INHIBITIONS OF RETURN: TWO INHIBITORY ORIENTING BIASES

by

Ralph S. Redden

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy

at

Dalhousie University
Halifax, Nova Scotia
August 2021

© Copyright by Ralph S. Redden, 2021

DEDICATION PAGE

This dissertation is dedicated to the Mi'kmaq people, on whose unceded land this work was conducted. The suffering and injustices induced by settler colonialism are still on-going. With each new step in my journey, I commit to using my position and privilege to promote recognition, reconciliation, and reparation.

To Owen:

"

If the moon stays up until morning one day,
or a ladybug lands and decides to stay,
or a little bird sits at your window awhile,
it's because they're all hoping to see you smile…

For never before in story or rhyme (not even once upon a time)
has the world ever known a you, my friend,
and it never will, not ever again…
"

- Nancy Tillman

To Marilyn:

"

Love has no other desire but to fulfil itself.
But if you love and must needs have desires, let these be your desires:
To melt and be like a running brook that sings its melody to the night.
To know the pain of too much tenderness.
To be wounded by your own understanding of love;
And to bleed willingly and joyfully.
To wake at dawn with a winged heart and give thanks for another day of loving;
To rest at the noon hour and meditate love's ecstasy;
To return home at eventide with gratitude;
And then to sleep with a prayer for the beloved in your heart and a song of praise upon your lips.
"

- Kahlil Gibran

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

Inhibition of return (IOR) is an inhibitory aftereffect of visuospatial orienting, typically resulting in slower responses to targets presented in an area that has been recently attended. This thesis begins by briefly summarizing the phenomenon's history, and describing work supporting the functional significance of IOR as a foraging facilitator. Discordances in the literature with respect to mechanism are highlighted—in particular the lack of theoretical constructs that can consistently explain innumerable dissociations. Three diagnostics (central arrow targets, locus of slack logic and the psychological refractory period, and performance in speed-accuracy space) are summarized, and positioned relative to the theory that there are two forms of IOR—the form which is manifest being contingent upon the activation state of the reflexive oculomotor system. The input form, which operates to decrease the salience of inputs, is generated when the reflexive oculomotor system is suppressed; the output form, which operates to bias responding, is generated when the reflexive oculomotor system is not suppressed. This theory is then advanced in three empirical chapters. In Chapter 2, findings from two experiments converge to support the hypothesis that the output form is operating at a post-perceptual stage of processing. Chapter 3 contrasts the two forms in a paradigm that intermixed two perceptual diagnostics. The findings support the conclusions from Chapter 2 for the output form, whereas the input form is shown to delay the rate of information accrual at the cued location. In Chapter 4, arrays of multiple cues were used to explore the effect of the centre of gravity of the cueing array when the two forms were generated. The findings suggest the time course of centre of gravity effects may be contingent upon task demands, in addition to the form of IOR that is generated. In Chapter 5, the findings presented in the thesis are summarized and synthesized with the literature, and contextualized relative to a computational process model representing the two forms. On balance, the research presented here provides strong support for the proposal that there are two forms of IOR: one affecting information accrual, the other affecting response thresholds.

LIST OF ABBREVIATIONS USED

AIC	Akaike Information Criterion
2AFC	2-Alternative Forced Choice
ACS	Attentional Control Setting
CND	Cumulative Normal Distribution
COG	Center of Gravity
CRT	Cathode Ray Tube
CTOA	Cue-target onset asynchrony
DNF	Dynamic Neural Field
DVA	Degrees of Visual Angle
ERP	Event-Related Potential
GLMER	Generalized Linear Mixed-Effect Model
IOR	Inhibition of Return
JND	Just-Noticeable Difference
KS	Kolmogorov-Smirnov
LBA	Linear Ballistic Accumulator
PRP	Psychological Refractory Period
PSS	Point of Subjective Simultaneity
ROMS	Reflexive Oculomotor System
RT	Response/Reaction time
SAT	Speed-Accuracy Tradeoff

SOA	Stimulus Onset Asynchrony
SC	Superior Colliculus
SNR	Signal-to-Noise Ratio
TOJ	Temporal Order Judgement
TTOA	Target-Target Onset Asynchrony
TTOA	Target-Tone Onset Asynchrony

ACKNOWLEDGEMENTS

Where does one start when acknowledging the community of people who helped them reach their dreams? It’s hard to imagine finding space to fit everyone into this document, because it could undoubtedly increase my page count by an order of magnitude if I nodded to everyone who ever made an impact on my life leading me to today. However, I shall try.

First, to my amazing partner, Marilyn—without you, this journey never would have even started. Nor would I have survived it. I have learned, and continue to learn, so much from you about life, and the person I wish to be. Your passion for, and devotion to, serving the underserved is extraordinary, all the while putting your family first in everything you do. Your life is entirely a mission in service to others, and for that I couldn’t be more grateful nor in more awe. Your wisdom influences the questions I ask in the lab, the space I try to create in the classroom, and the mentor I try to be. I couldn’t have dreamed for a more incredible teammate on this journey called life!

And of course, to my best bud, Owen—I thought I knew why I loved what I do before we met, but since you came into our lives it has become abundantly clear. I get to be curious, everyday—just like you! It’s the best! Before, I thought my work was fun. Now, it’s still fun, but now I realize I have the ultimate privilege of, everyday, living out a fundamental component of what makes us human: to try to understand our world! And when we learn something new we get to share that with others, so they can take it and make their own discoveries. Thank you for reminding me everyday how fortunate I am. I love you more than anything.

To my parents—it hasn’t always been an easy ride! But I appreciate you letting me find my own way, and supporting me unequivocally whether or not you liked my choices. You have provided the foundation, the beacon, the toolkit, and the reprieve that is necessary to face life on one’s own. If I accomplish nothing else beyond doing the same for Owen, then I will have lived a wonderful life. I couldn’t be prouder to be your son.

To my sister—from mortal enemy to best friend. There wasn’t a single time that you packed your bags to go live with Nan and Grandfather that I didn’t wish you were *actually* going lol! Not that I was a treat to have as a big brother (yes, I *was* kicking your seat). All it took was me leaving for university for things to turn around! Of all the lessons I’ve learned from you though, none have been as significant as the master class in courage and grace you’ve taught us all over the past two years. You’ve taken so much adversity in stride, and faced the worst fears a parent could imagine. It all really put things in perspective: unkind reviews on a journal article, or a mean comment on a teaching evaluation all of a sudden seemed pretty insignificant, and what matters is the time we have with the people we love.

To my fellow lab rats: Swasti Arora, Brett Feltmate, Colin McCormick, Austin Hurst, Richard Drake, and Jon Mulle—I cherish our shared interests, and the unique elements each of you brought to our friendship. You’ve each helped me in innumerable ways, both supporting and challenging me. You’ve each given me confidence in my strengths, and provided a safe space to expose my weaknesses. I'll be in exceptional company if my future colleagues are half as remarkable as each of you.

To Ray—words can't describe the impact your mentorship has had on me. You have been along for the ride on some of the most significant events in my life, whether it was the dark days of losing a loved one, the ultimate high of becoming a father, or the innumerable hills and valleys of life in between. You've always made sure I understood that life is the priority and the science will wait, regardless of the occasion. I am grateful for your wisdom, generosity, and kindness. I appreciate that your benchmark for personal success is when your trainees reach their own goals. Well, the twice-academically dismissed bartender who walked into your lab 10 years ago (almost exactly to the day of my defence) has come a looooong way. Whatever goals I may have had that day have been long realized, as have many set since. No one has done more blazing of trails on my behalf than you, and for that I am forever thankful.

Lastly, to Marilyn (again)—we did it! As challenging as this journey has been for me, it has been every bit and more for you. This is as much your win as it is anyone's. You are the rock, the inspiration, and the model I needed to get here today. I love you more than anything.

CHAPTER 1 - INTRODUCTION

Excerpts from this manuscript are presented below. Co-authors for this manuscript are Dr. W. Joseph MacInnes and Dr. Raymond M. Klein, respectively. In this manuscript, my contributions include, but aren't limited to, writing, theory development, data collection, data analysis, and project management.

Redden, R. S., MacInnes, W. J., & Klein, R. M. (2021). Inhibition of return: An information processing theory of its natures and significance. Cortex, 135, 30-48. This manuscript does not exactly replicate the final version published in Cortex. It is not a copy of the original published article and is not suitable for citation as such.

Inhibition of return (IOR) is described as an inhibitory after- effect of visuospatial orienting (Posner & Cohen, 1984; Posner, Rafal, Choate, & Vaughan, 1985). IOR is typically observed experimentally as slower responding to a target that appears in a location toward which an observer had previously oriented. This effect has stimulated a broad spectrum of research across various domains of cognitive neuroscience in attempts to determine its cause and effect, nature and function and neural implementation. Perhaps the most noteworthy stimulus to the myriad of research on this topic is the proposal about the functional significance of IOR. For some time, it has been believed that there are two forms of inhibition, one that operates on the accrual of information from the inhibited location and one that operates as a response bias. We will review the evidence for these two forms, and show that the form generated depends on the activation state of the reflexive oculomotor system. We conclude by implementing a drift diffusion model to simulate the reaction time distributions and accuracies of the two types of IOR. Our model suggests that input and output IOR are best represented by different functional parameter sets in the decision process, furthering the evidence that the two forms of IOR are two distinct inhibitory phenomena.

1.1 THE MODEL TASK

IOR is typically explored experimentally with the spatial cueing paradigm (Fig. 1.1 - left). The typical spatial cueing paradigm requires an observer to fixate some central stimulus on a screen. This central stimulus is usually flanked by two (or more) equidistant peripheral placeholder boxes. Some time after the start of a trial, a brief stimulus (typically 50-300 ms in duration) occurs at one of the placeholder boxes. This onset could be a brightening of one of the boxes or the appearance of some stimulus within one of these boxes. This onset (also referred to as a cue), is uninformative as to the location of any subsequent stimuli. Some duration after the cue (cue-target onset asynchrony; CTOA), a target appears in one of these boxes that requires some response. Various response types have been used to explore the effect, such as manual detection (e.g., Maylor & Hockey, 1985; Posner & Cohen, 1984), manual localization (e.g., Avery, Cowper-Smith, & Westwood, 2015; Fischer, Pratt, & Neggers, 2003; Rafal, Egly, & Rhodes, 1994), manual discrimination (e.g., Hartley & Kieley, 1995; Lupianez, Milan, Tornay, Madrid, & Tudela, 1997), temporal order judgments (e.g., Maylor, 1985; Posner et al., 1985), and saccadic eye movements (e.g., Lim, Eng, Janssen, & Satel, 2018; Taylor & Klein, 2000; Vaughan, 1984).

The results that are obtained from these paradigms are typically biphasic (Fig. 1.1 - right)—when the CTOA is short (<300 ms), responses to targets that appear at the same location as the cue tend to be fastest. This effect is suggested to occur as a result of attention being captured by the cue (Lupianez, Milliken, Solano, Weaver, & Tipper, 2001, but see; MacInnes & Bhatnagar, 2018). When the interval between the cue and target is longer (>300 ms) and there is sufficient time for attention to return from the periphery to central fixation, responses to targets that appear at the same location as the cue tend to be slowest. This effect has been shown to be long-lasting both temporally (lasting up to 2-3s post-cue; Samuel & Kat, 2003) and spatially (persisting across 6-7 sequentially cued locations; Snyder & Kingstone, 2000), and is suggested to occur as an attentional bias against previously attended locations. This inhibitory effect was first discovered by

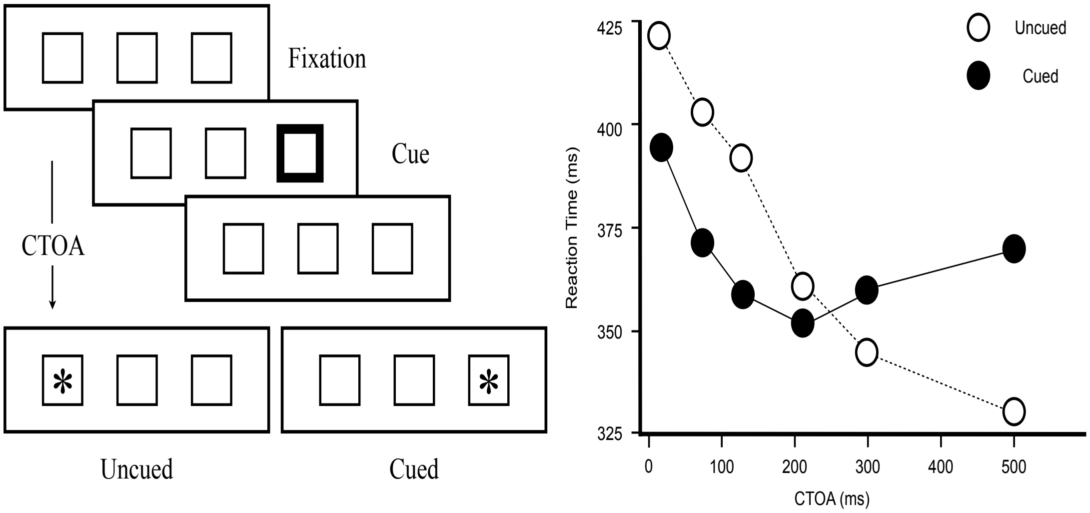

Figure 1.1. (Left) Sequence of events in the model spatial cueing task. (Right) Prototypical pattern of response times as a function of CTOA in the model task. Responses to targets at the cued location are represented with black circles, and responses to targets at uncued locations are represented with white circles (redrawn from Klein (2000)).

Posner and Cohen (1984), and subsequently named inhibition of return (IOR; Posner et al., 1985).

1.2 AN HISTORICAL REVIEW OF THE PHENOMENON

Posner and Cohen (1984) attributed the cause of the inhibition to be any peripheral stimulus and suggested that the effect was to reduce the efficiency of target detection in the vicinity of this stimulus. However, in different variants of the model task Posner et al. (1985) identified oculomotor activation as the cause and a motor bias as the effect. Different causes and resulting effects between these patterns of findings raise the question as to whether or not the two patterns of results are actually reflecting the same mechanism (for a more thorough review of the seminal distinctions, see Hilchey, Klein, & Satel, 2014; Klein & Redden, 2018).

1.2.1 - Functional Significance: Inhibition of Return During and After Search

Posner and colleagues (1984; 1985) proposed that IOR was a bias that favoured orienting to novel locations. Posner et al. (1985) extended this idea by proposing that "such a bias would have obvious advantages in scanning for visual targets." Klein (1988) tested this "foraging facilitator" proposal in a prototypical search task, where observers had to indicate the presence or absence of a particular target amongst homogenous distractors. Following the logic proposed by Treisman and Gelade (1980), Klein used popout and serial search trials, which according to Treisman and Gelade (1980) are done pre-attentively and attentively, respectively. Klein's hypothesis was that if inhibitory tags are left behind as a consequence of disengaging attention during serial search, then these tags ought to exist in serial, but not in popout, search. By presenting a simple detection probe "On" or "Off" a location where a search array item had just been, Klein could test whether having attended to items in an array was sufficient to generate IOR. The findings were precisely in line with this hypothesis: when probes were presented at a location where a serial search item had been presented, observers were slower to detect the probe relative to when the probe was presented at a location where no search stimulus had been. However, when probes were presented subsequent to a popout search trial, there was no difference in probe detection performance for probes that were presented "On" or "Off" a search item. This dissociation reinforces Klein's hypothesis, as IOR was operating in service of search when attention was necessary to perform the search task (vis in the serial search condition), but not when search is performed pre-attentively (vis in the popout search condition). Inhibitory effects on reorienting behaviour have since been observed in both covert (e.g., Muhler & von Muhlenen, 2000; Takeda & Yagi, 2000) and overt (e.g., Klein & MacInnes, 1999) search paradigms so long as the scene is not removed before the probe for IOR is delivered (for a review, see Wang & Klein, 2010). Although there are some who disagree with the foraging facilitator hypothesis for IOR (e.g., Hooge, Over, van Wezel, & Frens, 2005) the weight of the evidence strongly favors it.

1.2.2 - A phenomenon without consensus on definition or theoretical framework

Dukewich and Klein (2015) explored the diversity of perspectives researchers endorse for the phenomenon of inhibition of return. The results from a survey of over 60 IOR experts (defined as individuals with 4+ peer-reviewed publications with IOR as a keyword) revealed little agreement on the phenomenon, with varying perspectives on causes, mechanisms, effects and components. For instance, whereas 43% of the experts agreed that IOR is a general term for similar looking effects, the remaining 57% conversely supported the idea that IOR could not be explained as a single, unitary construct.

Much research has supported the proposal that a single mechanism can account for IOR effects across sensory and response modalities. For instance, adaptive gain models of IOR have been proposed to suggest the phenomenon operates as a supramodal habituation response to orienting (Dukewich, 2009; Rutherford, O'Brien, & Raymond, 2010; Sereno, Lehky, Patel, & Peng, 2010). Supporting its supramodal nature, IOR has been observed between successive responses to stimuli presented from all pairings of vision, audition and touch (Spence, Lloyd, McGlone, Nicholls, & Driver, 2000). Mathot, Dalmaijer, Grainger and van der Stigchel (2014) conceptualized IOR as a habituation to visual inputs, and demonstrated the size of an individual's IOR effect correlates with the pupillary light response. Differences in the time course of IOR between effectors (Briand, Larrison, & Sereno, 2000) have also been interpreted in the context of the premotor theory of attention (Klein, 1980; Rizzolati, Riggio, Dascola & Umilta, 1987). Cowper-Smith, Harris, Eskes, and Westwood (2013) argue in support of a unitary inhibitory mechanism across effectors when comparing pointing responses to oculomotor responses. Studies using multiple stimuli to measure the influence of the center of gravity of the cueing array have also suggested a unitary inhibitory phenomenon affecting both manual and saccadic responses (Christie, Hilchey, & Klein, 2013; MacInnes, 2017).

Alternatively, numerous dissociations support the hypothesis that there are (at least) two qualitatively different inhibitory cueing effects. Sumner, Nachev, Vora, Husain, and Kennard (2004) used a mixture of cue stimuli that did (S-Cone stimuli) or did not (traditional luminance stimuli) bypass the superior colliculus to examine the influence of the subcortical structure on IOR when measured with manual and saccadic responding. They found that when manual responding was required, both cue types generated IOR. However, when saccadic responses were required, typical luminance cues generated IOR whereas S-Cone cues did not. Investigating the effect of sequences of eye movements planned independently or in parallel, MacInnes, Kruger and Hunt (2015) found saccadic IOR only after sequences of independently programmed eye movements, but manual IOR after both independent and parallel sequence conditions (see also Sereno, Jeter, Pariyadath, & Briand, 2006). IOR seems to be readily encoded in an object-based reference when manual responses are required (Smith, Ball, Swalwell, & Schenk, 2016; Theeuwes, Mathot & Grainger, 2014; Tipper, Driver, & Weaver, 1991; Tipper, Jordan, & Weaver, 1999; Tipper, Weaver, Jerreat, & Burak, 1994; Weaver, Lupianez & Watson, 1998), but when saccadic responses are required, IOR may (Abrams & Dobkin, 1994; Swalwell, 2019; Tas, Dodd, & Hollingworth, 2012) or may not (Abrams & Pratt, 2000; Senturk, Greenberg & Liu, 2016; Redden, Hilchey, & Klein, 2018) be object-based. Contrary to Cowper-Smith et al. (2013), Pratt and Neggers (2008) argue that separate inhibitory mechanisms influence pointing and oculomotor responses. Hunt and Kingstone (2003) asked their participants to make manual or saccadic localization responses to targets in alternate blocks while orthogonally manipulating target luminance and whether or not the fixation stimulus was removed at the time of the target. Their experiment revealed a double-dissociation: when saccadic responses were required IOR interacted with fixation removal but not with luminance, and when manual responses were required IOR interacted with luminance but not with fixation removal (but see, Michalczyk & Bielas, 2019).

Regardless of whether one ascribes to a unitary or dualist theory to explain IOR, neuroscientific studies of non-human primates (Dorris, Klein, Everling, & Munoz, 2002; Dorris, Taylor, Klein, & Munoz, 1999; Mirpour, Arcizet, Ong, & Bisley, 2009), and patients with impairments in cortical (Bourgeois, Chica, Migliaccio, de Schotten, & Bartolomeo, 2012; Bourgeois, Chica, Valero-Cabre, & Bartolomeo, 2013) and sub-cortical (Gabay, Henik, & Gradstein, 2010; Smith, Rorden, & Jackson, 2004) brain regions support the hypothesis that the neural systems involved in eye movements play a key role in the nature of IOR's effect(s).

1.2.3 - Two Forms Proposed

Taylor and Klein (2000) conducted a parametric investigation into the phenomenon of IOR to systematically explore the relative contributions of input and output mechanisms. They did so by assessing the relationships between mixed central arrow and peripheral signals across response modalities (a more detailed description of motivation underlying the target types is outlined in Section 1.3.1). By counterbalancing the response required to a spatially uninformative first signal (ignore, manual, saccade) and second signal (manual, saccade) over six experimental sessions, 24 unique conditions were created (Fig. 1.2). The results showed that when observers were not required to generate eye movements during a trial (Fig. 1.2: cells 1 & 2), slower responding at cued locations was only found when the second signal was a peripheral event. Because this inhibitory effect could be generated by a centrally-presented arrow, these results preclude any sensory cause for the resulting effect as the central arrow and peripheral luminance signals do not overlap in location. Furthermore because inhibition can only be measured in responses to targets presented at the peripheral locations suggests that the effect exists at the spatial location and manifests on input mechanisms. In striking contrast, when observers were required to engage the oculomotor system at any point during a trial (Fig. 1.2: cells 3-6), equivalent inhibition to both types of second signal was found, suggesting that this inhibitory effect operates on output processes, and manifests as a bias against

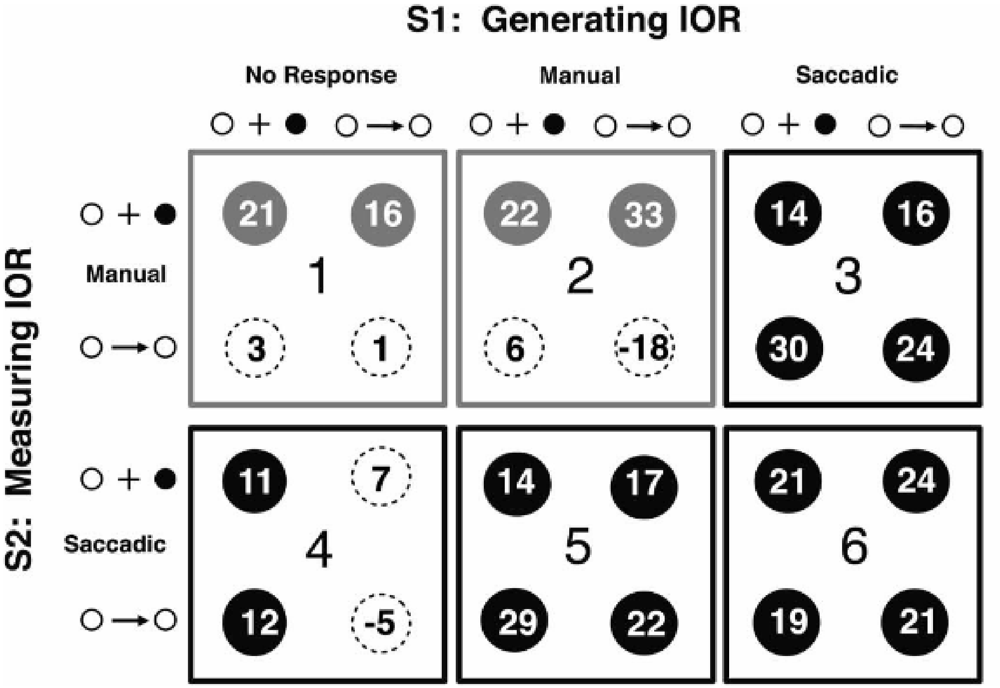

Figure 1.2. Results matrix (Cued RT - Uncued RT) from Taylor & Klein (2000). The large rows and columns represent the stimulus-response ensemble (numbered in the center of each panel). The small rows and columns within them represent the signal types. The conditions purported to elicit output-based effects are the black panels, whereas input-based effects are grey panels. Solid circles depict statistically significant effects.

responses spatially compatible with the first signal. Taylor and Klein posited the dissociation as evidence for two forms of IOR: when eye movements are made, the output form was generated; when eye movements are not made the input form was generated. As we will show, there are two forms, but the association (circa 2000) of the two forms of IOR simply with the response modality (an attribution made by many IOR scholars), was mistaken.

1.3 THREE DIAGNOSTICS FOR INPUT AND OUTPUT EFFECTS

As reviewed above, there is little empirical or theoretical agreement in the IOR literature. With that in mind, we propose that perhaps a major factor in the discordant literature is that researchers have, in fact, been measuring (and ultimately conflating) two distinct effects: an input form of IOR degrading the quality of information recruited from the inhibited location, and an output form of IOR influencing the criterion at which responses are executed. In the following section, we will describe three diagnostics used to determine whether an effect is operating at the input or output level. Furthermore, we will summarize the research using each of these diagnostics, including how the findings converge with the theory that there are two forms of inhibition of return, and how the form observed in any experimental context is contingent upon the activation state of the reflexive oculomotor system.

1.3.1 - 1—Central Arrow Targets

Inhibition measured by responses to a peripheral onset target can be attributed to output and/or input processes as the cue and target overlap spatially. Using centrally-presented arrows as targets was first pioneered independently by Rafal, Egly, and Rhodes (1994) and Abrams and Dobkin (1994). A central arrow target following a peripheral onset allows the evaluation of a motoric (output) contribution to the inhibition, as there is no repetition of the sensory pathway from the cue relative to the target (e.g., Fischer et al., 2003). When arrow targets are used in conjunction with peripheral onset targets, one can discern whether the inhibitory effect generated by the cue is operating on inputs or outputs with the following logic: If the inhibitory effect is influencing the quality of

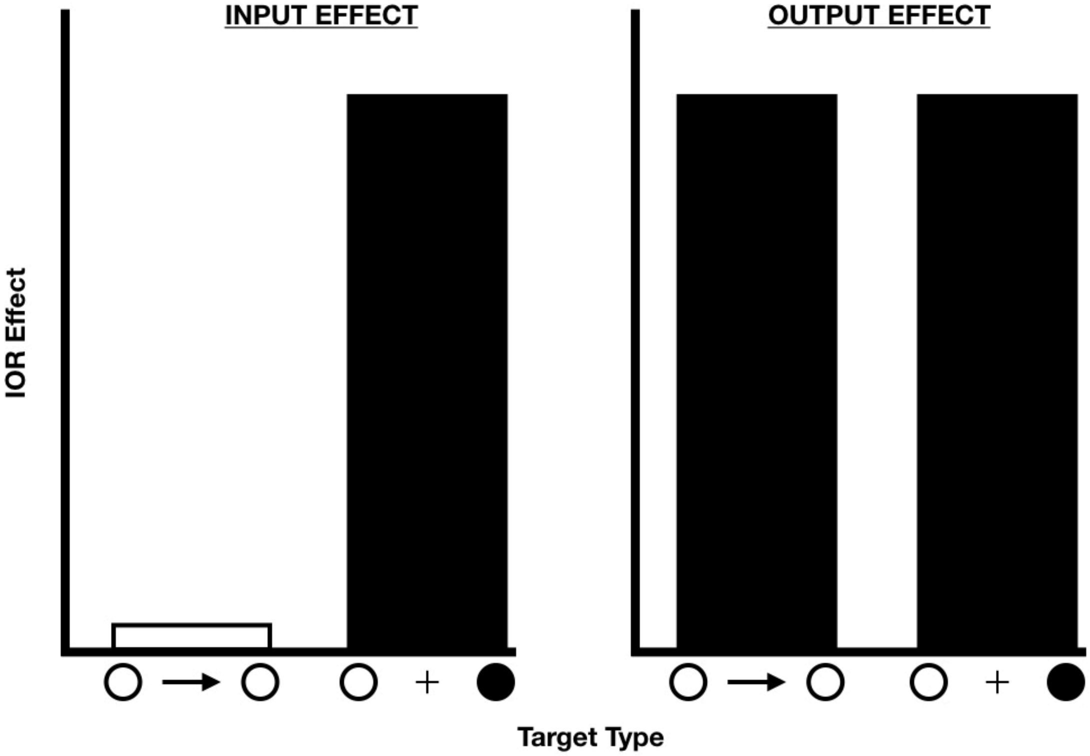

Figure 1.3. How the pattern of results when central arrow and peripheral onset targets (depicted schematically on the x-axis) are randomly intermixed can be used to infer whether an input effect (LEFT) or output effect (RIGHT) has been generated.

information at the cued location, then one would observe slower RTs for cued peripheral onsets (relative to uncued), but no effect on arrow targets since there is no spatial overlap between the cue and arrow target (as illustrated in the left panel of Fig. 1.3). However, if the inhibitory effect is only influencing responding, then RTs will be slower to cued peripheral onsets and to arrows in the direction of the cue (as illustrated in the right panel of Fig. 1.3). It is noteworthy that we have already highlighted this diagnostic in the previous “Two Forms” section when presenting Taylor and Klein’s (2000) methods and findings (see Fig. 1.2).

Rafal et al. (1994; Experiment 4) had observers execute an eye movement in the direction of a central arrow presented after a spatially uninformative peripheral onset. Supporting a motoric effect, they concluded that “because the go-signal never appeared at the cued location, any inhibition of return in Experiment 4 would not be attributable to an inhibition of detection (pp. 293).” Abrams and Dobkin (1994) found uninformative peripheral onsets inhibited saccadic responses for both central arrows and peripherally-presented targets, however to a greater extent for peripheral onset targets than central arrows. They proposed that two additive components comprise IOR: an attentional component, that under our interpretation would be an effect that would delay input processing, and a motoric component. They suggest central arrow targets were only sensitive to measuring the motoric component, whereas both components contributed to the inhibition for peripheral targets.

Abrams and Dobkin’s pattern of results was contradicted by Taylor and Klein’s (2000) findings in the same condition (see Fig. 1.2 - Cell 4). Klein and Hilchey (2011, p. 480) noted that a key methodological difference between Abrams and Dobkin’s (1994) static display experiments and those of Taylor and Klein (2000) was that Abrams and Dobkin implemented a between-subjects design to evaluate the effect of the different target diagnostics, whereas Taylor and Klein implemented a within-subjects design with cue and target types completely intermixed. Whereas Abrams and Dobkin showed different magnitudes of inhibition between target types, Taylor and Klein (2000) showed statistically indistinguishable inhibition between target types in each cell in which the output form of IOR was generated (Fig. 1,2, cells 3-6). Hilchey, Klein, and Ivanoff

(2012) hypothesized that when target types are blocked, observers could instantiate different spatial attentional control settings (ACS)—observers might more effectively filter out the uninformative peripheral cue in blocks with central arrow targets because the cue never appears in task-relevant space. This would not be the case when targets are only peripheral onsets, or randomly intermixed between peripheral onsets and central arrows because the uninformative cue appears in a task-relevant location. Hilchey et al. directly compared this within/between blocks distinction and showed that when target types were randomly intermixed, there was equivalent inhibition for each target type. However, when target types were implemented in separate blocks, IOR only occurred with peripheral targets. Hilchey et al. thusly refute Abrams and Dobkin's two component model for IOR and explain the pattern in Abrams and Dobkin's static display experiments as caused by a difference in spatial ACS.

Hilchey, Dohmen, Crowder, and Klein (2016) measured manual localization responses to central arrow targets following either a pro- or anti-saccade to a peripheral cue. On prosaccade trials, they found that responses to central arrows in the direction compatible with the cued location were slower than those compatible with the uncued location—consistent with an output effect in the cued direction. However, when antisaccades were required to the cue there were no discernible differences between responses to central arrows compatible with either the location of the cue or the direction of the eye movement—inconsistent with an output effect in either direction, but rather precisely the result that would be predicted for an input effect at the cued location. These results suggest that the output form is not elicited by an eye movement per se, because both groups were required to make eye movements. Rather, the important distinction is whether the system responsible for reflexive eye movements is suppressed or not, since one must suppress the reflexive eye movement tendency in order to generate an antisaccade (Everling, Dorris, Klein, & Munoz, 1999; Ignashchenkova, Dicke, Haarmeier, & Thier, 2004).

1.3.2 - 2—Psychological Refractory Period & Locus of Slack

The psychological refractory period refers to the delay in the processing of the second of two sequentially presented targets when the interval between them is relatively short (cf Pashler, 1998). When "structural" interference can be eliminated as the cause (because there is no overlap in the input modality or responding effectors for the two tasks), it is generally assumed that the delay is caused by a central processing stage (often called a central bottleneck) that can only perform some mental operation(s) for one task at a time. When processing of one task reaches this central bottleneck, processing of the subsequent task cannot usurp the central bottleneck until processing of the first event has completed this stage. As illustrated in Fig. 1.4, when the second task is delayed by some manipulation, such as IOR, it is possible to use the pattern of performance when the interval between the stimuli for the two tasks is varied to distinguish whether the manipulation was affecting early (pre-bottleneck) or later (bottleneck or post-bottleneck) stages of processing.

Klein and colleagues (Kavyani, Farsi, Abdoli, & Klein, 2017; Klein, Kavyani, Farsi, & Lawrence, 2020) applied this logic to discern the stage at which the two forms of

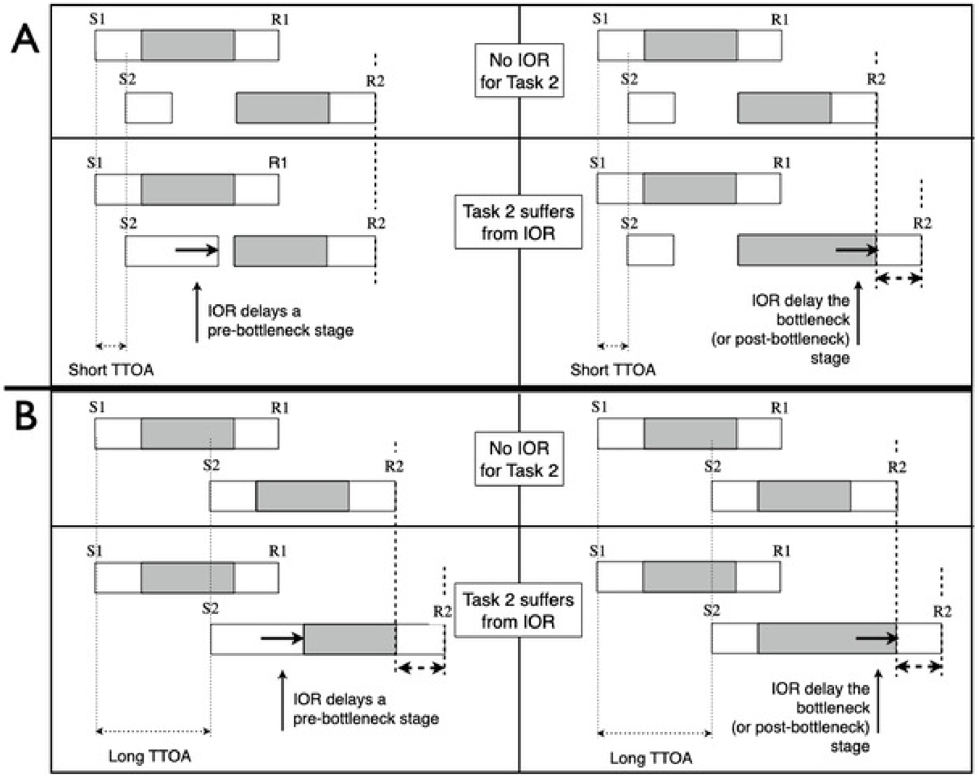

Figure 1.4. Prediction of a central bottleneck model (as described by Pashler, 1998) when the interval (TTOA) between the onset of two stimuli, S1 and S2 requiring speeded responses, is relatively short (A) and relatively long (B). The bottleneck stages of processing for both tasks are shaded in grey. A baseline condition, for which Task 2 does not suffer from IOR, is presented in the upper portion of each panel, whereas in the lower portion, Task 2 suffers from IOR. This results in the prolongation of a Task 2 stage of processing (represented by the solid horizontal arrows). On the left side a pre-bottleneck stage of processing has been prolonged. Compared to the baseline condition, so long as the TTOA is long enough to avoid any overlap of the bottleneck stages of processing (B) the full effect of the manipulation that prolonged processing in Task 2 is manifest in RT2 (dashed double-sided arrow); but when the TTOA is short (A), the prolonged early stage of processing is absorbed into the period of slack and no effect of IOR is seen in RT2. On the right side it is assumed that the manipulation slows the bottleneck stage of processing (the predictions are the same if a post-bottleneck stage is delayed). In this case, the response time to complete Task 2 is prolonged by identical amounts (dashed double-sided arrows) whether the TTOA is short (A) or long (B). (Redrawn from Klein et al., 2020).

IOR may be operating. They found that when participants were forbidden to make eye movements (Kavyani et al., 2017), the effect of the cue on T2 interacted with the delay between tasks—there was no observable IOR when the TTOA was short, but the effect emerged at longer TTOAs—suggesting that the inhibition was operating at a pre-bottleneck (sensory; input) stage, and thus would have been absorbed into the slack period when the delay between tasks was sufficiently short. However, when eye movements were required (Klein et al., 2020), the inhibitory effect was additive with the delay between tasks, suggesting that the inhibition on T2 was operating at a post-bottleneck (response selection/execution; output) stage of processing.

1.3.3 - 3—Speed-Accuracy Space

Ivanoff, Klein, & Lupianez (2002) hypothesized how an inhibitory after-effect operating at early or late stages in information processing would influence performance in terms of speed and accuracy. Using this approach, a genuine decrease in performance (i.e., slower and less accurate response; Fig. 1.5 - dotted arrows) would indicate input-based inhibition at the cued location, whereas output-based inhibition would be represented by a speed-accuracy trade off (Fig. 1.5 - solid arrow). Ivanoff and Klein (2001) used a Go/No-Go task and found a speed-accuracy tradeoff (SAT) function at cued versus uncued locations when participants were instructed to keep their eyes at fixation during the experiment. Later Hilchey, Hashish, et al. (2014) examined the role of eye movements using the Ivanoff and Klein (2001) paradigm. They used a go/no-go task like Ivanoff and Klein (2001), however, they provided real-time eye movement feedback on each trial in order to suppress oculomotor activation. They showed that actively

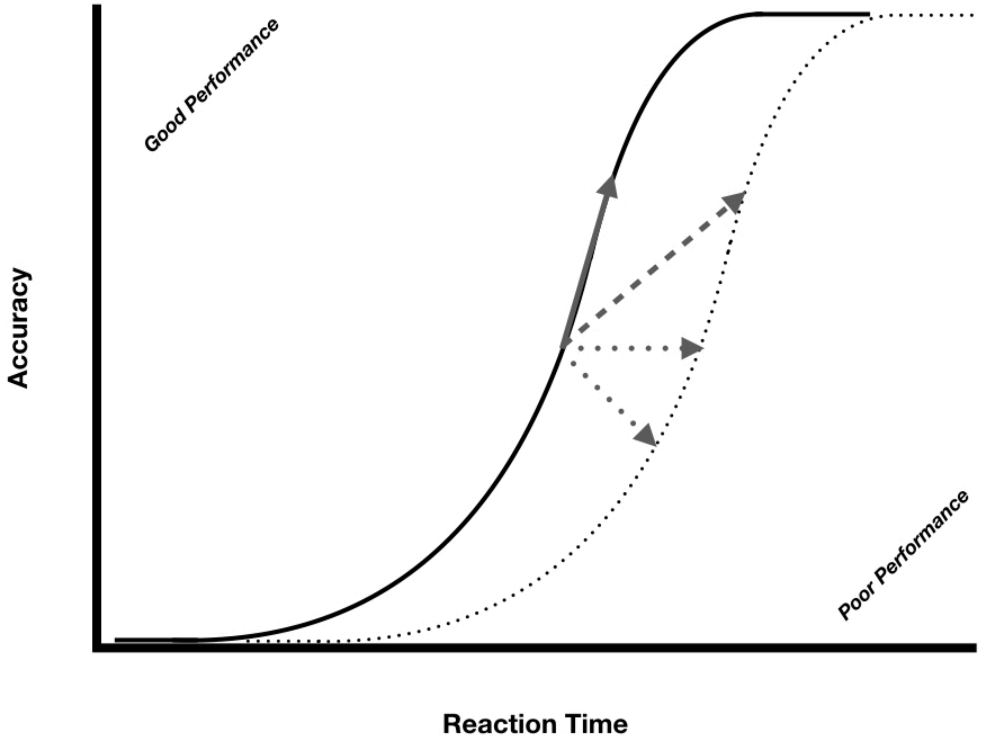

Figure 1.5. Illustration of the two theories of IOR in speed-accuracy space. All functions represent hypothetical rates of information accrual where improvement in the accuracy of performance as response time increases. If the RT delay associated with IOR is caused a genuine deterioration in performance this would result in a rightward shift of the function (as shown by a shift from the solid function to the less efficient dotted function, represented by the dotted arrows) or a change in slope of the function. Alternatively, observers may demonstrate slower responding, with the additional response time eliciting an improvement in accuracy, also referred to as a speed-accuracy trade-off or criterion shift (represented by the solid arrow). The dashed arrow illustrates the ambiguity of a pattern with increase RT and increased accuracy (see text for explanation).

suppressing the oculomotor system in a manual response task changed the form of IOR from output-based (e.g., Ivanoff & Klein, 2001; Ivanoff & Klein, 2003; Taylor & Ivanoff, 2003; Ivanoff & Klein, 2003; Prime & Jolicœur, 2009; Prime & Ward, 2006) to input-based (e.g., Cheal, Chastain, & Lyon, 1998).

In addition, Chica, Taylor, Lupianez & Klein (2010) examined the two effects underlying IOR using a non-spatial two-alternative forced choice (2AFC) task with coloured targets. The results showed that when participants were instructed to maintain fixation, and on-line feedback on unwanted eye movements was provided, the inhibition was input-based (degraded performance). On the contrary, when subjects successfully made a saccade to the cue and back to the original fixation before the final target, output-based inhibition (a speed-accuracy tradeoff) was observed. Chica, Taylor, Lupianez, and Klein (2010) concluded that the motoric flavor of IOR elicited by oculomotor activation does not affect perceptual processing but instead delays response time allowing for more information accrual and hence more accurate responses.

This dissociation was reinforced by Redden, Hilchey, and Klein (2016), whereby subjects were required to make either a prosaccade or antisaccade (instead of ignore, as in Chica et al.) in response to the cue, and execute a subsequent manual 2AFC response to targets. Observers who made a prosaccade to the cue showed slower but more accurate target responses (output effect), whereas observers who made an antisaccade to the cue showed slower and less accurate target responses (input effect). This pattern also reinforces the causal relationship between reflexive oculomotor activation and inhibitory effect reported by Hilchey et al. (2016) using the central arrow diagnostic. Furthermore, generating full SAT functions using the response-window method (Wickelgren, 1977) has clarified a possible ambiguity in SAT patterns. It is possible that an upward pointing arrow representing an SAT could be accompanied by a genuine change in performance (Fig. 1.5 - dashed arrow), whereby the arrow starts on one function and ends on a less efficient one. However, using the response-window method has shown that when eye movements are discouraged (Ivanoff & Klein, 2006) performance for cued targets exist on a less efficient function than do uncued targets, but when eye movements are required to the cue (Redden, Hilchey, Aslam, Ivanoff, & Klein, 2020; Chapter 2) performance for cued and uncued targets exist on a single performance function (Fig. 1.5 - solid arrow).

1.4 OVERVIEW OF DISSERTATION

This dissertation will serve to build upon the framework outlined above, and advance our understanding of the nature of the two forms of inhibition of return. Throughout the dissertation, I will refer to a computational model in Appendix A, as reported in Redden, MacInnes and Klein (2021). The model shows that the two forms of inhibition of return are best accounted for by different parameters in the decision process: the input form is best modelled as an effect on the rate of information accrual, and the output form is best modelled as an increase in the threshold of information required to make decisions. In Chapter 2, I outline an ambiguity in the model results for the output form, and test a hypothesis pertaining to this ambiguity wherein two parameters modelled the data well, but only one fit within the framework proposed above. In Chapter 3, I explore the perceptual consequences of the two forms of IOR — explicitly whether the input form is manifest as an effect on the rate or on the quality of information. In Chapter 4, I use a method from the oculomotor priming literature to explore whether the effect at the center of gravity of an array of stimuli is affecting input or output processes. Lastly, in Chapter 5, I synthesize my contributions in context with the framework proposed above, outline limitations in the theoretical account, and propose avenues for future research.

CHAPTER 2: USING SPEED AND ACCURACY AND THE SIMON EFFECT TO EXPLORE THE OUTPUT FORM OF INHIBITION OF RETURN

The (unpublished) manuscript based on this submission is presented below. Co-authors for this manuscript are Dr. Matthew D. Hilchey, Mr. Sinan Aslam, Dr. Jason Ivanoff and Dr. Raymond M. Klein, respectively. In this manuscript, my contributions include, but aren't limited to, writing, theory development, data collection, data analysis, project management, and experimental design.

A revision of this manuscript is currently under review at the Quarterly Journal of Experimental Psychology (QJE-STD-19-441.R1).

In spatial cueing paradigms, responses are usually slower to targets at previously cued relative to uncued locations at cue-target onset asynchronies between .2 and 3 seconds (Posner & Cohen, 1984; Posner, Rafal, Choate & Vaughan, 1985). The finding is often referred to as inhibition of return (IOR; Klein, 2000; Lupianez, Klein & Bartolomeo, 2006) and it is thought to promote efficient visual search by discouraging re-inspection of previously processed stimulus locations (Wang & Klein, 2010; Klein & Redden, 2018).

Some studies show that the nature of IOR is different depending on whether eye movements are made. When eye movements are not made, slower responses are observed only when the target appears within the vicinity of the cue (Taylor & Klein, 2000; Fischer, Pratt & Neggers, 2003; Hilchey, Klein & Satel, 2014; Wang, Hilchey, Cao & Wang, 2014). Because this effect depends on repeated stimulation of an input pathway, it is commonly described as attentional/perceptual (Taylor & Klein, 2000) or, synonymously, as occurring nearer the input end of the information processing continuum (Hilchey et al., 2014). When eye movements are made, responses toward the cued location are slower than responses away from it, even when responding to a centrally presented arrow (Taylor & Klein, 2000; Rafal, Egly & Rhodes, 1994; Posner, Rafal, Choate and Vaughan, 1985). Because this effect does not depend on repeated stimulation of an input pathway, it is commonly described as motoric/decisional (Taylor & Klein, 2000) or, synonymously, as occurring nearer the output end of the information processing continuum (Hilchey, Klein & Ivanoff, 2012).

Ivanoff, Klein and Lupianez (2002) suggested two different effects of IOR (Figure 1.5; Section 1.3). An inhibitory effect closer to the input end of the information processing continuum would be reflected by a rightward shift in the information accrual function for cued relative to uncued targets (i.e., a genuine decrease in information processing efficiency at the cued location). An inhibitory effect closer to the output end of the information processing continuum would be reflected by a response bias or criterion shift (i.e., a greater evidential threshold for responding). Recent investigative work has focused on experimentally testing when and whether to expect these kinds of effects of IOR.

For example, Chica, Taylor, Lupianez and Klein (2010) administered a spatial cueing paradigm in which they manipulated, between-subjects, whether a participant made an eye movement toward a spatially uninformative peripheral cue and back to fixation (pro-saccade condition) or not (no-saccade condition). Shortly thereafter, participants had to discriminate the color of a target with a keypress response. Two qualitatively different patterns were obtained, each corresponding to the theoretical constructs posited by Ivanoff, Klein and Lupianez (2002). In the no-saccade condition, responding was slower and less accurate to targets at the cued location, suggesting an effect on input processes. In the pro-saccade condition, responding was slower but also more accurate at the cued location, suggesting an effect on output processes.

Redden, Hilchey and Klein (2016) extended Chica et al (2010)'s findings by replacing the no-saccade condition with an anti-saccade condition (i.e., an eye movement was to be made to the location opposite the cue). They did this to test the hypothesis from Klein and Hilchey (2011; see also Klein & Redden, 2018) that the critical factor in determining the form of IOR is not whether eye movements are involved (i.e., Taylor & Klein, 2000) but rather whether eye movements are permitted toward cues and targets in peripheral vision. Redden et al. replicated Chica et al.'s output-based IOR effect when eye movements were made to the cue. However, anti-saccades to the cue led to input-based IOR effects. According to Klein and Hilchey, this dissociation occurred because it was necessary to suppress the reflexive oculomotor machinery in order to make anti-saccades (but not in order to make pro-saccades). They theorized that it is whether this machinery is in a tonically suppressed or active state that determines whether input- or output-based forms of IOR will be observed (see also Hilchey, Klein & Satel, 2014; Hilchey, Dohmen, Crowder & Klein, 2016).

Recently, drift diffusion modelling (e.g., Ratcliff, 2008) has provided converging evidence that the input- and output-based forms of IOR arise from dissociable mechanisms. Redden, MacInnes and Klein (2021) applied diffusion modelling to the data from Redden, Hilchey and Klein (2016) and found that the input-based form of IOR generated in the anti-saccade condition was best accounted for by a reduction in the drift rate parameter (Appendix A). The drift rate parameter represents the average slope at which information accrues toward the "Correct" response in a random walk model of a 2AFC, and thus a reduction of this slope would produce slower RTs (because information isn't accruing as fast) and less accurate responses overall (since reducing slope away from "Correct" necessitates it is more sloped toward "Incorrect")—a pattern consistent with an effect on the quality of perceptual information processing. The output-based form of IOR generated in the pro-saccade condition was well fit by either (1) an increase in the response threshold parameter or (2) a reduction in the drift noise parameter. The response threshold parameter represents the distance between the start point and each of the two responses, or how far the random walk has to travel to reach one of the two response options. An increase in threshold would result in slower RTs (because more information is needed) and more accurate responses overall (because more information results in greater accuracy). The drift noise parameter represents the magnitude of variability in the accrual of information as time progresses within a trial, representing the signal-to-noise ratio (SNR) within a trial. A reduction in drift noise would also result in slower RTs (because both signal and noise contribute to overall RT) and more accurate responses (because less noise is interfering with signal). Whereas an increased threshold account would be consistent with an output effect, it is ambiguous whether reducing drift noise would suit such a theoretical construct. The conceptualization of an output effect is explicitly that there is a reluctance to respond as a consequence of the cue (change in threshold), not that the information at the cued location has been influenced (change in SNR). The main goal here is to resolve this ambiguity by determining experimentally which behavioural effect best characterizes output-based IOR (Experiment 1) and to help clarify the nature of this phenomenon (Experiment 2).

2.1 EXPERIMENT 1

In Experiment 1, we aimed to resolve this ambiguity by intermixing centrally-presented arrow targets (see Figure 1.3 and Section 1.3.1) with the peripheral targets in

Redden, Hilchey and Klein (2016)'s pro-saccade condition. The arrows at fixation pointed to the left or right and required left and right keypress responses, respectively. The peripheral targets appeared to the left or right of fixation and their shapes were discriminated with left and right keypress responses. All targets were preceded by a transient, spatially uninformative cue on the left or right side of fixation that called for a pro-saccade. In this design, we expect to replicate the pattern that discrimination responses will be slower but more accurate to cued as compared to uncued targets in peripheral vision, which is diagnostic of output-based IOR (Chica et al., 2010; Redden, Hilchey & Klein, 2016). More important, the design provides two additional diagnostics useful for resolving the aforementioned ambiguity.

The first diagnostic as to whether the response threshold or drift noise better accounts for output-based IOR concerns whether the arrow-elicited responses are affected by their compatibility with the cued location (e.g., left cue, left arrow response = compatible/cued; left cue, right arrow response = incompatible/uncued). If IOR results from an increase in the response threshold toward the cued location, then arrow responses compatible with the cue location should be slower than arrow responses incompatible with it. If IOR results from a decrease in drift noise at the cued location, then responses to these arrows should be unaffected by the cueing, since there is no spatial overlap between the cues and central arrow targets.

A second diagnostic is provided by the Simon effect (Simon, 1969), which refers to the observation that responses are faster and more accurate for effectors nearest the target. The Simon effect has long been thought to be the result of response conflict (Lu & Proctor, 1995). According to additive factors logic (Sternberg, 1969), factors that interact share a processing locus ought to interact in a factorial design. Ivanoff, Klein, & Lupiáñez (2002) observed an interaction between IOR and the Simon effect such that the Simon was greater for targets at locations impacted by IOR. Wang, Fuentes, Vivas, and Chen (2013) have also observed this interaction, noting that the neural activity in precentral cortex (i.e., primary sensorimotor cortex) may be the source of the interaction between IOR and the Simon effect. As discussed earlier, the absence of eye monitoring (and not knowing whether the reflexive oculomotor system is in a continuously suppressed state) poses a problem for interpreting the locus (input, output, or both) of IOR. To put this in perspective, Redden, Hilchey and Klein (2016 - Supplemental Materials) demonstrated that the input and output forms of IOR have opposite interactions with the Simon effect, with input- and output forms increasing and decreasing Simon effects, respectively[1]. An enhancement of the Simon effect is consistent with an increased tendency toward the prepotent response when target signal quality is reduced (viz an input effect). An attenuation of the Simon effect is precisely what would be expected if there was a reluctance to make responses compatible with the location of the cue (vis an output effect), consistent with an effect on response threshold, but not drift noise.

2.2 METHOD

---

[1] Evaluation of the Simon effect is a convenient but incidental consequence of response mappings ('z' & '/' keys) on the same spatial axis as target locations (left & right).

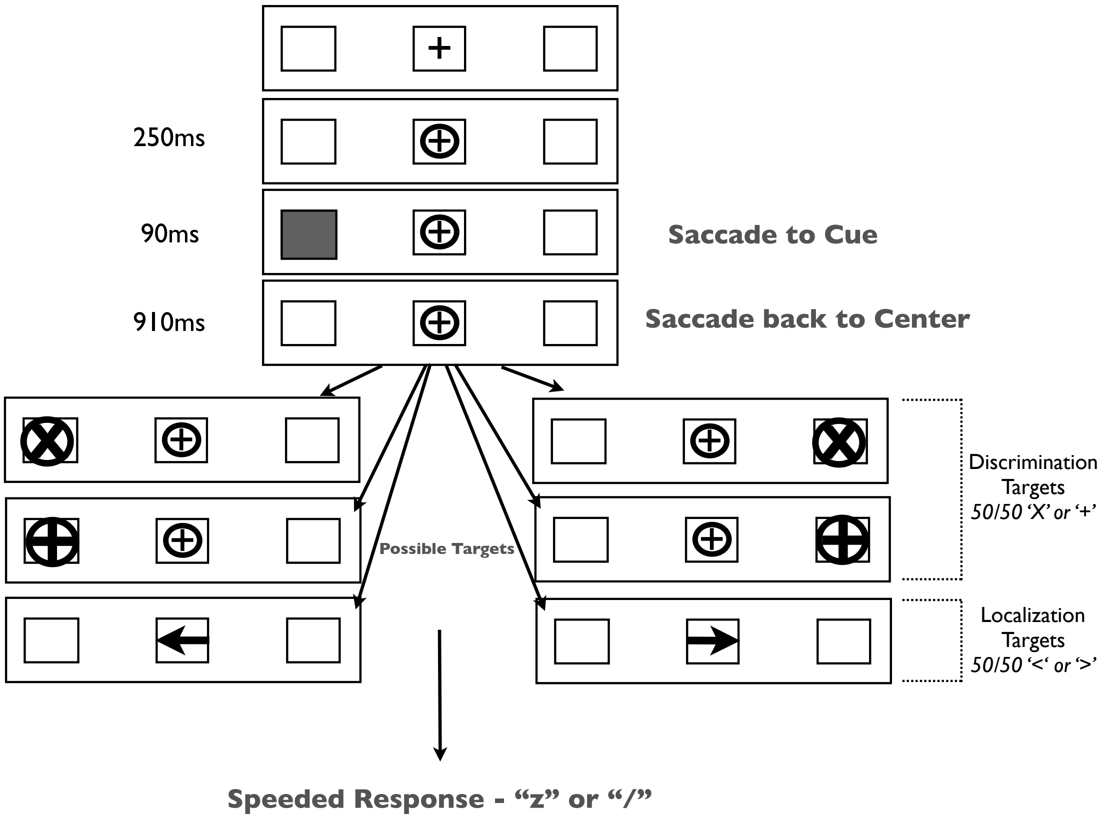

Figure 2.1. Methods figure depicting the time course of a trial. Duration of each subsequent event is depicted to the left of the image. Discrimination targets were counter-balanced for location (left or right placeholder) and identity ('X' or '+'). Localization targets were counter-balanced for direction (left or right pointing). The relative frequency of target type was manipulated between groups. This image is not to scale, so the relative size of features may be misrepresented.

2.2.1 Participants

24 (1 left-handed; 5 male) naive participants ranging in age from 17-24 (M = 19.8) participated in the study in a 60-minute session. Participants were compensated at a rate of either 1.0 course credits or $12 per hour. All participants were recruited from the undergraduate subject pool at Dalhousie University.

2.2.2 Apparatus and Procedure

The experiment was run in a dimly lit room on a 19” CRT monitor. Gaze position was monitored continuously by EyeLink II head mounted eye tracking equipment. Our stimuli and procedure (Figure 2.1) were identical to those in Redden, Hilchey and Klein (2016) except for a single change: participants were presented with one of two target types on each trial—peripheral x/+ discrimination or central left/right arrow targets.

Trials began with the presentation of three black outline placeholder boxes [1.5 x 1.5 degrees visual angle (DVA)] separated horizontally by 6.2 DVA on a grey background. The center box contained a black ‘+’ (0.5 DVA) as a fixation stimulus. Trials began with a drift correction that required the participant to fixate the central stimulus and to press the spacebar. If the participant was not fixating the central stimulus, then a tone alerted them to refixate. Upon fixation, a circle subtending 0.9 DVA encircled the fixation stimulus and remained onscreen for the duration of the trial. Two hundred fifty milliseconds (ms) after the appearance of the circle, one of the lateral placeholder boxes was cued by filling in the empty space with grey. This stimulus lasted for 90 ms and did not predict the target location. Participants were required to generate a saccade to the cued placeholder box and back to the fixation stimulus. Trials on which inaccurate (>3.0 DVA from the target or centre location) or early eye movements occurred (i.e., prior to cue onset) were terminated and recycled. After the successful eye movements, participants were instructed to maintain fixation for the duration of the trial. Target type was randomly selected on each trial. On discrimination trials, a target was presented in one of the lateral placeholder boxes (50% left, 50% right) 1000 ms after the onset of the cue. These targets were equally likely to be either an ‘X’ or a ‘+’ within a circle (1.3

DVA). Participants were required to identify the target by pressing either the ‘z’ or ‘/’ key, respectively. Arrow targets (1.0 DVA) were presented in the central placeholder box pointing either left or right. These targets required a speeded response indicating the direction of the arrow (‘Z’ for left; ‘/’ for right). All targets remained on screen until response. Participants completed 32 practice trials, followed by 320 experimental trials.

2.3 RESULTS

Trials on which inaccurate eye movements occurred prior to target onset were excluded (0.9%). Based on visual inspection of the overall RT histograms, anticipatory target responses (<250ms arrow: 0.1%; <300ms discrimination: 0.2%) and slow target responses (>800ms arrow: 3.5%; >1200ms discrimination: 4.5%) were excluded from analysis. After these exclusions, one participant was removed due to performance being close to chance in the discrimination task (Accuracy < 60%), leaving an N of 23 for analysis.

Generalized linear mixed effects models were used (GLMER - lme4 R package; Bates, Maechler, Bolker & Walker, 2015) to examine the trial-by-trial hierarchical relationship between predictor variables — Cueing and Simon — and the outcome variables — Reaction Time and Proportion Correct. The model did not converge when each predictor was treated as both a fixed and random effect, however removing Simon as a random effect afforded convergence. The interaction model was run first, followed by the main effect model, with AICs computed via the drop1 method in the {stats}

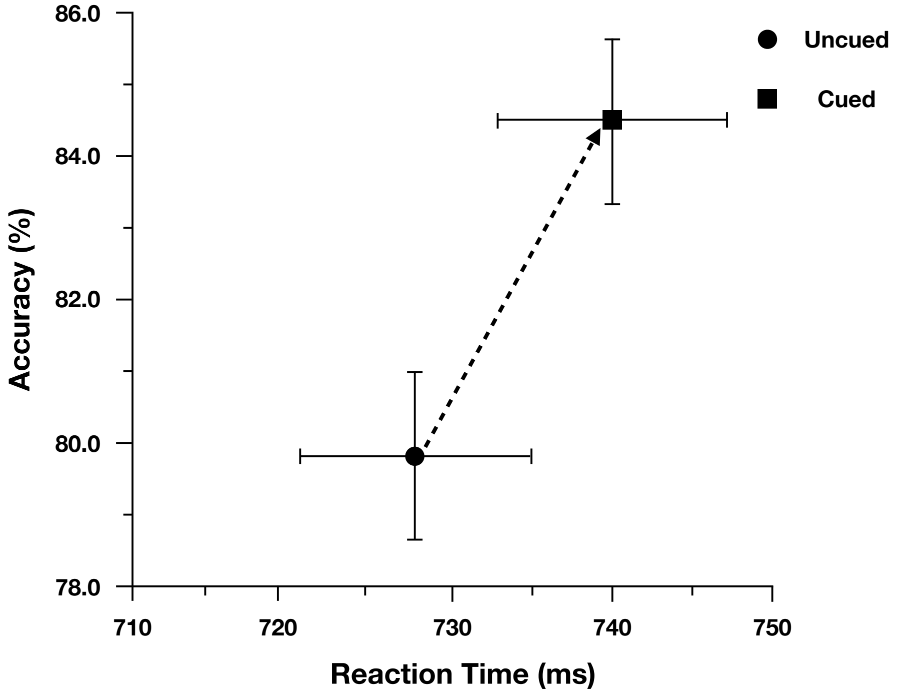

Figure 2.2. Overall performance for discrimination targets plotted in speed-accuracy space. Error bars represent Fisher's least significant difference.

package. Effect sizes for parameter estimates are reported as bootstrapped 95% confidence intervals, generated via confint.

2.3.1 Peripheral Discrimination Task

When examining Correct Reaction Time[2], participants were slower (13ms) and more accurate (4.5%) when responding to cued peripheral discrimination targets (Figure 2.2). Moreover, as predicted, the Simon effect was reduced for Cued targets (24ms) relative to Uncued targets (47ms; Figure 2.3). There was evidence to support the two-way interaction, Cueing x Simon, b = 0.029, CI95% = [-0.007, 0.059], as the model performed

---

[2] We performed the analysis on log(Correct RT), and report effect estimates in log space. However, for ease of interpretation we plotted Correct RT.

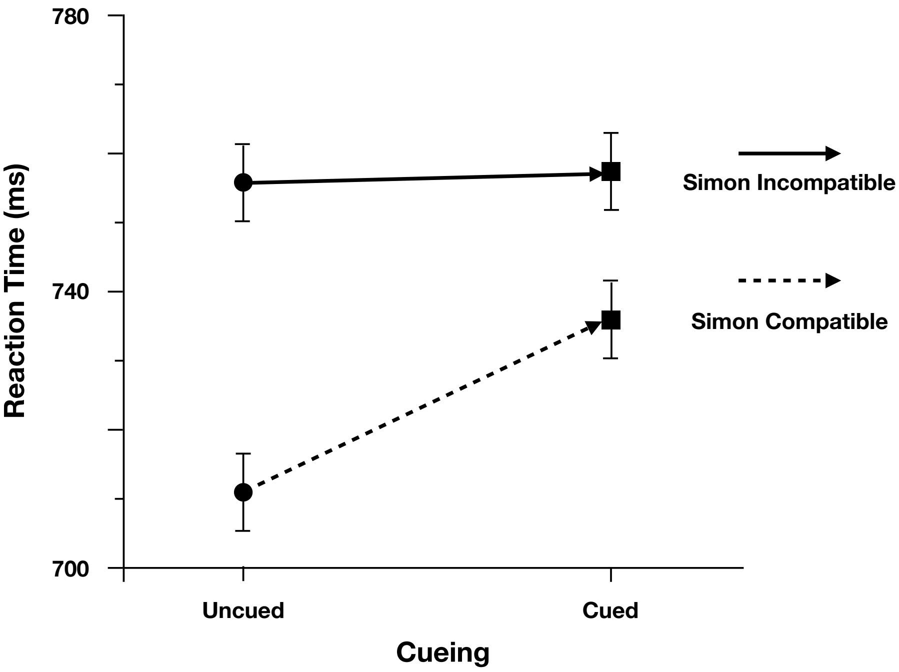

Figure 2.3. Reaction time to discrimination targets, split by Simon Compatibility (solid line = Simon Compatible; dashed line = Simon Incompatible). Performance to cued and uncued targets is represented on the x-axis. Error bars represent Fisher's least significant difference.

worse with the interaction term dropped (AIC = -971) than when the term was included (AIC = -973).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = -971) with models where each term was dropped. The model performed worse (△AIC = +1) when dropping the main effect of Cueing, b = -0.018, CI95% = [-0.038, 0.006]. The model also performed worse (△AIC = +36) when dropping the main effect of Simon, b = 0.047, CI95% = [0.032, 0.063].

When examining Proportion Correct, there was no evidence to support the two-way interaction, Cueing x Simon, b = -0.168, CI95% = [-0.560, 0.190], as the model

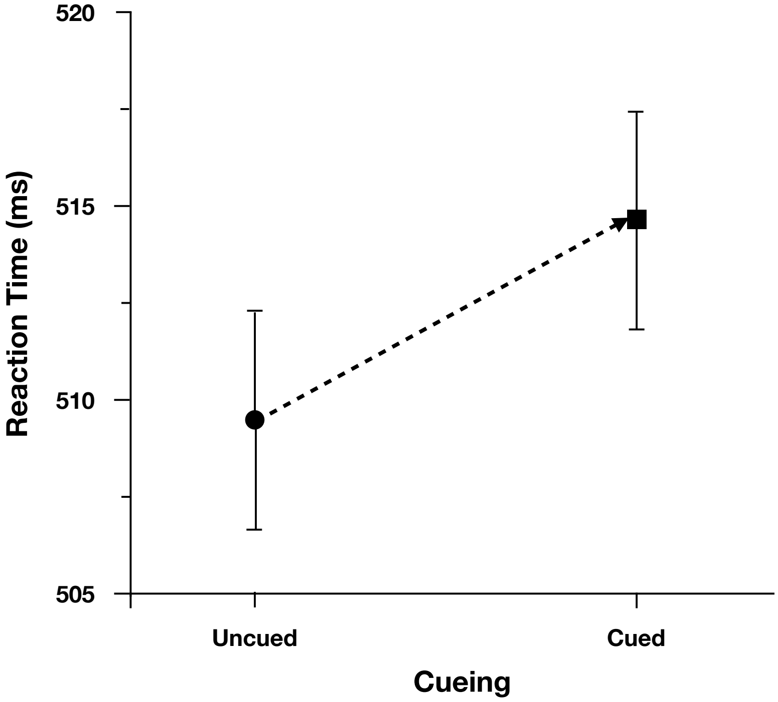

Figure 2.4. Reaction time to arrow targets. Performance to cued and uncued targets is represented on the x-axis. Error bars represent Fisher's least significant difference.

performed better with the interaction term dropped (AIC = 3013) than when the term was included (AIC = 3014).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 3013) with models where each term was dropped. The model performed worse (△AIC = +11) when dropping the main effect of Cueing (Cued = 84.5%, Uncued = 79.7%), b = 0.422, CI95% = [0.179, 0.649]. The model also performed worse (△AIC = +92) when dropping the main effect of Simon (Simon Compatible = 87.9%, Simon Incompatible = 75.9%), b = 0.911, CI95% = [0.693, 1.132].

2.3.2 Central Arrow Task.

Accuracy was not analyzed due to so few errors recorded to these targets (Cued = 98.2%, Uncued = 98.4%). Participants were slower (5ms) to respond to cued central arrow targets (Figure 2.4). There was evidence to support the main effect of Cueing, b = -0.012, CI95% = [-0.028, 0.001], as the model performed worse with the effect term dropped (AIC = -2359) than when the term was included (AIC = -2360).

2.4 DISCUSSION

This experiment yields three basic findings with respect to the nature of IOR caused by a pro-saccade: (1) Participants showed slower and more accurate responses for cued discrimination targets, consistent with an output-based form of IOR that could be generated by an increase in either response threshold or drift rate noise. (2) Arrow responses in the direction compatible with the cue were slower than arrow responses in the direction incompatible with the cue, and (3) the Simon effect was reduced at cued locations. The latter two effects are unambiguously consistent with an effect of IOR on response threshold.

The main point here is that we have provided empirical support for one of the two tenable mechanisms underlying the output form of IOR, as proposed by a computational model for the two forms of IOR (Redden, MacInness, & Klein, 2021). The findings do not rule out possible effects on the drift noise parameter, but they do ensure effects on the response threshold parameter.

As a secondary point, it is worth noting that whereas a meta-analysis of the literature reveals that the Simon effect is enhanced by IOR (Ivanoff, Klein & Lupianez, 2002), it can also clearly be reduced by IOR, depending on the kind of IOR that is generated (Redden, Hilchey & Klein, 2016). The result of the meta-analysis likely reflects an amalgam of different kinds of IOR from a literature that, on balance, just happens to succeed more often at generating the input form.

2.5 EXPERIMENT 2

Experiment 1 provides strong evidence that output-based forms of IOR can affect response thresholds but there is still ambiguity about the exact nature of the effect. Namely, when the inhibition from the cue expresses itself as a speed-accuracy tradeoff (SAT), it is not possible to tell unambiguously whether the effect is a shift along one or to a different accrual function (see Figure 1.5 - solid and dashed arrows). No such ambiguity exists when the inhibition from the cue is not expressed as a speed accuracy tradeoff (SAT, see Figure 1.5 - dotted arrows; e.g., Ivanoff & Klein, 2006; Chica et al., 2010; Hilchey, Hashish, MacLean, Ivanoff, Satel, & Klein, 2014; Redden, Hilchey & Klein, 2016). SATs can be due to a criterion shift alone, with slower but more accurate responses to cued as compared to uncued targets on a single accrual function (Figure 1.5 - solid arrow), or due to a criterion shift and a change in performance, with slower but more accurate responses to cued as compared to uncued targets on separate accrual functions (Figure 1.5 - dashed arrow).

The goal of Experiment 2 is to further clarify the nature of the SAT elicited in Experiment 1, Chica et al. (2010) and Redden, Hilchey and Klein (2016) by way of the response signal methodology (Wickelgren, 1977; Ivanoff & Klein, 2006; Zhao, Heinke, Ivanoff, Klein & Humphreys, 2011; Hilchey, Ivanoff, Taylor, & Klein, 2011). By forcing responses to targets at different times, this methodology generates functions in SAT space that represent accuracy as a function of response (or processing) time.

If the effect of IOR generated by an eye movement to a peripheral cue is best represented by an increase in the threshold parameter, ergo is operating purely on output stages of information processing, then performance from cued and uncued targets will belong to a single accrual function (representing the same accrual of information with shifts along the function as a result of IOR - Figure 1.5 solid arrow). Moreover, if the response-signal methodology allows for complete control over the speed of responding, no performance difference will be observed across cueing conditions. If there is some concomitant inhibitory effect on inputs, then the accrual function for cued (inhibited) targets will be shifted to the right (Figure 1.5 - dashed arrow).

2.6 METHOD

2.6.1 Participants

Eleven naive participants ranging in age from 19-32 participated in the study over five 60 minute sessions, one of whom was excluded for an inordinately high rate of target fixations (70% of trials; all others < 18%). Participants were compensated at a rate of $12 per session. All participants were recruited from the undergraduate subject pool at Dalhousie University.

2.6.2 Apparatus and Procedure

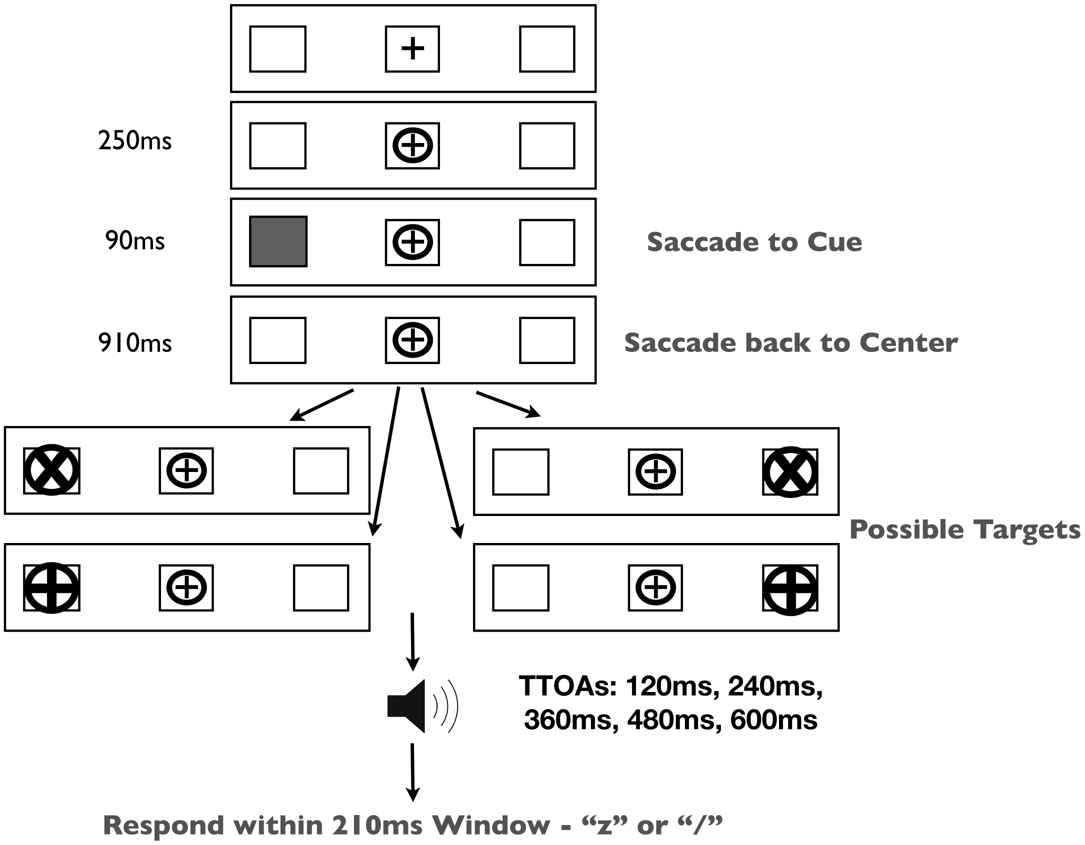

Figure 2.5. Methods figure depicting the time course of a trial. Duration of each subsequent event is depicted to the left of the image. Participants were required to execute their response within the response window indicated by a tone presented at a single TTOA in a given block. This image is not to scale, so the relative size of features may be misrepresented.

Our stimuli and procedure (Figure 2.5) were identical to those in Experiment 1 except for two changes: the response window method (described below) was added (i.e., responses were constrained to a predetermined experimental criterion), and participants were only presented with peripheral x/+ discrimination targets (i.e., arrow targets were removed).

Keypress responses were to be enacted after the onset of a tone. The target-tone onset asynchronies (TTOA) were 120, 240, 360, 480 or 600ms and the response window was 210ms. The target remained present until a response was made or until the response window had closed. Feedback was given onscreen when anticipatory ("Too Early!"), late responses ("Miss!") or untoward eye movements ("Inaccurate eye movement detected") were made. TTOA was blocked within sessions and the order within a session was random. In each session, each participant completed 5 blocks of 80 trials, one for each TTOA, for a total of 400 trials per participant, per session. Across the five sessions, there were thus 2000 trials (1000 cued and 1000 uncued) per participant.

2.7 RESULTS

Session One was considered practice and was excluded from analysis. Trials with inaccurate eye movements prior to target onset were excluded (7.1%). Analyses were performed on trials for which a response was recorded within the response window, resulting in the exclusion of 19.2% of trials[3]. Trials for which an eye movement was recorded during target presentation were excluded (8.1%)

Generalized linear mixed effects models were used (GLMER - lme4 R package; Bates, Maechler, Bolker & Walker, 2015) to examine the trial-by-trial hierarchical relationship between predictor variables — Processing Time[4], and Cueing — and the outcome variable — Proportion Correct. The model did not converge when each predictor was treated as both a fixed and random effect, however removing Processing

---

[3] A paired samples t-test was conducted on the number of responses falling outside the response window as a function of Cueing. These tests showed that there was no influence of cueing on the average frequency of early (Cued = 53, Uncued = 49; t(9) = 1.25, p = 0.24) or late (Cued = 82, Uncued = 83; t(9) = -0.33, p = 0.77) responses. Additionally, inclusion of these trials in subsequent models changes no statistical patterns or conclusions.

[4] Where Processing Time = Time from Tone to Response + TTOA

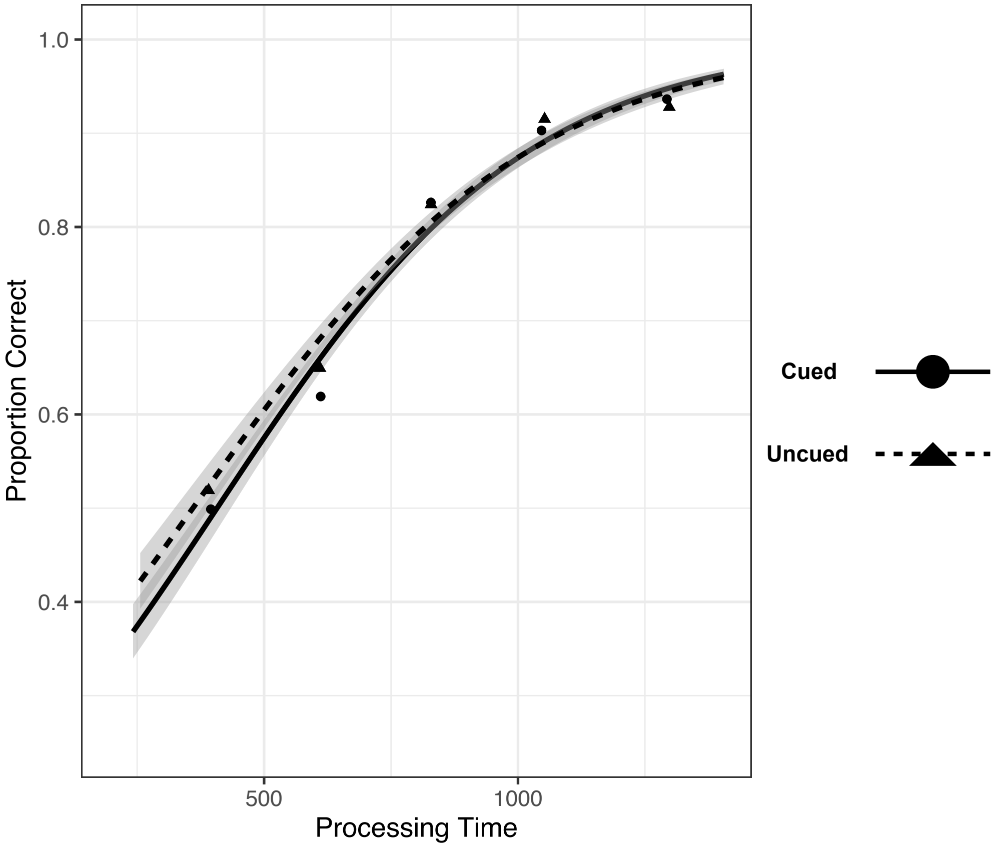

Figure 2.6. Information processing functions for cued (solid; circles) and uncued (dashed; triangles) performance. Processing time was calculated as the sum of the mean tone RT within a TTOA plus the TTOA. Grey shaded area reflects the 95% confidence interval for the fit of each cueing condition. Points represent the mean Processing Time and Proportion Correct for each Cueing condition for each of the five TTOAs.

Time as a random effect afforded convergence. The interaction model was run first, followed by the main effect model, with AICs computed via the drop1 method in the {stats} package. Effect sizes for parameter estimates are reported as bootstrapped 95% confidence intervals, generated via confint.

There was little evidence to support the two-way interaction (Figure 2.6), Processing Time x Cueing, b = 0.0002, CI95% = [-0.0001, 0.0005], as the model performed only slightly worse with the interaction term dropped (AIC = 11461) than when the term was included (AIC = 11460).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 11461) with models where each term was dropped. The model performed substantially worse (△AIC = +1730) when dropping the main effect of Processing Time, b = 0.003, CI95% = [0.015, 0.029]. Model performance was unchanged (△AIC = 0) when dropping the main effect of Cueing, b = 0.069, CI95% = [-0.017, 0.176].

2.8 DISCUSSION

These results suggest that input processing at cued and uncued locations is quite similar, with the proportion of correct discrimination responses being accounted for by a single information accrual function. This result would be expected if output-IOR were expressed purely as a criterion shift or post-perceptual effect. Further, since the same overt response was required at the same time to the same cue as in Experiment 1, the lack of a statistical cueing effect cannot be taken easily to suggest that there was no inhibitory consequence of the cue (i.e., IOR). That is, without the response signal methodology, these procedures robustly reveal that responses are slower when the target location is cued as compared to uncued. This pattern is qualitatively different from those reported by Ivanoff and Klein (2006). Eye movements to the cue were discouraged in their experiments, and they found an effect consistent with an input attribution (see Ivanoff & Klein, 2006 - Figure 5), whereby cued performance existed on a less efficient function than uncued (represented theoretically in the shift from solid to dotted functions in the present Figure 1.5).

2.9 GENERAL DISCUSSION

We have shown that when combining multiple measures of the output form of IOR in a single task (vis SAT pattern, arrow targets, relationship with the Simon effect), the effect is robust across the type of diagnostic. We observed inhibited performance for cued peripheral onset targets consistent with an SAT, inhibited performance for centrally-presented arrows in the direction compatible with the cued location, and the attenuation of the Simon effect for cued targets. Each of these effects are consistent with the patterns predicted for a cue-elicited inhibitory mechanism that reduces the propensity to make responses in the cued direction.

We have utilized the response-signal methodology as a method for investigating the time course of information processing dynamics in the aftermath of inhibitory cueing effects. This method has allowed us to determine how full information processing functions in a 2-AFC task are affected by IOR when generated by overt, prosaccadic orienting, and contrast these findings with prior studies utilizing this methodology to evaluate the effect of the input form of IOR. This reinforces the theory that the output form of IOR is not affecting the quality of information accrued at the cued location and converges with that theory's proposal that the output form is essentially a response bias.

When the present results are combined with the findings of the literature, it becomes clear that at least two qualitatively different mechanisms underlie IOR. Through the lens of drift diffusion modelling, the two forms of IOR can definitely be captured by differences in the drift rate (input form) and response threshold (output form) parameters. The parameter(s) that will best capture the nature of IOR depends critically on the activation state of the oculomotor machinery for reflexive eye movements.

These findings provide converging evidence for the theory that the critical factor determining the type of IOR observed is the activation state of the reflexive oculomotor system. Furthermore, both forms of IOR show behavioural effects that would accomplish the novelty-seeking function attributed in the seminal papers by Posner and Cohen (1984) and Posner et al. (1985)—however by altogether different mechanisms. The input form doing so by decreasing the salience of recently attended inputs, whereas the output form doing so by biasing orienting behaviors against previously attended locations.

CHAPTER 3: USING TEMPORAL ORDER JUDGMENTS AND COLOUR PERCEPTION TO DISSOCIATE INHIBITORY CUEING EFFECTS

The (unpublished) manuscript based on this study is presented below. The sole co-author for this manuscript is Dr. Raymond M. Klein. In this manuscript, my contributions include, but aren't limited to, writing, theory development, data collection, data analysis, project management, experiment design, and experiment programming.

This manuscript is not currently under review.

3.1 INTRODUCTION

Inhibition of return (IOR) is an inhibitory aftereffect of visuospatial orienting, typically observed in the spatial cueing paradigm by way of slower responses to cued rather than uncued targets. Early work on IOR using temporal order judgments (TOJ; Posner et al., 1985; Maylor, 1985; Kwak, 1992; Gibson & Egeth, 1994; for a review, see Klein, Schmidt & Müller, 1998) showed no effect on arrival time judgments, suggesting IOR was acting at a post-perceptual information processing stage, although at precisely what stage has been quite contested (Taylor & Klein, 2000; Abrams & Dobkin, 1994; Ivanoff & Klein, 2006; Ivanoff et al., 2002; de Vries et al., 2020).

Recent work suggests that there are in fact two forms of IOR; one operating nearer the input end of the information processing continuum affecting the quality of inputs and the other nearer the output end affecting the threshold for responding (Klein & Redden, 2018; Redden, MacInnes & Klein, 2021). Several studies using a variety of methods converge on this premise. Taylor and Klein (2000; see Section 1.2.3 & Figure 1.2) conducted a parametric investigation using a combination of central arrow and peripheral onset targets requiring localization responses, showing a dissociation in the nature of IOR between overt and covert orienting. Two-alternative forced choice tasks (see Section 1.3.3, Figure 1.5) have shown dissociations in the trading relationship between speed and accuracy for the two forms (Chica et al., 2010; Redden et al., 2016), whereby the input form manifests as a genuine cost in performance (slower RT with a cost in accuracy), whereas the output form manifests as a speed-accuracy tradeoff (SAT; slower RT with improved accuracy). Generating full SAT functions for each of the two forms using the SAT method also converges with this theory (Ivanoff & Klein, 2006; Hilchey et al., 2011; Redden, Hilchey, Aslam, Ivanoff & Klein, 2020; see Chapter 2 Experiment 2). The Psychological Refractory Period method (see Section 1.3.2, Figure 1.4) has shown that the two forms exist at different processing stages relative to a "central bottleneck" (Kavyani et al., 2017; Klein et al., 2018); and drift diffusion modelling of behaviour has implicated different parameters in decision tasks for the two forms (Redden, MacInnes, & Klein, 2021; see Appendix A). Importantly, these various methods converge with the ascription that the type of effect that is manifest is contingent upon the activation state of the reflexive oculomotor system at the time the effect is generated: the input form is generated when the reflexive oculomotor system is suppressed and the output form is generated when the reflexive oculomotor system is not suppressed.

The original notion that IOR is acting on a post-perceptual stage of information processing is rooted in findings from studies conducted prior to the conceptualization of two forms of IOR. Since generating the input form of IOR requires both the monitoring of eye movements and assiduous feedback about untoward oculomotor behaviour in order to ensure the reflexive oculomotor system is in fact suppressed (Hilchey, Hashish, MacLean, Ivanoff, Satel, & Klein, 2014), it is highly probable that most, if not all, of the studies that tested the effect of IOR on TOJ performance were testing the output form (Posner et al., 1985; Maylor, 1985; Kwak, 1992; Gibson & Egeth, 1994; Klein et al., 1998). This raises the question as to whether the input form is acting at a stage of processing early enough to affect perceptual arrival time. Following a strategy we have

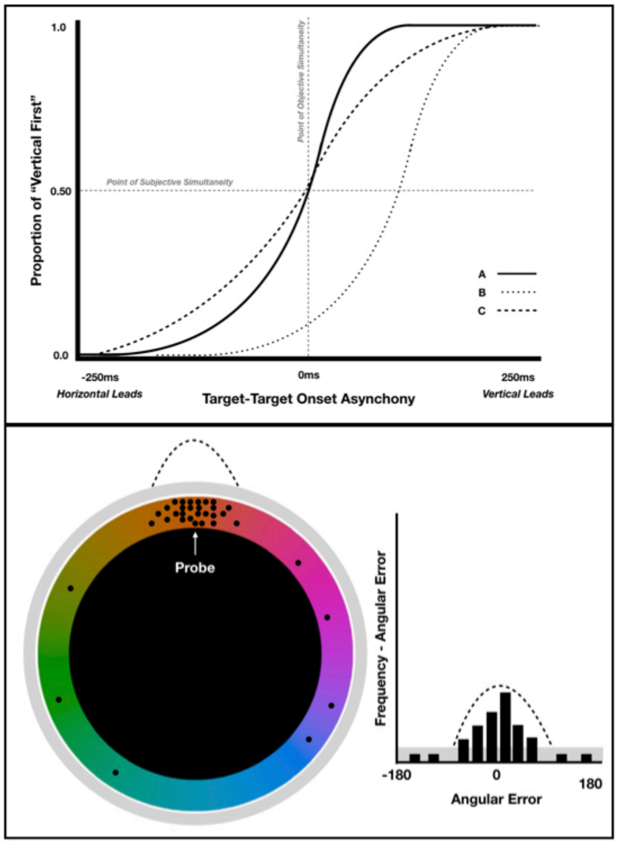

Figure 3.1. (TOP) Hypothetical results for the effect of an attentional manipulation on TOJs. The x-axis represents the target-target asynchrony (TTOA), or the amount of time between the presentation of the two targets to be judged. The y-axis represents the proportion of responses where the vertical stimulus is reported as appearing first. The point of objective simultaneity (not pertinent to our investigation) is represented by a vertical dotted line at TTOA = 0ms. The point of subjective simultaneity (PSS; integral to our investigation) is represented by a horizontal dotted line at Proportion of “Vertical First” = 0.50. Function A represents a prototypical performance function, where Vertical is rarely selected as appearing first when the Horizontal stimulus leads by a great deal (extreme left end of x-axis). The rate of choice for Vertical appearing first increases as the Horizontal stimulus leads by less, and becomes the more likely response when the Vertical stimulus starts to lead. Vertical is primarily reported as appearing first when the Vertical stimulus leads by a great deal (extreme right end of x-axis). Attentional manipulations could have two possible effects on performance. A shift from Function A to Function B would indicate a change in the PSS, whereby the perceptual experience is such that now Horizontal would need to lead by more (~100ms given the approximate PSS for Function B) in order for the attentional manipulation to be overcome (i.e, to counteract the attentional effect that shifted perception from A to B). A shift from

Figure 3.1. (Continued) Function A to C would indicate a change in the slope of the function, or the Just- Noticeable Difference (JND), representing the degree of sensitivity to the TTOA manipulation. As slopes increase, performance becomes more veridical (smaller JNDs), and as slopes decrease, performance approaches chance responding (larger JNDs). (BOTTOM) Hypothetical responses on the colour wheel are represented on the left. Each dot represents a trial; indicating the location of the mouse click relative to the probe location. Dots are clustered near the actual probe location, suggesting these responses are approximating the actual probe identity (better performance would be clustered more tightly). Others appear randomly distributed about the wheel relative to the actual probe. The informed responses are depicted by the dashed line above the colour wheel, and represented as such in the histogram on the right by a (circular) normal distribution. The random guess responses are depicted by the grey bar surrounding the colour wheel, and represented as such on the histogram on the right by a uniform distribution. Mixture modelling will afford computation of the relative contributions of these two distributions as the probability (Guesses vs Non-Guesses) and fidelity (variability in the circular normal) of encoding.

used in previous work (Redden, Hilchey & Klein, 2016), in the present study we generated IOR by requiring participants to make a pro-saccade or an anti-saccade to a peripheral stimulus (Rafal, Egly, & Rhodes, 1994). Because the reflexive oculomotor system is active when pro-saccades are required but must be suppressed in order to make a successful anti-saccade we expected that the peripheral stimuli in these two conditions would generate different forms of IOR at the cued location: the output form following pro-saccades; the input form following anti-saccades. We measured the effects of these two forms using the following procedures.

3.2 PRESENT STUDY

3.2.1 Measuring Attention with Temporal Order Judgements

Our TOJ trials require an unspeeded judgment of the temporal order of a vertical and horizontal line. These lines vary in their onset time (target-target onset asynchrony; TTOA), from easy (250ms) to very difficult (16.67ms). Performance is summarized in psychometric functions (Figure 3.1 - TOP), wherein the proportion of selecting one of the two target identities is plotted as a function of TTOA. Two parameters of interest are rendered from these functions: the point of subjective simultaneity (PSS) and the just-noticeable difference (JND). Mathematically, the PSS is the point on the psychometric function at which the probability of perceiving either target as first is .50. The PSS represents the onset asynchrony at which the two targets are perceived to be simultaneous; an effect on PSS would indicate a change in the speed of perceptual detection as a consequence of IOR. The JND is a measure of the slope of the psychometric function, as indexed by the change in onset asynchrony required to improve performance by some amount. The JND is proposed to represent temporal sensitivity, whereby smaller JNDs (steeper slopes) indicate greater sensitivity to the TTOA manipulation (for a review of the TOJ method, see Spence & Parise, 2010). TOJ stimuli will be presented either in the same placeholder box (unilateral trials) or in opposite boxes (bilateral trials), as a means of assessing the effect of IOR on sensitivity to temporal information (JND), and on the rate of information accrual (PSS), respectively. Subjects will be asked to report the temporal order of the stimuli by reporting which orientation appeared first. By reporting the orientation of the stimulus instead of location, the method conforms to the orthogonality criterion (Spence & Driver, 1994) that is important in spatial cuing studies to avoid the possibility of the location of the cue directly affecting response choices. Moreover, to minimize first- and second-order response biases, respectively (see Spence, Shore & Klein, 2001) different groups of participants will be asked to report either “Which First?” or “Which Second?”.

3.2.2 Measuring Attention with Colour Probes

Probe trials required a speeded detection response to a brief visual stimulus followed by an unspeeded colour matching (Figure 3.1 - BOTTOM). Detection responses were included as a manipulation check to see if we observed the canonical inhibited RT effect for both forms. The colour matching task is an orienting diagnostic pioneered by Prinzmetal et al. (1998), and is a relatively novel approach for assessing how the orienting of attention might affect information processing. By using a continuous measure of error in this task, we can model two components of behaviour that might be affected by IOR: probability of encoding the stimulus and, given it had been encoded, the fidelity of encoding (Zhang & Luck, 2008; Lawrence, 2010). Probability of encoding is represented as the proportion of non-guess responses, whereas fidelity of encoding is represented as the inverse of the spread of non-guess responses (see Figure 3.1 for description and visual representation of these parameters).

To our knowledge, only one other study has used the continuous measurement of error to explore performance at a late cue-target onset asynchrony (CTOA) following an uninformative peripheral onset. Hurst, Lawrence and Klein (2019) found IOR for speeded detection responses at an 800ms CTOA, and observed an associated cost in the probability of encoding. However, their findings are ambiguous with respect to our two forms theory: their subjects were instructed to maintain fixation, but no eye monitoring was employed to ensure this instruction was followed. Their lack of eye monitoring is problematic because the literature shows conflicting evidence for whether an input or output effect will be generated when eye monitoring is not employed. Using the SAT method (Wickelgren, 1977) in the absence of eye monitoring, Ivanoff and Klein (2006) showed that when IOR was generated by an ignored cue, performance shifted to a less efficient function (i.e., an input effect). However, Hilchey, Hashish, MacLean, Ivanoff and Klein (2014) obtained an input effect in a task that had repeatedly produced an output effect in the literature by introducing eye monitoring and assiduous feedback about erroneous eye movements. Importantly though, in spite of this ambiguity regarding the ability to accurately diagnose the form IOR, Hurst, Lawrence and Klein (2019) do demonstrate that IOR (of some form) is in fact sensitive to measurement with colour wheel responses.

3.2.3 Predictions

For TOJs, when the output form is generated, we predict no inhibitory effect on performance on either parameter (PSS or JND) since it is theorized to be affecting a post-perceptual stage of processing (Maylor, 1985; Posner et al., 1985; Kwak, 1992; Redden, MacInnes & Klein, 2021; Redden et al., 2020). However, when the input form of IOR is generated, which is theorized to be affecting a perceptual stage of processing, two tenable predictions arise. If this effect serves to impair the sensitivity to temporal information, we will observe an effect on JND in unilateral trials. If this effect serves to impair the rate at which information is processed, we will observe an effect on PSS in bilateral trials.

Furthermore, it is predicted that speeded detection responses to probes presented at the cued location will be slowed relative to probes presented at the uncued location. We consider colour response analysis to be exploratory in nature, as we have no a priori predictions with respect to the two forms of IOR. However, at least one of the two forms ought to influence the probability of responses (i.e., greater guess rate for cued probes), as first reported by Hurst et al. (2019).

3.3 EXPERIMENT 1

Here we examine the effect of the output form of IOR on perception. As described above, the output form of IOR is thought to be influencing post-perceptual processes, and is manifest as an increase in the threshold at which responses are executed to cued targets (Redden, MacInnes, & Klein, 2021). Furthermore, we believe this is the form of IOR that had been generated when the TOJ method was used in the past to examine the effect of IOR on perception—repeatedly demonstrating no inhibitory effect of cueing on performance.

3.4 METHOD

3.4.1 Participants

26 people participated in the study; 18 females, and 8 males (Median age = 22 yrs; 4 left-handed, 1 ambidextrous). Participants received financial compensation ($6/ 30 min) or class credit for their participation, and provided informed consent as approved by the Research Ethics Board, Dalhousie University.

3.4.2 Apparatus

Stimuli were presented on an iMac running OS X 10.11.4. Eye position was monitored by a head-mounted EyeLink II system. Images were displayed on a 27-in. monitor with a resolution of 2560 × 1440 pixels. Participants sat approximately 57 cm from the screen. All stimuli were presented on a dark grey background. The colour probes had a diameter of approximately 1.0 degree of visual angle (DVA). The fixation stimulus was a "+" sign of equal width and height (1.0 DVA) presented at the midpoint of the two peripheral placeholder boxes. The cue was presented by filling in one of the two peripheral placeholder boxes entirely with white for 50 ms. The boxes had equal sides of 4.0 DVA and were displayed 11.3 DVA to the left and right of the fixation stimulus. The horizontal and vertical TOJ stimuli had a length of 3.0 DVA. The colour wheel had a diameter of 10.6 DVA, with an annulus width of 2.7 DVA. Participants indicated their TOJ responses by pressing either the "8" or "2" key on a keyboard, to signal vertical or horizontal responses, respectively. Colour wheel responses were recorded via a mouse click.

3.4.3 Procedure

The sequence of events and trial types is illustrated in Figure 3.2. Each trial began with the fixation stimulus centred on the screen, with a hollow white box on either side of it. Participants launched each trial with a drift correction. Participants focused their gaze on the fixation stimulus and pressed space bar. Inaccurate fixation was indicated by a brief tone, and participants were instructed to try again. The fixation stimulus changed into an asterisk to indicate the trial was started successfully, and subsequent events would soon occur. 500 ms after the appearance of the asterisk, a cue appeared in one of the two placeholder boxes. Participants were required to make a fast and accurate eye movement to the cued placeholder and back to fixation. Trials with too slow or inaccurate eye movements (not executed by the time of target onset; landing > 3.0 DVA from centre of

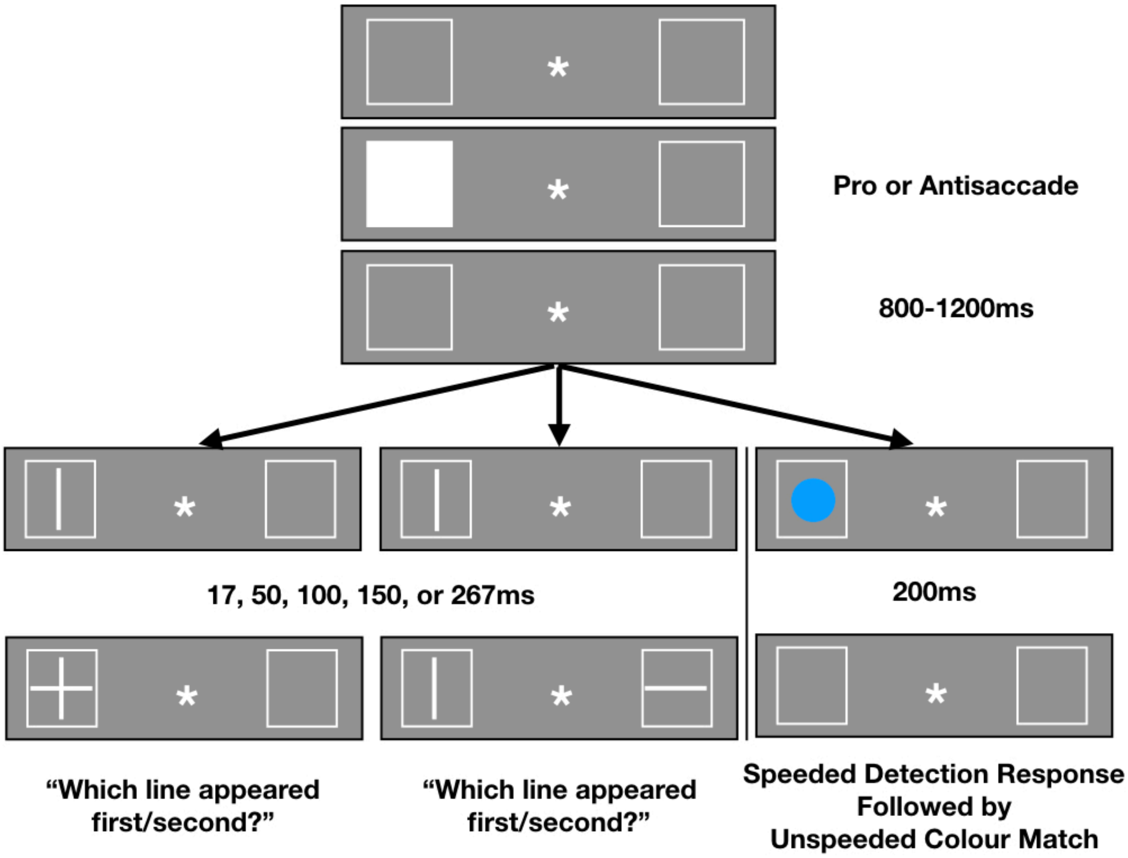

Figure 3.2 Sequence of events in a trial. Cues were equally likely to appear in the left or right placeholder box, requiring either an eye movement to the cued location and back to fixation (prosaccade group) or to the uncued location and back to fixation (antisaccade group). Target conditions were equiprobable: unilateral TOJ, bilateral TOJ, or colour probe. On unilateral TOJ trials, a horizontal and vertical line were drawn in the same placeholder box, with one of five possible onset asynchronies. On bilateral trials, a horizontal and vertical line were drawn in opposite placeholder boxes, with one of the same five onset asynchronies. Regardless of whether the stimuli were unilateral or bilateral, the response demands were the same: the subject reported as accurately as possible the perceived temporal order of the stimuli; either "Which line appeared first?" or "Which line appeared second?" using the '8' and '2' keys on the number pad to report either 'Vertical' or 'Horizontal', respectively. On colour probe trials, the subject was instructed to press the space bar as quickly as possible when the probe was detected. Then a constant-luminance colour wheel (see Figure 3.1 - Bottom) was presented, and subjects were required to match the colour of the probe as accurately as possible with a mouse click on the colour wheel.

appropriate placeholder; landing > 3.0 DVA from fixation upon gaze return), or trials where fixation was not maintained after returning gaze prior to target onset were aborted and recycled. On-screen feedback indicated the nature of the error to the participant.

After a random CTOA ranging from 800 to 1200 ms, one of two target types was presented: TOJ (2/3 trials) or colour probe (1/3 trials). Participants were instructed to maintain fixation when the target stimuli were presented, and to use their peripheral vision to the best of their abilities to process the targets.

On TOJ trials, a vertical or horizontal line was presented in either the left or the right box. The other line was equally likely to be presented in either the same box (unilateral trials) or opposite box (bilateral trials) after a given SOA (Figure 3.2). Each participant was required to make one of two TOJ judgments: “which line appeared first?” or “which line appeared second?”. TOJ judgment type was a between-subjects factor, with the levels “which first” (N = 15) and “which second” (N = 11).

On colour probe trials, the target was presented in the centre of one of the two peripheral placeholder boxes for 200ms. On these trials, participants were required to make a speeded detection response to the onset of the probe by pressing space bar. Then, given a detection response was successfully executed, they were then required to indicate the colour of the disk by using a constant-luminance colour wheel (Figure 3.1 - BOTTOM). The colour wheel was randomly rotated from trial to trial so as to avoid response biases to the positions of certain colours with respect to the computer screen and/or the mouse cursor. Colour wheel responses were recorded as angular deviations from the actual probe colour. Both trial types terminated after 3500ms if no response was recorded.

Participants were instructed that neither the cue, nor the location to which they had moved their eyes, were informative about subsequent target locations. Participants were also asked to emphasize accurate responding on TOJ trials and on colour wheel responses, and to emphasize speed when making detection responses to probes. Participants completed a block of 32 practice trials (not included in analysis), and three blocks of 96 experimental trials, each containing 32 bilateral TOJ trials, 32 unilateral TOJ trials and 32 probe trials.

3.5 RESULTS

Parameters for all dependent variables of interest for both experiments are summarized in Table 3.1.

3.5.1 TOJ Analysis

Two participants in the "Which Second" condition were excluded from analysis due to misunderstanding the instructions and reporting the first stimulus. Additionally, four other participants were excluded due to poor performance (JND > 125ms), leaving 13 and 7 participants in the Which First and Which Second conditions, respectively.

Psychometric functions were generated by fitting a generalized linear model using a binomial cumulative normal distribution (CND), predicting the proportion of responses indicating "Horizontal First" choices as a function of SOA (as in Redden, d'Entremont, & Klein, 2017a; 2017b). Data was fit for each participant, and included the factors, Cueing and Unilateral/Bilateral. Two parameters were extracted from each function for each condition: the PSS (the inferred SOA value when responding to each target is 0.50), and

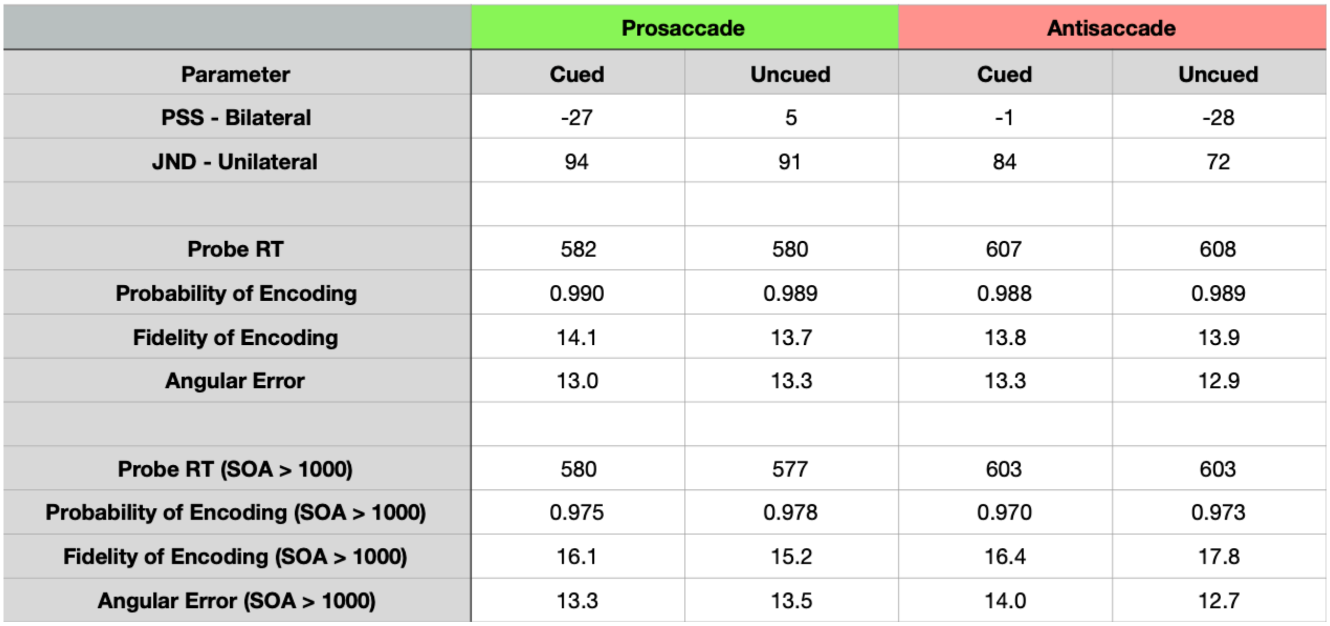

	Prosaccade		Antisaccade	
Parameter	**Cued**	**Uncued**	**Cued**	**Uncued**
PSS - Bilateral	-27	5	-1	-28
JND - Unilateral	94	91	84	72
Probe RT	582	580	607	608
Probability of Encoding	0.990	0.989	0.988	0.989
Fidelity of Encoding	14.1	13.7	13.8	13.9
Angular Error	13.0	13.3	13.3	12.9
Probe RT (SOA > 1000)	580	577	603	603
Probability of Encoding (SOA > 1000)	0.975	0.978	0.970	0.973
Fidelity of Encoding (SOA > 1000)	16.1	15.2	16.4	17.8
Angular Error (SOA > 1000)	13.3	13.5	14.0	12.7

Table 3.1. Table of mean values for each of the dependent variables of interest. Parameter names are listed in the leftmost column. Mean data for Experiment 1 (Prosaccade Condition) is listed in the middle columns. Mean data for Experiment 2 (Antisaccade Condition) is listed in the rightmost columns.

the JND (SOA difference from 0.33 to 0.67). Aggregate functions are plotted in Figure 3.3.

A t-test on the effect of Cueing on PSS in the Bilateral condition showed significant facilitation, t(19) = 3.87, p = 0.0010, CI95% = [15, 49], whereby the uncued stimulus needed to lead the cued stimulus by 32ms in order to be perceived as simultaneous. A t-test on the effect of Cueing on JND in the Unilateral condition showed no significant effect, t(19) = 0.24, p = 0.8156, CI95% = [-26, 33], whereby sensitivity to the SOA manipulation was uninfluenced by Cueing.

3.5.2 Colour Probe Analysis

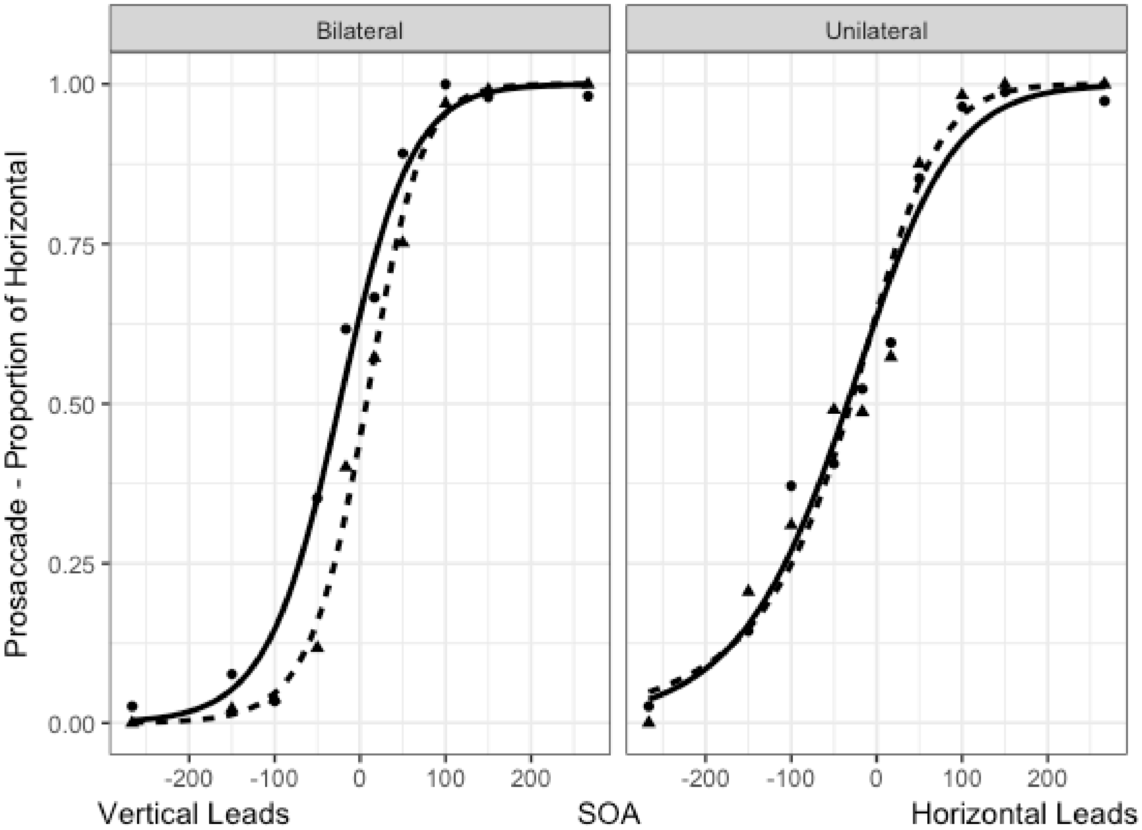

Figure 3.3 Prosaccade TOJ data plotted as the proportion of ‘Horizontal First’ choices reported as a function of Cueing (Cued = Solid; Uncued = Dashed) and SOA. Negative SOAs represent trials when the vertical stimulus was presented first. (LEFT) Trials in which the stimuli were presented in opposite placeholder boxes (one stimulus cued, one stimulus uncued). (RIGHT) Trials in which the stimuli were presented in the same placeholder box (both stimuli cued or both stimuli uncued).

The four participants who were excluded from the TOJ analysis for poor performance were also excluded from the colour probe analysis. RT cutoffs were established by binning RTs and assessing the total count of data within each 50ms bin. The lower bound was set at 300ms, which contained 21 trials (69 trials in 350ms bin, 3 trials in 250ms bin), resulting in trimming 8 trials. The upper bound was set at 850ms, which contained 23 trials (41 trials in 800ms bin, 12 trials in 900ms bin), resulting in trimming 34 trials. Three additional participants were excluded due to too few probe trials (24, 28 & 42; next fewest = 77, Mean = 92)[5].

A t-test on the effect of Cueing on RT showed no effect on Cued (582ms) vs. Uncued (580ms) probes, t(18) = 0.34, p = 0.7371, CI95% = [-7, 10]. An exploratory analysis was conducted wherein SOA was bifurcated into Short (800-1000ms) and Long (1000-1200ms). At Long SOAs, RTs to Cued (580ms) probes were again mathematically (but not significantly) slower than to Uncued (577ms) probes, CI95% = [-11, 18].

To examine colour accuracy, a mixture model consisting of a uniform and a Von Mises (circular normal) distribution was used. The mixture model produced two parameters: the probability and the fidelity of stimulus encoding. Probability of encoding represents the relative fit of Von Mises distribution over uniform distribution, whereas fidelity of encoding represents the narrowness of the Von Mises distribution. A t-test on the effect of Cueing on Probability of Encoding showed no effect on Cued (0.990) vs. Uncued (0.989) probes, t(18) = 1.356, p = 0.1919, CI95% = [-0.0004, 0.002]. A t-test on the effect of Cueing on Fidelity of Encoding showed no effect on Cued (14.1) vs. Uncued (13.7) probes, t(18) = 0.327, p = 0.7521, CI95% = [-2.22, 3.02]. Examining performance exclusively at Long SOAs showed no evidence that Cueing was influencing Probability (Cued = 0.975, Uncued = 0.978; CI95% = [-0.008, 0.003) or Fidelity (Cued = 16.1, Uncued = 15.2; CI95% = [-4.43, 6.15]).

---

[5] RT Cueing Effect CI95: Included = +/-9ms; Removed = +/-8ms

Lastly, ignoring the mixture modelling and comparing overall angular error also showed no effect on Cued (13.0 deg) vs. Uncued (13.3 deg) probes, t(18) = 0.5326, p = 0.6008, CI95% = [-1.29, 0.77]. Examining performance exclusively at Long SOAs showed no evidence that Cueing was influencing performance (Cued = 13.3 deg, Uncued = 13.5 deg; CI95% = [-2.09, 1.79]).

3.6 DISCUSSION

We have shown across multiple measures that when IOR is generated by a prosaccade, the resulting effect does not impair perception. As shown by TOJ performance, the output form of IOR is clearly not impairing temporal perception: Cueing has no effect on JNDs on unilateral trials, suggesting the output form is not influencing the sensitivity to temporal information. Cueing did have an effect on PSS on bilateral trials, however this effect was not inhibitory. Rather it appears that output IOR may be generating a form of prior entry (Shore, Spence & Klein, 2001; Spence & Parise, 2010), in that the uncued stimulus must lead the cued stimulus in order for them to be perceived as simultaneous. Additionally, we observed no inhibitory effect on colour perception; neither Probability nor Fidelity of encoding were convincingly impaired by cueing, nor was overall angular error. Although, mathematically speaking cued targets were responded to more slowly than uncued targets this effect was unusually small and not significant. Moreover, because detection RTs were unusually slow (no doubt to the complicated nature of the task), and previous findings suggest that the time course of IOR may be delayed when a task is more complex (Lupianez et al., 1997), we computed this effect separately for just the longer CTOAs and this made no statistical difference.

3.7 EXPERIMENT 2

Here we examine the effect of the input form of IOR on perception. As described above, the input form of IOR is thought to be influencing the quality of information accrued from the cued location, and is manifest as a genuine cost in performance. Furthermore, we believe this is the form of IOR that has heretofore been untested using the TOJ method, as generating the input form of IOR requires suppression of the reflexive oculomotor system, which as noted in the introduction we do not believe has been done in studies of the effect of IOR on TOJs. In Experiment 2, we accomplished this by requiring antisaccades to the uncued location (rather than prosaccades as in Experiment 1), which has not been done in any previously published IOR+TOJ studies.

3.8 METHOD

3.8.1 Participants

29 people participated in the study; 18 females, and 11 males (Median age = 19 yrs; 4 left-handed). Participants received financial compensation ($6/ 30 min) or class credit for their participation, and provided informed consent as approved by the Research Ethics Board, Dalhousie University.

3.8.2 Stimuli and Procedure

Stimuli and Procedure were identical to those used in Experiment 1, except for the eye movement instruction. In this experiment, participants were required to make a fast and accurate eye movement to the uncued placeholder (antisaccade) and back to fixation.

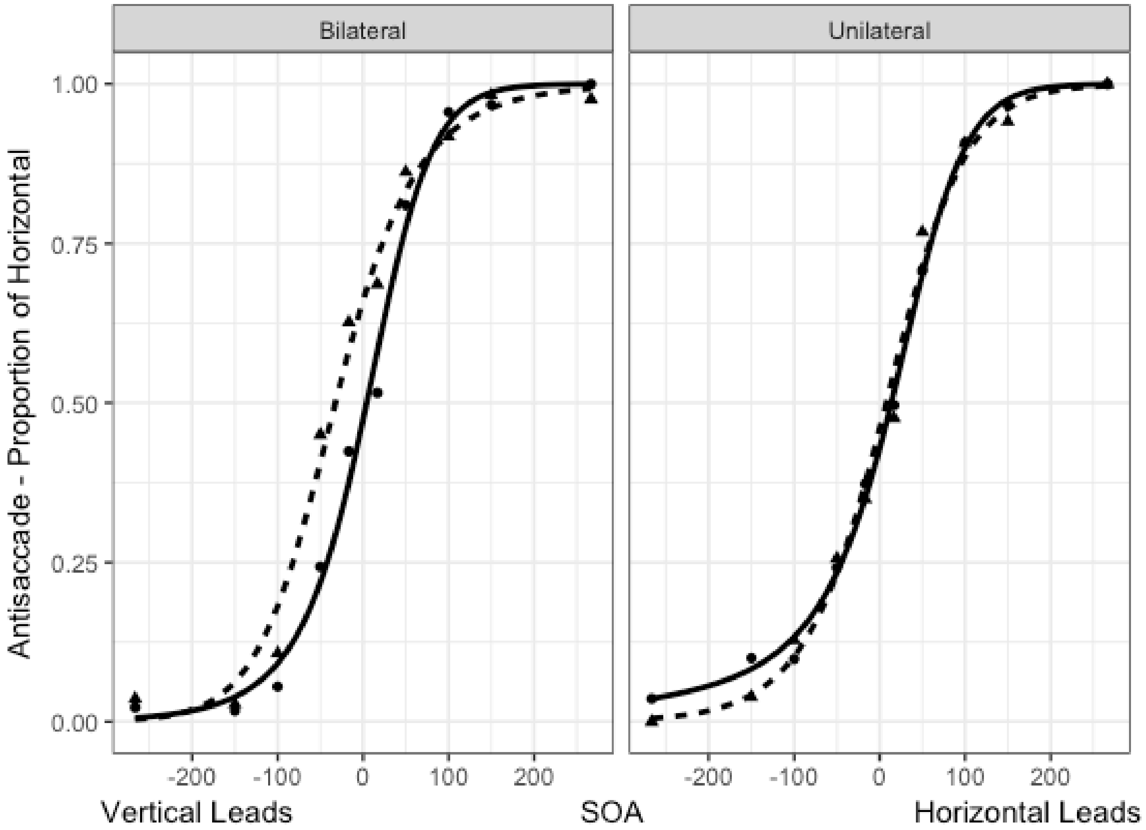

Figure 3.4 Antisaccade TOJ data plotted as the proportion of 'Horizontal First' choices reported as a function of Cueing (Cued = Solid; Uncued = Dashed) and SOA. Negative SOAs represent trials when the vertical stimulus was presented first. (LEFT) Trials in which the stimuli were presented in opposite placeholder boxes (one stimulus cued, one stimulus uncued). (RIGHT) Trials in which the stimuli were presented in the same placeholder box (both stimuli cued or both stimuli uncued).

All other aspects of the experiment remained the same as outlined above in the Method of Experiment 1.

3.9 RESULTS

Analyses were conducted using the same models as used in Experiment 1.

3.9.1 TOJ Analysis

Two participants in the “Which Second” condition were excluded from analysis due to misunderstanding the instructions and reporting the first stimulus. Additionally, four other participants were excluded due to poor performance (JND > 125ms), leaving 12 and 11 participants in the Which First and Which Second conditions, respectively. Aggregate functions are plotted in Figure 3.4.

A t-test on the effect of Cueing on PSS in the Bilateral condition showed significant inhibition, t(22) = 3.04, p = 0.0059, CI95% = [9, 46], whereby the Cued stimulus needed to lead the Uncued stimulus by 27ms in order to be perceived as simultaneous. A t-test on the effect of Cueing on JND in the Unilateral condition showed no significant effect, t(22) = 1.62, p = 0.1187, CI95% = [-3, 27], whereby sensitivity to the SOA manipulation was uninfluenced by Cueing.

3.9.2 Colour Probe Analysis

The four participants who were excluded from the TOJ analyses for poor performance were also excluded from the colour probe analyses. RT cutoffs were established by binning RTs and assessing the total count of data within each 50ms bin. The lower bound was set at 350ms, which contained 34 trials (65 trials in 400ms bin, 7 trials in 300ms bin), resulting in trimming 13 trials. The upper bound was set at 950ms, which contained 17 trials (20 trials in 900ms bin, 6 trials in 1000ms bin), resulting in trimming 38 trials. Two additional participants were excluded due to too few probe trials (25 & 28; next fewest = 75, M = 91 trials)[6].

---

[6] RT Cueing Effect CI95: Included = +/-10ms; Removed = +/-7ms

A t-test on the effect of Cueing on RT showed no effect on Cued (607ms) vs. Uncued (608ms) probes, t(22) = 0.576, p = 0.5701, CI95% = [-9, 5]. Following suit with Experiment 1, exploratory analysis wherein SOA was bifurcated into Short (800-1000ms) and Long (1000-1200ms) was conducted. Again at Long SOAs, RTs to Cued (603ms) probes were not significantly slower than to Uncued (603ms) probes, CI95% = [-11, 11].

A t-test on the effect of Cueing on Probability of Encoding showed no effect on Cued (0.988) vs. Uncued (0.989) probes, t(22) = 1.382, p = 0.1809, CI95% = [-0.003, 0.0006]. Additionally, a t-test on the effect of Cueing on Fidelity of Encoding showed no effect on Cued (13.8) vs. Uncued (13.9) probes, t(22) = 0.053, p = 0.9585, CI95% = [-2.46, 2.34]. Examining performance exclusively at Long SOAs showed no statistical evidence that Cueing was influencing Probability (Cued = 0.970, Uncued = 0.973; CI95% = [-0.02, 0.01]) or Fidelity (Cued = 16.4, Uncued = 17.8; CI95% = [-9.2, 6.2]).

Lastly, comparing overall angular error also showed no effect on Cued (13.3 deg) vs. Uncued (12.9 deg) probes, t(22) = 0.871, p = 0.3929, CI95% = [-0.49, 1.19]. Examining performance exclusively at Long SOAs showed color matching of Cued (14.0 deg) probes was less accurate than Uncued (12.7 deg), t(22) = 1.995, p = 0.029 (one-tailed), CI95% = [-0.05, 2.63].

3.10 DISCUSSION

This experiment is the first demonstration of an inhibitory cueing effect on TOJs (Figure 3.5). The pattern wherein bilateral trials are affected by inhibitory cueing but unilateral trials are not, suggests that the input form of IOR is operating to delay the

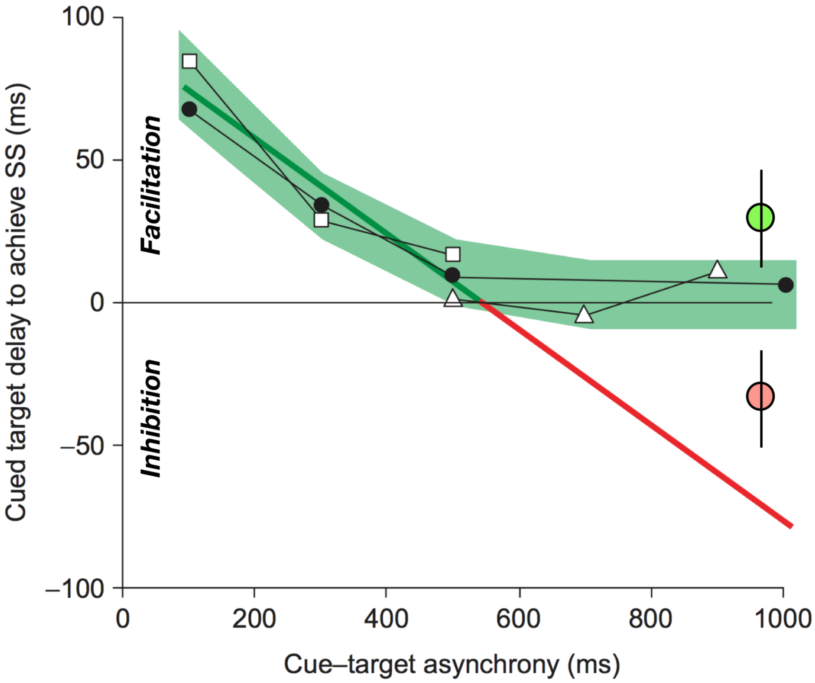

Figure 3.5. Graphical representation of cueing effects in TOJ experiments as a consequence of cue-target onset asynchrony. Values above the 0 line represent cue-elicited facilitation, and values below the line represent inhibitory cueing effects. The three studies depicted in black and white shapes (Closed circles - Maylor, 1985; Open squares - Gibson and Egeth, 1994 Exp 1; Open triangles - Gibson and Egeth, 1994 Exp 2) show cue-elicited facilitation at short CTOAs that decays to no effect (trend highlighted in green) at longer CTOAs when one would expect IOR to be influencing behaviour (trend delineated by red line). Our data (plotted in green and red circles; error bars = CI95%) show that when a prosaccade is generated to the cue, performance does not look different from what was observed previously. However, when an antisaccade is generated to the cue, the prediction is confirmed. (Redrawn from Klein, 2000).

accrual of information at the cued location, rather than impairing the quality of that information. Moreover, as reported in Experiment 1, the evidence for IOR influencing probes was nearly nonexistent. However, performance on cued probes was worse than performance on uncued probes as indicated by angular error of colour matching responses when the interval between cue and target was sufficiently long, but this effect was small and based entirely on criteria established post hoc.

3.11 GENERAL DISCUSSION

We found no inhibitory effect on TOJs in the group that elicited the output form, supporting our hypothesis that output IOR is affecting a late stage of information processing. However, we did find inhibited perceptual processing in the TOJ task when the input form was generated whereby cued stimuli needed to be presented earlier in order to be perceived as simultaneous with uncued stimuli, supporting our hypothesis that input IOR is affecting an early stage of information processing.

Pertaining to our manipulation check, in neither experiment was IOR convincingly observed in RT to probe targets. Neither form elicited a convincing version of the canonical RT effect (the output form did show mathematically slower RTs to cued targets), but when considerations for the time course of the effect were included post hoc, the input form of IOR appeared to be impairing the accuracy of colour matching responses — consistent with the hypothesis that this form of IOR operates to impair the quality of information at the cued location. That we didn't obtain robust IOR effects on our probe trials is not unexpected; previous studies have shown that IOR is weakened or non-existent for rare targets. For instance, Ivanoff and Klein (2004) showed reduced IOR for low probability targets relative to more frequent targets in both a Go/No-Go task and 2AFC task. Moreover, when stimulus-response probabilities were instantiated across blocks, Taylor (2007) also found reduced IOR for unexpected targets requiring a 2AFC response. Lupianez, Ruz, Funes, & Milliken (2007) found no IOR for rare targets (but robust IOR for frequent targets) when also using a 2AFC task. Importantly, none of these studies employed a task-switch between rare and frequent targets — each simply manipulated whether a target stimulus within a specific task was frequent or infrequent. In our study, not only were the probes rare, but they were a qualitatively different task (requiring a speeded detection response followed by an unspeeded colour match) relative to the frequently-occurring unspeeded temporal order judgement. As Lupianez et al. (2007) argue, the lack of an IOR effect on low probability targets "cannot be attributed to differences in attentional capture or to differences in orienting–reorienting of attention after capture" since any modulation of the resulting effect by these factors would influence both low and high probability targets because the target type is not known until the target appears. This suggests, instead, that perhaps under conditions that make the task difficult (i.e., when various different S-R probability rules are in place), observers employ a mental set to prepare for the most likely scenario (i.e., the TOJ task in the present study; the frequent targets in the other examples). When the unlikely target occurs, a task switch is required. The degree of complexity of this task switch is likely commensurate with the amount of reduction of IOR in the infrequent task. For instance, in Ivanoff and Klein (2004), there was only a small reduction of IOR as a consequence of frequency, as the complexity of the switch between frequent and infrequent targets was quite simple ('X' vs 'O' discrimination). The same logic applies to Taylor (2007), who performed an ostensible replication of Ivanoff and Klein (2004). Lupianez et al. (2007) employed a more complex design with various S-R contingencies and found IOR was completely abolished when rare targets occurred. This is consistent with our findings, where a complete shift in task demands was required when probes appeared. Further research is required to test this ad hoc hypothesis accounting for the frequency effect.

In spite of the lack of a convincing IOR effect on probes, our findings provide converging evidence that there are two forms of IOR: an input form influencing what we will attend and become aware of, and an output form influencing the behaviours (including orienting) in which we are likely to engage. Both forms can bias orienting and hence can perform the novelty-seeking function attributed to the inhibitions in the seminal papers; they just do so differently (Klein & Redden, 2018; Redden, MacInnes & Klein, 2021). This theoretical framework is of significance to cognitive science in general because IOR has been shown to influence both overt and covert orienting (Taylor & Klein, 2000), search behaviour (Klein, 1988; Muller & von Muhlenen, 2000; Takeda & Yagi, 2000), and has been shown to play a critical role in computational models of search (Itti & Koch, 2001). Specifically, these findings are an important advancement in the understanding of inhibitory effects on orienting, as we are the first in many attempts to show that (one form of) IOR influences perceptual processing by way of inhibited temporal perception.

CHAPTER 4: IS INHIBITION OF RETURN AT THE MIDPOINT OF SIMULTANEOUS CUES “ INPUT-BASED” OR “ OUTPUT-BASED”?

The (unpublished) manuscript based on this study is presented below. Co-authors for this manuscript are Dr. Maryam Kavyani, Dr. John Christie, and Dr. Raymond M. Klein. In this manuscript, my contributions include, but aren’t limited to, writing, theory development, data collection, data analysis, project management, experiment design, experiment programming.

This manuscript is not currently under review.

4.1 INTRODUCTION

The term inhibition of return (IOR) refers to slower reaction times (RTs) to targets presented at previously stimulated or inspected locations (Klein, 2000). IOR is generated by both overt and covert orienting of attention (i.e., when the eyes move to a location or remain at fixation during a cue, respectively; Posner et al., 1985). A standard conceptualization of the cause of IOR is that it is generated by the initial capture of attention to the cued location (Klein & Taylor, 1994; Berlucchi, 2006). An alternative view is that the sudden onset of a cue in the visual periphery generates IOR as a result of the activation of an oculomotor programme (e.g., Rafal et al., 1989). Under this view, even if no eye movements are required, a peripheral cue can produce the automatic activation of an eye movement to the cued location. This hypothesis is consistent with the idea that the superior colliculus (SC) plays a role in the generation of IOR (see Dorris et al. 2002; Posner et al. 1985; Rafal et al. 1989; Sapir et al. 1999), as the SC is a component of the oculomotor pathways involved in the programming of eye movements (Schiller, 1977).

4.1.1 Inhibition of Return Following Multiple Element Cueing Arrays

Research investigating the existence of IOR at the midpoint of simultaneous stimuli or centre of gravity (COG) supports the role of the SC in generating IOR. Saccadic averaging, also known as the ‘global effect’ or the ‘centre of gravity effect,’ occurs when two or more targets are presented simultaneously within the peripheral region of the visual field. The resulting saccade, rather than landing at one of the stimulated locations, often lands at the midpoint or ‘centre of gravity’ of the separate elements (Findlay, 1982). Using manual responses to onset targets following a single cue or the simultaneous presentation of two, three, or four cues around an imaginary circle, Klein, Christie, & Morris (2005; see also Langley, Gayzur, Saville, Morlock, & Bagne, 2011) inferred that the effect of IOR following multiple cues may be operating at the centre of gravity of the grouped array rather than the local elements, due to seeing little to no inhibitory effect at any of the locations of the individual elements of the array. They suggested that this IOR effect is due to population coding in the oculomotor pathways (e.g., the superior colliculus) directing the eye movement system toward the centre of gravity of the cue array. Christie et al. (2013) replicated this method of Klein et al (2005) using saccadic localization responses rather than manual detection. If the population coding hypothesized by Klein et al. (2005) were implemented within the oculomotor pathway, then Christie et al proposed that they ought to see the same, or similar, pattern of IOR when saccades rather than manual responses were made; which they did. Christie et al posed this question to the data from both the manual and saccadic experiments: "When the visual system is confronted with multiple, simultaneously occurring cues, does IOR occur at each cue location (local), or does it occur at the midpoint of the cues (global), as might be expected if IOR is generated by the machinery responsible for orienting?" Using regression methods, they found that the IOR effect measured by both manual and saccadic responses is best explained by the distance between the target and the centre of gravity of the cue array.

4.1.2 Two Forms of Inhibition of Return in Speed-Accuracy Space

Each previous chapter in this dissertation has outlined the framework and theoretical evidence supporting the proposal that there are two forms of IOR; the form of which is contingent upon the state of the reflexive oculomotor system (for a review on the SAT diagnostic, see Chapter 1 Section 1.3.3). There is much empirical support for the conclusion that the inhibitory effect reported in Christie et al. (2013) is the output form of IOR, as oculomotor responses were explicitly required to the target. The findings from Klein, Christie and Morris (2005) are ambiguous due to the fact that manual detection responses were required (which afford no indication of accuracy of responding), and oculomotor behaviour was monitored, but no feedback was provided when eye movements were detected. Indeed, neither experiment is conclusive regarding the form of IOR, since neither employed a diagnostic sensitive to detecting the different effects. Here, we seek to address this gap, and determine explicitly whether the effect at the centre of gravity of multiple cues is of the input or output form.

4.2 EXPERIMENT 1

We will employ a 2AFC task, which will afford measurement of speed and accuracy, and manipulate between-groups whether online eye movement feedback is provided to observers. Hilchey, Hashish, MacLean, Satel, Ivanoff & Klein (2014) showed that actively suppressing the oculomotor system in a manual response task changed the form of IOR from output-based to input-based. As in Hilchey et al. (2014), when observers are not given feedback on their untoward eye movements, we expect to elicit the output form of IOR. However, when eye movements are strictly forbidden due to trial-by-trial feedback, we expect to elicit the input form of IOR. The “Single Cue” trials will be exceptionally analytic to our hypothesis: we should see slower RTs for cued targets relative to uncued targets. The group that receives feedback ought to have a corresponding cost in response accuracy, whereas the group that does not receive feedback ought to have a corresponding improvement in accuracy, as observed by Chica et al. (2010), Hilchey et al. (2014), Redden et al. (2016), and the present Chapter 2 Experiment 1. The novel empirical insight will come from trials with multiple cues. Christie et al. (2013) observed inhibitory contributions from both the centre of gravity of multiple cues and the local elements (albeit smaller for the latter) when saccadic responses were required. Although we expect to replicate this pattern in reaction time in our No Feedback (output form) condition, it is to be determined what will result for accuracy. The Feedback (input form) condition is heretofore untested in the multiple-cueing array paradigm, so if there is IOR it should not be accompanied by an SAT. If there is in fact IOR, the question remains as to whether it will be at the CoG, at the locations of the individual elements, or at both loci.

4.3 METHODS

4.3.1 Participants

Forty-eight students (36 females and 12 males) from Dalhousie university volunteered to participate in this study. Their ages ranged from 18 to 43 years, and five of them were left-handed. They were compensated one point of course credit. The research was approved by the relevant Dalhousie Research Ethics Board.

4.3.2 Apparatus and stimuli

Participants were tested in a dimly lit room, and all stimuli were presented against a black background, at a viewing distance of 57cm on a 19” ViewSonic Optiquest Q95 connected to an Apple Mac mini with an Intel Core Due processor. An EyeLink II head-based eye monitoring system tracked eye movements from the initiation of the trial until the manual discrimination response. After participants had performed a nine point calibration procedure to calibrate and validate the precision of eye tracker, the Eye-link sampled gaze position once every 4ms.

The stimuli were presented in this experiment were as similar as possible to a combination of those administered by Klein, Christie and Morris (2005) and Chica, Taylor, Lupianez and Klein (2010). The fixation stimulus was a white cross measuring 0.5*0.5 degrees visual angle (DVA), unfilled white squares measuring 2 DVA were used as cues; and the target was either a green or red filled circle, 1 DVA in diameter. The cues and targets were presented at eight equally spaced locations on an imaginary circle with a radius of 8.5 DVA, demarcated by grey placeholders 2 DVA in size.

4.3.3 Procedure and Design

The sequence of events in every trial is shown in Figure 4.1. After the drift correction, the fixation display was presented for 200 ms. The participant was instructed to fixate the white fixation cross for the entire trial. A manual spacebar response was required when the participant had successfully fixated the fixation cross. When stable fixation had been achieved, one, two, or four cues appeared simultaneously for 200 ms at any of eight randomly selected locations along the circumference of an imaginary circle.

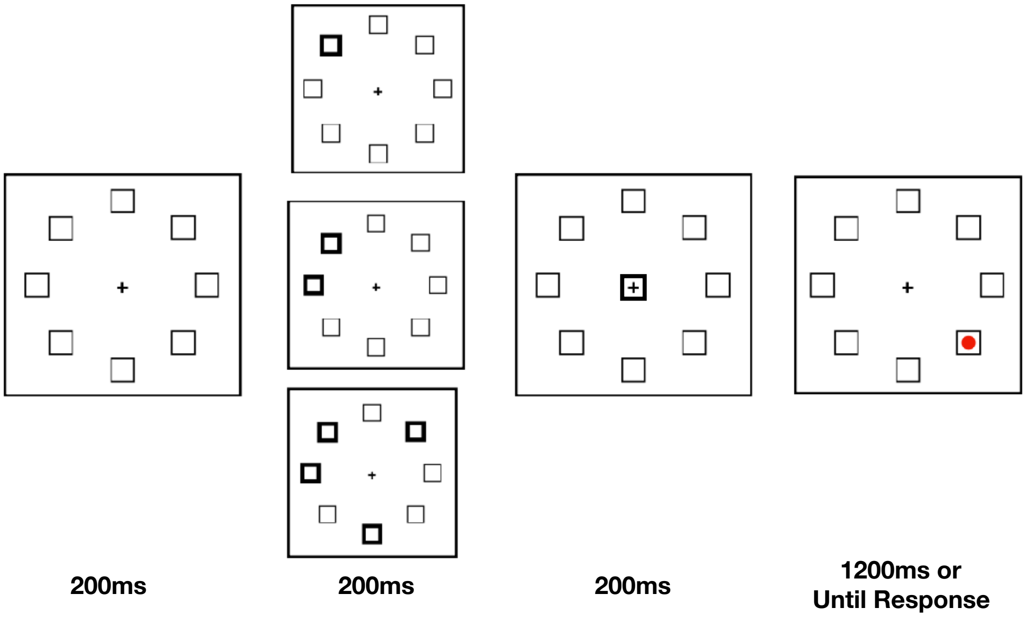

Figure 4.1: Sequence of events in a trial. Trials start with a fixation display prior to a cueing array (1, 2, or 4 cues; 200 ms display). Cues were uninformative regarding subsequent target locations. After an SOA of 400ms, a target was drawn at one of the eight placeholder locations, requiring a speeded discrimination response. Whether feedback was given on untoward eye movements was manipulated between groups.

The participants were told that these cue(s) were irrelevant, and did not predict the forthcoming target's location. Once the peripheral cue display was terminated, a central cue (cue-back) appeared at the centre of the screen for another 200 ms. The cue-back display was immediately followed by a single target stimulus presented at any of the eight candidate locations, to which participants were instructed to make a speeded and accurate manual response. Participants were instructed to perform a colour discrimination task by pressing one of the keys (a or ‘) for a green target and other one for a red target; key designation was counterbalanced across participants. Cue-target onset asynchrony (CTOA) was 400 ms. The target appeared onscreen until response or for 1200 ms. The experiment consisted of a total of 750 trials.

There were two groups: feedback and no-feedback. For both groups, a trial abruptly terminated if an eye movement was detected at any time after drift correction, but before response execution. For the feedback group, such termination was followed by the following feedback presented on screen for 1000ms: “Your eyes moved. Please do not do that". No feedback was presented after untoward eye movements to the no-feedback group.

4.4 RESULTS

4.4.1 Data Pre-Processing

The data were prepared for analysis in several steps. Firstly the number of eye movements was assessed and any individual with more than 40% eye movements more than 1° from fixation after trial start and prior to responding were excluded from further analysis (N = 8; 6 in No Feedback Condition). The RTs and accuracies of the remaining 40 participants were assessed after removing the remaining 17.4% of trials that were eye movements. The feedback group averaged eye movements on 14.1% of trials, whereas the no feedback group showed on average 20.8%, for a difference of 6.7%, CI95% = [5.8, 7.6]. An initial assessment of RT showed an average decrease over the first 5 trials of about 36 ms / trial, whereas trials outside those first 5 decreased at about 0.06 ms / trial (assessed through a regression). Therefore, those first 5 trials were excluded as practice. Finally, there were a very small proportion of trials excluded for which there was no response at all (0.002). These exclusions resulted in leaving between 396 and 730 trials per participant (M = 604) for analysis.

Some of the remaining responses were deemed anticipatory. The anticipatory criterion was determined by looking at how accuracy changed across RT. Mean accuracy was calculated for 50ms RT bins. RTs below 350 ms were deemed anticipations as this was the first bin at which mean accuracy indicated performance was stable (75.5% vs. 51.9% at preceding bin; 87.5% in next bin). The same method was used to determine the upper cutoff of the RT distribution. At longer RTs there was a drop in accuracy at 1050 ms (92.9% vs. 99.1% at preceding bin) suggesting that responses at that time and later were not to the onset of the target. These cutoffs removed 2.6% of trials.

4.4.2 Analysis

The data were analyzed with mixed effects models, linear analysis of correct trials for response speed and logistic regression for accuracy. Again, the rate of responding per second was analyzed rather than the typical RT in ms, and back transformations to RT were made for figures and RT summary values. Regression coefficients are negative rates in order to keep direction of effects the same as RT in ms.

The effects of Cue Array (Single; Multiple Cued; Multiple Uncued), Feedback (Feedback; No Feedback), and CoG (Continuous) were included in the model as predictors. The most complex models were run first, with AICs computed via the drop1 method in the {stats} package (Table 4.1). Effect sizes for parameter estimates are reported as bootstrapped 95% confidence intervals, generated via confint. A summary of overall results follows the model reporting.

4.4.3 Rate of Responding

	Model	Rate of Responding	Accuracy
Experiment 1	**Three-Way**	6285	9132
	Cue Array x Feedback x CoG	-4	-4
	Two-Way	6281	9128
	Cue Array x CoG	6	-4
	Cue Array x Feedback	5	-2
	CoG x Feedback	-2	-2
	Main Effect	6290	9120
	Cue Array	21	0
	CoG	-2	-2
	Feedback	-2	6
Experiment 2	**Three-Way**	27230	7868
	Cue Array x Feedback x CoG	-3	-4
	Two-Way	27227	7864
	Cue Array x CoG	-2	-4
	Cue Array x Feedback	-3	-4
	CoG x Feedback	-1	-2
	Main Effect	27222	7855
	Cue Array	-3	-3
	CoG	2	-2
	Feedback	-2	-2

Table 4.1: Table of AIC scores for model comparison for both dependent variables across both experiments. The row describing the model level (three-way, two-way, main effect) represents the AIC score with all terms included. The change in AIC is represented for each individual term at that level. Negative change represents improved model performance when the term is dropped, while positive change represents poorer model performance when the term is dropped.

There was no evidence to support the three-way interaction (Figure 4.2) between Cue Array, Feedback, and CoG, (Multi Cued - Multi Uncued) x Feedback x CoG, b = 0.000, CI95% = [-0.008, 0.009], (Multi Cued - Single) x Feedback x CoG, b = 0.000,

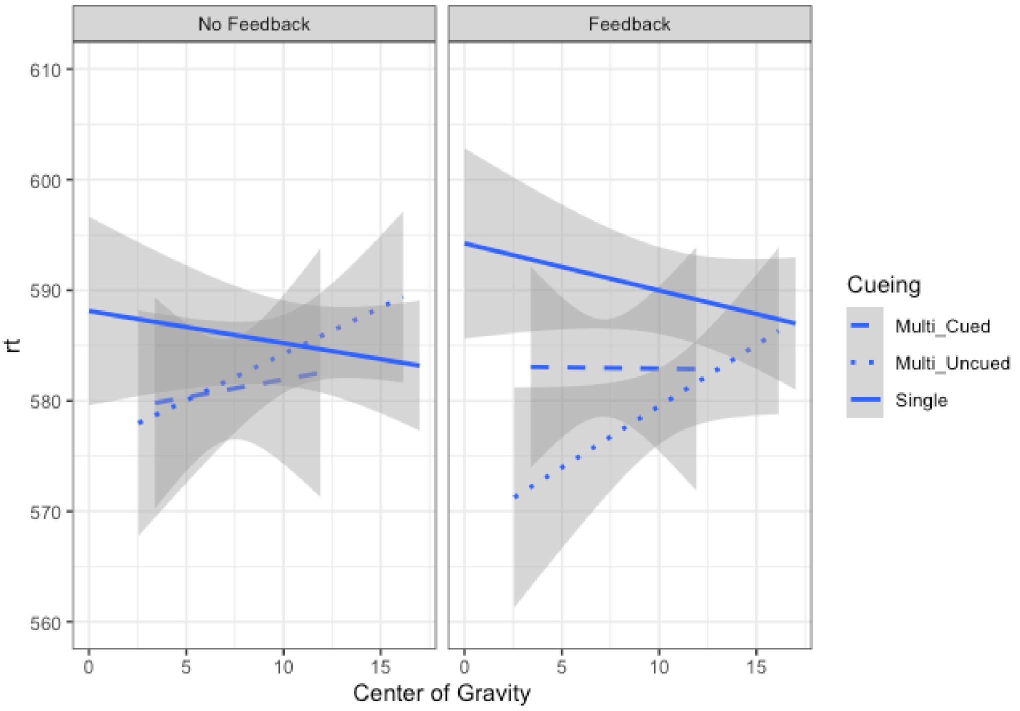

Figure 4.2: The effect of the centre of gravity of the cueing array on reaction time as a function of the cue array (solid line = Single Cue; dashed line = Multiple Cues + Cued Target; dotted line = Multiple Cues + Uncued Target). The group that did not receive feedback about untoward eye movements is depicted on the left, whereas the group that received feedback is depicted on the right. Error bars represent 95% confidence intervals.

CI95% = [-0.007, 0.008], with greater support for the model with the interaction term dropped (AIC = 6281) than when the term was included (AIC = 6285).

To evaluate the two-way interactions, we contrasted the model with all two-way interaction terms included (AIC = 6281) with models where each term was dropped. The model performed worse (△AIC = +6) when dropping the two-way interaction term between Cue Array and CoG, (Multi Cued - Multi Uncued) x CoG, b = -0.003, CI95% = [-0.007, 0.001], (Multi Cued - Single) x CoG, b = 0.001, CI95% = [-0.002, 0.005]. The model performed worse (△AIC = +5) when dropping the two-way interaction term between Cue Array and Feedback, (Multi Cued - Multi Uncued) x Feedback, b = 0.019, CI95% = [-0.001, 0.038], (Multi Cued - Single) x Feedback, b = -0.006, CI95% = [-0.027, 0.013]. The model performed better (△AIC = -2) when dropping the two-way interaction term between Feedback x CoG, b = -0.001, CI95% = [-0.002, 0.002].

To evaluate the main effects, we contrasted the model with all main effect terms included (AIC = 6290) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the main effect of CoG, b = 0.000, CI95% = [-0.001, 0.001]. The model performed worse (△AIC = +21) when dropping the main effect of Cue Array, (Multi Cued - Multi Uncued), b = 0.000, CI95% = [-0.010, 0.011], (Multi Cued - Single), b = -0.022, CI95% = [-0.031, -0.010]. The model performed better (△AIC = -2) when dropping the main effect of Feedback, b = -0.004, CI95% = [-0.136, 0.123].

4.4.4 Accuracy

There was no evidence to support the three-way interaction (Figure 4.3) between Cue Array, Feedback, and CoG, (Multi Cued - Multi Uncued) x Feedback x CoG, b = 0.036, CI95% = [-0.055, 0.096], (Multi Cued - Single) x Feedback x CoG, b = 0.037, CI95% = [-0.043, 0.095], with greater support for the model with the interaction term dropped (AIC = 9128) than when the term was included (AIC = 9132).

To evaluate the two-way interactions, we contrasted the model with all two-way interaction terms included (AIC = 9128) with models where each term was dropped. The

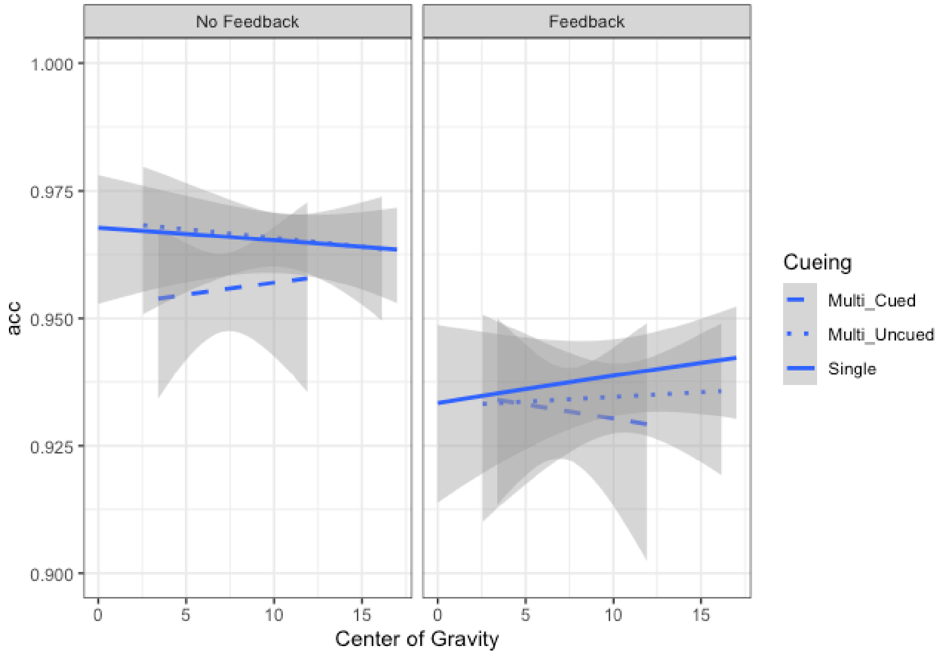

Figure 4.3: The effect of the centre of gravity of the cueing array on accuracy as a function of the cue array (solid line = Single Cue; dashed line = Multiple Cues + Cued Target; dotted line = Multiple Cues + Uncued Target). The group that did not receive feedback about untoward eye movements is depicted on the left, whereas the group that received feedback is depicted on the right. Error bars represent 95% confidence intervals.

model performed better (△AIC = -4) when dropping the two-way interaction term between Cue Array and CoG, (Multi Cued - Multi Uncued) x CoG, b = 0.000, CI95% = [-0.065, 0.067], (Multi Cued - Single) x CoG, b = 0.004, CI95% = [-0.045, 0.063]. The model performed better (△AIC = -2) when dropping the two-way interaction term between Cue Array and Feedback, (Multi Cued - Multi Uncued) x Feedback, b = -0.209, CI95% = [-0.601, 0.122], (Multi Cued - Single) x Feedback, b = -0.141, CI95% = [-0.490, 0.195]. The model performed better (△AIC = -2) when dropping the two-way interaction term between Feedback x CoG, b = 0.010, CI95% = [-0.026, 0.042].

To evaluate the main effects, we contrasted the model with all main effect terms included (AIC = 9120) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the main effect of CoG, b = 0.001, CI95% = [TBD, TBD]. The model performed equivalently well (△AIC = 0) when dropping the main effect of Cue Array, (Multi Cued - Multi Uncued), b = 0.130, CI95% = [-0.025, 0.303], (Multi Cued - Single), b = 0.174, CI95% = [-0.011, 0.342]. The model performed worse (△AIC = +6) when dropping the main effect of Feedback, b = -0.614, CI95% = [-1.065, -0.235].

For more nuanced interpretation, we treat the two (between-subjects) Feedback conditions separately for consideration relative to our theoretical predictions in Appendix B.

4.4.5 Summary of Results

Here we found evidence suggesting the effect of CoG was contingent upon the number of cues in the cueing array: the CoG effect appears to slow the rate of responding when only a single cue was presented, and appears to facilitate the rate of responding when multiple cues are presented. Rate of responding was facilitated in general when multiple cues were presented relative to only a single cue. The facilitatory CoG effects of multiple cues and of the type of array on rate of responding were larger when feedback was presented, and feedback created a cost in overall accuracy.

When feedback on eye movements was presented, the reversal of CoG effect in the omnibus model occurring as a consequence of the cueing array persisted. There was a large facilitatory effect on rate of responding when multiple cues were presented. When no feedback on eye movements was presented, the reversal of CoG effect in the omnibus model as a consequence of the cueing array also persisted, as did the facilitatory main effect of multiple cues. There were no effects or interactions on accuracy of responding for either group.

4.5 DISCUSSION

Numerous intriguing patterns were uncovered when testing the effect of multiple element cueing arrays on performance in a 2AFC task. Trials with multiple cues were responded to faster than trials with only a single cue, suggesting the increased stimulation was also influencing temporal attention (Lawrence & Klein, 2013; McCormick et al., 2018). The Feedback group showed no effect on accuracy on single cue trials, but the No Feedback group showed the predicted SAT on single-cue trials (see Appendix B). However, the CoG effects did not manifest as predicted. In previous studies using the multiple element cueing array requiring manual responses (Klein, Christie, & Morris, 2005; Langley et al., 2011), there was an inhibitory effect of the CoG. Here we find evidence that the CoG is priming responses, with the possibility of an inhibitory local effect cancelling out the priming effect in the Feedback group since the function representing cued targets following multiple cues appears to show no CoG effect relative to all other multiple cue functions (Figure 4.2). The lack of an inhibitory CoG effect — and instead finding evidence for priming — when discrimination responses are required is noteworthy: if the CoG effects are in fact a result of oculomotor activation in the superior colliculus (as proposed by Klein et al., 2005, and reinforced by Christie et al., 2013), then these results suggest that activity in the superior colliculus is susceptible to top-down, task-based influences (vis, stimulus-response demands), and is not solely responding to sensory inputs.

Although noteworthy, it is the purpose of this investigation to examine the performance consequences of inhibitory CoG effects. In the following experiment, we seek to address this question by modifying the time course of the stimuli in order to generate these effects.

4.6 EXPERIMENT 2

As first reported by Lupianez et al. (1997), the time course of IOR effects are highly contingent upon task demands. They found reliable IOR effects for detection responses as early as 400ms post-cue onset — akin to the time course employed by previous CoG studies. However, to measure IOR in a discrimination task, an SOA of at least 700ms was required. Indeed, in a comparison of meta-analyses, Habibnezhad, Lawrence, and Klein (2019) showed the cross-over point from cue-elicited facilitation to an inhibitory cueing effect could be nearly 500ms later for discrimination tasks relative to detection tasks. To this end, we will use the same stimuli from Experiment 1, but we will employ an SOA shown to be long enough to reliably induce IOR in a discrimination task.

4.7 METHODS

4.7.1 Participants

Forty-six students (39 females and 7 males) from Dalhousie university volunteered to participate in this study. Their ages ranged from 18 to 23 years (M = 19.6), and four of them were left-handed. They were compensated one point of course credit. The research was approved by the relevant Dalhousie Research Ethics Board.

4.7.2 Procedure and Design

All aspects of this experiment were precisely the same as Experiment 1, except the cue-target onset asynchrony. Whereas in Experiment 1 the target array was presented simultaneously with the offset of the central cue-back stimulus (400ms CTOA), here a 400ms wait period was instantiated between central cue-back offset and the onset of the target array, resulting in an 800ms CTOA.

4.8 RESULTS

The same analytic approach that was used to examine the data in Experiment 1 was used here.

4.8.1 Data Pre-Processing

The data were prepared for analysis in several steps. Firstly the number of eye movements was assessed and any individual with more than 40% eye movements more than 1° from fixation after trial start and prior to responding were excluded from further analysis (N = 10; 8 in No Feedback Condition). The RTs and accuracies of the remaining 36 participants were assessed after removing the remaining 20.8% of trials that were eye movements. The feedback group averaged eye movements on 18.0% of trials, whereas the no feedback group showed on average 24.7%, for a difference of 6.7%, CI95% = [5.7, 7.7]. An initial assessment of RT showed an average decrease over the first 5 trials of about 43 ms / trial, whereas trials outside those first 5 decreased at about 0.06 ms / trial (assessed through a regression). Therefore, those first 5 trials were excluded as practice. Finally, there were a very small proportion of trials excluded for which there was no response at all (0.004). These exclusions resulted in one participant being excluded for too few trials (279), leaving between 429 and 709 trials per participant (M = 563) for analysis.

Some of the remaining responses were deemed anticipatory. The anticipatory criterion was determined by looking at how accuracy changed across RT. Mean accuracy was calculated for 50ms RT bins. RTs below 250 ms were deemed anticipations as this was the first bin at which mean accuracy indicated performance was stable (85.4% vs. 65.5% at preceding bin; 91.7% in next bin). The same method was used to determine the upper cutoff of the RT distribution. At longer RTs there was a drop in accuracy at 1050 ms (87.5% vs. 98.6% at preceding bin) suggesting that responses at that time and later were not to the onset of the target. These cutoffs removed 3.0% of trials.

4.8.2 Analysis

To examine the main hypotheses, the effects of Cue Array (Single; Multiple Cued; Multiple Uncued), Feedback (Feedback; No Feedback), and CoG (Continuous) were included in the model as predictors. The most complex models were run first, with AICs computed via the drop1 method in the {stats} package (Table 4.1). Effect sizes for parameter estimates are reported as bootstrapped 95% confidence intervals, generated via confint. A summary of overall results follows the model reporting.

4.8.3 Rate of Responding

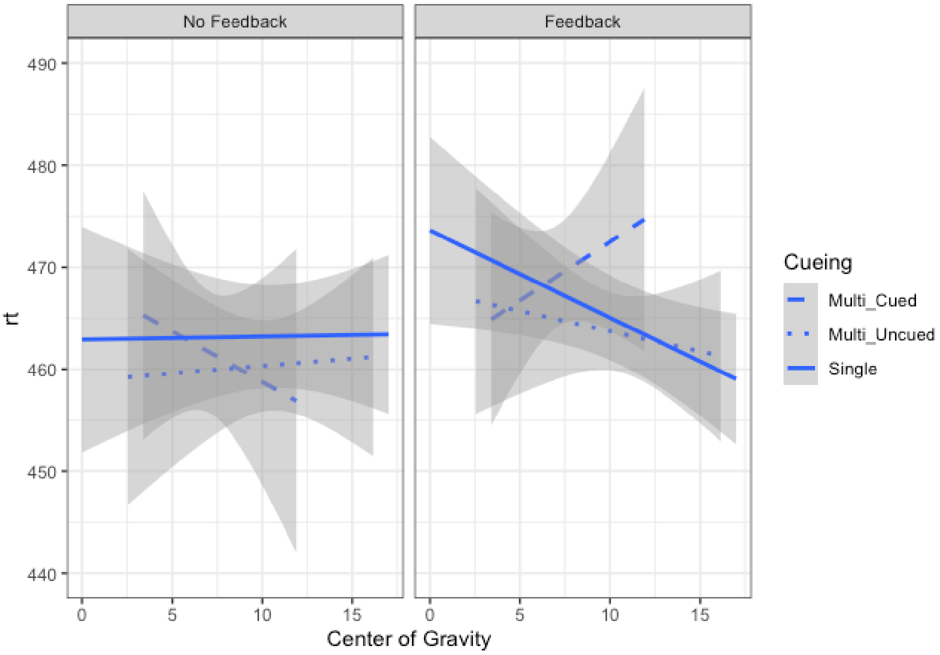

Figure 4.4: The effect of the centre of gravity of the cueing array on reaction time as a function of the cue array (solid line = Single Cue; dashed line = Multiple Cues + Cued Target; dotted line = Multiple Cues + Uncued Target). The group that did not receive feedback about untoward eye movements is depicted on the left, whereas the group that received feedback is depicted on the right. Error bars represent 95% confidence intervals.

There was no evidence to support the three-way interaction (Figure 4.4) between Cue Array, Feedback, and CoG, (Multi Cued - Multi Uncued) x Feedback x CoG, b = 0.007, CI95% = [-0.008, 0.023], (Multi Cued - Single) x Feedback x CoG, b = 0.010, CI95% = [-0.005, 0.024], with greater support for the model with the interaction term dropped (AIC = 27227) than when the term was included (AIC = 27230).

To evaluate the two-way interactions, we contrasted the model with all two-way interaction terms included (AIC = 27227) with models where each term was dropped. The model performed better (ΔAIC = -2) when dropping the two-way interaction term between Cue Array and CoG, (Multi Cued - Multi Uncued) x CoG, b = -0.001, CI95% = [-0.009, 0.008], (Multi Cued - Single) x CoG, b = 0.002, CI95% = [-0.006, 0.009]. The model performed better (△AIC = -3) when dropping the two-way interaction term between Cue Array and Feedback, (Multi Cued - Multi Uncued) x Feedback, b = 0.014, CI95% = [-0.024, 0.052], (Multi Cued - Single) x Feedback, b = 0.021, CI95% = [-0.020, 0.060]. The model performed better (△AIC = -1) when dropping the two-way interaction term between Feedback x CoG, b = 0.002, CI95% = [-0.002, 0.006].

To evaluate the main effects, we contrasted the model with all main effect terms included (AIC = 27222) with models where each term was dropped. The model performed worse (△AIC = +2) when dropping the main effect of CoG, b = 0.002, CI95% = [0.000, 0.004]. The model performed better (△AIC = -3) when dropping the main effect of Cue Array, (Multi Cued - Multi Uncued), b = 0.009, CI95% = [-0.001, 0.028], (Multi Cued - Single), b = 0.002, CI95% = [-0.002, 0.020]. The model performed better (△AIC = -2) when dropping the main effect of Feedback, b = -0.016, CI95% = [-0.187, 0.177].

4.8.4 Accuracy

There was no evidence to support the three-way interaction (Figure 4.5) between Cue Array, Feedback, and CoG, (Multi Cued - Multi Uncued) x Feedback x CoG, b = -0.010, CI95% = [-0.088, 0.067], (Multi Cued - Single) x Feedback x CoG, b = -0.016, CI95% = [-0.092, 0.051], with greater support for the model with the interaction term dropped (AIC = 7864) than when the term was included (AIC = 7868).

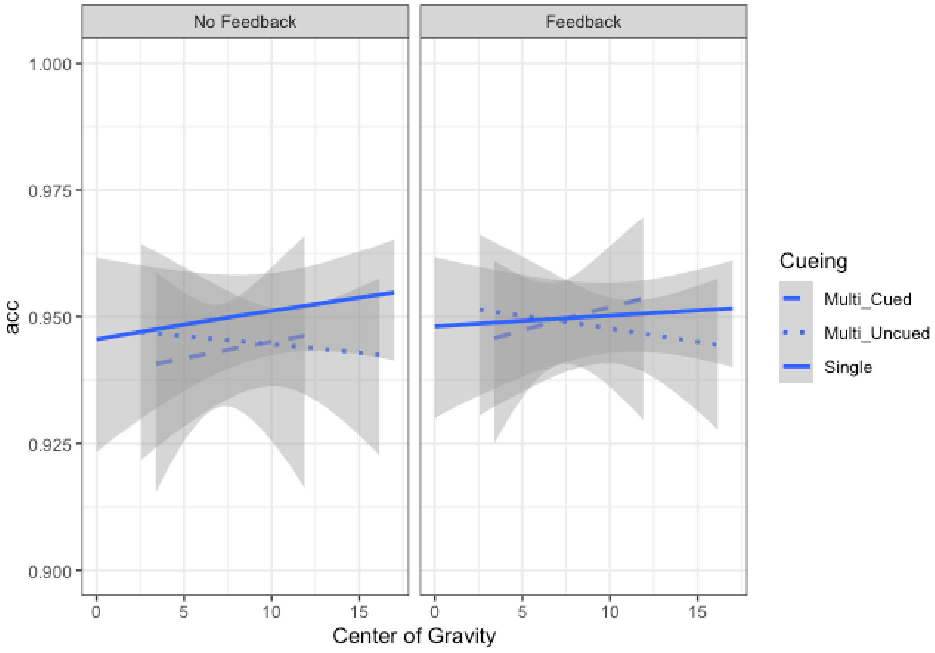

Figure 4.5: The effect of the centre of gravity of the cueing array on accuracy as a function of the cue array (solid line = Single Cue; dashed line = Multiple Cues + Cued Target; dotted line = Multiple Cues + Uncued Target). The group that did not receive feedback about untoward eye movements is depicted on the left, whereas the group that received feedback is depicted on the right. Error bars represent 95% confidence intervals.

To evaluate the two-way interactions, we contrasted the model with all two-way interaction terms included (AIC = 7864) with models where each term was dropped. The model performed better (△AIC = -4) when dropping the two-way interaction term between Cue Array and CoG, (Multi Cued - Multi Uncued) x CoG, b = -0.018, CI95% = [-0.073, 0.033], (Multi Cued - Single) x CoG, b = -0.004, CI95% = [-0.063, 0.049]. The model performed better (△AIC = -4) when dropping the two-way interaction term between Cue Array and Feedback, (Multi Cued - Multi Uncued) x Feedback, b = -0.058, CI95% = [-0.426, 0.311], (Multi Cued - Single) x Feedback, b = -0.119, CI95% =

[-0.578, 0.238]. The model performed better (△AIC = -2) when dropping the two-way interaction term between Feedback x CoG, b = -0.004, CI95% = [-0.037, 0.035].

To evaluate the main effects, we contrasted the model with all main effect terms included (AIC = 7855) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the main effect of CoG, b = 0.003, CI95% = [-0.014, 0.021]. The model performed better (△AIC = -3) when dropping the main effect of Cue Array, (Multi Cued - Multi Uncued), b = -0.028, CI95% = [-0.174, 0.163], (Multi Cued - Single), b = 0.064, CI95% = [-0.094, 0.267]. The model performed better (△AIC = -2) when dropping the main effect of Feedback, b = -0.107, CI95% = [-0.2233, 0.503].

For more nuanced interpretation, we treat the two (between-subjects) Feedback conditions separately for consideration relative to our theoretical predictions in Appendix B.

4.8.5 Summary of Results

Here we found evidence suggesting a main effect of CoG operating to slow the rate of responding. This main effect did not interact with any other predictors, but only existed for the Feedback group when the two groups were modelled separately (see Appendix B). When testing the groups separately, the main effect of cueing was only supported for the Feedback group (see Appendix B). In fact, there was no evidence to support any main effects or interactions on rate of responding when specifically examining the No Feedback group.

When examining the full scope of predictors on accuracy no main effects or interactions were supported, nor were any main effects or interactions supported when examining the two groups separately.

4.9 DISCUSSION

Here we found evidence that the CoG showed the predicted inhibitory effect on RT. However, when testing each Feedback group separately, this main effect existed for the Feedback group only. Although there was no statistical support for the interaction between the cueing array and the CoG, it appears from Figure 4.4 that the nature of the CoG effect for the input form may be contingent upon whether the target appears at a locally stimulated location. However, this graphical interpretation must be taken with caution as the Cued Multiple Element function is the most variable given the width of the CI95, so this pattern may in fact be spurious. Regardless, the resulting pattern for this group suggest that the CoG is inhibiting RT.

Disappointingly, contrary to our primary motivation for utilizing a discrimination task to assess the influences of these effects on accuracy, no effects on accuracy were found. That we didn’t see convincing evidence for IOR at all in the No Feedback group casts doubt on the degree or extent of comparative conclusions that can be rendered from the present investigation. Either of two possibilities likely lead this outcome. It is possible that the two forms of IOR have different time courses, and the 800ms SOA was long enough to measure the CoG when inducing the input form, but too short to reliably detect the CoG for the output form. Another aspect that may have lead to this result is that trials terminated following an untoward eye movement regardless of whether feedback was provided. For those observers randomly assigned to the No Feedback condition, perhaps the conditions were such that observers were at a times suppressing the reflexive tendency to make eye movements even though no feedback was provided to this end. This would result in a conflation of effects, wherein either between-subjects or from trial-to-trial both the input and output form of IOR were operating, akin to what was posited regarding the interaction between the Simon effect and the two forms of IOR in Chapter 2, Section 2.4. If this conflation increased variability in the measure or even changed the direction of the effect for any of the factors we were manipulating—from trial to trial or across subjects—as was specifically predicted for accuracy, then it would be no surprise that none of the predictors significantly influenced behaviour.

4.10 CONCLUSIONS

In Experiment 1, the CoG showed evidence of priming at an SOA that had previously reliably shown inhibitory CoG effects (Klein, Christie, & Morris, 2005; Christie, Hilchey, & Klein, 2013). When the SOA was extended to a duration that had previously demonstrated to be long enough to observe IOR in a discrimination task (Lupianez et al., 1997), we found inhibitory CoG effects under conditions that ought to have elicited the input form of IOR, but no effects were reliably detected when we predicted we would be inducing the output form. Across both experiments, the pattern of slower responses to targets appearing at previously cued locations was observed, although this was only statistically supported when the SOA was sufficiently long in Experiment 2 (see Appendix B).

That there were no effects or interactions involving the CoG on accuracy in either experiment suggests three possibilities. The first is that the CoG effects are only manifest in speed of responding. The priming in Experiment 1 and the inhibition in Experiment 2 induce genuine changes in performance since neither RT effect is accompanied with a tradeoff in accuracy (Chapter 1: Figure 1.5 — upper dashed arrow). This would suggest the effect is operating to either facilitate (Experiment 1) or delay (Experiment 2 - Feedback group) the orienting of attention in the direction of the CoG, as any post-perceptual effect on the threshold of responding would likely be associated with a concomitant trade in accuracy. The second possibility is that observers were at the asymptote of performance, and the CoG was simply shifting them along the horizontal span at the tail end of the function shown in Figure 1.5 (Chapter 1). This result would be ambiguous relative to our initial hypothesis regarding the effect of the CoG in speed-accuracy space, but wouldn't rule out a post-perceptual response-based account. This is because either a criterion shift or a genuine change in performance appear identical at the asymptote of performance. This is not implausible as accuracies were very high in both experiments (~95%+ overall). The third possibility is not unrelated to the previous one: since accuracy was so high, it is possible observers committed too few errors for us to detect any reliable effects on accuracy. Observers completed an average of 604 trials in Experiment 1, and 563 in Experiment 2. Since accuracy was nearly 95% overall, that would result in only 30 error trials per participant on average in Experiment 1 (604 x 0.05), and only 28 error trials per participant on average in Experiment 2 (563 x 0.05). Presuming an even distribution of errors across cue array conditions, this count would afford us on average only 10 trials per cue array condition in Experiment 1 and 9 trials per cue array condition in Experiment 2 to fit our hierarchical models. Future work ought to employ a target that is more difficult to discriminate in order to increase the number of errors generated by participants. This would address the hypothesis that there were too few errors to generate reliable estimates, and likely resolve the discrepancy between the first two possible explanations for the accuracy patterns.

Together, the present study presents numerous important findings pertaining to the nature of CoG effects. Increasing target response demands (i.e., discrimination vs. detection) delays the time course of the effect, resulting in priming effects when inhibitory effects are commonly observed. Using a target that affords measurement of accuracy shows that CoG effects (may) influence primarily the speed of responding, suggesting the CoG is operating on orienting mechanisms rather than the threshold of responding. Moreover, when reliably generating the input form of IOR, the inhibitory effect of the CoG may interact with local effects, but more work is needed to further evaluate this relationship. Lastly, counter to Hilchey, Hashish, et al. (2014), there is little evidence that the output form of IOR is generated in the absence of explicit eye movement feedback (vis the No Feedback condition), suggesting observers may learn implicitly to suppress the tendency to produce untoward eye movements when trials are aborted to this end.

CHAPTER 5: GENERAL DISCUSSION

In addition to previously unpublished material, excerpts from this manuscript are presented below. Co-authors for this manuscript are Dr. W. Joseph MacInnes and Dr. Raymond M. Klein, respectively. In this manuscript, my contributions include, but aren't limited to, writing, theory development, data collection, data analysis, and project management.

Redden, R. S., MacInnes, W. J., & Klein, R. M. (2021). Inhibition of return: An information processing theory of its natures and significance. Cortex, 135, 30-48. This manuscript does not exactly replicate the final version published in Cortex. It is not a copy of the original published article and is not suitable for citation as such.

5.1 SUMMARY OF DISSERTATION

In an overview of the literature on inhibition of return, we argue that there is evidence that two distinct inhibitory orienting effects have been conflated by the literature. These two forms of IOR operate at different stages of the information processing continuum: an input form operating to inhibit the quality of information arising from a previously cued location, and an output form operating to bias responding in a previously cued direction. Importantly, the form of IOR that is manifest is contingent upon the activation state of the reflexive oculomotor system (ROMS): the input form arising when the ROMS is suppressed, and the output form arising when the ROMS is not suppressed.

We characterized three diagnostics (Chapter 1, Section 1.3) that have been used in the literature which provide converging empirical evidence for our theoretical position: central arrow targets, the locus of slack logic as it pertains to the Psychological Refractory Period (PRP), and consideration of performance in speed-accuracy space. The evidence across the literature using these diagnostics strongly support our two forms framework, which will be briefly summarized next.

When the ROMS is suppressed, central arrow targets tend to show no effect of cueing, suggesting the inhibitory effect exists at the stimulated location. In a PRP task, an inhibitory cueing effect on Task 1 is absorbed into the central bottleneck when two tasks are sufficiently close in time, suggesting a perceptual locus of the effect. When considering the effect of cueing on both speed and accuracy in tandem, the effect tends to result in a genuine cost in performance whereby cued targets are responded to more slowly and less accurately, suggesting the effect is operating to impair target processing.

When the ROMS is not suppressed, central arrow targets tend to show an effect of cueing equivalent to that of intermixed peripheral targets, suggesting the inhibitory effect is biasing responses in the direction compatible with the cued location. In a PRP task, an inhibitory cueing effect on Task 1 delays responding to Task 2 regardless of the temporal delay between the two tasks, suggesting a post-bottleneck (i.e., decisional, or response execution) locus of the effect. When considering the effect of cueing on both speed and accuracy in tandem, the effect tends to result in a speed-accuracy tradeoff whereby cued targets are responded to more slowly but more accurately, suggesting the effect is operating post-perceptually.

Across three empirical chapters in this dissertation, evidence is presented to support and advance this theoretical framework, and contextualized with respect to a computational process model (see Section 5.2, and Appendix A). In Chapter 2, two experiments explore the output form explicitly. In Experiment 1, we intermixed two diagnostics outlined above: central arrow targets requiring a localization response, and peripheral targets requiring a discrimination response (affording measurement of speed and accuracy). Here we show that when IOR is caused by a reflexive eye movement (vis, a prosaccade), both diagnostics are sensitive to the effect — slower localization responses to arrows pointing in the direction compatible with the prosaccade, and an SAT for cued discrimination targets. In Experiment 2, we resolved an ambiguity in the criterion-shift account for the effect, wherein slower but more accurate responses could exist in SAT space on either a single performance function or manifest as a shift from one function to a less efficient one. Using the response-signal method and again causing IOR with a prosaccde, we reinforce the proposal that the output form of IOR is indeed operating at a post-perceptual stage of processing, as performance for cued and uncued targets exist on a single performance function.

In Chapter 3, across two experiments we sought to compare and contrast the two forms of IOR in tasks primarily perceptual in nature. Temporal order judgements required an observer to discriminate the relative onsets of two targets, and in our design afforded insight on the rate of perceptual information accrual (PSS effect in bilateral TOJs) and sensitivity to temporal information (JND effect in unilateral TOJs). Probes were presented instead of TOJ stimuli on a third of trials, and required a speeded detection response followed by an unspeeded colour matching task. Detection responses afforded measurement of the canonical RT effect, whereas applying mixture models to performance in the colour matching task afforded insight into specific perceptual consequences of attentional effects: the probability that a stimulus was encoded, and given it was encoded successfully, the fidelity of that encoding. Again reinforcing the post-perceptual nature of the output form, in the experiment wherein IOR was caused by a prosaccade to a peripheral cue, no inhibitory effects were observed on any of the perceptual tasks. However, in the experiment wherein input IOR was caused by an antisaccade relative to a peripheral cue, we demonstrate for the first time inhibited bilateral TOJs (but not unilateral TOJs) and impaired colour matching responses as a consequence of cueing, suggesting the input form of IOR is operating to delay the accrual of perceptual information from the previously stimulated location. The finding of no significant RT effects for either form was contextualized relative to other reports of no IOR for rare targets, where we suggest switching mental sets between complex tasks may delay processing or responding beyond the time course of any IOR effect.

In Chapter 4, we sought to determine whether the effect at the centre of gravity of multiple cues is of the input- or output form. Previous work found that the effect of IOR for both manual and saccadic responses following multiple cues may be operating at the centre of gravity of the grouped array rather than the local elements, suggesting that IOR is due to population coding in the oculomotor pathways (e.g., the superior colliculus). Across two experiments, we probed responding with a discrimination target to measure speed and accuracy (Chapter 1, Section 1.3.3) following single- and multiple-item cue arrays. In Experiment 1, we found priming of responses as a result of the CoG at a cue-target onset interval that has previously shown inhibition, suggesting that top-down, task-based contexts can modulate the nature of these effects heretofore presumed to be reflexive in nature. In Experiment 2, we found weak evidence for inhibitory CoG effects when eye movements were expressly forbidden, but no convincing evidence for IOR at all when eye movement feedback was not provided. We conclude that since no reliable effects on accuracy were found, CoG effects (may) influence primarily the speed of responding. This finding suggests the CoG is operating on orienting mechanisms rather than the threshold of responding, but the lack of convincing evidence that the output form of IOR was generated at all suggests more work is needed to further evaluate whether these patterns apply to only the input form, or to both forms.

In addition to these empirical findings, we present a computational process model that accounts for empirical dissociations in discrimination task performance, and advances the theoretical framework. The implications of this model are discussed next.

5.2 DRIFT DIFFUSION MODEL

We have presented a diffusion model (AppendixA) with excellent fit to human error and reaction time data in a choice reaction task. Further, supporting our hypothesis, we have demonstrated that the two forms of IOR are clearly best fit by adjusting different parameters from the baseline response model. Diffusion parameters are often thought to reflect different underlying cognitive effects (Smith and Ratcliff, 2004). We, therefore, propose that this adds to the growing evidence for the two forms—input and output—of IOR. As expected, IOR generated with an antisaccade was best modelled by a drift rate decrease reflecting a pure performance decrement. IOR generated by a prosaccade, however, was equally fit by two parameter shifts, though it should be noted that both of these possibilities are quite different from the drift rate result of the antisaccade condition. The hypothesized increase in the threshold parameter did match human results as predicted, but so too did a decrease in the trial noise parameter. Although there was little difference in these parameters in simulated results, we believe the bulk of the evidence, as we describe next, currently supports the threshold parameter as the cause of the speed accuracy tradeoff of the output form of IOR.

Whereas a reduction in trial noise provided an equally good fit to these SAT data as the threshold parameter, we believe the trial noise cannot explain the various output effects observed using the other diagnostics: central arrow targets and the PRP method. A reduction in trial noise can be considered an improvement in signal-to-noise ratio (SNR) as a consequence of signal strength remaining unaffected while noise is reduced. This effect results in an SAT in speed-accuracy space since both noise and signal contribute to RT (vis slower RT with equivalent “signal” but reduced “noise”), and reduced noise elicits fewer errors. It is unknown to us how such an influence on the decision process (reduced SNR) could also account for inhibited RT to central arrows presented at fixation (i.e., a target sharing no stimulus overlap with the cue). Nor is it clear how a reduction in SNR as a consequence of the cue could account for the additive-with-TTOA pattern observed in the PRP design when output IOR is generated, thought to be a result of prolonged queueing of T2 response demands. However, the threshold parameter account neatly converges with the patterns observed in these other paradigms: an increase in response threshold in the direction compatible with the cue would also inhibit localization responses to central arrows[7], and produce the additive-with- TTOA pattern in the PRP task.

The information processing theory we are proposing for the two forms of IOR is represented in Fig. 5.1. Here we have embedded, into an information flow diagram adapted from Attneave (1959), illustrations of the diffusion parameters we assume to be

---

[7] A pattern observed directly by Redden et al. (2020) and the present Chapter 2 Experiment 1, when intermixing central arrow and peripheral discrimination targets following a prosaccade.

responsible for the ubiquitous delays in reaction time that are the hallmark of IOR. As reviewed above, the reduced drift rate associated with the input form will result in slower and often less accurate responding whereas the increase in threshold associated with the output form will result in slower but more accurate responding (a speed-accuracy tradeoff).

Computational models can be seen as mathematical instantiations of theories and have the advantage of forcing us to be specific about theory components that are often missing in verbal theories (Broadbent, 1987; Guest & Martin, 2020). Generative models in particular are well suited since they are able to recreate (vis simulate) distributions of responses that are similar to those produced by human participants. The diffusion model, for example (see Appendix A), was able to match human responses in distribution of correct responses, distributions of incorrect responses, percentage of errors, the overall vector in Speed-Accuracy trade-off space and (to a degree) the response variability. Relative to other computational models of IOR, ours is the only accumulation approach to focus on error rates in addition to reaction time distributions. Whereas RT distribution shape is an important factor in determining parameter contribution (Ludwig, 2009), the rate and latencies of errors on 2AFC tasks is often a defining feature of key theories.

5.3 FUTURE WORK

Perhaps the most interesting aspect of a good model is its ability to suggest new hypotheses for testing (van Rooij & Baggio, 2020). In this regard, whereas our model results converge theoretically and empirically with a substantial amount of IOR research, the findings are not precisely consistent with modelling work on neurophysiological phenomena. Reviews of neural models of speed-accuracy tradeoffs (i.e., when the observer opts to trade speed for accuracy) implicate a change in baseline—rather than a change in decision threshold—as the mechanism underlying the SAT in several neural substrates shown to be associated with IOR, such as the lateral intraparietal (Hanks, Kiani, & Shadlen, 2014), frontal eye fields (Heitz & Schall, 2012), and the superior colliculus (Ratcliff et al., 2003). Computationally, a reduction in baseline as a result of IOR would equate to an increased decision threshold since both result in a change in distance from starting point (z) to the response (A or O), however it is believed that there are differences in how the start and decision boundary shifts are represented neurally (Bogacz, Wagenmakers, Fortsmann & Nieuwenhuis, 2010). We chose to focus on threshold changes since our task manipulation of interest was non-informative as to the correctness of response. While a change in threshold can increase the information needed for either (correct or incorrect) response equally, the change in starting point (z) is a true bias with a reduction in distance for one response coming at the expense of an increase in the other.

Whereas we have argued in the Introduction for our two forms conceptualization over various theoretical and empirical accounts of IOR, work remains to determine the extent to which our constructs can account for other expert's frameworks for the phenomenon. In particular, Lupianez and colleagues (Lupianez, 2010; Lupianez, Martin-Arevalo, & Chica, 2013) have shown inhibition can be observed in the absence of facilitation at short SOAs, and as such describe IOR as an onset-detection cost—

manifesting as a function of similarity between cue and target stimuli, with increasing similarity eliciting greater costs—rather than a consequence of the attentional disengagement from the cue stimulus. The prosaccade/antisaccade paradigm that provides the foundation of our model and much of the present work does not afford the measurement of attentional facilitation, since the time it takes to generate the requisite eye movements to the periphery and back to fixation in this task (~300 ms) is longer than the time course of most facilitatory effects of an uninformative peripheral onset (Lupianez, Milliken, Solano, Weaver, & Tipper, 2001). As such, the degree to which either or both forms of IOR co-exist with attentional facilitation is heretofore unanswered (but see Hilchey, Klein, & Satel, 2014). Additionally, in the empirical work leading to Lupianez and colleagues' framework, a shift of attention as a result of the cue would have been implied (since there were no explicit oculomotor instructions relative to the cue), whereas in our experiments, the shifts of attention were explicit (as overt eye movements were required). Lupianez's experiments did not employ eye monitoring, so the state of the reflexive oculomotor system for observers in those studies, the importance of which to the form of IOR was emphasized by Hilchey, Hashish, et al. (2014), is ambiguous. Lastly, as described above, the onset-detection cost framework proposes IOR is manifest as a function of similarity between cue and target stimuli. In our paradigm, the same cue and target stimuli were used to elicit qualitatively different effects: a genuine performance cost, and an SAT. As such, whereas the onset-detection cost framework may or may not map on to either form of IOR (see Hilchey, Klein, & Satel, 2014, for further discussion of this point), it is unclear how this framework would completely account for the present dissociation and various other dissociations across our three diagnostics (e.g., Hilchey et al., 2016) without further conditionals to the theory[8].

Additionally, our theoretical representation of two dissociable inhibitory phenomena provides a framework to begin addressing other areas lacking consensus—the neural implementation, and the object-based representation of IOR. Whereas patient data has been informative in implicating the brain regions involved in some form of IOR, to our knowledge no patient studies have employed a non-spatial decision task to measure both RT and accuracy. Thus, the results are ambiguous as to the particular form of IOR that was generated, or impaired, in these special populations. Moreover, Martin-Arevalo, Chica, and Lupianez reviewed the literature on IOR using event-related potentials (ERPs), and found no single component reliably indexed the various behavioural effects reported across studies. However, perhaps the lack of consensus stems from the insensitivity to the two forms construct, and researchers may have been measuring both forms simultaneously—either between trials or between subjects (see, e.g., Satel, Hilchey, Wang, Story, & Klein, 2013, for an ERP study sensitive to the two forms), just as we propose as an account for literature described in Chapter 2 and results reported in Chapter 4. By applying the prosaccade/antisaccade methodological injunction to generate the two forms separately, it is conceivable that more reliable patterns may be gleaned due to a reduction in inter- or intra-subject variability resulting from the conflation of two inhibitory processes.

---

[8] This same logic can serve to critique the habituation account of IOR (Dukewich, 2009). Whereas habituation may serve as an explanatory mechanism of one of the two forms we propose, without additional parameters it cannot explain the numerous dissociations we have summarized across our three diagnostics.

This same approach could serve to address the lack of consensus on the degree to which IOR is encoded in an object-based reference. Inhibitory cueing effects appear to be most readily encoded in object-based reference under conditions akin to when the input form of IOR would have been generated: when observers are instructed to ignore the cue and generate a manual target response (Fig. 1.2 - Cell 1; Tipper et al., 1991; Tipper et al., 1994; Weaver, Lupianez & Watson, 1998; Tipper et al., 1999; Theeuwes, Mathot & Grainger, 2014; Smith et al., 2016). However, under conditions wherein the output form is thought to be generated (i.e., when saccadic responses are required; Fig. 1.2 - Cells 3-6), it is less certain whether IOR is (Abrams & Dobkin, 1994; Swalwell, 2019; Tas et al., 2012) or is not (Abrams & Pratt, 2000; Senturk et al., 2016; Redden et al., 2018) object-based. Moreover, whereas the input form is sensitive to the presence of placeholder objects in the traditional cueing paradigm, the output form is not (Hilchey, Pratt, & Christie, 2018). Regardless of which form may have been generated, object-based IOR effects tend to be smaller than those observed in more traditional cueing paradigms (e.g., Abrams & Dobkin, 1994; Tipper et al., 1991). To this end, it has been posited that IOR (at least in dynamic displays) may in fact not be encoded in an object-based frame of reference, rather that object-based effects are artefacts of a commonplace methodological flaw (Redden & Klein, 2019): "There is a potential confound implemented in many investigations employing the moving boxes paradigm—the motion (if/when it occurs) always occurs in the same direction. Even though the spatial cue is uninformative regarding subsequent target locations, the direction in which the cued object will move is typically predictable. If observers are using the cue and/or motion in

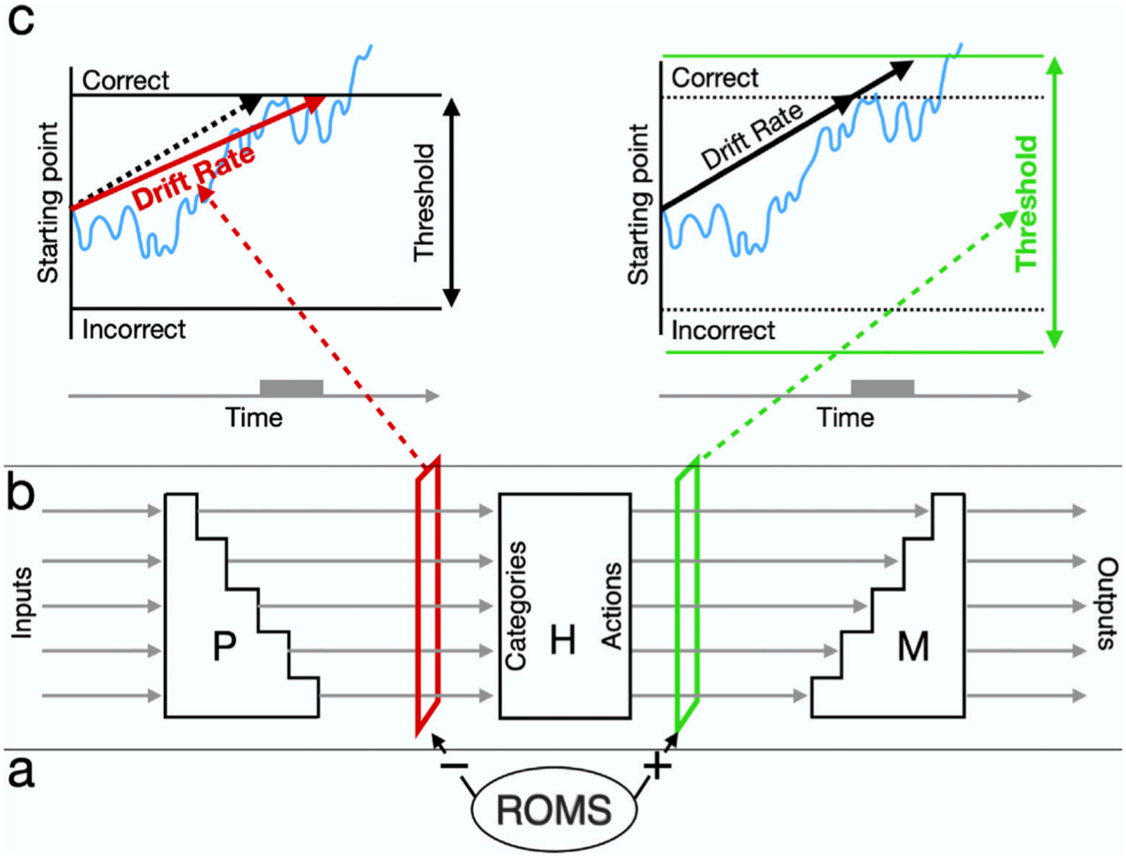

Figure 5.1 An information processing theory of the two forms of inhibition return described in the text (with the input form represented in red and the output form represented in green). a) Which form is generated depends on the state of the reflexive oculomotor system (ROMS: - indicates it is suppressed; + indicates it is not suppressed). b) A classic information flow diagram simplified from Attneave, 1959. Inputs flow through the perceptual machinery (P) and lead to perceptual experiences at varying levels of abstraction at a central processing system (H). Outputs from this system, corresponding to voluntary behavior, give rise to actions at various levels of complexity. The red and green polygons represent the locus of operation of the two forms of inhibition of return (this component of the figure is adapted from Klein & Redden, 2018). It is important to note that for present purposes we have greatly simplified Attneave's original figure which included feedforward and feedback signals, reflexes and habits that bypass (H) and the possibility for affective (evaluative) modulation. c) How the two forms of IOR affect processing as represented in a diffusion model, with the input form slowing the accumulation of evidence (decrease in drift rate) and the output form operating as a bias against responding (an increase in the threshold of evidence required to initiate a response). The rectangles above the "time" lines represent the effect upon RT generated by the two forms of IOR. The possible effects upon accuracy are not illustrated here (see text for explanation).

some temporal attentional control setting in order to prepare for the target onset, then the shift of spatial attention elicited by the cue onset may be biased (albeit slightly) toward the region of space in which the cued object would eventually terminate. This bias in the allocation of spatial attention might actually result in an inhibitory, location-based gradient ... spreading disproportionately into the location of the cued object relative to the uncued object due to its asymmetry with the cued axis, which would masquerade as an object-based effect." Explicitly generating the two forms in a variant of the moving boxes paradigm wherein the motion direction is not predictable would provide clarity about this proposal.

Future research ought to discern which of the two forms of IOR is operating in search—in particular because IOR in search is one of the more ubiquitous findings in the literature (see Wang & Klein, 2010, for a review), and is thought to empirically represent the functional significance of the phenomenon. Whereas, as suggested by Klein and Redden (2018), both forms of IOR could operate in service of search—one reducing the salience of previously attended inputs, the other biasing responses from previously attended locations—currently the evidence in the literature is ambiguous. Klein and MacInnes (1999) argued that the inhibitory tags generated during overt search are encoded in the visual representation of the scene (IOR encoded at the input level?), since there was no inhibitory effect on performance when the scene was removed; a pattern reinforced in the spatial cueing paradigm by Redden, Klages, and Klein (2017). However, Boot, McCarley, Kramer, and Peterson (2004) demonstrated that saccadic inhibition persists at a location, even when previously fixated stimuli have been removed (IOR encoded at the output level?). The relative activation level of the reflexive oculomotor system may in fact offer some insight into these different patterns. In the study conducted by Klein and MacInnes (1999), observers were instructed to freely search a complex visual scene for a target, a search strategy that may or may not be guided by the reflexive oculomotor system. Boot et al. (2004), however, had observers track onsets as they were presented sequentially, an injunction that is certain to engage the reflexive eye movement system. Determining the degree to which the reflexive oculomotor system is involved in any given search behaviour ought to clarify this ambiguity, and may serve to explain the various patterns observed when different tasks (and objectives) are employed (Dodd, van der Stigchel, & Hollingworth, 2009).

5.4 CONCLUSION

In conclusion, this dissertation presents a framework that can account for numerous divergent patterns observed in the IOR literature, and several new empirical findings that fit within this framework. Although work remains to address outstanding components of our framework (i.e., object-based representations, mechanisms influencing visual search, neural implementation), empirical dissociations across three converging diagnostics—central arrow targets, PRP method and locus of slack logic, and consideration of performance in speed-accuracy space—and computational evidence support our hypothesis that there are two distinct forms of inhibition of return, with the activation state of the reflexive oculomotor system influencing which form is observed. Importantly, both forms perform the functional role of biasing the organism toward novelty, as ascribed originally by Posner—they just do so differently. We propose the input form operates on a saliency map (Fig. 5.1 - Red), serving to influence the information received from previously cued objects/locations, whereas the output form operates on a priority map (Fig. 5.1 - Green), serving to influence the selection of responses relative to prior orienting behaviours.

APPENDIX A - DRIFT DIFFUSION MODELLING

Excerpts from this manuscript are presented below. Co-authors for this manuscript are Dr. W. Joseph MacInness and Dr. Raymond M. Klein, respectively. In this manuscript, my contributions include, but aren't limited to, writing, theory development, data collection, data analysis, and project management.

Redden, R. S., MacInnes, W. J., & Klein, R. M. (2021). Inhibition of return: An information processing theory of its natures and significance. Cortex, 135, 30-48. This manuscript does not exactly replicate the final version published in Cortex. It is not a copy of the original published article and is not suitable for citation as such.

SMA.1 - COMPUTATIONAL MODELS OF IOR

Inhibition of return has been the subject of numerous computational models with each exploring various combinations of neural, cognitive and functional aspects. For example, salience models (Itti & Koch, 2001; Walther, Itti, Riesenhuber, Poggio, & Koch, 2002; Park, Shin & Lee, 2002; for review, see; Krasovskaya & MacInnes, 2019) have used IOR in its functional role by reducing activation at previously attended salient locations. In these models, bottom up salience is computed using gaussian pyramids and separate feature maps inspired by feature integration theory (Treisman & Gelade, 1980), but a mechanism is needed to prevent repeated fixations at the point of maximal salience. IOR is implemented on top of a leaky integrate and fire layer (Burkitt, 2006) to depress activation on the salience map to allow selection of new locations. The Superior Colliculus plays a role in IOR, and has been the source of multiple approaches to neural models of its expression. Dynamic neural field (DNF) models for example can represent IOR as short-term depression within the superior colliculus and have been used to simulate spatial cuing (Satel, Wang, Trappenberg, & Klein, 2011), biphasic patterns of RTs at short and long CTOAs (Lim et al., 2018), and in overt search (Wang, Satel, Trappenberg, & Klein, 2011).

Accumulator models have been a popular approach since they provide accurate simulation of response time distributions (Ratcliff & McKoon, 2008; Ratcliff & Van Dongen, 2011), but they also mirror neural activation in areas involved in IOR such as the lateral intraparietal area (Gold & Shadlen, 2003), the frontal eye fields (Hanes & Schall, 1996) and the superior colliculus (Ratcliff, Hasegawa, Hasegawa, Smith, &

Segraves, 2007). Many accumulator models have also been used to infer underlying cognitive mechanisms through parameter changes needed to satisfy experimental conditions. For example, Ludwig, Farrell, Ellis, and Gilchrist (2009) applied a Linear Ballistic Accumulator (LBA) to sequential saccades, and found inhibition of saccadic return was best accounted for with a reduction in information accumulation at previously fixated locations and they interpreted this as shift in target desirability. However, Prinzmetal, Taylor, Myers, and Nguyen-Espino (2011) applied the LBA to a spatial cueing task, and found IOR was best described as a change in decision threshold. Likewise, MacInnes (2017) used drift diffusion to model the spatial gradient of IOR (e.g., Vaughan, 1984; Hooge & Frens, 2000; Bennett & Pratt, 2001; Wang, Yan, Klein, & Wang, 2018), and found that the gradient of inhibition for both manual and saccadic responses could be accounted for by a change in the amount of information needed to make a response. This was interpreted as evidence for similar underlying mechanisms for the phenomenon across effectors.

SMA.1.1 - Present Work

Here we describe a computational model of the findings from Redden et al. (2016; see Fig. SMA.1). As described previously, in this study two groups of participants completed the same cueing task wherein an uninformative peripheral cue preceded a non-spatial discrimination target. The key experimental manipulation was the eye movement instruction at the time of the cue: one group generated an eye movement to the cued location and back to fixation (prosaccade) prior to target onset, and the other generated an eye movement to the uncued location and back to fixation (antisaccade) prior to target onset (Fig. SMA.1 - Left). Importantly, whereas both groups showed slower responses to targets presented at the cued location, they showed markedly different performance with respect to ac- curacy. The antisaccade group showed a cost in accuracy at the previously cued location (i.e., genuine cost in performance; Fig. 1.5 - dotted arrows), whereas the prosaccade group showed an improvement in accuracy at the previously cued location (i.e., speed-accuracy tradeoff; Fig. 1.5 - solid arrow). These patterns are consistent, respectively, with the input attribution ascribed to IOR when the reflexive oculomotor system is suppressed, and the output attribution ascribed to IOR when the reflexive oculomotor system is not suppressed. By implementing drift diffusion modelling of these data, we hope to specify how various underlying computational decision components may be affected by these two inhibitory phenomena. Specifically, we hypothesize that different decision parameters will be affected by the eye movement instruction manipulation.

SMA.2 DRIFT DIFFUSION MODEL: INPUT AND OUTPUT IOR

SMA.2.1 - Parameters

The drift diffusion model was implemented in Matlab (2016a) and based on parameter learning from MacInnes (2017). The model parameters specified a 2AFC with separate but equal thresholds (A and 0) for the possible responses—the correct or the incorrect one for each trial. Starting bias and variance (Z; s(z)) were not enabled in the model since the cues were not informative of the non-spatial response. The mean drift rate v ensured accumulation toward the correct threshold, but within trial variability s(v) around that mean could result in accumulation to either threshold and at variable times (see Fig. SMA.2 for model parameters). Non-decision parameters u and w were combined to a single parameter. Of the parameters included in the model, single best fit parameter effects that would result in increased RTs are limited to: increased non-decision time (u + w), reduced drift rate (v), reduced within-trial variability (s(v)), and increased threshold (A/0). A reduction in drift rate (v) would elicit an associated cost in accuracy, whereas reducing within-trial variability (s(v)) or increasing threshold (A/0) would result in an associated improvement in accuracy. Increased non-decision time (u + w) would elicit no change in accuracy, as the evidence accumulation process is independent of these parameters.

SMA.2.2 - Method

We report how we determined our sample size, all data exclusions (if any), all data inclusion/exclusion criteria, whether inclusion/exclusion criteria were established prior to data analysis, all manipulations, and all measures in the study. Since our modelling is a re-analysis of previously reported data (Redden, 2020[9]), the sample size, data inclusion/exclusion criteria, and all experimental manipulations/measures were determined by the original authors (Redden et al., 2016). As such, whereas no part of the present study was pre-registered in a time-stamped, institutional registry, all decision criteria were based precisely on the methods and analysis of the original study.

Best fit values for this model were learned using a genetic algorithm from MacInnes (2017) and matched to data from Redden et al. (2016). Since we were interested in the change in parameter space for the different experiment conditions, the

---

[9] Model code and data available at https://osf.io/3wqtr/.

model was initially trained for a baseline condition with threshold fixed and freedom for all remaining parameters, then restricted to a single best fit parameter to describe the model change for experiment manipulations. The primary question was whether information accrual at cued locations would be best described as a change in accumulation rate or as a bias (change in threshold) as compared to uncued locations, and whether this best fit differed in the prosaccade and antisaccade condition. To this end, the full set of parameters of the initial model were initially trained to fit uncued prosaccade trials (uncued antisaccade trials could have also been used since these baselines were identical) using the process described next.

Genetic algorithms begin with multiple sets of random values for parameter selection, and use some 'fitness function' to quantify how well the resulting parameters match against real world observation. We will call each set of random parameters 'contenders', and the initial set of contenders the first 'generation'. For each contender, a set of responses were generated so that the contender's performance could be statistically compared to the true human response distribution. For all tests and comparisons, the human data for each condition was treated in its entirety, meaning that individuals were not modelled separately. Comparison initially used three statistical tests: a t-test of the percentage errors found for the contender and for the human participants, a z-test of human vs model RTs, and two Kolmogorov-Smirnov (KS) tests (Weber, Leemis & Kinkaid, 2006) to compare the human and contender RT distributions for both correct and incorrect responses. The t-test and z-test required multiple samples, so means for multiple simulated subjects were generated for a given parameter set. Optimizing multiple values for a single fitness function is a known problem in machine learning (see Konak, Coit, & Smith, 2006 for a review) and in practice often combined to a single ‘fitness score’ to evaluate each contender and provide the best fit for both correct and incorrect responses. We did try a number of weighted formulas to combine speed and accuracy optimization, but our best results were found with the following two stage process using *pareto optimality* (Konak et al., 2006). A vector of three separate statistics [KS, t, z] were kept for the parameter set belonging to every contender. A contender’s vector was said to dominate that of another if any of the statistics improved without the others in the vector getting worse. A pareto optimal solution (or solutions) is one that is not dominated by any other solution. For each generation, our fitness function attempted to find the best known pareto optimal set for that generation, meaning contenders that were not dominated within that generation. We additionally placed a restriction rule that any pareto optimal solution must not differ significantly for the z and t test when comparing human and model data. So the minimal KS statistics would dominate, but only if the other tests dismissed the hypothesis that the human and model rates were different.

*If T(acc) < 2.0 & Z(corrRT) < 1.7, F() 1⁄4 KS(inc) * KS(corr)*

Else F() 1⁄4 1.0

The parameter space was optimized through 10,000 iterations in stage one, but was followed by a second optimization regime to prioritize the correct accuracy KS test, since the relatively fewer incorrect responses may have biased the fitness toward that statistic in the first stage. Using the parameter space from stage one as the starting point, the genetic algorithm was run a second time with 5000 iterations with only accuracy and KS test for correct responses in the fitness function.

If T(acc) < 2.0, F() 1/4 KS(corr)
Else F() 1/4 1.0

This allowed a refinement of the stage one parameter space toward a better fit on correct responses, but hopefully with parameters still near the incorrect KS-test space if such a solution existed.

For each iteration and contender, the vector of values listed above were recorded as a similarity score for fitness function that reflected the goodness of that contender as compared to human data. Contenders in the first generation were then sorted according to whether they were dominated by other contenders, with up to five contenders satisfying pareto optimality being selected to move on to the next generation. If fewer than 5 solutions were currently pareto optimal, the best KS statistic among dominated solutions were added to keep the number consistent at each generation. In addition to the best contenders, the next generation was also populated with variations of these winners including: mutations that had a single parameter randomly changed; and crossovers that received parameters from two different top contenders from the previous generation. Five additional contenders with random parameters rounded out the set of contenders for the next generation. This process was repeated recursively for following generations until the similarity score was seen to asymptote. KS statistic reached asymptote after 400 generations, but the algorithm was tested up to 10000/5000 to ensure that a global maximum was reached.

Once best-fit parameters were obtained for the uncued pro-saccade condition, a modified search algorithm was run to determine which single parameter best produced the change in results caused by a cued target, separately for the prosaccade and antisaccade conditions. The genetic search only allowed a single parameter to mutate from the uncued prosaccade parameter set to determine which single parameter best fit that experiment manipulation. Although allowing multiple parameters to mutate would have certainly produced a better mathematical fit, our objective for this model was to determine the single best fitting parameter to inform how the experiment manipulation changed underlying information processing. Since the presence/absence of speed accuracy tradeoffs are critical in the input/output theory of IOR, the fitness function was modified to test the accuracy of the resulting vector in speed-accuracy (SAT) space (Redden et al., 2016). In particular, the fitness function was the inverse of the difference in angular vector between human and model data in SAT space. This allowed us to promote RT and accuracy equally and test for the parameters predicted by the speed accuracy tradeoff. Note that while we did not include KS statistic as part of the fitness function for the cued data, we still tested these results for model evaluation.

SMA.2.3 - Results

We compared the resulting simulated distribution of the initial parameter (after the complete two-stage process) set against human data from the uncued/pro-saccade condition. With the KS test, we could not dismiss the null hypothesis that the human and model data are from the same distribution (K = .03, p = .57, and see Fig. SMA.3). Likewise for incorrect responses, we could not dismiss that distributions were the same (K =.14, p = .31). Note that the small number of incorrect trials makes fitting the KS statistic easier despite the better visual fit of the correct distribution.

Likewise, human-model t-tests of the reaction times (t < 1, p = .84) and errors (t < 1,p = .63) do not reject the null that human and model are the same.

For results testing the changes needed for the cued data, our primary hypothesis focused on the direction of the speed accuracy tradeoff. Cued trials in the antisaccade condition were consistently well modelled by a 7.5% reduction in drift mean suggesting a pure performance cost at the location of the cue that slows RTs and reduces accuracy rates. Model simulations with the drift mean adjusted for antisaccade cue were not seen as different from human data for neither RT (t < 1, p = .59) nor accuracy (t < 1, p = .99). Prosaccade trials, however, were equally well modelled by two different parameter changes. As predicted, simulations with a 6% increase in threshold for cued as compared to uncued prosaccade trials (Fig. SM.4) resulted in RTs (t < 1, p = .94) and accuracies (t < 1, p = .99) similar to human data. But we saw equally well-fitting simulations with a 10% reduction in trial noise for cued pro-saccade data with RT (t < 1, p = .98) and accuracy (t < 1, p = .85) not dissimilar from human results. Both trial noise (Fig. SMA.4 - TOP) and threshold (Fig. SMA.4 - BOTTOM) match the critical speed accuracy trade-off (slower, more accurate) that we observe in human data.

Although cued data were not trained specifically on reaction time distributions, we decided to test the two contending parameters for the pro-saccade findings to see if either held an advantage (Fig. SMA.5). Simulated results for both parameters were indistinguishable from human data according to the KS test (Threshold KS = .027; p = .70; Noise KS = .033, p = .41). The only minor diagnostic that preferred one parameter fit over the other was comparing the standard deviations of 20 simulated results compared to standard deviations of human data. Threshold results were seen as different from human data (t(19) = 5.7, p < .001) whereas Noise data was not (t(19) = 1.4, p = .18). Whereas this difference may point to diagnostics that may disentangle these two parameters in the future, we will not read too much into this post-hoc result given the excellent fit of both parameters on the other key metrics. Moreover, there are logical reasons based on the other diagnostics described earlier that favour the change in threshold (see Ch 6).

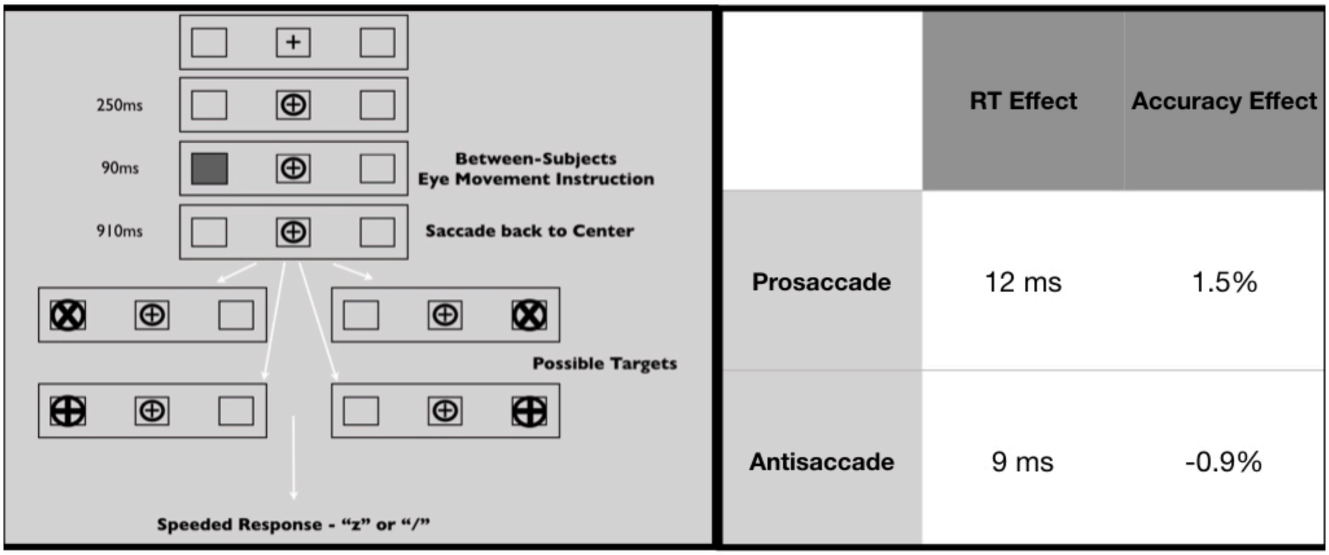

	RT Effect	Accuracy Effect
Prosaccade	12 ms	1.5%
Antisaccade	9 ms	-0.9%

Figure SMA.1 (LEFT) Methods figure redrawn from Redden et al. (2016). (RIGHT) Table of cueing effects (RT and Accuracy) as a function of whether observers were required to generate a prosaccade or antisaccade relative to the onset of the spatially- uninformative cue.

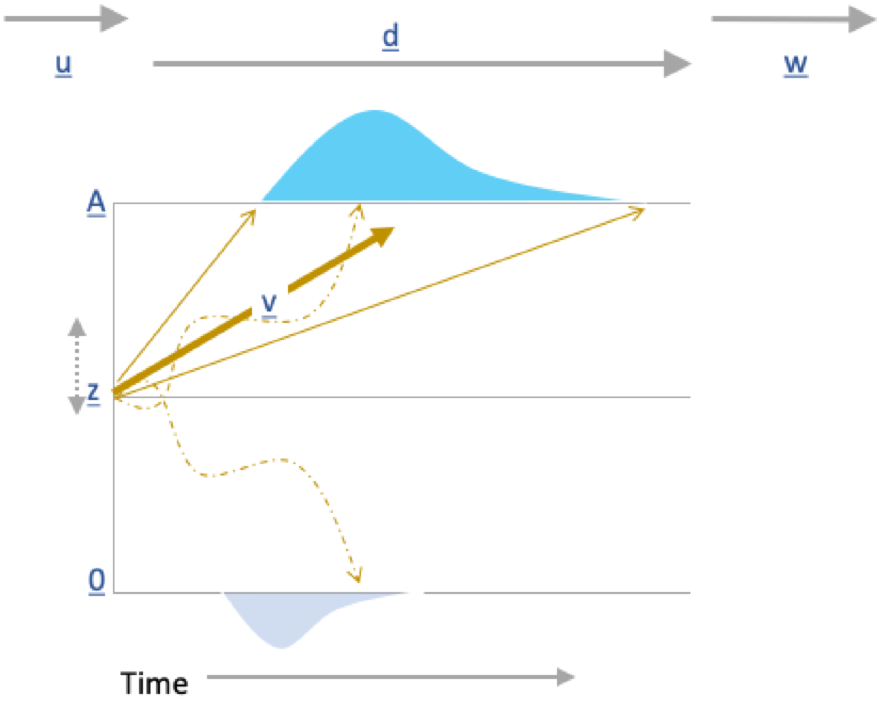

Figure SMA.2 Decision parameters for the accumulation of evidence over time toward one of two possible choice thresholds (A & 0). Starting points z and starting variances s(z) represent bias to one of the choices. Drift rate v is a constant signal toward one (correct) of the two choices while within trial variance around that drift rate (s(v)) add further response choice and response time uncertainty. A response is made when accumulation crosses either threshold and the time to reach that threshold is recorded. Non decision components (u; w) are also added.

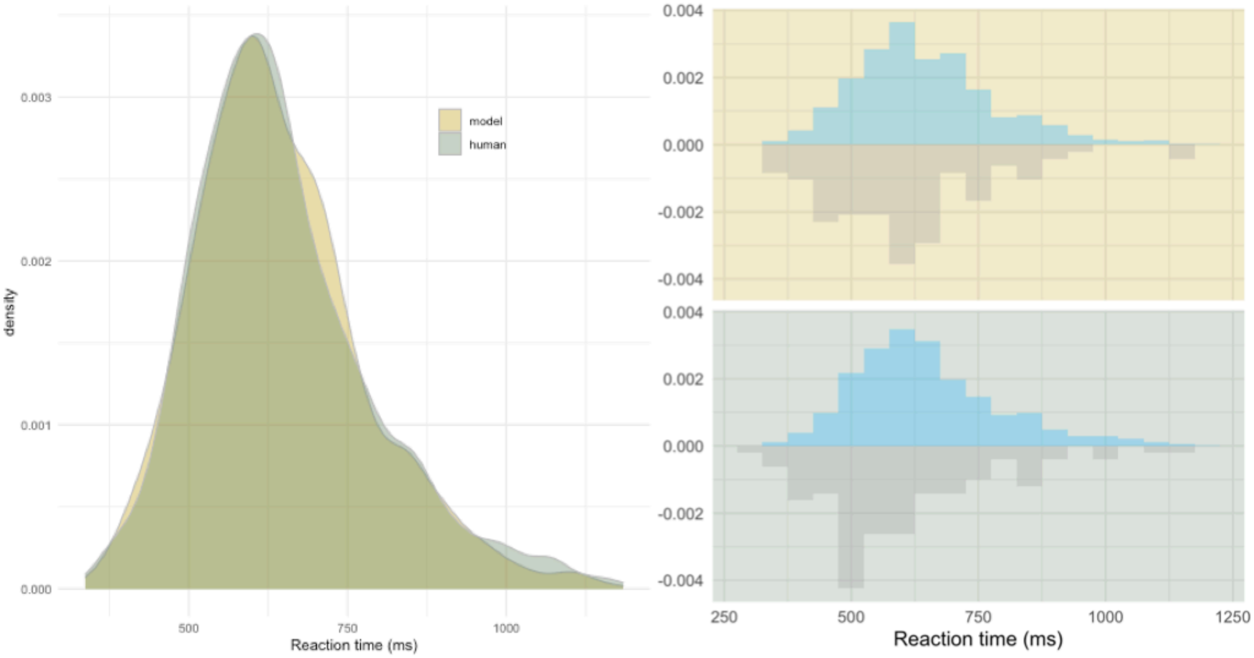

Figure SMA.3 Human and model density distributions of reaction time for correct responses (LEFT). On the right, correct (blue) and incorrect (dark grey) distributions together, for model (TOP; yellow background) and human (BOTTOM; light grey background) performance. All densities are set for the area under the curve to integrate to one separately for the correct and incorrect responses.

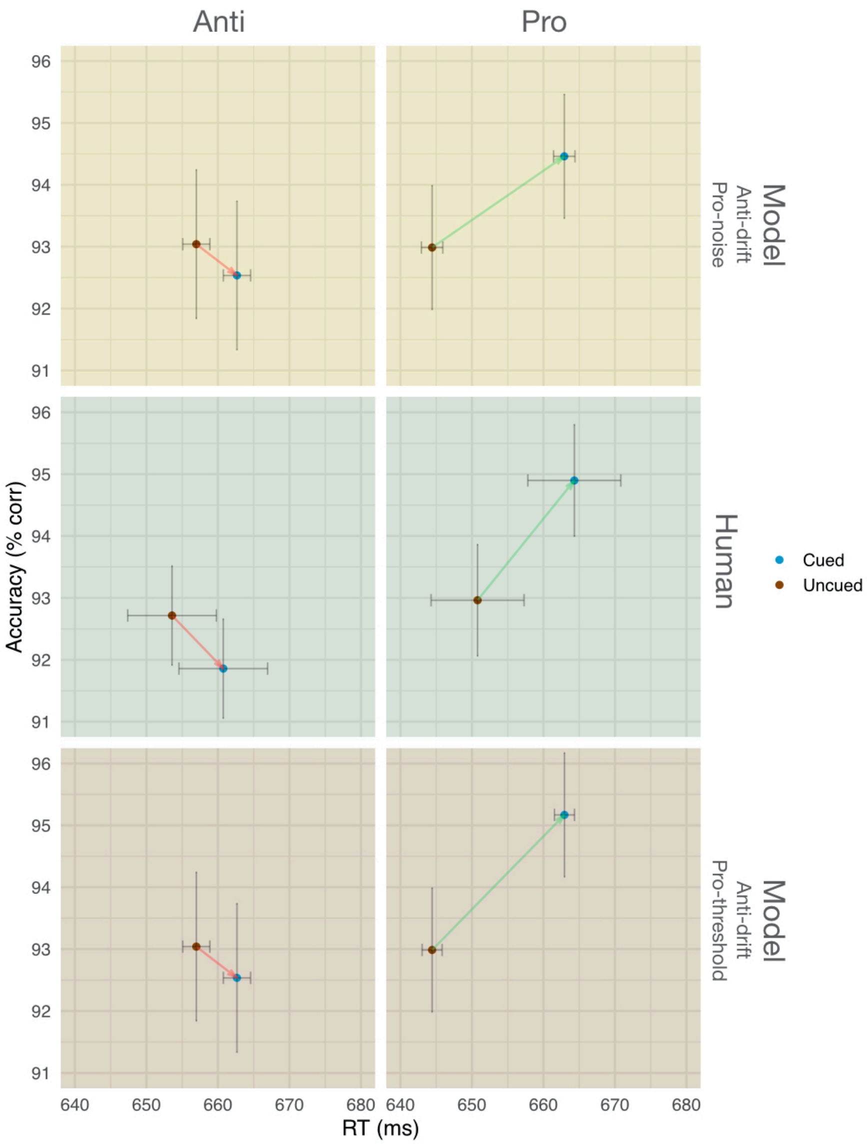

Figure SMA.4 Results from drift diffusion modelling (top and bottom panels) of the human data from Redden et al. (2016; middle panels). Model results from two separate simulations of the model. Both simulations are fit by a 7.5% reduction in drift mean (top left and bottom left) for antisaccade data and display the expected drop in performance (slower and less accurate; red arrows). Prosaccade data, however was equally fit by a 6% increase in threshold (bottom right) and a 10% reduction in trial noise (top right) with both resulting in the observed speed accuracy tradeoff (slower but more accurate; green arrows). Error bars represent Fisher's least significant difference for the effect of cueing.

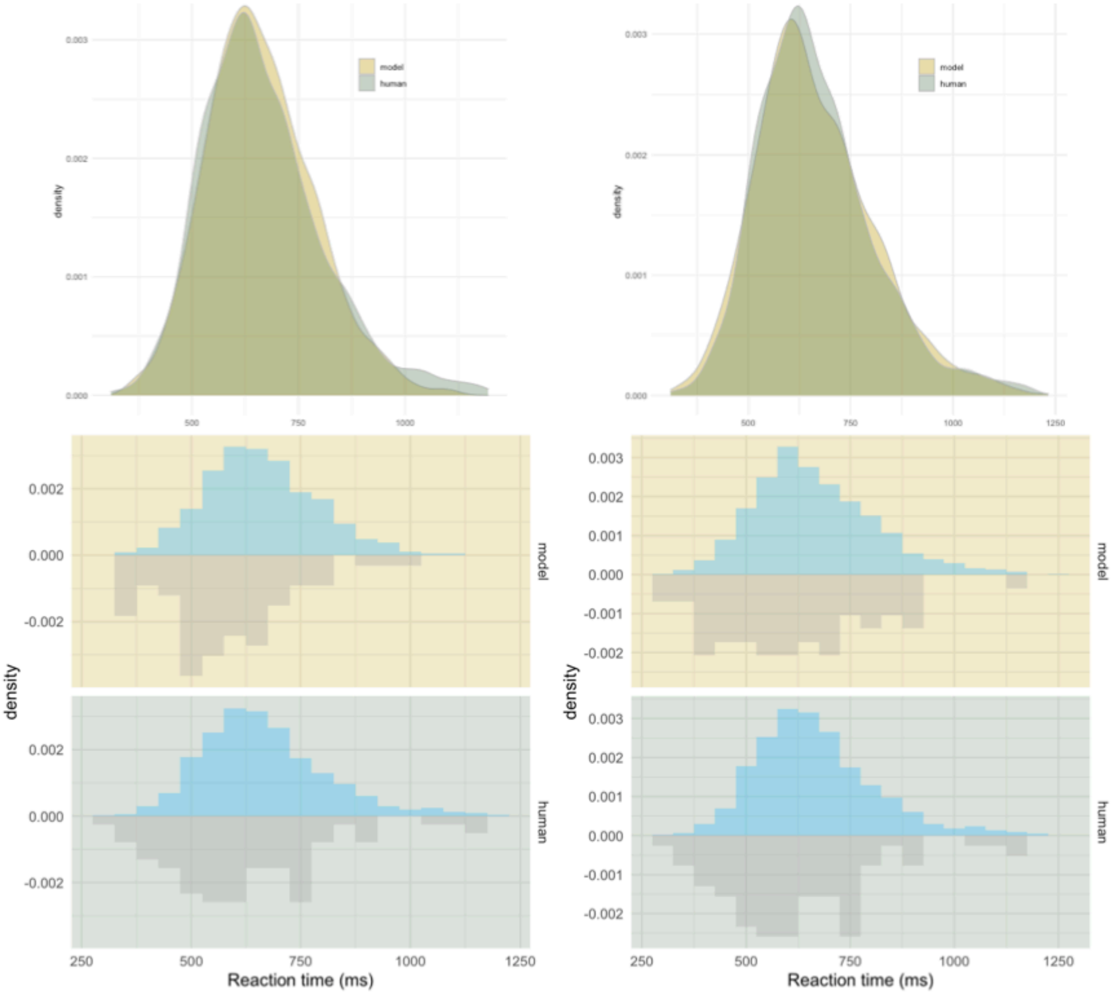

Figure SMA.5 Density distributions for correct responses for prosaccade cued trials when shifting noise parameter (top left) and threshold parameter (top right). Likewise, model (yellow) and human (light grey) distributions for correct (blue) and incorrect (dark grey) for Noise (bottom left) and threshold (bottom right) adjustments. All densities are set for the area under the curve to integrate to one separately for the correct and incorrect responses.

APPENDIX B - SUPPLEMENTAL ANALYSES

Supplemental analyses to those reported in Chapter 4 are reported here. Again, co-authors for this manuscript are Dr. Maryam Kavyani, Dr. John Christie, and Dr. Raymond M. Klein.

SMB.1 - SUPPLEMENTAL ANALYSES

The analyses that follow are exploratory in nature. For both experiments, the primary analyses from the main body of the dissertation are bifurcated into two separate conditions: one for each level of Feedback. This post-hoc approach affords examination of the nature of the effects of interest within each group.

Additionally, a second set of exploratory post-hoc analyses are presented for both experiments (and subsequently bifurcated as above) to explore the extent to which local effects influenced targets appearing at cued locations. For these analyses, the main effect of CoG is removed, and instead replaced with the main effects of Cueing (Cued; Uncued) and Number of Cues (1, 2, and 4).

SMB 1.1 Experiment 1 Bifurcated Main Analysis - No Feedback

For rate of responding, there was evidence to support the two-way interaction (Figure 4.2 - Left) between Cue Array and CoG, (Multi Cued - Multi Uncued) x CoG, b = -0.003, CI95% = [-0.009, 0.003], (Multi Cued - Single) x Feedback, b = 0.001, CI95% = [-0.004, 0.007], with greater support for the model with the interaction term included (AIC = 3038) than when the term was dropped (AIC = 3040).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 3040) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the main effect of CoG, b = -0.000, CI95% = [-0.001, 0.001]. The model performed worse (△AIC = +1) when dropping the main effect of Cue

Array, (Multi Cued - Multi Uncued), b = -0.009, CI95% = [-0.025, 0.006], (Multi Cued - Single), b = -0.017, CI95% = [-0.032, -0.001].

For accuracy, there was no evidence to support the two-way interaction (Figure 4.3 - Left) between Cue Array and CoG, (Multi Cued - Multi Uncued) x CoG, b = -0.022, CI95% = [-0.136, 0.080], (Multi Cued - Single) x Feedback, b = -0.018, CI95% = [-0.123, 0.079], with greater support for the model with the interaction term dropped (AIC = 3520) than when the term was included (AIC = 3523).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 3520) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the main effect of CoG, b = -0.006, CI95% = [-0.031, 0.021]. The model performed equivalently well (△AIC = 0) when dropping the main effect of Cue Array, (Multi Cued - Multi Uncued), b = 0.264, CI95% = [0.011, 0.536], (Multi Cued - Single), b = 0.264, CI95% = [0.022, 0.549].

SMB 1.2 Experiment 1 Bifurcated Main Analysis - Feedback

For rate of responding, there was evidence to support the two-way interaction (Figure 4.2 - Right) between Cue Array and CoG, (Multi Cued - Multi Uncued) x CoG, b = -0.002, CI95% = [-0.009, 0.003], (Multi Cued - Single) x Feedback, b = 0.001, CI95% = [-0.004, 0.006], with greater support for the model with the interaction term included (AIC = 3250) than when the term was dropped (AIC = 3251).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 3251) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the main effect of CoG, b = 0.000, CI95% = [-0.002, 0.001]. The model performed worse (△AIC = +25) when dropping the main effect of Cue Array, (Multi Cued - Multi Uncued), b = 0.009, CI95% = [-0.005, 0.023], (Multi Cued - Single), b = -0.023, CI95% = [-0.035, -0.007].

For accuracy, there was no evidence to support the two-way interaction (Figure 4.3 - Right) between Cue Array and CoG, (Multi Cued - Multi Uncued) x CoG, b = 0.014, CI95% = [-0.067, 0.092], (Multi Cued - Single) x Feedback, b = 0.018, CI95% = [-0.056, 0.090], with greater support for the model with the interaction term dropped (AIC = 5606) than when the term was included (AIC = 5610).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 5606) with models where each term was dropped. The model performed better (△AIC = -1) when dropping the main effect of CoG, b = 0.005, CI95% = [-0.014, 0.023]. The model performed better (△AIC = -2) when dropping the main effect of Cue Array, (Multi Cued - Multi Uncued), b = 0.054, CI95% = [-0.122, 0.252], (Multi Cued - Single), b = 0.122, CI95% = [-0.057, 0.323].

SMB.1.3 Analysis - Experiment 1

As in the main body of the manuscript, the data were analyzed primarily with mixed effects models, linear analysis of correct trials for response speed and logistic regression for accuracy. Again, the rate of responding per second was analyzed rather than the typical RT in ms, and back transformations to RT were made for figures and RT summary values. Regression coefficients are again negative rates in order to keep direction of effects the same as RT in ms.

To examine the more traditional effect of cueing, the effects of Cueing (Cued; Uncued), Number of Cues (1, 2 , 4), and Feedback (Feedback; No Feedback) were included in the model as predictors. The most complex models were run first, with AICs computed via the drop1 method in the {stats} package.

SMB.1.4 Rate of Responding - Experiment 1

The there was no evidence to support the three-way interaction (Figure SMB.1), Cueing x Number of Cues x Feedback, b = 0.005, with greater support for the model with the interaction term dropped (AIC = 6295) than when the term was included (AIC = 6296).

To evaluate the two-way interactions, we contrasted the model with all two-way interaction terms included (AIC = 6295) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the two-way interaction term between Cueing and Number of Cues, b = 0.001. The model performed equivalently well (△AIC = 0) when dropping the two-way interaction term between Cueing and Feedback, b = 0.014. The model performed worse (△AIC = +2) when dropping the two-way interaction term between Number of Cues and Feedback, b = 0.006.

To evaluate the main effects, we contrasted the model with all main effect terms included (AIC = 6293) with models where each term was dropped. The model performed better (△AIC = -1) when dropping the main effect of Cueing, b = 0.002. The model performed worse (△AIC = +19) when dropping the main effect of Number of Cues, b = 0.007. The model performed better (△AIC = -2) when dropping the main effect of Feedback, b = -0.004.

SMB.1.5 Accuracy - Experiment 1

The there was evidence to support the three-way interaction (Figure SMB.2), Cueing x Number of Cues x Feedback, b = -0.226, with greater support for the model with the interaction term included (AIC = 9123) than when the term was dropped (AIC = 9125).

To evaluate the two-way interactions, we contrasted the model with all two-way interaction terms included (AIC = 9125) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the two-way interaction term between Cueing and Number of Cues, b = 0.038. The model performed better (△AIC = -2) when dropping the two-way interaction term between Cueing and Feedback, b = -0.108. The model performed better (△AIC = -3) when dropping the two-way interaction term between Number of Cues and Feedback, b = -0.025.

To evaluate the main effects, we contrasted the model with all main effect terms included (AIC = 9120) with models where each term was dropped. The model performed equivalently well (△AIC = 0) when dropping the main effect of Cueing, b = 0.107. The model performed better (△AIC = -2) when dropping the main effect of Number of Cues, b = -0.011. The model performed worse (△AIC = +6) when dropping the main effect of Feedback, b = -0.614.

Furthermore, we treated the two (between-subjects) Feedback conditions separately for consideration relative to our theoretical predictions.

SMB.1.6 Experiment 1 - No Feedback Cueing

For rate of responding, there was no evidence to support the two-way interaction (Figure SMB.1 - Left) between Cueing and Number of Cues, b = -0.001, with greater support for the model with the interaction term dropped (AIC = 3038) than when the term was included (AIC = 3040).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 3038) with models where each term was dropped. The model performed better (△AIC = -1) when dropping the main effect of Cueing, b = -0.005. The model performed worse (△AIC = +1) when dropping the main effect of Number of Cues, b = 0.004.

For accuracy, there was evidence to support the two-way interaction (Figure SMB.2 - Left) between Cueing and Number of Cues, b = 0.184, with greater support for the model with the interaction term included (AIC = 3518) than when the term was dropped (AIC = 3520).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 3520) with models where each term was dropped. The model performed equivalently well (△AIC = 0) when dropping the main effect of Cueing, b = 0.175. The model performed better (△AIC = -2) when dropping the main effect of Number of Cues.

SMB.1.7 Experiment 1 - Feedback Cueing

For rate of responding, there was no evidence to support the two-way interaction (Figure SMB.1 - Right) between Cueing and Number of Cues, b = 0.004, with greater support for the model with the interaction term dropped (AIC = 3258) than when the term was included (AIC = 3259).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 3258) with models where each term was dropped. The model performed equivalently well (△AIC = 0) when dropping the main effect of Cueing, b = 0.009. The model performed worse (△AIC = +19) when dropping the main effect of Number of Cues, b = 0.010.

For accuracy, there was no evidence to support the two-way interaction (Figure SMB.2 - Right) between Cueing and Number of Cues, b = -0.042, with greater support for the model with the interaction term dropped (AIC = 5605) than when the term was included (AIC = 5607).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 5605) with models where each term was dropped. The model performed better (△AIC = -1) when dropping the main effect of Cueing, b = 0.069. The model performed better (△AIC = -2) when dropping the main effect of Number of Cues, b = -0.020.

SMB 1.8 Experiment 2 Bifurcated Main Analysis - No Feedback

For rate of responding, there was no evidence to support the two-way interaction (Figure 4.4 - Left) between Cue Array and CoG, (Multi Cued - Multi Uncued) x CoG, b = -0.005, CI95% = [-0.018, 0.007], (Multi Cued - Single) x Feedback, b = -0.004, CI95% = [-0.015, 0.008], with greater support for the model with the interaction term dropped (AIC = 11463) than when the term was included (AIC = 11466).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 11463) with models where each term was dropped. The model performed better (△AIC = -1) when dropping the main effect of CoG, b = 0.001, CI95% = [-0.002, 0.004]. The model performed better (△AIC = -3) when dropping the main effect of Cue Array, (Multi Cued - Multi Uncued), b = 0.001, CI95% = [-0.032, 0.034], (Multi Cued - Single), b = -0.010, CI95% = [-0.042, 0.022].

For accuracy, there was no evidence to support the two-way interaction (Figure 4.5 - Left) between Cue Array and CoG, (Multi Cued - Multi Uncued) x CoG, b = -0.013, CI95% = [-0.118, 0.097], (Multi Cued - Single) x Feedback, b = 0.004, CI95% = [-0.094, 0.101], with greater support for the model with the interaction term dropped (AIC = 3404) than when the term was included (AIC = 3407).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 3404) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the main effect of CoG, b = 0.005, CI95% = [-0.022, 0.030]. The model performed better (△AIC = -3) when dropping the main effect of Cue Array, (Multi Cued - Multi Uncued), b = 0.006, CI95% = [-0.241, 0.261], (Multi Cued - Single), b = 0.133, CI95% = [-0.121, 0.427].

SMB 1.9 Experiment 2 Bifurcated Main Analysis - Feedback

For rate of responding, there was no evidence to support the two-way interaction (Figure 4.4 - Right) between Cue Array and CoG, (Multi Cued - Multi Uncued) x CoG, b = 0.002, CI95% = [-0.008, 0.013], (Multi Cued - Single) x Feedback, b = 0.006, CI95% = [-0.004, 0017], with greater support for the model with the interaction term dropped (AIC = 15765) than when the term was included (AIC = 15766).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 15765) with models where each term was dropped. The model performed worse (△AIC = +2) when dropping the main effect of CoG, b = 0.003, CI95% = [0.000, 0.005]. The model performed better (△AIC = -3) when dropping the main effect of Cue Array, (Multi Cued - Multi Uncued), b = 0.015, CI95% = [-0.009, 0.040], (Multi Cued - Single), b = 0.010, CI95% = [-0.016, 0.036].

For accuracy, there was no evidence to support the two-way interaction (Figure 4.5 - Right) between Cue Array and CoG, (Multi Cued - Multi Uncued) x CoG, b =

-0.022, CI95% = [-0.117, 0.074], (Multi Cued - Single) x Feedback, b = -0.011, CI95% = [-0.099, 0.076], with greater support for the model with the interaction term dropped (AIC = 4455) than when the term was included (AIC = 4459).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 4455) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the main effect of CoG, b = 0.002, CI95% = [-0.022, 0.025]. The model performed better (△AIC = -3) when dropping the main effect of Cue Array, (Multi Cued - Multi Uncued), b = -0.055, CI95% = [-0.303, 0.156], (Multi Cued - Single), b = 0.013, CI95% = [-0.257, 0.241].

SMB.1.10 Analysis - Experiment 2

As in the main body of the manuscript, the data were analyzed primarily with mixed effects models, linear analysis of correct trials for response speed and logistic regression for accuracy. Again, the rate of responding per second was analyzed rather than the typical RT in ms, and back transformations to RT were made for figures and RT summary values. Regression coefficients are again negative rates in order to keep direction of effects the same as RT in ms.

To examine the more traditional effect of cueing, the effects of Cueing (Cued; Uncued), Number of Cues (1, 2 , 4), and Feedback (Feedback; No Feedback) were included in the model as predictors. The most complex models were run first, with AICs computed via the drop1 method in the {stats} package.

SMB.1.11 Rate of Responding - Experiment 2

The there was no evidence to support the three-way interaction (Figure SMB.3), Cueing x Number of Cues x Feedback, b = 0.018, with greater support for the model with the interaction term dropped (AIC = 27223) than when the term was included (AIC = 27224).

To evaluate the two-way interactions, we contrasted the model with all two-way interaction terms included (AIC = 27223) with models where each term was dropped. The model performed better (△AIC = -1) when dropping the two-way interaction term between Cueing and Number of Cues, b = 0.004. The model performed worse (△AIC = +1) when dropping the two-way interaction term between Cueing and Feedback, b = 0.028. The model performed better (△AIC = +2) when dropping the two-way interaction term between Number of Cues and Feedback, b = -0.003.

To evaluate the main effects, we contrasted the model with all main effect terms included (AIC = 27221) with models where each term was dropped. The model performed worse (△AIC = +4) when dropping the main effect of Cueing, b = 0.020. The model performed better (△AIC = -1) when dropping the main effect of Number of Cues, b = 0.002. The model performed better (△AIC = -1) when dropping the main effect of Feedback, b = -0.016.

SMB.1.12 Accuracy - Experiment 2

The there was no evidence to support the three-way interaction (Figure SMB.4), Cueing x Number of Cues x Feedback, b = -0.078, with greater support for the model with the interaction term dropped (AIC = 7855) than when the term was included (AIC = 7856).

To evaluate the two-way interactions, we contrasted the model with all two-way interaction terms included (AIC = 7855) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the two-way interaction term between Cueing and Number of Cues, b = -0.030. The model performed better (△AIC = -2) when dropping the two-way interaction term between Cueing and Feedback, b = -0.056. The model performed better (△AIC = -2) when dropping the two-way interaction term between Number of Cues and Feedback, b = 0.039.

To evaluate the main effects, we contrasted the model with all main effect terms included (AIC = 7850) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the main effect of Cueing, b = -0.049. The model performed worse (△AIC = +3) when dropping the main effect of Number of Cues, b = -0.061. The model performed better (△AIC = -2) when dropping the main effect of Feedback, b = 0.107.

Furthermore, we treated the two (between-subjects) Feedback conditions separately for consideration relative to our theoretical predictions.

SMB.1.13 No Feedback Cueing - Experiment 2

For rate of responding, there was no evidence to support the two-way interaction (Figure SMB.3 - Left) between Cueing and Number of Cues, b = -0.001, with greater support for the model with the interaction term dropped (AIC = 11462) than when the term was included (AIC = 11463).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 11462) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the main effect of Cueing, b = 0.003. The model performed better (△AIC = -2) when dropping the main effect of Number of Cues, b = 0.004.

For accuracy, there was no evidence to support the two-way interaction (Figure SMB.4 - Left) between Cueing and Number of Cues, b = 0.013, with greater support for the model with the interaction term dropped (AIC = 3399) than when the term was included (AIC = 3401).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 3399) with models where each term was dropped. The model performed better (△AIC = -2) when dropping the main effect of Cueing, b = -0.018. The model performed worse (△AIC = +2) when dropping the main effect of Number of Cues, b = -0.083.

SMB.1.14 Feedback Cueing - Experiment 2

For rate of responding, there was no evidence to support the two-way interaction (Figure SMB.3 - Right) between Cueing and Number of Cues, b = 0.012, with equivalent support for the model with the interaction term dropped (AIC = 15763) than when the term was included (AIC = 15763).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 15763) with models where each term was dropped. The model performed worse (△AIC = +6) when dropping the main effect of Cueing, b = 0.032. The model performed better (△AIC = -2) when dropping the main effect of Number of Cues, b = 0.001.

For accuracy, there was no evidence to support the two-way interaction (Figure SMB.4 - Right) between Cueing and Number of Cues, b = -0.064, with greater support for the model with the interaction term dropped (AIC = 4452) than when the term was included (AIC = 4453).

To evaluate the main effects, we contrasted the model with both main effect terms included (AIC = 4452) with models where each term was dropped. The model performed better (△AIC = -1) when dropping the main effect of Cueing, b = -0.073. The model performed equivalently well (△AIC = 0) when dropping the main effect of Number of Cues, b = -0.044.

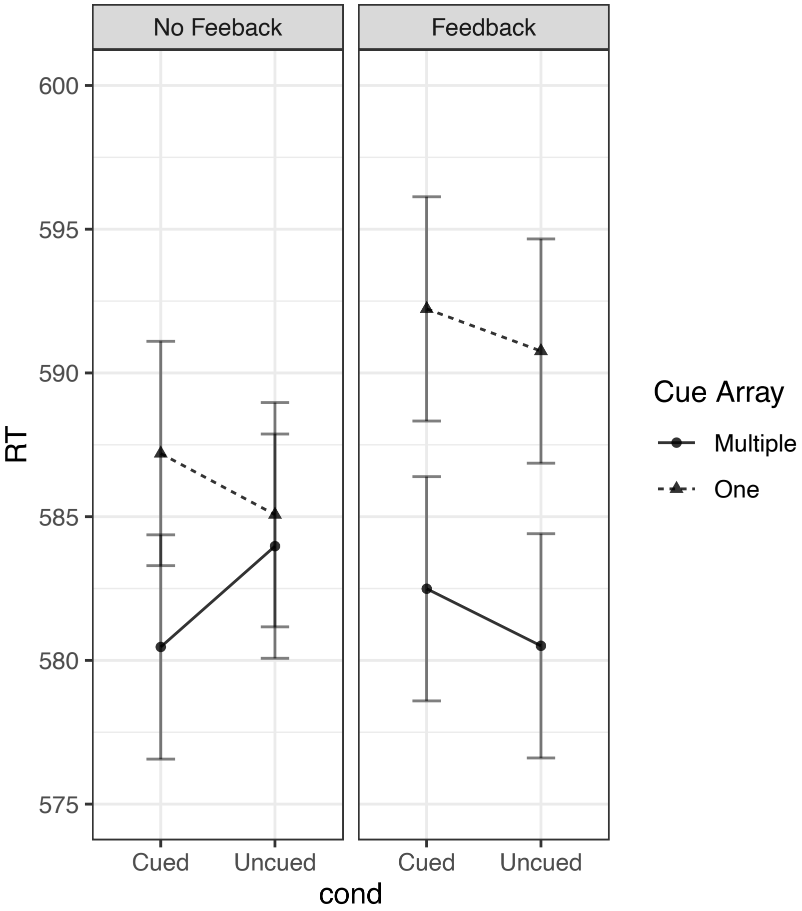

Figure SMB.1: The effect of cueing on reaction time as a function of the cue array (solid line = Multiple Cues; dashed line = Single Cue). The group that did not receive feedback about untoward eye movements is depicted on the left, whereas the group that received feedback is depicted on the right. Error bars represent Fisher's least significant difference.

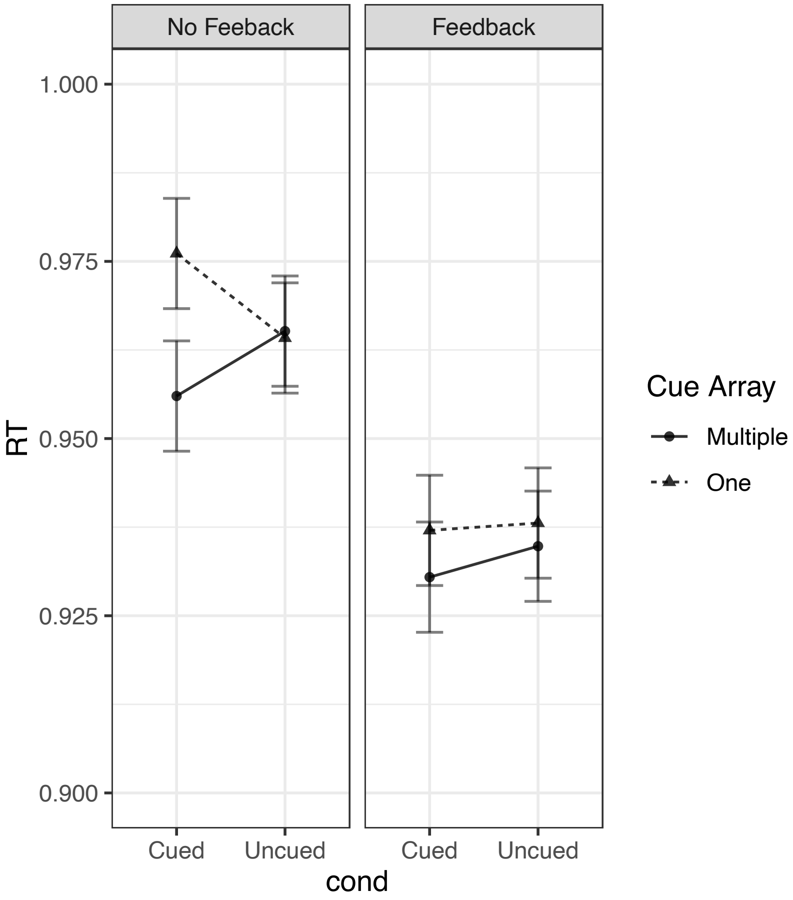

Figure SMB.2: The effect of cueing on accuracy as a function of the cue array (solid line = Multiple Cues; dashed line = Single Cue). The group that did not receive feedback about untoward eye movements is depicted on the left, whereas the group that received feedback is depicted on the right. Error bars represent Fisher's least significant difference.

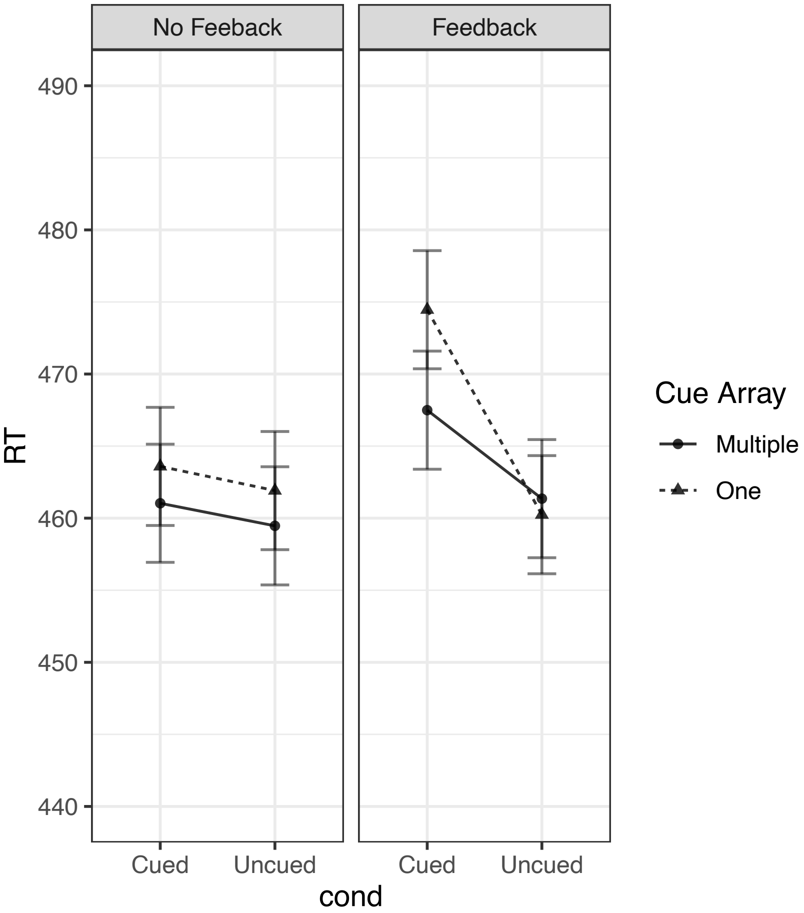

Figure SMB.3: The effect of cueing on reaction time as a function of the cue array (solid line = Multiple Cues; dashed line = Single Cue). The group that did not receive feedback about untoward eye movements is depicted on the left, whereas the group that received feedback is depicted on the right. Error bars represent Fisher's least significant difference.

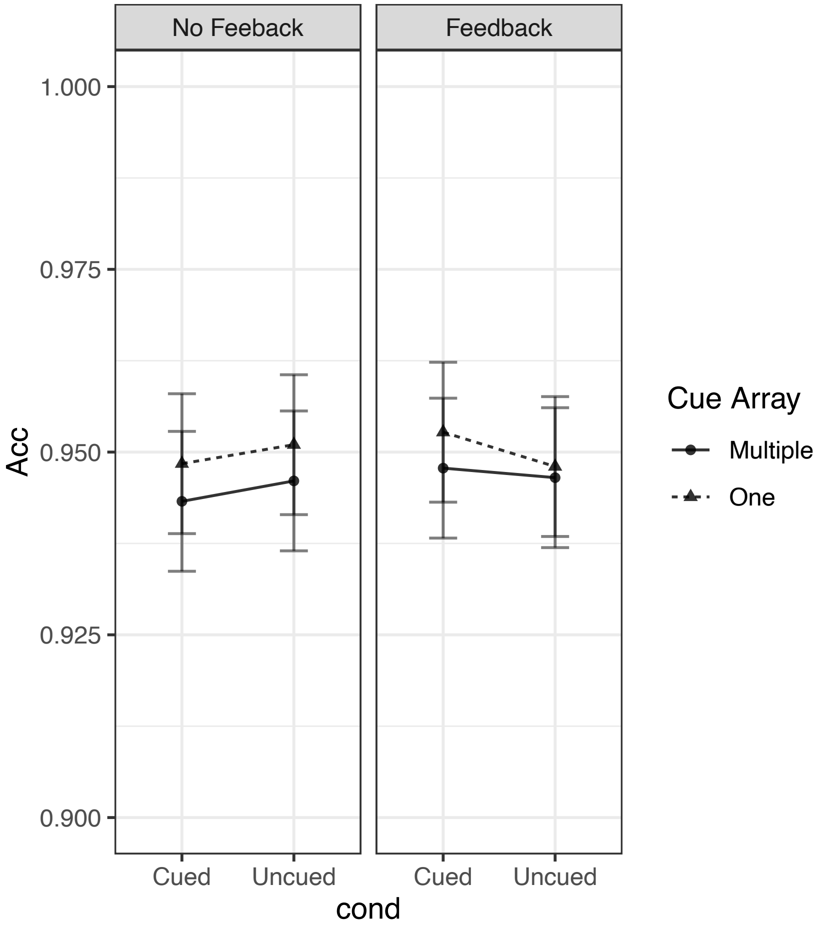

Figure SMB.4: The effect of cueing on accuracy as a function of the cue array (solid line = Multiple Cues; dashed line = Single Cue). The group that did not receive feedback about untoward eye movements is depicted on the left, whereas the group that received feedback is depicted on the right. Error bars represent Fisher’s least significant difference.

APPENDIX C - COPYRIGHT AGREEMENT LETTER

Copyright agreement letter for the subset of materials appearing in Chapters 1, 5, and Appendix A, from the following published manuscript:

Redden, R. S., MacInnes, W. J., & Klein, R. M. (2021). Inhibition of return: An information processing theory of its natures and significance. Cortex, 135, 30-48.

May 3, 2021

Attn: Cortex Editors

Drs. Della Sala and Grafman,

I am preparing my PhD thesis for submission to the Faculty of Graduate Studies at Dalhousie University, Halifax, Nova Scotia, Canada. I am seeking your permission to include a manuscript version of the following paper(s) as a chapter in the thesis:

Redden, R. S., MacInnes, W. J., & Klein, R. M. (2021). Inhibition of return: An information processing theory of its natures and significance. Cortex, 135, 30-48.

Canadian graduate theses are reproduced by the Library and Archives of Canada (formerly National Library of Canada) through a non-exclusive, world-wide license to reproduce, loan, distribute, or sell theses. I am also seeking your permission for the material described above to be reproduced and distributed by the LAC(NLC). Further details about the LAC(NLC) thesis program are available on the LAC(NLC) website (www.nlc-bnc.ca).

Full publication details and a copy of this permission letter will be included in the thesis.

Yours sincerely,

Ralph S. Redden

Permission is granted for:

a) the inclusion of the material described above in your thesis.

b) for the material described above to be included in the copy of your thesis that is sent to the Library and Archives of Canada (formerly National Library of Canada) for reproduction and distribution.

Name:_Sergio Della Sala__________ Title:_Professor__________

Signature:________________ Date:_14/5/2021__________

BIBLIOGRAPHY

Abrams, R. A., & Dobkin, R. S. (1994). Inhibition of return: effects of attentional cuing on eye movement latencies. *Journal of Experimental Psychology: Human Perception and Performance*, 20(3), 467.

Abrams, R. A., & Pratt, J. (2000). Oculocentric coding of inhibited eye movements to recently attended locations. *Journal of Experimental Psychology: Human Perception and Performance*, 26(2), 776.

Attneave, F. (1959). Comments: In defense of homunculi. In W. A. Rosenblith (Ed.), Sensory communication: Contributions to the symposium on principles of sensory communication, July 19–August 1 (pp. 777–781). Cambrdidge, MA: MIT Press.

Avery, B., Cowper-Smith, C. D., & Westwood, D. A. (2015). Spatial interactions between consecutive manual responses. *Experimental Brain Research, 233*(11), 3283-3290.

Bates D, Maechler M, Bolker B and Walker S (2014). _lme4: Linear mixed-effects models using Eigen Statistical Computing, Vienna, Austria. URL http://www.R-project.org/.

Bennett, P. J., & Pratt, J. (2001). The spatial distribution of inhibition of return. *Psychological Science, 12*(1), 76-80.

Berlucchi, G. (2006). Inhibition of return: A phenomenon in search of a mechanism and a better name. *Cognitive Neuropsychology, 23*(7), 1065-1074.

Bogacz, R., Wagenmakers, E. J., Forstmann, B. U., & Nieuwenhuis, S. (2010). The neural basis of the speed–accuracy tradeoff. *Trends in Neurosciences, 33*(1), 10-16.

Boot, W. R., Mccarley, J. S., Kramer, A. F., & Peterson, M. S. (2004). Automatic and intentional memory processes in visual search. *Psychonomic Bulletin & Review, 11*(5), 854-861.

Bourgeois, A., Chica, A. B., Migliaccio, R., de Schotten, M. T., & Bartolomeo, P. (2012). Cortical control of inhibition of return: evidence from patients with inferior parietal damage and visual neglect. *Neuropsychologia, 50*(5), 800-809.

Bourgeois, A., Chica, A. B., Valero-Cabré, A., & Bartolomeo, P. (2013). Cortical control of inhibition of return: causal evidence for task-dependent modulations by dorsal and ventral parietal regions. *Cortex, 49*(8), 2229-2238.

Briand, K. A., Larrison, A. L., & Sereno, A. B. (2000). Inhibition of return in manual and saccadic response systems. *Perception & Psychophysics, 62*(8), 1512-1524.

Broadbent, D. (1987). Simple models for experimentable situations.

Burkitt, A. N. (2006). A review of the integrate-and-fire neuron model: I. Homogeneous synaptic input. *Biological Cybernetics, 95*(1), 1-19.

Cheal, M., Chastain, G., & Lyon, D. R. (1998). Inhibition of return in visual identification tasks. *Visual Cognition, 5*(3), 365-388.

Chica, A. B., Taylor, T. L., Lupiañez, Y., & Klein, R. M. (2010). Two mechanisms underlying inhibition of return. *Experimental Brain Research, 201,* 25–35. doi:10.1007/s00221-009-2004-1

Christie, J., Hilchey, M. D., & Klein, R. M. (2013). Inhibition of return is at the midpoint of simultaneous cues. *Attention, Perception, & Psychophysics, 75*(8), 1610-1618.

Christie, J., Hilchey, M. D., Mishra, R., & Klein, R. M. (2015). Eye movements are primed toward the center of multiple stimuli even when the interstimulus distances are too large to generate saccade averaging. *Experimental Brain Research, 233*(5), 1541-1549.

Cowper-Smith, C. D., Harris, J., Eskes, G. A., & Westwood, D. A. (2013). Spatial interactions between successive eye and arm movements: signal type matters. *Plos One, 8*(3).

de Vries, J. P., Verstraten, F. A., Hooge, I. T., Fabius, J. H., & Van der Stigchel, S. (2020). Inhibition of return in the oculomotor decision process: Dissociating visual target discrimination from saccade readiness delays. *Journal of Experimental Psychology: Human Perception and Performance, 47*(1), 140–160.

Dorris, M. C., Klein, R. M., Everling, S., & Munoz, D. P. (2002). Contribution of the primate superior colliculus to inhibition of return. *Journal of Cognitive Neuroscience, 14*, 1256–1263.

Dorris, M. C., Taylor, T. L., Klein, R. M., & Munoz, D. P. (1999). Influence of previous visual stimulus or saccade on saccadic reaction times in monkey. *Journal of Neurophysiology, 81*, 2429 –2436.

Dukewich, K. R. (2009). Reconceptualizing inhibition of return as. *Psychonomic Bulletin & Review, 16*(2), 238-251.

Dukewich, K. R., & Klein, R. M. (2015). Inhibition of return: A phenomenon in search of a definition and a theoretical framework. *Attention, Perception, & Psychophysics, 77*(5), 1647-1658.

Everling, S., Dorris, M. C., Klein, R. M., & Munoz, D. P. (1999). Role of primate superior colliculus in preparation and execution of anti-saccades and pro-saccades. *Journal of Neuroscience, 19,* 2740–2754.

Findlay, J. M. (1982). Global visual processing for saccadic eye movements. *Vision Research, 22*(8), 1033-1045.

Fischer, M. H., Pratt, J., & Neggers, S. F. (2003). Inhibition of return and manual pointing movements. *Perception & Psychophysics, 65*(3), 379-387.

Forbes, K., & Klein, R. M. (1996). The magnitude of the fixation offset effect with endogenously and exogenously controlled saccades. *Journal of Cognitive Neuroscience, 8*(4), 344–352.

Frey, R. D. (1990). Selective attention, event perception and the criterion of acceptability principle: Evidence supporting and rejecting the doctrine of prior entry. *Human Movement Science, 9*, 481–530.

Gabay, S., Henik, A., & Gradstein, L. (2010). Ocular motor ability and covert attention in patients with Duane Retraction Syndrome. *Neuropsychologia, 48*(10), 3102-3109.

Gibson, B. S., & Egeth, H. (1994). Inhibition and disinhibition of return: Evidence from temporal order judgments. *Perception & Psychophysics, 56*(6), 669-680.

Gold, J. I., & Shadlen, M. N. (2003). The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *Journal of Neuroscience, 23*(2), 632-651.

Guest, O., & Martin, A. E. (2020, April 25). How computational modeling can force theory building in psychological science. https://doi.org/10.31234/osf.io/rybh9

Habibnezhad, M., Lawrence, M. A., & Klein, R. M. (2019). Using Rescorla's truly random control condition to measure truly exogenous covert orienting. *Psychonomic Bulletin & Review, 26*(2), 569-575.

Handy, T. C., Jha, A. P., & Mangun, G. R. (1999). Promoting novelty in vision: Inhibition of return modulates perceptual-level processing. *Psychological Science, 10*(2), 157-161.

Hanes, D. P., & Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science, 274*(5286), 427-430.

Hanks, T., Kiani, R., & Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *Elife, 3*, e02260.

Hartley, A. A., & Kieley, J. M. (1995). Adult age differences in the inhibition of return of visual attention. *Psychology and Aging, 10*(4), 670.

Heitz, R. P., & Schall, J. D. (2012). Neural mechanisms of speed- accuracy tradeoff. *Neuron, 76*(3), 616e628.

Hilchey, M. D., Dohmen, D., Crowder, N. A., & Klein, R. M. (2016). When is inhibition of return input-or output-based? It depends on how you look at it. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale, 70*(4), 325.

Hilchey, M. D., Hashish, M., MacLean, G. H., Ivanoff, J., Satel, J., & Klein, R. M. (2014). On the role of eye movement monitoring and discouragement on inhibition of return in a go no-go task. *Vision Research, 96,* 133–139. doi:10.1016/j.visres.2013.11.008

Hilchey, M. D., Ivanoff, J., Taylor, T. L., & Klein, R. M. (2011). Visualizing the temporal dynamics of spatial information processing responsible for the Simon effect and its amplification by inhibition of return. *Acta Psychologica, 136*(2), 235-244.

Hilchey, M. D., Klein, R. M., & Ivanoff, J. (2012). Perceptual and motor IOR: Components or flavors? *Attention, Perception, & Psychophysics, 74*, 1416–1429. doi:10.3758/s13414-012-0332-x

Hilchey, M. D., Klein, R. M., & Satel, J. (2014). Returning to “inhibition of return” by dissociating long-term oculomotor ior from short-term sensory adaptation and other nonoculomotor “inhibitory” cueing effects. *Journal of Experimental Psychology: Human Perception and Performance, 40*(4), 1603–1616.

Hilchey, M. D., Pratt, J., & Christie, J. (2018). Placeholders dissociate two forms of inhibition of return. *Quarterly Journal of Experimental Psychology, 71*(2), 360-371.

Hommel, B. (1994). Spontaneous decay of response-code activation. *Psychological Research, 56*(4), 261-268.

Hooge, I. T. C., & Frens, M. A. (2000). Inhibition of saccade return (ISR): Spatio-temporal properties of saccade programming. *Vision research, 40*(24), 3415-3426.

Hooge, I. T. C., Over, E. A., van Wezel, R. J., & Frens, M. A. (2005). Inhibition of return is not a foraging facilitator in saccadic search and free viewing. *Vision Research, 45*(14), 1901-1908.

Hunt, A. R., & Kingstone, A. (2003). Inhibition of return: Dissociating attentional and oculomotor components. *Journal of Experimental Psychology: Human Perception and Performance, 29*(5), 1068.

Hurst, A. J., Lawrence, M. A., & Klein, R. M. (2019). How Does Spatial Attention Influence the Probability and Fidelity of Colour Perception?. *Vision, 3*(2), 31.

Ignashchenkova, A., Dicke, P. W., Haarmeier, T., & Thier, P. (2004). Neuron-specific contribution of the superior colliculus to overt and covert shifts of attention. *Nature Neuroscience, 7*, 56–64.

Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience, 2*(3), 194-203.

Ivanoff, J., Klein, R. M., & Lupiañez, J. (2002). Inhibition of return interacts with the Simon effect: An omnibus analysis and its implications. *Perception & Psychophysics, 64*, 318–327. doi:10.3758/ BF03195794

Ivanoff, J., & Klein, R. M. (2001). The presence of a nonresponding effector increases inhibition of return. *Psychonomic Bulletin & Review, 8*(2), 307-314.

Ivanoff, J., & Klein, R. M. (2003). Orienting of attention without awareness is affected by measurement-induced attentional control settings. *Journal of Vision, 3*(1), 4-4.

Ivanoff, J., & Klein, R. M. (2004). Stimulus-response probability and inhibition of return. *Psychonomic Bulletin & Review, 11*(3), 542-550.

Ivanoff, J., & Klein, R. M. (2006). Inhibition of return: Sensitivity and criterion as a function of response time. *Journal of Experimental Psychology: Human Perception and Performance, 32*(4), 908.

Kavyani, M., Farsi, A., Abdoli, B., & Klein, R. M. (2017). Using the locus-of-slack logic to determine whether inhibition of return in a cue–target paradigm is delaying early or late stages of processing. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale, 71*(1), 63.

Klein, R. M. (1980). Does oculomotor readiness mediate cognitive control of visual attention? In R. S. Nickerson (Ed.), Attention and performance VIII (pp. 259 –276). Hillsdale, NJ: Erlbaum.

Klein, R. (1988). Inhibitory tagging system facilitates visual search. *Nature, 334*(6181), 430-431.

Klein, R. M. (2000). Inhibition of return. Trends in Cognitive Sciences, 4, 138–147. doi:10.1016/S1364-6613(00)01452-2

Klein, R. M., Christie, J., & Morris, E. P. (2005). Vector averaging of inhibition of return. *Psychonomic Bulletin & Review, 12*(2), 295-300.

Klein, R. M., & Hilchey, M. D. (2011). Oculomotor inhibition of return. In S. Liversedge, I. D. Gilchrist & S. Everling (Eds.), The Oxford handbook of eye movements (pp. 471– 492). Oxford, UK: Oxford University Press.

Klein, R. M., Kavyani, M., Farsi, A., & Lawrence, M. A. (2020). Using the locus of slack logic to determine whether the output form of inhibition of return affects an early or late stage of processing. *Cortex, 122*, 123-130.

Klein, R. M., & Ivanoff, J. (2011). The components of visual attention and the ubiquitous Simon effect. Acta Psychologica, 136(2), 225-234.

Klein, R. M. & Redden, R S. (2018) Two "inhibitions of return" bias orienting differently. In T. Hubbard (Ed.) Spatial Biases in Cognition (pp. 295-306), Cambridge University Press.

Klein, R. M., Schmidt, W. C., & Müller, H. J. (1998). Disinhibition of return: Unnecessary and unlikely. *Perception & Psychophysics, 60(*5), 862-872.

Konak, A., Coit, D. W., & Smith, A. E. (2006). Multi-objective optimization using genetic algorithms: A tutorial. *Reliability Engineering & System Safety, 91*(9), 992-1007.

Krasovskaya, S., & MacInnes, W. J. (2019). Salience Models: A Computational Cognitive Neuroscience Review. *Vision, 3*(4), 56.

Kwak, H.-W.(1992). Inhibitory and facilitatory components of orienting attention to locations and to features. Unpublished doctoral dissertation, Johns Hopkins University.

Langley, L. K., Gayzur, N. D., Saville, A. L., Morlock, S. L., & Bagne, A. G. (2011). Spatial distribution of attentional inhibition is not altered in healthy aging. *Attention, Perception, & Psychophysics, 73*(3), 766-783.

Lawrence, M. A. (2010). Estimating the probability and fidelity of memory. Behavior Research Methods, 42(4), 957-968.

Lawrence, M. A., & Klein, R. M. (2013). Isolating exogenous and endogenous modes of temporal attention. *Journal of Experimental Psychology: General, 142*(2), 560.

Lim, A., Eng, V., Janssen, S. M., & Satel, J. (2018). Sensory adaptation and inhibition of return: Dissociating multiple inhibitory cueing effects. *Experimental Brain Research, 236*(5), 1369-1382.

Lu, C. H., & Proctor, R. W. (1995). The influence of irrelevant location information on performance: A review of the Simon and spatial Stroop effects. *Psychonomic Bulletin & Review, 2*(2), 174-207.

Ludwig, C. J. (2009). Temporal integration of sensory evidence for saccade target selection. *Vision Research, 49*(23), 2764-2773.

Ludwig, C. J., Farrell, S., Ellis, L. A., & Gilchrist, I. D. (2009). The mechanism underlying inhibition of saccadic return. *Cognitive Psychology, 59*(2), 180-202.

Lupiáñez, J. (2010). Inhibition of Return. In A. C. Nobre & J. T. Coull (Eds.), Attention and Time: Oxford University Press.

Lupiañez, J., Klein, R. M., & Bartolomeo, P. (2006). Inhibition of return: Twenty years after. *Cognitive Neuropsychology*, 23, 1003–1014. doi: 10.1080/02643290600588095.

Lupiáñez, J., Martín-Arévalo, E., & Chica, A. B. (2013). Is Inhibition of Return due to attentional disengagement or to a detection cost? The Detection Cost Theory of IOR. *Psicologica: International Journal of Methodology and Experimental Psychology, 34*(2), 221-252.

Lupiáñez, J., Milán, E. G., Tornay, F. J., Madrid, E., & Tudela, P. (1997). Does IOR occur in discrimination tasks? Yes, it does, but later. Perception & Psychophysics, 59(8), 1241-1254.

Lupiáñez, J., Milliken, B., Solano, C., Weaver, B., & Tipper, S. P. (2001). On the strategic modulation of the time course of facilitation and inhibition of return. *The Quarterly Journal of Experimental Psychology: Section A, 54*(3), 753-773.

Lupiañez, J., Ruz, M., Funes, M. J., & Milliken, B. (2007). The manifestation of attentional capture: Facilitation or IOR depending on task demands. *Psychological Research, 71*, 77– 91.

MacInnes, W. J. (2017). Multiple diffusion models to compare saccadic and manual responses for inhibition of return. *Neural Computation, 29*(3), 804-824.

MacInnes, W. J., & Bhatnagar, R. (2018). No supplementary evidence of attention to a spatial cue when saccadic facilitation is absent. *Scientific Reports, 8*(1), 1-13.

MacInnes, W. J., Krüger, H. M., & Hunt, A. R. (2015). Just passing through? Inhibition of return in saccadic sequences. *Quarterly Journal of Experimental Psychology, 68*(2), 402-416.

Mathôt, S., Dalmaijer, E., Grainger, J., & Van der Stigchel, S. (2014). The pupillary light response reflects exogenous attention and inhibition of return. *Journal of Vision, 14*(14), 7-7.

Maylor, E. A. (1985). Facilitatory and inhibitory components of orienting in visual space. Attention and performance XI, 189-204.

Maylor, E. A., & Hockey, R. (1985). Inhibitory component of externally controlled covert orienting in visual space. *Journal of Experimental Psychology: Human Perception and Performance, 11*(6), 777.

McCormick, C. R., Redden, R. S., Lawrence, M. A., & Klein, R. M. (2018). The independence of endogenous and exogenous temporal attention. *Attention, Perception, & Psychophysics, 80*(8), 1885-1891.

Michalczyk, L., & Bielas, J. (2019). The gap effect reduces both manual and saccadic inhibition of return (IOR). *Experimental Brain Research, 237*(7), 1643-1653.

Mirpour, K., Arcizet, F., Ong, W. S., & Bisley, J. W. (2009). Been there, seen that: A neural mechanism for performing efficient visual search. *Journal of Neurophysiology, 102*, 3481–3491.

Müller, H. J., & von Mühlenen, A. (1996). Attentional tracking and inhibition of return in dynamic displays. *Perception & Psychophysics*, 58, 224–249.

Park, S. J., Shin, J. K., & Lee, M. (2002, November). Biologically inspired saliency map model for bottom-up visual attention. In International Workshop on Biologically Motivated Computer Vision (pp. 418-426). Springer, Berlin, Heidelberg.

Pashler, H. E. (1998). The psychology of attention. Cambridge, MA: MIT Press.

Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. In H. Bouma & D. G. Bouwhuis (Eds.), Attention and performance X: Control of language processes (pp. 531–556). Hove, UK: Erlbaum Ltd.

Posner, M. I., Rafal, R. D., Choate, L., & Vaughan, J. (1985). Inhibition of return: Neural basis and function. Cognitive Neuropsychology, 2, 211– 228. doi:10.1080/02643298508252866

Pratt, J., & Neggers, B. (2008). Inhibition of return in single and dual tasks: Examining saccadic, keypress, and pointing responses. *Perception & Psychophysics, 70*(2), 257-265.

Prime, D. J., & Jolicœur, P. (2009). On the relationship between occipital cortex activity and inhibition of return. *Psychophysiology, 46*(6), 1278-1287.

Prime, D. J., & Ward, L. M. (2006). Cortical expressions of inhibition of return. *Brain Research, 1072*(1), 161-174.

Prinzmetal, W., Amiri, H., Allen, K., & Edwards, T. (1998). Phenomenology of attention: I. Color, location, orientation, and spatial frequency. *Journal of Experimental Psychology: Human Perception and Performance, 24*, 261–282.

Prinzmetal, W., Taylor, J. A., Myers, L. B., & Nguyen-Espino, J. (2011). Contingent capture and inhibition of return: A comparison of mechanisms. *Experimental Brain Research, 214*(1), 47.

R Core Team (2015). R: A language and environment for statistical computing. R Foundation for and S4_. R package version 1.1-7, <URL: http://CRAN.R-project.org/package=lme4>.

Rafal, R., Egly, R., & Rhodes, D. (1994). Effects of inhibition of return on voluntary and visually guided saccades. Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale, 48, 284–300. doi:10.1037/1196-1961.48.2.284

Rafal, R. D., Calabresi, P. A., Brennan, C. W., & Sciolto, T. K. (1989). Saccade preparation inhibits reorienting to recently attended locations. *Journal of Experimental Psychology: Human Perception and Performance, 15(*4), 673.

Ratcliff, R. (2008). The EZ diffusion method: Too EZ?. *Psychonomic Bulletin & Review, 15*(6), 1218-1228.

Ratcliff, R., Cherian, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *Journal of Neurophysiology, 90*(3), 1392-1407.

Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Smith, P. L., & Segraves, M. A. (2007). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *Journal of Neurophysiology, 97*(2), 1756-1774.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation, 20*(4), 873-922.

Ratcliff, R., & Van Dongen, H. P. (2011). Diffusion model for one-choice reaction-time tasks and the cognitive effects of sleep deprivation. *Proceedings of the National Academy of Sciences, 108*(27), 11285-11290.

Redden, R. S. (2020, August 31). Diffusion model and two forms of inhibition of return. Retrieved from https://osf.io/3wqtr/?view_only1⁄4c931ff2a0fff418785a319706574f35b.

Redden, R. S., d’Entremont, G., & Klein, R. M. (2017a). Further evidence in favor of prior entry from endogenous attention to a location in space. *Attention, Perception, & Psychophysics, 79*(4), 1027-1038.

Redden, R. S., d'Entremont, G., & Klein, R. M. (2017b). Safe or out: Does the location of attention affect judgments at first base in baseball?. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale, 71*(2), 146.

Redden, R.S., Hilchey, M.D., Aslam, S., Ivanoff, J., & Klein, R.M. (2020). Using speed and accuracy and the Simon effect to explore the output form of inhibition of return. Manuscript under review.

Redden, R. S., Hilchey, M. D., & Klein, R. M. (2016). Peripheral stimuli generate different forms of inhibition of return when participants make prosaccades versus antisaccades to them. *Attention, Perception, & Psychophysics, 78*(8), 2283-2291.

Redden, R. S., Hilchey, M. D., & Klein, R. M. (2018). Oculomotor inhibition of return: Evidence against object-centered representation. *Visual Cognition, 26*(9), 719-733.

Redden, R. S., Klages, J., & Klein, R. M. (2017). The effect of scene removal on inhibition of return in a cue-target task. *Attention, Perception, & Psychophysics, 79*(1), 78-84.

Redden, R. S., & Klein, R. M. (2019). Object-based dissociations for two forms of inhibition of return. https://doi.org/10.17605/OSF.IO/ HPQJX. Retrieved from https://osf.io/hpqjx/.

Redden, R. S., MacInnes, W. J., & Klein, R. M. (2021). Inhibition of return: An information processing theory of its natures and significance. *Cortex, 135*, 30-48.

Rizzolatti, G., Riggio, L., Dascola, I., & Umiltá, C. (1987). Reorienting attention across the horizontal and vertical meridians: evidence in favor of a premotor theory of attention. *Neuropsychologia, 25*(1), 31-40.

Rutherford, H. J., O'Brien, J. L., & Raymond, J. E. (2010). Value associations of irrelevant stimuli modify rapid visual orienting. *Psychonomic Bulletin & Review, 17*(4), 536-542.

Samuel, A. G., & Kat, D. (2003). Inhibition of return: A graphical meta-analysis of its time course and an empirical test of its temporal and spatial properties. *Psychonomic Bulletin & Review, 10*(4), 897-906.

Sapir, A., Soroker, N., Berger, A., & Henik, A. (1999). Inhibition of return in spatial attention: Direct evidence for collicular generation. *Nature Neuroscience, 2*(12), 1053-1054.

Satel, J., Hilchey, M. D., Wang, Z., Story, R., & Klein, R. M. (2013). The effects of ignored versus foveated cues upon inhibition of return: An event-related potential study. *Attention, Perception, & Psychophysics, 75*(1), 29-40.

Satel, J., Wang, Z., Trappenberg, T. P., & Klein, R. M. (2011). Modeling inhibition of return as short-term depression of early sensory input to the superior colliculus. *Vision Research, 51*(9), 987-996.

Schiller, P. H. (1977). The effect of superior colliculus ablation on saccades elicited by cortical stimulation. *Brain Research*.

Şentürk, G., Greenberg, A. S., & Liu, T. (2016). Saccade latency indexes exogenous and endogenous object-based attention. *Attention, Perception, & Psychophysics, 78*(7), 1998-2013.

Sereno, A. B., Lehky, S. R., Patel, S. S., & Peng, X. (2010). A neurophysiological correlate and model of reflexive spatial attention. *Advances in Cognitive Science, 2*, 104-131.

Sereno, A. B., Jeter, C. B., Pariyadath, V., & Briand, K. A. (2006). Dissociating sensory and motor components of inhibition of return. *The Scientific World Journal, 6*, 862-887.

Simon, J. R. (1969). Reactions toward the source of stimulation. Journal of Experimental Psychology, 81(1), 174.

Smith, D. T., Ball, K., Swalwell, R., & Schenk, T. (2016). Object-based attentional facilitation and inhibition are neuropsychologically dissociated. *Neuropsychologia, 80*, 9-16.

Smith, D. T., Rorden, C., & Jackson, S. R. (2004). Exogenous orienting of attention depends upon the ability to execute eye movements. *Current Biology, 14*(9), 792-795.

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences, 27*(3), 161-168.

Snyder, J. J., & Kingstone, A. (2000). Inhibition of return and visual search: How many separate loci are inhibited?. *Perception & Psychophysics, 62*(3), 452-458.

Spence, C. J., & Driver, J. (1994). Covert spatial orienting in audition: Exogenous and endogenous mechanisms facilitate sound localization. *Journal of Experimental Psychology: Human Perception and Performance, 20*(3), 555.

Spence, C., Lloyd, D., McGlone, F., Nicholls, M. E., & Driver, J. (2000). Inhibition of return is supramodal: a demonstration between all possible pairings of vision, touch, and audition. *Experimental Brain Research, 134*(1), 42-48.

Spence, C., & Parise, C. (2010). Prior-entry: A review. *Consciousness and Cognition, 19*, 364–379. doi:10.1016/j.concog.2009.12.001

Spence, C., Shore, D. I., & Klein, R. M. (2001). Multisensory prior entry. *Journal of Experimental Psychology: General, 130*, 799–832. doi: 10.1037/0096-3445.130.4.799

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Acta Psychologica, 30*, 276-315.

Sumner, P., Nachev, P., Vora, N., Husain, M., & Kennard, C. (2004). Distinct cortical and collicular mechanisms of inhibition of return revealed with S cone stimuli. *Current Biology, 14*(24), 2259-2263.

Swalwell, R. (2019). Attentional facilitation for faces in dynamic spatial cueing tasks (Doctoral dissertation, Durham University).

Tas, A. C., Dodd, M. D., & Hollingworth, A. (2012). The role of surface feature continuity in object-based inhibition of return. *Visual cognition*, 20(1), 29-47.

Takeda, Y., & Yagi, A. (2000). Inhibitory tagging in visual search can be found if search stimuli remain visible. *Perception & Psychophysics, 62*(5), 927-934.

Taylor, T. L. (2007). Inhibition of return for expected and unexpected targets. *Acta Psychologica, 124*(3), 257-273.

Taylor, T. L., & Ivanoff, J. (2003). The interplay of stop signal inhibition and inhibition of return. *The Quarterly Journal of Experimental Psychology Section A, 56*(8), 1349-1371.

Taylor, T. L., & Klein, R. M. (2000). Visual and motor effects in inhibition of return. *Journal of Experimental Psychology: Human Perception and Performance, 26*, 1639–1656.

Theeuwes, J., Mathôt, S., & Grainger, J. (2014). Object-centered orienting and IOR. *Attention, Perception, & Psychophysics, 76*(8), 2249-2255.

Tipper, S. P., Driver, J., & Weaver, B. (1991). Short report: Object-centred inhibition of return of visual attention. *The Quarterly Journal of Experimental Psychology*, 43(2), 289-298.

Tipper, S. P., Jordan, H., & Weaver, B. (1999). Scene-based and object-centered inhibition of return: Evidence for dual orienting mechanisms. *Perception & Psychophysics*, 61(1), 50-60.

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

van Rooij, I., & Baggio, G. (2020). Theory before the test: How to build high-verisimilitude explanatory theories in psychological science. *Perspectives on Psychological Science,* 1745691620970604.

Vaughan, J. (1984) Saccades directed at previously attended locations in space. In: Gale AJ, Johnson CW (eds) Theoretical and applied aspects of eye movement research. Elsevier, North Holland, pp. 143-150.

Walther, D., Itti, L., Riesenhuber, M., Poggio, T., & Koch, C. (2002, November). Attentional selection for object recognition—a gentle way. In International workshop on biologically motivated computer vision (pp. 472-479). Springer, Berlin, Heidelberg.

Wang, P., Fuentes, L. J., Vivas, A. B., & Chen, Q. (2013). Behavioral and neural interaction between spatial inhibition of return and the Simon effect. *Frontiers in Human Neuroscience, 7*, 572.

Wang, B., Hilchey, M. D., Cao, X., & Wang, Z. (2014). The spatial distribution of inhibition of return revisited: No difference found between manual and saccadic responses. Neuroscience Letters, 578, 128-132.

Wang, Z., & Klein, R. M. (2012). Focal spatial attention can eliminate inhibition of return. *Psychonomic Bulletin & Review*, 19(3), 462-469.

Wang, Z., Satel, J., Trappenberg, T. P., & Klein, R. M. (2011). Aftereffects of saccades explored in a dynamic neural field model of the superior colliculus.

Wang, B., Yan, C., Klein, R. M., & Wang, Z. (2018). Inhibition of return revisited: Localized inhibition on top of a pervasive bias. *Psychonomic Bulletin & Review, 25*(5), 1861-1867.

Weaver, B., Lupiáñez, J., & Watson, F. L. (1998). The effects of practice on object-based, location-based, and static-display inhibition of return. *Perception & Psychophysics*, 60, 993– 1003.

Weber, M. D., Leemis, L. M., & Kincaid, R. K. (2006). Minimum Kolmogorov–Smirnov test statistic parameter estimates. *Journal of Statistical Computation and Simulation, 76*(3), 195-206.

Wickelgren, W. (1977). Speed–accuracy tradeoff and information processing dynamics. Acta Psychologica, 41, 67–85.

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature, 453*, 233–235.

Zhao, Y., Heinke, D., Ivanoff, J., Klein, R. M., & Humphreys, G. W. (2011). Two components in IOR: evidence for response bias and perceptual processing delays using the SAT methodology. Attention, Perception, & Psychophysics,73(7), 2143-2159.