppression benefit when the ignore color appears as a distractor. Alternatively, suppression may not occur in this experiment, which would provide evidence that suppression only occurs when there is no possibility of engaging a gain mechanism, which would severely limit the situations when suppression would likely influence the allocation of attention to rare cases where features of the target were completely unknown.

4.3 Methods

4.3.1 Participants

Sample size was estimated a priori based on Cohen's recommended small effect size of d = 0.2 (Cohen, 1992). Using G*Power version 3.1.9.7 with power = .95 and $\alpha = .05$ for repeated measures ANOVA with between-subjects interaction for the cross-experiment analyses, this gave a minimum sample size of 56.

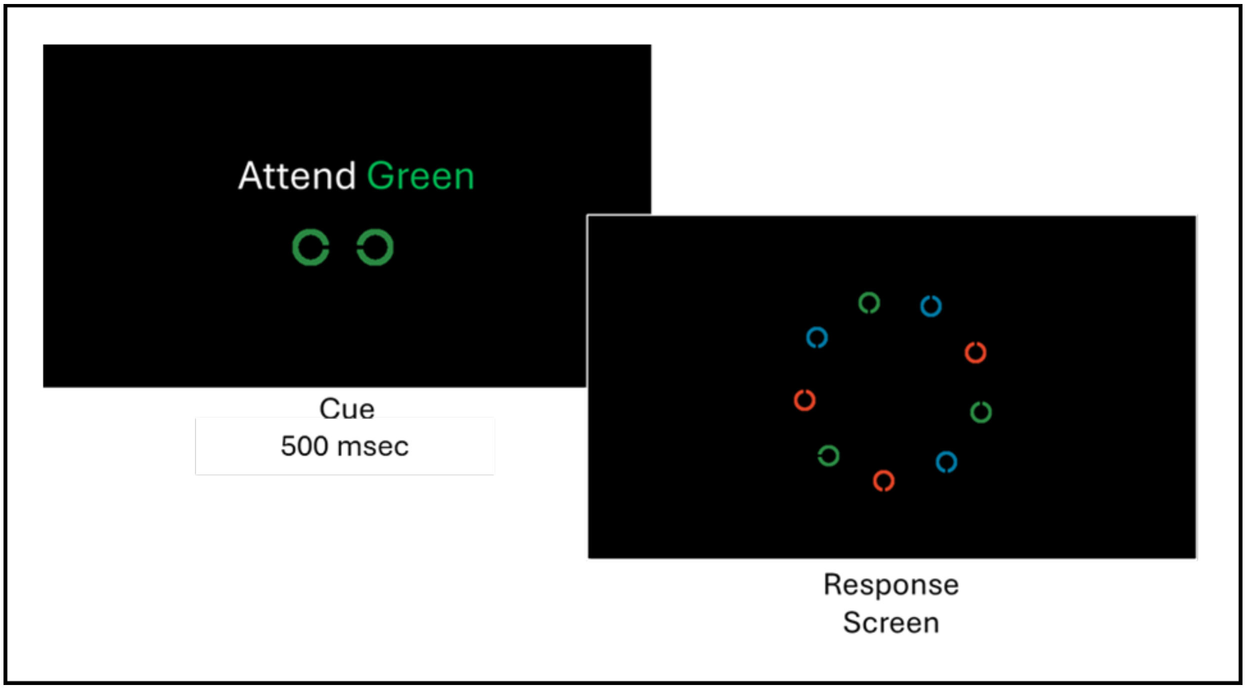

Figure 26. An example of a possible trial in the training task. Each trial begins with a positive cue appearing for 500 ms. The cue instructs which color to attend to and shows examples of what the target will look like. After the cue disappears, the search display appears, consisting of nine Landolt Cs, and stays on screen until a button press is made. Eight Cs are distractors with vertical gaps (either upwards or downward facing) and the target is the C with the horizontal gap (either rightward or leftward facing). Participants indicate the location of horizontal gap via button press. On a given trial, three colors appear, with three Cs in each color. The ignore color (the red Cs) appeared on 80% of trials and was never the target while the attend color (the green Cs) appeared on 40% of trials and always as the target.

As such, 74 undergraduate students from Michigan State University were recruited to participate in the experiment. Data from three participants were excluded from further analysis for having error rates greater than 3 SDs from the overall mean, leaving 71 participants whose data was analyzed. All the participants had normal or corrected-to-normal acuity. Participants gave written informed consent under the study protocol approved by the Institutional Review Board and were compensated with research credits.

4.3.2 Stimulus, apparatus, and procedure

All stimuli and procedures were identical to those in Exp. 2, with the following exceptions. Participants were explicitly instructed that one color (attend color) would appear on a large proportion of trials and when it appears the target would be in that color so they should actively attend to this color. However, participants were not made aware of the ignore color. In the training task, each trial began with the presentation of a positive cue for 500 ms followed by the search array (figure 26). This cue was 40% valid, matching the frequency the attend color appeared as the target. Participants performed 250 trials of the training task split into two blocks of 125 trials, with a rest in between.

For the test block, both the gain and suppression contingencies from the training task were removed. Additionally, the positive cue from the training task was also removed. Participants performed 60 trials of test block.

4.4 Results

4.4.1 Training Task

Trials with reaction times (RTs) more than 3 SDs away from the overall mean and with incorrect responses were excluded from the analyses, resulting in removal of 1.21% of total trials. Accuracy was generally high in the training block (*M* = 98.91%; *SE* = .13%), likely indicating that participants were performing at ceiling. A repeated measures ANOVA found no significant difference in accuracy between the attend (*M* = 98.73%, *SE* = .26%), ignore (*M* = 99.12%, SE = .14%), attend+ignore (*M* = 98.63%, *SE* = .20%), and neutral (*M* = 98.97%, *SE* = .21%) conditions, $F(3, 210) = 1.45$, $p = .229$, partial eta squared = .020.

A repeated measures ANOVA on RTs for the four conditions (attend, ignore, attend+ignore, and neutral conditions) revealed a significant difference (figure 27), $F(3, 210) = 155.36$, $p < .001$, partial eta squared = .689. Follow-up pairwise comparisons were performed to better understand the source of this difference. The attend condition (*M* = 1256.52 ms, *SE* = 39.76 ms) had significantly faster RTs than the ignore condition (*M* = 1681.39 ms, *SE* = 39.68

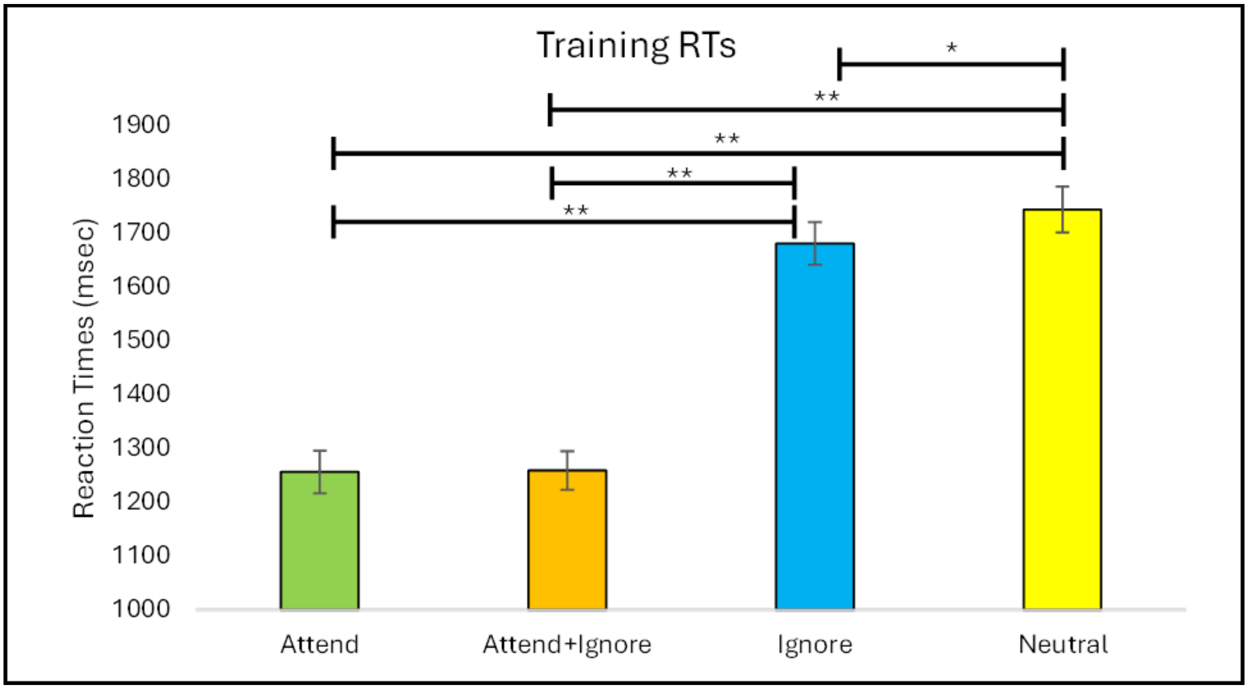

Figure 27. Search times for the training task by condition (* $p < .005$; ** $p < .001$).

ms), $t(70) = 13.69$, $p < .001$, and neutral condition ($M = 1744.68$ ms, $SE = 42.73$ ms), $t(63) = 13.25$, $p < .001$, indicative of feature-based gain benefit. Critically, the ignore condition also had significantly faster RTs than the neutral condition, $t(70) = 3.17$, $p = .002$, showing there is also a benefit of feature-based suppression in addition to the gain benefit. Lastly, the attend+ignore condition ($M = 1259.51$ ms, $SE = 35.80$ ms) also had significantly faster RTs than the ignore, $t(70) = 13.41$, $p < .001$, and neutral conditions, $t(70) = 13.09$, $p < .001$, but was not statistically different from the attend condition, $t(70) = .17$, $p = .863$, indicating a lack of an additive effect from the gain and suppression mechanisms.

4.4.2 Time Course Analysis for Training Task

We conducted a follow-up analysis splitting the training block into halves to examine if there is a suppression benefit (difference in RTs by subtracting the ignore trials from the neutral trials) that increases in the second half of the training block. A paired samples t-test revealed no difference in suppression benefit between the first half ($M = 67.09$ ms, $SE = 30.04$ ms) and

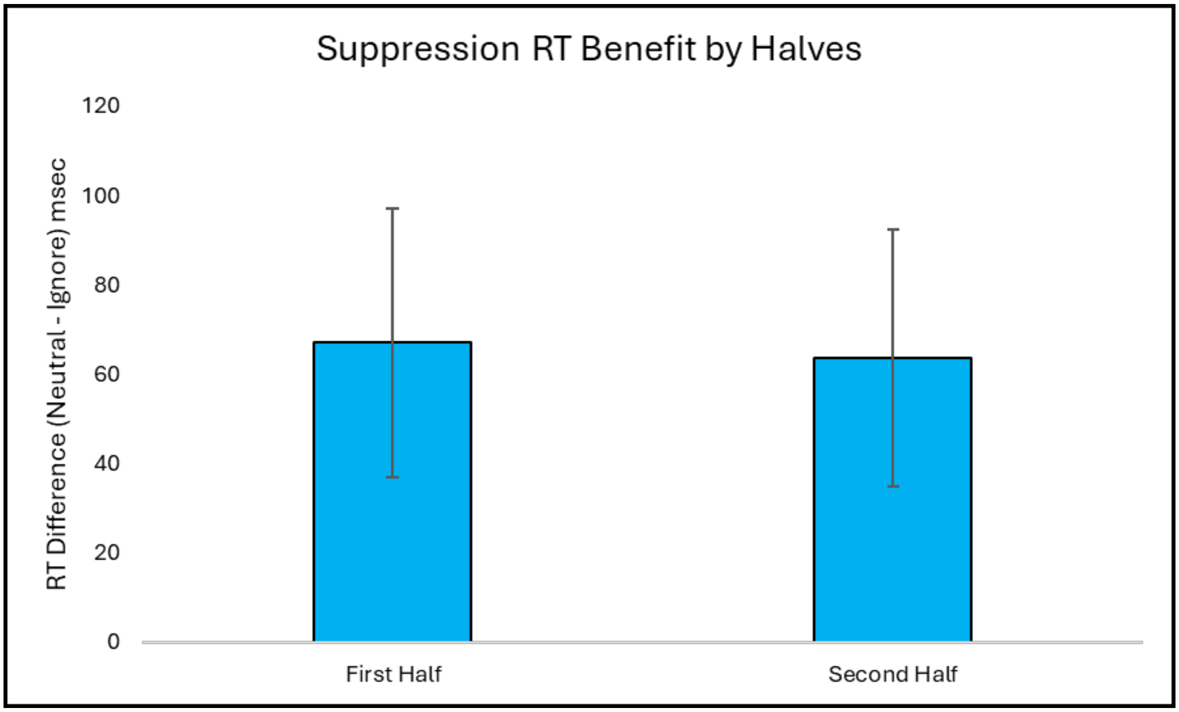

Figure 28. The suppression benefit (subtracting the RTs in ignore trials from the RTs in neutral trials) in the training task split by half.

second half (*M* = 63.62 ms, *SE* = 28.79 ms) of the training block, $t(70) = .08$, $p = .935$, indicating there is not a greater benefit of suppression from prolonged periods of learning (figure 28).

We also decided to investigate the suppression benefit by mini-blocks, each mini-block consisting of 50 trials. However, this analysis still found no differences in suppression benefit between the first (*M* = 66.03 ms, *SE* = 50.45 ms), second (*M* = 9.68 ms, *SE* = 57.59 ms), third (*M* = 80.41 ms, *SE* = 53.37 ms), fourth (*M* = 21.40 ms, *SE* = 61.81 ms), and fifth (*M* = 116.31 ms, *SE* = 48.71 ms) mini-blocks of the training block (figure 29), $F(4, 280) = .62$, $p = .645$, partial eta squared = .009. As we found in Exp. 1, the suppression effect is beneficial after a few trials and did not increase as the amount of training increased.

4.4.3 Test Block

To determine if the suppression effect persists in the test block without the inter-trial priming from suppression contingency in the training block, we conducted a repeated measures

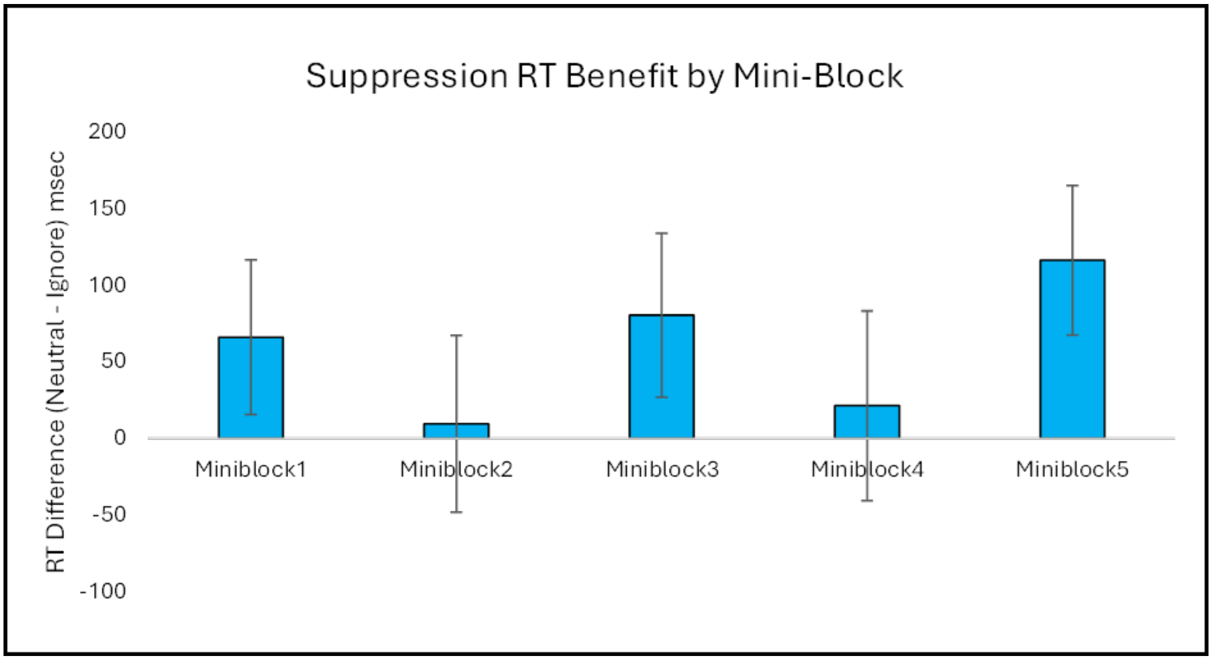

Figure 29. The suppression effect (subtracting the RTs in ignore trials from the RTs in neutral trials) in the training task split by mini-blocks.

ANOVA on RTs for the three target conditions (attend, ignore, and neutral target colors). This analysis revealed that there is a significant difference in RTs between the attend (M = 1460.95 ms, SE = 50.85 ms), ignore (M = 1833.34 ms, SE = 51.22 ms), and neutral (M = 1775.83 ms, SE = 46.06 ms) target trials (figure 30), F (2, 140) = 27.79, $p < .001$, partial eta squared = .284. Follow-up pairwise comparisons were performed to investigate the source of this difference.

Attend target RTs were significantly faster than the neutral target, t(70) = 6.47, $p < .001$, and the ignore target RTs, t(63) = 5.73, $p < .001$. While the RTs for the ignore target condition were numerically slower than the neutral target condition, they were not statistically different, t(70) = 1.26, p = .212. These results indicate the gain effect from the training task still persists into the test block, while the suppression effect may not materialize in the test block.

However, we also conducted a repeated measures ANOVA on RTs for the four distractor color conditions (attend was a distractor, attend+ignore were distractors, ignore was a distractor, and neutral was s distractor). This analysis found a statistical difference in RTs between the

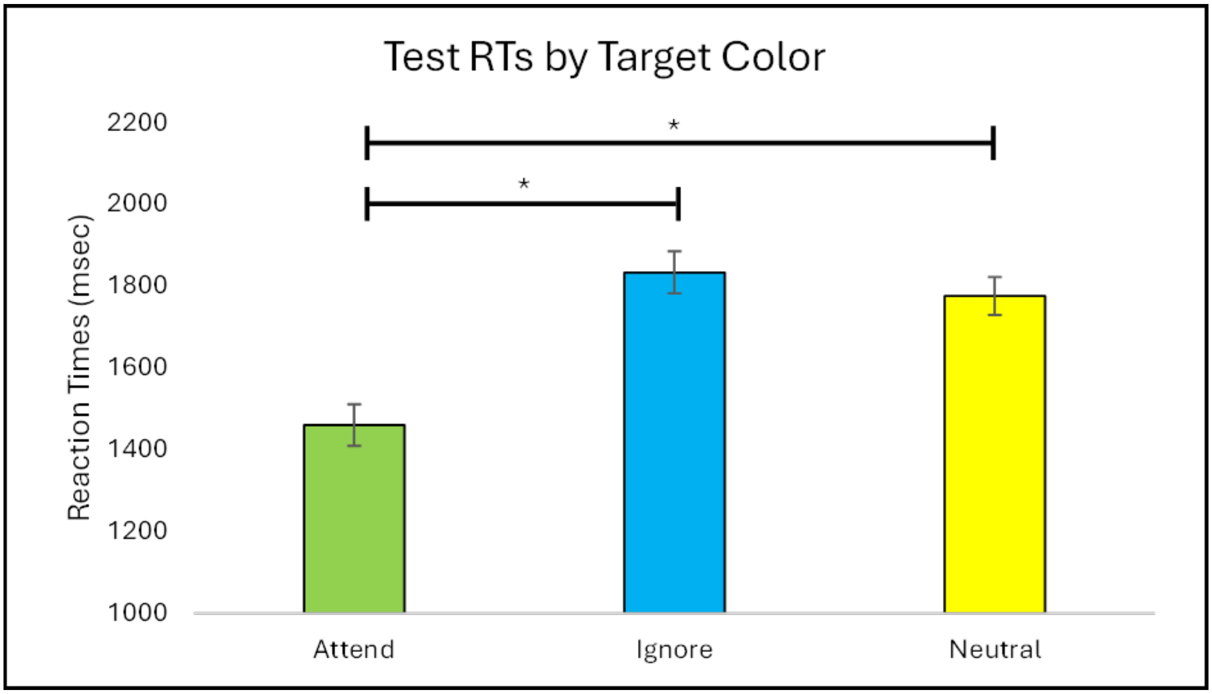

Figure 30. Search times for the test block by target color (* $p < .05$).

attend (M = 1945.33 ms, SE = 54.57 ms), attend+ignore (M = 1896.27 ms, SE = 65.84 ms), ignore (M = 1590.36 ms, SE = 42.72 ms), and neutral (M = 1660.57 ms, SE = 43.42 ms) distractor conditions (figure 31), $F(3, 210) = 23.20$, $p < .001$, partial eta squared = .249. Follow-up pairwise comparisons were performed to investigate the source of this difference.

Ignore distractor condition RTs were significantly faster than the neutral, $t(70) = 2.22$, $p = .03$, attend, $t(70) = 8.37$, $p < .001$, and attend+ignore, $t(70) = 4.82$, $p < .001$, conditions. While the suppression effect is not apparent when examining the RTs by target colors, there is a suppression benefit when it appears as a distractor, resulting in fastest RTs. The neutral condition RTs were significantly faster than attend, $t(70) = 7.31$, $p < .001$, and attend+ignore, $t(70) = 3.88$, $p < .001$, conditions. However, there were no differences in RTs between attend and attend+ignore conditions, $t(70) = .81$, $p = .419$.

Like the training task, accuracy was generally high in the test block (M = 98.26%; SE = .32%), but for completeness, we also examined accuracy by the target and distractor conditions.

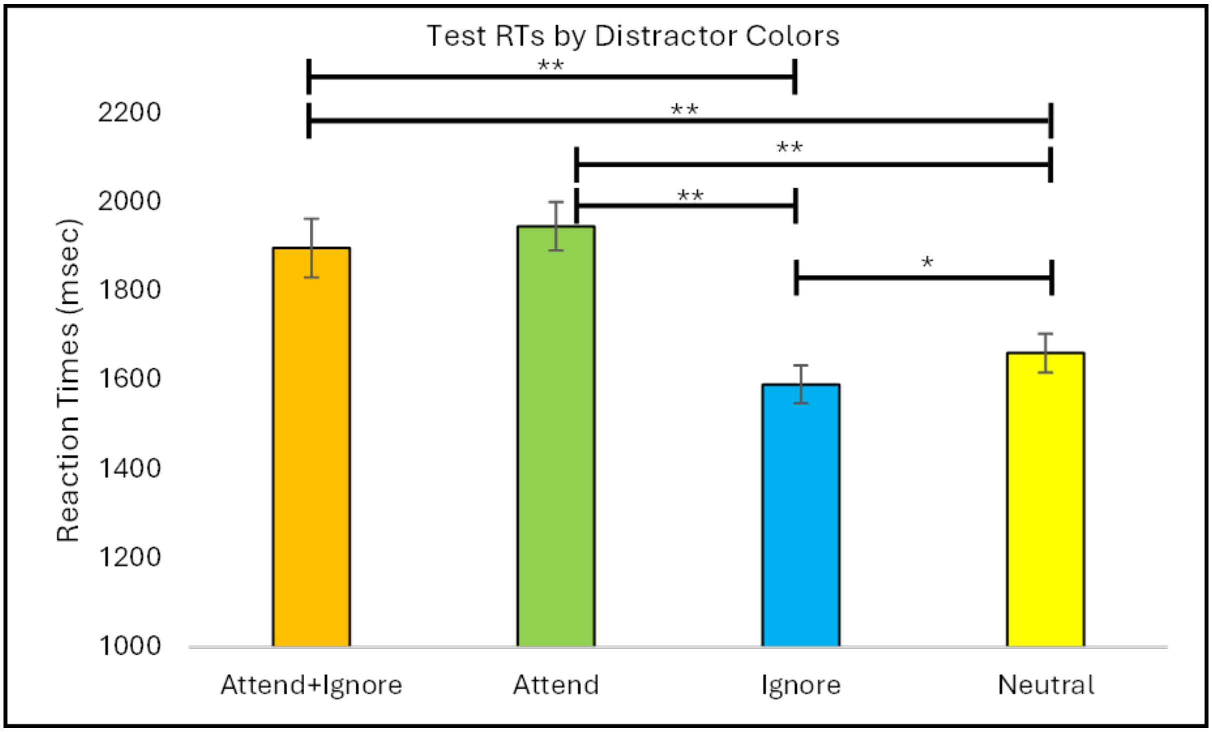

Figure 32. Search times for the first half of the test block by target color (* $p < .05$; ** $p < .001$).

A repeated measures ANOVA for the target conditions showed no difference in accuracy between the attend (*M* = 98.17%, *SE* = .46%), ignore (*M* = 98.17%, *SE* = .50%), and neutral (*M* = 98.31%, *SE* = .38%) colors, $F(2, 140) = .04$, $p = .960$, partial eta squared = .001. A repeated measures ANOVA on accuracy for the distractor conditions showed no difference in accuracy between the attend (*M* = 98.15%, *SE* = .47%), attend+ignore (*M* = 98.94%, *SE* = .60%), ignore (*M* = 98.24%, *SE* = .47%), and neutral (*M* = 98.24%, *SE* = .34%) colors, $F(3, 210) = .88$, $p = .454$, partial eta squared = .012. Overall, we find evidence of a gain effect in RT data when the attend color is both the target, benefiting RTs, and the distractor, resulting in a cost for RTs. The suppression effect is only evident when the ignore color is the distractor, also benefiting RTs, but not when it appears as a target, where we would expect it to slow RTs. However, there could be a suppression effect when the ignore color appears as the target, but it quickly disappears. As such,

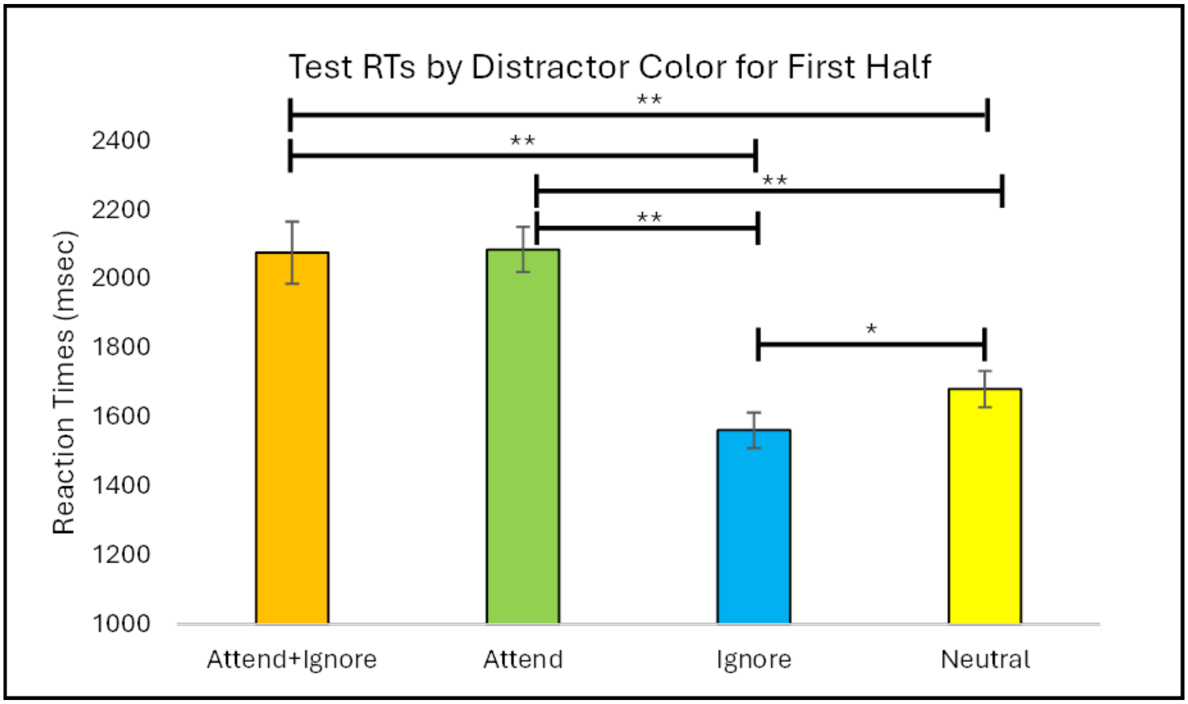

Figure 33. Search times for the first half of the test block by distractor color (* $p < .05$; ** $p < .001$).

we also performed time course analyses for the test block, looking at the first half of the test block.

4.4.4 Time Course Analysis for Test Block

We repeated the above RTs analyses but only in the first half of the test block as that is where the suppression effect should be the strongest. A repeated measures ANOVA revealed a significant difference in RTs for the three target conditions (attend, ignore, and neutral target colors; figure 32), $F(2, 140) = 26.03$, $p < .001$, partial eta squared = .271. Follow-up pairwise comparisons revealed that the attend target (M = 1467.60 ms, SE = 66.82 ms) condition had significantly faster RTs compared to the neutral (M = 1827.02 ms, SE = 50.89 ms), $t(70) = 5.83$, $p < .001$, and ignore target (M = 1971.84 ms, SE = 68.85 ms), $t(70) = 6.06$, $p < .001$, conditions. Importantly, the ignore target condition had significantly slower RTs to the neutral conditions,

$t(70) = 2.09$, $p = .041$. Thus, we find a suppression effect in the first half of the test block when the target appears in the ignore color that disappears over time.

A repeated measures ANOVA for the four distractor color conditions also revealed a significant difference in RTs (figure 33), $F(3, 198) = 21.69$, $p < .001$, partial eta squared = .247. Like the omnibus analysis, follow-up pairwise analyses found that the ignore distractor condition RTs were significantly faster than the neutral, $t(70) = 2.31$, $p = .024$, attend, $t(70) = 9.15$, $p < .001$, and attend+ignore, $t(66) = 5.01$, $p < .001$, conditions, providing evidence for a suppressive effect. In addition, the neutral condition RTs were significantly faster than attend, t(70) = 6.68, $p < .001$, and attend+ignore, $t(66) = 3.91$, $p < .001$, conditions, providing evidence for a gain effect. However, there were no differences in RTs between attend and attend+ignore conditions, $t(66) = .20$, $p = .841$. Overall, we find both a gain and suppression benefit in the training and test blocks.

4.4.5 Comparing the Suppression Effect between Experiments 1, 2, and 3

While it appears that by making the feature-based gain component explicit returned the suppression benefit, we compared the magnitude of suppression benefits across Exp. 1, 2, and 3 to confirm this result. To begin, we performed an One-Way ANOVA comparing the suppression benefit (subtracting the RTs in neutral trials from the RTs in ignore trials) in training task for Exp. 1, 2, and 3. While the suppression benefit in Exp. 3 ($M = 63.29$ ms, $SE = 19.97$ ms) is numerically larger than the suppression benefit in both Exp. 1 ($M = 35.92$ ms, $SE = 17.50$ ms) and Exp. 2 ($M = 34.42$ ms, $SE = 22.48$ ms), this analysis found no difference in suppression benefit, $F(2, 190) = .67$, $p = .517$.

We also compared the difference in suppression benefit (subtracting the RTs in ignore target trials from the RTs in ignore distractor trials) in test block for Exp. 1, 2, and 3. Here we found a significant difference in suppression effect, $F(2, 190) = 10.78$, $p < .001$. Follow-up comparisons confirm a stronger suppression benefit in Exp. 3 (M = 242.99 ms, SE = 52.28 ms) compared to Exp. 2 (M = -30.32 ms, SE = 42.75 ms), $t(133) = 4.00$, $p < .001$, and Exp. 1 (M = 47.40 ms, SE = 38.08 ms), $t(127) = 2.91$, $p = .002$. These results confirm that making the feature-based gain component explicit returned the suppression benefit in Exp. 3 and showed a larger suppression benefit in Exp. 3 compared to either Exp. 1 and 2.

4.5 Discussion

When an implicit feature-based gain component was added to the Landolt C search task in Exp. 2, there was a gain effect, but a lack of a suppression effect. While this could indicate that feature-based suppression has limited real-world applicability in guiding attention, the lack of a suppression effect also could have resulted due to limitations in the implicit visual mechanism's ability to simultaneous learn both a gain and suppression contingency. Since FBG is a more straightforward and robust mechanism (Arita, Carlisle, & Woodman, 2012; Beck & Hollingworth, 2012; Beck, Luck, & Hollingworth, 2018; Becker, Hemsteger, & Peltier, 2015; Olivers, Meijer, & Theeuwes, 2006; van Moorselaar, Theeuwes, & Olivers, 2014), there was a possibility that participants were only encoding the implicit FBG mechanism and not the implicit FBS mechanism.

Here, we replicated the task paradigm from Exp. 2 but made the feature-based gain component explicit via a positive cue, essentially moving it from the implicit to explicit domain. We reproduced the gain benefit from Exp. 2, with faster search times when the attend color appeared as the target in both the training task and the test block. More critically, there was now evidence of a suppression benefit, with faster search times when the ignore color appeared as a distractor in both the training task and the test block. Thus, moving the gain mechanism from the implicit to explicit domain seems to have allowed for the learning and implementation of the implicit suppression mechanism. Furthermore, only 5.63% of participants had conscious knowledge of the to-be-ignored feature, confirming that the feature-based suppression component remained in the implicit domain.

The findings here mirror those in Stilwell & Vecera (2020) showing a suppression benefit via implicit learning even when providing a 100% valid positive cue. Unlike Stilwell & Vecera (2020), we did not observe an additive effect of the gain and suppression mechanism as search performance across the attend and attend+ignore conditions were similar, indicating that the addition of the suppression component did not improve search performance beyond the positive cue. It is worth noting that Stilwell & Vecera used a two-color search array instead of the three-color array used here. This difference could explain the lack of an additive effect in the current study. While we equalized the utility of the gain and suppression components by having the ignore color appear twice as often as the attend color (80% to 40%) to compensate that the attend color reduces the search array by six items compared to the three item reduction from the ignore color, it is possible that the ignore color did not provide enough utility on the trials when they both colors appeared together. Perhaps if we used a two-color search array, we could find an additive effect of the attend and ignore colors.

Importantly, the results across Exps. 1, 2, and 3 indicate that the resources available for encoding implicit mechanisms are severely limited and only one implicit mechanism can be represented at a given time. In each experiment, participants could only learn and implement one implicit feature mechanism, and when two implicit feature mechanisms were provided in Exp. 2, they only encoded one (implicit FBG) and not the other (implicit FBS). Only by making the FBG explicit could participants use the implicit FBS mechanism.

Research on implicit mechanism representation is limited, requiring more work to better characterize how this limitation extends to mechanisms and cognitive processes outside of feature-based attention. For example, it is not clear if this limitation would extend to spatial-based attention. If observers are provided with both implicit spatial-based gain and implicit spatial-based suppression mechanisms, would it also restrict the ability to encode both mechanisms? Furthermore, would encoding an implicit spatial mechanism prevent or limit the encoding of an implicit feature mechanism or are there separate implicit resources for feature- and spatial-based attention? Future studies should address these questions.

Nevertheless, we found a suppression benefit even with the inclusion of a gain mechanism, showing that FBS has real-world applicability and is potentially an important mechanism for allocating attention. However, there is one *caveat* to our findings: across the three experiments, we have examined the limitations of *implicit* feature-based suppression. However, there are several previous studies that indicate that explicit feature-based suppression (for example, providing a negative cue that indicates the color that the target will not appear in) is weaker than implicit FBS (Addleman & Stormer, 2023; Lien, Ruthruff, & Tolomeo, 2024; Stilwell & Vecera, 2019) and takes time to overcome the initial bias to attend to the to-be-ignored feature (Cunningham & Egeth, 2016; Moher & Egeth, 2012; Vatterott & Vecera, 2012). To determine the importance of feature-based suppression in allocating attention, it is important to characterize the limitations of explicit FBS.

To do so, we decided to replicate Exp. 1 (FBS without an FBG component), but make the FBS component explicit via a negative cue. If we find a suppression benefit, we can compare the magnitudes of the implicit and explicit suppression benefit and examine if implicit FBS is indeed a more robust mechanism than explicit FBS. Additionally, we can conduct further follow-up studies to see if the explicit suppression benefit persists if there is feature-based gain component. However, if there is no benefit of explicit FBS, then it indicates the overall limitations of feature-based suppression and that it is only beneficial when implicitly learned, a stark contrast to feature-based gain that is beneficial if it is explicitly cued or implicitly learned as observed in both Exps. 2 and 3.

CHAPTER 5:

WEAK EFFECT OF FEATURE-BASED SUPPRESSION WHEN EXPLICITLY CUED

5.1 Abstract

Previous studies have found that explicit feature-based suppression, the active biasing of attention away from a feature, is a weaker mechanism compared to implicit feature-based suppression, the unconscious ignoring of a feature that appears as a distractor with high regularity. While we have found evidence that feature-based suppression can improve search performance, those experiments involved implicit learning of suppression based on historical context, leaving uncertain whether suppression can still be beneficial if it involves an explicit, top-down attentional mechanism. The goal of this study is to see if explicit suppression can benefit search performance, and to ultimately address the significance of feature-based suppression as real-world mechanism for guiding attention. We replicated the Landolt C search task from Exp. 1 but made the FBS component explicit. To achieve this, during the training task we explicitly cued participants that the target would not appear in a specific color. We found very little evidence of explicit suppression benefiting search performance. While the explicit cue produced a small increase in accuracy in the training task, there was no reaction time benefit, and the accuracy effect did not maintain into a test block where the suppression contingency is removed. While the results of our previous experiments found a robust benefit of implicit suppression that persisted even with a feature-based gain mechanism, there is weak to no evidence of a benefit for explicit suppression, even without a FBG component present. Our findings suggest that feature-based suppression is useful for guiding attention when implicitly learned, but the same may not be true when suppression is actively implemented, limiting the utility of FBS as a mechanism for allocating visual attention.

5.2 Introduction

Across Experiments 1 – 3, there was evidence of feature-based suppression improving search performance, even when participants had the ability to engage in feature-based gain. However, all three experiments examined the utility of *implicit* FBS, which has been suggested to be a stronger mechanism to explicit FBS (Addleman & Stormer, 2023; Lien, Ruthruff, & Tolomeo, 2024; Stilwell & Vecera, 2019). To better understand the importance of FBS in allocating attention in everyday life, we also investigated the effectiveness of *explicit* FBS.

Previous research has shown that explicit negative cues (cues that indicate the feature in which the target will not appear) initially bias attention towards the to-be-ignored feature before rapidly disengaging their attention away from the to-be-ignored feature, resulting in "search & destroy" search mode (Moher & Egeth, 2012). When observers engage in "search & destroy," search performance is hindered by to-be-ignored feature, leading to slower search times when the to-be-ignored feature is present compared to when it is absent. Only after prolonged exposure and training to the same to-be-ignored feature do participants learn to suppress it, resulting in faster search times when the to-be-ignored feature is present in the search array (Cunningham & Egeth, 2016; Vatterott & Vecera, 2012). This contrasts with implicit FBS, which requires minimal training to benefit search when the to-be-ignored feature is present in the search array (Golan & Lamy, 2022). Furthermore, some studies find either minimal or no improvement in search performance from the explicit cuing of the to-be-ignored feature (Beck, Luck, & Hollingworth, 2015; Becker, Hemsteger, & Peltier, 2015). Given the inconsistency of results concerning the effectiveness of explicit FBS in guiding attention, it is important to address this concern to better characterize the overall importance of FBS, regardless of whether it is explicit or implicit

To address the effectiveness of explicit FBS, in Experiment 4 we replicated the Landolt C search task, a task that produced a robust suppression benefit with implicit learning, but here we made the FBS component explicit via a negative cue that instructed participants to suppress the to-be-ignored feature. There were three possible patterns of results. One, an explicit suppression cue might produce a benefit with a similar magnitude as implicit suppression. Two, the explicit suppression cue might produce a benefit, but the benefit may be of a smaller magnitude than the implicit suppression benefit found in Exp. 1. Three, the explicit suppression cue might show no reliable benefit.

Given that prior evidence of an explicit suppression benefit tends to be weaker than implicit suppression (Addleman & Stormer, 2023; Lien, Ruthruff, & Tolomeo, 2024; Stilwell & Vecera, 2019), we predicted that there would be an explicit suppression benefit when the to-be-ignored feature is present, but the effect would be weaker than the implicit suppression benefit we found in Exp. 1. However, if we find no benefit here, then it limits the utility and benefit of FBS as it can only improve search efficiency when it is implicitly learned.

5.3 Methods

5.3.1 Participants

Sample size was estimated a priori based on Cohen's recommended small effect size of d = 0.2 (Cohen, 1992). Using G*Power version 3.1.9.7 with power = .95 and $\alpha = .05$ for repeated measures ANOVA with between-subjects interaction for the cross-experiment analyses, this gave a minimum sample size of 56.

As such, 63 undergraduate students from Michigan State University were recruited to participate in the experiment. Data from one participant was excluded from further analysis for having error rates greater than 3 SD from the overall mean, leaving 62 participants whose data

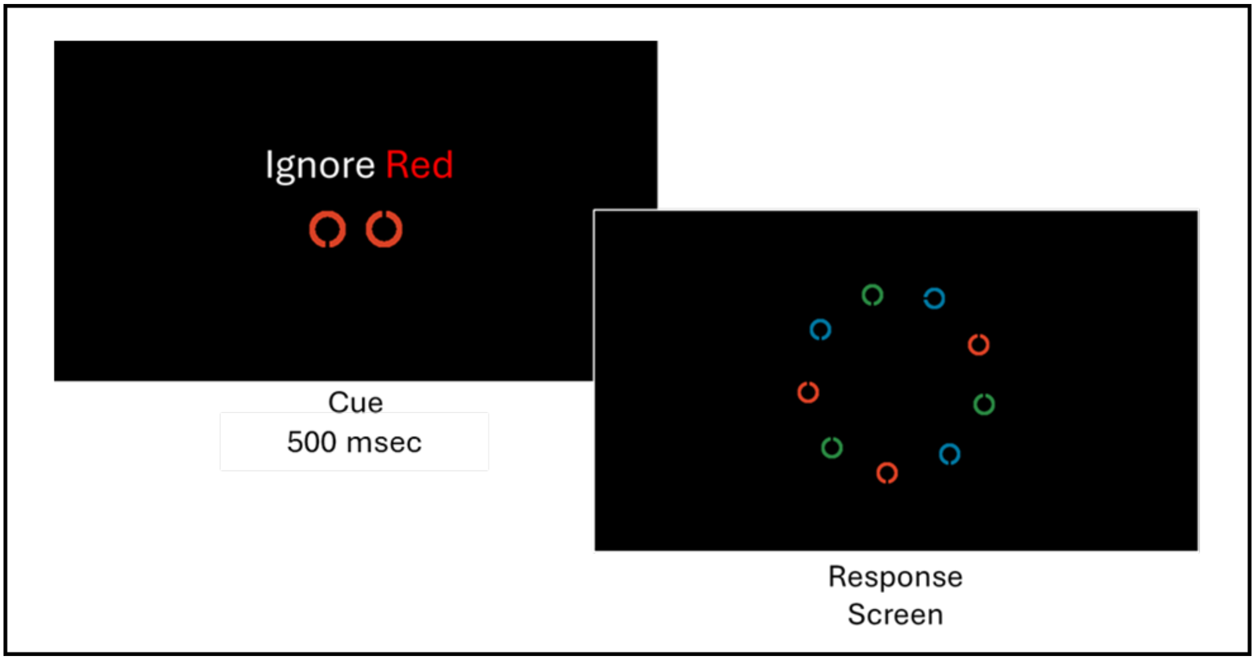

Figure 34. An example of a possible trial in the training task. Each trial begins with a negative cue appearing for 500 ms. The cue instructs which color to ignore to and shows examples of what the distractors will look like. After the cue disappears, the search display appears, consisting of nine Landolt Cs, and stays on screen until a button press is made. Eight Cs are distractors with vertical gaps (either upwards or downward facing) and the target is the C with the horizontal gap (either rightward or leftward facing). Participants indicate the location of horizontal gap via button press. On a given trial, three colors appear, with three Cs in each color. The ignore color (the red Cs) appeared on 80% of trials and the target never appeared in this color.

was analyzed. All the participants had normal or corrected-to-normal acuity. Participants gave written informed consent under the study protocol approved by the Institutional Review Board and were compensated with research credits.

5.3.2 Stimulus, apparatus, and procedure

All stimuli and procedures were identical to those in Exp. 1, with the following exceptions. Participants were explicitly instructed that one color (ignore color) would appear on a large proportion of trials, and when it appears, the target would not be in that color so they should actively ignore to this color. The ignore color (either red or green) remained constant across the experiment for a given participant. In the training task, each trial began with the

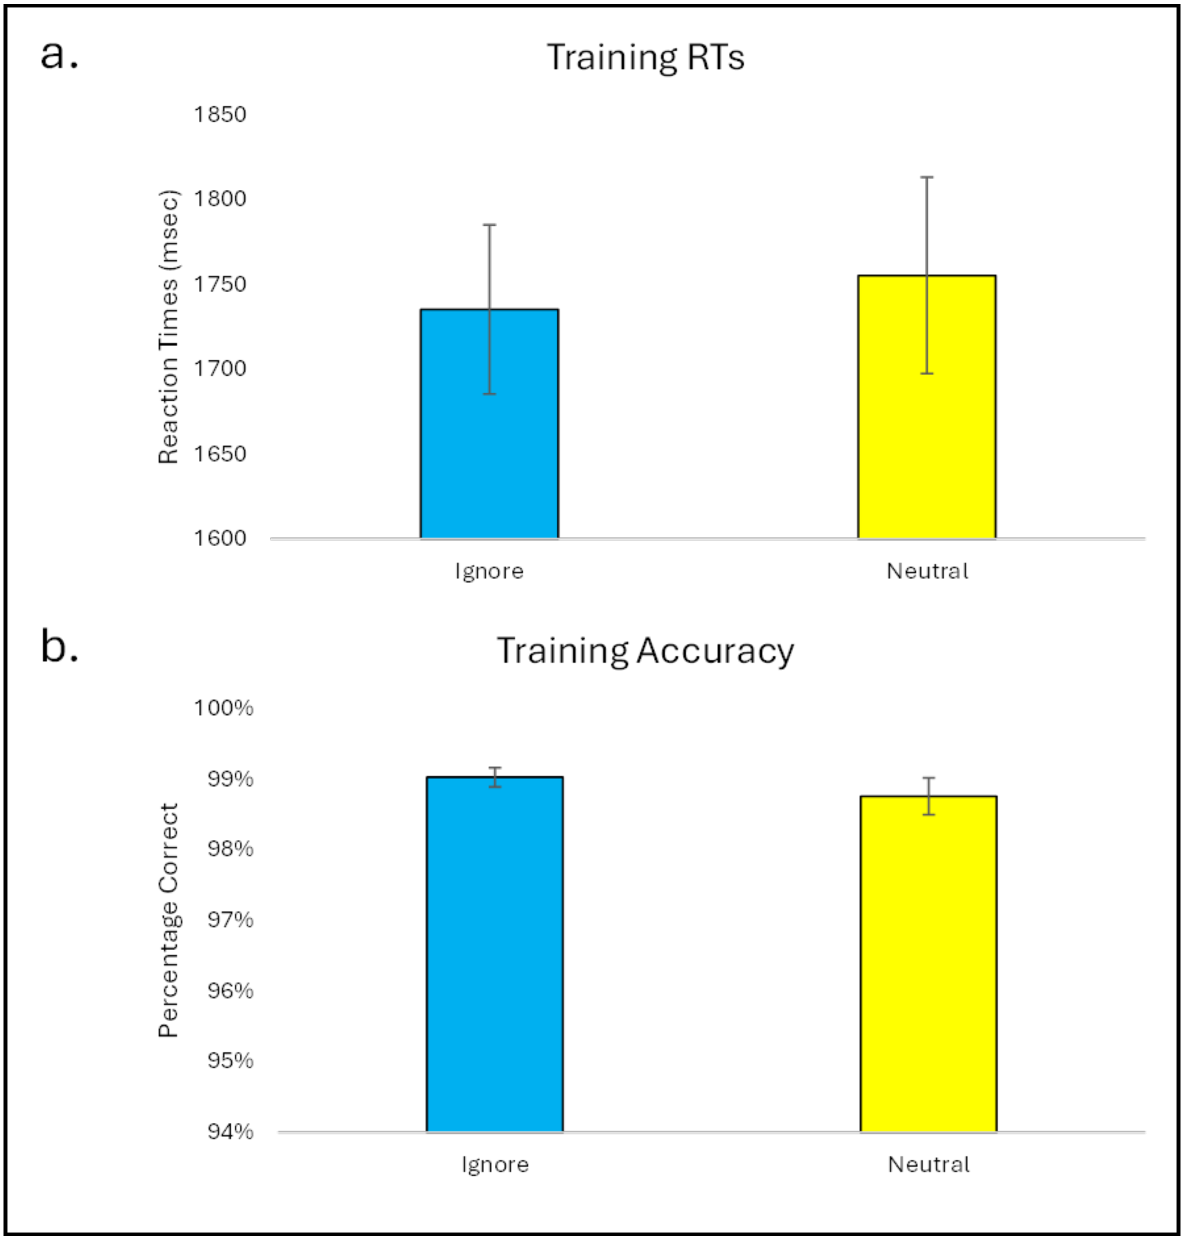

Figure 35. Search times (a.) and accuracy (b.) for the training task by condition.

presentation of a negative cue for 500 ms followed by the search array (figure 34). The cue matched the ignore color that appeared in the search array on 80% of trials, and when the ignore color appeared, the target was never in that color. Participants performed 240 trials of the training task split into two blocks of 120 trials, with a rest in between.

For the test block, the suppression contingency from the training task was removed. Additionally, the negative cue from the training task was also removed. Participants performed 60 trials of test block.

5.4 Results

5.4.1 Training Block

Trials with reaction times (RTs) more than 3 standard deviations away from the overall mean and with incorrect responses were excluded from the analyses, resulting in removal of 1.09% of total trials. To examine if there is a benefit of explicit feature-based suppression, we conducted a paired samples t-test on RTs between the ignored (trials where the cued color appeared in the array) and neutral trials (when the cued color did not appear in the search array). This analysis revealed no difference in RTs between ignored trials (M = 1735.38 ms, SE = 49.83 ms) and the neutral (M = 1755.34 ms, SE = 57.98 ms) trials (figure 35a), $t(61) = 1.03$, $p = .306$. Additionally, a Bayesian paired samples t-test implemented in JASP (Version 0.19.1) compared the ignore and neutral trials on RTs. This analysis provided moderate evidence for the null hypothesis with a Bayes Factor (BF_{10}) = .23.

We also performed a paired samples t-test on accuracy, but this analysis also showed no difference between the ignore (M = 99.03%, SE = .13%) and neutral (M = 98.76%, SE = .26%) trials (figure 35b), $t(61) = 1.19$, $p = .241$. A Bayesian paired samples t-test reconfirmed the above result, finding moderate evidence for the null hypothesis with a BF_{10} = .27. Across both the RT and accuracy data, there is no evidence of a benefit of explicit FBS. However, we decided to conduct time course analyses to see if the suppression benefit appears later in the training task as product of learning and repetition of a constant to-be-ignored feature, as has been found by previous research (Cunningham & Egeth, 2016; Vatterott & Vecera, 2012).

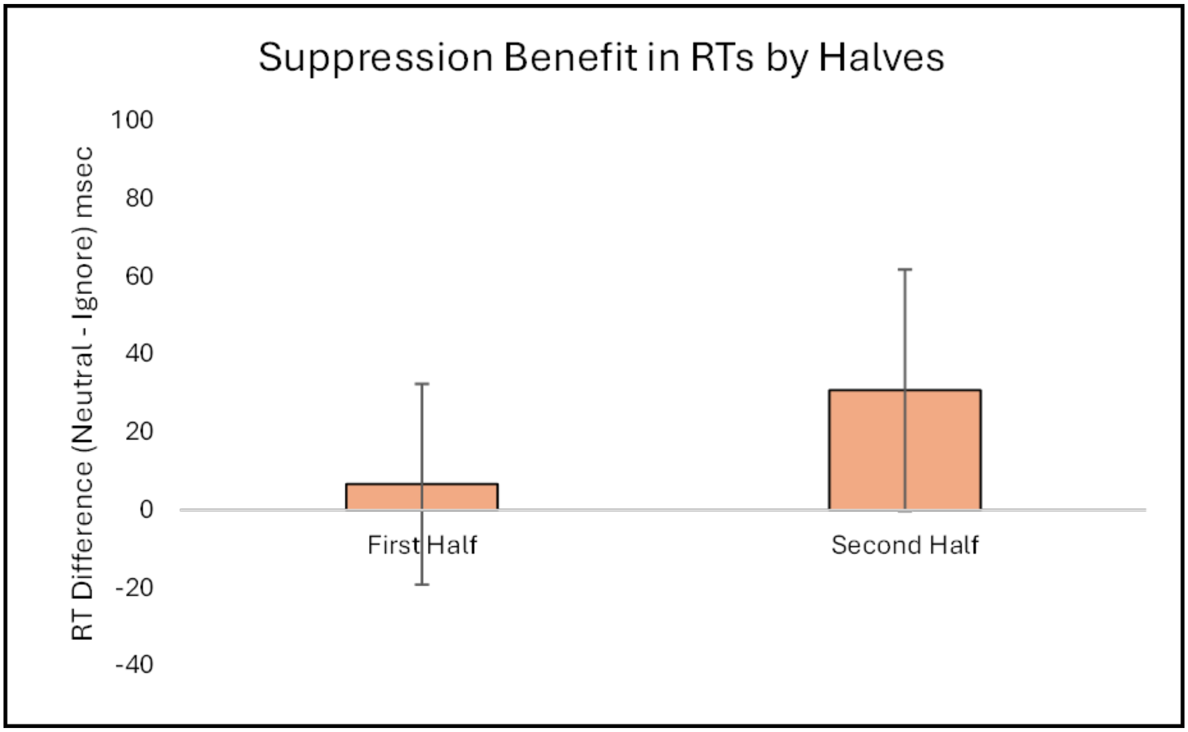

Figure 36. The suppression benefit (subtracting the RTs in ignore trials from the RTs in neutral trials) in the training task split by half.

5.4.2 Time Course Analysis for Training Task

We conducted follow-up block analysis splitting the training task into halves (first half vs. second half) to examine whether a suppression effect emerges after prolonged learning that a given color is never associated with the target. For this analysis, we calculated the suppression benefit (neutral trials - ignore trials). A paired samples t-test revealed no difference in suppression effect between the first half (*M* = 6.65 ms, *SE* = 25.78 ms) and second half (*M* = 30.79 ms, *SE* = 31.79 ms) of the training block, $t(61) = .59$, $p = .560$, indicating there is not a greater benefit of suppression from prolonged periods of learning (figure 36). Furthermore, we performed one samples t-tests on the first half and second half suppression effects to see if either is significantly different from zero as this would indicate that there is a suppression benefit. The one sample t-test analysis revealed no difference for both the first half, $t(61) = .29$, $p = .791$, and second half, $t(61) = .99$, $p = .325$, indicating that the explicit suppression cuing did not benefit

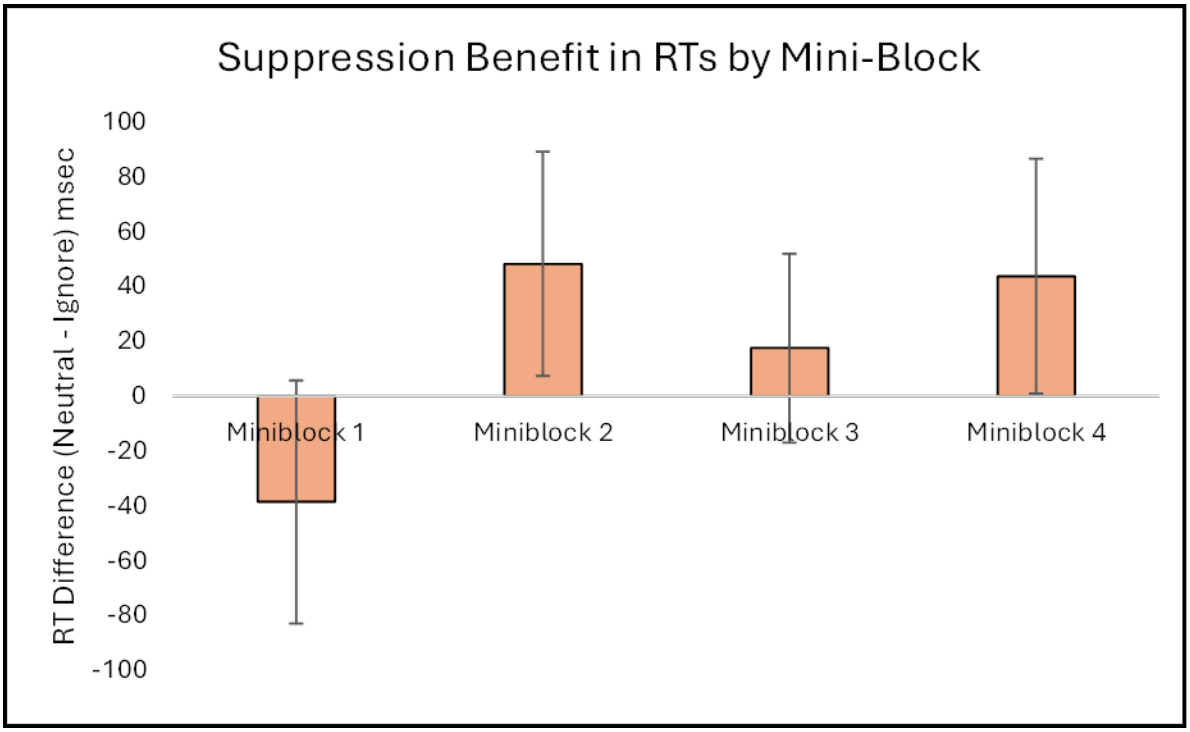

Figure 37. The suppression effect (subtracting the RTs in ignore trials from the RTs in neutral trials) in the training task split by quarters.

search performance. For completeness, we conducted Bayesian one sample t-tests on the suppression effects for the first half and second half of the training task. This analysis provided moderate evidence for the null hypothesis with a $BF_{10} = .14$ for the first half and a $BF_{10} = .22$ for the second half.

We also investigated the suppression effect by quarters as it may reveal more nuanced learning that is lost by examining the suppression benefit by halves. While there is a trend towards a suppression benefit developing in RT data (figure 37), this analysis still found no differences in suppression effect across the first quarter (M = -38.53 ms, SE = 44.31 ms), second quarter (M = 48.33 ms, SE = 40.90 ms), third quarter (M = 17.56 ms, SE = 34.43 ms), and fourth quarter (M = 43.78 ms, SE = 42.83 ms) of trials, $F(3, 183) = .92$, $p = .431$, partial eta squared = .015. A Bayesian repeated-measures ANOVA confirmed the above finding, with strong evidence

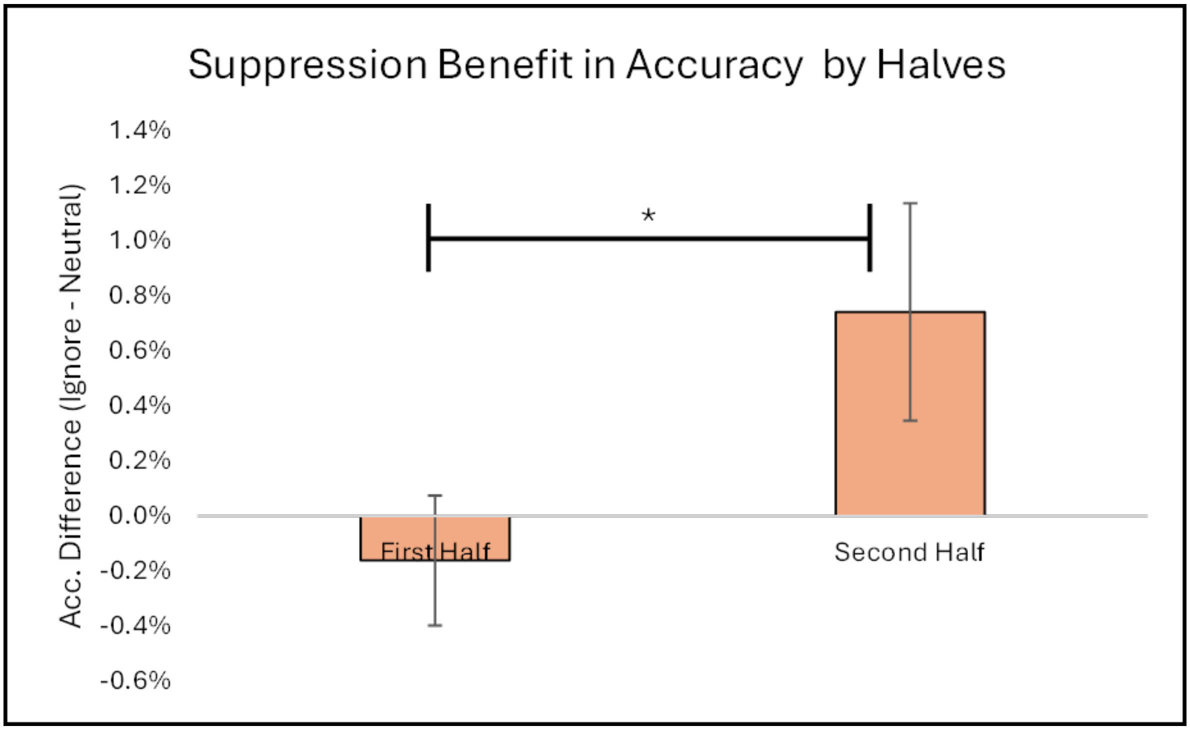

Figure 38. The suppression benefit (subtracting the accuracy of neutral trials from the accuracy in ignore trials) in the training task split by half (* $p < .05$).

for the null hypothesis with a $BF_M = .07$. Thus, there is little evidence of explicit suppression benefiting RTs as product of learning to suppress the same to-be-ignored feature.

Similar analysis was conducted to examine if there was a difference in the suppression effect for accuracy (by subtracting the accuracy of the neutral trials from the ignore trials). A paired samples t-test found a significant difference in suppression benefit between the first half (*M* = -.16%, *SE* = .23%) and second half (*M* = .74%, *SE* = .40%) of the training task, $t(61) = 2.01$, $p = .049$, indicating that there was an improvement in accuracy as from explicitly learning the to-be-ignore feature (figure 38).

A repeated measures ANOVA also found a significant difference in suppression benefit for accuracy across the first quarter (*M* = .46%, *SE* = .37%), second quarter (*M* = -.84%, *SE* = .37%), third quarter (*M* = .60%, *SE* = .50%), and fourth quarter (*M* = .98%, *SE* = .52%) of trials (figure 39), $F(3, 183) = 3.48$, $p = .017$, partial eta squared = .054, with a general trend of

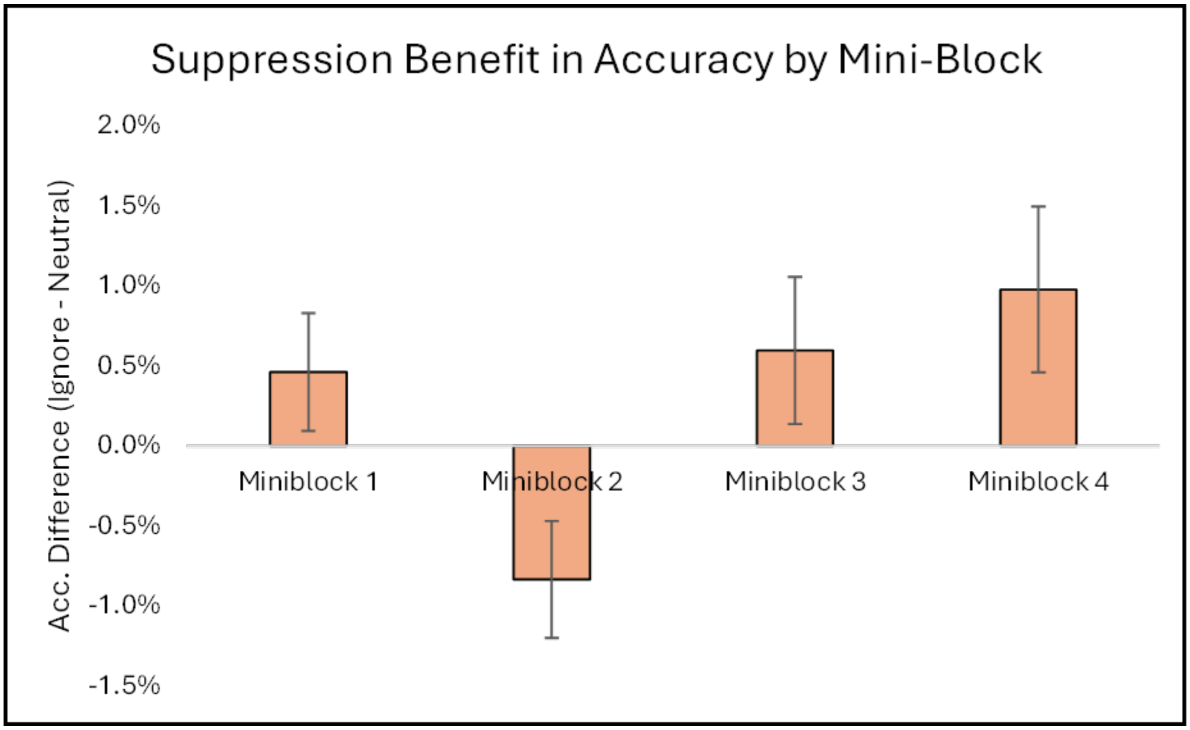

Figure 39. The suppression benefit (subtracting the accuracy of neutral trials from the accuracy in ignore trials) in the training task split by quarters.

accuracy improving as participants learn to suppress the same to-be-ignored feature. While explicit suppression does not appear to benefit search times, there appears to be a suppression benefit for accuracy.

5.4.3 Test Block

We conducted a repeated measures ANOVA on RTs for the three test block conditions (ignore target, ignore distractor, and neutral conditions). This analysis revealed that there were no differences in RTs between the ignore target (M = 1713.47 ms, SE = 61.74 ms), ignore distractor (M = 1797.43 ms, SE = 58.06 ms), and neutral (M = 1743.18 ms, SE = 60.58 ms) conditions (figure 40a), $F(2, 122) = 1.65$, $p = .196$, partial eta squared = .026. This null finding was confirmed by a Bayesian repeated-measures ANOVA, with moderate evidence for the null hypothesis with a $BF_M = .23$.

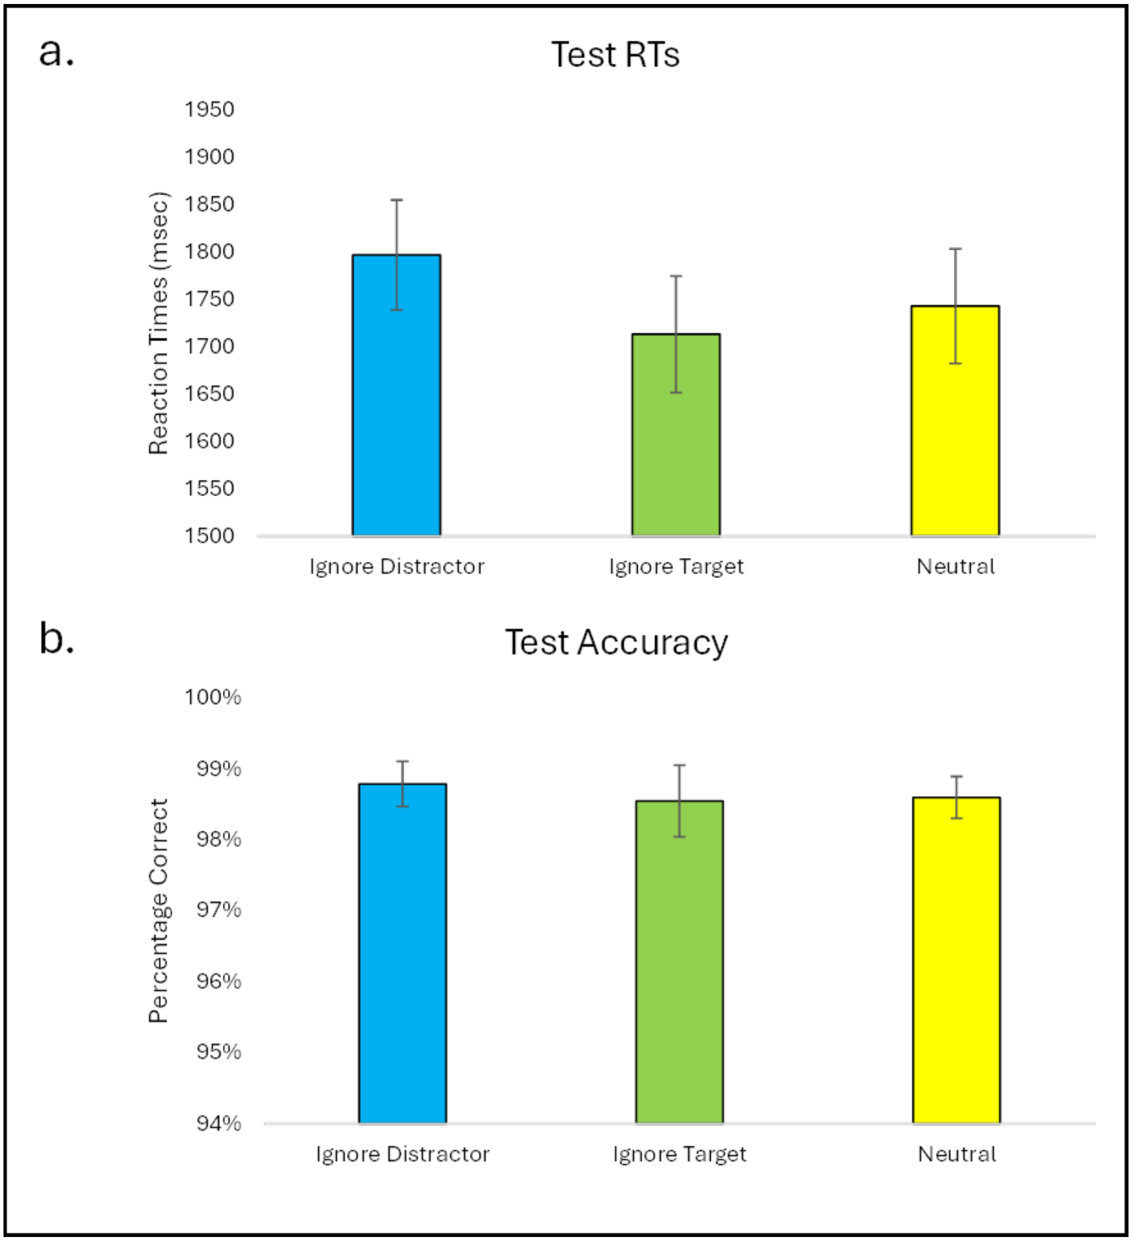

Figure 40. Search times (a.) and accuracy (b.) for the test block by condition.

We also performed a repeated measures ANOVA on accuracy, but also found no difference between ignore target (*M* = 98.55%, *SE* = .51%), ignore distractor (*M* = 98.79%, *SE* = .32%), and neutral (*M* = 98.60%, *SE* = .29%) conditions (figure 40b), $F(2, 122) = .11$, $p = .892$, partial eta squared = .002. This null result was also confirmed by a Bayesian repeated-measures ANOVA, with strong evidence for the null hypothesis with a $BF_M = .08$. Like the training task,

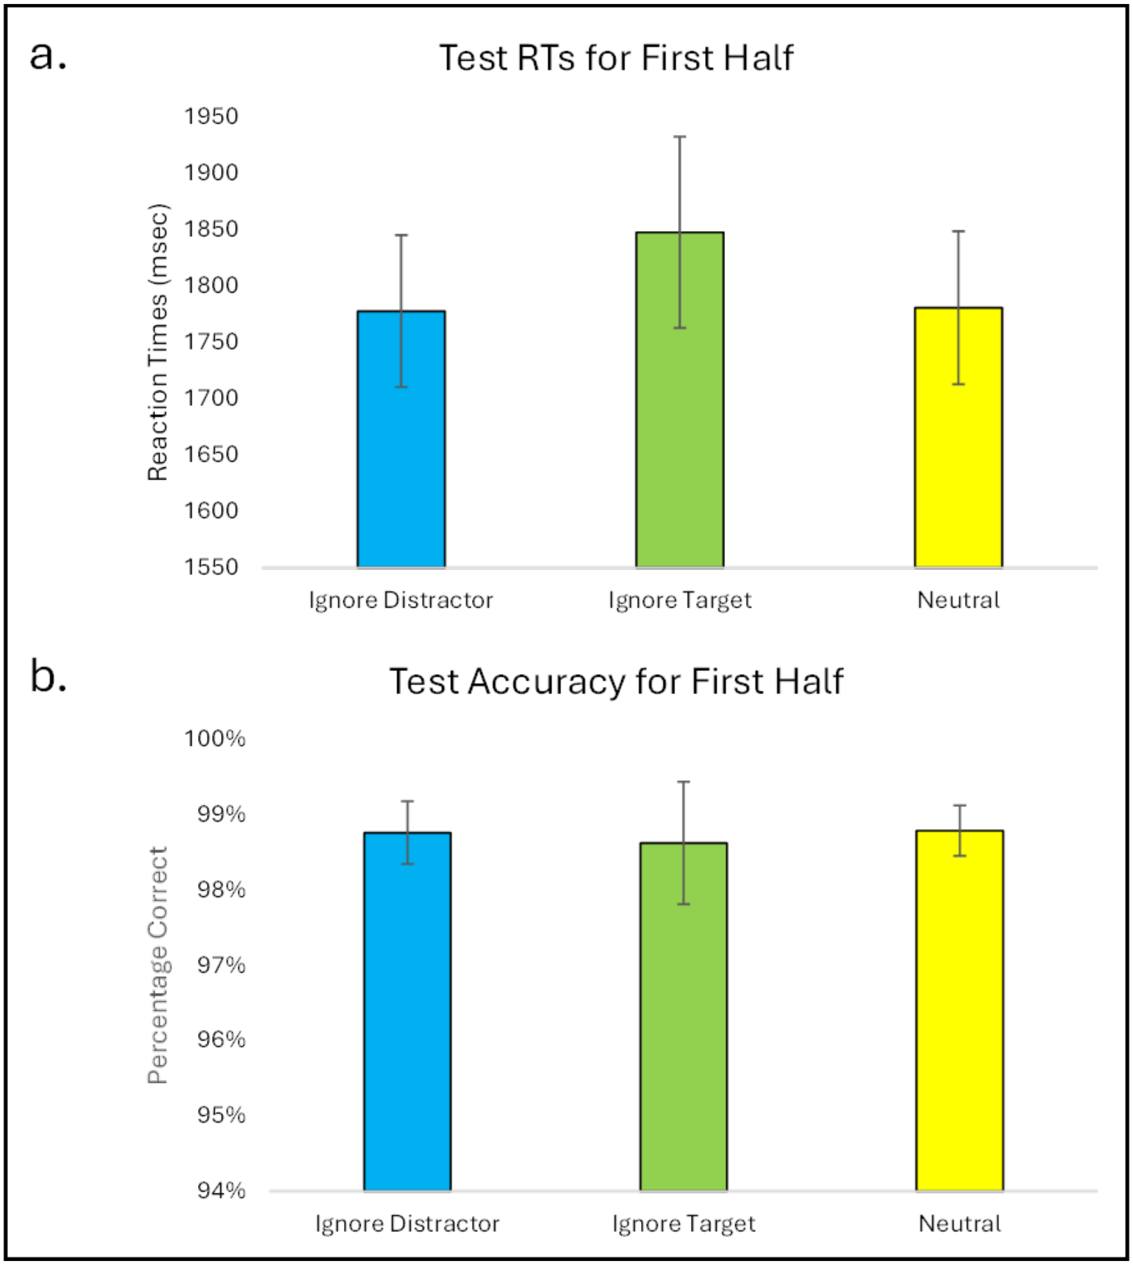

Figure 41. Search times (a.) and accuracy (b.) for the first half of test block by condition.

there does not appear to be a suppression benefit in either the RT or accuracy data.

However, we also conducted a time course analysis for the first half of the test block to see if there is evidence of a suppression effect that quickly disappears. A repeated measures ANOVA revealed that there were no differences in RTs between the ignore target (*M* = 1847.36 ms, *SE* = 84.82 ms), ignore distractor (*M* = 1777.59 ms, *SE* = 67.44 ms), and neutral (*M* =

1780.51 ms, *SE* = 67.85 ms) conditions (figure 41a), *F*(2, 122) = .72, *p* = .488, partial eta squared = .012, and was further confirmed by a Bayesian repeated-measures ANOVA, with moderate evidence for the null hypothesis with a BF_M = .10.

We also performed a repeated measures ANOVA on accuracy for the first half of the test block, but also found no difference between ignore target (*M* = 98.63%, *SE* = .81%), ignore distractor (*M* = 98.77%, *SE* = .42%), and neutral (*M* = 98.79%, *SE* = .34%) conditions (figure 41b), *F*(2, 122) = .03, *p* = .975, partial eta squared < .001, and was further confirmed by a Bayesian repeated-measures ANOVA, with strong evidence for the null hypothesis with a BF_M = .06. There is no evidence of a suppression effect in the test block, and overall, a very weak benefit of explicit feature-based suppression that is only observed in a small improvement of accuracy over time in the training task.

5.4.4 Comparing the Suppression Effect between Experiments 1, 2, 3, and 4

While that explicit feature-based used in Exp. 4 produced the weakest suppression effects, we compared the magnitude of suppression benefits across Exp. 1, 2, 3, and 4 to confirm this result. To begin, we performed an One-Way ANOVA comparing the suppression benefit (subtracting the RTs in neutral trials from the RTs in ignore trials) in training task for Exp. 1, 2, 3, and 4. While the suppression effect in Exp. 4 (*M* = 19.97 ms, *SE* = 19.34 ms) is numerically smaller than the suppression effects in Exp. 1 (*M* = 35.92 ms, *SE* = 17.50 ms), Exp. 2 (*M* = 34.42 ms, *SE* = 22.48 ms), and Exp. 3 (*M* = 63.29 ms, *SE* = 19.97 ms), this analysis found no difference in suppression benefit, *F*(3, 251) = .86, *p* = .463.

We also compared the difference in suppression benefit (subtracting the RTs in ignore target trials from the RTs in ignore distractor trials) in test block for Exp. 1, 2, 3, and 4. Here we found a significant difference in suppression effect, *F*(3, 251) = 9.79, *p* < .001. Follow-up comparisons confirm a weaker suppression effect in Exp. 4 (*M* = -83.96 ms, *SE* = 51.16 ms) compared to Exp. 1 (*M* = 47.40 ms, *SE* = 38.08 ms), $t(118) = 2.04$, $p = .022$, and Exp. 3 (*M* = 242.99 ms, *SE* = 52.28 ms), $t(131) = 4.44$, $p < .001$. However, there was no difference in suppression effect between Exp. 4 and Exp. 2 (*M* = -30.32 ms, *SE* = 42.75 ms), $t(124) = .81$, $p = .211$. These comparisons confirm that the weakest suppression effects occur in Exp. 2 when there is an implicit feature-based gain mechanism and in Exp. 4 with an explicit feature-based suppression component, while the strongest suppression effects occur in Exp. 1 without a feature-based gain component and in Exp. 3 when there is an explicit feature-based gain mechanism.

5.5 Discussion

In Experiments 1 and 3, there was robust evidence of feature-based suppression benefitting search performance. However, those experiments relied on implicit learning of suppression through historical context. Here, we used the same task but explicitly cued the to-be-ignored color and found very little evidence for FBS. Previous research has found that explicit cuing of feature-based suppression is weaker than implicit learning of suppression (Addleman & Stormer, 2023; Lien, Ruthruff, & Tolomeo, 2024; Stilwell & Vecera, 2019), but others have suggested that explicit feature-based suppression can still impact attention (Arita, Carlisle, & Woodman, 2012; Carlisle & Nitka, 2019; Zhang, Gaspelin, & Carlisle, 2020). By contrast, we find very little evidence that explicitly cuing the to-be-ignored color produced suppression.

However, it is worth noting that it has been suggested that explicit suppression may only develop after many, many trials with a consistent to-be-suppressed feature. For example, Cunningham & Egeth (2016) explicitly cued people about a to-be-ignored color, and found the cue produced a cost rather than benefit, early in the task. This cost slowly diminished and eventually became a suppression benefit in their fourth block of trials. Importantly, they had subjects perform far more trials than we did – each of their blocks were 180 trials – so their fourth block started after 540 trials of training with the same color being the distractor. By contrast, our participants performed only 240 trials – 120 per block. Thus, it is possible that if we had more extensive training, participants may have eventually be able to suppress attention to the to-be-ignored color. Consistent with this view, our results did appear to show a trend of a suppression effect developing in the RT data as the training became more extensive. Early in the training task, participants appeared to bias their attention towards the to-be-ignored feature, but this bias to attend to the to-be-ignored feature weakened (see Figure 39).

Even so, we believe our results suggest that explicit FBS probably plays a limited role in the allocation of attention in most real-world tasks. That is, if explicit suppression only occurs when a given feature is always associated with distractions (never associated with a target) for hundreds of consecutive searches, it is difficult to imagine it occurring often in the real-world.

Our claim that explicit FBS likely plays a minimal role in the allocation of attention is inconsistent with some authors' (Arita, Carlisle, & Woodman, 2010) claims that negative cuing can rapidly create a "template for rejection" (another name for feature-based suppression). However, the methods used in those studies are also unlikely to occur in the real-world. In those experiments, there were only two colors presented and the displays had each color appear in separate hemi-fields. These highly structured displays may have allowed people to rapidly use the negative cue to create a positive cue for the alternate color or hemifield. Experiments that removed the possibility of remapping to a positive cue for the opposite color (Becker, Hemsteger, & Peltier, 2015) or location (Beck, Luck, & Hollingworth, 2015) found no evidence for suppression. Taken together, this pattern of results suggests that one might find evidence for explicit suppression only in highly contrived situations that might rarely occur in the real world, thereby limiting the likely importance of explicit feature-based suppression in controlling the allocation of attention in real world tasks.

In summary, while the findings of Exps. 1 and 3 indicate that implicit learning of FBS may play a critical role in guiding attention, Exp. 4 found only very weak evidence (or arguably, no evidence) of an explicit suppression benefit, suggesting that the real-world benefit of suppression might be confined to circumstances where search history effects support the implicit learning of FBS. Further, we note that the apparent inability to implement explicit FBS is in stark contrast to feature-based gain, as we observed a robust benefit of both implicit (Exp. 2) and explicit (Exp. 3) FBG.

CONCLUSION AND DISCUSSION

The results from the four experiments now allow us to better characterize and contextualize the importance of feature-based suppression in allocating attention. Implicit FBS is a robust and effective mechanism for guiding attention. In the absence of top-down guidance towards target features (as in Experiment 1), implicit FBS can direct attention away from features that are commonly associated with distractors and towards the target, leading to efficient searches in the ignore condition (when a set of distractors appears in the color frequently associated with distractors) compared to neutral condition. The suppression benefit is quickly learned, appearing within the first 60 trials of the training task, and remains at the same magnitude throughout the training task. While the rapid appearance of a suppression benefit was unexpected, previous studies have shown that implicit suppression is learned very quickly (Golan & Lamy, 2022), illustrating the strength of this mechanism in guiding attention towards the appropriate target. While the suppression benefit in Experiment 1 was created by having the ignore color appear as distractors with high statistical regularity, the benefit still persists even when the statistical regularity was removed in the test block.

However, it is worth noting that Experiment 2 suggested some limitations to the situations under which implicit FBS will be effective. Specifically, in Experiment 2 the design included the same implicit FBS contingencies as Experiment 1, but also included contingencies that supported an implicit feature-based gain component. The inclusion of two implicit feature-based attentional mechanisms seemingly stymied the effectiveness of the suppression mechanism while the gain mechanism was effective. This result potentially indicates that the resources available to implement multiple implicit mechanisms are extremely limited and only one implicit mechanism can be represented. Consistent with this view that difficulty in implementing FBS in

Experiment 2 was due to two mechanisms competing for implementation within the implicit system, in Experiment 3 when the feature-based gain mechanism was moved to the explicit realm, the implicit FBS benefit returned. Like Experiment 1, in Experiment 3 the FBS benefit was quickly learned, appearing within the first 50 trials of the training task, and carried over into the test block where the statistical regularities from the training task were removed. Thus, our results suggest that implicit feature-based suppression is effective as long as there are not multiple implicit mechanisms competing for implementation.

Even so, in Experiment 2 when there was the ability to implement both a feature-based gain and suppression mechanism, why did the participants encode the implicit gain mechanism in lieu of the implicit suppression mechanism? There are two potential explanations. First, feature-based gain may be a more powerful and accessible mechanism for allocating attention than feature-based suppression. Previous studies that have compared the utility of gain and suppression mechanisms find that the inclusion of the gain component interferes with suppression (Kawashima & Amano, 2022; Kawashima & Matsumoto, 2018; Kugler et al., 2015; Rajsic, Carlisle, & Woodman, 2020). Thus, it could be that participants learn and implement the more accessible and straightforward mechanism (feature-based gain).

The other possible explanation is that the gain component provides more perceived utility than the suppression component. While we attempted to equalize the amount of information each feature mechanism provided, with the ignore color appearing on 80% of trials but only reducing the search array by 33.33% and the attend color appearing only on 40% of trials but reducing the search array by 66.66%, participants may unconsciously prefer the reduction in search array over the frequency of appearance. The advantage provided by the gain component by reducing the search array to 3 items may outweigh the frequency advantage of the suppression component.

Future work that systematically manipulates the number of items to be searched as function of implementing a gain or suppression mechanism and manipulates the frequency with which a given mechanism is beneficial might be able to tease apart the perceived utility of each factor to determine if implicit FBS may be adopted over implicit feature-based gain. Even with this limitation, our findings suggest that *implicit* feature-based suppression can be a critical mechanism for allocating visual attention.

Given the robust nature of implicit feature-based suppression, this mechanism should be better incorporated into models of feature-based attention. Traditional feature-based attention models have mostly focused on feature-based gain mechanism, with suppression occurring either in intermediate (Hopf et al. 2006; Stormer & Alvarez, 2014), far (Trueu & Trujillo, 1999), or both intermediate and far (Fang, Becker, & Liu, 2019) distances in feature space. However, this suppression occurs as a byproduct of feature-based gain, not as a separate mechanism, as the to-be-attended feature determines which features are suppressed. Our findings strongly suggest that suppression of features can occur as a direct result of implicit feature-based suppression mechanism, indicating its importance as one of the putative mechanisms underlying feature-based attention.

As mentioned above, models of feature-based attention have described the attentional profile of feature-based gain, finding neural upweighting of the to-be-attended feature (Martinez-Trujillo & Treue, 2004) along with aforementioned areas of suppression. Unlike feature-based gain, the documentation for the attentional profile of feature-based suppression is not as extensive. Given that our findings suggest that implicit feature-based suppression is a separate mechanism that is centered on the to-be-ignored feature, theoretically it should have an attentional profile that can be measured. However, there has been little work documenting this attentional profile for feature-based suppression.

Additionally, the areas of the visual cortex that can be modulated by feature-based gain are also well researched, with evidence showing that feature-based gain can influence areas of early- (Liu & Hou, 2011), intermediate- (Saenz, Buracas, & Boynton, 2002; 2003), and late-stage visual processing (Chelazzi, Duncan, Miller, & Desimone, 1998). The areas of the visual system that can be influenced and modulated by feature-based suppression are not as well documented as feature-based gain. While addressing those limitations are outside the scope of this dissertation, future models of feature-based attention should document the attentional profile of feature-based suppression and examine what areas of visual system can be modulated by implicit feature-based suppression.

By contrast, our results suggest that *explicit* FBS is an ineffective and weak mechanism for guiding attention. Experiment 4 replicated Experiment 1 (in which the design supported only a feature-based suppression mechanism) while providing an explicit cue of the to-be-ignored color. Providing this explicit cue to engage a top-down guidance mechanisms *away* from a consistent distractor feature produced very weak to no effect of explicit FBS. Overall, search times and accuracy did not differ between the ignore and neutral conditions. When examining RTs and accuracy in the training task by mini-block, we found that explicit suppression initially hindered search performance as attention was biased towards the to-be-ignored feature. But with additional training with the same to-be-ignored feature, the cost associated with explicit suppression is reduced and it appears participants no longer bias their attention towards the to-be-ignore feature. This pattern of explicit feature-based suppression initially leading to a cost rather than a benefit, but with the cost reducing over exposure to the task is consistent with prior findings on explicit feature-based suppression (Cunningham & Egeth, 2016; Vatterott & Vecera, 2012). However, in those studies the initial costs associated with an explicit suppression cue eventually turned into a benefit, with faster RTs in the negatively cued trials than neutral trials. While we did not find that explicit suppression eventually improved search times this may have been due to a failure to have enough trials for the pattern to completely reverse. Nevertheless, the need to build up an explicit suppression benefit through ample training is in stark contrast to implicit FBS, which benefits search performance almost immediately and does not require excessive training to become beneficial.

Given that the explicit suppression effect in Experiment 4 is weak to non-existent, it is difficult to justify examining if the weak effects here can persist if there is also a feature-based gain mechanism to guide attention, as those suppression effects would almost certainly be nullified by the much stronger gain benefit. Nevertheless, the results across these studies confirm that implicit FBS is a more robust and powerful mechanism for allocating attention than explicit FBS (Addleman & Stormer, 2023; Lien, Ruthruff, & Tolomeo, 2024; Stilwell & Vecera, 2019). Implicit FBS produces larger suppression benefits that emerge rapidly without copious training compared to explicit FBS. While some studies (Moher et al., 2014) have claimed that FBS, and not FBG, is the main mechanism responsible for allocating visual attention, the results across these four experiments do not support this claim. Even though implicit FBS is a useful and important mechanism in guiding attention, our findings confirm that it is not as effective or as useful as FBG (Beck & Hollingworth, 2012; Beck, Luck, & Hollingworth, 2018; Becker, Hemsteger, & Peltier, 2015; Olivers, Meijer, & Theeuwes, 2006; van Moorselaar, Theeuwes, & Olivers, 2014). Unlike FBS, FBG is robust and successful at guiding attention both when explicitly cued (as in Experiment 3) and implicitly learned (as in Experiment 2), making it a more flexible mechanism than FBS.

In conclusion, feature-based suppression has some limitations that reduces its utility as a real-world mechanism for allocating visual attention. In particular, *explicit* feature-based suppression has very little real-world applicability as it requires considerable repetition to the same to-be-ignore feature to overcome the initial bias to attend to it and to develop into a suppression effect. It is incredibly unlikely that observers would partake in the same visual search repeatedly while consciously ignoring the same to-be-ignored feature.

However, *implicit* feature-based suppression can play an important role guiding attention in real-world visual searches, even in the face of strong explicit target guidance. It requires very little repetition of the same to-be-ignore feature to become effective and requires no conscious effort to suppress the to-be-ignore feature. In real-world searches, like grocery shopping where many items in the store are never purchased, observers may be able to unconsciously ignore the undesirable items' features to direct their attention towards features matching their desired product without having to engage in the same search hundreds of times. While implicit feature-based suppression may not be as powerful as feature-based gain, it can still assist and improve search performance in every day life.

BIBLIOGRAPHY

Addleman, D. A., & Störmer, V. S. (2023). Distractor ignoring is as effective as target enhancement when incidentally learned but not when explicitly cued. *Attention, Perception, & Psychophysics*, *85*(3), 834-844.

Arita, J. T., Carlisle, N. B., & Woodman, G. F. (2012). Templates for rejection: configuring attention to ignore task-irrelevant features. *Journal of experimental psychology: human perception and performance*, *38*(3), 580.

Beck, V. M., & Hollingworth, A. (2015). Evidence for negative feature guidance in visual search is explained by spatial recoding. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(5), 1190.

Beck, V. M., Hollingworth, A., & Luck, S. J. (2012). Simultaneous control of attention by multiple working memory representations. *Psychological science*, *23*(8), 887-898.

Beck, V. M., Luck, S. J., & Hollingworth, A. (2018). Whatever you do, don't look at the...: Evaluating guidance by an exclusionary attentional template. *Journal of Experimental Psychology: Human Perception and Performance*, *44*(4), 645.

Becker, M. W., Hemsteger, S., & Peltier, C. (2015). No templates for rejection: A failure to configure attention to ignore task-irrelevant features. *Visual Cognition*, *23*(9-10), 1150 - 1167.

Britton, M. K., & Anderson, B. A. (2020). Specificity and persistence of statistical learning in distractor suppression. *Journal of Experimental Psychology: Human Perception and Performance*, *46*(3), 324.

Carlisle, N. B., & Nitka, A. W. (2019). Location-based explanations do not account for active attentional suppression. *Visual Cognition*, *27*(3-4), 305-316.

Chang, S., & Egeth, H. E. (2019). Enhancement and suppression flexibly guide attention. *Psychological science*, *30*(12), 1724-1732.

Chang, S., & Egeth, H. E. (2021). Can salient stimuli really be suppressed? *Attention, Perception, & Psychophysics*, *83*, 260-269.

Chelazzi, L., Duncan, J., Miller, E. K., & Desimone, R. (1998). Responses of neurons in inferior temporal cortex during memory-guided visual search. *Journal of neurophysiology*, *80*(6), 2918-2940.

Cohen, J. (1992). Statistical power analysis. *Current directions in psychological science*, 1(3), 98-101.

Conci, M., Deichsel, C., Müller, H. J., & Töllner, T. (2019). Feature guidance by negative attentional templates depends on search difficulty. *Visual Cognition*, *27*(3-4), 317-326.

Cunningham, C. A., & Egeth, H. E. (2016). Taming the white bear: Initial costs and eventual benefits of distractor inhibition. *Psychological science*, *27*(4), 476-485.

De Waard, J., Van Moorselaar, D., Bogaerts, L., & Theeuwes, J. (2023). Statistical learning of distractor locations is dependent on task context. *Scientific Reports*, *13*(1), 11234.

Egeth, H. E., Virzi, R. A., & Garbart, H. (1984). Searching for conjunctively defined targets. *Journal of Experimental Psychology: Human Perception and Performance*, *10*(1), 32.

Fang, M. W., Becker, M. W., & Liu, T. (2019). Attention to colors induces surround suppression at category boundaries. *Scientific reports*, *9*(1), 1443.

Failing, M., Feldmann-Wüstefeld, T., Wang, B., Olivers, C., & Theeuwes, J. (2019). Statistical regularities induce spatial as well as feature-specific suppression. *Journal of experimental psychology: human perception and performance*, *45*(10), 1291.

Forschack, N., Andersen, S. K., & Müller, M. M. (2017). Global enhancement but local suppression in feature-based attention. *Journal of Cognitive Neuroscience*, *29*(4), 619-627.

Gaspelin, N., Leonard, C. J., & Luck, S. J. (2015). Direct Evidence for Active Suppression of Salient-but-Irrelevant Sensory Inputs. *Psychological Science, 26*(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, & Psychophysics*, *79*, 45-62.

Golan, A., & Lamy, D. (2022). Is Statistical Learning of a Salient Distractor's Color Implicit, Inflexible and Distinct From Inter-Trial Priming? *Journal of cognition*, *5*(1).

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*, *103*(4), 1053-1058.

Kawashima, T., & Amano, K. (2022). Can enhancement and suppression concurrently guide attention? An assessment at the individual level. *F1000Research*, *11*.

Kawashima, T., & Matsumoto, E. (2018). Negative cues lead to more inefficient search than positive cues even at later stages of visual search. *Acta psychologica*, *190*, 85-94.

Kerzel, D., & Burra, N. (2020). Capture by context elements, not attentional suppression of distractors, explains the PD with small search displays. *Journal of Cognitive Neuroscience*, *32*(6), 1170-1183.

Kerzel, D., Huynh Cong, S., & Burra, N. (2021). Do we need attentional suppression?. *Visual cognition*, *29*(9), 580-582.

Kugler, G., ‘t Hart, B. M., Kohlbecher, S., Einhäuser, W., & Schneider, E. (2015). Gaze in visual search is guided more efficiently by positive cues than by negative cues. *PloS one*, *10*(12), e0145910.

Lien, M. C., Ruthruff, E., & Hauck, C. (2022). On preventing attention capture: Is singleton suppression actually singleton suppression? *Psychological research*, 1-14.

Lien, M. C., Ruthruff, E., & Tolomeo, D. (2024). Evidence that proactive distractor suppression does not require attentional resources. *Psychonomic Bulletin & Review*, *31*(3), 1376 - 1386.

Liu, T., & Hou, Y. (2011). Global feature-based attention to orientation. *Journal of Vision*, *11*(10), 8-8.

Liu, T., Larsson, J., & Carrasco, M. (2007). Feature-based attention modulates orientation selective responses in human visual cortex. *Neuron*, *55*(2), 313-323.

Martinez-Trujillo, J. C., & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current biology*, *14*(9), 744-751.

Maunsell, J. H., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in neurosciences*, *29*(6), 317-322.

Moher, J., & Egeth, H. E. (2012). The ignoring paradox: Cueing distractor features leads first to selection, then to inhibition of to-be-ignored items. *Attention, Perception, & Psychophysics*, *74*, 1590-1605.

Moher, J., Lakshmanan, B. M., Egeth, H. E., & Ewen, J. B. (2014). Inhibition drives early feature-based attention. *Psychological science*, *25*(2), 315-324.

Navalpakkam, V., & Itti, L. (2006). Top–down attention selection is fine grained. *Journal of Vision*, *6*(11), 4-4.

Olivers, C. N., Meijer, F., & Theeuwes, J. (2006). Feature-based memory-driven attentional capture: visual working memory content affects visual attention. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(5), 1243.

Oxner, M., Martinovic, J., Forschack, N., Lempe, R., Gundlach, C., & Müller, M. (2023). Global enhancement of target color—not proactive suppression—explains attentional deployment during visual search. *Journal of Experimental Psychology: General*, *152*(6), 1705.

Peirce, J. W. (2007). PsychoPy—psychophysics software in Python. *Journal of neuroscience methods*, *162*(1-2), 8-13.

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., ... & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior research methods*, *51*, 195-203.

Rajsic, J., Carlisle, N. B., & Woodman, G. F. (2020). What not to look for: Electrophysiological evidence that searchers prefer positive templates. *Neuropsychologia*, *140*, 107376.

Reeder, R. R., Olivers, C. N., & Pollmann, S. (2017). Cortical evidence for negative search templates. *Visual Cognition*, *25*(1-3), 278-290.

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.

Saenz, M., Buraĉas, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision research*, *43*(6), 629-637.

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, & Psychophysics*, *72*(6), 1455-1470.

Serences, J. T., & Boynton, G. M. (2007). Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron*, *55*(2), 301-312.

Stilwell, B. T., & Vecera, S. P. (2019a). Cued distractor rejection disrupts learned distractor rejection. *Visual Cognition*, *27*(3-4), 327-342.

Stilwell, B. T., & Vecera, S. P. (2019b). Learned and cued distractor rejection for multiple features in visual search. *Attention, Perception, & Psychophysics*, *81*, 359-376.

Stilwell, B. T., & Vecera, S. P. (2020). Learned distractor rejection in the face of strong target guidance. *Journal of experimental psychology: human perception and performance*, *46*(9), 926.

Stilwell, B. T., & Vecera, S. P. (2023). Learned distractor rejection persists across target search in a different dimension. *Attention, Perception, & Psychophysics*, *85*(3), 785-795.

Störmer, V. S., & Alvarez, G. A. (2014). Feature-based attention elicits surround suppression in feature space. *Current Biology*, *24*(17), 1985-1988.

Töllner, T., Conci, M., & Müller, H. J. (2015). Predictive distractor context facilitates attentional selection of high, but not intermediate and low, salience targets. *Human Brain Mapping*, *36*(3), 935-944.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive psychology*, *12*(1), 97-136.

Treue, S., & Martinez-Trujillo, J. C. (2007). Attending to features inside and outside the spotlight of attention. *Neuron*, *55*(2), 174-176.

Treue, S., & Trujillo, J. C. M. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, *399*(6736), 575-579.

van Moorselaar, D., & Theeuwes, J. (2024). Transfer of statistical learning between tasks. *Journal of Experimental Psychology: Human Perception and Performance*, 1-12.

van Moorselaar, D., Theeuwes, J., & Olivers, C. N. (2014). In competition for the attentional template: Can multiple items within visual working memory guide attention? *Journal of Experimental Psychology: Human Perception and Performance*, *40*(4), 1450.

Vatterott, D. B., & Vecera, S. P. (2012). Experience-dependent attentional tuning of distractor rejection. *Psychonomic bulletin & review*, *19*, 871-878.

Vickery, T. J., King, L. W., & Jiang, Y. (2005). Setting up the target template in visual search. *Journal of vision*, *5*(1), 8-8.

White, A. L., & Carrasco, M. (2011). Feature-based attention involuntarily and simultaneously improves visual performance across locations. *Journal of vision*, *11*(6), 15-15.

Williams, L. G. (1966). The effect of target specification on objects fixated during visual search. *Perception & Psychophysics*, *1*(5), 315-318.

Won, B. Y., & Geng, J. J. (2020). Passive exposure attenuates distraction during visual search. *Journal of experimental psychology: general*, *149*(10), 1987.

Woodman, G. F., & Luck, S. J. (2003). Serial deployment of attention during visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(1), 121.

Woodman, G. F., & Luck, S. J. (2007). Do the contents of visual working memory automatically influence attentional selection during visual search? *Journal of Experimental Psychology: Human Perception and Performance*, *33*(2), 363.

Yantis, S. (1993). Stimulus-driven attentional capture and attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, *19*(3), 676-681.

Yeshurun, Y., & Carrasco, M. (1998). Attention improves or impairs visual performance by enhancing spatial resolution. *Nature*, *396*(6706), 72-75.

Zhang, W., & Luck, S. J. (2009). Feature-based attention modulates feedforward visual processing. *Nature neuroscience*, *12*(1), 24-25.

Zhang, Z., Gaspelin, N., & Carlisle, N. B. (2020). Probing early attention following negative and positive templates. *Attention, Perception, & Psychophysics*, *82*, 1166-1175.