

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Capabilities and consequences of feature-based attention

Permalink

<https://escholarship.org/uc/item/2fj7w3fq>

Author

Chapman, Angus

Publication Date

2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Capabilities and consequences of feature-based attention

A dissertation submitted in partial satisfaction of the  
requirements for the degree Doctor of Philosophy

in

Experimental Psychology

by

Angus Frederick King Chapman

Committee in charge:

Professor John Serences, Chair  
Professor Viola Störmer, Co-Chair  
Professor Tim Brady  
Professor Steve Hillyard  
Professor Anastasia Kiyonaga

Copyright

Angus Frederick King Chapman, 2022

All Rights Reserved

The dissertation of Angus Chapman is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2022

TABLE OF CONTENTS

Dissertation Approval Page..... iii

Table of Contents..... iv

List of Figures.....v

List of Tables.....vi

Acknowledgements..... vii

Vita..... x

Abstract of the Dissertation..... xi

Introduction..... 1

Chapter 1.....8

Chapter 2..... 32

Chapter 3..... 68

Chapter 4..... 96

Conclusion..... 126

## LIST OF FIGURES

Figure 1.1. Example of behavioral task.....	15
Figure 1.2. Mean hit rates for speed change detection in Experiment 1.....	18
Figure 1.3. Mean hit rate for speed changes in Experiment 2.....	21
Figure 2.1. Example trial structure for Experiment 1a & 1b.....	40
Figure 2.2. Results of Experiment 1a.....	42
Figure 2.3. Results of Experiment 2.....	49
Figure 2.4. Results of Experiment 3a.....	53
Figure 2.5. Example of trial conditions in Experiment 4.....	58
Figure 3.1. Structure of an example experiment trial.....	74
Figure 3.2. Results from Experiment 3, where there was no delay between the search array and similarity judgment.....	78
Figure 3.3. Hypothetical neural population representing a circular feature space.....	79
Supplementary Figure 3.1. Demonstration of $d'$ calculation for signal detection model simulations.....	89
Supplementary Figure 3.2. Visual search response time in Experiment 1 as a function of target-distractor similarity.....	89
Supplementary Figure 3.3. Estimated similarity judgment performance in Experiment 1 as a function of visual search RT.....	90
Figure 4.1. Example of the stimulus display for the non-selective task in Experiment 1.....	101
Figure 4.2. Average model-based reconstruction for stimuli in the non-selective task in Experiment 1.....	109
Figure 4.3. Average model-based reconstructions of target and distractor stimuli in the selective attention task in Experiment 1.....	111
Figure 4.4. Example of the stimulus display in Experiment 2.....	113
Figure 4.5. Average model-based reconstructions of target and distractor stimuli in Experiment 2.....	116

LIST OF TABLES

Supplementary Table 1.1. Mean (SD) performance on each task in Experiment 1..... 27

Supplementary Table 1.2. Mean (SD) performance on each task for each group in  
Experiment 2..... 27

Supplementary Table 3.1. Best fitting parameters for the signal detection model for each target-distractor  
similarity level..... 90

## ACKNOWLEDGEMENTS

This dissertation is the culmination of ideas formed through a series of conversations, debates, off-hand remarks, thought experiments, but (mostly) focused discussions with many smart and interesting people. First, my advisors John and Viola, who invited me to UCSD in the first place, and have made my time in this graduate program rewarding both personally and professionally. John, you have a skill for cultivating a group of people who are constantly doing their best to advance our understanding of the brain, but knowing that the best way to do so is to do it alongside people you care about. I've learned a lot from you on how to be a better scientist, while keeping an eye on the bigger picture of what it means to have a life and career in science. Viola, it has been such an honor to be one of your first graduate students. You have always made so much time for me, even for me to spend multiple hours explaining an idea to you that ultimately just loops back to what you had said in the first place. I'm so excited to see the direction that your lab at Dartmouth goes in next, and hopefully I can still stay a part of that. Thank you to my dissertation committee for their advice and mentorship. Tim, who has basically been my third advisor, and has sat through every single lab meeting I've ever given (all of which were on attention, the best research topic). Your advice and perspective on academia has been invaluable. Steve, thank you for being so supportive and interested in my research even before I asked you to be on my committee. And Anastasia, I was so glad when I found out you were joining the Cog Sci department here, and I'm looking forward to seeing the work that comes out of your lab in the future.

I have been so lucky to have so many amazing labmates along this journey: Dr. Jonathan Keefe, Jamal Williams, Dr. Anna Shafer-Skelton, Dr. Kirsten Adam, Sunyoung Park, Janna Wennberg, Tim Sheehan, Holly Kular, Leah Etensohn, Stella Yuan, Dr. Chaipat Chunharas, Dr. Maggie Henderson, Dr. Nuttida Rungratsameetaweemana, Dr. Rosanne Rademaker, Dr. Doug Addleman, Yong Hoon Chung, Kevin Ortego, Mert Ozkan, Dr. Steph Nelli, Dr. Vy Vo, Frederik Geweke, Dr. Maria Robinson, Dr. Mark Schurgin, Isabella Destefano, Hayden Schill, Michael Allen, Will McCarthy, Yang Wang, Dr. Malinda McPherson. None of this would have been nearly as fun or rewarding without you. Thank you to the terrific undergrads who I worked with on some of the projects that made up this dissertation: Audrey



Barszcz, Young-jin Choi, Lora Hsu, Xinwen Wang, and Ashley Williams. As well, my graduate cohort have been an irreplaceable source of support: Chloe, Aubrey, Katie, Tanushree, Xinqi, Liz, Ethan, and Derek. There are too many other grads in our department to name, but thanks especially to Lauren, Alex, Erik, Andrew, and Holly. I couldn't have wished for a cooler bunch of people to spend five years hanging out in San Diego with. Thank you to everyone involved in GPSA who work exceptionally hard to make sure student voices are heard on campus: Breana, Rachel, Quynh, Jack, Burgundy, Victor, Mia, Laura, Andy, Hannah, and everyone I got to serve with on various committees over the years. Thank you especially to my friends outside of campus, who were there when I needed social support, memes, or someone to play video games with: Rowan, Steph, Elyse, Cameron, Danielle, Callum, Keegan, Jordan, Sara, Lauren, Matthew, Nici, Faith, Alice, Julie, Rory, and the many others who tolerate me.

I was extremely fortunate to receive a Fulbright New Zealand Science and Innovation Award that provided additional financial support during my time in grad school. My thanks to the entire team at Fulbright New Zealand who were amazingly helpful and supportive at preparing me in my move to the US, and special thanks to Pip Climo for her extra help during my time here. Thank you to my colleagues in New Zealand who helped me make it to UCSD in the first place: Dr. Gina Grimshaw, Dr. Tirta Susilo, Dr. Josh Foster, Dr. Michael Tooley, Dr. Hazel Godfrey, Dr. Sophie Hedley, Dr. Amy Walsh, Dr. Kealagh Robinson, Laura Kranz, Lauren Bell, Daniel Jenkins, Sumaya Lamb, and many others.

Lastly, thank you to my family. My brothers, Oli and Hunter, and my parents, Donna and Brett, who have supported my academic journey for over 30 years now. You have always given me the freedom to explore my interests, and I couldn't have made it this far without your help.

Chapter 1, in full, is a reprint of the material as it appears in *Psychonomic Bulletin & Review*, 28, 1252-1260, 2022, Chapman, Angus F.; Störmer, Viola S., by Springer Nature. Reproduced with permission from Springer Nature. The dissertation author was the primary investigator and author of this paper.

Chapter 2, in full, has been submitted for publication. Chapman, Angus F.; Störmer, Viola S. The dissertation author was the primary investigator and author of this paper.

Chapter 3, in full, has been submitted for publication. Chapman, Angus F.; Chunharas, Chaipat; Störmer, Viola S. The dissertation author was the primary investigator and author of this paper.

Chapter 4 is currently being prepared for submission for publication. Chapman, Angus F.; Geweke, Frederik; Serences, John T.; Störmer, Viola S. The dissertation author was the primary investigator and author of this material.

## VITA

- 2014 Bachelor of Science, Victoria University of Wellington
- 2015 Bachelor of Science (Hons), Victoria University of Wellington
- 2020 Master of Arts, University of California San Diego
- 2022 Doctor of Philosophy, University of California San Diego

## PUBLICATIONS

- Chapman, A. F., & Störmer, V. S. (2022). Feature similarity is non-linearly related to attentional selection: evidence from visual search and sustained attention tasks. *Journal of Vision*, 22(8:4). <https://dx.doi.org/10.1167/jov.22.8.4>
- Chapman, A. F. & Störmer, V. S. (2021). Feature-based attention is not confined by object boundaries: spatially global enhancement of irrelevant features. *Psychonomic Bulletin & Review*, 28, 1252-1260. <https://dx.doi.org/10.3758/s13423-021-01897-x>
- Chapman A., Devue, C., & Grimshaw, G. M. (2019). Fleeting reliability in the dot-probe task. *Psychological Research*, 83(2), 308-320. <https://dx.doi.org/10.1007/s00426-017-0947-6>
- Chapman, A. F.\*, Hawkins-Elder, H.\*, & Susilo T. (2018). How robust is familiar face recognition? A repeat detection study of more than 1000 faces. *Royal Society Open Science*, 5, 170634. <https://dx.doi.org/10.1098/rsos.170634> \*co-first authors
- Tooley, M. D., Carmel, D., Chapman, A., & Grimshaw, G. M. (2017). Dissociating physiological components of unconscious emotion. *Neuroscience of Consciousness*, 3(1). <https://dx.doi.org/10.1093/nc/nix021>

## ABSTRACT OF THE DISSERTATION

Capabilities and consequences of feature-based attention

by

Angus Frederick King Chapman

Doctor of Philosophy in Experimental Psychology

University of California San Diego, 2022

Professor John Serences, Chair  
Professor Viola Störmer, Co-Chair

Attention—the process of selecting and prioritizing relevant stimuli in our environment—has long been considered a highly important cognitive process that enables us to deal with the large volume of information that we encounter at any moment. While we still do not have a complete understanding of the mechanisms that guide selection, different theories have carved out how attention might function with respect to distinct stimulus domains, namely spatial locations, bound objects, or basic visual features. A large body of research has emerged that investigates the behavioral implications of attentional selection within these domains, as well as the neural processes that underlie this selection. To extend upon this work, in this dissertation I present a series of studies exploring fundamental aspects of feature-based attention: how attention spreads between different features within a single object, as well as globally across locations in the visual field; how broader and narrower ranges of relevant colors can be attended

simultaneously and efficiently; how attention warps the perception of colors in service of efficient processing, and; how information about target and distractor colors are carried within a neural marker of visual processing, and how representations of those colors are affected by attention. In each of these studies, I tie the main findings back to broader theories connecting feature-based attention with object-based and spatial theories of attention. By considering how these theories interact and make claims about the nature of attentional selection, I argue for a framework in which attention is thought of as a process by which the representational content of information is shaped in support of adaptive behavior, as opposed to a mechanism which highlights particular types of stimuli.

## INTRODUCTION

The role of the visual system, and of sensory systems more generally, is to extract behaviorally relevant information from the vast array of raw input we receive from our sensory organs. Ultimately, vision is a process of refining the basic characteristics of the light that hits our retina (e.g., wavelength, intensity, position) into a representation of the world around us. To assist in the parsing of this information—at least according to current theories—attention allows for the prioritization of aspects of the visual input that are either particularly salient or relevant to our ongoing behavior (so called “bottom-up” and “top-down” attention, respectively; Carrasco, 2011; Desimone & Duncan, 1995). Research over the past several decades has illuminated our understanding of how attention can dramatically affect our experience of the world. For example, attending to a particular stream of information, such as auditory information presented in one ear (Deutsch & Deutsch, 1963; Treisman, 1964) or to a visual stimulus shown in one location on a display (Mack & Rock, 1998), can result in little to no awareness of information presented outside of the focus of attention. This *selective* aspect of attention—enhancing processing of attended information at the detriment of unattended information—has become a key factor in our understanding of how processing unfolds during perception.

One major theoretical division in research on attention is what type of perceptual representation is actually targeted by attentional processes. The major domains of attention that have been extensively studied include spatial locations (Itti & Koch, 2000; Posner, 1980; Wolfe, 1994), basic visual features (e.g., color, orientation, motion direction; Maunsell & Treue, 2006; Sàenz, Buraças, & Boynton, 2003), bound objects (O’Craven, Downing, & Kanwisher, 1999; Shomstein, 2012), and points in time (Denison, Carrasco, & Heeger, 2021). While some of the understanding garnered in specific domains has provided insight into more general principles that underlie attention, most research still focuses on investigating aspects that are specific to one domain or another. Indeed, it could be inferred from the literature that some researchers believe that a fully domain-general theory of attention is impossible, and that each domain has its own unique characteristics and mechanisms that can only be understood in isolation of others. Some attempts at more general theories of attention have tried to argue that one domain or another

is the “true” domain of attentional selection, and therefore attempt to subsume the other domains into them (Cavanagh, Caplovitz, Taissa, Maechler, & Tse, preprint; Chen, 2012; Scholl, 2001). In my research, I have taken a perspective somewhere in the middle ground: different domains likely have unique attentional mechanisms given that attention in each case is acting upon different underlying representational structures, however we should strive to find commonalities across domains where possible, in the hope that a more general understanding of the role of attention on information processing is possible. Therefore, even when investigating a specific domain, researchers should keep an eye to how their findings might reflect a broader understanding of attention.

In this dissertation, I aim to uncover principles of attentional selection in the visual system through an in-depth investigation into one of these domains: feature-based attention. Feature-based attention describes the way that basic visual properties of an image—colors, orientations, motion directions, shape—can be enhanced during visual processing. Attention to features differs from other domains of attention: while spatial attention describes how attention can be directed to specific locations in the visual field, often conceptualized as a “spotlight” moving from one location to another, feature-based attention acts globally and enhances the processing of the attended feature throughout the visual field (Saenz et al., 2003; White & Carrasco, 2011); object-based attention describes the way in which we focus on bound object representations (e.g., a car driving past you on the street, or a tennis ball hit towards you by your opponent), where selection enhances all the features contained by that object, rather than a single feature in isolation (O’Craven et al., 1999; Schoenfeld, Hopf, Merkel, Heinze, & Hillyard, 2014). Of course, these domains are not fully independent: objects are composed of a variety of basic features and exist at particular positions in space. Why then is it useful to study feature-based attention? First, feature-based attention has arguably received less dedicated research in the literature compared to both spatial and object-based attention, meaning that many basic aspects feature-based selection remain unclear or unknown. Second, a rich body of research in neuroscience has paved a solid understanding of how different types of features are encoded in visual cortex, meaning that there are direct predictions we can make about the way representations of those features might be modulated by attention. Indeed, we

have learned a significant amount about how attention towards a specific feature affects the response of populations of neurons that are tuned to different features within a feature space (Martinez-Trujillo & Treue, 2004; McAdams & Maunsell, 1999). Finally, I believe that by understanding feature-based attention more deeply, we can begin to piece together the bigger picture of how selection works across multiple domains: are there general mechanisms that support attentional selection in all contexts, or does the brain use distinct processes depending on the relevant stimulus? To answer this question, we need to have a solid grasp on what is similar and different about attention in each domain.

In Chapter 1, I investigate how feature- and object-based attention might interact. Previous studies have demonstrated that feature-based attention enhances the processing of a target feature throughout the visual field, while studies of object-based attention have shown that attention can spread from one feature of an object to another, allowing for selection of the entire bound object. Critically, these theories make distinct predictions about how attention will enhance the processing of secondary features that are part of an attended object: feature-based theories argue that all attended features should be enhanced globally, while object-based theories instead predict that attention should be confined to the boundaries of the attended object. In two experiments, I demonstrate that attention enhances all object features in a global manner, consistent with the predictions of feature-based theories. This finding places a significant challenge on theories that argue that objects are the primary units of attention and suggests that object boundaries may not be as relevant for attentional selection as once thought.

In Chapter 2, I shift focus to understanding a particular aspect of feature-based attention, specifically whether attention can be directed to a range of features at once or is limited to selection of one feature at a time. The ability to broaden or narrow the scope of attention has been well established in the spatial domain, where attention has been demonstrated to adjust efficiently and flexibly when different spatial extents are relevant. Despite this, there is limited evidence for such flexibility towards ranges of relevant features. Across four experiments, I provide clear evidence in support of the idea that the breadth of feature-based attention can be flexibly adjusted, allowing for simultaneous selection of a range of feature values. Further, these findings show that selection of a broad range of colors occurs



relatively efficiently, with relatively little decrease in performance as the range of attended features increases. This argues in support of the idea that feature-based attention, like spatial attention, is flexible and adaptable in the response to changes in the relevant information in our environment.

In Chapter 3, I demonstrate a novel aspect of how feature-based attention influences our perception. Previous research has shown that attention can affect perception of a target feature, by increasing its perceived contrast or saturation, for example. Other research has shown that attention is sensitive to the context of selection, and that when target and distractor features are more similar, enhancement of a non-target feature is optimal for selection. I hypothesized that this optimal tuning might therefore result in distortions in the perceived target color. By asking participants to make judgments about the specific color they had selected moments before, I found that they had systematic biases in the target color they perceived, consistent with the idea that attention warps features in a way that is optimal for selection. This evidence for warping challenges previous work that frames attention as a process that directly enhances the processing of relevant information, and instead suggests that attention can act to exaggerate the difference between similar features, allowing for more efficient selection. Additionally, to account for the effects of attention on the underlying feature space, I introduce a new framework for understand attentional warping based on representational geometry.

In Chapter 4, I extend upon the idea of perceptual warping by attention and examine how attention changes the neural representation of visual features using electroencephalography (EEG). Specifically, I used steady-state visual evoked potentials (SSVEPs) to “tag” the processing of colored stimuli and examine how attention directed towards or away from particular colors affects SSVEP responses. In two experiments, I first use inverted encoding models (IEMs) to show that SSVEPs carry information about target features (in both experiments) and may potentially carry information about distractors (in Experiment 2). This builds on research that has used multivariate decoding models to show that neural measures can recover the identity of a specific visual feature. Further, when targets and distractors were more similar, the representation of those features within the SSVEP response showed evidence of repulsion, such that they were systematically biased away from one another. This provides

converging evidence with the behavioral findings of Chapter 3, suggesting that attention directly acts on the underlying representation of target and distractor features in a way which supports better separation of relevant and irrelevant information.

To conclude the dissertation, I will attempt to bring together the ideas from each chapter with findings from the literatures on feature-based, object-based, and spatial attention more broadly. Based on my reading of the relevant theories and empirical findings I disagree that one type of stimulus is the true “unit” of selection, as well as the idea that each of these domains are truly distinct. Instead, as I will argue through the following chapters, attention can be better understood as a mechanism that shapes the processing of information by transforming representational spaces. Under this way of thinking, attention is not a mechanism by which the brain picks out discrete pieces of information from our environment, but a process that shapes the flow of information dynamically across cortex.

## References

- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*(13), 1484–1525. <https://doi.org/10.1016/j.visres.2011.04.012>
- Cavanagh, P., Caplovitz, G. P., Taissa, K., Maechler, M. R., & Tse, P. U. (n.d.). Object-Based Attention. *PsyArxiv*. <https://doi.org/10.31234/osf.io/2bsn7>
- Chen, Z. (2012). Object-based attention: A tutorial review. *Attention, Perception, and Psychophysics*, *74*(5), 784–802. <https://doi.org/10.3758/s13414-012-0322-z>
- Denison, R. N., Carrasco, M., & Heeger, D. J. (2021). A dynamic normalization model of temporal attention. *Nature Human Behaviour*, *5*(12), 1674–1685. <https://doi.org/10.1038/s41562-021-01129-1>
- Desimone, R., & Duncan, J. (1995). Neural Mechanisms of Selective Visual Attention. *Annual Review of Neuroscience*, *18*(1), 193–222. <https://doi.org/10.1146/annurev.ne.18.030195.001205>
- Deutsch, J. A., & Deutsch, D. (1963). Attention: Some Theoretical Considerations. *Psychological Review*, *70*(1), 80–90. <https://doi.org/10.1037/h0039515>
- Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, *40*(10–12), 1489–1506. [https://doi.org/10.1016/S0042-6989\(99\)00163-7](https://doi.org/10.1016/S0042-6989(99)00163-7)
- Mack, A., & Rock, I. (1998). *Inattentional Blindness*. MIT Press. <https://doi.org/10.7551/mitpress/3707.001.0001>
- 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. <https://doi.org/10.1016/j.cub.2004.04.028>
- Maunsell, J. H. R., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, *29*(6), 317–322. <https://doi.org/10.1016/j.tins.2006.04.001>
- McAdams, C. J., & Maunsell, J. H. R. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience*, *19*(1), 431–441. <https://doi.org/10.1523/jneurosci.19-01-00431.1999>
- O’Craven, K. M., Downing, P. E., & Kanwisher, N. (1999). fMRI evidence for objects as the units of attentional selection. *Nature*, *401*(6753), 584–587. <https://doi.org/10.1038/44134>
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*(1), 3–25. <https://doi.org/10.1080/00335558008248231>
- Sàenz, M., Buraças, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision Research*, *43*(6), 629–637. [https://doi.org/10.1016/S0042-6989\(02\)00595-3](https://doi.org/10.1016/S0042-6989(02)00595-3)
- Schoenfeld, M. A., Hopf, J.-M., Merkel, C., Heinze, H.-J., & Hillyard, S. A. (2014). Object-based attention involves the sequential activation of feature-specific cortical modules. *Nature Neuroscience*, *17*(4), 619–624. <https://doi.org/10.1038/nn.3656>

- Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, *80*, 1–46. [https://doi.org/10.1016/S0010-0277\(00\)00152-9](https://doi.org/10.1016/S0010-0277(00)00152-9)
- Shomstein, S. (2012). Object-based attention: Strategy versus automaticity. *Wiley Interdisciplinary Reviews: Cognitive Science*, *3*(2), 163–169. <https://doi.org/10.1002/wcs.1162>
- Treisman, A. (1964). Monitoring and storage of irrelevant messages in selective attention. *Journal of Verbal Learning and Verbal Behavior*, *3*(6), 449–459. [https://doi.org/10.1016/S0022-5371\(64\)80015-3](https://doi.org/10.1016/S0022-5371(64)80015-3)
- White, A. L., & Carrasco, M. (2011). Feature-based attention involuntarily and simultaneously improves visual performance across locations. *Journal of Vision*, *11*(6), 1–10. <https://doi.org/10.1167/11.6.15>
- Wolfe, J. M. (1994). Guided Search 2.0 A revised model of visual search. *Psychonomic Bulletin & Review*, *1*(2), 202–238. <https://doi.org/10.3758/BF03200774>

CHAPTER 1: Feature-based attention is not confined by object boundaries: spatially global enhancement  
of irrelevant features

## Abstract

Theories of visual attention differ in what they identify as the core unit of selection. Feature-based theories emphasize basic visual features (e.g., color, motion), demonstrated through enhancement of attended features throughout the visual field, while object-based theories propose that attention enhances all features belonging to the same object. These theories make distinct predictions about the processing of features that are not attended primarily: Object-based theories predict that such secondary, task-irrelevant features are enhanced within object boundaries, while feature-based theories predict enhancement of irrelevant features across locations, regardless of objecthood. To test these two accounts, we had participants attend a set of colored dots among distractor dots (moving coherently upwards or downwards) to detect brief luminance decreases, while simultaneously detecting speed changes in other sets of dots in the opposite visual field. In the first experiment, we demonstrate that participants have higher speed detection rates in the dot array that matched the motion direction of the attended color array, although motion direction was task-irrelevant. In a second experiment, we manipulated the probability that speed changes occurred in the matching motion direction and found that enhancement of the irrelevant motion direction persisted even when it was detrimental for task performance, suggesting that spatially global effects of feature-based attention cannot easily be flexibly adjusted. Overall, these results indicate that features that are not primarily attended are enhanced globally, surpassing object boundaries.

## Introduction

When confronted with a crowded visual scene, we can selectively process relevant information by attending to a particular feature (e.g., the color red) or object (e.g., a red car moving leftward; Carrasco, 2011; Desimone & Duncan, 1995; Maunsell & Treue, 2006). Different theories have advanced the importance of either basic visual features ('feature-based' attention) or bound objects ('object-based' attention) for non-spatial selection. These theories differ in what they consider to be the primary unit of attentional selection and propose different mechanisms for how such selection occurs.

Feature-based theories are primarily supported by the finding that attention to a feature at one location enhances processing of that feature throughout the visual field (Andersen, Hillyard, & Müller, 2013; Martinez-Trujillo & Treue, 2004; Rossi & Paradiso, 1995; Sàenz, Buraças, & Boynton, 2002; Wegener et al., 2008; White & Carrasco, 2011). For example, selection of the color red at one location enhances processing of that color across the visual field. Neural data indicates that enhancement of the attended feature occurs even at task-irrelevant locations (Andersen et al., 2013; Martinez-Trujillo & Treue, 2004; Sàenz et al., 2002; Treue & Martinez-Trujillo, 1999; Wannig, Stanisor, & Roelfsema, 2011), and even in the absence of visual stimulation (Serences & Boynton, 2007). This global effect of feature-based attention also has behavioral consequences, such as higher task performance when attending to the same feature in two locations than to opposing features (Andersen, Hillyard, & Müller, 2008; Andersen et al., 2013; Sàenz, Buraças, & Boynton, 2003; Störmer & Alvarez, 2014; Xiao et al., 2014). Similarly, attending to a specific feature in one task (a single direction of motion or orientation) can improve performance for matching features on a secondary task (White & Carrasco, 2011). Other results showed that a subthreshold motion prime improved participants' detection of the coherent motion *only* when the prime dots' color was the same as the color they attended elsewhere, suggesting that attending to a feature at one location can have secondary influences on perception of another feature at a different location through global enhancement of the attended feature (Melcher, Papathomas, and Vidnyánszky, 2005).

In contrast, object-based theories of attention state that objects are the core unit of attentional selection. Seminal studies on object-based attention showed advantages for shifting attention within an

object (an outlined rectangle) relative to between objects (two outlined rectangles), even when the spatial distance between these attentional shifts was fixed (Egley, Driver, & Rafal, 1994), which was interpreted as attention spreading seemingly automatically within an object (Davis, Driver, Pavani, & Shepherd, 2000; Vecera & Farah, 1994). Later studies showed that this within-object spreading might depend on several factors, such as the spatial and configural uncertainty about the display, indicating that some of these object-based attention effects are likely dependent on task context and may be better explained with an attentional priority account, according to which under certain conditions attention prioritizes enhancement within objects (Yeary & Goldsmith, 2010; Shomstein & Yantis, 2002; 2004; Chen & Cave, 2019; for a recent review, see Shomstein, 2012). In all of these studies, however, object benefits are defined as enhanced processing within an object relative to between objects, assuming that object-benefits are confined – at least to some degree – by object boundaries (Duncan, 1984; Scholl, 2001; Baylis & Driver, 1993)<sup>1</sup>.

Other studies of object-based attention have focused on investigating whether attention to one feature spreads to a ‘secondary’ feature that is part of the same object. In particular, when objects have multiple features (e.g., an array of red dots moving upwards), selecting one feature (e.g., the color red) can modulate processing of the second feature (e.g., upward motion). Evidence for this within-object spreading of attention comes from single-cell recordings in primates (Katzner, Busse, & Treue, 2009), as well as human neuroimaging studies (Ernst, Boynton, & Jazayeri, 2013; O’Craven, Downing, & Kanwisher, 1999; Schoenfeld, Hopf, Merkel, Heinze, & Hillyard, 2014; Schoenfeld et al., 2003). For example, it was found that hemodynamic responses related to the motion of an attended surface were increased even when participants were cued to attend to color, suggesting a spread of attention from the object’s color to its motion direction (Ernst et al., 2013). In all of these studies, the primarily attended feature is cued initially and participants are instructed to select that feature, while the secondary feature is

---

<sup>1</sup> Spreading of attention within an object has been shown to occur even when object boundaries are obscured or require perceptual completion (Behrmann et al., 1998; Davis & Driver, 1997; Moore et al., 1998) and for objects defined by shared surface properties (Adamian et al., 2019; Ernst et al., 2012).



considered task-irrelevant, as participants are not cued to attend to it, nor is that secondary feature directly related to the task participants are performing. The fact that processing of the secondary feature is nonetheless amplified solely through the voluntary selection of the primary feature has been taken as evidence for object-based theories of attention. Together, these results support a view in which all features of an object are efficiently and relatively automatically selected, while there are costs for selection beyond object boundaries (Duncan, 1984; Scholl, 2001).

Thus, theories of object- and feature-based attention make distinct predictions about the processing of information that is not selected via top-down goals. According to a strict object-based view, attention should enhance processing of features that are part of the same object (but not between objects), while feature-based theories predict that features are enhanced across locations regardless of objecthood. That is, object-based theories imply that features of an object are attended *conjointly*, while feature-based theories imply that features are attended *independently*. The feature-based account has some initial empirical support from studies showing neural enhancement of unattended features at task-irrelevant location (Adamian, Andersen, & Hillyard, 2019; Bartsch et al., 2018; Boehler, Schoenfeld, Heinze, & Hopf, 2011; Lustig & Beck, 2012), however, the perceptual consequences of this neurally assessed attentional spreading are not known, as these studies did not measure behavioral performance for the secondary, irrelevant features. Here, we examined how selecting a single feature of an object (i.e., its color) affects visual processing of a secondary feature of the same object (i.e., its motion direction) at another location. We find that processing of the secondary feature – that is part of the attended multi-feature object but never cued to be selected – is enhanced at another location. This processing benefit for the secondary feature persists even when it is detrimental for task performance, consistent with a relatively obligatory account of attentional spreading. We discuss these findings in the context of object- and feature-based theories of attention, and argue that behavioral effects of attentional spreading within objects and across locations do not necessarily belong solely to one theory alone, as often suggested, but can provide important insights into the nature of attentional selection when considered together.

## Experiment 1

In Experiment 1, participants monitored two dot-motion displays, one in each visual hemifield, and completed separate, independent tasks in each of them (Figure 1.1): a color selection task, in which they detected decreases in the luminance of dots in the attended color; and a motion task, in which they detected a speed increase in the other set of dots, regardless of motion direction. Here, we define objects as the spatially-overlapping surfaces of dots in each visual field, as in previous research (e.g., Melcher et al., 2005; Schoenfeld et al., 2014, 2003; Katzner et al., 2009). Importantly, participants were only instructed to attend to one particular color (amongst another), and although the dots were moving in two distinct directions, the direction of motion was irrelevant to the task. Thus, a particular color was cued at the beginning of each trial, and can thus be considered the primarily attended feature, while motion direction was never cued or mentioned to participants as a relevant feature, and thus served as a secondary feature of the object. For the motion task, participants were instructed to attend to all of the dots, regardless of their motion direction, to detect a speed increase. The speed increase could occur with the same probability in the dots moving in the same direction as those in the attended color (feature ‘match’ trials), or in the opposite direction (feature ‘non-match’ trials). Thus, attending to one particular motion direction was not helpful to perform well in this task. If attention to a particular color spreads to the motion direction of the same dot field and subsequently across the visual field, we expect that participants will better detect speed changes that happen to match the direction of the attended color than those that do not.

## Method

### Participants

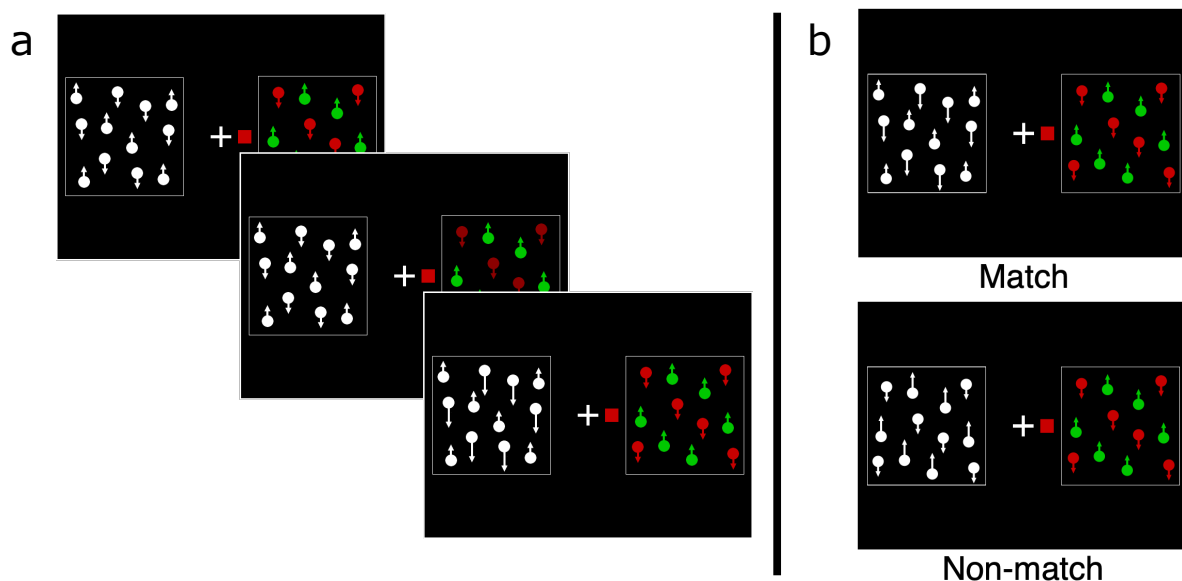
Twenty undergraduate students (18 women) from the University of California, San Diego subject pool participated for course credit. Data from five participants whose performance on one of the tasks was extremely low were excluded from the main analysis ( $d' < 0.5$  across all conditions in one task; in Experiment 2, we preregistered this exclusion criteria). Note that we obtained individual target detection

thresholds prior to the experiment, but only while participants performed one of the two tasks at a time (see Procedure). Thus, it seems as if those five participants with poor performance had overall difficulties in dividing attention between the two tasks in the main experiment. The final sample ( $n = 15$ , 14 women) were aged 18-28 years ( $M = 20.9 \pm 2.6$  years). This sample size provides 80% power to detect an effect of  $d_z > 0.778$ , smaller than previously reported for behavioral effects of feature-based attention (Sàenz et al., 2003; White & Carrasco, 2011). All participants provided written informed consent in accordance with the Institutional Review Board at UC San Diego.

### **Stimuli**

Participants were seated with a viewing distance of approximately 57 cm from the display. Two overlapping fields of dots ( $9.0^\circ$  by  $9.0^\circ$  visual angle) were presented on a black background and centered  $6.02^\circ$  either side of fixation. One comprised two sets of colorful dots (red and green), the other comprised white dots. Each dot field contained 200 dots (each  $\sim 0.2^\circ$  diameter) moving either upwards or downwards at  $2.25^\circ/\text{s}$ . To prevent observers from tracking single dots, each dot had a limited lifetime and was redrawn at a new random location every 300 ms.

In one of the dot displays, half of the dots were green [RGB: 20, 200, 20; luminance:  $37.6 \text{ cd/m}^2$ ], and the remaining half were red [RGB: 200, 20, 20; luminance:  $12.5 \text{ cd/m}^2$ ], and each set of colored dots moved either upwards or downwards (i.e., 100% overlap between color and motion direction thus comprising a dual-feature object). At the beginning of each trial, participants were cued to attend to one of the colors (red or green) with a small color cue (a square in the color of the to-be-attended dots,  $0.4^\circ$  by  $0.4^\circ$ ). This cue was presented slightly to the left or right of fixation (by  $0.4^\circ$ ), indicating the side of the display that the colored dots would be presented (see Figure 1.1). The color cue remained on the screen throughout each trial. After 800 ms, the two dot displays appeared and remained on the screen for 2 s. Dots in the attended color were always drawn on the display first and were occluded by the unattended colored dots if they overlapped. Participants were instructed to attend to the dots in the cued color to detect a brief luminance decrement (300 ms). At the end of each trial, participants had to indicate whether or not this change occurred in the attended dots. The change could occur in the attended dots (50% of



*Figure 1.1.* Example of behavioral task. (a) Participants attended two displays, each containing two sets of dots moving in opposite directions (as indicated by the arrows, which were not present on the display in the experiment; also note that the dots are larger and fewer than in the actual experiment, and a box depicting the border of each dot display is included for illustration purposes). On each trial, a colored square in the center indicated which color of dots to attend to on that side of the display (in this example, attend red) to detect a brief decrease in the dots' luminance (middle display). Simultaneously, participants were monitoring the non-colored (white) dots to detect a brief increase in speed (indicated by the longer arrows in the bottom display) of either set of white dots. (b) The main manipulation of interest concerned whether the speed change matched the direction of the attended color (top display, where both the speed change and the attended red dots were moving downwards) or did not match (bottom display).

trials), the unattended dots (25%), or neither set of dots (25%). The magnitude of the luminance decrement was determined for each participant using an independent thresholding task (see *Procedure*).

In the other dot display, all of the dots were presented in the same color (white [RGB: 200, 200, 200]); half of the dots moved upwards, and the remaining half of the dots moved downwards. Participants were instructed to identify a brief speed increase (300 ms) that occurred on half of the trials, and could occur in either motion direction. Importantly, half of the time the speed change occurred in the group of dots that were moving in the *same* direction as the attended colored dots presented in the other visual half-field (hereafter referred to as 'match' trials), and the remaining half of the time the speed change occurred in the dots moving the *opposite* direction as the attended colored dots (hereafter referred to 'non-match'

trials). Just like the luminance change, the magnitude of the speed change was determined individually for each participant using a thresholding task prior to the main task. The luminance and speed events were determined randomly and independently in each dot display. Each event lasted 300 ms and occurred randomly within one of three time-windows (300-700 ms, 800-1200 ms, or 1300-1700 ms after stimulus onset). If both events occurred on the same trial, they could not occur in the same time window and it was determined randomly which one occurred first.

## **Procedure**

Participants first completed separate thresholding tasks for the color and the motion task to adjust task difficulty. During the thresholding procedure, participants were shown the same displays with two groups of overlapping dots presented left and right of fixation (see Stimuli), but were instructed to either only focus on the color task or the motion task. During the color thresholding, no speed changes occurred, and vice-versa. The magnitudes of the luminance decrement (color task) and speed increase (motion task) were varied using a staircase procedure such that the change became smaller (less detectable) after two consecutive hits, and larger (more detectable) after a miss. For the color task, the luminance decrement was initially set at 40% of the maximum luminance and was adjusted additively by 2% each step. For the motion task, the speed increase was initially set at 1.9x the base speed and was adjusted by 0.04 at each step. Participants completed 64 trials of each task separately (32 target events; the color thresholding was always completed first). Hit rates were fit with a logistic curve (guess rate = 0%) using the Palamedes toolbox (Prins & Kingdom, 2009), and thresholds were selected as the magnitude corresponding to 80% (color task) or 70% (motion task) hit rate. One participant completed the color thresholding twice, and two participants completed the motion thresholding twice, as performance was not fit well after one run. In the main task (Figure 1.1), participants performed the color and motion tasks simultaneously. In particular, they were instructed to attend to the cued colored dots in one of the dot displays (e.g., in the right visual half-field) to detect luminance decrements in the attended color, and at the same time monitor for speed increases in the non-colored (i.e., white) dot display (e.g., in the left visual half-field). A speed increase occurred on half of the trials. Critically, participants were instructed to detect *any* speed increase,

regardless of the motion direction of the dots (upwards or downwards). Note that this response format does not allow us to calculate false alarms in addition to hit rates, because participants reported only if they detected a change in speed, not its direction. This ensured that motion direction remained irrelevant for the motion task. At the end of each trial, participants were prompted to first indicate whether they saw a luminance decrement of the attended colored dots by pressing one of two keys on the keyboard ('m' for a detected change, 'n' for no change) and were then asked to indicate whether they saw a speed increase of the non-colored (white) dots using the same keys. Participants completed 256 trials, which consisted of two full counterbalances of display side (color task left or right of fixation), attended color (green, red), color direction (upwards, downwards), luminance change (50% attended color, 25% unattended color, 25% no change), speed change (50% present, 50% absent), and speed change direction (upwards, downwards; also determines match or non-match trials). Thus, for both tasks the correct response was 'change' on half of the trials. Match and non-match trials occurred equally often, and all trial types were randomly intermixed. Note that participants were always trained and thresholded on the color task before the motion task. In the main task they also responded to events in the colored dots first. This procedure ensured that both tasks were treated as independent.

## Results

Average accuracy was within the expected range on both the color task ( $M = 73.4\%$ ,  $SD = 7.9$ ) and the motion task ( $M = 74.7\%$ ,  $SD = 7.2$ ), and did not differ across the two tasks,  $t(14) = 0.53$ ,  $p = 0.606$ ,  $d_z = 0.14$ . Average hit and false alarm rates for each task are presented in Supplementary Materials. The main question of interest was whether the detection of speed increases in the motion task was influenced by which groups of dots were attended in the color task. We found that participants detected more speed changes in the group of dots that matched the motion direction of the attended color ( $M = 67.2\%$ ) than in the dots that did not match the attended color ( $M = 55.7\%$ ; Figure 1.2),  $t(14) = 5.89$ ,  $p <$

only enhanced visual processing of the primarily attended feature (color) but also processing of the overlapping secondary feature (motion direction); critically, this enhancement of the secondary feature




Figure 1.2. Mean hit rates for speed change detection in Experiment 1. Grey points and lines correspond to individual participants, while the overlaid black line is the group mean with 95% within-subjects confidence intervals.

was not confined to the attended group of colored dots, but instead propagated across the visual field and also enhanced processing of motion direction of the other non-colored dot group. This suggests that selecting a single feature can facilitate processing of secondary features that belong to the same object across the visual field.

## Experiment 2

In Experiment 2 we tested whether the global spreading of attention to a secondary feature is sensitive to how beneficial or detrimental it is for task performance. In Experiment 1, enhanced processing of the motion direction of the colored dots (in addition to their color) may enable better performance in the color task, by decreasing the similarity between targets and distractors (Duncan & Humphreys, 1989; McLeod, Driver, & Crisp, 1988; Wolfe, Cave, & Franzel, 1989). However, attention to

motion direction is not advantageous for the motion task because the speed change occurs equally often in each direction. There are two explanations for spreading of attention in this context: 1) spreading (within and between objects) is obligatory and occurs regardless of perceptual or task demands, or; 2) the spread of attention is flexible and can be adjusted to exploit situations when it is advantageous or avoided when it is disadvantageous.

To test this, we manipulated the frequency of trials in which the speed change matched the direction of the attended color. For half of the participants, 80% of the trials matched, encouraging an “attend to direction” strategy; for the other half of participants, 20% of the trials matched, such that a strategy of “ignore direction” (or even “attend to opposite direction”) would be most beneficial. If the spread of feature-based attention can be flexibly controlled, the size of the effect (match minus non-match accuracy) should be greater when matches are more frequent and should decrease when matches are unlikely. On the other hand, if attentional spreading cannot be adjusted flexibly, the two groups should show an equivalent effect.

## Method

### Participants

We registered the predictions and analysis of this experiment on AsPredicted (<https://aspredicted.org/xb8nn.pdf>) and planned to run 48 participants after exclusions in two experimental conditions (24 per group). This sample size provides 80% power to detect a difference between the groups of at least  $d_s > 0.83$  (which corresponds approximately to a reduction in one group by 50% of the effect observed in Experiment 1). Fifty-six undergraduate students participated in Experiment 2 for course credit. Eight participants were excluded with  $d' < 0.5$  across all conditions on one of the tasks. The remaining 48 participants (38 women) were randomly assigned to the high (80%) or low (20%) match condition. The final sample of participants were between 18-23 years of age ( $M = 20.0 \pm 1.3$  years) and had normal or corrected-to-normal vision.



## Procedure

The procedure and stimulus details were identical to Experiment 1, with the exception that we manipulated the frequency of trials in which the speed change matched the motion direction of the attended color. For half of the participants, 80% of the trials were match-trials, encouraging an “attend to motion direction” strategy; for the other half of participants, 20% of the trials were match-trials, such that a strategy of “ignore motion direction” (or even “attend to opposite direction”) would be most beneficial. To avoid making motion direction explicitly task-relevant, we did not instruct participants about this manipulation.

Participants completed 320 trials consisting of one full counterbalance of display side (color task left or right of fixation), attended color (green, red), color direction (upwards, downwards), luminance change (50% attended color, 25% unattended color, 25% no change), speed change (present, absent), and speed change direction (match to attended color, or non-match; 4-to-1 ratio dependent on condition). Just like in Experiment 1, participants’ individual luminance and speed change thresholds were obtained prior to the main experiment. Out of the 48 participants included in the final analysis, 14 completed more than one run of thresholding before fits were acceptable (8 motion, 5 color, 1 both motion and color).

## Results

The two groups did not differ in terms of their overall thresholds and performance, which we report in detail in Supplementary Materials.

To assess whether participants could flexibly adjust the amount of spreading depending on the percentage of match/non-match trials, we conducted a 2 (trial type: match, non-match)  $\times$  2 (group: 20% or 80% matches) ANOVA on hit rates in the motion task. There was a main effect of trial type (match vs. non-match),  $F(1, 46) = 18.21$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.284$ , replicating the main finding of Experiment 1 across the two groups on average. There was no main effect of group,  $F(1, 46) = 0.38$ ,  $p = 0.543$ ,  $\eta_p^2 = 0.008$ , and crucially no group by trial type interaction,  $F(1, 46) = 0.41$ ,  $p = 0.527$ ,  $\eta_p^2 = 0.009$ , revealing that the

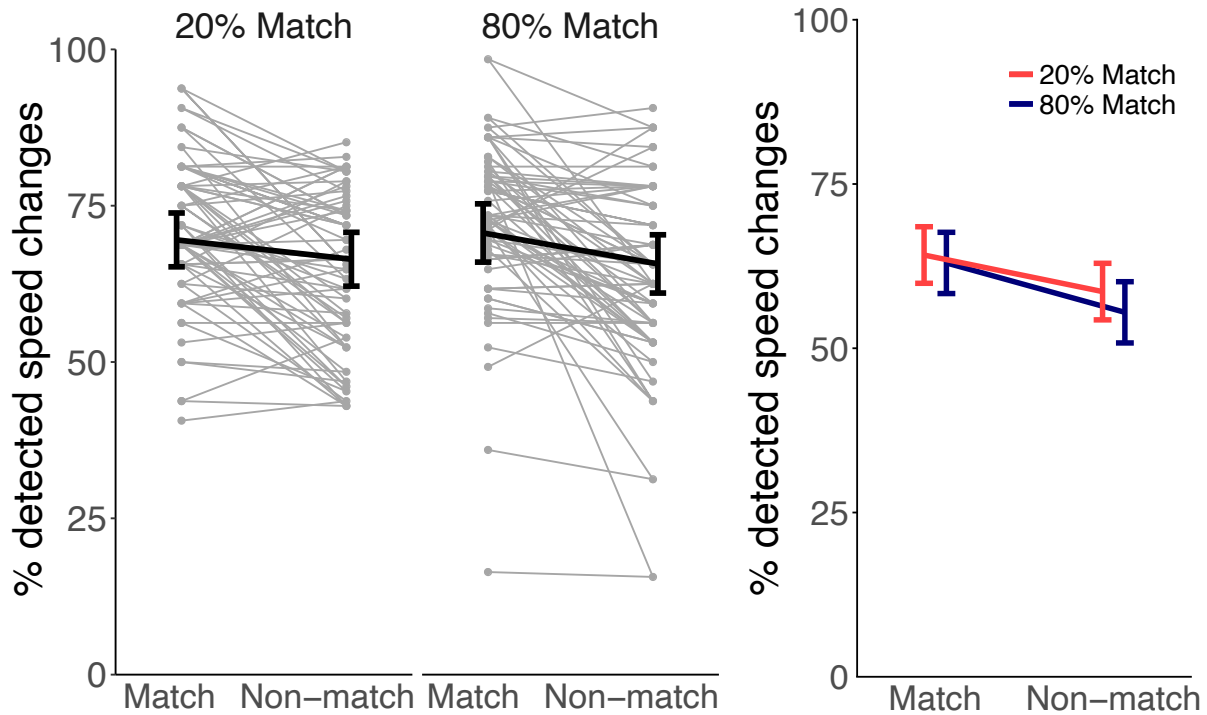


Figure 1.3. Mean hit rate for speed changes in Experiment 2. In the left panel, grey points and lines correspond to individual participants, while the overlaid black line is the group mean. In the right panel, group means are plotted separately. Error bars are 95% within-subjects confidence intervals for each condition.

two groups did not differ in detection rates of the speed changes on match and non-match trials (Figure 1.3). In each group, there was an advantage in detecting speed changes when they matched the direction of the attended colored dots ( $M_{20\%} = 5.6\%$ ,  $t(23) = 2.68$ ,  $p = 0.014$ ,  $d_z = 0.55$ ;  $M_{80\%} = 7.5\%$ ,  $t(23) = 3.34$ ,  $p = 0.003$ ,  $d_z = 0.68$ ). Thus, even when the task was designed so that spreading of feature-based attention from color to motion direction would be more or less advantageous, the magnitude of the behavioral effect was unaffected.

## Discussion

We assessed whether the enhancement of a secondary feature that overlaps spatially with an attended feature (and thus presumably belongs to the same object) results in spatially global facilitation of that feature. Participants performed a color-based selective attention task on a dot array on one side of the

visual field while at the same time monitoring for speed increases in separate dot arrays on the other side of the visual field.

Participants better detected speed changes in the dot array that matched the motion direction of the attended color array at the other location (i.e., opposite hemifield). Importantly, this was the case although motion direction was irrelevant to both tasks and never cued directly. We show that attentional enhancement of the secondary feature was not confined to the attended object, but that it was enhanced throughout the visual field.

How is processing of the secondary feature (motion direction) enhanced, given that participants were never cued to directly select it? Presumably, enhancement of the cued feature (color) occurred quickly and in a directed or voluntary manner, and only subsequently spread to the secondary feature (motion direction), likely in an incidental way, as many previous studies have shown (Ernst et al., 2013; O’Craven, et al., 1999; Schoenfeld et al., 2003; Schoenfeld et al., 2014). One general concern that pertains to all studies using this sort of design to investigate attentional spreading to a secondary feature is that, in principle, participants could attend to the secondary feature – in our case motion direction – voluntarily. We think this is unlikely for several reasons. First, the color cue appeared prior to the onset of the dot array, enabling participants to bias visual processing towards the target color in advance. In contrast, motion direction was randomly assigned to one of the colors on a trial-by-trial basis, and so participants were not able to select the relevant motion direction before the trial began, so color selection had to precede motion selection. Second, participants were asked to detect a change in dots closely related to color (luminance decrement), rather than motion. Third, the target speed change in the secondary task could occur in either motion direction (Exp. 1) or even more often in the opposite motion direction (Exp. 2), discouraging participants from paying attention to motion direction. Thus, we believe that our task design allows a distinction in how the processing of color and motion direction was enhanced.

Our findings are broadly consistent with studies demonstrating that feature-based attention increases neural responses to features at irrelevant locations (Bartsch et al., 2018; Boehler et al., 2011; Katzner et al., 2009; Lustig & Beck, 2012), and behavioral studies reporting that aftereffects (Arman,

Ciaramitaro, & Boynton, 2006; Liu & Hou, 2011; Sohn, Chong, Papathomas, & Vidnyánszky, 2005) and priming (Melcher et al., 2005) are affected by feature-based attentional spreading. Here we demonstrate that even when a feature is never cued to be selected but is part of a multi-feature object that comprises an attended feature, attention spreads to the secondary feature across the visual field, seemingly crossing object boundaries. Our findings suggest that attention is not confined to the attended object as often discussed in the literature on object-based attention (Adamian et al., 2019; Boehler et al., 2013; Egly et al., 1994; Lustig & Beck, 2012; Scholl, 2001), but instead spreads throughout feature-maps, regardless of what objects these features belong to. These findings can be understood as a combination of object-based (i.e., within-object) and feature-based (i.e., across-object) spreading of attention. However, we believe that feature-based theories alone may be able to explain these effects. For example, consistent with the ‘feature-similarity gain hypothesis’ (Martinez-Trujillo & Treue, 2004; Maunsell & Treue, 2006; Treue & Martinez-Trujillo, 1999), neurons tuned to an irrelevant feature might be enhanced as a result of their shared spatial receptive field with the attended, currently behaviorally relevant feature. Specifically, at a local scale, the attended colored dots overlapped spatially with one particular motion direction (e.g., upwards motion) in our tasks. Thus, it could be the case that the attended feature (color) is linked with the task-irrelevant feature (motion direction) through these small shared receptive fields. Under this framework, the spreading of attention from one feature to the other could be accounted for by attention at the level of independent features, with no additional assumptions about objects needed (Schoenfeld et al., 2014). Of course, such enhancement would not necessarily occur for every overlapping feature; attentional prioritization, for example, likely limits the spreading of attention between features (i.e., spreading would not occur from the unattended color to its concurrent motion direction). However, behaviorally we cannot separate attention to each distinct feature, and so further evidence, assessing attentional modulation of multiple features independently (e.g., electrophysiological recordings; Andersen, Müller, & Hillyard, 2015; Bartsch et al., 2018; Painter, Dux, Travis, & Mattingley, 2014; Störmer & Alvarez, 2014), is necessary to test this hypothesis. Such studies using neural measures could

also test whether attentional spreading across separate objects occurs when the second object (e.g., white moving dots) is entirely task-irrelevant and participants do not attend to it at all.

The present finding raises important questions about how this spreading is modulated by other task factors, such as the perceptual similarity between dot fields, perceived objecthood, or top-down control. Both feature-based (Martinez-Trujillo & Treue, 2004) and object-based (Shomstein & Berhmann, 2008) selection are known to be modulated by the similarity across features/objects, but it is an open question whether this spreading occurs across all shared features or only those that have some form of perceptual grouping (e.g., common fate through motion). However, because attentional spreading relies on different sets of objects with shared features, grouping by similar features likely contributes to the overall effect. This explanation would predict that the size of the observed effect should decrease as the number of shared features between different objects decreases (if the dots in the two tasks were different shapes, e.g., circles and triangles), the extent of which will be important for understanding the nature of attentional spreading in real-world situations. Additionally, such manipulations may provide further insights into the relationship between perceptual grouping, and feature- and object-based attentional selection. Furthermore, apparent objecthood may potentially limit the extent of location spreading. In our task, objects were separated spatially, but other studies have used boundary boxes to manipulate perceived objecthood (e.g., Egly & Driver, 1994). Interestingly, feature-spreading of directly attended features appears to occur robustly regardless of such boundary boxes (Xiao et al., 2014), in general agreement with the dominance of feature-based attentional spreading. Finally, while we show that attentional spreading persists even when it is detrimental to task performance (Exp. 2), consistent with an obligatory account of attentional spreading, it could still be the case that in other task contexts participants can exert some control over the spreading to other locations (for example, if informed verbally about the probabilities). This would be an interesting future study to inform the debate on object- and feature-based attention, as many feature-based studies have claimed obligatory spreading across locations (Serences & Boynton, 2007; Andersen et al., 2013), while spreading of attention within an object has been argued to be under some strategic control (Shomstein, 2012).

Overall, our findings suggest that attentional enhancement is not confined to features belonging to the same object but can spread to features at another location. This suggests that attention is not confined by objecthood but that features play an important role in selection independently of the objects they constitute. Critically, these features do not need to be attended primarily, but can simply ‘tag along’ an attended feature (by being part of the same object or perceptual group) to receive a global boost in visual processing. Such spreading of feature-based attention to currently task-irrelevant features may increase sensitivity to features that will potentially become relevant in the near future.

### **Acknowledgements**

Thank you to Audrey Barszcz for assistance with data collection in Experiment 2, and Tim Brady and John Serences for comments on an earlier version of this manuscript. This research was supported by a grant from the National Science Foundation (BCS-1850738), and AFC is supported by a Science and Innovation Graduate Award from Fulbright New Zealand. Chapter 1, in full, is a reprint of the material as it appears in *Psychonomic Bulletin & Review*, 28, 1252-1260, 2022, Chapman, Angus F.; Störmer, Viola S., by Springer Nature. Reproduced with permission from Springer Nature. The dissertation author was the primary investigator and author of this paper.

## Supplementary Materials

### Experiment 2

There was no difference between the two groups on their thresholded levels for the color task ( $M_{20\%} = 30.3\%$ ,  $M_{80\%} = 30.9\%$  luminance decrease),  $t(44.49) = 0.32$ ,  $p = 0.753$ ,  $d_s = 0.09$ , or motion task ( $M_{20\%} = 1.86$ ,  $M_{80\%} = 1.92$  times baseline speed),  $t(46.00) = 1.09$ ,  $p = 0.281$ ,  $d_s = 0.31$ . Overall, participants performed well on both the color task ( $M = 75.6\%$ ,  $SD = 7.4$ ) and the motion task ( $M = 75.4\%$ ,  $SD = 8.2$ ), and overall accuracy did not differ across the tasks,  $t(47) = 0.19$ ,  $p = 0.851$ ,  $d_z = 0.03$ . There was no overall difference in mean accuracy between the groups on the color task,  $t(44.43) = 1.20$ ,  $p = 0.238$ ,  $d_s = 0.35$ , or the motion task,  $t(44.53) = 0.93$ ,  $p = 0.355$ ,  $d_s = 0.27$ . Supplementary Table 1.2 gives the breakdown of task performance in Experiment 2 separated by group: in the color task, hit rates were greater in the 80% match group,  $t(44.20) = 2.10$ ,  $p = .042$ ,  $d_s = 0.60$ , however there was no difference between the groups in false alarms for both non-target changes,  $t(40.61) = 1.26$ ,  $p = .216$ ,  $d_s = 0.36$ , and no change trials,  $t(44.59) = 0.31$ ,  $p = .755$ ,  $d_s = 0.05$ . In the motion task, there was no difference between the groups on hit rates,  $t(43.93) = 0.47$ ,  $p = .639$ ,  $d_s = 0.14$ , or false alarm rates,  $t(44.59) = 1.08$ ,  $p = .284$ ,  $d_s = 0.31$ .

*Supplementary Table 1.1.* Mean (SD) performance on each task in Experiment 1. Hits are correct detections of changes, while false alarms (FA) are incorrect responses to non-changes. In the color task, false alarms could occur when the luminance decrement occurred in the non-target color, or when there was no change in either set of dots.

Color task			Motion task	
Hits	FA (non-target)	FA (no change)	Hits	FAs
62.4% (12.4)	20.3% (15.3)	10.9% (11.3)	61.5% (14.9)	12.0% (12.3)

*Supplementary Table 1.2.* Mean (SD) performance on each task for each group in Experiment 2. Hits are correct detections of changes, while false alarms (FA) are incorrect responses to non-changes. In the color task, false alarms could occur when the luminance decrement occurred in the non-target color, or when there was no change in either set of dots.

Condition	Color task			Motion task	
	Hits	FA (non-target)	FA (no change)	Hits	FAs
20% Match	61.4% (12.2)	18.0% (12.6)	7.3% (6.3)	59.7% (11.3)	11.1% (9.2)
80% Match	69.7% (15.0)	23.8% (18.5)	8.0% (7.5)	61.5% (14.1)	8.5% (7.7)



## References

- Adamian, N., Andersen, S. K., & Hillyard, S. A. (2019). Parallel attentional facilitation of features and objects in early visual cortex. *Psychophysiology*, e13498. <https://doi.org/10.1111/psyp.13498>
- Andersen, S. K., Hillyard, S. A., & Müller, M. M. (2008). Attention Facilitates Multiple Stimulus Features in Parallel in Human Visual Cortex. *Current Biology*, 18(13), 1006–1009. <https://doi.org/10.1016/j.cub.2008.06.030>
- Andersen, S. K., Hillyard, S. A., & Müller, M. M. (2013). Global Facilitation of Attended Features Is Obligatory and Restricts Divided Attention. *Journal of Neuroscience*, 33(46), 18200–18207. <https://doi.org/10.1523/JNEUROSCI.1913-13.2013>
- Andersen, S. K., Müller, M. M., & Hillyard, S. A. (2015). Attentional Selection of Feature Conjunctions Is Accomplished by Parallel and Independent Selection of Single Features. *The Journal of Neuroscience*, 35(27), 9912–9919. <https://doi.org/10.1523/JNEUROSCI.5268-14.2015>
- Arman, A. C., Ciaramitaro, V. M., & Boynton, G. M. (2006). Effects of feature-based attention on the motion aftereffect at remote locations. *Vision Research*, 46(18), 2968–2976. <https://doi.org/10.1016/j.visres.2006.03.003>
- Bartsch, M. V., Donohue, S. E., Strumpf, H., Schoenfeld, M. A., & Hopf, J.-M. (2018). Enhanced spatial focusing increases feature-based selection in unattended locations. *Scientific Reports*, 8, 16132. <https://doi.org/10.1038/s41598-018-34424-5>
- Baylis, G. C., & Driver, J. (1993). Visual attention and objects: evidence for hierarchical coding of location. *Journal of Experimental Psychology: Human Perception and Performance*, 19(3), 451–470.
- Behrmann, M., Zemel, R. S., & Mozer, M. C. (1998). Object-based attention and occlusion: evidence from normal participants and a computational model. *Journal of Experimental Psychology: Human Perception and Performance*, 24(4), 1011–1036. <https://doi.org/10.1037/0096-1523.24.4.1011>
- Boehler, C. N., Schoenfeld, M. A., Heinze, H.-J., & Hopf, J.-M. (2011). Object-based selection of irrelevant features is not confined to the attended object. *Journal of Cognitive Neuroscience*, 23(9), 2231–2239. <https://doi.org/10.1162/jocn.2010.21558>
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, 51(13), 1484–1525. <https://doi.org/10.1016/j.visres.2011.04.012>
- Chen, Z., & Cave, K. R. (2019). When is object-based attention not based on objects? *Journal of experimental psychology: human perception and performance*, 45(8), 1062–1082. <https://doi.org/10.1037/xhp0000657>
- Davis, G., & Driver, J. (1997). Spreading of visual attention to modally versus amodally completed regions. *Psychological Science*, 8(4), 275–281. <https://doi.org/10.1111/j.1467-9280.1997.tb00438.x>

- Davis, G., Driver, J., Pavani, F., & Shepherd, A. (2000). Reappraising the apparent costs of attending to two separate visual objects. *Vision research*, *40*(10-12), 1323-1332. [https://doi.org/10.1016/S0042-6989\(99\)00189-3](https://doi.org/10.1016/S0042-6989(99)00189-3)
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222. <https://doi.org/10.1146/annurev.ne.18.030195.001205>
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of experimental psychology: General*, *113*(4), 501-517. <https://doi.org/10.1037/0096-3445.113.4.501>
- Duncan, J., & Humphreys, G. W. (1989). Visual Search and Stimulus Similarity. *Psychological Review*, *96*(3), 433–458. <https://doi.org/10.1037/0033-295X.96.3.433>
- Egly, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General*, *123*(2), 161-177. <https://doi.org/10.1037/0096-3445.123.2.161>
- Ernst, Z. R., Boynton, G. M., & Jazayeri, M. (2013). The spread of attention across features of a surface. *Journal of Neurophysiology*, *110*(10), 2426–2439. <https://doi.org/10.1152/jn.00828.2012>
- Ernst, Z. R., Palmer, J., & Boynton, G. M. (2012). Dividing attention between two transparent motion surfaces results in a failure of selective attention. *Journal of Vision*, *12*(12:6), 1–17. <https://doi.org/10.1167/12.12.6>
- Katzner, S., Busse, L., & Treue, S. (2009). Attention to the color of a moving stimulus modulates motion-signal processing in macaque area MT: evidence for a unified attentional system. *Frontiers in Systems Neuroscience*, *3*, 12. <https://doi.org/10.3389/neuro.06.012.2009>
- Liu, T. & Hou, Y. (2011). Global feature-based attention to orientation. *Journal of Vision*, *11*(10:8), 1-8. <https://doi.org/10.1167/11.10.8>
- Lustig, A. G., & Beck, D. M. (2012). Task-relevant and Task-irrelevant Dimensions Are Modulated Independently at a Task-irrelevant Location. *Journal of Cognitive Neuroscience*, *24*(9), 1884–1895. [https://doi.org/10.1162/jocn\\_a\\_00249](https://doi.org/10.1162/jocn_a_00249)
- 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. <https://doi.org/10.1016/j.cub.2004.04.028>
- Maunsell, J. H. R., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, *29*(6), 317–322. <https://doi.org/10.1016/j.tins.2006.04.001>
- McLeod, P., Driver, J., & Crisp, J. (1988). Visual search for a conjunction of movement and form is parallel. *Nature*, *332*, 154–155. <https://doi.org/10.1038/332154a0>
- Melcher, D., Pappathomas, T. V., & Vidnyánszky, Z. (2005). Implicit attentional selection of bound visual

- features. *Neuron*, 46(5), 723–729. <https://doi.org/10.1016/j.neuron.2005.04.023>
- Moore, C. M., Yantis, S., & Vaughan, B. (1998). Object-based visual selection: Evidence from perceptual completion. *Psychological Science*, 9(2), 104–110. <https://doi.org/10.1111/1467-9280.00019>
- O’Craven, K. M., Downing, P. E., & Kanwisher, N. (1999). fMRI evidence for objects as the units of attentional selection. *Nature*, 401, 584–587. <https://doi.org/10.1038/44134>
- Painter, D. R., Dux, P. E., Travis, S. L., & Mattingley, J. B. (2014). Neural Responses to Target Features outside a Search Array Are Enhanced during Conjunction but Not Unique-Feature Search. *Journal of Neuroscience*, 34(9), 3390–3401. <https://doi.org/10.1523/JNEUROSCI.3630-13.2014>
- Prins, N., & Kingdom, F. A. A. (2009). *Palamedes: Matlab routines for analyzing psychophysical data*. Retrieved from <http://www.palamedestoolbox.org>
- Rossi, A. F. & Paradiso, M. A. (1995). Feature-specific effects of selective visual attention. *Vision Research*, 35(5), 621–634. [https://doi.org/10.1016/0042-6989\(94\)00156-G](https://doi.org/10.1016/0042-6989(94)00156-G)
- Sàenz, M., Buraças, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience*, 5(7), 631–632. <https://doi.org/10.1038/nn876>
- Sàenz, M., Buraças, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision Research*, 43(6), 629–637. [https://doi.org/10.1016/S0042-6989\(02\)00595-3](https://doi.org/10.1016/S0042-6989(02)00595-3)
- Schoenfeld, M. A., Hopf, J.-M., Merkel, C., Heinze, H.-J., & Hillyard, S. A. (2014). Object-based attention involves the sequential activation of feature-specific cortical modules. *Nature Neuroscience*, 17(4), 619–624. <https://doi.org/10.1038/nn.3656>
- Schoenfeld, M. A., Tempelmann, C., Martinez, A., Hopf, J.-M., Sattler, C., Heinze, H.-J., & Hillyard, S. A. (2003). Dynamics of feature binding during object-selective attention. *Proceedings of the National Academy of Sciences*, 100(20), 11806–11811. <https://doi.org/10.1073/pnas.1932820100>
- Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, 80, 1–46. [https://doi.org/10.1016/S0010-0277\(00\)00152-9](https://doi.org/10.1016/S0010-0277(00)00152-9)
- Shomstein, S. (2012). Object-based attention: strategy versus automaticity. *Wiley Interdisciplinary Reviews: Cognitive Science*, 3(2), 163–169. <https://doi.org/10.1002/wcs.1162>
- Shomstein, S., & Yantis, S. (2002). Object-based attention: Sensory modulation or priority setting? *Perception & Psychophysics*, 64(1), 41–51. <https://doi.org/10.3758/BF03194556>
- Shomstein, S., & Yantis, S. (2004). Configural and contextual prioritization in object-based attention. *Psychonomic Bulletin & Review*, 11(2), 247–253. <https://doi.org/10.3758/BF03196566>
- Serences, J. T., & Boynton, G. M. (2007). Feature-Based Attentional Modulations in the Absence of Direct Visual Stimulation. *Neuron*, 55(2), 301–312. <https://doi.org/10.1016/j.neuron.2007.06.015>

- Sohn, W., Chong, S. C., Papathomas, T. V., & Vidnyánszky, Z. (2005). Cross-feature spread of global attentional modulation in human area MT+. *NeuroReport*, *16*(12), 1389–1393. <https://doi.org/10.1097/01.wnr.0000174059.57144.62>
- Störmer, V. S., & Alvarez, G. A. (2014). Feature-based attention elicits surround suppression in feature space. *Current Biology*, *24*(17), 1985–1988. <https://doi.org/10.1016/j.cub.2014.07.030>
- Treue, S., & Martinez-Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, *399*, 575–579. <https://doi.org/10.1038/21176>
- Vecera, S. P., & Farah, M. J. (1994). Does visual attention select objects or locations? *Journal of Experimental Psychology: General*, *123*(2), 146-160.
- Wannig, A., Stanisor, L., & Roelfsema, P. R. (2011). Automatic spread of attentional response modulation along Gestalt criteria in primary visual cortex. *Nature Neuroscience*, *14*(10), 1243–1244. <https://doi.org/10.1038/nn.2910>
- Wegener, D., Ehn, F., Aurich, M. K., Galashan, F. O., & Kreiter, A. K. (2008). Feature-based attention and the suppression of non-relevant object features. *Vision Research*, *48*, 2696-2707. <https://doi.org/10.1016/j.visres.2008.08.021>
- White, A. L., & Carrasco, M. (2011). Feature-based attention involuntarily and simultaneously improves visual performance across locations. *Journal of Vision*, *11*(6), 1–10. <https://doi.org/10.1167/11.6.15>
- Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *15*(3), 419–433.
- Yeari, M., & Goldsmith, M. (2010). Is object-based attention mandatory? Strategic control over mode of attention. *Journal of Experimental Psychology: Human Perception and Performance*, *36*(3), 565-579. <https://doi.org/10.1037/a0016897>
- Xiao, G., Xu, G., Liu, X., Xu, J., Wang, F., Li, L., Itti, L., & Lu, J. (2014). Feature-based attention is independent of object appearance. *Journal of Vision*, *14*(3), 1-11. <https://doi.org/10.1167/14.1.3>

CHAPTER 2: Efficient tuning of attention to narrow and broad ranges of task-relevant feature values

## Abstract

Feature-based attention is the ability to select relevant information on the basis of visual features, such as a particular color or motion direction. In contrast to spatial attention, where the attentional focus has been shown to be flexibly adjustable to select small or large regions, it is unclear whether feature-based attention can be efficiently tuned to different feature ranges. Here, we establish that the focus of feature-based attention can be adjusted more broadly or narrowly to select currently relevant features. Participants attended to a set of target-colored dots among distractor dots to detect brief decreases in luminance (Experiments 1a, 1b, 2) or bursts of coherent motion (Experiments 3a, 3b, 4). To vary the size of the attentional focus, we manipulated the range of colors that the target dots spanned and found that while participants' performance decreased with larger feature ranges to select, it remained at a relatively high level even at the largest color range, suggesting that broadening the focus of feature-based attention comes only at a small cost and that large feature ranges can be selected relatively efficiently at once. Overall, our findings argue against the idea that feature-based attention is limited to a single feature value at a time and demonstrate that selecting large swaths of feature space is surprisingly efficient. Broadly, these results are consistent with accounts that propose a flexible and generalized set of attentional mechanisms that act across both spatial and feature-based domains.

## Introduction

When navigating and interacting with the world around us, sensory information that is highly salient or relevant to our goals is prioritized by attention (Carrasco, 2011; Desimone & Duncan, 1995; Reynolds & Chelazzi, 2004). Seminal theories have attempted to explain how information processing is shaped by attention, and the resulting perceptual and behavioral outcomes. For example, attention to a particular location in space increases neural processing in cortical regions selective for that location, and improves perception of items that appear there (Kastner, Pinsk, De Weerd, Desimone, & Ungerleider, 1999; Pestilli & Carrasco, 2005; Posner, 1980). Likewise, attention can select information based on visual features, such as a particular color or motion direction (Andersen, Hillyard, & Müller, 2008; Sàenz, Buraças, & Boynton, 2003). Several studies have shown that feature-based attention enhances processing of the selected feature (Liu, Larsson, & Carrasco, 2007; Martinez-Trujillo & Treue, 2004; Sàenz et al., 2003; Störmer & Alvarez, 2014; White & Carrasco, 2011), similar to what has been shown for spatial attention. While spatial attention has received a particularly prominent focus in the literature (Eriksen & Yeh, 1985; Itti & Koch, 2000; Lamy & Tsal, 2001; Posner, Snyder, & Davidson, 1980; Treisman & Gelade, 1980; Wolfe, 1994)—often treated as separate and distinct from feature-based attention—it is undetermined whether spatial and feature-based attention rely on largely separate processes or may instead share common processes that act on different mental representations and neural substrates. As such, understanding under what circumstances attentional selection is comparable or different across these modes of attention is necessary for establishing a more general theory of attentional selection.

In research on spatial attention, significant effort has been dedicated to understanding the structure and size of the attentional focus. Many influential models of spatial attention, such as the spotlight (Posner, 1980; Posner et al., 1980), gradient (Downing & Pinker, 1985), zoom lens (Eriksen & St. James, 1986; Eriksen & Yeh, 1985), and normalization model of attention (Reynolds & Heeger, 2009), propose different characteristics of the attentional focus. For example, some have hypothesized that spatial attention exhibits an excitatory peak at the attended location that gradually falls off with increasing distance (e.g., Dori & Henik, 2006; Downing & Pinker, 1985; LaBerge & Brown, 1989;

Shulman, Sheehy, & Wilson, 1986), while others have proposed a center-surround profile where locations nearby an attended region are inhibited (Cutzu & Tsotsos, 2003; Hopf et al., 2006; Mounts, 2000; Müller, Mollenhauer, Rösler, & Kleinschmidt, 2005). Another important property of the attentional focus is its size: whether it can be adjusted flexibly or is fixed at a particular spatial scale. The zoom-lens model, initially conceived of by Eriksen and St. James (1986), proposes that the size of the attention field can be changed (i.e., one can “zoom in” or “zoom out” to attend to small or big regions in the visual field), but that these changes come at a cost, such that processing efficiency decreases as the size of the attentional focus increases. The more recently developed normalization model of attention also suggests that the attentional field size is a critical component to understand selection (Reynolds & Heeger, 2009). For example, it has been shown that variations in the effects of attention can be explained by different attentional field sizes (Herrmann, Montaser-Kouhsari, Carrasco, & Heeger, 2010; Itthipuripat, Garcia, Rungratsameetaweemana, Sprague, & Serences, 2014). In sum, seminal models of attention—and spatial attention in particular—have made explicit assumptions about the flexibility of the attentional focus and how this relates to the pattern of attentional modulations and task performance.

While ample research has investigated the limits of spatial attention, the selection limits of feature-based attention are relatively poorly understood. In particular, it is unknown whether attention to features is limited to select single, individual features (i.e., the color red, the motion direction upwards), or whether the focus of feature-based attention, like spatial attention, can be flexibly adjusted, such as by tuning more broadly to select multiple feature values (e.g., colors ranging from red to purple). Alternatively, it might be the case that spatial attention is unique in its ability to change the scale over which it operates. Research investigating the limits of feature-based attention have predominantly focused on examining whether multiple distinct features can be selected at once and have suggested a stark and fixed limit of feature-based attention. For example, Boolean Map theory (Huang & Pashler, 2007) proposes that attentional selection requires the division of the visual field into regions based on whether a given feature (e.g., the color red) is present or absent at that location, and argues that only a single feature value can be used for selection at a time, while multiple locations can be selected simultaneously.



Consistent with this account, studies have shown that cuing two directions of motion (Liu, Becker, & Jigo, 2013) or colors (Liu & Jigo, 2017) benefits perceptual processing less than cuing to a single feature, suggesting a stark selection limit. Furthermore, Liu and Jigo (2017) found evidence to suggest that any benefits in the two-cue condition could be explained by attention towards a single feature, and proposed that attention can only select a single feature value.

However, previous studies have mostly assessed situations in which feature-based attention is divided between distinct feature values (e.g., the colors red and blue; Andersen, Hillyard, & Müller, 2013; but see Störmer & Alvarez, 2014), but an unanswered question is whether attention can be effectively distributed across a range of feature values (e.g., colors ranging from red to orange); in other words, can feature-based attention change its scope to select broader (or narrower) parts in feature space, similar to how spatial attention can increase its attentional field size to select larger (or smaller) spatial regions? A few studies indirectly speak to this: Herrmann and colleagues (2012) manipulated the uncertainty of a precue in an orientation discrimination task, and found that when precues were highly uncertain (covering 60° of orientation), the behavioral benefit of valid cues was equivalent regardless of the actual orientation of the target stimulus, providing suggestive evidence that participants were distributing their attention across the entire range of potential target orientations during the cue-target interval. Another study using a similar approach found that sensitivity to a motion target decreased as a function of the reliability of the cued motion direction (Ball & Sekuler, 1981), suggesting that spreading attention to larger ranges comes at a behavioral cost. While these studies provide some support for the idea that feature-based attention—like spatial attention—may be relatively flexible in its focus, their manipulation of cue reliability does not require that participants maintain attention to a wide range of features at once, and ultimately the target feature selected was only a single value. Thus, current research has not directly assessed whether and how multiple feature values can be selected continuously.

If feature-based attention can be flexibly adjusted to select more broad or narrow ranges of feature values, then many models, based predominantly on explaining spatial attention, could incorporate this shared aspect. Some models of attention already assume similar principles for selecting locations and

features (Reynolds & Heeger, 2009), predicting that the flexibility of spatial attention should extend to the selection of visual features. Some have even construed spatial dimensions as simply being “features” of the visual world in the same way that colors or orientations are (Maunsell & Treue, 2006), implying that the traditional division between the two modes of attention are unnecessary and that selection for features and spatial locations should follow analogous principles. In this case, the size of the attentional field should not only be able to adjust in location space, as previous work has shown (Castiello & Umiltà, 1990), but should also adjust in feature space, and be thus able to select a range of features with reasonable efficiency. Although there are many differences between spatial and feature-based attention (e.g., Ling, Liu, & Carrasco, 2009; Liu, Stevens, & Carrasco, 2007), a major goal of attention research is to characterize general selection mechanisms that bridge across stimulus spaces. Determining the flexibility of the attentional focus to features therefore can help us understand to what extent spatial and feature-based attention are largely supported by distinct processes or share common mechanisms for selection.

Here, across five experiments, we examine whether feature-based attention can be flexibly tuned to narrow and broad ranges of colors. We adapted a sustained feature-based attention task (Andersen et al., 2008; Sàenz et al., 2003; Störmer & Alvarez, 2014), where participants attend to a set of target colored dots amongst differently colored distractor dots to detect brief target events (luminance decreases or coherent motion). To assess the efficiency of narrowly and more broadly tuned attention, we varied the range of colors that the target dots spanned from trial to trial, allowing us to measure performance as a function of the distribution of target colors attended. Our findings demonstrate that increasing the size of the attentional focus in feature space results in only a relatively small decrease in processing efficiency, similar to studies of spatial attention. In particular, Experiment 1a (which used a fixed-luminance circular color space) and Experiment 1b (which used a fixed-luminance 2D color space) show that when attention is tuned broadly to a large range of color space, performance decreases only slightly compared to a narrow focus. Experiment 2 shows that the decrease in performance is much lower than would be expected if participants were attending to only a narrow, fixed range of colors. In Experiment 3a and 3b,

we confirm that this effect is driven by early attentional selection of colors, rather than later decision-related processes by assessing performance on both target and distractor-colored dots. Finally, Experiment 4 demonstrates that attention was allocated across the entire range of colors relatively uniformly, as participants were just as good at detecting changes at the edges of the color range as those in the center of the range. Together, these experiments demonstrate that feature-based attention can be efficiently adjusted as necessitated by task conditions.

## **Experiment 1a**

### **Method**

#### **Preregistration**

We preregistered the predictions and analysis of this experiment on AsPredicted, with a predetermined minimum sample size of 30 (<http://aspredicted.org/blind.php?x=mj2dh7>), providing 80% power to detect a significant effect of  $\eta_p^2 > .308$ .

#### **Open Data**

Raw data and analysis code for all experiments are available on the Open Science Framework (<https://osf.io/uf4k6/>).

#### **Participants**

Thirty-six undergraduate students from the University of California, San Diego subject pool participated for course credit and gave written informed consent prior to starting the experiment as approved by the Institutional Review Board at UC San Diego. Six participants with an average  $d' < 0.5$  across all conditions in the main task were excluded. The remaining 30 participants (24 women, 6 men) were between 18-25 years of age ( $M = 19.9 \pm 1.7$  years) and had normal or corrected-to-normal vision.

#### **Stimuli**

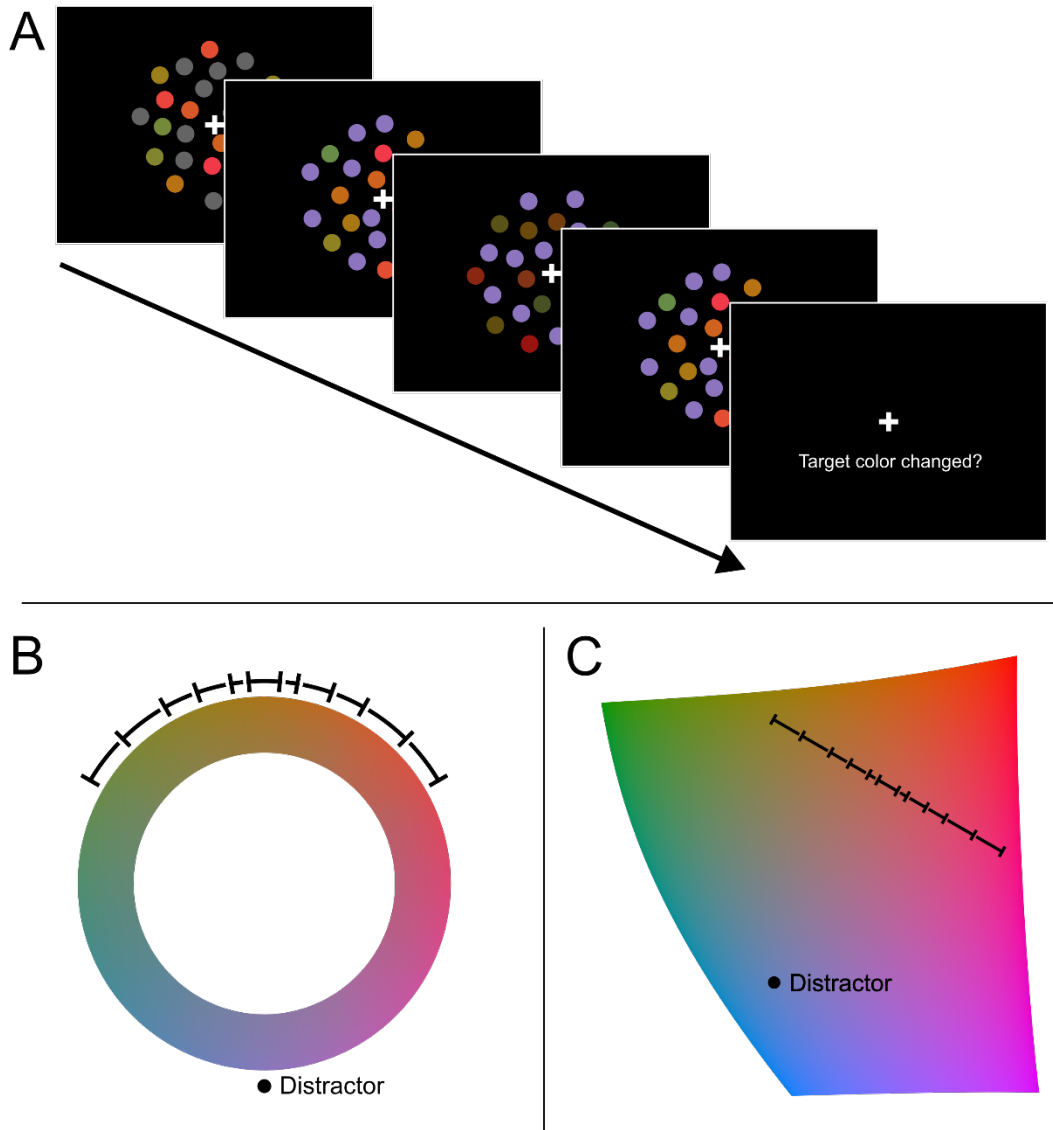
The experiment was presented on a 22" CRT monitor that was calibrated to linearize RGB output. Participants were seated at approximately 57" from the monitor during the experiments. All stimuli were generated using MATLAB (R2016b) with the Psychophysics Toolbox (Brainard, 1997). A centrally

presented circular field of dots ( $5.8^\circ$  visual angle radius) was presented on a black background. This field contained 200 dots moving independently and randomly at  $2.25^\circ/\text{s}$ . To prevent participants from tracking single dots, each dot had a limited lifetime and was redrawn at a new random location every 300 ms. A white fixation cross (each bar  $0.4^\circ$  long) was presented in the center of the dot field.

The dot field was separated into two groups: a set of target dots and a set of distractor dots. The target and distractor contained 100 dots each and differed in color. All colors were selected from a set of 360 equally spaced equiluminant colors in the CIE Lab color space, drawn from a circle with radius 49 units, centered at  $L = 54$ ,  $a = 21.5$ ,  $b = 11.5$  (see *Figure 2.1B*). On each trial, the “center” of the target colors was randomly selected from the color wheel and the distractor was set  $180^\circ$  away, on the opposite side of the color wheel. Across trials, we manipulated the range that the target colors varied in by uniformly drawing colors symmetrically around the “center” target color, such that the width of this range spanned either  $10^\circ$ ,  $20^\circ$ ,  $40^\circ$ ,  $60^\circ$ ,  $90^\circ$ , or  $120^\circ$  around the color wheel (*Figure 2.1B*).

### **Procedure**

In the task, participants attended to a set of target-colored dots to detect brief decreases in luminance (*Figure 2.1A*). On each trial, the dot field was presented for 800 ms during a cueing period, in which the target dots were presented in color, while the distractor dots were presented in a neutral grey (RGB: [100, 100, 100]). This allowed us to present the exact target colors to participants without having to use specific cues or labels to differentiate targets and distractors. After the cueing period, target and distractor dots (now both colored) were presented for 2000 ms, during which participants attended the target dots to detect brief (300 ms) decreases in luminance. The luminance decrease occurred at a random time with the constraint that it could not happen in the 300 ms period at the beginning or end of each trial. At the end of each trial, participants indicated whether this change occurred in the target dots by responding on the keyboard (‘m’ for a target change, ‘n’ for no target change). The luminance change could occur in the target dots (50% of trials), the distractor dots (25%), or neither set of dots (25%). Sensitivity ( $d'$ ) was calculated as the normalized hit rate minus normalized false alarm rate (across distractor change and no-change trials). Participants completed 288 trials of this task (48 per condition).



*Figure 2.1.* A) Example trial structure for Experiment 1a & 1b. During the pre-trial period (800 ms), only the target dots were presented in color (in this example spanning  $120^\circ$  of the color space), while the distractor dots were shown in grey. Participants then had to attend to the target dots among colored distractor dots throughout the trial to detect a brief (300 ms) decrease in luminance in the target colors. Displays are not shown to scale. B) Circular color space used in Experiment 1a. Target colors were centered at a random position around the space on each trial and spanned 6 ranges as shown by the arc segments ( $10^\circ$ ,  $20^\circ$ ,  $40^\circ$ ,  $60^\circ$ ,  $90^\circ$ , or  $120^\circ$ ). C) 2D color space used in Experiment 1b. Target colors were randomly chosen on each trial and spanned 6 ranges along a line as shown by line segments. The distractor color was at a fixed relative position from this line.

The magnitude of the luminance decrease was determined for each participant individually at the beginning of the experiment session by a thresholding task. Participants completed 32 trials per

thresholding run, in which the luminance decrease was adjusted using a staircasing method: the change became smaller (less detectable) after two consecutive correct responses, and larger (more detectable) after an incorrect response. The luminance decrement was initially set at 40% of the maximum luminance of the dots and was adjusted by 2% with each step. During the thresholding task, all target and distractor dots were colors 180° away from each other on the color wheel (i.e., single colors, no range of colors), and these colors varied randomly from trial-to-trial. Accuracy was fit with a logistic curve using the Palamedes toolbox (Prins & Kingdom, 2009) with a guess rate of 50%, and thresholds were selected as the luminance decrement corresponding to 70% accuracy. Participants completed 1-3 runs of the thresholding task until performance was adequately estimated ( $M = 1.87$  runs,  $SD = 0.73$ ).

At the end of the experiment, participants also completed a triad color similarity task, the data from which we do not include here.

## Results

Sensitivity ( $d'$ ) to detect the luminance change was computed separately for all color range conditions. To assess how the range of colors present in the target dots affected performance in the task, we conducted a repeated-measures ANOVA on  $d'$  across target color range (10°, 20°, 40°, 60°, 90°, or 120° around the color wheel). There was a main effect of target color range,  $F(5,145) = 3.42$ ,  $p = .006$ ,  $\eta_p^2 = .106$ , such that  $d'$  decreased as the range of the target colors increased, further evidenced by a significant linear trend,  $F(1,29) = 15.31$ ,  $p < .001$ ,  $\eta_p^2 = .346$ , as can be seen in *Figure 2.2A*. Planned follow-up comparisons (FDR corrected) revealed that this effect was driven by higher  $d'$  at the lowest levels of color range (10° & 20°) relative to the highest levels of color range (90° & 120°). Performance at 10° was marginally better than at 90°,  $p = .052$ , and 120°,  $p = .052$ , but not different from 20°, 40°, or 60°,  $ps > .17$ . Performance at 20° was significantly better than at 90°,  $p = .026$ , and 120°,  $p = .005$ , marginally better than 40°,  $p = .064$ , but not different from 60°,  $p = .135$ . There were no differences in performance between the higher levels of target color range,  $ps > .45$ . Changes in  $d'$  appeared to primarily




Figure 2.2. A) Results of Experiment 1a. Detection of target luminance changes ( $d'$ ) decreased as the range of target colors increased as a function of the angular distance around the color wheel. B) Results of Experiment 1b. Discrimination decreased as the range of target colors increase as a function of distance in the 2D CIELab color space. Error bars correspond to within-subject SEM.

be driven by a reduction in hit rates (correct detection of target changes) as color range increased,

$F(5,145) = 5.75, p < .001, \eta_p^2 = .165$ , as there was no significant effects on false alarms to distractor

changes,  $F(5,145) = 0.40, p = .849, \eta_p^2 = .014$ , or no-change trials,  $F(5,145) = 0.46, p = .807, \eta_p^2 = .016$ .

These changes are reflected in a significant increase in (i.e., more conservative) response criterion with higher target color ranges,  $F(5,145) = 2.34, p = .044, \eta_p^2 = .075$ .

This experiment demonstrates that increasing the size of the focus of feature-based attention leads to decreased efficiency in selecting target features, consistent with findings from spatial attention (Castiello & Umiltà, 1990; Maringelli & Umiltà, 1998). The observed decrease in performance, from  $d' = 1.76$  at  $10^\circ$  of target color range to  $d' = 1.44$  at  $120^\circ$ , is smaller than might be anticipated given the twelve-fold increase in the range of target colors to be attended. Specifically, assuming a strict capacity model where only a single color is selected at a time, there should be a steep, linear drop off as the range of target colors to be attended increases. Instead, our findings show that attentional resources are not strictly fixed but flexibly adapt and can be allocated almost equally across large parts of the feature space.

## Experiment 1b

Experiment 1a suggests that participants can focus feature-based attention broadly across a range of colors quite efficiently, with relatively small costs in performance. However, one potential alternative explanation for these findings is that the decrease in performance we found is driven not just by the range of target colors, but by the correlated change in target-distractor similarity with different ranges, which affects performance in this type of feature-based attention task (Chapman & Störmer, 2022). Because we used a circular color space, increasing the range of to-be-attended colors necessarily increases the similarity between targets and distractors (as can be seen in *Figure 2.1B*). At the lowest level of target color range, the distractor color was  $175^\circ$  away from each end of the target distribution, while at the highest level the distractor color was  $120^\circ$  away. To rule out the alternative that (at least some of) the decrease in performance was due to changes in target-distractor similarity across conditions, we conducted a second experiment that ensured target color range was not confounded with target-distractor distance. To do this, in Experiment 1b, we chose target and distractor colors from a two-dimensional plane in the CIELab color space (*Figure 2.1C*). Target colors were selected along a line drawn in this space at a fixed distance from a distractor color; thus, increased color range here actually *decreases* the average distance between target and distractor colors (see *Figure 2.1C*). It is also worth noting that distractors now differed from targets in both hue and saturation. While this might complicate a direct comparison between Experiments 1a and 1b, because distractors are more distinct from targets in two feature dimensions, this experiment provides a strong test of the hypothesis that a broad range of feature values can be selected by attention, regardless of distractor similarity.

## Method

### Preregistration

Just like for Experiment 1a, we preregistered the predictions and analysis of this experiment on AsPredicted, with a predetermined minimum sample size of 30



(<http://aspredicted.org/blind.php?x=24r4aw>), providing 80% power to detect a significant effect of  $\eta_p^2 > .308$ .

## **Participants**

Thirty-nine undergraduate students participated in this experiment and gave written informed consent as approved by the Institutional Review Board at UC San Diego. Nine participants with an average  $d' < 0.5$  across conditions in the main task were excluded. The remaining 30 participants (21 women, 9 men) were between 18-25 years of age ( $M = 20.4 \pm 1.9$  years) and had normal or corrected-to-normal vision.

## **Stimuli**

Target and distractor colors were selected from a fixed-luminance 2D-plane in the CIELab color space. For a particular luminance level ( $L = 54$ ), we generated all colors within the sRGB gamut at steps of 0.25 units along the a and b dimensions. A ‘d50’ white point was assumed for converting colors from Lab to RGB values. We selected points in this space such that the distance between target and distractor colors, and between points along the variable target line, were comparable in distance to those selected from the color wheel in Experiment 1a. To generate target and distractor colors, we sampled three points in the color space that formed an equilateral triangle with sides 85 units long. Two points were randomly selected as the maximal ends of the target color line and the remaining point was the distractor color. The midpoint of the target line was the “average” target color, and was restricted to be at least 15 units from the grey point of the color space ( $Lab = [54, 0, 0]$ ) to avoid target colors that were highly desaturated. Target color range was manipulated by uniformly selecting colors that spanned particular distances along the target line (widths of 7.1, 14.2, 28.3, 42.5, 63.8, and 85 units; these distances correspond to the length of the chords on the color wheel for the angles used in Experiment 1a). This process was repeated until all target and distractor colors fell within the sRGB gamut.

## **Procedure**

The details of the procedure are similar to Experiment 1a. Participants first completed a thresholding task using colors drawn from the 2D color space (target color range = 0; 1-3 runs,  $M = 1.60$ ,

$SD = 0.77$ ). They then completed 336 trials of the main task. Because the colors were slightly less saturated overall, due to the way we selected points in the color space, we lowered the brightness of the baseline grey dots (RGB: [80, 80, 80]). All other parameters were the same.

## Results

As in Experiment 1a, we conducted a repeated-measures ANOVA on  $d'$  across target color ranges. There was a main effect of target color range,  $F(5,145) = 5.03, p < .001, \eta_p^2 = .148$ , such that  $d'$  decreased as the target color range increased, as shown in *Figure 2.2B*. Similar to the previous experiment, there was a highly significant linear trend,  $F(1,29) = 17.08, p < .001, \eta_p^2 = .371$ , and planned follow-up comparisons (FDR corrected) revealed that  $d'$  was higher at the lowest levels of color range (7.1 and 14.2 units) relative to the highest levels of color range (42.5, 63.8, and 85 units),  $ps < .05$ . Performance at 28.3 units did not differ from any the other levels,  $ps > .063$ , nor was there a difference between 7.1 and 14.2 units,  $p = .888$ , and performance at the highest levels of target color range (42.5, 63.8, and 85 units) did not differ amongst themselves,  $ps > .49$ . As for Experiment 1a, changes in  $d'$  appeared to primarily be driven by a reduction in hit rates as target color range increased,  $F(5,145) = 8.90, p < .001, \eta_p^2 = .235$ , as there was no significant effects on false alarms to distractors,  $F(5,145) = 1.28, p = .275, \eta_p^2 = .042$ , or no-change trials,  $F(5,145) = 0.23, p = .948, \eta_p^2 = .008$ . As in Experiment 1a, these changes are reflected in a significant increase in response criterion with higher target color ranges,  $F(5,145) = 2.31, p = .047, \eta_p^2 = .074$ .

This experiment confirmed that increasing the range of to-be-attended target colors results in decreased performance, even when target-distractor similarity does not concurrently increase. Importantly, the overall decrease in performance for higher ranges of colors is likely not due to factors introduced by the 2D color space. While target colors were less saturated than in Experiment 1b, because the 2D space constrained how the target colors could be chosen, colors at the edge of high range targets were further from the distractor and higher saturation on average compared to low target color ranges,

demonstrating that attentional selection was similar despite differences between the circular and full 2D CIELab color spaces. However, although the pattern of results was similar across these two experiments, targets and distractors could differ across two dimensions in Experiment 1b (hue and saturation) which may limit the direct comparability of these findings.

## **Experiment 2**

In Experiment 1a and 1b, we demonstrated that participants could attend to a broad range of colors at a relatively low cost. These results appear inconsistent with accounts that assume a strict capacity limit of feature-based attention, for example those that propose only a single feature value (or a very small range of feature values) can be selected at once (e.g., Huang & Pashler, 2007). According to such accounts, the decrease in performance should be inversely proportional to the number of features attended—and thus much more pronounced than the performance decrease we observed. To directly quantify how strongly performance would decrease if participants only selected a small and fixed feature range across all conditions, in Experiment 2, we directly manipulated the number of target dots and used the observed decrease in performance to predict the range of attended colors in Experiment 1b.

## **Method**

### **Preregistration**

We preregistered the predictions and analysis of this experiment on AsPredicted, with a predetermined minimum sample size of 20 ([https://aspredicted.org/Q9S\\_CDK](https://aspredicted.org/Q9S_CDK)). Based on pilot experiments, we anticipated a larger effect than in Experiments 1a and 1b, so we estimated fewer participants were needed to detect a significant effect.

### **Participants**

Twenty-four undergraduate students participated in the experiment for course credit. Data from four participants was excluded from the final data set, as their average  $d'$  in the main task was below 0.5.

The final 20 participants (13 women, 7 men) were between 18-23 years of age ( $M = 19.8 \pm 1.2$  years), and had normal or corrected-to-normal color vision.

### **Stimuli & Procedure**

The task proceeded similarly to Experiment 1b, however on each trial the target and distractor colors were fixed at a single value. Colors were selected from the 2D CIELab space, as in Experiment 1b, assuming a color range of 85 units, however the target color was always selected as the midpoint of this range. This ensured that target and distractor colors would be as similar as possible to those in Experiment 1b, to enable comparison of performance across the tasks.

We manipulated the number of target dots in the display across trials. In our baseline condition, equivalent to previous experiments, there were 100 target dots and 100 distractor dots. We manipulated the proportion of target dots relative to this baseline across 6 conditions (100%, 84%, 68%, 52%, 35%, 20%). When the number of target dots was reduced, the number of distractor dots was simultaneously increased to maintain a display of 200 dots. To allow for comparison across conditions, when the distractor dots changed in luminance this occurred in only 100 dots. Participants completed 336 trials of the main task, preceded by a luminance thresholding task with equal numbers of target and distractor dots (equivalent to thresholding in Experiment 1b).

### **Data Analysis**

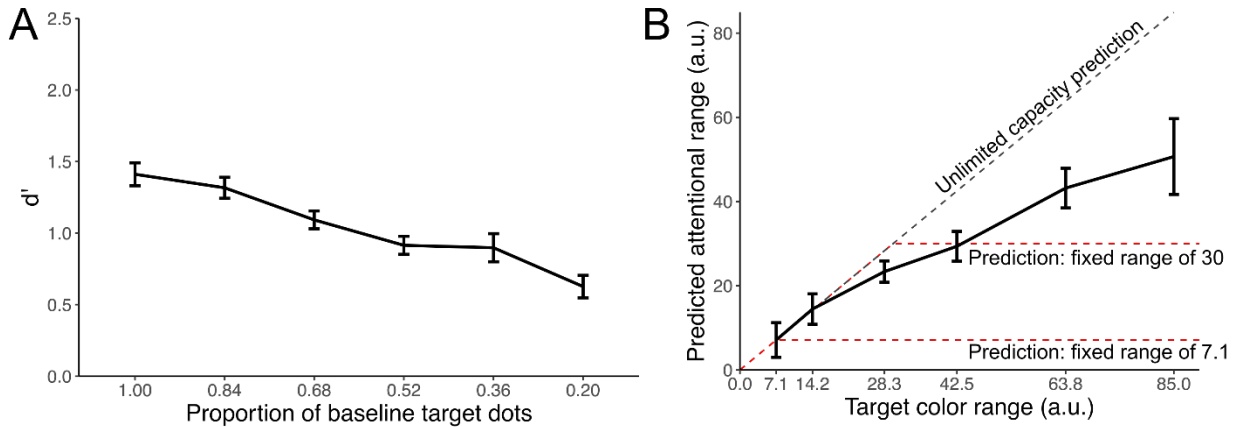
We used this data to predict the range of attended colors in Experiment 1b. To do this, we first extracted the slope of the relationship between the number of target dots in the display and  $d'$  from a linear mixed-effects model with 'lmer' from the R package *lme4* (v1.1-26; Bates, Maechler, Bolker, & Walker, 2015). We then inverted this relationship, estimating the change in the proportion of attended target dots based on performance in Experiment 2. To allow an estimate of the attended color range from this relative proportion estimate, we assumed that all dots were attended with a color range of 7.1 units (the lowest target color range used in Experiment 1b). Finally, we multiplied the predicted proportion of

attended target dots by the width of the target color range, resulting in a measure of the predicted attentional range.

## Results

We first examined the effect of the number of target dots with a repeated-measures ANOVA, which revealed a significant main effect on performance,  $F(5,95) = 14.33, p < .001, \eta_p^2 = 0.430$ . Similar to previous experiments, there was a significant linear trend downwards in performance as the number of target dots decreased,  $F(1,19) = 51.70, p < .001, \eta_p^2 = .731$ . This strong linear effect can be seen in *Figure 2.3A*. As in previous experiments, this change in  $d'$  was driven by lower hit rates as the proportion of target dots decreased,  $F(5,95) = 20.38, p < 0.001, \eta_p^2 = 0.517$ , while there was no effect on false alarms to distractors,  $F(5,95) = 1.05, p = .394, \eta_p^2 = .052$ , or no-change trials,  $F(5,95) = 0.19, p = .964, \eta_p^2 = .010$ .

This data suggests that the results from Experiment 1b cannot be explained by a simple “subsampling” strategy: the target color range increases by 12 times across conditions in Experiment 1b (7.1 to 85 units), while the proportion of target dots in Experiment 2 decreases by only 5 times (100% to 20% of baseline target dots), yet the decrease in performance in Experiment 2 is much greater. To quantify this difference, we used the performance in Experiment 2 to estimate how broadly participants tuned their attention in Experiment 1b. If participants attend to a strictly fixed range of colors, then increasing the target color range (as in Experiment 1b) should be equivalent to decreasing the number of target dots on the display (as in Experiment 2). Thus, we can predict the range of attended colors in Experiment 1b based on how much performance drops in each condition, relative to the observed changes in Experiment 2. As demonstrated in *Figure 2.3B*, if attention was unlimited to any range of colors, the predicted range would follow the diagonal, while if there was a fixed attentional focus, the predicted range would flatten out once this fixed range was reached. However, we find an intermediate effect, revealing that attention cannot select broader ranges of colors without limits, but that attention is also clearly not fixed or particularly narrow in its scope. This analysis assumes that attention inside a fixed



*Figure 2.3.* A) Results of Experiment 2. Detection of target luminance changes ( $d'$ ) decreased with fewer target dots to attend to. B) Predicted range of feature-based attention in Experiment 1b as a function of the true target color range. Predicted attentional range was based on the slope of the performance in Experiment 1b compared to the slope observed in Experiment 2. Error bars correspond to within-subject SEM.

range would be all-or-none, and changes in the target color range would map directly onto adding or removing target dots from the display—which is likely not true in practice. Regardless, the much larger decrease in performance in Experiment 2 clearly demonstrates that participants did not attend to a fixed, narrow range of target colors in Experiments 1a and 1b; instead, attentional resources were flexibly allocated across the different color ranges following task demands.

### Experiment 3a

In Experiment 1a and 1b, we demonstrated that increasing the range of to-be-attended colors resulted in a decrease in sensitivity for detecting changes in the luminance of target dots. However, it is possible that participants were not attending to only the target colors, but instead distributed their attention across target and distractor colors to detect *any* luminance change, and only at the end of the trial, when making a response, determined what kind of change it was (target, distractor, or no-change). In this case, decreased performance at larger target color ranges might then be attributable to greater confusability between target and distractor colors. We think such a strategy is unlikely for several reasons. First, we thresholded participants to a performance level where, in our view, it would be difficult

to detect luminance changes without focusing attention onto the target colors. Second, across both experiments, we found that increases in the target color range exclusively affected hit rates (detection of changes in the targets) and not false alarm rates (false attribution of distractor changes as target changes). Third, we found a similar pattern of performance in Experiment 1b, despite increases in the target color range increasing the separability of target and distractor colors.

Nonetheless, we sought to get direct evidence to argue against this alternate account. In Experiment 3, participants performed a task similar to Experiments 1a and 1b, but this time were cued to prioritize the target color to discriminate the direction of a brief interval of coherent motion (either left, right, up, or down), which was more likely to occur in the target dots (80% of changes) relative to distractor dots. However, participants were asked to report the motion direction also when they occurred in the distractor-colored dots. If performance remains higher for target than distractor dots, while still decreasing as a function of the range of target colors attended, this suggests that early attentional selection of colors, and not a decision-making process, is responsible for the effects we observed in Experiment 1.

## **Method**

### **Participants**

Sixty-two undergraduate students at UC San Diego participated in this experiment for course credit. Sixteen participants were excluded with accuracy below 60% on attention check trials (motion coherence: 100%; target color range: 0°; see *Procedure*). The remaining 46 participants (31 women, 15 men) were between 18-29 years of age ( $M = 21.0 \pm 2.4$  years). All participants gave informed consent prior to starting the experiment as approved by the Institutional Review Board at UC San Diego. This sample size provides 80% power to detect a significant main effect of target color range and interaction with coherent motion color type of  $\eta_p^2 > .196$ .

### **Stimuli**

This experiment was conducted online, and participants completed the study in a web browser on devices they provided. Stimulus presentation for the experiment was managed using jsPsych (de Leeuw,

2015) with customized plugins based on previous RDK code (Rajananda, Lau, & Odegaard, 2018). A central, circular RDK (400 by 400 px) was presented on a black background and consisted of two sets of 100 independently moving dots (radius 5 px). Each dot had a limited lifetime and was redrawn at a new random location every 500 ms. A white fixation cross (each bar 20 px by 1px) was presented in the center of the dot field. Target and distractor dot colors were selected from a circular CIELab color space and colors were determined randomly on each trial as in Experiment 1a. The range of target colors was manipulated across trials, spanning either 20°, 40°, 60°, or 120° around the color wheel, similar to Experiment 1a.

### **Procedure**

Trials proceeded similar to Experiment 1a, but with a change from detection of luminance decreases to periods of coherent motion. First, for 1200 ms, target dots were presented in color while distractor dots were presented in grey ([RGB: 100, 100, 100]). Following this, target and distractor dots were presented in color for 2000 ms, and participants had to attend both sets of dots for the duration of the trial to detect a brief period of coherent motion (300 ms, 40% coherence) in one of the cardinal directions (left, right, up, or down) that occurred in either the target or distractor dots. The coherent motion could occur at any time within the window 300-1800 ms after the main trial period began. At the end of the trial, participants had to report the correct direction using the arrow keys regardless of which dots moved coherently. Target dots moved coherently on 80% of trials, and participants were instructed to attend to them primarily, while still reporting the direction of the distractor dots if they moved coherently. Participants completed 504 trials of this task, of which 24 had 100% coherent motion in the target dots (which were presented in a single color) and were considered “attention check” trials, to determine whether participants were performing the task as intended. These trials were only used to exclude participants from the main analysis (see *Participants*) and were excluded from the main data analysis. The remaining 480 trials were distributed evenly across the four target color range conditions (20°, 40°, 60°, or 120°) and coherent motion color type (target or distractor dots, 80%/20% respectively).



## Results

To assess the effect of target color range and attention (coherent motion in attended vs. unattended color) on performance, we performed a repeated-measures ANOVA on participants' accuracy in reporting the direction of the coherent motion. This revealed a main effect of attention,  $F(1,45) = 65.65, p < .001, \eta_p^2 = .593$ , such that participants were on average more accurate at reporting coherent motion in the target dots than the distractor dots. There was also a main effect of target color range,  $F(3,135) = 3.11, p = .028, \eta_p^2 = .065$ , such that accuracy decreased as the range of target colors increased, replicating the findings from Experiment 1a and 1b. There was no interaction between these factors,  $F(3,135) = 0.09, p = .967, \eta_p^2 = .002$ , however when assessing the effect of target color range for each set of dots separately, we did find that the range of target colors had a detrimental effect on accuracy for coherent motion in the target dots,  $F(3,135) = 5.87, p < .001, \eta_p^2 = .115$ , but not distractor dots,  $F(3,135) = 1.00, p = .354, \eta_p^2 = .022$ . Follow-up pairwise comparisons (FDR corrected) revealed that accuracy was generally lower for coherent motion in the target dots at  $120^\circ$  compared to smaller target color ranges ( $20^\circ$  vs  $120^\circ, p < .001$ ;  $40^\circ$  vs  $120^\circ, p = .061$ ;  $60^\circ$  vs  $120^\circ, p = .003$ ), while there was no difference in accuracy among the lower target color ranges ( $\leq 60^\circ; ps > .2$ ). Additionally, there were no differences in accuracy for coherent motion in the distractor dots for any levels of target color range (all  $ps > .4$ ). For all levels of target color range, coherent motion was detected better in the target dots compared to distractor dots ( $ps < .001$ ). These patterns are summarized in *Figure 2.4A*. These results indicate that the effects of increasing the range of target colors are driven by attentional selection during perceptual processing, as participants were worse in detecting changes in the distractor color compared to the target color. Furthermore, they show a small performance decrease as the target color range increases, replicating the pattern observed in Experiment 1.

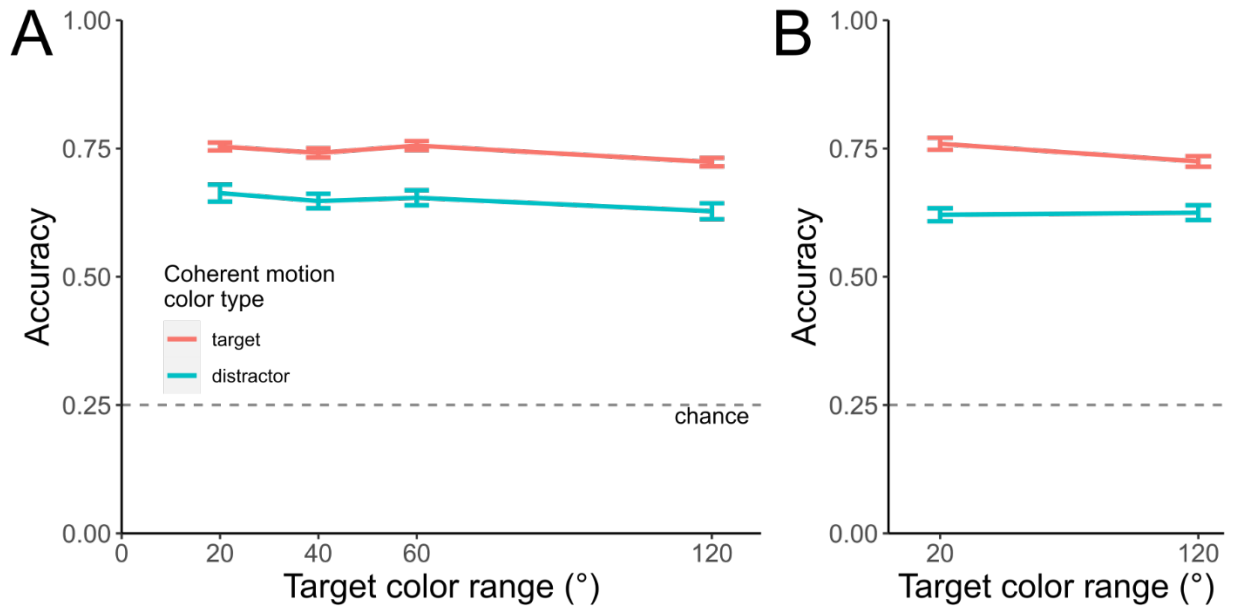


Figure 2.4. A) Results of Experiment 3a. Accuracy in reporting the direction of coherent motion was higher for target than distractor dots. Further, a small but reliable decrease in accuracy was observed as a function of the range of target colors. B) Results of Experiment 3b with two, shorter presentation windows. Accuracy for coherent motion direction was higher for target than distractor dots. There was a significant decrease in accuracy at higher target color ranges, but only for coherent motion in the target dots. Error bars are within-subjects SEM.

### Experiment 3b

In Experiment 3b, we attempted to replicate the basic findings of Experiment 3a in an alternate version of the task that used two shorter time windows instead of one long presentation window. Specifically, participants were presented with two consecutive 500 ms windows with only one window containing the coherent motion. They were instructed to attend to both consecutive displays to detect motion in one window or the other.

### Method

#### Participants

Sixty-eight undergraduate students at UC San Diego participated in this experiment for course credit. Eight participants were excluded with accuracy below 60% on attention check trials (motion coherence: 100%; target color range: 0°; see *Procedure*). The remaining 60 participants (45 women, 11

men, one non-binary, 3 did not report their gender) were between 18-32 years of age ( $M = 20.14 \pm 2.2$  years). This sample size provides 80% power to detect a significant main effect of target color range and interaction with coherent motion color type of  $\eta_p^2 > .119$ .

### **Stimuli & Procedure**

Trials proceeded similar to Experiment 3a, but with the presentation of two shorter stimulus windows rather than one longer window. First, the target dots were presented alone for 1000 ms to cue participants which colors to attend to. Following this, participants were presented with two consecutive 500 ms displays consisting of target and distractor dots, and had to attend to both sets of dots in both windows to detect coherent motion (50% coherence). The coherent motion could occur in either the target or distractor dots, in either the first or second presentation window, and in one of the cardinal directions (left, right, up, or down). At the end of the trial, participants had to report which window the motion was present in (responding with '1' or '2') as well as the direction of the motion (using the arrow keys), regardless of which dots moved coherently. Target dots moved coherently on 80% of trials, and participants were instructed to attend to them primarily, while still reporting the direction of the distractor dots if they moved coherently. Participants completed 264 trials of this task, including 24 "attention check" trials which had 100% coherent motion in the target dots in one window and all target dots were presented in a single color. The remaining 240 trials were distributed evenly across the two target color range conditions (20° or 120°) and coherent motion color type (target or distractor dots, 80%/20% respectively).

### **Results**

To assess the effect of target color range (20° vs. 120°) and attention (coherent motion in the target vs. distractor dots) on performance, we performed a repeated-measures ANOVA on participants' accuracy in reporting the direction of the coherent motion. There was a significant effect of attention,  $F(1,59) = 47.41, p < .001, \eta_p^2 = .446$ , as participants had higher accuracy for reporting the direction of

coherent motion in target dots relative to distractors. There was no main effect of target color range,  $F(1,59) = 2.56, p = .115, \eta_p^2 = .042$ , however there was a significant interaction,  $F(1,59) = 5.37, p = .024, \eta_p^2 = .083$ . Follow-up analyses revealed that accuracy for detecting the target motion was significantly lower for target color ranges of  $120^\circ$  relative to  $20^\circ$ ,  $t(59) = 4.23, p < .001, d_z = 0.55$ . In contrast, accuracy for detecting distractor motion was not affected by the target color range,  $t(59) = 0.26, p = .793, d_z = 0.03$ . These patterns are summarized in *Figure 2.4B*.

Overall, the findings from Experiments 3a and b indicate that the effects of increasing the range of target colors are driven by attentional selection during perceptual processing and not later decision-related processing. Even when participants reported changes that occurred in either set of dots, they still had lower performance as the target color range increased, suggesting that selection of the target colors was less efficient as the size of their attentional focus increased. Despite this, performance was still relatively high overall, consistent with the interpretation that feature-based attention can be tuned relatively efficiently to a wide range of colors. Additionally, because participants attended to targets within relatively short presentation windows in Experiment 3b, these findings cannot straightforwardly be attributed to strategic shifts of attention throughout the range of target colors during the trial.

#### **Experiment 4**

Thus far, our data are consistent with the interpretation that participants can adjust the size of the focus of feature-based attention according to task demands, with only small costs in performance. This interpretation is similar to results in spatial attention where studies have shown that participants adjust the size of their spatial focus of attention in response to different cues (Castiello & Umiltà, 1990; Eriksen & St. James, 1986). However, the decrease in performance we observed is also consistent with an account in which the size of the attentional focus is fixed and relatively narrow, in which case lower performance at wider ranges could simply be driven by the fact that participants only attend to a subset of target colors, effectively not processing colors outside of this narrow focus to the same extent. Because the target

events in our experiments (i.e., luminance decreases or coherent motion) occur throughout the entire range of colors, signals occurring outside the scope of a fixed focus would be much weaker, resulting in overall lower performance as the range of target colors increases. While Experiment 2 rules out a strong version of such fixed capacity models, which would predict a much stronger decrease in performance than the one we find, we aimed to test more directly how attention is distributed across broad ranges of color in Experiment 4.

We instructed participants to attend to broad ranges of feature values (120° around the color wheel), while coherent motion events occurred either for the 50% of target dots in the center of the color distribution ( $\pm 30^\circ$  of the central target color) or at the edges of this color distribution (outer 50%; i.e., colors 30-60° either side of the target range; see *Figure 2.5A*). If participants attend primarily to a narrow range of colors in the center of the color distribution, we expect performance to be higher for trials in which the motion coherence appeared in the center, relative to at the edges, of the distribution. Alternatively, if attention is tuned broadly and uniformly across the entire color range, we would expect no difference in performance between these conditions.

## Method

### Participants

Forty-six undergraduate students at UC San Diego participated in this experiment for course credit and gave informed consent prior to starting the experiment as approved by the Institutional Review Board at UC San Diego. Six participants were excluded with accuracy below 60% on attention check trials (motion coherence: 100%). The remaining 40 participants (31 women, 8 men, 1 did not report their gender) were between 18-30 years of age ( $M = 20.5 \pm 2.0$  years). This sample size provides 80% power to detect a significant main effect of  $\eta_p^2 > .200$  and significant paired-comparisons of  $d_z > 0.45$ .

### Stimuli & Procedure

The experiment was presented similarly to Experiment 3a, however on each trial the range of target colors was 120° and motion coherence was 50% for 500 ms. Participants completed 264 trials of

the task, divided into 3 main conditions determined by which dots moved coherently: distractor events (48 trials), central target events (96 trials), and edge target events (96 trials). On distractor trials, a random 50% of the distractor-colored dots moved coherently; on central target trials, the 50% of dots at the center of the target color range ( $\pm 30^\circ$  of the mean color) moved coherently; on edge target trials, the 50% of dots outside the center of the target color range ( $30\text{--}60^\circ$  clockwise and counter-clockwise from the mean color, see *Figure 2.5A*) moved coherently. Distractor trials made up 20% of the main trials of the experiment, as in Experiment 3a. The remaining 24 trials were attention check trials, in which 100% of the target dots moved coherently. Just like in Experiment 3a, these trials only served the purpose to exclude data from participants that were not doing the task as instructed and were not used in the main analysis.

## Results

The effect of the color type (target-center, target-edge, distractor) revealed that performance was significantly affected by which dots moved coherently,  $F(2,78) = 19.05$ ,  $p < .001$ ,  $\eta_p^2 = .328$ .

Comparisons between the conditions showed that performance was significantly lower for motion coherence in the distractor dots compared to either of the target conditions (target-center:  $t(39) = 4.70$ ,  $p < .001$ ,  $d_z = 0.74$ ; target-edge:  $t(39) = 4.36$ ,  $p < .001$ ,  $d_z = 0.69$ ), indicating that participant selectively attended to the target colors. However, performance did not differ between motion coherence probed at the center or the edge of the target dot range,  $t(39) = 0.04$ ,  $p = .972$ ,  $d_z = 0.01$ , and Bayesian analysis revealed support for a null effect,  $BF_{01} = 5.85$  (JZS prior, with scale  $r = .707$ ; Rouder, Speckman, Sun, Morey, & Iverson, 2009). That is, participants detected coherent motion just as accurately when it occurred in colors at the edge of the target color range—split across two segments of the range—as when it occurred in the center of the target color range. This suggests that participants selected all colors uniformly.

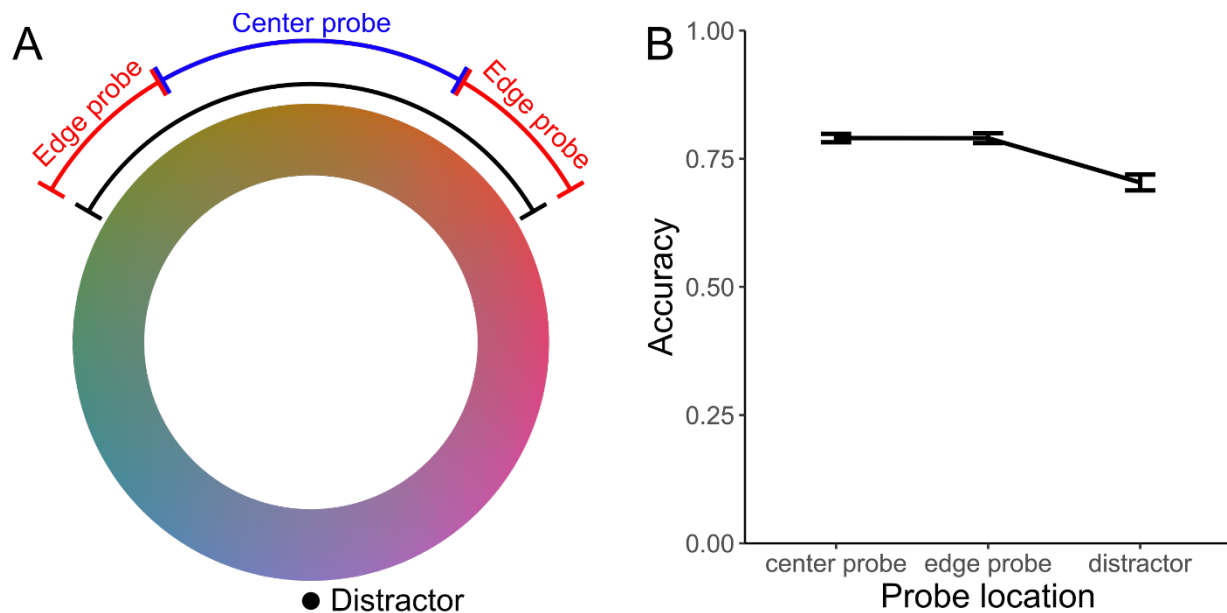


Figure 2.5. A) Example of trial conditions in Experiment 4. The target colors always spanned a range of 120° (black line), but the coherent motion occurred in only 50% of the dots: either the dots in the central 60° of the target range ( $\pm 30^\circ$  from the target mean, blue line), or the dots at the edge of the target range (30-60° either side of the target mean, red line). B) Results of Experiment 4. Accuracy in reporting the direction of coherent motion was just as high for coherence in the center of the range as at the edge of the target color range. However, accuracy was higher for both compared to coherence in the distractor dots.

## Discussion

A critical question for models of attention is whether selection operates in a similar way across different domains, such as location and feature spaces, or whether selection relies on distinct mechanisms—and exhibits distinct limits—depending on the nature of the domain in question. Here, we tested if one proposed characteristic of spatial attention generalizes to the feature domain: whether the attentional focus can be adjusted in size and be flexibly tuned more narrowly or broadly to select different ranges of target features. We used a task which required participants to attend to sets of dots in particular colors to detect brief changes in their luminance or motion direction and varied the range of target colors to assess the efficiency of concurrent attentional selection for multiple colors. We found that as target color ranges increased, performance decreased, as would be expected if selection of multiple colors was less efficient; however, this performance decrease was rather small and participants still performed well above chance at the highest levels of target color range in all

experiments, even when target colors spanned 120° of a circular color space (e.g., from greens through to reds; see *Figure 2.1B*). Interestingly, we found little-to-no cost in efficiency for attending colors spanning approximately 20° in color space, suggesting that this might be the default span of attention focus in this feature space. Experiment 3 confirmed that the effects of color range on performance were driven by changes in the efficiency of early target selection, rather than by later decision-related processes that might be based on the distractors as well. Finally, Experiment 4 showed that attention was distributed across the entire range of target colors, suggesting that selection of a range of colors seems to occur relatively uniformly, given that participants were just as good at detecting a change when it happened to colors within the center of the distribution as when it happened for perceptually more distinct colors at the edge of the color distribution (see *Figure 2.5A*).

Our findings are inconsistent with theories that assume strict capacity limits where only a fixed, narrow range of features can be selected at once—in the extreme, just a single color value. Under such accounts, as the target color range increases, fewer target dots would fall inside this fixed attentional range, resulting in a steep, linear decrease in performance like we observed in Experiment 2 where we decreased the number of target dots presented. Instead, our results suggest that feature-based attentional resources are not strictly fixed, but can instead be allocated relatively flexibly across large ranges of feature values as required by the task at hand.

Previous research has shown a restricted focus of attention when multiple discrete and categorically distinct features are selected (e.g., blue and red, Liu & Jigo, 2017). How then, in our study, do participants select large ranges of color values spanning distinct color categories relatively efficiently? Cueing multiple distinct features as in previous work (Liu et al., 2013; Liu & Jigo, 2017) likely does not affect attention field size, but instead appears more comparable to manipulations assessing “multi-focal” attention, similar to what has been done in spatial attention (Alvarez & Cavanagh, 2005; Awh & Pashler, 2000; Castiello & Umiltà, 1992; McMains & Somers, 2004). In our experiments we always cued entire ranges of colors, and these colors were always chosen from contiguous regions in feature space. Thus, it could be the case that selecting perceptually contiguous color values is easier than selecting two distinct



colors. That is, intermediate colors may help span disconnected regions of the feature space, allowing attention to “bridge the gap” between distinct parts—and thus categorically different colors—of the feature space. Recent work of ours is consistent with this interpretation. By manipulating the similarity between target and distractor items, we found that attentional selection was equally and maximally effective for distractors beyond approximately 40-50° from a target color (Chapman & Störmer, 2022), suggesting that such colors can be easily ignored by the visual system when irrelevant for a task. Compared to the current study, where colors spanning up to 120° of color space were efficiently selected, and previous research arguing that distinct targets could not be jointly selected (Liu et al., 2013; Liu & Jigo, 2017), these findings collectively suggest that a contiguous set of target features can provide a scaffold that aids attentional selection.

Our results indicate that feature-based attention can select broad ranges of feature values relatively efficiently. This is in agreement with previous studies that showed that participants could prepare for an upcoming target feature even when cues were unreliable (e.g., spanning a range of possible target orientation), suggesting that attention was broadly tuned (Ball & Sekuler, 1981; Herrmann et al., 2012). In one study, Herrmann and colleagues (2012) argued in support of the normalization model of attention, in which increasing the attentional field size (i.e., the range of features that are enhanced by attention) resulted in enhancement of relevant feature values, despite a decrease in the overall response of populations tuned to those values. This is consistent with the results of the current study, where wider ranges of target colors lead to overall worse performance, though target colors were still enhanced relative to distractors. Notably, the normalization model does not have separate mechanisms for features and spatial dimensions but treats them as fundamentally similar (Reynolds & Heeger, 2009), and thus allows for both dimensions to have a flexible focus. That is, the differences between feature-based and spatial modes of selection might not be driven by different mechanisms, but simply by the properties of the stimulus space under investigation. Importantly, the normalization model distinguishes between changes in the stimulation field size (i.e., how broad the stimulus is in location or feature space) and changes in the attentional field size, which allows the model to account for a wide range of experimental findings. While

Herrmann and colleagues (2012) kept the stimulus field fixed (e.g., the target was only a single feature value), within the context of the normalization model, our experiments varied the stimulus field by manipulating the range of the target colors and assumed that participants adjusted their attentional focus accordingly. Further investigations that comprehensively manipulate the size of the stimulation and attention fields separately will be necessary to directly test the predictions of the normalization model within different feature spaces.

Models of spatial attention take advantage of the known map-like structure of spatial representations. For example, neighboring spatial positions in the visual field are represented nearby each other in visual cortex, while positions further apart are represented at larger cortical distances (Gardner, Merriam, Movshon, & Heeger, 2008; Mountcastle, 1997), providing a neurological basis for accounts such as the zoom-lens model of attention. Presumably, this retinotopic organization allows attention to modulate visual processing in a targeted way across different parts of the cortical map, for example enhancing smaller or larger contiguous regions within the map, which directly correspond to small and large regions in the external world. Indeed, researchers have argued that this organizing principle can explain cognitive capacity limits, such as those found in attention and working memory (Franconeri, Alvarez, & Cavanagh, 2013). Such map-like structures also exist for visual features: for example, motion direction is organized by similarity in MT (Albright, Desimone, & Gross, 1984), and color maps that are based on perceptual similarity have been found in regions of the ventral processing stream such as V4 (Bohon, Hermann, Hansen, & Conway, 2016; Brouwer & Heeger, 2009; Conway & Tsao, 2009). Thus, there may be a fundamental similarity in how attention selects ranges of features within these feature maps and locations within retinotopically-organized spatial maps. Together, this would suggest that separate mechanisms for feature-based and spatial attention are not necessary (as in Boolean Map theory, Huang & Pashler, 2007, e.g.); instead, a general set of attentional mechanisms can be used across any representations that are appropriately organized. Therefore, we might anticipate that any stimulus space that has a map-like organizational structure would follow similar principles and allow for flexibility in the

focus of attention. Possible examples include shape (Tanaka, 1996, 2003), motion direction (Albright et al., 1984), and sound frequency (Merzenich & Brugge, 1973; Moerel, De Martino, & Formisano, 2012).

Overall, our study demonstrates that feature-based attention is not as restricted as previous research has suggested, and that attention can be tuned relatively efficiently more broadly or narrowly as required by the range of currently relevant feature values. This seems appropriate for an adaptive visual system, given that the quality of visual input is rarely constant. For color in particular, viewing and lighting conditions can dramatically impact perception (Brainard et al., 2006; Lafer-Sousa, Hermann, & Conway, 2015; McDermott, Malkoc, Mulligan, & Webster, 2010), and so flexibility in the attentional system can allow for some stability across these possible changes.

As we have argued, our results mirror previous research in spatial attention, where a flexible size of the attentional focus is a core aspect of many models. Our study thus adds to the evidence suggesting that similar selection principles underlie spatial and feature-based attention. For example, it has recently been shown that feature-based attention elicits surround-suppression in feature space (Navalpakkam & Itti, 2007; Störmer & Alvarez, 2014; Wang, Miller, & Liu, 2015), similar to the selection profile of spatial attention (Hopf et al., 2006; Larsson, Solomon, & Kohn, 2016; Müller et al., 2005). Our results demonstrate that attention can “zoom in” or “zoom out” to select narrow or broad parts of feature values, analogous to selecting small or large spatial regions. These findings demonstrate that at least some of the mechanisms underlying spatial and feature-based attention may be shared, indicating that more effort should be given to understanding the intersection of these two modes of attention.

## **Acknowledgements**

Thank you to Audrey Barszcz, Youngjin Choi, Lora Hsu, and Ashley Williams for assistance with data collection. This research was supported by a grant from the National Science Foundation (BCS-1850738). Chapter 2, in full, has been submitted for publication. Chapman, Angus F.; Störmer, Viola S. The dissertation author was the primary investigator and author of this paper.

## References

- Albright, T. D., Desimone, R., & Gross, C. G. (1984). Columnar organization of directionally selective cells in visual area MT of the macaque. *Journal of Neurophysiology*, *51*(1), 16–31. <https://doi.org/10.1152/jn.1984.51.1.16>
- Alvarez, G. A., & Cavanagh, P. (2005). Independent Resources for Attentional Tracking in the Left and Right Visual Hemifields. *Psychological Science*, *16*(8), 637–643. <https://doi.org/10.1111/j.1467-9280.2005.01587.x>
- Andersen, S. K., Hillyard, S. A., & Müller, M. M. (2008). Attention Facilitates Multiple Stimulus Features in Parallel in Human Visual Cortex. *Current Biology*, *18*(13), 1006–1009. <https://doi.org/10.1016/j.cub.2008.06.030>
- Andersen, S. K., Hillyard, S. A., & Müller, M. M. (2013). Global Facilitation of Attended Features Is Obligatory and Restricts Divided Attention. *Journal of Neuroscience*, *33*(46), 18200–18207. <https://doi.org/10.1523/JNEUROSCI.1913-13.2013>
- Awh, E., & Pashler, H. (2000). Evidence for split attentional foci. *Journal of Experimental Psychology: Human Perception and Performance*, *26*(2), 834–846. <https://doi.org/10.1037/0096-1523.26.2.834>
- Ball, K., & Sekuler, R. (1981). Cues reduce direction uncertainty and enhance motion detection. *Perception & Psychophysics*, *30*(2), 119–128. <https://doi.org/10.3758/BF03204469>
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models using lme4. *Journal of Statistical Software*, *67*(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Bohon, K. S., Hermann, K. L., Hansen, T., & Conway, B. R. (2016). Representation of perceptual color space in macaque posterior inferior temporal cortex (The V4 complex). *ENeuro*, *3*(4). <https://doi.org/10.1523/ENEURO.0039-16.2016>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436.
- Brainard, D. H., Longère, P., Delahunt, P. B., Freeman, W. T., Kraft, J. M., & Xiao, B. (2006). Bayesian model of human color constancy. *Journal of Vision*, *6*, 1267–1281. <https://doi.org/10.1167/6.11.10>
- Brouwer, G. J., & Heeger, D. J. (2009). Decoding and Reconstructing Color from Responses in Human Visual Cortex. *The Journal of Neuroscience*, *29*(44), 13992–14003. <https://doi.org/10.1523/JNEUROSCI.3577-09.2009>
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*(13), 1484–1525. <https://doi.org/10.1016/j.visres.2011.04.012>
- Castiello, U., & Umiltà, C. (1990). Size of the attentional focus and efficiency of processing. *Acta Psychologica*, *73*(3), 195–209. [https://doi.org/10.1016/0001-6918\(90\)90022-8](https://doi.org/10.1016/0001-6918(90)90022-8)
- Castiello, U., & Umiltà, C. (1992). Splitting Focal Attention. *Journal of Experimental Psychology:*

- Human Perception and Performance*, 18(3), 837–848. <https://doi.org/10.1037/0096-1523.18.3.837>
- Chapman, A. F., & Störmer, V. S. (2022). Feature similarity is non-linearly related to attentional selection: evidence from visual search and sustained attention tasks. *Journal of Vision*, 22(8), 4. <https://doi.org/10.1167/jov.22.8.4>
- Conway, B. R., & Tsao, D. Y. (2009). Color-tuned neurons are spatially clustered according to color preference within alert macaque posterior inferior temporal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 106(42), 18034–18039. <https://doi.org/10.1073/pnas.0810943106>
- Cutzu, F., & Tsotsos, J. K. (2003). The selective tuning model of attention : psychophysical evidence for a suppressive annulus around an attended item. *Vision Research*, 43, 205–219.
- de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, 47(1), 1–12. <https://doi.org/10.3758/s13428-014-0458-y>
- Desimone, R., & Duncan, J. (1995). Neural Mechanisms of Selective Visual Attention. *Annual Review of Neuroscience*, 18(1), 193–222. <https://doi.org/10.1146/annurev.ne.18.030195.001205>
- Dori, H., & Henik, A. (2006). *Indications for two attentional gradients in endogenous visual-spatial attention*. *Visual Cognition* (Vol. 13). <https://doi.org/10.1080/13506280500277504>
- Downing, C. J., & Pinker, S. (1985). The spatial structure of visual attention. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and Performance XI*. Routledge.
- Eriksen, C. W., & St. James, J. D. (1986). Visual attention within and around the field of focal attention. *Perception & Psychophysics*, 40(4), 225–240.
- Eriksen, C. W., & Yeh, Y. Y. (1985). Allocation of Attention in the Visual Field. *Journal of Experimental Psychology: Human Perception and Performance*, 11(5), 583–597. <https://doi.org/10.1037/0096-1523.11.5.583>
- Franconeri, S. L., Alvarez, G. A., & Cavanagh, P. (2013). Flexible cognitive resources: Competitive content maps for attention and memory. *Trends in Cognitive Sciences*, 17(3), 134–141. <https://doi.org/10.1016/j.tics.2013.01.010>
- Gardner, J. L., Merriam, E. P., Movshon, J. A., & Heeger, D. J. (2008). Maps of visual space in human occipital cortex are retinotopic, not spatiotopic. *Journal of Neuroscience*, 28(15), 3988–3999. <https://doi.org/10.1523/JNEUROSCI.5476-07.2008>
- Herrmann, K., Montaser-Kouhsari, L., Carrasco, M., & Heeger, D. J. (2010). When size matters: Attention affects performance by contrast or response gain. *Nature Neuroscience*, 13(12), 1554–1561. <https://doi.org/10.1038/nn.2669>
- 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.  
<https://doi.org/10.1073/pnas.0507746103>
- Huang, L., & Pashler, H. (2007). A Boolean Map Theory of Visual Attention. *Psychological Review*, *114*(3), 599–631. <https://doi.org/10.1037/0033-295X.114.3.599>
- Itthipuripat, S., Garcia, J. O., Rungratsameetaweemana, N., Sprague, T. C., & Serences, J. T. (2014). Changing the spatial scope of attention alters patterns of neural gain in human cortex. *Journal of Neuroscience*, *34*(1), 112–123. <https://doi.org/10.1523/JNEUROSCI.3943-13.2014>
- Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, *40*(10–12), 1489–1506. [https://doi.org/10.1016/S0042-6989\(99\)00163-7](https://doi.org/10.1016/S0042-6989(99)00163-7)
- Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1999). Increased Activity in Human Visual Cortex during Directed Attention in the Absence of Visual Stimulation. *Neuron*, *22*, 751–761. [https://doi.org/10.1016/S0896-6273\(00\)80734-5](https://doi.org/10.1016/S0896-6273(00)80734-5)
- LaBerge, D., & Brown, V. (1989). Theory of Attentional Operations in Shape Identification. *Psychological Review*, *96*(1), 101–124. <https://doi.org/10.1037/0033-295X.96.1.101>
- Lafer-Sousa, R., Hermann, K. L., & Conway, B. R. (2015). Striking individual differences in color perception uncovered by “the dress” photograph. *Current Biology*, *25*(13), R545–R546.  
<https://doi.org/10.1016/j.cub.2015.04.053>
- Lamy, D., & Tsal, Y. (2001). On the status of location in visual attention. *European Journal of Cognitive Psychology*, *13*(3), 305–342. <https://doi.org/10.1080/09541440125789>
- Larsson, J., Solomon, S. G., & Kohn, A. (2016). fMRI adaptation revisited. *Cortex*, *80*, 154–160.  
<https://doi.org/10.1016/j.cortex.2015.10.026>
- Ling, S., Liu, T., & Carrasco, M. (2009). How spatial and feature-based attention affect the gain and tuning of population responses. *Vision Research*, *49*(10), 1194–1204.  
<https://doi.org/10.1016/j.visres.2008.05.025>
- Liu, T., Becker, M. W., & Jigo, M. (2013). Limited featured-based attention to multiple features. *Vision Research*, *85*, 36–44. <https://doi.org/10.1016/j.visres.2012.09.001>
- Liu, T., & Jigo, M. (2017). Limits in feature-based attention to multiple colors. *Attention, Perception, and Psychophysics*, *79*(8), 2327–2337. <https://doi.org/10.3758/s13414-017-1390-x>
- Liu, T., Larsson, J., & Carrasco, M. (2007). Feature-Based Attention Modulates Orientation-Selective Responses in Human Visual Cortex. *Neuron*, *55*(2), 313–323.  
<https://doi.org/10.1016/j.neuron.2007.06.030>
- Liu, T., Stevens, S. T., & Carrasco, M. (2007). Comparing the time course and efficacy of spatial and

- feature-based attention. *Vision Research*, 47(1), 108–113.  
<https://doi.org/10.1016/j.visres.2006.09.017>
- Maringelli, F., & Umiltà, C. (1998). The control of the attentional focus. *European Journal of Cognitive Psychology*, 10(3), 225–246. <https://doi.org/10.1080/713752276>
- 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.  
<https://doi.org/10.1016/j.cub.2004.04.028>
- Maunsell, J. H. R., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, 29(6), 317–322. <https://doi.org/10.1016/j.tins.2006.04.001>
- McDermott, K. C., Malkoc, G., Mulligan, J. B., & Webster, M. A. (2010). Adaptation and visual salience. *Journal of Vision*, 10(13), 1–32. <https://doi.org/10.1167/10.13.17>
- McMains, S. A., & Somers, D. C. (2004). Multiple Spotlights of Attentional Selection in Human Visual Cortex. *Neuron*, 42, 677–686. [https://doi.org/10.1016/S0896-6273\(04\)00263-6](https://doi.org/10.1016/S0896-6273(04)00263-6)
- Merzenich, M. M., & Brugge, J. F. (1973). Representation of the cochlear partition on the superior temporal plane of the macaque monkey. *Brain Research*, 50(2), 275–296.  
[https://doi.org/10.1016/0006-8993\(73\)90731-2](https://doi.org/10.1016/0006-8993(73)90731-2)
- Moerel, M., De Martino, F., & Formisano, E. (2012). Processing of natural sounds in human auditory cortex: Tonotopy, spectral tuning, and relation to voice sensitivity. *Journal of Neuroscience*, 32(41), 14205–14216. <https://doi.org/10.1523/JNEUROSCI.1388-12.2012>
- Mountcastle, V. B. (1997). The columnar organization of the neocortex. *Brain*, 120(4), 701–722.  
<https://doi.org/10.1093/brain/120.4.701>
- Mounts, J. R. W. (2000). Attentional capture by abrupt onsets and feature singletons produces inhibitory surrounds. *Perception & Psychophysics*, 62(7), 1485–1493. <https://doi.org/10.3758/BF03212148>
- Müller, N. G., Mollenhauer, M., Rösler, A., & Kleinschmidt, A. (2005). The attentional field has a Mexican hat distribution. *Vision Research*, 45(9), 1129–1137.  
<https://doi.org/10.1016/j.visres.2004.11.003>
- Navalpakkam, V., & Itti, L. (2007). Search Goal Tunes Visual Features Optimally. *Neuron*, 53(4), 605–617. <https://doi.org/10.1016/j.neuron.2007.01.018>
- Pestilli, F., & Carrasco, M. (2005). Attention enhances contrast sensitivity at cued and impairs it at uncued locations. *Vision Research*, 45(14), 1867–1875. <https://doi.org/10.1016/j.visres.2005.01.019>
- Posner, M. I. (1980). Orienting of attention. *The Quarterly Journal of Experimental Psychology*, 32(1), 3–25. <https://doi.org/10.1080/00335558008248231>
- Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of*

- Experimental Psychology: General*, 109(2), 160–174. <https://doi.org/10.1037/0096-3445.109.2.160>
- Prins, N., & Kingdom, F. A. A. (2009). Palamedes: Matlab routines for analyzing psychophysical data. Retrieved from <http://www.palamedestoolbox.org>
- Rajananda, S., Lau, H., & Odegaard, B. (2018). A random-dot kinematogram for web-based vision research. *Journal of Open Research Software*, 6(1). <https://doi.org/10.5334/jors.194>
- Reynolds, J. H., & Chelazzi, L. (2004). Attentional modulation of visual processing. *Annual Review of Neuroscience*, 27, 611–647. <https://doi.org/10.1146/annurev.neuro.26.041002.131039>
- Reynolds, J. H., & Heeger, D. J. (2009). The Normalization Model of Attention. *Neuron*, 61(2), 168–185. <https://doi.org/10.1016/j.neuron.2009.01.002>
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin and Review*, 16(2), 225–237. <https://doi.org/10.3758/PBR.16.2.225>
- Sàenz, M., Buraças, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision Research*, 43(6), 629–637. [https://doi.org/10.1016/S0042-6989\(02\)00595-3](https://doi.org/10.1016/S0042-6989(02)00595-3)
- Shulman, G. L., Sheehy, J. B., & Wilson, J. (1986). Gradients of spatial attention. *Acta Psychologica*, 61, 167–181. [https://doi.org/10.1016/0001-6918\(86\)90029-6](https://doi.org/10.1016/0001-6918(86)90029-6)
- Störmer, V. S., & Alvarez, G. A. (2014). Feature-based attention elicits surround suppression in feature space. *Current Biology*, 24(17), 1985–1988. <https://doi.org/10.1016/j.cub.2014.07.030>
- Tanaka, K. (1996). Inferotemporal Cortex and Object Vision. *Annual Review of Neuroscience*, 19, 109–139. <https://doi.org/10.1146/annurev.neuro.19.1.109>
- Tanaka, K. (2003). Columns for complex visual object features in the inferotemporal cortex: Clustering of cells with similar but slightly different stimulus selectivities. *Cerebral Cortex*, 13(1), 90–99. <https://doi.org/10.1093/cercor/13.1.90>
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97–136. [https://doi.org/10.1016/0010-0285\(80\)90005-5](https://doi.org/10.1016/0010-0285(80)90005-5)
- Wang, Y., Miller, J., & Liu, T. (2015). Suppression effects in feature-based attention. *Journal of Vision*, 15(5), 1–16. <https://doi.org/10.1167/15.5.15>
- White, A. L., & Carrasco, M. (2011). Feature-based attention involuntarily and simultaneously improves visual performance across locations. *Journal of Vision*, 11(6), 1–10. <https://doi.org/10.1167/11.6.15>
- Wolfe, J. M. (1994). Guided Search 2.0 A revised model of visual search. *Psychonomic Bulletin & Review*, 1(2), 202–238. <https://doi.org/10.3758/BF03200774>



## CHAPTER 3: Optimal tuning of feature-based attention warps the perception of visual features

## Abstract

Selective attention improves sensory processing of relevant information but can also impact the quality of perception. For example, attention increases visual discrimination performance and at the same time boosts apparent stimulus contrast of attended relative to unattended stimuli. Can attention also lead to perceptual distortions of visual representations? Optimal tuning accounts of attention suggest that processing is biased towards “off-tuned” features to maximize the signal-to-noise ratio in favor of the target, especially when targets and distractors are confusable. Here, we tested whether such optimal tuning gives rise to phenomenological changes of visual features. We instructed participants to select a color among other colors in a visual search display and subsequently asked them to judge the appearance of the target color in a 2-alternative forced choice (2-AFC) task. Participants consistently judged the target color to appear more dissimilar from the distractor color in feature space. Critically, the magnitude of these perceptual biases varied systematically with the similarity between target and distractor colors during search, indicating that attentional tuning quickly adapts to current task demands. In control experiments we rule out possible non-attentional explanations such as color contrast or memory effects. Overall, our results demonstrate that selective attention warps the representational geometry of color space, resulting in profound perceptual changes across large swaths of feature space. Broadly, these results indicate that efficient attentional selection can come at a perceptual cost by distorting our sensory experience.

## Introduction

Our perceptual experience is strongly influenced by attention. At any given moment, our senses are confronted with an abundant amount of information, exceeding our mind's capacity to process that information effectively. Selective attention enables us to prioritize relevant information over irrelevant information. Attention can be allocated to specific regions in the visual field or to visual features, such as a color or orientation, and in both cases results in increased performance for attended relative to unattended information (Carrasco, 2011; Desimone & Duncan, 1995; Duncan & Humphreys, 1989; Posner, 1980). These behavioral improvements are accompanied by increases in the activity of neural populations associated with the attended spatial location (Bisley & Goldberg, 2003; Kastner, Pinsk, De Weerd, Desimone, & Ungerleider, 1999; Luck, Chelazzi, Hillyard, & Desimone, 1997; Sprague & Serences, 2013; Williford & Maunsell, 2006) or visual feature (Martinez-Trujillo & Treue, 2004; Sàenz, Buraças, & Boynton, 2002; Störmer & Alvarez, 2014; Treue & Martinez-Trujillo, 1999), consistent with theories proposing that attention increases the gain of neurons most selective for the relevant stimulus (Desimone & Duncan, 1995; Martinez-Trujillo & Treue, 2004).

More recent evidence indicates that selective attention not only improves sensory processing, but can also have profound effects on our subjective visual experience, that is, how an attended visual stimulus appears to the observer. Several studies have demonstrated that spatially attending to a visual object increases the apparent intensity of that stimulus, making it appear higher contrast or more saturated than an unattended stimulus (Carrasco, Ling, & Read, 2004; Fuller & Carrasco, 2006; Liu, Abrams, & Carrasco, 2009; Störmer & Alvarez, 2016; Störmer, McDonald, & Hillyard, 2009), in line with neurophysiological findings showing that attention alters the strength of a stimulus by increasing its effective contrast (Martínez-Trujillo & Treue, 2002; Reynolds, Pasternak, & Desimone, 2000; Treue, 2001; Williford & Maunsell, 2006). In these cases, gain modulation models can explain the changes in both perception and performance related to attended stimuli.

Critically, theories of feature-based attention in particular argue that selection via gain modulation of target-tuned neurons is not always optimal, especially in complex visual scenes where

distractor features may be easily confusable with the target feature. Optimal tuning accounts instead propose that, when targets and distractors are similar, attention may enhance neurons that are tuned away from the target feature to maximize the relative saliency of the target stimulus (Navalpakkam & Itti, 2007; Scolari, Byers, & Serences, 2012; Scolari & Serences, 2009). Specifically, when selecting a target among similar distractors (e.g., an orange item among yellow items), increasing the signal of an exaggerated target feature that is less similar to the distractor (e.g., red-orange) would provide a higher signal-to-noise ratio relative to selecting the exact target feature. While previous studies have found that human observers enhance processing of the most informative feature value given the experimental context, not necessarily the target feature itself (Scolari et al., 2012; Scolari & Serences, 2009), it is unknown how such enhancements would impact the perceptual experience of target and nontarget features. We predicted that if attention-induced modulations of neural responses are linked to changes in appearance, then selection would distort feature representations in systematic ways such that the perceived target would be biased away from the distractor feature. Indeed, the current study demonstrates profound distortions in color perception due to attention: if the goal is to select an orange item among yellow distractors, the selected target color appears more red-like than it actually is. Importantly, we also show that these distortions affect large swaths of feature space, indicating that attentional enhancement to an off-tuned color in one part of the feature space alters the underlying representational geometry globally. We are the first to show such dramatic warping in representational spaces due to attention, discovering that our subjective experience of the world is often distorted to optimize information processing.

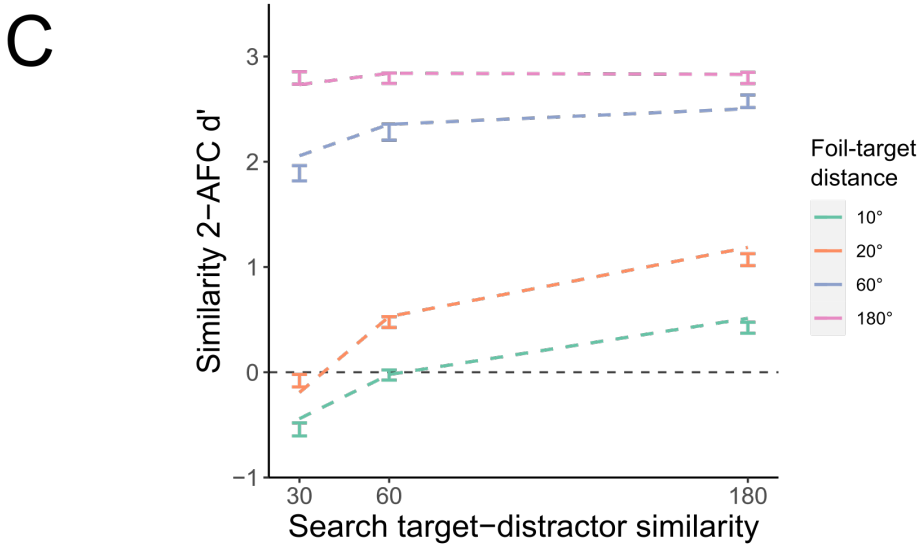
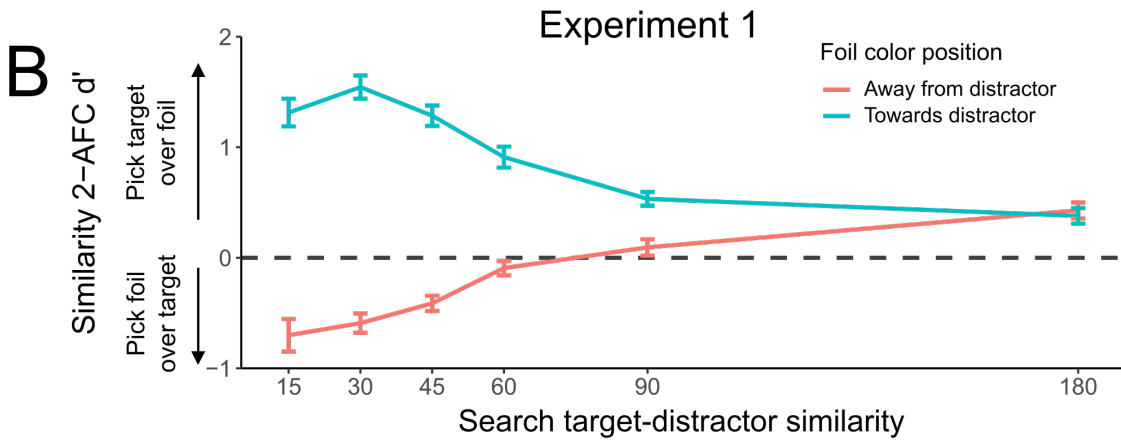
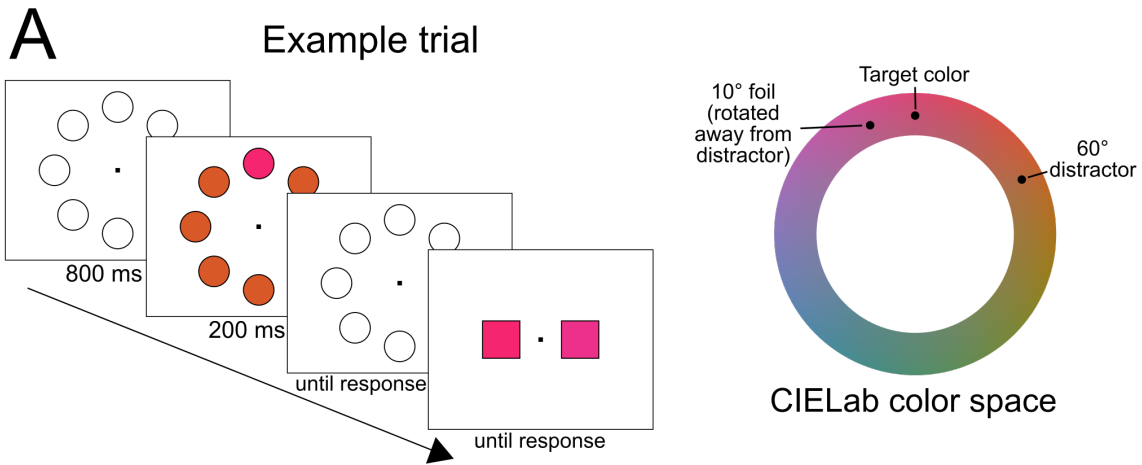
## Results

**Experiment 1: Perception of an attended color is biased away from distractor colors.** We investigated how feature-based attention impacts the perception of target features across a series of experiments in which participants first searched through an array of briefly (200 ms) presented colored circles to locate an oddball target (see *Figure 3.1A*). The target color was chosen randomly on each trial from a 360° CIELab color space, while the distractor was chosen at specific feature distances from this

target to be more or less perceptually similar. Immediately after reporting the target item's location, participants performed a 2-alternative forced choice (2AFC) task reporting which of two squares was most similar in color to the previously selected target item. Of the two squares, one was always presented in the exact target color, while the other was a foil color chosen to be 10° away from the target on the color wheel, either in the direction towards the search distractor, or away from it (i.e., more similar to the distractor or less similar, *Figure 3.1A*).

The main question was whether participants' color reports during the 2AFC were influenced by the relation between the distractor color during search and the foil color in the 2AFC. If selective attention is tuned away from the distractor color to more effectively select the target, as predicted by optimal tuning accounts, then the target color itself might appear shifted away from the distractor color. In this case, participants should choose the foil color more often than the target when it is more distant from the distractor color (i.e., rotated away from the distractor on the color wheel). Conversely, when the foil is less distant from the distractor (rotated towards the distractor), participants should choose the target color more often than the foil. Our results demonstrate that this was the case: Participants' ( $n = 30$ ) responses in the 2AFC were strongly influenced by the relative feature distance between the foil color of the 2AFC and the distractor color during search ( $F = 182.06$ ,  $P < .001$ ,  $\eta_p^2 = .863$ ). Participants chose the foil color more often than the target color when the foil was rotated away from the distractor, apparent as a negative  $d'$  in 2AFC performance; and they chose the correct target color more often when the foil color was rotated towards the search distractor color, evident as a positive (and quite large)  $d'$  in the 2AFC (*Figure 3.1B*). This pattern implies that the appearance of the target color was distorted and shifted away from distractor colors. Importantly, the magnitude of this bias was affected by the similarity of targets and distractors during search ( $F = 58.37$ ,  $P < .001$ ,  $\eta_p^2 = .668$ ), with the largest bias present for trials with

*Figure 3.1.* A) Structure of an example experiment trial. Participants first searched for an oddball target item among distractors. Following report of the target's location, they then reported which of two colors was most similar to the previous visual search target. All colors were selected from a 360° CIELab color space, and we systematically manipulated target-distractor similarity during visual search as well as the target-foil distance during the 2-AFC. Here an example is shown with a 60° color distance between target and distractor during visual search, and a foil color 10° different from the target, rotated in the direction away from the distractor color. B) Results from Experiment 1. 2-AFC similarity judgment performance ( $d'$ ) was significantly affected by the foil color and target-distractor similarity: Participants chose the target color more often when the foil was rotated towards the distractor (positive  $d'$ ), but chose the foil color more often when it was rotated away for the distractor (negative  $d'$ ) especially for similar target-distractor conditions. C) Results from Experiment 2. 2-AFC similarity judgment performance increased with less similar visual search distractors and greater target-foil distances. All foils in this experiment were rotated in the direction away from the distractor color. Dashed lines show fits from a signal detection model that allowed for variations in bias and strength of the target representation. Best fitting models suggested that 30° and 60° distractors resulted in a bias of 11.5° and 5.3°, respectively.



highly similar distractors (15-45°, *Figure 3.1B*). However, biases were significant for all target-distractor distances below 180° around the color wheel ( $p$ 's < .001). This suggests that feature-based attention alters color perception across a range of difficult search tasks, with particularly strong distortions when target-distractor similarity is high.

**Experiment 2: The magnitude of induced biases depends on target-distractor similarities.** To

estimate the magnitude of attention-induced perceptual biases, in Experiment 2, we manipulated the feature distance between the target and foil colors during the 2AFC similarity judgment as well as the similarity between targets and distractors during visual search ( $n = 50$ ). To ensure a sufficient number of trials, in this experiment, all foils were rotated away from the visual search distractor on the color wheel at 10°, 20°, 60°, or 180° from the target. We found a highly significant interaction between target-foil distance during the 2AFC and target-distractor similarity during search ( $F = 23.6, P < .001, \eta_p^2 = .325$ ).

When the foil color was rotated 10° away from the target color, we replicated the pattern of performance found in the first experiment, with participants selecting the foil more often when the visual search distractor was more similar to the target (30° away in color space, *Figure 3.1C*). Increasing the target-foil distance resulted in better discrimination of the target color, however we found that smaller target-distractor similarities during visual search still affected performance when the foil was up to 60° from the target color ( $p$ 's < .001). These findings indicate that distractors that are similar to targets during visual search can have large biasing effects on the perception of colors across the feature space.

Notably, this experiment revealed two conditions in which similarity judgments, on average, were at near-chance performance ( $d' = 0$  for 60° target-distractor similarity during search, 10° target-foil distance; and 30° target-distractor similarity during search, 20° target-foil distance; *Figure 3.1C*). Such conditions, where participants choose the target and foil color equally often, suggest that participants perceived the two choices to be equidistant from their in-mind target representations. Thus, using these conditions we can infer the magnitude of the bias induced by the distractor color for different target-distractor similarities: the results suggest that for high target-distractor similarities of 30°, the perceptual



bias is approximately  $10^\circ$  in color space, while it is about half ( $5^\circ$ ) when target-distractor similarity is relatively low ( $60^\circ$  on the color wheel). In support of this interpretation, we fit the full experimental data with a model based on signal detection theory (see Supplementary Methods) (Schurgin, Wixted, & Brady, 2020), and found that this model fit the data best when biases of about these magnitudes were included in the target representation for target-distractor distances of  $30^\circ$  (bias =  $11.5^\circ$ ) and  $60^\circ$  (bias =  $5.3^\circ$ ).

**Experiment 3: Biases are not driven by differences in temporal delay.** Delays between encoding and report can exaggerate biases in memory for visual features (Bae & Luck, 2017; Chunharas, Rademaker, Brady, & Serences, 2022; Scotti, Hong, Leber, & Golomb, 2021), and although including RT as a covariate in a reanalysis of Experiment 1 data did not explain the perceptual biases we found (see Supplementary Materials), we wanted to confirm this experimentally by fixing the delay across conditions. Thus, we ran an experiment in which participants ( $n = 30$ ) completed a similar task, however following the presentation of the search array they either responded with the location of the target or immediately made a 2AFC similarity judgment. Despite the fact that there was no delay between the search presentation and the similarity response, and that the timing was now the same across all conditions, we still found a large bias for trials with target-distractor similarities of  $30^\circ$  ( $t = 13.75$ ,  $P < .001$ ,  $d = 2.51$ , *Figure 3.2A*), while there was no bias with  $180^\circ$  similarity ( $t = 1.01$ ,  $P = .321$ ,  $d = 0.18$ ; interaction:  $F = 102.61$ ,  $P < .001$ ,  $\eta_p^2 = .780$ ). Importantly, the magnitude of this bias did not differ when compared with the same condition in the first experiment ( $F = 2.52$ ,  $P = .118$ ). Thus, the perceived color distortions cannot be explained by longer RTs in some conditions than others, but they occur robustly even when there is no temporal delay between selection (visual search) and report (2AFC).

**Experiment 4: Biases are induced by attention and not low-level visual properties.** In each of the previous experiments, visual search arrays contained a single target surrounded by distractors in another color. Thus, low-level visual properties of these arrays, like those involved in simultaneous contrast effects (MacLeod, 2003), rather than feature-based selection of the target itself, might influence the biases

we have observed. To account for this possibility, we conducted a version of the experiment in which participants ( $n = 44$ ) searched for an oddball target color amongst a set of distractors that were rotated  $30^\circ$  in each direction from the target (6 distractors total, 3 rotated clockwise, 3 counterclockwise, see *Figure 3.2B*), negating any potential low-level perceptual effects. To assess the influence of attention on perceptual biases, on each trial distractors rotated in one direction were made more salient by increasing the size of the items' border (Chen, Leber, & Golomb, 2019). When similarity judgments were compared based on whether the foil was rotated towards the salient or non-salient distractor color, we found that the salient distractor induced a clear bias: participants chose the target more often when foils were rotated towards the salient distractor ( $t = 6.57$ ,  $P < .001$ ,  $d = 0.99$ , *Figure 3.2C*). Therefore, attention drives the perceptual biases observed in these experiments, rather than low-level visual properties.

**Simulation: Warping can be explained by changes in representational geometry.** How might the observed behavioral biases be accounted for at the level of neural representations? To explore this question, we took inspiration from the literature on representational geometry (Kriegeskorte & Kievit, 2013; Kriegeskorte & Wei, 2021). Representational geometry aims to understand how information about a stimulus is carried within neural activity, primarily through assessing the similarity or dissimilarity in activity patterns for different stimuli. To that end, we first simulated a population of neurons, with each neuron tuned to a different color within our circular feature space, and then performed multidimensional scaling (MDS) based on the Euclidian distance between each pair of hypothetical neurons. As expected, this MDS resulted in a circular geometry (*Figure 3.3A-B*), demonstrating that the simulated population encoded the circular feature space. We then simulated the effects of attention by modulating the gain of each neuron based on its preference for a particular feature and repeating the MDS analysis, which resulted in a distortion of the feature space, expanding the representation near the target feature, and compressing the representation at the opposite end of the feature space (*Figure 3.3B*).

This change in the representational geometry relates to our findings when considering gain applied to an off-tuned feature: for a given target and distractor,  $30^\circ$  apart in feature space, enhancement

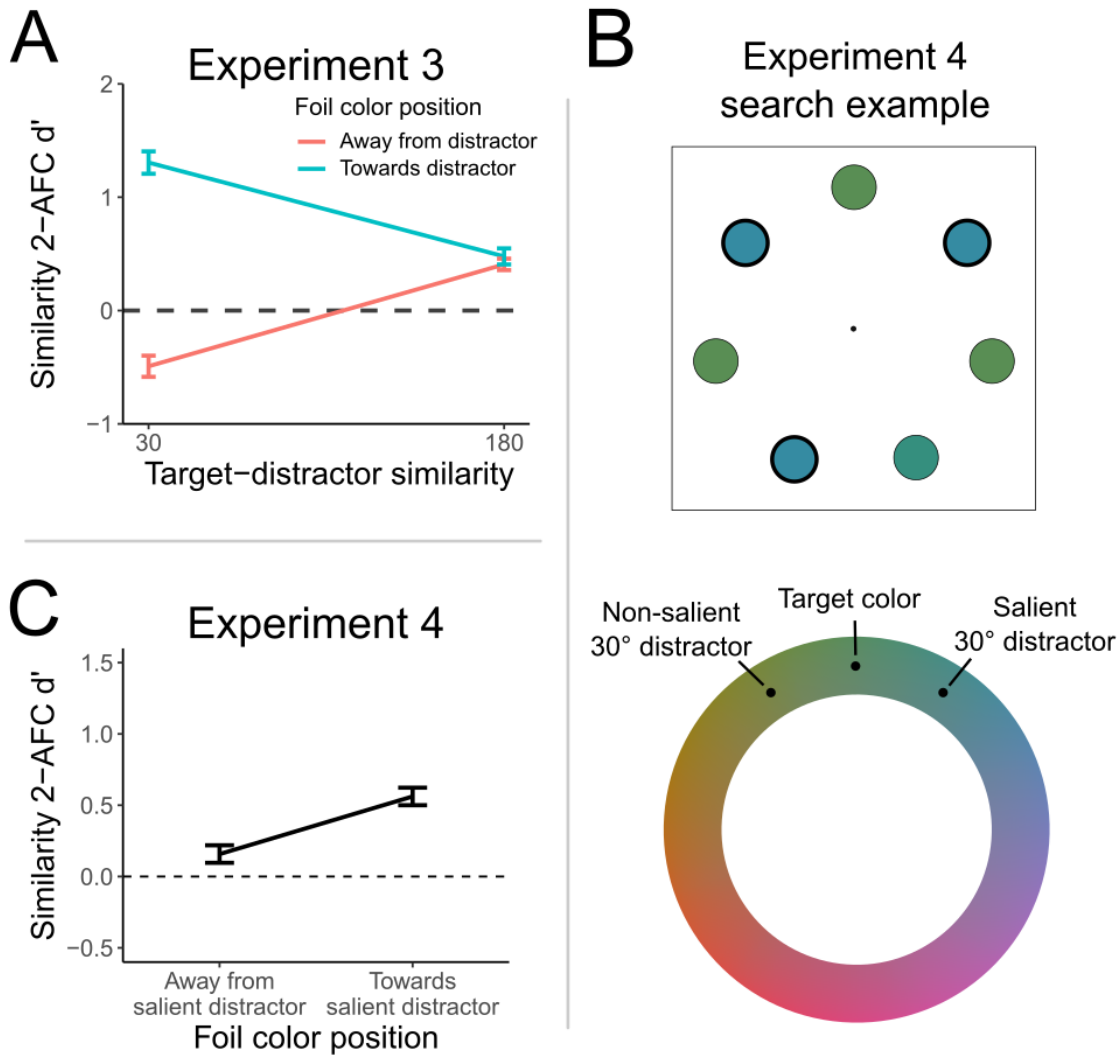



Figure 3.2. A) Results from Experiment 3, where there was no delay between the search array and similarity judgment. Color judgment performance in the 2AFC task was comparable to Experiment 1 for 30° target-distractor similarity during search. B) Example of a visual search trial in Experiment 4. Participants searched for an oddball target color (bottom right in this example) among distractors that were evenly chosen 30° clockwise and counterclockwise relative to the target. One distractor color was made more salient by increasing the thickness of its border. C) Results from Experiment 4, where we manipulated the salience of distractors from one direction around the color wheel. Perception of the target color was more strongly biased by the salient distractor.

of neurons tuned away from the distractor will result in an expansion of the representation centered on that off-tuned feature. Subsequently, this results in an asymmetry in the representational space, with the distance between the target and a 10° foil depending on the position of the foil relative to the distractor. In

the representational space, there is a greater distance between the target and a foil rotated away from the distractor, consistent with the repulsion effects we find behaviorally. Thus, changes in the representational geometry of color space provide a novel explanation for the biases in perception driven by salient distractors.

## Discussion

It is well established that attention enhances our perceptual capacities by improving the detection or discrimination of relevant information in our environment. Our findings show that attention can also lead to quite dramatic distortions in perceptual processing, possibly by causing asymmetric shifts in the representational geometry due to the enhancement of off-tuned neurons. Specifically, our results demonstrate that the appearance of a target color is distorted and biased away from a perceptually similar visual distractor, and that these distortions occur across large parts of the feature space. These findings add significantly to our understanding of how attention supports most effective target selection, namely by rapidly stretching and compressing different parts of the feature space to increase target-distractor dissimilarity. Broadly, the present results relate to other recent research that has shown attention-induced



*Figure 3.3.* A) Hypothetical neural population representing a circular feature space. B) Multidimensional scaling (MDS) of the population results in a circular representational geometry (small circles overlaid on dashed line). Points in MDS space were subsampled to represent colors  $10^\circ$  in the feature space. C) Neural population with simulated attentional gain applied at  $20^\circ$  from the true target position ( $0^\circ$ ), resulting in distortions in the representational geometry with expansion of the representation that is maximal around the center of the attentional gain (large circles in B).

changes in perception, for example studies showing that attention increases the apparent contrast of attended relative to unattended stimuli (Barbot & Carrasco, 2018; Carrasco & Barbot, 2019; Carrasco et al., 2004; Liu et al., 2009; Störmer et al., 2009). Critically, we here show that changes in stimulus appearance due to attention are not necessarily improvements in perception but can also reduce perceptual fidelity.

The perceptual biases observed in our study also have strong connections to previous models of (predominantly spatial) attention. Specifically, a number of attentional mechanisms have been identified that can account for the present effects. For example, recent work indicates that feature-based attention operates via center-surround selection in feature space, such that feature representations similar to the attended feature (i.e., nearby in feature space) are suppressed (Fang, Becker, & Liu, 2019; Hopf et al., 2006; Störmer & Alvarez, 2014). Other studies suggest that feature tuning can be sharpened around the focus of attention (Desimone, Wessinger, Thomas, & Schneider, 1990; Moran & Desimone, 1985), or that additional neural populations are recruited to allocate additional resources towards the selected feature (i.e., receptive field recruitment) (Compte & Wang, 2006; Connor, Gallant, Preddie, & Van Essen, 1996; Connor, Preddie, Gallant, & Van Essen, 1997; McAdams & Maunsell, 1999). In all three cases, the result would be an overall skewed neural population response, effectively attenuating color representations nearby the selected color, which could cause the perceptual distortions observed here. Importantly all these changes in neural populations can lead to similar changes in representational geometry, as the same geometry can be observed from different underlying tuning properties (Kriegeskorte & Wei, 2021). Additionally, these mechanisms are similar to what has previously been proposed in spatial attention tasks that induced perceived position shifts of targets (Suzuki & Cavanagh, 1997), suggesting spatial and feature-based attention bias perception of inputs in comparable ways. Because visual cortex contains topographic representations of features such as color (Bohon, Hermann, Hansen, & Conway, 2016; Brouwer & Heeger, 2009; Conway & Tsao, 2009) in addition to spatial retinotopic maps, the same proposed mechanisms could readily explain the findings of our experiments.

Thus, similar to perceived location distortions in spatial attention, feature-based attention can distort the perception of features and warp feature-based representational maps, such as color.

A number of studies have shown that biases exist in working memory, such that items in memory can be attracted towards or repelled from each other (Bae & Luck, 2017; Chen et al., 2019; Chunharas et al., 2022; Scotti et al., 2021). These reported biases in memory are proposed to be the consequence of actively maintaining multiple representations at once, with some models demonstrating that representations naturally drift when held in memory for extended durations (Bouchacourt & Buschman, 2019; Panichello, DePasquale, Pillow, & Buschman, 2019). In contrast to these memory effects, our findings demonstrate that biases in perceived target color can occur even when only a single item is task-relevant, needs to be selected among another nontarget color in a briefly presented display, and is reported right after with only a minimal temporal delay (Exp. 3). This shows that perceptual representations are distorted at even the earliest stages of visual processing, before they enter into working memory and compete with other representations during storage. Additionally, recent work suggests that attentional templates of target colors can become biased by learning within a consistent distractor context (Yu & Geng, 2019), suggesting that the perceptual effects we observed in our study may be consolidated in more statistically regular situations. In working memory studies, such biases have been attributed to adaptive distortions that improve memory for an overall set of items at the cost of precision for individual items (Chunharas et al., 2022). A similar explanation can be made regarding our findings, as the bias increases the representational distance between target and distractor colors, effectively improving attentional efficiency. The cost of this improved efficiency is that perception of the target is warped from its veridical feature value. This interpretation implies that the role of attention is not necessarily to directly enhance the relevant features, but to shape representational spaces in a way that separates relevant from irrelevant information for most optimal selection. Such warping may also serve to exaggerate and increase the detectability of changes in the environment (Mehrpour, Martinez-Trujillo, & Treue, 2020). Enhancement of targets is an effective mechanism when features are distinct enough from one another, but when neural

representations are more similar and overlapping, other mechanisms such as off-target enhancement can provide more efficient target selection.

Importantly, our findings cannot be explained by basic perceptual effects, such as aftereffects or simultaneous contrast, which can also result in altered perception of visual features. Aftereffects typically take several seconds of stimulus presentation to emerge (Hershenson, 1989; Magnussen & Johnsen, 1986), whereas our visual search arrays were presented for only 200-400 ms at a time. Although simultaneous contrast can occur much more quickly than aftereffects, it is typically observed when one color is fully surrounded by another (Kingdom, 1997; MacLeod, 2003). Our color stimuli were physically separated on the display and delineated by borders (Ekroll & Faul, 2012), meaning low-level interactions are likely negligible. Furthermore, we showed that with a balanced stimulus display, where the contrast between the target and distractors was effectively “cancelled out” (Exp. 4), perception was still biased by whichever distractors were most salient, demonstrating that attention is the driving factor behind the warping effects we observed.

In sum, our study provides strong evidence that attention warps representational spaces of visual features, resulting in biased perception of target features. Further, our simulation shows that these perceptual changes can be explained by changes in the underlying representational geometry. Such findings support models of attention that emphasize the relationship between target and distractor features (Becker, Folk, & Remington, 2013), such as optimal tuning accounts (Navalpakkam & Itti, 2007; Scolari et al., 2012), and argue against those that focus primarily on the target feature itself. Along with comparable findings in spatial attention (Suzuki & Cavanagh, 1997), this suggests that warping may be a general principle of attention, with perceptual representations undergoing rapid and adaptive changes when attention acts to separate items that are similar.

## Methods

### Participants

All experiments were conducted online, with participants recruited from the University of California, San Diego undergraduate subject pool and compensated with course credit. Participants gave informed consent before starting the experiment as approved by the Institutional Review Board at UC San Diego. Experiment 1 consisted of 30 participants (25 women, 5 men; age  $20.4 \pm 1.6$  years). Experiment 2 consisted of 50 participants (42 women, 5 men, 3 did not report their gender; age  $20.6 \pm 1.7$  years), with an additional 2 participants who were excluded with  $<70\%$  accuracy in visual search. Experiment 3 consisted of 30 participants (20 women, 10 men; age  $20.3 \pm 2.6$  years). Experiment 4 consisted of 50 participants (33 women, 7 men, 4 did not report their gender; age  $20.5 \pm 2.1$  years), with an additional 9 participants excluded with  $<40\%$  accuracy during visual search (this exclusion criteria was lower than in the other experiments because the more difficult task resulted in lower overall accuracy).

### Stimuli

Participants completed the experiments online on their own personal computer. Stimulus sizes are given in pixels, given that we had limited control over the display size and viewing distance. However, the display size was restricted to a minimum of 800x600 pixels, and participants were instructed to complete the experiment in full screen.

All colors were selected from a set of 360 equally spaced colors in a fixed-luminance plane of CIELab color space, drawn from a circle with radius 49 units, centered at  $L = 54$ ,  $a = 21.5$ ,  $b = 11.5$ . The visual search array consisted of 8 circles (80px diameter) arranged evenly in a ring centered on fixation (10px diameter filled black dot), except for Experiment 4 in which there were 7 circles evenly arranged around fixation. Each search array item was positioned 260px from fixation. On each trial, the target color was selected randomly from the full set of colors, and the distractor was chosen relative to the target based on the experimental condition. In Experiments 1-3, distractor colors were equally often chosen by rotating the color wheel clockwise and counterclockwise from the target position. In Experiment 4, half of the distractors were chosen in the clockwise direction and half in the counterclockwise direction, relative



to the target color. For the similarity judgment, two colored squares (80px by 80px) were presented on the left and right of the fixation point. One square was always presented in the target color, while the color of the foil square was determined by the experimental condition.

### **Procedure**

Each trial of the task was divided into two steps (see *Figure 3.1A* for an example trial sequence). Participants first searched for an oddball target item among 7 distractor items (6 distractors in Experiment 4). Blank placeholders were initially presented in a circular array around the fixation point for 800 ms. Then the search array items were shown in color for 200 ms, one in the target color and the remaining in the distractor color. The target item's position was randomly determined on each trial. Participants were instructed to report the location of the briefly presented oddball target item as quickly as possible by clicking on the placeholder in the location at which they saw the target. Following their report of the target location, two colored squares were presented on either side of the fixation point, and participants were instructed to select which of the two colors was most similar to the target color from that search trial. One item was always presented in the target color, and the other color was a foil color that varied in similarity depending on experiment and condition. The similarity judgment items remained on the screen until participants made a response. We opted for a 2-AFC judgment as opposed to continuous report to minimize any possible strategic responses by participants based on their knowledge of the distractor color.

For Experiment 4, the distractors in one color (either the clockwise or counterclockwise color on the color wheel and relative to the target) were made more salient by presenting them with a thicker border (5px width compared to the 1px width of the other items). Participants were instructed that the size of the border was irrelevant to the task and to do their best to ignore it. Because of the added task difficulty, we increased visual search presentation time to 400 ms.

In Experiment 1, participants completed 300 trials of this task, with 25 repetitions of each combination of search target-distractor distance (15°, 30°, 45°, 60°, 90°, or 180°) and similarity foil color (10° rotated towards or away from the distractor color on the color wheel). In Experiment 2, participants

completed a total of 432 trials, consisting of 36 trials for each combination of search target-distractor distance (30°, 60°, or 180°) and similarity foil distance (10°, 20°, 60°, or 180° rotated in the direction away from the search distractor color). In Experiment 3, participants completed a total of 288 trials, consisting of 144 visual search responses and 144 similarity judgment responses on separate trials. Visual search trials were evenly divided across the two target-distractor distances (30° or 180°), while similarity judgment trials were divided across the two target-distractor distances and similarity foil colors (foil rotated 10° towards or away from the distractor). In Experiment 4, participants completed a total of 240 trials, divided evenly across the foil color positions (10° towards or away from the salient distractor).

### **Acknowledgements**

Chapter 3, in full, has been submitted for publication. Chapman, Angus F.; Chunharas, Chaipat; Störmer, Viola S. The dissertation author was the primary investigator and author of this paper.

## Supplementary Materials

**Signal detection model of perceptual warping.** To estimate how perceptual representations of the target color were distorted by visual search, we implemented a version of the Target Confusability Competition (TCC) memory model (Schurgin et al., 2020), which takes a signal detection approach to estimating the representational strength of items based on the similarity between target and non-target items. TCC uses a behaviorally measured psychological similarity function to determine how a familiarity signal spreads among features based on their similarity/feature distance. Instead of measuring this similarity function, which limits the resolution to only the distances measured in the behavioral task, we constructed a version of the function reported by (Schurgin et al., 2020) that allowed for infinite resolution. Specifically, we assumed a population of neural tuning curves spanning the stimulus space ( $-180^\circ$  to  $180^\circ$  of color space, centered on a hypothetical target at  $0^\circ$ ). Each tuning curve was a normal distribution ( $SD = 11^\circ$ ) and they were spaced at  $2^\circ$  intervals across the stimulus space. To account for the circular color space, we included tuning curves out to  $-240^\circ$  and  $240^\circ$ , although the final function was restricted to the actual range of the color space ( $-180^\circ$  to  $180^\circ$ ). We then scaled the tuning curves based on their distance from  $0^\circ$  by multiplying them with an exponential function:  $\exp(-L \cdot \mu_d / 180)$ , where  $\mu_d$  was the absolute distance of the distribution from 0, and  $L$  scales the steepness of the exponential and was estimated from data. This roughly corresponds to the psychophysically assessed similarity function.

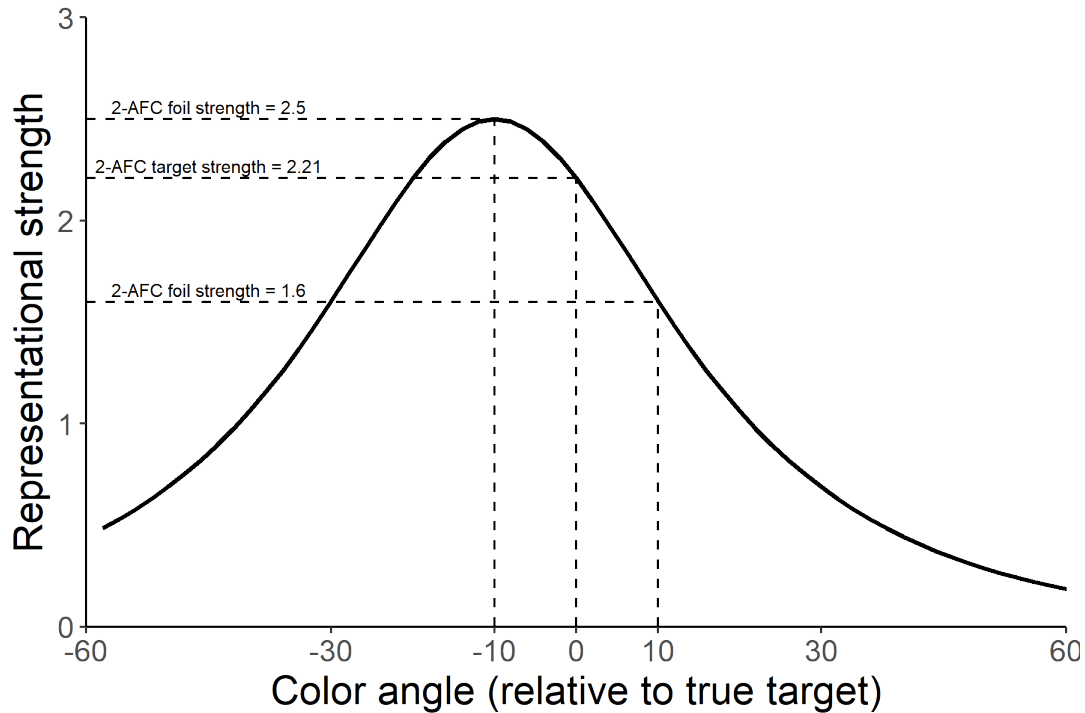
With this similarity function in place, we then proceeded to fit this model to the data from Experiment 2. For each target-distractor distance condition, we estimated two parameters: representational strength for the target color (measured in  $d'$ ), which scales the height of the similarity function; bias (in degrees), which shifts the peak of the similarity function. We also estimated the steepness parameter  $L$ , which was shared across all target-distractor distances. Performance in the 2-AFC similarity judgment was simulated by taking the difference in the height of the similarity function between the target ( $0^\circ$ ) and foil colors, after applying the transformation by the two parameters (*Supplementary Figure 3.1*). Functions were fit to the behavioral data for each subject using the “optim”

function in R (version 4.0.4, R Core Team, 2021) with the Nelder-Mead algorithm. We then averaged the best fitting parameters for each condition across all subjects (see *Supplementary Table 3.1* for summary statistics). Final results of the fitted function are shown in *Figure 3.2*.


**Representational geometry analysis.** We assessed the changes in representational geometry as a function of attention by first simulating a hypothetical neural population. Forty simulated neurons were defined, with a firing rate modeled by a von Mises distribution ( $\kappa = 5$ ) and a preferred feature chosen for each neuron to span the feature space ( $9^\circ$  spacing; see *Figure 3.3A*). Population responses were normalized by dividing by the mean neural response. To simulate the effects of attention on this neural population, we defined an “attention filter” with a gain envelope modeled by another von Mises distribution ( $\kappa = 2$ ), centered at  $-20^\circ$  to reflect a bias away from a hypothetical  $30^\circ$  distractor position. This attention filter was scaled (all values divided by the peak response), weighted (all values multiplied by 0.4), and shifted (a constant amount of 1 added to all values, resulting in increased gain for neurons selective for this feature, but no change in the response of neurons tuned at the other side of the feature space). The response of the neural population was then multiplied with this filter and normalized, giving a transformed population response under the simulated attention conditions (*Figure 3.3C*). To examine the geometry of the population, we performed multidimensional scaling (MDS) using principal components analysis (PCA) on the response properties of the simulated neural population. For these purposes, this is equivalent to using classical MDS with a Euclidian distance metric. PCA was first done on the untransformed population, and then the responses of the population modulated by attention were projected into this PC space using the component weights (*Figure 3.3B*). For plotting purposes, we used the first 2/40 PCs, which accounted for 57.1% of the variance in the simulated neural population.

**Visual search RTs as a function of target-distractor similarity.** In Experiment 1, participants reported the location of the visual search target prior to the presentation of the 2AFC similarity judgment options. RTs during visual search are generally longer the more similar the targets and distractors are (Chapman &

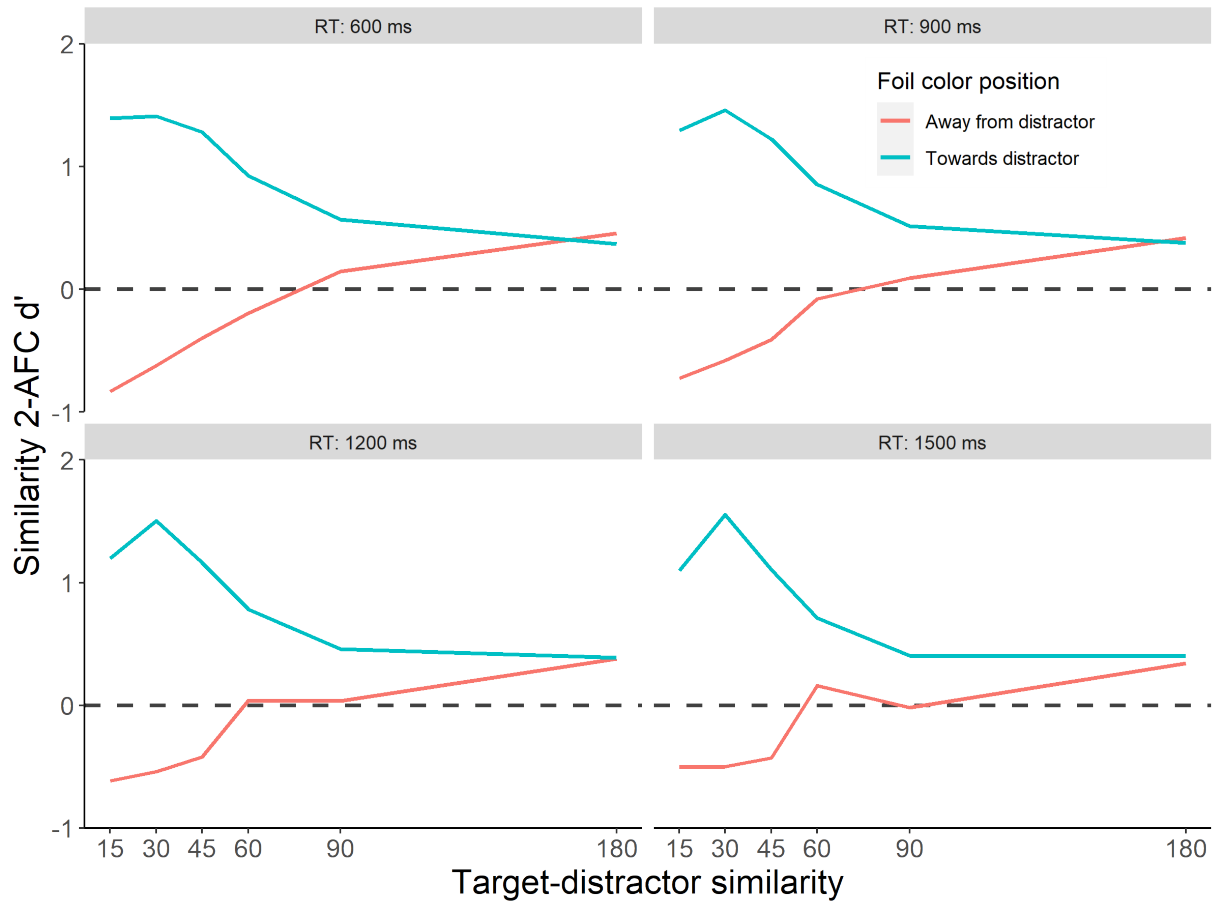
Störmer, 2022), and we observed this in our data as well (*Supplementary Figure 3.2*). Differences in RT might propagate as errors in the similarity judgment task, since biases in working memory studies have been shown to increase with delays between encoding and recall (Chunharas et al., 2022). To test whether this was the case in Experiment 1, we reanalyzed our data while including visual search RT as a covariate in the analysis. The results of this additional analysis are summarized in *Supplementary Figure 3.3*, across a subset of simulated visual search RTs that roughly spanned the range observed in the data. These results suggest that the perceptual biases in this experiment were largely unaffected by differences in RTs from trial-to-trial.



*Supplementary Figure 3.1.* Demonstration of  $d'$  calculation for signal detection model simulations. The representational strength varies for different colors, and in this example peaks at  $-10^\circ$  (negative distances represent colors away from the visual search distractor). To calculate  $d'$ , we compare the representational strength of the colors in the similarity judgment: the target at  $0^\circ$ , and a foil at  $\pm 10^\circ$ . In this example, the  $d'$  for a  $10^\circ$  foil away from the distractor is  $2.21 - 2.5 = -0.29$ , whereas the  $d'$  for a  $10^\circ$  foil towards the distractor is  $2.21 - 1.6 = 0.61$ . See main manuscript for the actual model fits.



*Supplementary Figure 3.2.* Visual search response time in Experiment 1 as a function of target-distractor similarity. In general, response times increase for less similar targets and distractors.



*Supplementary Figure 3.3.* Estimated similarity judgment performance in Experiment 1 as a function of visual search RT.

*Supplementary Table 3.1.* Best fitting parameters for the signal detection model for each target-distractor similarity level. The steepness parameter  $L$ , which was shared across all similarity levels, was 7.30 (SD = 2.46).

Target-distractor similarity	Representational strength ( $d'$ )	Bias
30°	3.18 (0.50)	11.52° (3.82)
60°	2.94 (0.42)	5.30° (4.71)
180°	2.89 (0.44)	-3.95° (5.80)

## References

- Bae, G.-Y., & Luck, S. J. (2017). Interactions between visual working memory representations. *Attention, Perception, and Psychophysics*, *79*(8), 2376–2395. <https://doi.org/10.3758/s13414-017-1404-8>
- Barbot, A., & Carrasco, M. (2018). Emotion and anxiety potentiate the way attention alters visual appearance. *Scientific Reports*, *8*(1), 1–10. <https://doi.org/10.1038/s41598-018-23686-8>
- Becker, S. I., Folk, C. L., & Remington, R. W. (2013). Attentional Capture Does Not Depend on Feature Similarity, but on Target-Nontarget Relations. *Psychological Science*, *24*(5), 634–647. <https://doi.org/10.1177/0956797612458528>
- Bisley, J. W., & Goldberg, M. E. (2003). Neuronal Activity in the Lateral Intraparietal Area and Spatial Attention. *Science*, *299*(5603), 81–86. <https://doi.org/10.1126/science.1077395>
- Bohon, K. S., Hermann, K. L., Hansen, T., & Conway, B. R. (2016). Representation of perceptual color space in macaque posterior inferior temporal cortex (The V4 complex). *ENeuro*, *3*(4). <https://doi.org/10.1523/ENEURO.0039-16.2016>
- Bouchacourt, F., & Buschman, T. J. (2019). A Flexible Model of Working Memory. *Neuron*, *103*(1), 147–160.e8. <https://doi.org/10.1016/j.neuron.2019.04.020>
- Brouwer, G. J., & Heeger, D. J. (2009). Decoding and Reconstructing Color from Responses in Human Visual Cortex. *The Journal of Neuroscience*, *29*(44), 13992–14003. <https://doi.org/10.1523/JNEUROSCI.3577-09.2009>
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*(13), 1484–1525. <https://doi.org/10.1016/j.visres.2011.04.012>
- Carrasco, M., & Barbot, A. (2019). Spatial attention alters visual appearance. *Current Opinion in Psychology*, *29*, 56–64. <https://doi.org/10.1016/j.copsyc.2018.10.010>
- Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience*, *7*(3), 308–313. <https://doi.org/10.1038/nn1194>
- Chapman, A. F., & Störmer, V. S. (2022). Feature similarity is non-linearly related to attentional selection: evidence from visual search and sustained attention tasks. *Journal of Vision*, *22*(8), 4. <https://doi.org/10.1167/jov.22.8.4>
- Chen, J., Leber, A. B., & Golomb, J. D. (2019). Attentional Capture Alters Feature Perception. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(11), 1443–1454.
- Chunharas, C., Rademaker, R. L., Brady, T. F., & Serences, J. T. (2022). An adaptive perspective on visual working memory distortions. *Journal of Experimental Psychology: General*. <https://doi.org/10.1037/xge0001191>
- Compte, A., & Wang, X. J. (2006). Tuning curve shift by attention modulation in cortical neurons: A



- computational study of its mechanisms. *Cerebral Cortex*, *16*(6), 761–778.  
<https://doi.org/10.1093/cercor/bhj021>
- Connor, C. E., Gallant, J. L., Preddie, D. C., & Van Essen, D. C. (1996). Responses in area V4 depend on the spatial relationship between stimulus and attention. *Journal of Neurophysiology*, *75*(3), 1306–1308. <https://doi.org/10.1152/jn.1996.75.3.1306>
- Connor, C. E., Preddie, D. C., Gallant, J. L., & Van Essen, D. C. (1997). Spatial attention effects in macaque area V4. *Journal of Neuroscience*, *17*(9), 3201–3214. <https://doi.org/10.1523/jneurosci.17-09-03201.1997>
- Conway, B. R., & Tsao, D. Y. (2009). Color-tuned neurons are spatially clustered according to color preference within alert macaque posterior inferior temporal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(42), 18034–18039.  
<https://doi.org/10.1073/pnas.0810943106>
- Desimone, R., & Duncan, J. (1995). Neural Mechanisms of Selective Visual Attention. *Annual Review of Neuroscience*, *18*(1), 193–222. <https://doi.org/10.1146/annurev.ne.18.030195.001205>
- Desimone, R., Wessinger, M., Thomas, L., & Schneider, W. (1990). Attentional control of visual perception: Cortical and subcortical mechanisms. *Cold Spring Harbor Symposia on Quantitative Biology*, *55*, 963–971. <https://doi.org/10.1101/SQB.1990.055.01.090>
- Duncan, J., & Humphreys, G. W. (1989). Visual Search and Stimulus Similarity. *Psychological Review*, *96*(3), 433–458. <https://doi.org/10.1037/0033-295X.96.3.433>
- Ekroll, V., & Faul, F. (2012). Basic Characteristics of Simultaneous Color Contrast Revisited. *Psychological Science*, *23*(10), 1246–1255. <https://doi.org/10.1177/0956797612443369>
- Fang, M. W. H., Becker, M. W., & Liu, T. (2019). Attention to colors induces surround suppression at category boundaries. *Scientific Reports*, *9*, 1443. <https://doi.org/10.1038/s41598-018-37610-7>
- Fuller, S., & Carrasco, M. (2006). Exogenous attention and color perception: Performance and appearance of saturation and hue. *Vision Research*, *46*(23), 4032–4047.  
<https://doi.org/10.1016/j.visres.2006.07.014>
- Hershenson, M. (1989). Duration, time constant, and decay of the linear motion aftereffect as a function of inspection duration. *Perception & Psychophysics*, *45*(3), 251–257.  
<https://doi.org/10.3758/BF03210704>
- 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.  
<https://doi.org/10.1073/pnas.0507746103>
- Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1999). Increased Activity

- in Human Visual Cortex during Directed Attention in the Absence of Visual Stimulation. *Neuron*, 22, 751–761. [https://doi.org/10.1016/S0896-6273\(00\)80734-5](https://doi.org/10.1016/S0896-6273(00)80734-5)
- Kingdom, F. A. A. (1997). Simultaneous contrast: the legacies of Hering and Helmholtz. *Perception*, 26(6), 673–677. <https://doi.org/10.1068/p260673>
- Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: Integrating cognition, computation, and the brain. *Trends in Cognitive Sciences*, 17(8), 401–412. <https://doi.org/10.1016/j.tics.2013.06.007>
- Kriegeskorte, N., & Wei, X. X. (2021). Neural tuning and representational geometry. *Nature Reviews Neuroscience*, 22(11), 703–718. <https://doi.org/10.1038/s41583-021-00502-3>
- Liu, T., Abrams, J., & Carrasco, M. (2009). Voluntary attention enhances contrast appearance. *Psychological Science*, 20(3), 354–362. <https://doi.org/10.1111/j.1467-9280.2009.02300.x>
- Luck, S. J., Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, 77(1), 24–42. <https://doi.org/10.1152/jn.1997.77.1.24>
- MacLeod, D. I. A. (2003). New dimensions in color perception. *Trends in Cognitive Sciences*, 7(3), 97–99. [https://doi.org/10.1016/S1364-6613\(03\)00022-6](https://doi.org/10.1016/S1364-6613(03)00022-6)
- Magnussen, S., & Johnsen, T. (1986). Temporal Aspects of Spatial Adaptation. A Study of the Tilt Aftereffect. *Vision Research*, 26(4), 661–672. [https://doi.org/10.1016/0042-6989\(86\)90014-3](https://doi.org/10.1016/0042-6989(86)90014-3)
- 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. <https://doi.org/10.1016/j.cub.2004.04.028>
- Martínez-Trujillo, J. C., & Treue, S. (2002). Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron*, 35(2), 365–370. [https://doi.org/10.1016/S0896-6273\(02\)00778-X](https://doi.org/10.1016/S0896-6273(02)00778-X)
- McAdams, C. J., & Maunsell, J. H. R. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience*, 19(1), 431–441. <https://doi.org/10.1523/jneurosci.19-01-00431.1999>
- Mehrpour, V., Martínez-Trujillo, J. C., & Treue, S. (2020). Attention amplifies neural representations of changes in sensory input at the expense of perceptual accuracy. *Nature Communications*, 11(1). <https://doi.org/10.1038/s41467-020-15989-0>
- Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, 229(4715), 782–784. <https://doi.org/10.1126/science.4023713>
- Navalpakkam, V., & Itti, L. (2007). Search Goal Tunes Visual Features Optimally. *Neuron*, 53(4), 605–617. <https://doi.org/10.1016/j.neuron.2007.01.018>

- Panichello, M. F., DePasquale, B., Pillow, J. W., & Buschman, T. J. (2019). Error-correcting dynamics in visual working memory. *Nature Communications*, *10*(1), 1–11. <https://doi.org/10.1038/s41467-019-11298-3>
- Posner, M. I. (1980). Orienting of attention. *The Quarterly Journal of Experimental Psychology*, *32*(1), 3–25. <https://doi.org/10.1080/00335558008248231>
- R Core Team. (2021). R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.r-project.org/>
- Reynolds, J. H., Pasternak, T., & Desimone, R. (2000). Attention Increases Sensitivity of V4 Neurons. *Neuron*, *26*, 703–714.
- Sàenz, M., Buraças, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience*, *5*(7), 631–632. <https://doi.org/10.1038/mn876>
- Schurgin, M. W., Wixted, J. T., & Brady, T. F. (2020). Psychophysical scaling reveals a unified theory of visual memory strength. *Nature Human Behaviour*, *4*(11), 1156–1172. <https://doi.org/10.1038/s41562-020-00938-0>
- Scolari, M., Byers, A., & Serences, J. T. (2012). Optimal deployment of attentional gain during fine discriminations. *Journal of Neuroscience*, *32*(22), 7723–7733. <https://doi.org/10.1523/JNEUROSCI.5558-11.2012>
- Scolari, M., & Serences, J. T. (2009). Adaptive allocation of attentional gain. *Journal of Neuroscience*, *29*(38), 11933–11942. <https://doi.org/10.1523/JNEUROSCI.5642-08.2009>
- Scotti, P. S., Hong, Y., Leber, A. B., & Golomb, J. D. (2021). Visual Working Memory Items Drift Apart Due to Active, Not Passive, Maintenance. *Journal of Experimental Psychology: General*, *150*(12), 2506–2524. <https://doi.org/10.1037/xge0000890>
- Sprague, T. C., & Serences, J. T. (2013). Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. *Nature Neuroscience*, *16*(12), 1879–1887. <https://doi.org/10.1038/mn.3574>
- Störmer, V. S., & Alvarez, G. A. (2014). Feature-based attention elicits surround suppression in feature space. *Current Biology*, *24*(17), 1985–1988. <https://doi.org/10.1016/j.cub.2014.07.030>
- Störmer, V. S., & Alvarez, G. A. (2016). Attention Alters Perceived Attractiveness. *Psychological Science*, *27*(4), 563–571. <https://doi.org/10.1177/0956797616630964>
- Störmer, V. S., McDonald, J. J., & Hillyard, S. A. (2009). Cross-modal cueing of attention alters appearance and early cortical processing of visual stimuli. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(52), 22456–22461. <https://doi.org/10.1073/pnas.0907573106>

- Suzuki, S., & Cavanagh, P. (1997). Focused attention distorts visual space: An attentional repulsion effect. *Journal of Experimental Psychology: Human Perception and Performance*, 23(2), 443–463. <https://doi.org/10.1037/0096-1523.23.2.443>
- Treue, S. (2001). Neural correlates of attention in primate visual cortex. *Trends in Neurosciences*, 24(5), 295–300. [https://doi.org/10.1016/S0166-2236\(00\)01814-2](https://doi.org/10.1016/S0166-2236(00)01814-2)
- Treue, S., & Martinez-Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, 399, 575–579. <https://doi.org/10.1038/21176>
- Williford, T., & Maunsell, J. H. R. (2006). Effects of spatial attention on contrast response functions in macaque area V4. *Journal of Neurophysiology*, 96(1), 40–54. <https://doi.org/10.1152/jn.01207.2005>
- Yu, X., & Geng, J. J. (2019). The attentional template is shifted and asymmetrically sharpened by distractor context. *Journal of Experimental Psychology: Human Perception and Performance*, 45(3), 336–353. <https://doi.org/10.1037/xhp0000609>

CHAPTER 4: Biased representations of attended and unattended colors recovered from human  
scalp electroencephalography

## Abstract

Human electroencephalography (EEG) studies have shown that attention enhances relevant over irrelevant visual features. Such amplitude modulations have been a major focus of feature-based attention research, but recent computational approaches have demonstrated how neural signals can also carry information about stimulus features. Here, we test whether color can be reconstructed from the scalp activity pattern using steady-state visual evoked potentials (SSVEPs) and inverted encoding modeling (IEM). Furthermore, we test how these color reconstructions vary across attentional contexts. In two experiments, participants performed a color-based selective attention task in which they detected changes in target-colored dots among distractors. The target and distractor colors were either distinct ( $180^\circ$  away on a CIELab colorwheel) or similar ( $60^\circ$  or  $63^\circ$  away). Stimuli flickered at different frequencies to elicit separable SSVEPs that we used to generate color-selective response profiles based on their spatially distributed patterns. In both experiments, we found that information about the target colors was reliably recovered from single trials with IEMs, confirming the usefulness of this technique for investigating feature representations. Model-based reconstructions were stronger for more distinct colors and showed evidence for biased representations for more similar colors, which may act to better separate targets from similar distractors. Broadly, these findings demonstrate that SSVEPs along with IEMs can be used to investigate non-spatial visual features that may not be as well represented as spatial information in event-related potentials or alpha-band signals.

## Introduction

Selective attention is the process by which specific salient or goal-relevant information in the environment is highlighted and prioritized. Many research findings indicate that attention is not solely a function of a specific target stimulus, but that distractor information also affects the way attention operates (Duncan & Humphreys, 1989; Geng & Witkowski, 2019; Navalpakkam & Itti, 2007). Models of attention, such as the feature-similarity gain and biased competition models, propose that attention acts to enhance the processing of target items or features over other, task-irrelevant items (Desimone & Duncan, 1995; Treue & Martinez-Trujillo, 1999). However, precisely which items are enhanced by attention depends on the degree of similarity between targets and distractors, because perceptually more similar features are represented nearby each other on the cortical map; thus, assuming a gradient of attentional gain modulation, those features that have greater representational overlap with the target will also receive some enhancement that monotonically falls off as feature similarity decreases (Martinez-Trujillo & Treue, 2004; Treue & Martinez-Trujillo, 1999). Indeed, behavioral findings clearly demonstrate that attentional selection is less efficient when the similarity between targets and distractors increases (Arun, 2012; Chapman & Störmer, 2022; Duncan & Humphreys, 1989; Nagy & Sanchez, 1990). Despite this, most studies investigating the neural effects of attention on target and nontarget representations have used contexts where target and distractor features are largely distinct (e.g., attend red vs green; upwards vs. downwards motion, etc.; Andersen, Hillyard, & Müller, 2013; Müller et al., 2006; Sàenz, Buraças, & Boynton, 2002). While this work has confirmed that attention leads to increased processing of targets relative to distinct distractors (Andersen et al., 2013; Sàenz et al., 2002), it remains unknown exactly how the neural processes of feature-based attention are impacted as target-distractor similarity changes.

Alongside research investigating the mechanisms of feature-based attention, computational methods developed over the past 20 years have illuminated our understanding of how different visual features are represented in the brain using a range of neural measures. For example, Kamitani and Tong (2005) first demonstrated that the pattern of activations across voxel in different regions of visual cortex, as recorded by fMRI, could be used to decode the orientation of a stimulus. Brouwer and Heeger (2009)

expanded on this work, using novel model-based reconstructions to examine how color is represented in different regions of visual cortex. Their model showed that fMRI responses in V1 carried more opponent-like color representations, while V4 carried a representation that was more consistent with the perceptual experience of color. More recent work has demonstrated that the pattern of activity in EEG or MEG recordings also carries information about a perceived or remembered non-spatial feature, such as color or orientation (Bocincova & Johnson, 2019; Hajonides, Nobre, van Ede, & Stokes, 2021; Hermann, Singh, Rosenthal, Pantazis, & Conway, 2022; Rosenthal, Singh, Hermann, Pantazis, & Conway, 2021; Sutterer, Coia, Sun, Shevell, & Awh, 2021). These computational methods provide a unique opportunity to look beyond changes in neural signals under attention or other cognitive processes, and to understand more deeply how stimulus representations are represented and transformed across different task contexts.

Here, we aimed to bring these two strands of research together, to examine how the representation of visual features is transformed by attention. To this end, we recorded EEG activity while participants performed a feature-based attention task where they selectively attended to a stimulus array of a particular color, while ignoring an array in a different color (*Figure 4.1B*). To assess the mechanisms underlying attention as similarity between targets and distractors changes, we also manipulated whether the distractor color was more distinct ( $180^\circ$  distant from the target color around a color wheel) or more similar ( $\sim 60^\circ$ ) in color to the target. To independently track visual processing of the different colors, we used steady-state visual evoked potentials (SSVEPs), with each set of colored dots flickering at a distinct frequency. SSVEPs are the oscillatory response of the visual cortex to incoming flickering stimuli (Norcia, Appelbaum, Ales, Cottareau, & Rossion, 2015; Regan, 1989) and have commonly been used to investigate how attention affects the processing of features in these paradigms (e.g., Andersen & Müller, 2010; Andersen, Müller, & Hillyard, 2009; Störmer & Alvarez, 2014).

In addition to examining the amplitude of the SSVEP responses, our goal was to also employ inverted encoding models (Brouwer & Heeger, 2009; Sprague & Serences, 2015) to reconstruct the stimulus colors and examine how different attentional contexts (i.e., similar vs. dissimilar targets and distractors) modulate these population-level response profiles for color. Previous work has shown that



attention can increase the strength of reconstructed stimulus responses (Sprague & Serences, 2013; Vo, Sprague, & Serences, 2017), as well as demonstrating that SSVEPs carry information about attended orientations (Garcia, Srinivasan, & Serences, 2013). We were particularly interested in how feature-based attention would affect the shape of the color-selective response profiles in different attention conditions. Given recent behavioral evidence that attention may amplify off-tuned target features to increase separability between targets and distractors in some situations (Chapman, Chunharas, & Störmer, under review), one possibility was that during the similar target-distractor condition, we would observe a shift in position, rather than—or in addition to—a change in the strength of the stimulus reconstruction.

## **Experiment 1**

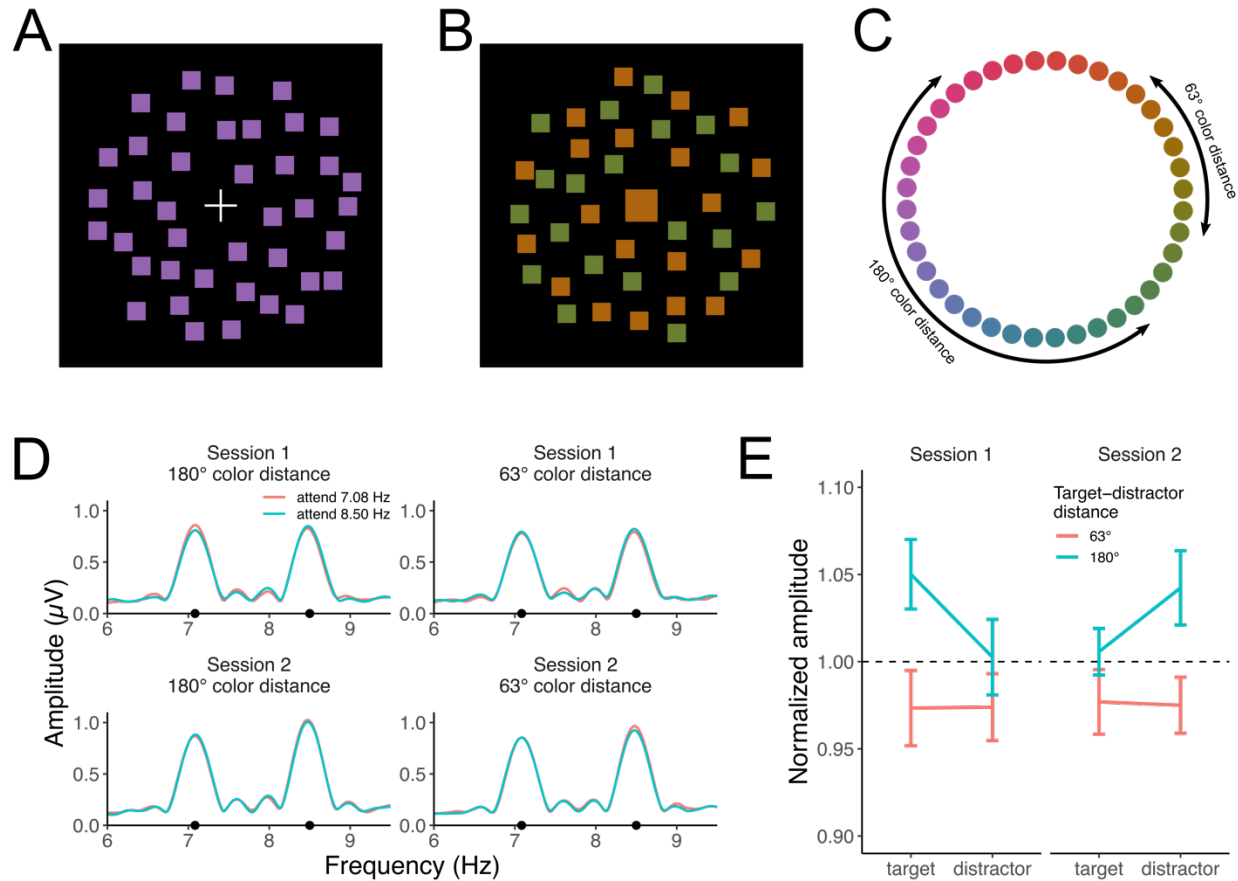
### **Method**

#### **Participants**

Twenty-two adults recruited from the UCSD community participated in two sessions of the experiment. Data from four participants were excluded due to excessive EEG artifacts, resulting in <75% of trials remaining in the selective attention task after artifact rejection. The final sample consisted of 18 participants (10 women and 8 men; 15 right-handed), aged 18-31 years ( $M = 21.3$ ,  $SD = 3.5$ ), with normal or corrected-to-normal vision, and normal color vision as assessed using Ishihara plates. Informed consent was collected prior to the experiment, and all experimental and EEG procedures were conducted as approved by the Institutional Review Board at UCSD. All participants were compensated at a rate of \$10-20/hour.

#### **Stimuli and Procedure**

Participants completed the experiment on an iiyama Vision Master Pro514 22" CRT monitor and were seated approximately 57 cm from the display. All stimuli were presented on a black background. Stimulus arrays consisted of two random-dot kinematograms (RDKs) that were presented within a circular frame ( $5.8^\circ$  visual degrees radius; *Figure 4.1A-B*). Each RDK was made up of 100 square dots (each dot  $0.27^\circ$  by  $0.27^\circ$ ) in different colors, that moved in a random direction at a speed of  $0.37^\circ/\text{sec}$ .



*Figure 4.1.* A) Example of the stimulus display for the non-selective task in Experiment 1. All dots were presented in the same color and participants had to detect a brief coherent motion event. Stimuli are not drawn to scale. B) Example stimulus display for the selective attention task in Experiment 1, with a 63° target-distractor distance. Participants attended the target-colored dots (indicated by the color of the fixation square) to detect a brief decrease in the luminance of the targets. C) Stimulus colors used in Experiment 1. All colors were drawn from a circular CIELab space and were 9° apart. D) SSVEP responses separately for each target-distractor distance for each session. E) Normalized amplitudes for each session. The attention modulation of the amplitudes was only significant for 180° target-distractor distance in Session 1.

The direction of all dots was switched randomly on average every 118 ms. To generate SSVEPs, the target and distractor RDKs were flickered on the display at two distinct frequencies (7.08 Hz or 8.5 Hz; 50% on-off duty cycle). Colors were selected from a circular CIELab color space, with  $L^* = 50$ , centered at coordinates  $a^* = 18$ ,  $b^* = 10$ , with radius 45. From this space, we selected 40 stimulus colors spanning 9° steps (*Figure 4.1B*). On each trial, the color of the target array was determined, and the distractor color was taken either 63° or 180° clockwise from the target.

Experiment 1 was conducted across two sessions, which were structured similarly except for the behavioral thresholding which was done only at the beginning of the first session. During each session participants completed two tasks: a non-selective task, and a selective feature-based attention task. The tasks were designed similarly as to allow for comparisons between the EEG responses collected during them. In the non-selective task (*Figure 4.1A*), participants were instructed to attend to the entire array to detect a motion coherence event. Trials began with the presentation of a fixation cross (each bar  $0.4^\circ$  long) for a jittered interval between 400-800 ms. Following this, the RDK was presented for 3500 ms. The coherent motion event occurred randomly between 500-3000 ms after the stimulus onset, where a proportion of the dots moved in one of the four cardinal directions (up, down, left, right) for 500 ms. Motion coherence was adjusted throughout the task to maintain accuracy levels of around 80% using a 3-down/1-up staircasing procedure. The initial motion coherence level for the task was set based on a thresholding task conducted prior to the EEG setup, consisting of 32 trials of the task with a 2-down/1-up staircase, starting at 50% coherence. Coherence thresholds were chosen as the motion coherence level corresponding to 80% accuracy. Participants completed two blocks of this task per session, consisting of 80 trials per block (320 trials total across both sessions). While color was not meaningful for the task itself, the color of the dots varied across trials for a combination of four repetitions of each of the 40 stimulus colors per session.

In the selective feature-based attention task (*Figure 4.1B*), participants were presented with two spatially intermixed RDKs in different colors. A centrally presented cue color ( $0.5^\circ$  by  $0.5^\circ$  of visual angle) indicated which set of dots to pay attention to on that trial, and also served as a fixation point for the task. Participants attended to this target color throughout the duration of the trial and were instructed to detect a brief (400 ms) decrease in the luminance of the target dots. At the end of the trial, they reported whether the target dots had decreased in luminance ('m') or not ('n'). Target dots decreased in luminance on 50% of trials, distractor dots on 25% of trials, and there was no change on the remaining 25% of trials. Trials began with the presentation of the cue square for a jittered interval of 400-800 ms, and then the RDK array for 3500 ms. The luminance decrease could occur randomly between 500-3000

ms after stimulus onset. The magnitude of the luminance decrease was adjusted during the task to maintain an accuracy level around 80%, as in the non-selective task, and was initially based on a threshold estimate from a separate behavioral task, starting at a decrease of 30% of the max luminance. Participants completed four blocks of the selective task per session, consisting of 80 trials per block (640 trials total across both sessions). On half of the trials the target and the distractor were 180° apart in the color space, and in the other half of trials were 63° apart. Each session consisted of a single repetition of each of the 40 stimulus colors as the target color, with a target-distractor distance of 63° or 180°, with distractors selected clockwise or counterclockwise relative to the target (this condition results in the same distractor when the distance is 180°), and with the target color flickered at either 7.08 Hz or 8.50 Hz. This counterbalancing also ensured that each of the 40 colors was equally often presented as a target or distractor.

### **EEG Recording**

EEG was recorded continuously from 32 Ag/AgCl electrodes embedded in an elastic headcap. Electrodes were arranged according to the 10-20 system, placed at Fp1, Fp2, F3, F4, FC1, FC2, C3, C4, Cz, T7, T8, CP1, CP2, CP5, CP6, P3, P4, P7, P8, PO3, PO4, PO7, PO8, POz, O1, O2, Oz, I3, I4, and Iz, as well as the left and right mastoids. Horizontal electrooculogram (HEOG) was recorded from additional electrodes placed on the left and right external ocular canthi and grounded with an electrode placed on the neck. Vertical electrooculogram was measured using electrodes Fp1 and Fp2. EEG signals were amplified by an ActiCHamp amplifier (BrainProducts, GmbH), with all electrodes digitized at 500 Hz and referenced online to the right mastoid.

### **EEG Preprocessing**

Continuous EEG was re-referenced offline to the average of the left and right mastoids and filtered with a Butterworth bandpass filter between 0.01 and 112.5 Hz. Data was then epoched from 500 ms before the main motion onset to 3500 ms after onset, when the dots disappeared from the display, and activity from each electrode was baselined to the average amplitude during the 500 ms baseline period. Each epoch was manually inspected for eyeblink and eye movement artifacts, and trials containing

artifacts were rejected from further analysis. This resulted in an average exclusion of 10.0% (SD = 5.7%) of trials per participant in the non-selective task and 10.4% (SD = 6.6%) in the selective attention task.

## Analysis

**SSVEP modulation by attention.** To assess the effect of attention on the stimulus-evoked SSVEP response, we calculated the average ERP across the main conditions of the selective attention task for each experiment session. For each participant, we separately averaged across trials in which the target-distractor distance was either 63° or 180°, and in which the target dots flickered at 7.08 or 8.50 Hz. ERPs were then restricted to an analysis window of 500-3500 ms following stimulus onset and linearly detrended. For each participant and condition, condition-averaged ERPs were transformed into the frequency domain using the Fast Fourier Transform (FFT), and SSVEP amplitudes were calculated as the absolute value of the complex Fourier coefficients at the two stimulus frequencies, within a band  $\pm 0.2$  Hz from the center frequency. The SSVEP response was then taken as the average of these amplitudes across a set of posterior electrodes (O1, O2, Oz, I3, I4, Iz, PO3, PO4, PO7, PO8, and POz), and was normalized within-subjects by dividing each condition response by the mean response across all conditions for each frequency bin (cf. Andersen, Fuchs, & Müller, 2011; Störmer & Alvarez, 2014).

**Single-trial SSVEP response.** We also assessed the SSVEP response evoked on single trials. EEG activity on each trial was restricted to the 500-3500 ms analysis window, linearly detrended, then FFT'd as before. For all frequencies within a  $\pm 0.2$  Hz band around each stimulation frequency, we calculated an SNR measure by dividing the power (squared amplitude) in each bin by the standard deviation in a set of neighboring bins spanning 0.4-1.0 Hz above and below that bin. The peak of this measure within the  $\pm 0.2$  Hz band was taken as the SNR for each trial.

**Inverted encoding models.** To assess the representation of color within the recorded SSVEP signal, we used inverted encoding models (IEMs) fit to the SNR across posterior electrodes. Specifically, we assumed the SNR in each electrode could be modelled as a weighted sum of tuned channels tiling the stimulus space (Brouwer & Heeger, 2009; Garcia et al., 2013). We used a set of channels modeled across color angle as half sinusoids raised to the seventh power:

$$R = \sin(0.5\theta)^5$$

where  $\theta$  is the position in color space ( $-180^\circ$  to  $179^\circ$ ), and  $R$  is the response of the channel in arbitrary units. We used eight channels centered evenly across color space ( $-180^\circ$ ,  $-135^\circ$ ,  $-90^\circ$ , ...,  $135^\circ$ ), from which a predicted channel response was calculated.

We fit IEMs to data from each subject using a leave-one-trial-out cross-validation routine. First, we trained the model using SNR from all but one trial to estimate the corresponding weights of each electrode to each channel response function. Specifically, let  $B_1$  be the SNR at each electrode on each trial ( $m$  electrodes  $\times$   $n_l$  trials),  $C_1$  be the predicted response of each of the color channels for each trial ( $k$  channels  $\times$   $n_l$  trials), and  $W$  be the weight matrix that maps between the channels and electrodes ( $m$  electrodes  $\times$   $k$  channels). The linear model describing this mapping is:

$$B_1 = WC_1$$

and the weight matrix can be obtained via a least-squares estimate as follows:

$$\hat{W} = B_1 C_1^T (C_1 C_1^T)^{-1}$$

After obtaining this weight matrix, we can then apply it to the held-out training data  $B_2$  ( $m$  electrodes  $\times$  1 trial) to obtain the estimated channel response,  $\hat{C}_2$  ( $k$  channels  $\times$  1 trial):

$$\hat{C}_2 = (\hat{W}^T \hat{W})^{-1} \hat{W}^T B_2$$

This procedure was repeated so that each trial was used once as the training data, resulting in predicted channel responses for each trial. The combined estimated channel responses were multiplied by the set of basis channels, projecting them back into the full stimulus space. Finally, the stimulus reconstruction on each trial was shifted so that the true stimulus color was  $0^\circ$  in this space, allowing for alignment across trials.

For the main analysis, we fit IEMs separately for each frequency. That is, we took the SNR for a single stimulus color (at 7.08 Hz, e.g.) regardless of whether that color corresponded to the target or distractor, and whether the target-distractor similarity was  $63^\circ$  or  $180^\circ$ . Following the alignment step, we averaged stimulus reconstructions based on the stimulus type (target or distractor) and experimental

condition (63° or 180° similarity). Thus, during training, the model had no information about stimulus type or attention condition. The condition-averaged reconstructions were then collapsed across stimulus frequency. For our primary analysis, we used the SNR from the same posterior electrodes that were used to assess the modulation of the SSVEP response.

**Quantification of stimulus reconstructions.** We first determined the strength of the stimulus reconstructions using a representation fidelity metric (Sprague, Ester, & Serences, 2016). This fidelity metric was calculated as the vector mean of the stimulus reconstruction as projected onto the aligned color angle (0°) using the formula:

$$F = \text{mean}(r(\theta) \cos \theta)$$

where  $r(\theta)$  is the strength of the reconstruction at each point along the stimulus space  $\theta$ . This metric measures the extent to which a stimulus reconstruction aligns with the true color angle.

Because fidelity can be reduced by both shifts in the position or decreases in the amplitude of the stimulus reconstruction, we also performed analyses considering the circular mean and resultant vector length of the reconstructions. To compare the average direction of two stimulus reconstructions (e.g., targets vs distractors), we computed a scaled distance metric using formula:

$$d = (\theta_1 - \theta_2)(R_1 + R_2)$$

where  $\theta_k$  are the circular means for the reconstructions being compared, and  $R_k$  are the resultant vector lengths. This metric allows for estimating the angular distance between two reconstructions, while dampening any influence of reconstructions that are close to uniform (i.e., near zero vector length).

**Statistical methods.** Normalized amplitudes were analyzed using standard repeated-measures ANOVAs. To calculate the significance of all statistical effects for stimulus reconstructions, we used permutation testing. All permutation tests were done by shuffling the labels for trials of interest several times and then recomputing relevant metrics.

For testing the uniformity of stimulus reconstructions, we randomly shuffled the angle of the true stimulus color on each trial before averaging across conditions and subjects, after which the fidelity of the reconstruction for each condition was computed. This process was repeated 10,000 times to obtain a null

distribution of fidelity metrics, and the  $p$ -value for each condition was calculated as the proportion of permutations in which the fidelity of the null distribution exceeded that of the observed data.

To assess the fidelity of reconstructions across conditions, we randomly shuffled the stimulus type (target or distractor) and target-distractor distance ( $60^\circ$  or  $180^\circ$ ) for each reconstruction. Main effects for each condition were calculated for each reshuffled dataset by taking the absolute value of the mean difference between conditions across subjects, ignoring the direction of the effects, and the interaction effect was calculated as the second-order difference between main effects. This process was repeated 10,000 times, giving null distributions for both main effects and the interaction, and the  $p$ -value for each effect was calculated as the proportion of permutations in which the null distribution exceeded the size of the difference observed in the data.

To quantify the difference in the direction of reconstructions, we randomly shuffled the labels between two conditions (e.g., targets and distractors) before averaging across conditions and participants and then calculating the scaled distance metric. This process was repeated 10,000 times, giving a null distribution of distance metrics, and the  $p$ -value was calculated as the proportion of permutations in which the scaled distance in the null distribution exceeded that in the observed data.

## Results

### Color reconstructions in non-selective task

Behavioral performance was maintained at a consistent level in the non-selective by adjusting the coherence level of the motion event. Accuracy in this task was high overall and stable across sessions (Session 1:  $M = 74.5\%$ ,  $SD = 10.5$ ; Session 2:  $M = 75.7\%$ ,  $SD = 12.0$ ),  $t(17) = 0.85$ ,  $p = .405$ ,  $d_z = 0.20$ . The coherence threshold improved slightly across sessions (Session 1:  $M = 65.8\%$  coherence,  $SD = 16.5$ ; Session 2:  $M = 63.1\%$ ,  $SD = 15.8$ ) but was not statistically different,  $t(17) = 1.63$ ,  $p = .122$ ,  $d_z = 0.38$ . We also fit IEMs to the SSVEP response in the non-selective task (*Figure 4.1A*). Model-based reconstructions were collapsed across stimulus frequency, since only a single stimulus color was present on each trial. These reconstructions revealed a highly reliable representation of the stimulus color within



the SSVEP signal (*Figure 4.2A*), which significantly deviated from uniformity,  $p < .001$ . Overall, there was high consistency between participants, as can be seen in the polar plot in *Figure 4.2B*, with 16/18 participants having reconstructions that projected along the true stimulus color. When we calculated fidelity for each of the stimulus colors used, we found a relatively consistent pattern across the entire color space (*Figure 4.2C*). Thus, SSVEP responses in the non-selective task appeared to closely track the specific stimulus color presented on the display.

### **Behavioral performance in the selective attention task**

To maintain consistent performance, we thresholded the luminance decrease in the target dots continuously during the selective attention task. Overall performance was near the expected level (63° distance:  $M = 79.3\%$ ,  $SD = 1.6$ ; 180° distance,  $M = 78.8\%$ ,  $SD = 2.0$ ) and did not differ by target-distractor distance,  $t(17) = 1.72$ ,  $p = .103$ ,  $d_z = 0.41$ . Mean performance was also consistent across sessions (Session 1:  $M = 78.6\%$ ,  $SD = 2.5$ ; Session 2:  $M = 79.5\%$ ,  $SD = 1.5$ ),  $t(17) = 1.65$ ,  $p = .117$ ,  $d_z = 0.39$ . Luminance thresholds were also similar across condition (63° distance:  $M = 34.0\%$  luminance decrease,  $SD = 6.4$ ; 180° distance:  $M = 34.4\%$ ,  $SD = 7.4$ ),  $t(17) = 0.78$ ,  $p = .447$ ,  $d_z = 0.18$ . Luminance thresholds were marginally higher in the second session (Session 1:  $M = 35.1\%$ ,  $SD = 7.6$ ; Session 2:  $M = 33.3\%$ ,  $SD = 6.8$ ),  $t(17) = 1.82$ ,  $p = .087$ ,  $d_z = 0.43$ , suggestive of an improvement in task performance over time. Overall, these findings are consistent with previous behavioral work of ours showing that attentional efficiency can be comparable when target-distractor distance is beyond approximately 45° in color space (Chapman & Störmer, 2022).

### **SSVEP responses in the selective attention task**

We then assessed how attention affected the SSVEP response to targets and distractors. SSVEP amplitudes are shown in *Figure 4.1D* for each experimental session, and normalized amplitudes for responses to targets and distractors are shown in *Figure 4.1E*. Both figures suggest that an effect of attention was observed, but only in the first session. Consistent with this, there was a significant 3-way

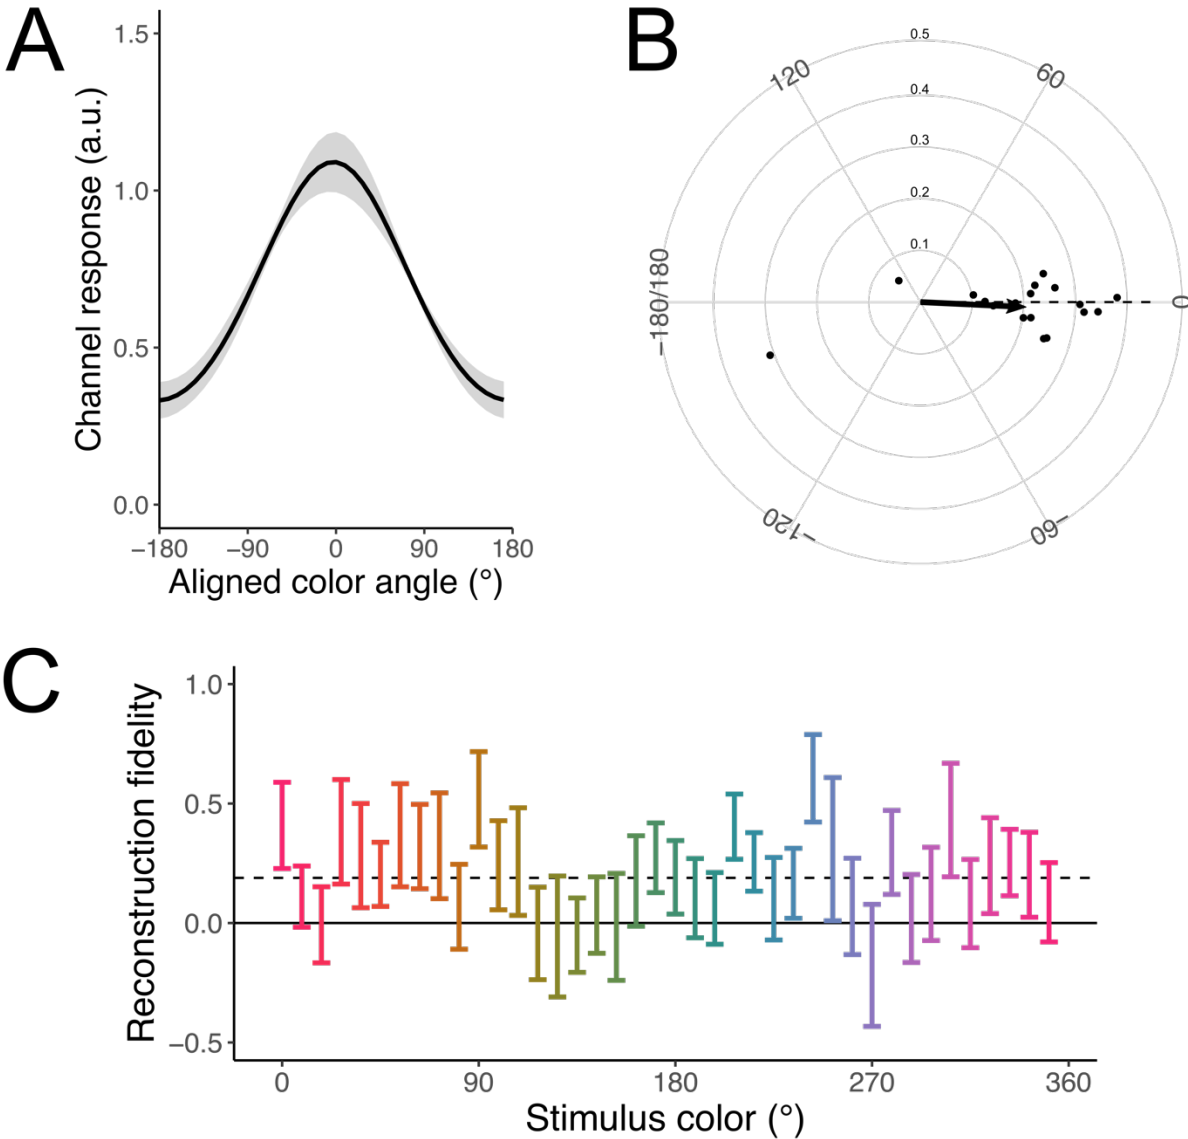


Figure 4.2. A) Average model-based reconstruction for stimuli in the non-selective task in Experiment 1. Shaded region corresponds to SEM. B) Circular statistics of reconstructions. Dots represent the circular mean and resultant vector length for individual participants, and the arrow represents the statistics of the average reconstruction across participants. C) Fidelity of the stimulus reconstructions across participants for each of the 40 stimulus colors used. Error bars are SEM, and the horizontal dashed line corresponds to the fidelity of the reconstruction across all colors.

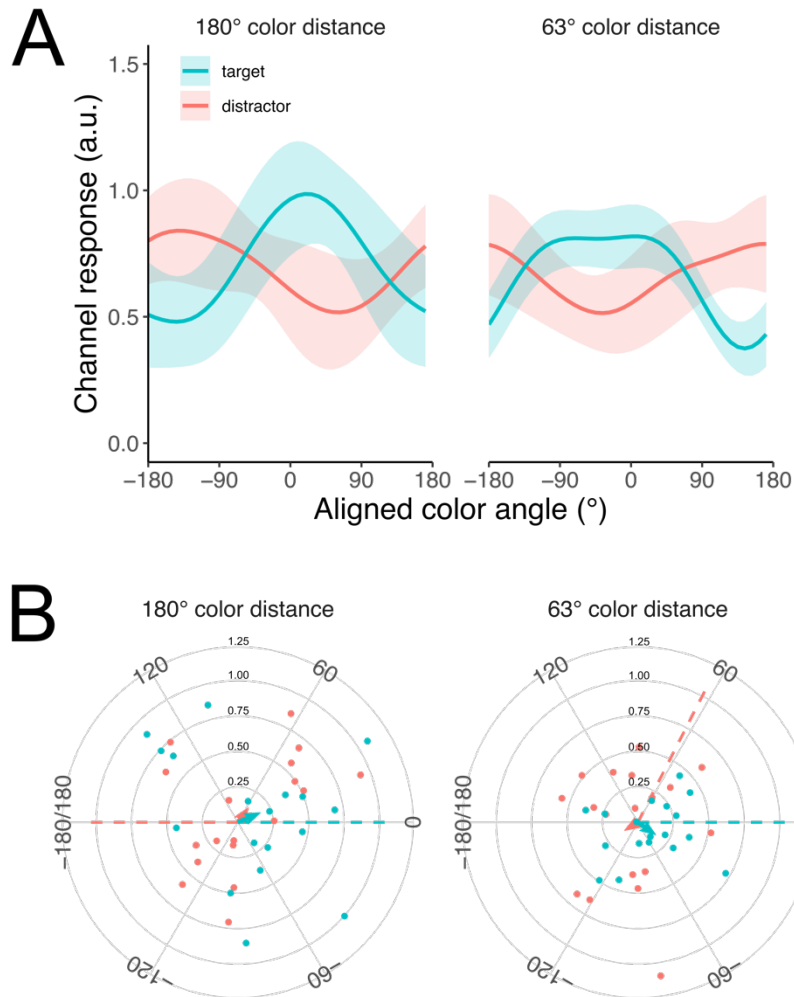
interaction in the normalized amplitudes between stimulus type, target-distractor distance, and session,  $F(1,17) = 4.50, p = .049, \eta_p^2 = .209$ . When assessed separately for each session, we found a significant interaction between stimulus type and target-distractor distance in the first session,  $F(1,17) = 4.58, p$

= .047,  $\eta_p^2 = .212$ , such that SSVEPs were stronger for targets than distractors when they were 180° apart,  $t(17) = 2.88$ ,  $p = .010$ ,  $d_z = 0.68$ , but there was no difference when colors were 63° apart,  $t(17) = 0.03$ ,  $p = .973$ ,  $d_z < 0.01$ . In the second session, there was no significant interaction,  $F(1,17) = 1.36$ ,  $p = .260$ ,  $\eta_p^2 = .074$ , suggesting that SSVEPs were not reliably modulated by attention.

### **Target and distractor reconstructions in selective attention task**

For the selective attention task, we fit inverted encoding models to assess the representational content in the SSVEP signal. The average model-based reconstructions for each condition are presented in *Figure 4.3A*. When we assessed whether the fidelity of the reconstructions in each condition deviated from uniformity, we found a trend towards significance for targets in both the 180°,  $p = .055$ , and 63° conditions,  $p = .099$ . The average of the fidelity for the two target stimuli was significantly greater than chance,  $p = .019$ , consistent with the idea that SSVEP responses carried some information about target colors. However, the model recovered no information about the distractor colors in either condition,  $p$ 's  $> .7$ . We also assessed how fidelity differed across the experimental conditions, finding a marginally significant effect of stimulus type, one-tailed  $p = .060$ , such that targets had higher fidelity than distractors on average, however there was no effect of target-distractor distance, one-tailed  $p = .424$ , and no interaction,  $p = .899$ . Overall, there was some evidence for information about target stimuli carried by the SSVEP signal, however the overall strength of the representations as recovered using IEMs was relatively weak.

There was some suggestion of a shift in the position of the target representations in the 63° condition (*Figure 4.3B*). For both distance conditions, there were large mean deviations from the true position between targets and distractors, as can be seen by the negative deflections near 0° in the reconstructions in *Figure 4.3A*, however the large variation across participants meant that neither of these position shifts reached statistical significance,  $p$ 's  $> .1$ . We also compared the position of the two target conditions directly, and while there was a numerical shift in the mean angle of 59.6°, this difference was



*Figure 4.3.* A) Average model-based reconstructions of target and distractor stimuli in the selective attention task in Experiment 1. Shaded area around each reconstruction denotes SEM. B) Circular statistics of reconstructions for individual participants. Each dot represents the reconstruction for a single subject in each condition, realigned so that the position of distractors is relative to targets as indicated by the dashed lines, and lines with arrows on each plot represent the mean angle and vector length for the reconstructions averaged across all participants.

also not significant,  $p = .248$ . Thus, while there was a shift in the position of target reconstructions when target-distractor similarity was higher, the consistency of this pattern across participants was not strong enough to establish a reliable effect.

## Experiment 2

Experiment 1 demonstrates that information about stimulus color can be recovered from the spatial topography of SSVEP responses as recorded using EEG. However, evidence for color representations in the selective attention task was relatively weak, and we found no statistical difference between the representations in the 63° and 180° target-distractor distance conditions. One limitation of Experiment 1 may be the large number of stimulus colors used, which reduced the number of trials per color that were available for training the IEMs. Thus, in Experiment 2, we aimed to replicate the previous results using a greatly reduced set of colors, which allowed for many more repetitions of the stimuli across trials.

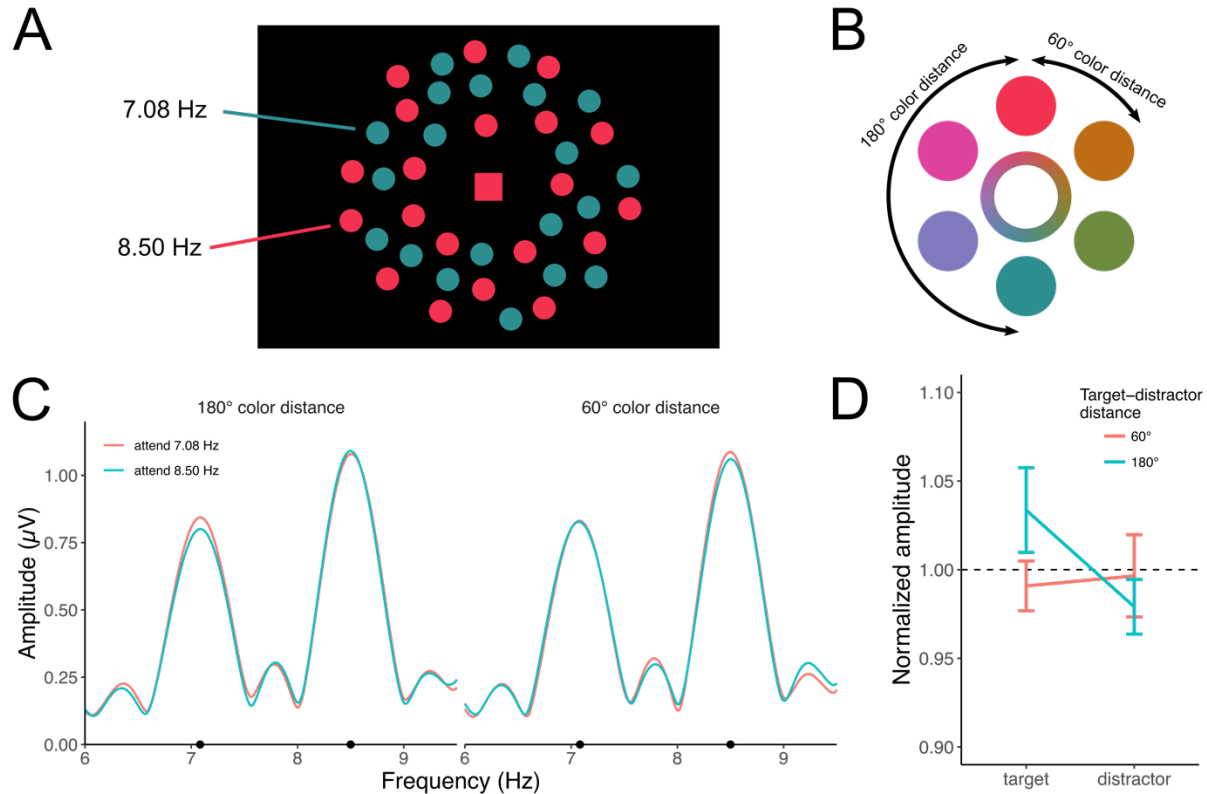
## Method

### Participants

Sixteen participants were recruited from the UCSD community. Participants were 18-22 years of age ( $M = 19.6$ ,  $SD = 1.3$ ), and had normal or corrected-to-normal vision, and normal color vision as assessed using Ishihara plates. As in Experiment 1, informed consent was collected prior to the experiment, and all experimental and EEG procedures were conducted as approved by the Institutional Review Board at UCSD. Participants were compensated for their participation at a rate of \$10 USD per hour.

### Stimuli and Procedure

The experiment was constructed similarly to Experiment 1, with a few important changes. For each experimental task, we generated two RDKs of 130 dots (7° visual angle radius, excluding the central 2.2°; each dot ~0.27° diameter; *Figure 4.4A*) that flickered at distinct frequencies (7.08 or 8.50 Hz; 50% on-off duty cycle). All dots drifted in a random direction on the display at a speed of 0.70°/sec and randomly changed direction every 60 ms on average. For Experiment 2, we selected a set of six colors from CIELab space spanning 60° steps, with  $L^* = 54$ , centered at  $a^* = 21.5$ ,  $b^* = 11.5$ , and a radius of 49 units (*Figure 4.4C*).



*Figure 4.4.* A) Example of the stimulus display in Experiment 2. Two RDKs were presented in different colors, each flickered at a different frequency. During the trial all dots moved in random directions, while participants attended to the target-colored dots (indicated by the central square) to detect a brief coherent motion event. Stimuli are not drawn to scale. B) Six stimulus colors used in Experiment 2, shown in large circles. All colors were drawn from a circular CIE Lab space (inset) and were 60° apart. C) SSVEP amplitudes plotted separately for each target-distractor distance when participants were attending to the stimulus flickering at 7.08 or 8.50 Hz. Dots along the x-axis indicate the two stimulation frequencies. D) Normalized amplitudes averaged for target and distractor stimuli as a function of the color distance between them. There was a greater amplitude for targets than distractors when they were 180° apart, but no difference when they were only 60° apart.

For the selective attention task, participants were instructed to attend to a set of target-colored dots to detect a coherent motion event. Each trial contained a cue (0.4° by 0.4° of visual angle) presented in the center of the array that was colored to indicate which set of dots participants should attend to. Both stimulus RDKs were first previewed on the display during which there was no motion. After a random interval between 800 and 1200 ms, dots in both RDKs started moving in random directions for 2600 ms. The coherent motion event occurred randomly between 400-2200 ms after the main motion onset, where

80% of the dots moved in one of the four cardinal directions (up, down, left, right) for 300 ms. At the end of the trial, participants reported whether coherent motion had occurred in the target dots (50% of trials) or not (25% distractor change, 25% no change). To set the task difficulty, participants completed thresholding runs to determine the speed of the coherent motion, in which the dot speed was adjusted using a 3-down/1-up staircasing procedure, starting at 0.74°/sec. Participants completed 48 trials of the thresholding task for each target-distractor distance, and coherence thresholds were chosen as the motion coherence level corresponding to 80% accuracy, as calculated based on fits to a logistic curve (guess rate = 50%) using the Palamedes toolbox (Prins & Kingdom, 2009). Participants completed 384 trials of the main task, balanced across each combination of target flicker frequency (7.08 or 8.50 Hz), target-distractor distance (60° or 180°), target color (one of six), and whether there was coherent motion in the target or distractor dots.

### **EEG Recording and Preprocessing**

EEG was recorded and preprocessed as for Experiment 1. For both tasks, data was epoched from 500 ms prior to stimulus onset to the end of the trial, 2500 ms after stimulus onset. Each epoch was manually inspected and trials containing eyeblink or eye movement artifacts were rejected from further analysis. This resulted in an average exclusion of 8.1% of trials per participant ( $SD = 5.9\%$ ).

### **Analysis**

Analysis details are as in Experiment 1. Because of the reduced number of colors used in Experiment 2, for IEMs we used a basis set of six channels, centered evenly across the color space (-180°, -120°, -60°, ... 120°). Stimulus reconstructions were quantified across conditions as in Experiment 1.

## **Results**

### **Behavioral performance in the attention task**

Task difficulty was adjusted for each target-distractor distance by setting the speed of the dots during the coherent motion event. Accuracy was within the expected range overall (60° distance:  $M = 76.9\%$ ,  $SD = 5.3$ ; 180° distance:  $M = 77.9\%$ ,  $SD = 6.6$ ) and did not differ by target-distractor distance,

$t(15) = 0.81, p = .430, d_z = 0.20$ . The average dot speed required to achieve this accuracy was higher for  $60^\circ$  ( $M = 0.83^\circ/\text{sec}, SD = 0.07$ ) than  $180^\circ$  target-distractor distance ( $M = 0.77^\circ/\text{sec}, SD = 0.06$ ),  $t(15) = 4.10, p < .001, d_z = 1.02$ , consistent with the expectation that task difficulty would be higher when the target and distractor colors were more similar. Overall, hit rates were high when coherent motion occurred in the target dots ( $M = 74.9\%, SD = 10.4$ ) and false alarm rates were relatively low when coherent motion occurred in distractor dots ( $M = 33.4\%, SD = 14.8$ ) or when there was no coherent motion ( $M = 7.0\%, SD = 7.5$ ).

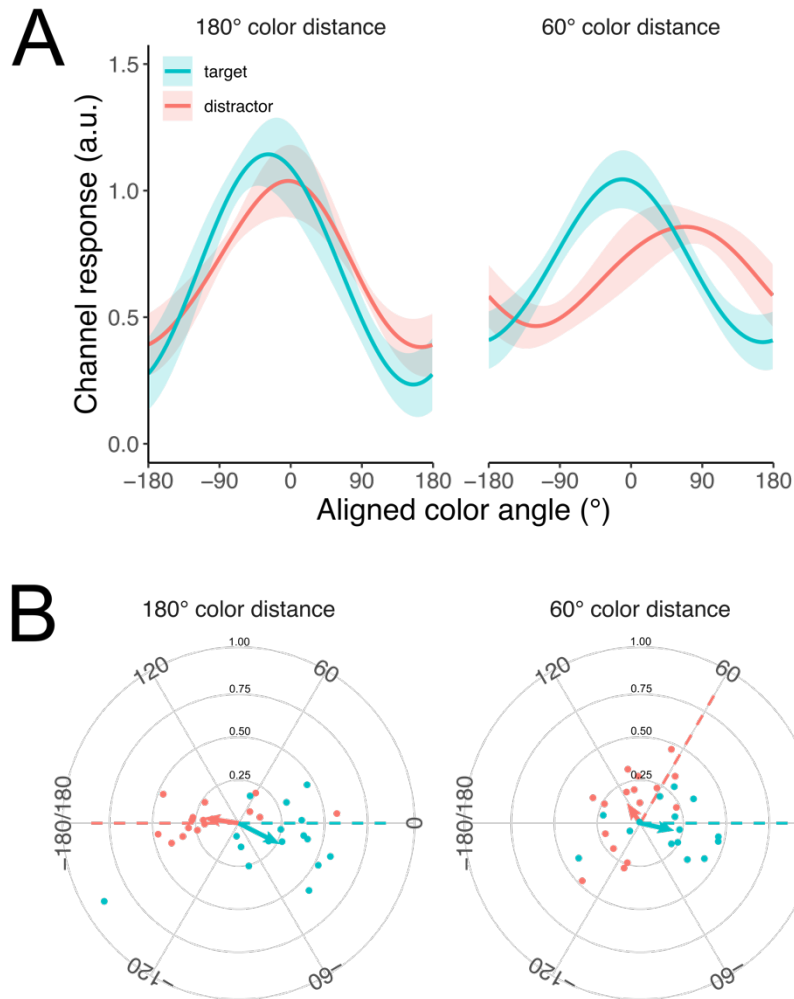
### **SSVEP responses to targets and distractors**

As in Experiment 1, we assessed how attention affected the SSVEP response to targets and distractors. As shown in *Figure 4.4C*, the amplitude of the SSVEP response was modulated by which stimulus participants were attending, but only when colors were  $180^\circ$  apart. To assess the magnitude of the modulation, we calculated a normalized amplitude measure for each participant at the two stimulus frequencies, which was then averaged across frequencies (i.e., SSVEP responses to the target averaged across 7.08 or 8.50 Hz stimulation) for each condition. The average normalized amplitudes are presented in *Figure 4.4D*, and demonstrate that attention increased the SSVEP response to targets only when colors were distinct, as shown by a significant interaction between the stimulus type and target-distractor distance,  $F(1,15) = 6.62, p = .021, \eta_p^2 = .306$ . Consistent with Experiment 1, when colors were  $180^\circ$  apart, SSVEPs were stronger for targets than distractors,  $t(15) = 2.24, p = .041, d_z = 0.56$ , whereas there was no difference in the SSVEP response when colors were only  $60^\circ$  apart,  $t(15) = 0.25, p = .806, d_z = 0.06$ .

### **Reconstruction of target and distractor colors**

We fit inverted encoding models to assess the representational content of color within the spatial topography of the SSVEP response. Average model-based reconstructions for each condition are shown in *Figure 4.5A*, and suggest that SSVEPs contained trial-by-trial information about both the target and distractor color. When we quantified these stimulus reconstructions using fidelity to test whether they





*Figure 4.5.* A) Average model-based reconstructions of target and distractor stimuli in Experiment 2. All reconstructions are aligned to a true color angle of  $0^\circ$ , so reconstructions that show a strong channel response near  $0^\circ$  reflect the ability to recover information about the true stimulus color from the inverted encoding model. Shaded area around each reconstruction denotes SEM. B) Circular statistics of reconstructions for individual participants plotted in polar coordinates. Angle around the circle represents the circular mean of the reconstruction and distance from the origin represents the resultant vector length. Each dot represents the reconstruction for a single subject in each condition, realigned so that the position of distractors is relative to targets as indicated by the dashed lines. Lines with arrows on each plot represent the mean angle and vector length for the reconstructions averaged across all participants.

deviated from uniformity, we found that both target conditions as well as distractors in the  $180^\circ$  condition had fidelity significantly above chance,  $p$ 's  $< .001$ , whereas the distractors in the  $60^\circ$  condition were not above chance,  $p = .156$ . When we assessed the difference in fidelity across conditions, we found a significant main effect of stimulus type,  $p = .039$ , where the average reconstruction for targets was

stronger than for distractors, as well as a marginally significant main effect of target-distractor distance,  $p = .059$ , with reconstructions in the  $180^\circ$  condition stronger than those in the  $60^\circ$  condition. There was no significant interaction between these factors,  $p = .281$ .

We also observed that the stimulus reconstructions in the  $60^\circ$  distractor condition appeared to be shifted in position. Given that distractor colors were always clockwise from the targets (i.e., in the positive direction along the x-axis in *Figure 4.5A*), these position shifts appear to reflect a repulsion of the stimulus reconstructions away from one another. This can be seen somewhat more clearly in the polar plots in *Figure 4.5B*, which show the true stimulus position relative to the stimulus reconstruction. To quantify these shifts, we compared the difference between the mean angle of the target and distractor reconstructions for each target-distractor distance condition. If both reconstructions are aligned to their true color angle, the difference between them should be near  $60^\circ$ . However, we observed an additional shift in the position of these reconstructions by an average of  $75.9^\circ$ , which was significantly different from what was expected by chance,  $p = .010$ . In contrast, the difference in position of the reconstructions for the  $180^\circ$  distractor conditions was only  $16.4^\circ$ ,  $p = .341$ . This suggests that attention distorted the representation of targets and distractors when they were more similar to one another, creating an exaggerated representational distance between them. This also accounts for the non-significant fidelity of the distractor reconstructions in the  $60^\circ$  condition, since the reconstructions were in a direction that was not aligned with the true color angle. Indeed, when we quantified the reconstructions using the resultant vector length, which isolates the amplitude from position offsets, we found that the  $60^\circ$  distractors were nearly significantly above chance,  $p = .067$ , while all other reconstructions remained highly significant,  $p$ 's  $< .001$ . This suggests that there was less information in the SSVEP signal for distractors at  $60^\circ$  distance, but is consistent with the finding that the representations are shifted away from the true color angle.

## Discussion

In this study, we assessed how attention modulates color representations, as measured by SSVEPs and inverted encoding models. Previous work has demonstrated that when colors are perceptually distinct, attention increases the SSVEP responses of target colors relative to distractors (Andersen & Müller, 2010; Andersen et al., 2009). We replicated this basic finding in both experiments when targets and distractors were 180° apart in color space. However, when target-distractor similarity was higher (60-63° in color space), we found no evidence for enhancement of the target SSVEP response. This lack of a modulation of the SSVEP response when target-distractor similarity was high is unlikely to reflect a lack of attention directed towards targets, given that participants performed well in our task even when targets and distractors were more similar, and in line with behavioral work that shows that feature-based attention is still relatively efficient at this level of color similarity (Chapman & Störmer, 2022). Thus, despite behavioral evidence suggesting that participants were effectively attending towards target stimuli in both conditions, we found that the SSVEP amplitude did not track this directly, suggesting that SSVEPs alone may be limited in their ability to measure how attention is deployed.

Thus, to gain a better understanding of how attention affects target and distractor color representations, we assessed stimulus reconstructions from the SSVEP responses for both targets and distractors using inverted encoding models. In Experiment 1, we found that SSVEPs carry information about stimulus color in a non-selective task, demonstrating the usefulness of this experimental approach investigating feature representations using EEG. Additionally, in the selective attention task, we found some evidence for target representations in the SSVEP signal, but not distractors. The overall representations were slightly stronger for the target colors among more distinct distractors. While there was no significant bias in the position of stimulus reconstructions for the more similar colors, there was a tendency for targets representations to be shifted away from the distractor position. In Experiment 2, we found that SSVEPs carried information about both targets and distractors, regardless of their similarity, though targets were represented relatively more strongly than distractors. Importantly, we found that when targets and distractors were perceptually more similar (60° in color space) there was a shift in the

model's predicted stimulus color, such that the representations were biased away from each other.

Overall, these findings suggest that not all effects of attention on underlying stimulus representations will bear out in the average SSVEP response to a given stimulus, since we found differences in the strength of similar target and distractor reconstructions despite no modulation of the overall SSVEPs. Thus, our work adds to research demonstrating the benefits of multivariate and model-based analyses for understanding visual processing and attention effects on visual representations.

Our findings support recent work that has shown that non-spatial features, such as color, can be decoded from EEG (Bocincova & Johnson, 2019; Hajonides et al., 2021; Sutterer et al., 2021) and MEG (Hermann et al., 2022; Rosenthal et al., 2021). Further research has established that neural activity can carry information not only about currently perceived stimuli, but also those that are maintained in working memory (Bae & Luck, 2018; Bettencourt & Xu, 2015; Ester, Sprague, & Serences, 2015; Foster, Sutterer, Serences, Vogel, & Awh, 2017; Rademaker, Chunharas, & Serences, 2019; Serences, Ester, Vogel, & Awh, 2009; Sprague et al., 2016), demonstrating the rich information available in fMRI and EEG signals that is often not captured in more traditional univariate analyses. However, this research often only assesses the representation of a single feature at a time. In contrast, feature-based attention research typically employs displays where features are spatially intermingled, so as to avoid any influence of spatial attention (Andersen et al., 2009). SSVEPs are a common measure in this line of research, allowing for separate, simultaneous measures of multiple stimuli. Our study leverages this method, demonstrating that these measures also carry information about the stimulus tagged by a specific SSVEP frequency. Previous work has either required stimuli to be presented in separate hemifields (Hajonides et al., 2021), or has implemented particular experimental designs or analysis techniques in an attempt to decompose a single neural response into measures for separate stimuli (Adam & Serences, 2021). In comparison, the use of SSVEPs, like in the current study, may provide practical advantages for experimenters looking to study the representation of multiple simultaneous stimuli.

However, our study also revealed limitations to using this method to study the representation of attended and unattended features. Notably, Experiment 1 provided relatively weak evidence for the

presence of information about target colors in the SSVEP signal, and no information about distractor colors. This is likely because we had fewer trials per stimulus color in Experiment 1: 40 colors with a maximum of 8 trials per color remaining per condition after artifact rejection, compared to six colors with a maximum of 32 trials per color in Experiment 2. Thus, variability at the level of single trials might have had a much larger impact, reducing model performance in Experiment 1 even though the SSVEP responses were relatively strong on average. However, model-based reconstructions were comparatively strong and reliable in the non-selective task, despite the overall number of trials per color being comparable to the selective attention task. This suggests there are limits to recovering information about multiple stimuli simultaneously, and that the ability to do so might be affected by competition or interactions between the neural populations that represent those stimuli. Indeed, most prior research investigating feature-based attention with SSVEPs has used a very limited stimulus space (see Adam, Chang, Rangan, & Serences, 2021, for a review), and so future studies should aim to test the generalizability of these findings with a broader set of features.

Evidence from both experiments suggests that attention may distort visual representations by shifting them further apart from one another to increase discriminability, particularly when target-distractor similarity is high. This idea is consistent with behavioral and neural findings suggesting that attention often enhances populations that are tuned away from the true target feature, because this allows for more optimal differentiation in processing between targets and distractors (Chapman, Chunharas, & Störmer, preprint; Navalpakkam & Itti, 2007; Scolar, Byers, & Serences, 2012). Other research has suggested that feature-based attention can also lead to shifts in the tuning of neural populations towards the attended feature, separating the representations of targets and distractors (David, Hayden, Mazer, & Gallant, 2008; Ibos & Freedman, 2014), and these shifts may serve to exaggerate relevant visual features (Mehrpour, Martinez-Trujillo, & Treue, 2020). Consistent with this account, IEMs in the current study revealed that the representation of targets and distractors can be shifted away from one another. We have proposed that such distortions can be understood as modulations in the representational geometry of a particular feature space, with attention distorting the overall representation of a given feature such that

more representational space is dedicated to target features (Chapman et al., preprint). In this way, attention may better separate targets and distractors when they are more similar to one another.

Overall, our findings demonstrate that the magnitude of SSVEP responses are not only modulated by attention, but that they can also track information about the features of the stimulus that elicits them. Using inverted encoding models, we found that the pattern of SSVEPs recorded by EEG carried information about the color of target (in both Experiments 1 and 2) and distractor stimuli (in Experiment 2 only). Further, while SSVEPs were not modulated by attention when target-distractor distance was low ( $60^\circ$  or  $63^\circ$ ), model-based reconstructions in Experiment 2 showed distortions in which the target and distractor representations were repulsed from one another, suggesting that attention can warp representations in a way that may exaggerate the difference between them, allowing for more effective processing. These findings fit into a growing literature that suggests that the role of attention is not simply to enhance the processing of target stimuli (e.g., through increases in gain; Martinez-Trujillo & Treue, 2004), but that attention may warp the processing of entire representational spaces in support of separating relevant from irrelevant information.

### **Acknowledgements**

Chapter 4 is currently being prepared for submission for publication. Chapman, Angus F.; Geweke, Frederik; Serences, John T.; Störmer, Viola S. The dissertation author was the primary investigator and author of this material.

## References

- Adam, K. C. S., Chang, L., Rangan, N., & Serences, J. T. (2021). Steady-state visually evoked potentials and feature-based attention: Preregistered null results and a focused review of methodological considerations. *Journal of Cognitive Neuroscience*, *33*(4), 695–724.  
[https://doi.org/10.1162/jocn\\_a\\_01665](https://doi.org/10.1162/jocn_a_01665)
- Adam, K. C. S., & Serences, J. T. (2021). History Modulates Early Sensory Processing of Salient Distractors. *The Journal of Neuroscience*, *41*(38), 8007–8022.  
<https://doi.org/10.1523/jneurosci.3099-20.2021>
- Andersen, S. K., Fuchs, S., & Müller, M. M. (2011). Effects of Feature-selective and Spatial Attention at Different Stages of Visual Processing. *Journal of Cognitive Neuroscience*, *23*(1), 238–246.  
<https://doi.org/10.1162/jocn.2009.21328>
- Andersen, S. K., Hillyard, S. A., & Müller, M. M. (2013). Global Facilitation of Attended Features Is Obligatory and Restricts Divided Attention. *Journal of Neuroscience*, *33*(46), 18200–18207.  
<https://doi.org/10.1523/JNEUROSCI.1913-13.2013>
- Andersen, S. K., & Müller, M. M. (2010). Behavioral performance follows the time course of neural facilitation and suppression during cued shifts of feature-selective attention. *Proceedings of the National Academy of Sciences*, *107*(31), 13878–13882. <https://doi.org/10.1073/pnas.1002436107>
- Andersen, S. K., Müller, M. M., & Hillyard, S. A. (2009). Color-selective attention need not be mediated by spatial attention. *Journal of Vision*, *9*(6), 1–7. <https://doi.org/10.1167/9.6.2>
- Arun, S. P. (2012). Turning visual search time on its head. *Vision Research*, *74*, 86–92.  
<https://doi.org/10.1016/j.visres.2012.04.005>
- Bae, G. Y., & Luck, S. J. (2018). Dissociable decoding of spatial attention and working memory from EEG oscillations and sustained potentials. *Journal of Neuroscience*, *38*(2), 409–422.  
<https://doi.org/10.1523/JNEUROSCI.2860-17.2017>
- Bettencourt, K. C., & Xu, Y. (2015). Decoding the content of visual short-term memory under distraction in occipital and parietal areas. *Nature Neuroscience*, *19*(1), 150–157.  
<https://doi.org/10.1038/nn.4174>
- Bocincova, A., & Johnson, J. S. (2019). The time course of encoding and maintenance of task-relevant versus irrelevant object features in working memory. *Cortex*, *111*, 196–209.  
<https://doi.org/10.1016/j.cortex.2018.10.013>
- Brouwer, G. J., & Heeger, D. J. (2009). Decoding and Reconstructing Color from Responses in Human Visual Cortex. *Journal of Neuroscience*, *29*(44), 13992–14003.  
<https://doi.org/10.1523/JNEUROSCI.3577-09.2009>
- Chapman, A. F., Chunharas, C., & Störmer, V. S. (n.d.). Optimal tuning of feature-based attention warps

the perception of visual features. *PsyArxiv*. <https://doi.org/10.31234/osf.io/hbfup>

- Chapman, A. F., & Störmer, V. S. (2022). Feature similarity is non-linearly related to attentional selection: evidence from visual search and sustained attention tasks. *Journal of Vision*, 22(8), 4. <https://doi.org/10.1167/jov.22.8.4>
- David, S. V., Hayden, B. Y., Mazer, J. A., & Gallant, J. L. (2008). Attention to Stimulus Features Shifts Spectral Tuning of V4 Neurons during Natural Vision. *Neuron*, 59(3), 509–521. <https://doi.org/10.1016/j.neuron.2008.07.001>
- Desimone, R., & Duncan, J. (1995). Neural Mechanisms of Selective Visual Attention. *Annual Review of Neuroscience*, 18(1), 193–222. <https://doi.org/10.1146/annurev.ne.18.030195.001205>
- Duncan, J., & Humphreys, G. W. (1989). Visual Search and Stimulus Similarity. *Psychological Review*, 96(3), 433–458. <https://doi.org/10.1037/0033-295X.96.3.433>
- Ester, E. F., Sprague, T. C., & Serences, J. T. (2015). Parietal and Frontal Cortex Encode Stimulus-Specific Mnemonic Representations during Visual Working Memory. *Neuron*, 87(4), 893–905. <https://doi.org/10.1016/j.neuron.2015.07.013>
- Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K., & Awh, E. (2017). Alpha-Band Oscillations Enable Spatially and Temporally Resolved Tracking of Covert Spatial Attention. *Psychological Science*, 28(7), 929–941. <https://doi.org/10.1177/0956797617699167>
- Garcia, J. O., Srinivasan, R., & Serences, J. T. (2013). Near-real-time feature-selective modulations in human cortex. *Current Biology*, 23(6), 515–522. <https://doi.org/10.1016/j.cub.2013.02.013>
- Geng, J. J., & Witkowski, P. (2019). Template-to-distractor distinctiveness regulates visual search efficiency. *Current Opinion in Psychology*, 29, 119–125. <https://doi.org/10.1016/j.copsyc.2019.01.003>
- Hajonides, J. E., Nobre, A. C., van Ede, F., & Stokes, M. G. (2021). Decoding visual colour from scalp electroencephalography measurements. *NeuroImage*, 237(April), 118030. <https://doi.org/10.1016/j.neuroimage.2021.118030>
- Hermann, K. L., Singh, S. R., Rosenthal, I. A., Pantazis, D., & Conway, B. R. (2022). Temporal dynamics of the neural representation of hue and luminance polarity. *Nature Communications*, 13(1), 1–19. <https://doi.org/10.1038/s41467-022-28249-0>
- Ibos, G., & Freedman, D. J. (2014). Dynamic integration of task-relevant visual features in posterior parietal cortex. *Neuron*, 83(6), 1468–1480. <https://doi.org/10.1016/j.neuron.2014.08.020>
- Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, 8(5), 679–685. <https://doi.org/10.1038/nn1444>
- 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.  
<https://doi.org/10.1016/j.cub.2004.04.028>
- Mehrpour, V., Martinez-Trujillo, J. C., & Treue, S. (2020). Attention amplifies neural representations of changes in sensory input at the expense of perceptual accuracy. *Nature Communications*, 11(1).  
<https://doi.org/10.1038/s41467-020-15989-0>
- Müller, M. M., Andersen, S. K., Trujillo, N. J., Valdés-Sosa, P., Malinowski, P., Hillyard, S. A., ... Hillyard, S. A. (2006). Feature-selective attention enhances color signals in early visual areas of the human brain. *Proceedings of the National Academy of Sciences*, 103(38), 14250–14254.  
<https://doi.org/10.1073/pnas.0606668103>
- Nagy, A. L., & Sanchez, R. R. (1990). Critical color differences determined with a visual search task. *Journal of the Optical Society of America A*, 7(7), 1209–1217.  
<https://doi.org/10.1364/JOSAA.7.001209>
- Navalpakkam, V., & Itti, L. (2007). Search Goal Tunes Visual Features Optimally. *Neuron*, 53(4), 605–617. <https://doi.org/10.1016/j.neuron.2007.01.018>
- Norcia, A. M., Appelbaum, L. G., Ales, J. M., Cottareau, B. R., & Rossion, B. (2015). The steady-state visual evoked potential in vision research: A review. *Journal of Vision*, 15(6), 1–46.  
<https://doi.org/10.1167/15.6.4>
- Prins, N., & Kingdom, F. A. A. (2009). Palamedes: Matlab routines for analyzing psychophysical data. Retrieved from <http://www.palamedestoolbox.org>
- Rademaker, R. L., Chunharas, C., & Serences, J. T. (2019). Coexisting representations of sensory and mnemonic information in human visual cortex. *Nature Neuroscience*, 22(8), 1336–1344.  
<https://doi.org/10.1038/s41593-019-0428-x>
- Regan, D. (1989). *Human Brain Electrophysiology: Evoked Potentials and Evoked Magnetic Fields in Science and Medicine*. New York: Elsevier.
- Rosenthal, I. A., Singh, S. R., Hermann, K. L., Pantazis, D., & Conway, B. R. (2021). Color Space Geometry Uncovered with Magnetoencephalography. *Current Biology*, 31(3), 515–526.e5.  
<https://doi.org/10.1016/j.cub.2020.10.062>
- Sàenz, M., Buraças, G. T., Boynton, G. M., 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.  
<https://doi.org/10.1038/nn876>
- Scolari, M., Byers, A., & Serences, J. T. (2012). Optimal deployment of attentional gain during fine discriminations. *Journal of Neuroscience*, 32(22), 7723–7733.  
<https://doi.org/10.1523/JNEUROSCI.5558-11.2012>
- Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human

primary visual cortex. *Psychological Science*, 20(2), 207–214. <https://doi.org/10.1111/j.1467-9280.2009.02276.x>

Sprague, T. C., Ester, E. F., & Serences, J. T. (2016). Restoring Latent Visual Working Memory Representations in Human Cortex. *Neuron*, 91(3), 694–707. <https://doi.org/10.1016/j.neuron.2016.07.006>

Sprague, T. C., & Serences, J. T. (2013). Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. *Nature Neuroscience*, 16(12), 1879–1887. <https://doi.org/10.1038/nn.3574>

Sprague, T. C., & Serences, J. T. (2015). Using human neuroimaging to examine top-down modulation of visual perception. In B. U. Forstmann & E.-J. Wagenmakers (Eds.), *An Introduction to Model-Based Cognitive Neuroscience* (pp. 245–274). New York, NY: Springer. [https://doi.org/10.1007/978-1-4939-2236-9\\_12\\_245](https://doi.org/10.1007/978-1-4939-2236-9_12_245)

Störmer, V. S., & Alvarez, G. A. (2014). Feature-based attention elicits surround suppression in feature space. *Current Biology*, 24(17), 1985–1988. <https://doi.org/10.1016/j.cub.2014.07.030>

Sutterer, D. W., Coia, A. J., Sun, V., Shevell, S. K., & Awh, E. (2021). Decoding chromaticity and luminance from patterns of EEG activity. *Psychophysiology*, 58(4), 1–21. <https://doi.org/10.1111/psyp.13779>

Treue, S., & Martinez-Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, 399, 575–579. <https://doi.org/10.1038/21176>

Vo, V. A., Sprague, T. C., & Serences, J. T. (2017). Spatial Tuning Shifts Increase the Discriminability and Fidelity of Population Codes in Visual Cortex. *The Journal of Neuroscience*, 37(12), 3386–3401. <https://doi.org/10.1523/JNEUROSCI.3484-16.2017>

## CONCLUSION

To experience a coherent and meaningful world, we must choose what information to prioritize and process. Selective attention refers to the cognitive function that allows us to select information from the constant stream of sensory inputs. Attention has been argued to operate in multiple domains, allowing selection of spatial locations (Carrasco, 2011; Posner, 1980), features (Carrasco, 2011; Driver & Baylis, 1989; Maunsell & Treue, 2006), or objects (Duncan, 1984; O’Craven, Downing, & Kanwisher, 1999). In this chapter, I will summarize core findings from the literature on feature-based attention and discuss them in the context of spatial and object-based attention, with the ultimate goal of integrating across findings to conceive of a more general framework of selective attention. I will focus in particular on the effects of attention on visual processing, arguing that the organization and architecture of the perceptual representations themselves are a major constraint underlying the limits and abilities to select information across domains. I propose that instead of dividing attention into distinct subtypes, it is useful to carve out the similarities of attentional operations across these domains, in the hope that this will reveal the canonical computations and processes of attention more generally. According to this view, mechanisms of selective attention are shared across different representational spaces, while the representational spaces themselves become distinctive factors that need to be understood. Moving forward, I thus believe it is important to characterize the architecture of the representational spaces over which attention operates, and understand not only how attention affects processing of information that is ‘attended’ or ‘unattended’, but also how attentional selection may alter and reshape their representational geometry.

### **What is feature-based attention?**

Feature-based attention refers to the way by which attention can select between different features within a specific feature dimension (e.g., between the colors red and blue). Investigations of this type of feature-based selection have provided useful insights into how visual features are processed in the brain, since many of the commonly used feature dimensions have been extensively researched and characterized in terms of their psychological (i.e., representational structure) and physiological (i.e., single-cell and

population tuning) properties. One of the most prominent tasks to study feature-based attention is visual search, in which participants are asked to find a target that is usually defined based on a particular visual feature, for example its color or orientation, among other nontargets (see *Figure 3.1A*, e.g.). These tasks nicely map on to how we often use feature-based attention in real-world situations (such as finding your keys on the desk or a friend in a crowd), and have provided important insights into what constraints feature-based selection. For example, early visual search studies demonstrated the importance of considering perceptual similarity between targets and nontargets as well as the feature variability of the nontargets (e.g., Duncan & Humphreys, 1989; Nagy & Sanchez, 1990; Wolfe & Horowitz, 2004). Specifically, finding a target amongst similar distractors takes longer than finding a target among distinct distractors, and higher nontarget variability also slows down visual search (Chapman & Störmer, 2022; Duncan & Humphreys, 1989; Wolfe & Horowitz, 2004). This has been demonstrated with simple features such as color, in which case search efficiency decreases exponentially with increasing target-distractor similarity (Chapman & Störmer, 2022; Nagy & Cone, 1996; Nagy & Sanchez, 1990), but also with other feature dimensions (Arun, 2012; Blough, 1988).

While the visual search paradigm has provided much evidence on the role of target-distractor similarity, it is not an ideal task for examining the mechanisms of attentional selection for features per se, because visual search also always includes a spatial component, with items of different features being presented at different locations on the display. Thus, visual search experiments cannot fully dissociate selection based on features alone from the role of features in guiding spatial attention to a location (Andersen, Müller, & Hillyard, 2009; Shih & Sperling, 1996; Vierck & Miller, 2008). Thus, researchers have developed other sustained feature-based attention tasks that use overlapping random dot kinematograms (RDKs) in which target and distractor features are presented at the same spatial location (Sàenz, Buraças, & Boynton, 2003; Sàenz et al., 2002; Treue & Martinez-Trujillo, 1999). In these tasks, groups of tightly intermingled stimuli are presented, and different groups of stimuli comprise a different feature (e.g., one group is red, another green; one moves upwards, the other moves downwards, etc., see *Figure 4.1A*). Because the stimuli move and overlap randomly and thus occupy the same area of space,

there is no spatial or depth separation to differentiate them. Such tasks provide a useful tool for investigating mechanisms of feature-based attention independently of their ability to guide spatial attention, and have been instrumental in showing that feature-based attention can operate in a spatially global way across the visual field. For example, attending to the same color in both the left and right visual fields results in enhanced perceptual discrimination performance compared to attending to different colors in each visual field (Sàenz et al., 2003), suggestive of competition between features despite their spatial separation. Further evidence comes from neural recordings demonstrating that attention to specific colors or motion directions results in increased activity in the visual cortex in the hemisphere that is not processing the currently attended feature (Lustig & Beck, 2012; Sàenz et al., 2002; Serences & Boynton, 2007; Treue & Martinez-Trujillo, 1999). Such findings suggest that feature-based attention acts to enhance processing in all neural populations that are tuned to the currently attended feature, regardless of spatial location. Importantly, similar to what has been shown for visual search, target-distractor similarity modulates the efficiency of attentional selection in these sustained attention tasks (Chapman & Störmer, 2022; Störmer & Alvarez, 2014), suggesting that both tasks index the overlap in the underlying target and distractor representations.

### **What is a visual feature?**

One of the most challenging aspects of studying feature-based attention is defining what constitutes a feature or feature dimension, and there have been several attempts to compile lists of proposed features, or more specifically which features guide spatial attention (Wolfe & Horowitz, 2017; Wolfe & Utochkin, 2019). While most researchers agree that certain stimuli make up a feature dimension straightforwardly (e.g., color, motion direction, orientation), a set of necessary and sufficient criteria for defining such features remains undetermined. However, some aspects of processing are at least shared among the commonly accepted feature dimensions. First, features have a psychologically relevant structure, in that they can be represented in terms of their dimensionality: one-dimension for orientation (0-180°) and motion direction (0-360°), or three-dimensions for color (hue, saturation, and luminance, for

one—though certainly not the only—representational space). Second, that this structure is reflected in behavior, such as through the effects of similarity on perceptual performance: features that are more similar in the representational space are responded to in more similar ways than features that are more distinct (i.e., stimulus generalization; Shepard, 1987; Treue & Martinez-Trujillo, 1999). Third, that the representational structure has associated neural underpinnings—such as the maps of orientation in V1 (Hubel & Wiesel, 1968; Roth, Kay, & Merriam, 2022), motion directions in hMT/V5 (Albright, Desimone, & Gross, 1984), or color in ventral visual cortex (Bohon, Hermann, Hansen, & Conway, 2016; Conway & Tsao, 2009)—and that the structure of these neural representations are similar to the psychological structure of the feature space. This could occur, for example, in map-based structures where similar features are represented nearby each other within a map (Franconeri, Alvarez, & Cavanagh, 2013). Alternatively, while color is represented in terms of opponent-processes in early visual regions, this representation becomes more reflective of the psychological representation (e.g., represented in terms of perceived hue) in later visual regions, such as V4 (Brouwer & Heeger, 2009).

While these aspects capture some of the shared properties of the accepted feature dimensions, they do not necessarily constitute a complete definition of what a feature is. Under this framework, it is unclear whether high-level visual properties such as animacy or real-world size would be considered feature dimensions, or whether object categories like faces, that have well characterized dimensionality in behavior (Leopold, O’Toole, Vetter, & Blanz, 2001; Todorov, Said, Engell, & Oosterhof, 2008) and neural representation (Chang & Tsao, 2017), should fall under this definition. Indeed, the distinction between features and objects is complicated. At what point does a collection of features become an object, and how is this reflected in the brain and behavior? The relationship between features and objects is also important for attention and will be discussed further below.

### **How is feature-based attention related to spatial attention?**

By far the most studied form of attention is spatial attention, which describes how different locations within the visual field can be prioritized for selection. Original studies investigated how

attention can be allocated towards the left or right side of space broadly (Posner, 1980; Posner & Petersen, 1990), while later work showed that the focus of spatial attention can be much more precise, though varies in its precision across the visual field (He, Cavanagh, & Intriligator, 1997; Intriligator & Cavanagh, 2001). Unlike feature-based attention, which is spatially global, spatial attention selects for specific, relevant locations within the visual field, and presumably selects all information within the attended spatial extent (i.e., spatial attention is global for features). That is, spatial and feature-based attention appear to be complementary. Like many visual features, space has a defined dimensionality that is relevant for behavior and is encoded within the visual system, and thus it could be argued that space is a basic visual feature (like color, orientation, etc.). Some researchers have stated this position (Maunsell & Treue, 2006), while others remain agnostic but imply in their model framework that similar principles might underlie both modes of attention (Boynton, 2009; Reynolds & Heeger, 2009). However, spatial and feature-based attention are more commonly portrayed in the literature as separate processes that may also share some general mechanisms.

In spatial attention, it has been extensively demonstrated that attentional resources can be allocated narrowly to small regions in space or relatively broadly to select large swaths of the visual field. The size of the spatial extent can be determined by the size of relevant stimuli, the anticipated location of target stimuli, or the predictiveness of spatial cues (Castiello & Umiltà, 1990; Eriksen & St. James, 1986; Herrmann, Montaser-Kouhsari, Carrasco, & Heeger, 2010; Itthipuripat, Garcia, Rungratsameetaweemana, Sprague, & Serences, 2014). As broader spatial regions are attended, the efficiency of processing decreases, as would be expected if there was a limited attentional resource that was distributed over varying regions of space, as is argued in theories such as the zoom-lens model of spatial attention (Eriksen & St. James, 1986; Eriksen & Yeh, 1985). In the feature-based domain, there is also evidence for flexible allocation of attention. In one study, participants were presented with cues that had high or low predictiveness with respect to an upcoming target orientation (Herrmann, Heeger, & Carrasco, 2012). While performance was lower with cues that were less predictive, a breakdown of performance by the exact target orientation demonstrated that participants were distributing attention relatively evenly

throughout the range of possible orientations. In **Chapter 2**, I extended upon this work to show that participants can sustain attention to broad ranges of target colors (up to at least 120° in a circular color space) while still maintaining relatively high levels of task performance. Like Herrmann and colleagues (2012), I observed that attention appeared to select the entire range of relevant colors relatively uniformly. These findings provide evidence that feature-based attention, like spatial attention, can be flexibly and efficiently adjusted as required based on the range of currently relevant features.

One of the most prominent theories of feature-based attention states that gain is applied to the target feature which then drops off monotonically as a function of the similarity between the neuron's preferred feature and the to-be-attended target feature (Treue & Martinez-Trujillo, 1999). This feature similarity gain model has received much support from both behavioral and neural data. For example, neural recordings in the middle temporal area (MT) to an unattended moving random dot pattern positioned inside a neuron's receptive field were differentially modulated depending on what type of motion a macaque was paying attention to at another location: when attending to the same motion (e.g. upward), the neural responses to the unattended stimulus matching this feature (i.e., upward) were enhanced; importantly, this enhancement dropped off as the feature similarity between attended and unattended stimuli decreased (Martinez-Trujillo & Treue, 2004; see also, Bichot, Rossi, & Desimone, 2005). While this feature gain similarity model seems to capture many properties of feature-based attention, it seems to not generalize well to all contexts. According to an optimal tuning account, for example, in situations where target and distractor features are highly similar, it may be instead optimal to apply gain to neurons that are tuned *away* from the target feature to increase the signal-to-noise ratio between targets and nontargets (Kerzel, 2020; Navalpakkam & Itti, 2007; Scolari, Byers, & Serences, 2012; Scolari & Serences, 2009). Several studies have now provided computational and empirical support for such optimal tuning in feature space. For example, Navalpakkam & Itti (2007) showed that after searching for a particular orientation in a difficult search task with highly similar distractors (55° among 50°, respectively), on randomly interspersed probe trials, participants reported a target orientation biased away from the distractor (i.e., 60°) more often than the actual target. In **Chapter 3**, I found that when



selecting an oddball color in a briefly presented visual search display, participants also have biased reports of the target color, as if their perception of the target was shifted away from the distractor color. The magnitude of this bias was greater when target and distractor colors were more perceptually similar, consistent with predictions of optimal tuning. Further, in **Chapter 4**, I found that this perceptual shift was encoded in the pattern of SSVEP responses recorded using EEG, suggesting that these behavioral biases could be the result of distortions in the representation of color in the visual system. Interestingly, these perceptual ‘repulsion’ effects due to attention resemble the perceived position shifts of targets in spatial attention tasks, where attending to one location induces the target location to be perceived shifted away from the distractor location (Arnott & Goodale, 2006; Pratt & Arnott, 2008; Pratt & Turk-Browne, 2003; Shim & Cavanagh, 2005; Suzuki & Cavanagh, 1997). Thus, attention to locations and visual features can both lead to perceptual distortions, such that the target is perceived to be pushed away from the distractor representation, ultimately increasing the feature distance between them, in support of optimal selection.

### **How is feature-based attention related to object-based attention?**

Numerous studies have reported that attention can operate not only over locations or visual features, but over full-fledged visual objects. Exactly what an object is in the context of attention is an important theoretical topic in its own right (Chen, 2012; Scholl, 2001), but here I will primarily focus on the relationship between features and objects, and how objects may differ from features when it comes to selection. Objects are naturally composed of basic features, and a number of major theories argue that the goal of attention is to bind these features together for object recognition (Treisman, 1986; Wolfe, 2021). Other research indicates that attention is not necessary to bind features into objects, but that objects are formed preattentively and subsequently constrain attentional selection. Initial evidence for this claim came from studies showing that participants were more accurate in reporting targets that were superimposed on the same object relative to separate objects (Duncan, 1984), or that a faint stimulus was detected more rapidly at an uncued location if that location happened to be part of an attended object relative to an unattended object (Egley, Driver, & Rafal, 1994; Francis & Thunell, 2022; Malcolm &

Shomstein, 2015). These studies have generally been taken as evidence that attention is object-based, meaning that it spreads faster or prioritizes locations more within an object than across objects (Lamy & Egeth, 2002; Shomstein & Yantis, 2002). Studies using electrophysiology have supported these behavioral effects, showing that the visual-evoked N1 component is enhanced more for stimuli appearing at a within-object location relative to a between-object location (Martínez et al., 2006; see also, Ekman, Roelfsema, & de Lange, 2020). In these studies, “objects” are usually operationalized as line drawings of rectangles (and in some cases also illusory objects; Martinez, Ramanathan, Foxe, Javitt, & Hillyard, 2007, e.g.), with edges fulfilling the role of object boundaries.

Other studies have used images of real-world objects (faces, houses, cars, or bodies, e.g.), arguing that attention can be directed to these real-world objects and enhance their processing. To differentiate object-based and spatial attention, these studies often overlay images of real-world stimuli on top of each other (similar to feature-based attention studies), and instruct participants to attend to one image or object category over the other (O’Craven et al., 1999; Tong, Nakayama, Vaughan, & Kanwisher, 1998). These studies showed that when participants attended to faces overlaid on houses to detect a brief movement of the attended stimulus, neural responses in category-selective visual cortex corresponding to the attended objects were enhanced relative to the unattended object (FFA and PPA respectively). In other cases, participants were asked to search for people or cars in real-world scene images, and neural processing was enhanced for these object categories across the visual field, in agreement with object-based selection (Peelen, Fei-Fei, & Kastner, 2009; Peelen & Kastner, 2011). Similarly, when participants attended to faces at one location in the visual field, processing of faces at another task-irrelevant location was selectively enhanced (Störmer, Cohen, & Alvarez, 2019), again suggesting that tuning attention to an object category (a face) can enhance visual processing of that category across the visual field, similar to what has been observed for the selection of basic visual features. Collectively, these types of studies have been taken as evidence that attention can operate over entire real-world objects or even well-defined object categories.

From the perspective of attention, is a distinction between object- and feature-based selection a useful theoretical dichotomy? Below I will outline why this distinction may be artificial the way it is defined in many experimental designs and will propose how feature- and object-based attention could be reunited under the same umbrella. First of all, as alluded to above, the main distinction often made between object- and feature-based attentional mechanisms is based on whether irrelevant, initially not selected (i.e., attentionally cued) features are also processed in a prioritized way when they are part of the same object, relative to another object. “Being part” of the same object is then either defined in terms of clear edges (Egly et al., 1994), spatial overlap between feature groups (Ernst, Boynton, & Jazayeri, 2013; Schoenfeld, Hopf, Merkel, Heinze, & Hillyard, 2014), or perceptual similarity (Adamian, Andersen, & Hillyard, 2020; Bartsch, Donohue, Strumpf, Schoenfeld, & Hopf, 2018; Boehler, Schoenfeld, Heinze, & Hopf, 2011; Chapman & Störmer, 2021; Lustig & Beck, 2012). For example, when participants are instructed to attend to the color of a moving stimulus, attention first enhances processing in brain regions that respond to color, and subsequently spreads to motion-related regions (Ernst et al., 2013; Schoenfeld et al., 2014; see also, LaBerge & Brown, 1989; Roelfsema, Lamme, & Spekreijse, 1998). Thus, attention is allocated to one individual feature volitionally first, which then leads to an enhancement of the secondary feature of the same “object”.

In **Chapter 1**, I demonstrated that spreading of attention between features within the same object can lead to enhancement of the secondary feature throughout the visual field, as if it was enhanced in a global feature-based way. This argues strongly against strict object-based theories, which claim that attention should be restricted to the boundaries of the attended object. However, it is worth considering what the object is in this case. An array of colored dots that move in a certain direction are effectively bundles of features that are grouped together based on the spatial overlap of those features (i.e., their joint color and motion direction). More generally, stimuli that move in the same direction or are of the same color can be grouped together into a coherent perceptual unit, something we may refer to as an “object”. The visual system constantly exploits these properties of the visual inputs (such as feature similarity or adjacency), because they are likely reflections of distinct objects in the real world (Wertheimer, 1923).

The question then becomes as to whether separate attentional mechanisms are needed to explain selection of the individual features (i.e., feature-based attention) compared to these grouped feature bundles (i.e., object-based attention). As argued in **Chapter 1**, I believe these effects can readily be explained by mechanisms typically thought of as feature-based: Neural populations tuned to the different features of an object are naturally linked through overlapping spatial receptive fields or similarity within a feature map, which can provide a scaffolding that allows for attention to spread between these populations. Reciprocal connections between populations tuned to the same feature at different locations, as well as with higher-level complex visual feature maps that code for faces, houses, and other categories, can then propagate these attention effects in a global manner.

Thus, I argue here that supposed “object-based” effects of attention can be readily explained within the framework of feature-based attention. Any attempt to revise object-based theories to better account for these findings should more explicitly define what an “object” is in terms of attentional selection: is a stimulus with multiple features (e.g., green and moving upwards) sufficiently “object-like”, or are other conditions necessary for objecthood? Such theoretical work would help delineate what is a feature and what is an object, and how each one constrains attentional selection. However, within current frameworks, I believe that feature-based theories offer a much more parsimonious, though perhaps not complete, explanation of attentional selection.

### **Towards an integrated model of attention across features, locations, and objects**

How do we bridge the divide between spatial, feature-based, and object-based theories of attention? As I’ve argued, a separation between object- and feature-based selection mechanisms is not useful as it does not provide additional explanatory power for the effects observed in the literature. Instead, it seems more useful to think about feature- and object-based attentional mechanisms along a continuum, where attention can select items based on a feature, or groups of items that are similar in terms of their features or spatial location (and thus have been segmented by the visual system

preattentively), and that attention operates on these segmented perceptual units (Scolari, Ester, & Serences, 2014).

If feature and object-based attentional mechanisms can largely be seen in combination, the question remains of whether spatial attention is unique, as many researchers have proposed in the past. As outlined above, several properties of spatial selection seem to map on to properties of feature-based attention quite well: flexibility in the size of the focus of attention, capacity limits of selection, or perceptual repulsion, are observed both in location space and feature space. This raises the question as to whether separate attentional mechanisms are necessary, or whether the computations and processes involved in attention are often shared but operate over different representational maps. Orientation and color are distinct visual features, with different dimensionality and representational structure, and are represented by neural populations in different visual cortical regions, but differences in the effects of attention on these features doesn't necessitate that we consider separate "orientation-based" or "color-based" attention systems. Likewise, differences in spatial and feature-based should not necessarily require separate theories, as long as the differences observed across paradigms can be explained in terms of the underlying representational spaces rather than how attention acts on those representations.

How could we conceive of a general theory of attentional selection? One possibility, that I will argue for here, is that attention acts on segmented *perceptual units*. By perceptual units, I mean any representation in which some stimulus has been "grouped" based on visual properties that determine what aspects belong to the group, while excluding other aspects that do not belong to the group. Perceptual grouping of visual arrays occurs somewhat preattentively, essentially providing a 'sketch' of different features present in the visual field for attention to act upon (Marr, 1982; Treisman, 1986). Although the basic representations formed by perceptual grouping are likely not completely discrete, they are spatially segmented (either directly, or through modal or amodal completion, e.g., Kanizsa, 1976; Leshner, 1995) and collate similar features, meaning that they consist of important spatial and feature information, and could also form the basis for object representations. Perhaps more concretely, consider the attentional "targets" from the sustained attention tasks reviewed earlier: although the RDKs are made up of many

individual dots, the overall impression is of a coherent stimulus defined by one feature or another. The similarity between dots based on their color or direction of motion provides a basis for the visual system to group them together, allowing for attention to act upon them simultaneously. In line with this, research has demonstrated that Gestalt cues, such as similarity, collinearity, and common fate can act as scaffolds to allow for the spreading of activation from one item to another (Wannig, Stanisor, & Roelfsema, 2011). This explanation also fits well with the effects of target-distractor similarity on attention since the overlap in representations naturally impacts the ability to restrict attention just to the target item. Likewise, research showing that attention can distort the perception of spatial locations (Pratt & Arnott, 2008; Suzuki & Cavanagh, 1997) and features (see **Chapter 3 & 4**) could be interpreted as revealing how attention can exaggerate differences between stimuli in a way that better separates highly similar target and non-target groups.

The idea that attention is dependent on perceptual grouping is not new (Driver, 2001; Driver, Davis, Russell, Turatto, & Freeman, 2001; Wolfe, 2021), particularly in how segmentation of the visual field via grouping mechanisms constrains attention. However, most research seems to have moved away from this way of thinking, towards consideration of attention within distinct stimulus domains (i.e., features vs objects vs space), which may ultimately lead us to miss the forest for the trees. Considering the nature of selection more broadly offers a way to direct research away from questions solely about how attention affects the processing of one target or distractor feature (e.g., “how does feature-based attention enhance targets or suppress distractors?”), and towards questions about how attention alters the flow of information in support of goal-directed behavior more generally (e.g., “how does attention transform representational spaces?”). In **Chapter 3**, for example, I found that selecting a target color distorts the representational feature space in a way that separates targets and distractors and, critically, these attention-induced shifts also had consequences across large swaths of the rest of the feature space. That is, selection was not just about the particular target and distractor colors that were shown on any given experimental trial but involved the warping of the entire representational space of color. Such warping of feature spaces due to attention have also been observed for more complex and higher-level dimensions,

such as semantic categories, where attention acts to separate neural population responses along the behaviorally relevant dimension, resulting in warping of the underlying representational structure (Çukur, Nishimoto, Huth, & Gallant, 2013; Nastase et al., 2017). Along these lines, recent computational approaches have explored how activity in neural systems follows dynamic and flexible trajectories during working memory and decision making (Meyers, 2018; Sreenivasan, Curtis, & D’Esposito, 2014; Stokes et al., 2013). Potentially, changes in the representational structure under attention may help shape these dynamics, essentially forging a pathway for relevant information to flow more efficiently. Thus, attention not only influences processing of information for the attended location, feature, or feature bundle (i.e., object), but can sculpt the entire representational geometry in a way to best support behavior. These changes can occur across representations of location and features, including groups of feature representations (perceptual units, or objects). Most broadly, this means that understanding these representational spaces—their organization, architecture, and flexibility—is a critical step towards explaining attentional selection and its limits.

## **Summary**

In this dissertation, I have explored a series of questions related to feature-based attention. In **Chapter 1**, this revealed that object-based theories may have a limited ability to explain findings related to how attention spreads between features within an object and across objects. In **Chapter 2**, I showed how the flexibility of feature-based attention is perhaps much greater than previous theories had suggested. Finally, in **Chapter 3 and 4**, I demonstrated with behavioral and neural evidence, respectively, that feature-based attention leads to distortions in the representational space that underlies selection. Collectively, these studies further our understanding of the nature of feature-based attention in the visual system, while also challenging the necessary division of attention into separate domains. While I have characterized this work in terms of feature-based attention, given the previous theoretical and empirical work that was foundational for my own research, over time I have come to consider that each of these studies tap into some broader understanding of attentional selection. Aspects of ideas I have presented in

each chapter are found easily in work on spatial attention or object-based attention, suggesting that these divisions might obscure a more general understanding of attention. Instead of debating which domain reflects the true “unit” of attentional selection, I believe that a productive direction for future research should focus on how attention affects representational structures, combining behavioral measures with neural recordings across different scales (from fMRI and EEG to multi-unit recordings in animal models). Insights from computational neuroscience, such as the framework provided by representational geometry (Kriegeskorte & Diedrichsen, 2019; Kriegeskorte & Wei, 2021), could be especially useful in helping describe how information is transformed as visual processing unfolds.



## References

- Adamian, N., Andersen, S. K., & Hillyard, S. A. (2020). Parallel attentional facilitation of features and objects in early visual cortex. *Psychophysiology*, *57*, e13498. <https://doi.org/10.1111/psyp.13498>
- Albright, T. D., Desimone, R., & Gross, C. G. (1984). Columnar organization of directionally selective cells in visual area MT of the macaque. *Journal of Neurophysiology*, *51*(1), 16–31. <https://doi.org/10.1152/jn.1984.51.1.16>
- Andersen, S. K., Müller, M. M., & Hillyard, S. A. (2009). Color-selective attention need not be mediated by spatial attention. *Journal of Vision*, *9*(6), 1–7. <https://doi.org/10.1167/9.6.2>
- Arnott, S. R., & Goodale, M. A. (2006). Distorting visual space with sound. *Vision Research*, *46*(10), 1553–1558. <https://doi.org/10.1016/j.visres.2005.11.020>
- Arun, S. P. (2012). Turning visual search time on its head. *Vision Research*, *74*, 86–92. <https://doi.org/10.1016/j.visres.2012.04.005>
- Bartsch, M. V., Donohue, S. E., Strumpf, H., Schoenfeld, M. A., & Hopf, J. M. (2018). Enhanced spatial focusing increases feature-based selection in unattended locations. *Scientific Reports*, *8*(1), 1–14. <https://doi.org/10.1038/s41598-018-34424-5>
- Bichot, N. P., Rossi, A. F., & Desimone, R. (2005). Parallel and Serial Neural Mechanisms for Visual Search in Macaque Area V4. *Science*, *529*(2005), 529–535. <https://doi.org/10.1126/science.1109676>
- Blough, D. S. (1988). Quantitative relations between visual search speed and target-distractor similarity. *Perception & Psychophysics*, *43*(1), 57–71. <https://doi.org/10.3758/BF03208974>
- Boehler, C. N., Schoenfeld, M. A., Heinze, H.-J., & Hopf, J.-M. (2011). Object-based selection of irrelevant features is not confined to the attended object. *Journal of Cognitive Neuroscience*, *23*(9), 2231–2239. <https://doi.org/10.1162/jocn.2010.21558>
- Bohon, K. S., Hermann, K. L., Hansen, T., & Conway, B. R. (2016). Representation of perceptual color space in macaque posterior inferior temporal cortex (The V4 complex). *ENeuro*, *3*(4). <https://doi.org/10.1523/ENEURO.0039-16.2016>
- Boynton, G. M. (2009). A framework for describing the effects of attention on visual responses. *Vision Research*, *49*(10), 1129–1143. <https://doi.org/10.1016/j.visres.2008.11.001>
- Brouwer, G. J., & Heeger, D. J. (2009). Decoding and Reconstructing Color from Responses in Human Visual Cortex. *Journal of Neuroscience*, *29*(44), 13992–14003. <https://doi.org/10.1523/JNEUROSCI.3577-09.2009>
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*(13), 1484–1525. <https://doi.org/10.1016/j.visres.2011.04.012>
- Castiello, U., & Umiltà, C. (1990). Size of the attentional focus and efficiency of processing. *Acta*

- Psychologica*, 73(3), 195–209. [https://doi.org/10.1016/0001-6918\(90\)90022-8](https://doi.org/10.1016/0001-6918(90)90022-8)
- Chang, L., & Tsao, D. Y. (2017). The Code for Facial Identity in the Primate Brain. *Cell*, 169(6), 1013–1028. <https://doi.org/10.1016/j.cell.2017.05.011>
- Chapman, A. F., & Störmer, V. S. (2021). Feature-based attention is not confined by object boundaries: Spatially global enhancement of irrelevant features. *Psychonomic Bulletin & Review*, 28(4), 1252–1260. <https://doi.org/10.3758/s13423-021-01897-x>
- Chapman, A. F., & Störmer, V. S. (2022). Feature similarity is non-linearly related to attentional selection: evidence from visual search and sustained attention tasks. *Journal of Vision*, 22(8), 4. <https://doi.org/10.1167/jov.22.8.4>
- Chen, Z. (2012). Object-based attention: A tutorial review. *Attention, Perception, and Psychophysics*, 74(5), 784–802. <https://doi.org/10.3758/s13414-012-0322-z>
- Conway, B. R., & Tsao, D. Y. (2009). Color-tuned neurons are spatially clustered according to color preference within alert macaque posterior inferior temporal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 106(42), 18034–18039. <https://doi.org/10.1073/pnas.0810943106>
- Çukur, T., Nishimoto, S., Huth, A. G., & Gallant, J. L. (2013). Attention during natural vision warps semantic representation across the human brain. *Nature Neuroscience*, 16(6), 763–770. <https://doi.org/10.1038/nn.3381>
- Driver, J. (2001). A selective review of selective attention research from the past century. *British Journal of Psychology*, 92(1), 53–78. <https://doi.org/10.1348/000712601162103>
- Driver, J., & Baylis, G. C. (1989). Movement and visual attention: the spotlight metaphor breaks down. *Journal of Experimental Psychology: Human Perception and Performance*, 15(3), 448–456. <https://doi.org/10.1037/h0090403>
- Driver, J., Davis, G., Russell, C., Turatto, M., & Freeman, E. (2001). Segmentation, attention and phenomenal visual objects. *Cognition*, 80(1–2), 61–95. [https://doi.org/10.1016/S0010-0277\(00\)00151-7](https://doi.org/10.1016/S0010-0277(00)00151-7)
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, 113(4), 501–517. <https://doi.org/10.1037/0096-3445.113.4.501>
- Duncan, J., & Humphreys, G. W. (1989). Visual Search and Stimulus Similarity. *Psychological Review*, 96(3), 433–458. <https://doi.org/10.1037/0033-295X.96.3.433>
- Egly, R., Driver, J., & Rafal, R. D. (1994). Shifting Visual Attention Between Objects and Locations: Evidence From Normal and Parietal Lesion Subjects. *Journal of Experimental Psychology: General*, 123(2), 161–177. <https://doi.org/10.1037/0096-3445.123.2.161>

- Ekman, M., Roelfsema, P. R., & de Lange, F. P. (2020). Object selection by automatic spreading of top-down attentional signals in V1. *Journal of Neuroscience*, *40*(48), 9250–9259. <https://doi.org/10.1523/JNEUROSCI.0438-20.2020>
- Eriksen, C. W., & St. James, J. D. (1986). Visual attention within and around the field of focal attention. *Perception & Psychophysics*, *40*(4), 225–240.
- Eriksen, C. W., & Yeh, Y. Y. (1985). Allocation of Attention in the Visual Field. *Journal of Experimental Psychology: Human Perception and Performance*, *11*(5), 583–597. <https://doi.org/10.1037/0096-1523.11.5.583>
- Ernst, Z. R., Boynton, G. M., & Jazayeri, M. (2013). The spread of attention across features of a surface. *Journal of Neurophysiology*, *110*(10), 2426–2439. <https://doi.org/10.1152/jn.00828.2012>
- Francis, G., & Thunell, E. (2022). Excess success in articles on object-based attention. *Attention, Perception, and Psychophysics*, *84*(3), 700–714. <https://doi.org/10.3758/s13414-022-02459-6>
- Franconeri, S. L., Alvarez, G. A., & Cavanagh, P. (2013). Flexible cognitive resources: Competitive content maps for attention and memory. *Trends in Cognitive Sciences*, *17*(3), 134–141. <https://doi.org/10.1016/j.tics.2013.01.010>
- He, S., Cavanagh, P., & Intriligator, J. (1997). Attentional resolution. *Trends in Cognitive Sciences*, *1*(3), 115–121. [https://doi.org/10.1016/s1364-6613\(97\)89058-4](https://doi.org/10.1016/s1364-6613(97)89058-4)
- Herrmann, K., Heeger, D. J., & Carrasco, M. (2012). Feature-based attention enhances performance by increasing response gain. *Vision Research*, *74*, 10–20. <https://doi.org/10.1016/j.visres.2012.04.016>
- Herrmann, K., Montaser-Kouhsari, L., Carrasco, M., & Heeger, D. J. (2010). When size matters: Attention affects performance by contrast or response gain. *Nature Neuroscience*, *13*(12), 1554–1561. <https://doi.org/10.1038/nn.2669>
- Hubel, D. H., & Wiesel, T. N. (1968). Receptive Fields and Functional Architecture of monkey striate cortex. *Journal of Physiology*, *195*, 215–243. <https://doi.org/10.1113/jphysiol.1968.sp008455>
- Intriligator, J., & Cavanagh, P. (2001). The Spatial Resolution of Visual Attention. *Cognitive Psychology*, *43*(3), 171–216. <https://doi.org/10.1006/cogp.2001.0755>
- Itthipuripat, S., Garcia, J. O., Rungratsameetaweemana, N., Sprague, T. C., & Serences, J. T. (2014). Changing the spatial scope of attention alters patterns of neural gain in human cortex. *Journal of Neuroscience*, *34*(1), 112–123. <https://doi.org/10.1523/JNEUROSCI.3943-13.2014>
- Kanizsa, G. (1976). Subjective Contours. *Scientific American*, *234*(4), 48–53.
- Kerzel, D. (2020). Direct Evidence for the Optimal Tuning of Attention. *Journal of Experimental Psychology: Human Perception and Performance*, *46*(7), 716–728. <https://doi.org/10.1037/xhp0000744>

- Kriegeskorte, N., & Diedrichsen, J. (2019). Peeling the Onion of Brain Representations. *Annual Review of Neuroscience*, 42, 407–432. <https://doi.org/10.1146/annurev-neuro-080317-061906>
- Kriegeskorte, N., & Wei, X. X. (2021). Neural tuning and representational geometry. *Nature Reviews Neuroscience*, 22(11), 703–718. <https://doi.org/10.1038/s41583-021-00502-3>
- LaBerge, D., & Brown, V. (1989). Theory of Attentional Operations in Shape Identification. *Psychological Review*, 96(1), 101–124. <https://doi.org/10.1037/0033-295X.96.1.101>
- Lamy, D., & Egeth, H. (2002). Object-based selection: The role of attentional shifts. *Perception and Psychophysics*, 64(1), 52–66. <https://doi.org/10.3758/BF03194557>
- Leopold, D. A., O’Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, 4(1), 89–94. <https://doi.org/10.1038/82947>
- Leshner, G. W. (1995). Illusory contours: Toward a neurally based perceptual theory. *Psychonomic Bulletin & Review*, 2(3), 279–321. <https://doi.org/10.3758/BF03210970>
- Lustig, A. G., & Beck, D. M. (2012). Task-relevant and Task-irrelevant Dimensions Are Modulated Independently at a Task-irrelevant Location. *Journal of Cognitive Neuroscience*, 24(9), 1884–1895. [https://doi.org/10.1162/jocn\\_a\\_00249](https://doi.org/10.1162/jocn_a_00249)
- Malcolm, G. L., & Shomstein, S. (2015). Object-based attention in real-world scenes. *Journal of Experimental Psychology: General*, 144(2), 257–263. <https://doi.org/10.1037/xge0000060>
- Marr, D. (1982). *Vision*. San Francisco: Walter Freeman and Co.
- 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. <https://doi.org/10.1016/j.cub.2004.04.028>
- Martinez, A., Ramanathan, D. S., Foxe, J. J., Javitt, D. C., & Hillyard, S. A. (2007). The role of spatial attention in the selection of real and illusory objects. *Journal of Neuroscience*, 27(30), 7963–7973. <https://doi.org/10.1523/JNEUROSCI.0031-07.2007>
- Martínez, A., Teder-Sälejärvi, W., Vazquez, M., Molholm, S., Foxe, J. J., Javitt, D. C., ... Hillyard, S. A. (2006). Objects are highlighted by spatial attention. *Journal of Cognitive Neuroscience*, 18(2), 298–310. <https://doi.org/10.1162/jocn.2006.18.2.298>
- Maunsell, J. H. R., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, 29(6), 317–322. <https://doi.org/10.1016/j.tins.2006.04.001>
- Meyers, E. M. (2018). Dynamic population coding and its relationship to working memory. *Journal of Neurophysiology*, 120(5), 2260–2268. <https://doi.org/10.1152/jn.00225.2018>
- Nagy, A. L., & Cone, S. M. (1996). Asymmetries in simple feature searches for color. *Vision Research*, 36(18), 2837–2847. [https://doi.org/10.1016/0042-6989\(96\)00046-6](https://doi.org/10.1016/0042-6989(96)00046-6)

- Nagy, A. L., & Sanchez, R. R. (1990). Critical color differences determined with a visual search task. *Journal of the Optical Society of America A*, 7(7), 1209–1217. <https://doi.org/10.1364/JOSAA.7.001209>
- Nastase, S. A., Connolly, A. C., Oosterhof, N. N., Halchenko, Y. O., Guntupalli, J. S., Di Oleggio Castello, M. V., ... Haxby, J. V. (2017). Attention selectively reshapes the geometry of distributed semantic representation. *Cerebral Cortex*, 27(8), 4277–4291. <https://doi.org/10.1093/cercor/bhx138>
- Navalpakkam, V., & Itti, L. (2007). Search Goal Tunes Visual Features Optimally. *Neuron*, 53(4), 605–617. <https://doi.org/10.1016/j.neuron.2007.01.018>
- O’Craven, K. M., Downing, P. E., & Kanwisher, N. (1999). fMRI evidence for objects as the units of attentional selection. *Nature*, 401(6753), 584–587. <https://doi.org/10.1038/44134>
- Peelen, M. V., Fei-Fei, L., & Kastner, S. (2009). Neural mechanisms of rapid natural scene categorization in human visual cortex. *Nature*, 460(7251), 94–97. <https://doi.org/10.1038/nature08103>
- Peelen, M. V., & Kastner, S. (2011). A neural basis for real-world visual search in human occipitotemporal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 108(29), 12125–12130. <https://doi.org/10.1073/pnas.1101042108>
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32(1), 3–25. <https://doi.org/10.1080/00335558008248231>
- Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience*, 13, 25–42. <https://doi.org/10.1146/annurev.ne.13.030190.000325>
- Pratt, J., & Arnott, S. R. (2008). Modulating the attentional repulsion effect. *Acta Psychologica*, 127(1), 137–145. <https://doi.org/10.1016/j.actpsy.2007.03.003>
- Pratt, J., & Turk-Browne, N. B. (2003). The attentional repulsion effect in perception and action. *Experimental Brain Research*, 152(3), 376–382. <https://doi.org/10.1007/s00221-003-1557-7>
- Reynolds, J. H., & Heeger, D. J. (2009). The Normalization Model of Attention. *Neuron*, 61(2), 168–185. <https://doi.org/10.1016/j.neuron.2009.01.002>
- Roelfsema, P. R., Lamme, V. A. F., & Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. *Nature*, 395, 376–381. <https://doi.org/10.1038/26475>
- Roth, Z. N., Kay, K., & Merriam, E. P. (2022). Natural scene sampling reveals reliable coarse-scale orientation tuning in human V1. *Nature Communications*, 13, 6469. <https://doi.org/10.1038/s41467-022-34134-7>
- Sàenz, M., Buraças, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision Research*, 43(6), 629–637. [https://doi.org/10.1016/S0042-6989\(02\)00595-3](https://doi.org/10.1016/S0042-6989(02)00595-3)
- Sàenz, M., Buraças, G. T., Boynton, G. M., 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.  
<https://doi.org/10.1038/nn876>
- Schoenfeld, M. A., Hopf, J.-M., Merkel, C., Heinze, H.-J., & Hillyard, S. A. (2014). Object-based attention involves the sequential activation of feature-specific cortical modules. *Nature Neuroscience*, 17(4), 619–624. <https://doi.org/10.1038/nn.3656>
- Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, 80, 1–46.  
[https://doi.org/10.1016/S0010-0277\(00\)00152-9](https://doi.org/10.1016/S0010-0277(00)00152-9)
- Scolari, M., Byers, A., & Serences, J. T. (2012). Optimal deployment of attentional gain during fine discriminations. *Journal of Neuroscience*, 32(22), 7723–7733.  
<https://doi.org/10.1523/JNEUROSCI.5558-11.2012>
- Scolari, M., Ester, E. F., & Serences, J. T. (2014). Feature- and Object-Based Attentional Modulation in the Human Visual System. In A. C. Nobre & S. Kastner (Eds.), *Oxford Handbook of Attention* (pp. 573–600). Oxford, UK: Oxford University Press.  
<https://doi.org/10.1093/oxfordhb/9780199675111.013.009>
- Scolari, M., & Serences, J. T. (2009). Adaptive allocation of attentional gain. *Journal of Neuroscience*, 29(38), 11933–11942. <https://doi.org/10.1523/JNEUROSCI.5642-08.2009>
- Serences, J. T., & Boynton, G. M. (2007). Feature-Based Attentional Modulations in the Absence of Direct Visual Stimulation. *Neuron*, 55(2), 301–312. <https://doi.org/10.1016/j.neuron.2007.06.015>
- Shepard, R. N. (1987). Toward a Universal Law of Generalization for Psychological Science. *Science*, 237(4820), 1317–1323.
- Shih, S. I., & Sperling, G. (1996). Is There Feature-Based Attentional Selection in Visual Search? *Journal of Experimental Psychology: Human Perception and Performance*. <https://doi.org/10.1037/0096-1523.22.3.758>
- Shim, W. M., & Cavanagh, P. (2005). Attentive tracking shifts the perceived location of a nearby flash. *Vision Research*, 45(25–26), 3253–3261. <https://doi.org/10.1016/j.visres.2005.05.029>
- Shomstein, S., & Yantis, S. (2002). Object-based attention: Sensory modulation or priority setting? *Perception and Psychophysics*, 64(1), 41–51. <https://doi.org/10.3758/BF03194556>
- Sreenivasan, K. K., Curtis, C. E., & D’Esposito, M. (2014). Revisiting the role of persistent neural activity during working memory. *Trends in Cognitive Sciences*, 18(2), 82–89.  
<https://doi.org/10.1016/j.tics.2013.12.001>
- Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., & Duncan, J. (2013). Dynamic coding for cognitive control in prefrontal cortex. *Neuron*, 78(2), 364–375.  
<https://doi.org/10.1016/j.neuron.2013.01.039>

- Störmer, V. S., & Alvarez, G. A. (2014). Feature-based attention elicits surround suppression in feature space. *Current Biology*, *24*(17), 1985–1988. <https://doi.org/10.1016/j.cub.2014.07.030>
- Störmer, V. S., Cohen, M. A., & Alvarez, G. A. (2019). Tuning attention to object categories: Spatially global effects of attention to faces in visual processing. *Journal of Cognitive Neuroscience*, *31*(7), 937–947. [https://doi.org/10.1162/jocn\\_a\\_01400](https://doi.org/10.1162/jocn_a_01400)
- Suzuki, S., & Cavanagh, P. (1997). Focused attention distorts visual space: An attentional repulsion effect. *Journal of Experimental Psychology: Human Perception and Performance*, *23*(2), 443–463. <https://doi.org/10.1037/0096-1523.23.2.443>
- Todorov, A., Said, C. P., Engell, A. D., & Oosterhof, N. N. (2008). Understanding evaluation of faces on social dimensions. *Trends in Cognitive Sciences*, *12*(12), 455–460. <https://doi.org/10.1016/j.tics.2008.10.001>
- Tong, F., Nakayama, K., Vaughan, J. T., & Kanwisher, N. (1998). Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron*, *21*, 753–759. [https://doi.org/10.1016/S0896-6273\(00\)80592-9](https://doi.org/10.1016/S0896-6273(00)80592-9)
- Treisman, A. (1986). Features and Objects in Visual Processing. *Scientific American*, *255*(5), 114–125. <https://doi.org/10.1038/scientificamerican1186-114B>
- Treue, S., & Martinez-Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, *399*, 575–579. <https://doi.org/10.1038/21176>
- Vierck, E., & Miller, J. (2008). Precuing benefits for color and location in a visual search task. *Perception & Psychophysics*, *70*(2), 365–373. <https://doi.org/10.3758/PP.70.2.365>
- Wannig, A., Stanisor, L., & Roelfsema, P. R. (2011). Automatic spread of attentional response modulation along Gestalt criteria in primary visual cortex. *Nature Neuroscience*, *14*(10), 1243–1244. <https://doi.org/10.1038/nn.2910>
- Wolfe, J. M. (2021). Guided Search 6. 0: An updated model of visual search. *Psychonomic Bulletin & Review*, *28*, 1060–1092. <https://doi.org/10.3758/s13423-020-01859-9>
- Wolfe, J. M., & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, *5*(6), 495–501. <https://doi.org/10.1038/nrn1411>
- Wolfe, J. M., & Horowitz, T. S. (2017). Five factors that guide attention in visual search. *Nature Human Behaviour*, *1*(3), 1–8. <https://doi.org/10.1038/s41562-017-0058>
- Wolfe, J. M., & Utochkin, I. S. (2019). What is a preattentive feature? *Current Opinion in Psychology*, *29*, 19–26. <https://doi.org/10.1016/j.copsyc.2018.11.005>