A Behavioural Perspective of Visual Search and the Low Prevalence Effect

A thesis

submitted in fulfilment

of the requirements for the degree

of

Doctor of Philosophy in Psychology

at

The University of Waikato

by

David M. Hollands

2024

Abstract

In visual search, rare items are missed disproportionately often. This Low Prevalence Effect (LPE) is a robust phenomenon with important societal consequences. We asked the question, “Is the schedule of signal presentation controlling eye movements and if so, how?” and hypothesised that visual search is an operant behaviour. To answer this question, we examined the LPE using eye-tracking and incorporated a rich schedule of signal presentation (medium-prevalence condition: 0.50 target probability across 200 trials) and a lean schedule (low-prevalence condition: 0.02 or 0.00 target probability across 1,000 trials) over three experiments. Experiment One was the control, in Experiments Two and Three we increased task difficulty by incorporating a staircase titration procedure, and in Experiment Three we also removed all target-present trials in the low-prevalence condition. We replicated the LPE in all experiments and observed increased levels of searching eye movements throughout the medium-prevalence condition due to an increase in behaviour for the first correct rejection immediately following a correct target-present response which we attribute to the local effects of reinforcement and the strengthening of search behaviour when we find what we are looking for. We also observed that visual attention was being guided towards the densest population of stimulus items (centroid) at the beginning of their search as a search strategy and conclude that this centroid zone is acting as a discriminative stimulus signalling the availability of reinforcement. We contend that eye movements are an operant behaviour controlled by environmental contingencies and reinforcement mechanisms are behind the LPE phenomenon.

Dedication

This thesis is dedicated to the loving memory of my late father

Michael Thomas Hollands

13th December 1939 – 30th August 1992

Those we love don't go away

They walk beside us every day

Acknowledgements

A PhD thesis is something that not only encompasses a significant portion of one's life but also those surrounding people without whom this would not have been possible. First and foremost, I would like to sincerely thank my friends and family for their undying love and support throughout this endeavour, I could not have endured without you. A special thanks also goes to my chief supervisor Professor John Perrone, your patience and support have been inspirational, and I could not have wished for a better mentor. Thank you also for helping me learn the foreign language that is MatLab coding, and a special thanks to my good friend Vance Farrow who walked me through the linguistics of this undertaking. To the Psychology technician Rob Bakker, thank you for helping make my dreams a reality. I must have driven you all insane!

Contents

List of Figures

List of Tables

Introduction

At first glance, the mundane task of searching for one's coffee cup on a kitchen bench may appear relatively easy. However, in reality, this seemingly simple visual task is an extremely complex set of behaviours. It involves the allocation and guidance of visual attention through eye movements via the oculomotor system, the visual processing of task-relevant information throughout the search task, memory templates for the target object, and learning and memory traces for the prioritisation of possible target locations. The search task also includes various task-relevant motor operations, such as head turning, body rotation, reaching and grasping, etc. The linchpin behind all these complex behaviours and processes is the visual system. As a result, this system has been extremely well studied, as one may learn a lot about these processes from an understanding of how the visual search system operates. However, there are still some questions which need to be answered. For example, we still do not have a full understanding of the consistency and complexity of eye-movement behaviour throughout the search process and how these behavioural patterns change over time, as a function of task difficulty (do you scan for your coffee cup differently when the bench is cluttered compared to when it is relatively clear?) and target schedule modulation (at what point do you abandon the search and go to the cupboard for a fresh coffee cup?). This thesis is aimed at filling in these missing parts.

1.1 Eye-movement behaviour

The observation of eye-movement behaviour has been a field of interest dating back to Aristotle more than 2000 years ago. For Aristotle, the fundamental features of eye movements were binocular and he theorised that the eyes were yoked together.

For most of recorded history, the eyes were thought to glide smoothly over scenes within the environment and focus on objects of interest, where they would be fixated with unmoved accuracy. It was not until the nineteenth century that the term saccade was first introduced to eye-movement research by Emile Javal, with saccade being the French word for “jerk” or “twitch”. These rapid, almost ballistic, eye movements were described and illustrated by Brown and Grant (1895). Brown provided an understanding that the eyes moved in a series of saccades which vary in size, and that visual perception only occurs during the brief intervals between these saccades, termed a fixation. Given the limited resolution of our peripheral vision, these rapid eye movements are necessary to bring the high-resolution fovea to bear upon items of interest within the environment, allowing us to examine in detail those objects. The development of objective eye-movement recording technology in the mid-twentieth century, illuminated characteristics of eye-movement behaviour which had, until this point in time, remained hidden from view, such as the small involuntary eye-movements which accompany a fixation (microsaccades) or the allocation of attention to multiple stimuli during a single fixation (covert attention)(Gegenfurtner, 2016).

The importance of our high-resolution fixation spot is due to the visual anatomy of the human eye and the fact that the neural areas for visual stimulation are not homogeneous. The distribution of cone photoreceptors for daylight vision has a sharp peak close to the centre of the retina, termed the fovea. Cone density in the fovea is extremely high, up to 200,000 cones/mm°, and as low as 2,000 at larger eccentricities, termed the visual periphery. The area of the fovea is extremely small, only about 50,000 of the six million cones in each eye are positioned there. The foveal area covers approximately 2 degrees of visual angle, which approximately equates to the size of a human thumbnail at arm’s length. Only within this tiny portion of the visual field can we achieve high-acuity resolution for such tasks as reading fine print or threading a needle.

Cone density is a limiting factor for visual acuity, and acuity decreases rapidly towards the visual periphery proportionally to cone density. The reason why we can cope with such a tiny area of high-acuity resolution is that we can move our foveal region to the things we want to investigate by making very precise voluntary eye movements (Wurtz, 2008).

Saccades direct the fovea sequentially to peripheral objects of interest by single, rapid, step-like eye movements. Saccades are easily identified from a record of eye position samples because they have a distinct velocity profile, such that the eye rapidly accelerates to a peak velocity of approximately 500 degrees per second. The disadvantage of eye movements is that the correspondences between retinal and environmental coordinates are dissociated (Gegenfurtner, 2016). For saccades, retinal positions get assigned to new positions within the visual environment and receptive fields in the visual cortex receive new input information.

Vision immediately before and during a saccade is extremely poor, due to a combination of factors (Holt, 1903; Volkmann, 1962). The retinal image gets blurred due to the high velocity of eye rotation during the saccade (Burr, Morrone, & Ross, 1994). Furthermore, the retinal image from the landing position masks the previous retinal image both before and during the saccade (Breitmeyer & Ganz, 1976; Diamond, Ross, & Morrone, 2000), and there is also an additional suppression mechanism which inhibits information which would otherwise be visible during the planning stage of an upcoming saccade before the eye has moved (Burr et al., 1994). As a result, the processing of visual information takes place largely during fixations, where the eye is relatively motionless. It is therefore assumed that the location and duration of fixations reflect what is being processed at a given moment in time. This assumption relies on an intricate relationship between overt attention (involving eye movement) and covert attention (not involving eye movement) but neglects the fact that saccades take at least 100ms to prepare, with larger saccades taking longer. This means that at least part of the time during a fixation is devoted to "planning" the upcoming saccade. However, it generally makes sense to talk about attention and fixation as synonymous, consistent with the Active Vision approach (Findlay & Gilchrist, 2003).

During fixation periods our visual system uses the limited information presented to recognise what it is observing. For complex objects, the problem of recognition may appear as some kind of visual jigsaw with missing and disappearing pieces. However, recognising an object is not enough, information about the object needs to be transferred into some form of visual memory. This was a problem that was first addressed by George Sperling (1960), who found that there is a visual representation with high capacity which can last for several 100ms. As long as the visual target image can be processed in an undisturbed manner for at least 70ms, we can recognise the image and transfer the corresponding visual information into visual memory (Gegenfurtner & Sperling, 1993).

While patterns of eye movements appear to be idiosyncratic, they are not random. Tatler and Vincent (2009) demonstrated that we can predict fixation locations by knowing how the eyes generally move. This is because participants show systematic biases towards certain locations and saccade patterns regardless of the image. Therefore, it is important to investigate how these biases are related to image context and why not if they are not systematically related.

When we examine a static image within a frame, we typically start with a fixation at, or towards, the centre of the frame, followed by exploratory saccades to various locations and objects of interest within that image. This is referred to as the "centre bias" phenomenon (Tatler & Vincent, 2009). This phenomenon may reflect the eyes' "orbital reserve", but is often masked by the practice of cueing participants with a central fixation cross at the beginning of each trial. Artists and photographers have long known of this tendency and often place objects and items of interest in the centre of the frame.

Saccades are also biased, occurring more often in horizontal directions. This is true even in square images and follows the perception of the layout when the scene is rotated (Foulsham, Kingstone, & Underwood, 2008). Another tendency is that we often execute saccades which land in locations between display items, termed the "centre-of-gravity" phenomenon (He & Kowler, 1989), with initial eye movements demonstrating greater centre-of-gravity behaviour and later saccades landing closer towards, or on, the target item. This pattern of behaviour demonstrates a coarse-to-fine strategy of eye-movement behaviour (Zelinsky, 2008).

The above findings were observed in laboratory settings when viewing static images. In the real world, we typically interact with the visual environment for much longer periods in a specific, informative surrounding. In recent times, it has been shown convincingly that eye movements go far beyond pure perception. It appears that as soon as we act in our environment the oculomotor system is engaged in providing information that guides our actions. The visual system is well adapted to visualize objects within a single fixation, even though that time is extremely brief. In the real world, we look around and make many fixations before arriving at a judgment decision. This is because we need more fixations to build up a picture of scenes and the objects within those scenes if we want to transfer this information into visual memory (M. Castelhano & Henderson, 2005; Hollingworth & Henderson, 2002; Huebner & Gegenfurtner, 2010). The world around us may be useful as an external memory but this only works as long as we are within a scene. If we move to the room next door, we can still remember things about the previous room, and this information representation is built up gradually over many fixations. Oculomotor behaviour constitutes a fundamental feature of visual search and the everyday exploration of our environment. Therefore, it is the visual search paradigm to which we now turn.

1.2 Visual Search

Long before Aristotle was theorising about eye-movement behaviour, humankind must have pondered over their visual abilities from the time they first looked out over the savannah, in search of food, shelter, or social interaction. Unsurprisingly, due to its ubiquitous nature in our everyday lives, visual search remains a topic of interest across varying fields of study. Adding to this interest is the extremely large body of visual search literature, incorporating psychologists, visual researchers, biologists, and neuroscientists. It may be said that the early interest in the visual search paradigm has arisen due to its ease, under certain circumstances in providing an understanding of how properties of our visual cortical areas may be used to explain complex perception without resorting to higher-order psychological or neurophysiological mechanisms.

We engage in visual search because the environment presents the visual system with an overwhelming array of information at any one point in time and we can only process a small proportion of this information presented (J. Duncan & Humphreys, 1989). To deal with this excess of input, the visual system involves cognitive mechanisms which prioritise a small subset of information for in-depth analysis while relegating the rest to only limited processing (Wolfe & Horowitz, 2004). In this respect, visual search involves sequential eye movements which bring the high-resolution fovea to locations of interest selected by peripheral vision (Wiecek, Pasquale, Fiser, Dakin, & Bex, 2012). For humans, and many animals, the light falling on the foveal region of the retina receives preferential analysis leading to higher spatial acuity (Wertheim, 1894), vernier acuity (Levi, Klein, & Aitsebaomo, 1985), and contrast sensitivity (Rovamo, Leinonen, Laurinen, & Virsu, 1984). These benefits are due to the higher density of foveal photoreceptors when compared to the visual periphery, and the relatively higher number of neurons in the primary visual cortex and other areas dedicated to the fovea (Daniel & Whitteridge, 1961; R. O. Duncan & Boynton, 2003).

The study of visual search provides an understanding of how the human brain coordinates a variety of functions involving oculomotor control through fixational eye movements (overt attention), covert attention (not involving eye movements), differences in visual processing across retinal locations, temporal integration of information across eye-movements as search progresses, memory for scene configurations, and decision strategies (Eckstein, 2011; Nakayama & Martini, 2011; Wolfe, 1998). It may be said that visual search is in itself a manifestation of attention due to our limited processing capacity. The concept of visual attention is not new (Estes & Taylor, 1964; Shiffrin & Gardner, 1972), however in the past four decades visual search studies have been driven by an effort to understand the processes behind visual attention and the attributes that guide the allocation of this attention. Below we discuss some theories of visual attention and their defining components.

1.2.1 Classic Visual Search

In the visual search laboratory, early search experiments involved the participant being presented with a display containing several items (set-size) randomly presented across a blank display. In each trial, the participant must determine if a specific target item is, or is not, present amongst a number of non-target items (distractors). These items were often artificial, so that the nature of visual differences between the target and distractor items may be precisely controlled, with the hope that this artificial array would transpose across to naturalistic search environments. In these early experiments, stimuli were often simple alphanumeric characters presented in orderly arrays (B. F. Green & Anderson, 1956; Neisser & Beller, 1965). However, alphanumeric characters are not simple stimuli, they are complex shapes, so when interest began to focus on the nature of visual attention processes, classic search displays became arrays of even simpler shapes (e.g., coloured bars and the like), often presented in quasi-random arrays or on a circle at a fixed distance from fixation. Much of this early research was conducted with stimuli that were deliberately made large enough and spaced widely enough so that acuity and crowding limits would not constrain the task.

There are typically two behavioural measures of interest in classic search tasks: response latency or reaction time, and detection accuracy. Accuracy measures convert search tasks into well-controlled, two-alternative forced-choice (2AFC) psychophysical experiments, making them amenable to analysis with the tools of signal detection theory (SDT), which break the response down into hits, misses, false alarms, and correct rejections. These responses may then be calculated as measures of sensitivity (an index of signal detectability) and the decision criterion (an index of the decision response bias). One drawback has been that these artificial laboratory methods and abstract stimuli are often quite removed from real-world search environments and experiences. However, these laboratory methods are imperative in shedding light on the intricate phases of the visual process during search and provide an understanding of how we allocate visual attention during this process, which would be extremely difficult, if not impossible, within a real-world setting.

In a typical latency study, stimuli are presented until the participant responds. The number of items in the display is varied from trial to trial, and the latency x set-size function is analysed (A. M. Treisman & Gelade, 1980). Accuracy in these tasks is tracked, but usually only to ensure the error rates are not so high that they would call the latencies into question. This is because, as error rates increase in these tasks, response latencies decline: a simple “speed-accuracy trade-off” (Heitz, 2014; Henmon, 1911) which may distort the shape of the latency x set-size function. In the classic analysis of these functions, the slope of the function provides the most interesting information about search; it is a measure of the rate at which items are processed. Whereas the intercept is a measure of the time required for non-search processes (e.g., the physical act of making a response) (Posner, Snyder, & Davidson, 1980). If all items can be processed in parallel without capacity limitations, the slope of the latency x set-size function would approximate zero ms/item. If items are processed in series, one item at a time, then the latency would increase linearly with the number of items in the display.

In the simplest case of a serial, self-terminating search, if items are sampled randomly, participants would on average, sample half of the items before finding the target on target-present trials. If target-absent trials required an exhaustive search through all items, the slope for target-absent trials would approximately be twice that for target-present trials, a pattern which has frequently been reported (A. M. Treisman & Gelade, 1980). Although a parallel, infinite-capacity search would indeed produce fairly flat slopes and a serial, self-terminating search would produce linear slopes in an approximate 2:1, absent/present ratio, it does not follow that this pattern of results proves this theoretical account of search. Different processes may produce similar patterns of results (Townsend, 1976; Townsend & Wenger, 2004). Furthermore, these functions are not always linear (Pashler, 1987) and, as noted, speed-accuracy trade-offs may complicate the analysis, as participants' error rates tend to increase with larger set-sizes, artificially curving the function and lowering the slopes (Dukewich & Klein, 2009). Given these concerns regarding parallel and serial search processes, researchers tend to use theory-neutral terminology, such as "efficiency" to describe the meaning of a search function. It should be noted, that the slopes of search tasks form a continuum from very efficient to very inefficient (Egeth, Virzi, & Garbart, 1984; Nakayama & Silverman, 1986; Wolfe, 1998). It would be poor practice to try to strictly define terms of efficiency using precise slope values, as there is no categorical boundary between efficient and inefficient search labels; search efficiency is a continuous measure. Furthermore, slopes tend to be quite variable, and it would be unhelpful to declare that one participant was categorically different from another based on a small slope difference. Today, it is generally accepted that both parallel and serial search processes may be activated simultaneously throughout the search process.

It was Treisman's original conceptualisation that search tasks could be divided into two classes: parallel searches with slopes approximating zero, indicating that all items may be processed in parallel and serial tasks with steeper slopes, indicating that items are selected in series at a rate of approximately 25-50 items/second (A. M. Treisman & Gelade, 1980). Early studies found that search was generally efficient when the target differed from all its distractors along a singleton feature dimension (e.g., colour, shape, orientation). This kind of search is termed a feature search. A second type of search involves a target that differs from the distractors in terms of a combination of two or more feature dimensions so that the target is not defined by any single feature. This kind of search is termed a conjunctive search and is fundamentally inefficient (A. Treisman & Sato, 1990; Wolfe, 1998). With this in mind, we now turn to the classic models of visual search and attention.

1.2.2 The Feature-integration Theory

The term "efficient" was initially proposed to argue that in feature searches, all items are processed in parallel because pre-attentive processing (not requiring attention) was sufficient for detecting a singleton feature, whereas all items in a conjunctive search are processed in series ("inefficient") because focal attention is required for identifying feature conjunctions. This concept of efficiency was at the heart of Treisman's "Feature Integration Theory" (A. M. Treisman & Gelade, 1980). At the core of this theory was the concept of distinct feature maps, two-dimensional arrays of detectors, any one of which could be activated in parallel. For example, there were feature maps for colour (red, green, blue, etc.), basic shapes (letters, geometric figures, etc.), and other dimensional properties. Search could occur in two modes, pre-attentive and attentive. In the pre-attentive mode, the activation of a singleton feature in the metaphorical feature map (e.g., a red bar amongst an array of green bars) would effortlessly "pop-out" to the participant by being the sole locus of activation in the red primitive feature map. Pop-out in a given search array occurs when performance (in terms of response latency) remains constant, no matter how many distractor items are presented in the array (latency function approximating zero). According to this theory, this occurs because pop-out is mediated by the unique occurrence of a feature in a retinotopic map, in which all items are processed in parallel. Therefore, there would be no cost in adding more distractor items to the array. Treisman claimed that these relatively flat latency functions were a diagnostic characteristic of parallel processing and could point to the existence of a retinotopic map for a given feature or dimension.

Differing from this form of search, are searches for conjunctions of features, in which parallel processing of a singleton feature map was not possible. Therefore, conjunctive searches could not be handled by the parallel operation of a single feature map and pre-attentive vision could not occur. Treisman postulated that attention was needed to bind features at some other more central locus and the slope of the search function or search rate (time per item) represented the amount of time it took focal attention to move from one item to another in a serial fashion. Moreover, the theory predicted that for target-absent trials, the overall search time would be twice that for target-present trials. This required the assumption that participants searched the array with full memory so that they would not be wasting time with repeated attentional visits to an already inspected item.

A defining feature of this theory was its compatibility with the search asymmetry phenomenon (Dosher, Han, & Lu, 2010; A. Treisman & Sato, 1990). In search asymmetry, search performance changes when the target and distractors change their roles. For example, searching for a Q amongst Os is efficient and insensitive to the number of distractors presented, whereas searching for an O amongst Qs is inefficient. According to Treisman, this was because efficient search is done by monitoring the presence of a feature (in this case the “tail” which exists on a Q) and not by its absence. In this respect, search asymmetry may serve as a test for the basic feature status of a given visual stimulus.

Unfortunately, as time passed, cracks began to form in the foundations of this theory. To account for criticism related to uncontrolled target-distractor differences and distractor-distractor similarity (J. Duncan & Humphreys, 1989), FIT underwent later refinements, in that visual attention was also needed for difficult feature searches (A. Treisman, 1991) and a feature inhibition mechanism was also proposed to account for the efficiency of highly discriminable conjunction searches (A. Treisman & Sato, 1990). Another stumbling block of FIT was it did not clarify how pre-attentive processing supports the localisation of a feature. It had become clear that when we detect a pop-out target, we are generally able to localise the target and it attracts focal attention (Nothdurft, 1999), suggesting that some pre-attentive mechanism must guide the deployment of focal attention. However, it was this theory’s strict dichotomy between parallel and serial processes, which was to become the downfall of the Feature Integration Theory.

1.2.3 The Guided Search Model

Wolfe, Cave, and Franzel (1989) modified Treisman and Gelade's Feature Integration Theory and introduced the concept of attentional guidance to a series of Guided Search (GS) Models, arguing that basic features could be processed in parallel and then used to guide the serial deployment of attention. The original GS model retains FIT's covert attention concept as a serial processor. However, unlike FIT, attention is not randomly assigned to items across the array but is guided by parallel processing across the visual field to elements in the display. A basic tenet of this model is a proposal that there is a pre-attentive mechanism that can provide information about the visual scene, so that focal attention may be intelligently allocated and guided.

The Guided Search 2.0 model (Wolfe, 1994) proposes that this guidance starts with the stochastic processing of items to produce a master map of locations, which represents the priority of attention allocated to each location. This master map receives inputs from feature maps, and the model proposes that these feature maps incorporate volitional goal-directed top-down (endogenous) attentional activation components related to the task relevance of the stimuli, or automatic stimulus-driven (exogenous) activation components related to the physical salience (feature contrast) of the stimuli. Items are then further processed if they exceed an activation threshold. Attention is then deployed serially from one item to another, starting with the item which has the highest activation and then progressing to the next item with the second highest activation and so forth. Importantly, if serial attention processes an item, it can accurately determine whether that item is the target or not. After a certain search time, the model quits searching, and if the serial attention processes have not identified the target, it guesses, allowing for false alarm decision errors. Increasing the target-distractor discriminability increases the probability that the target would be ranked highly in terms of activation, and thus, serial attention would reach the target more efficiently, reducing response latency.

In Guided Search 4.0 (Wolfe & Gray, 2007) the model includes a drift-diffusion process and noise which replaces error-free attentional guidance. In this respect, serial attention processes one item at a time, which then enters the diffusion process, in which evidence accumulates in a stochastic fashion for each item. This allows the model to generate errors, even if attention processes the item and allows for false alarm decisions in target-absent trials without resorting to guessing, as in GS 2.0. Multiple items may be incorporated in the diffusion process at any one time if it does not exceed a certain capacity limit. When the number of items exceeds this limit, an item has to be rejected before another can enter the diffusion process.

This version of the Guided Search Model has been successful because it accommodates variation in search performance. In this model, a search which can be accomplished based on an individual feature will be efficient because the target is salient (bottom-up activation) and contains target features (top-down activation). This combination would trigger a large master map activation and attention would be allocated with a high priority towards the target item. Conjunctive search may also be considered fairly efficient because the target contains more target features, resulting in more top-down activation. Search is least efficient for unknown identity conjunctive search, which is guided by neither top-down nor bottom-up activation.

In Guided Search 6.0 (Wolfe, 2021) the model was expanded to include the guiding factors of not only top-down and bottom-up processes, but also history, reward, and scene structure. These guiding sources are combined into a spatial "priority map" which evolves over the course of the search and selective attention is then guided to the most active locations on this map. The model states that guidance is not uniform across the visual field but favours items near the point of fixation and names three types of functional visual field (FVF) which describe the nature of these foveal biases; namely, a resolution FVF, an FVF governing eye-movements, and a FVF governing the covert deployment of attention. To identify a target or reject a distractor, items must be compared to a target template held in memory. As with GS2.0, this revised model retains a diffusion process allowing multiple items to undergo recognition simultaneously, although asynchronously. If a target is not found, the search terminates when an accumulating quitting threshold is reached. This threshold is adaptive, allowing feedback regarding performance which shapes subsequent searches.

We now expand on GS6.0's five factors which guide the deployment of visual attention, as well as discuss the concept of a priority map in some detail.

1.3 Attentional Visual Guidance

1.3.1 Bottom-up Attentional Guidance

Bottom-up guidance, also called exogenous or stimulus-driven control, refers to the deployment of attention prioritised by the physical salience of a stimulus relative to its immediate surroundings and this determines how likely that a stimulus captures attention (Itti & Koch, 2001; Theeuwes, 1991, 1992, 1994, 2010). In this respect, the salience of a stimulus increases with differences from its surrounding distractor stimuli (target-distractor heterogeneity) and with the homogeneity of the distractor items (distractor-distractor homogeneity) and is usually defined in terms of low-level feature dimensions, such as colour, orientation, luminance, size, or motion (Wolfe & Horowitz, 2004, 2017). Moreover, attentional capture may also occur via less obvious stimulus properties, such as lighting direction (Adams, 2008) or axis rotation (Schill, Cain, Josephs, & Wolfe, 2020). This attentional capture occurs quickly, and effortlessly, and is thought to occur independently of one's intentions or goals (involuntarily).

Following from the target-distractor heterogeneity theory, search efficiency is dependent upon the number of features shared by the target stimulus (Nordfang & Wolfe, 2014), and attention may be guided to multiple target features simultaneously (Friedman-Hill & Wolfe, 1995). However, Yu and Geng (2019) contends that when a target is close to its distractors in feature space, attention is guided by an internal representation or target template which is not centred on the actual target features. Instead, the template would be centred slightly away from distractors to avoid mistakenly categorising a distractor as the target. In signal detection terms, this would be equivalent to moving the decision criterion to a more conservative position to avoid false alarm errors.

Yantis and Jonides (1984) demonstrated that the abrupt onset of a new perceptual item captures attention even when the onset does not reliably predict the target location. Theeuwes (1991) demonstrated that for feature searches, the addition of a singleton colour distractor (task-irrelevant feature) slowed search times for a shape target by capturing attention, suggesting that bottom-up guidance influenced search priority, and was not subject to top-down control. However, a shape singleton did not slow a colour search down (Theeuwes, 1992). Assuming the colour singleton had a greater saliency than the shape singleton, this suggests that search is prioritised according to levels of visual saliency. Theeuwes (1994) showed that the interference effect caused by the irrelevant distractor remained present even after 1800 trials. This suggests that even extensive practice could not induce sufficient top-down control to overcome the interference caused by the salient distractor. Theeuwes concluded that for feature search, bottom-up attentional guidance was dominant and automatic, which may override top-down control.

1.3.2 Top-down Attentional Guidance

Top-down guidance, also termed endogenous or goal-driven attentional control, is driven by perceptual goals. Due to the very nature of visual search, it could be said that all search behaviour may be categorised as goal-oriented. When an individual searches for a particular item, feature, or location they can voluntarily direct overt or covert attention to the task-relevant item, feature, or location (Posner et al., 1980; Theeuwes, 2010; Wolfe & Gray, 2007; Yantis & Egeth, 1999). According to this view, known as the contingent capture hypothesis (Folk, Remington, & Johnston, 1992) selection depends critically on the explicit or implicit perceptual goals held by the individual at any given time. Since top-down processes are under volitional control, the individual must intentionally activate this control form, which may take time and effort. However, once this deployment of attention has been activated it has been shown to increase response speed and accuracy (Pashler, Johnston, & Ruthruff, 2001).

Folk, Remington, and Wright (1994) demonstrated this form of guidance using a spatial cueing procedure, in which a cue display was followed, in rapid succession, by a target display. They found that only when the search display was preceded by a cue which matched the singleton feature for which participants were searching, the cue captured attention. These results suggest that the top-down attentional set determines selection priority. When set for a particular feature singleton, the participant would select each element which matches this top-down set, feature singletons which did not match the top-down attentional sets would simply be ignored. Folk et al. concluded that the attentional readiness adopted by the participant determined selection.

The dichotomy between stimulus-driven and goal-driven attentional capture represents opposing viewpoints; stimulus-driven capture proposes that capture is always bottom-up, whereas contingent capture proposes that capture is always dependent on top-down processes. A viewpoint which is somewhere in the middle has been presented by Yantis and Egeth (1999) who contend that only abrupt onsets can capture attention in a truly bottom-up fashion, with other properties being dependent on top-down settings. In this respect, selection is under top-down control and a feature singleton (such as a unique colour or brightness) is not automatically selected. Only when the element is presented with abrupt onset, constituting a new object, it receives attentional priority. Whatever your viewpoint, it seems plausible that both top-down and bottom-up processes work together at any one point in time to guide the deployment of attention.

1.3.3 Scene Structure and Meaning

Scene structure and meaning are when attributes of the scene guide attention to areas more likely to contain the target item or object, irrespective of whether that target and/or its features are present. For example, when searching for milk, attention is often directed towards items inside the refrigerator, whether or not there is any pre-attentive feature evidence for its presence there. We certainly would not start our search by looking in the oven!

Scene guidance is not a new concept (Biederman, Glass, & Stacy, 1973; Kingsley, 1932), but has taken on greater prominence in recent times, in part because the study of visual search has largely been built upon searching for targets within arbitrary 2D arrays of items, whereas most real-world search tasks take place in structured scenes, and this structure provides a source of guidance. Extending the study of guidance from controlled arrays of distinct artificial items to structured scenes poses some difficulties. For example, how do we define the set size of a naturalistic scene? To bridge this gap, maybe the best solution is to talk about the number of items/locations which may be treated as candidate targets in a scene given a specific task or “effective set-size” (Neider & Zelinsky, 2008; Rensink, 2000). While the effective set size may not be a perfect measure, it remains clear that the effective set size for most tasks would be much smaller than the set of all possible items (Wolfe, Alvarez, Rosenholtz, Kuzmova, & Sherman, 2011).

When a scene is first encountered, scene guidance appears coarse and evolves during the search (Ehinger, Hidalgo-Sotelo, Torralba, & Oliva, 2009). Brief exposure to a scene (50-75 ms) is enough to guide the deployment of eye movements once a search begins (M. S. Castelhano & Krzyś, 2020). Importantly, this preview need not contain the target to be effective (Hollingworth, 2009). Search appears to be strongly guided by an initial understanding of the scene which does not rely on recognising specific objects (Greene & Oliva, 2009). Knowledge about the target provides an independent source of guidance (M. S. Castelhano & Heaven, 2011; Malcolm & Henderson, 2010). These sources provide useful information on where the target might be, which is more powerful than memory for where a target might have been seen in the past, in terms of search guidance (Torralba, Oliva, Castelhano, & Henderson, 2006; Le-Hoa Võ & Wolfe, 2015). This preview is limited to the first couple of fixations (Hillstrom, Scholey, Liversedge, & Benson, 2012). Once the search begins guidance is updated, in that the content of the current fixation in a scene influences the location of the next fixation, rendering the preview obsolete (Hwang, Wang, & Pomplun, 2011).

While feature guidance appears to be based upon the output of a fast, preattentive process, scene guidance must be learned. This learning evolves from the initial information about the spatial environment of the scene, to guidance based on a rich understanding of the scene, informed by attention to objects within that scene (Oliva & Torralba, 2001). Contextual cueing (Chun & Jiang, 1998) may be interpreted as a basic form of scene guidance in which repeated exposure to the same random arrays of arbitrary stimuli leads to learning of the association between the random array and the target location. This association appears to be largely implicit, even in situations where real-world stimuli are presented (Y. V. Jiang, Won, & Swallow, 2014).

1.3.4 Prior History of Search Over Time

Learning from the participant's prior history of search has also been shown to modulate the guidance of attention. As such, search performance is influenced by performance on previous searches and these effects may be organised across their timescales. Within a trial, when half of the search array is briefly presented a few hundred milliseconds before the rest of the array, performance was enhanced by the reduction of the set size, either because attention was guided away from the "marked" preview items (D. G. Watson & Humphreys, 1997) and/or towards the latter new items (Donk & Theeuwes, 2003).

On a longer timescale, the priming phenomenon may also be observed from trial to trial within an experiment. In a colour pop-out task, Maljkovic and Nakayama (1994) found that, even though the search for the colour target was highly efficient, participants were faster if the preceding target was the same as the current target. Performance was not enhanced if the participant was informed of the colour of the next target; rather, priming was based on the colour of the previous target. These results have been replicated and expanded to multiple feature dimensions (Hillstrom, 2000; Kristjansson, 2008) and a greater priming effect has been found if the task was more difficult (Lamy, Zivony, & Yashar, 2011). This priming by target features takes approximately 200ms to develop (Wolfe, Horowitz, Kenner, Hyle, & Vasan, 2004) and priming for conjunctive stimuli may last longer than for singleton feature pop-out priming (Kruijne & Meeter, 2015). Distractor properties may also be primed (Lamy, Yashar, & Ruderman, 2013), however, these effects were smaller than those found for target stimuli (Wolfe, Butcher, Lee, & Hyle, 2003). Importantly, the take-home point here is that by simply repeating the target dimensions from trial to trial enables this phenomenon to occur (Wolfe et al., 2003).

1.3.5 Value Modulated Attentional Capture

The value placed upon a search item is an extremely strong modulator of voluntary attention allocation (Wolfe & Horowitz, 2017). For example, if you reinforce a feature (e.g., red) and/or punish another (e.g., green), items with reinforced features will attract more attention and items with punished features will attract less attention (Anderson, Laurent, Yantis, & Smith, 2011). Sceptics may argue that the effect of reinforcement is to speed responses once the target has been found, and not to guide attention towards the target stimulus. However, Lee and Shomstein (2013) varied item set sizes and found that reinforcement learning could make reaction time slopes shallower which indicates that reinforcement had its effects on the search process and not just on the response once a target had been found. Moreover, the reinforcing effects of reward may be measured using real-world scenes (Hickey, Kaiser, & Peelen, 2015), an indication that value may be a factor in everyday search. We cover this topic in more detail in the Search Behaviour section 1.7.

1.4 Priority Maps

The term priority map has often been used as a simple catch-all concept encompassing some mechanism responsible for the allocation of attention in visual space. However, Zelinsky and Bisley (2015) argue that the function of a priority map is largely one of a filter for reducing the amount of incoming sensory information for processing, by selecting goal-relevant information in visual space, to control and coordinate activity across different effector systems for the achievement of some behavioural goal. In this respect, Zelinsky and Bisley define a priority map as a neural representation of topographic space, in which activity codes for the priority of locations within that space contribute to this prioritisation, irrespective of the bottom-up or top-down factors (p.155).

In this view, both parallel and serial search behaviours may be explained as products of a common mechanism, in which the locations of goal-modulated feature information compete on a priority map. Converging evidence suggests that this common mechanism constitutes the selective weighting of the low-level features, matching the features representing a goal state (Chelazzi, Duncan, Miller, & Desimone, 1998; Desimone & Duncan, 1995; Motter, 1993; Navalpakkam & Itti, 2005). From a physiological perspective, this may be implemented by changing synaptic weightings on pathways, such that the features corresponding to the goal are given greater weight. As such, features across low-level retinotopic feature maps would give priority to the locations of these features in downstream priority maps. Target locations would be disproportionally weighted on these maps relative to distractors which do not have these goal features, resulting in an effortless, so-called, parallel search. To the extent that distractors share these goal features, their locations would also be prioritised, resulting in a lower signal-to-noise ratio and less efficient search. In the extreme, the priority features of the target and distractors may not be discriminable, producing unguided overt and covert attentional movements, resulting in serial search (Zelinsky & Bisley, 2015).

Importantly, the neural weighting of features in visual space is organised across retinotopic space, reflecting the degree to which patterns falling within the receptive fields of retinotopically organised neurons match the features of the goal state. In this respect, Zelinsky and Bisley contend that features and locations are intricately bound together by activity over a priority map and are reflected in the motor system, with features being bound to achieve some behavioural goal, such as an upcoming saccade. In this view, if an eye movement is to be made to a visual stimulus at a given location on the retina, this same location may be used to coordinate the motor maps used by the oculomotor system, resulting in the quick and largely effortless generation of saccades.

The concept of a filter is not new to the study of attention, due to the need to reduce the amount of sensory information which must be processed by higher-level neural mechanisms (Broadbent, 1957). Importantly, the function of an information filter is essentially one of priority control. The visual system receives an overwhelming amount of information at any point in time, so some of this information must be filtered out (A. M. Treisman, 1960), giving priority to the information allowed through for processing. These filters may be set in both feature space (feature-based attention) and topographic space (spatial attention) (Gray & Wedderburn, 1960). In this respect, a priority map is the neural weighting of feature importance at each location in the topographic space, and priority control is the competitive process through which the most prominent locations are selected.

Priority maps may also function in the fluid coordination and control of actionable behaviour by interfacing with motor systems. The features of a goal are selected and identified in visual space, providing a point or region in this space to mediate a behavioural action via multiple visual-to-motor projections to the effector systems which enable various interactions with our environment. In this sense, the metaphor of "shifting attention" is just navigating the priority map by shifting or moving effector systems to the next most prioritised point on the map. Attention is simply the ability to prioritise information, and an attentional state would be the prioritising pattern of neuronal firing at any one point in time, expressed in multiple motor systems. For example, eye movements are simply the expression of this prioritised pattern by the oculomotor system; reaching, grasping, or head rotation would be expressions of other motor systems.

This prioritised pattern of behaviour suggests a relationship between the goal state maintaining the priority map and visual working memory (VWM). In this view, VWM embodies the goal state from which priority maps are constructed by changing the filter settings for information entering the priority map, effectively biasing responses on these maps and weighting those features and locations for satisfying the goal state. In visual search, a VWM representation of a search template would be compared to incoming visual information, to create a target map. The activation landscape of this map may then be navigated, in order of priority to make confident search decisions (Zelinsky, 2008).

VWM may also be important in sustaining the goal state for the upkeep of active and ongoing goal-directed behaviour over time. Disruptions are fleeting behavioural events which interrupt the current goal state and hence seize control of the ongoing behaviour, by changing the priority map. In these moments, the current goal-oriented behaviours stop and only resume again when the goal state in working memory can restore the previous sequence of priority settings needed to reinstate the original goal state for a behavioural task. Zelinsky and Bisley suggest that a central function of working memory is not only to re-establish a goal state following a disruption but also to prevent these disruptions from gaining priority in the first place. In the absence of the goal-maintaining function of working memory, new and salient patterns in the continuous stream of sensory information would gain priority and seize behavioural control, making the fluid execution of a goal-directed behaviour difficult, if not impossible.

Visual search offers some idea into the problems which may arise when a goal state in VWM is not correctly reflected on a priority map owing to some goal degradation or interference. When a participant searches for two or more targets simultaneously in a visual display, the VWM goal state must either be broadened to include features from both targets. This would lower the signal-to-noise ratio on the priority map, or the participant may "gamble" and choose to represent only one of the targets in VWM. This would result in strong search guidance for features or spatial locations on the corresponding priority map when the target representation on that map was chosen, but weak guidance when this gamble does not pay off. This conceptualisation may help explain the decrement in performance accuracy seen in multi-target tasks (Beck, Hollingworth, & Luck, 2012).

1.5 Memory

What has become clear, is that memory and learning are an integral component in the modulation of attentional guidance throughout visual search. Prior knowledge regarding the physical characteristics of a target is arguably an important factor for efficient search and will often improve detection performance. Conversely, uncertainty about the attributes of a target may slow search times and lead to search errors (Eckstein, 2011). The concept of a perfect search memory was an underlying assumption for Treisman and Gelade's Feature-Integration Theory (1980) to explain why search latencies tend to increase linearly with set sizes and why target-absent and target-present trials have an approximate 2:1 latency ratio for serial search. Therefore, it came as quite a surprise when Horowitz and Wolfe (1998) published a manuscript titled "Visual search has no memory." In this article, the authors examined serial search performance while randomly shuffling display items every 111 ms, so that spatial memory for rejected distractors was made null and void. They showed that search efficiency was not reduced with this manipulation and concluded that the search may not depend on memory at all. This caused quite a stir and inspired a flurry of research. As a result, several replications of this study have shown that the search was made inefficient by shuffling display items (Müller & Mühlenen, 2000; Kristjánsson, 2000). Hulleman (2010) suggests this discrepancy in the literature may be due to differences in search difficulty, in that more difficult search tasks (>80 ms/item) may be more influenced by shuffling. Furthermore, a simple tendency to avoid resampling from a just visited location, termed the “Inhibition of return” (IOR) phenomenon (R. Klein, 1988; R. M. Klein & MacInnes, 1999), can closely mimic the performance of an ideal observer (Najemnik & Geisler, 2005), suggesting that what is needed for optimal search performance may be some form of memory.

The existence of the IOR phenomenon demonstrates that attention avoids revisiting a previously examined location. However, this is not the whole story. Eye-movement studies have shown that, in visual search, participants do revisit already inspected items, and that the majority of these revisits typically occur immediately after the eyes first leave an already inspected item (Wienrich, Heße, & Müller-Plath, 2009). These so-called “*Lag-2* revisits” are unlikely to be the result of the participant having forgotten that they had just examined an item, rather it is generally accepted that these revisits occur because a fixation may be prematurely terminated before the item had been fully processed. Furthermore, the brief fixation on item n+1 is generally brief because it would be terminated rapidly to allow the participant’s search system to revisit an item, suggesting that the search system made a corrective saccade. For the most part, these *lag-2* revisits occur for the target item, further suggesting that they are wilful revisits, rather than random errors of memory or errors in saccadic trajectory (Godwin, Reichle, & Menneer, 2017; Rayner, 1998).

Within the timeframe of an experiment, a method for investigating the role of search memory is by adopting a dual-task methodology. In this procedure, the rationale has been to fill a participant’s working memory to capacity while simultaneously administering a search task, so that, if the search requires memory, performance should be reduced. This was what Oh and Kim (2004) found when participants had to keep several spatial locations in working memory (also see Woodman & Luck, 2004)(Woodman & Luck, 2004).

Also, within this timeframe, if participants remembered when they had previously searched for a target within the same context before, the memory of that context should improve performance. Chun and Jiang (1998) first investigated whether contextual memory could facilitate search by repeating some search displays within an experimental session. They found that the search was more efficient if the same display had been presented previously; indicating that participants were capable of associating certain display configurations with target locations. This phenomenon has been referred to as "Contextual cueing". It has been suggested that participants are not explicitly remembering each display; rather they are extracting local and holistic information from the display configurations and associating this information with target locations (Y. Jiang & Wagner, 2004). It has also been demonstrated that contextual cueing is not just limited to display configurations but also to display item shapes and movements (Y. Jiang, Chun, & Olson, 2004), and naturalistic scenes (Brockmole, Castelhano, & Henderson, 2006; Brockmole & Henderson, 2006).

Importantly, Chun and Nakayama (2000) contend that visual processes retain very little conscious information across time. They propose that visual information retention is accomplished by a set of independent "implicit memory" mechanisms that allow specific information from scenes to persist across image changes and over time. These memory traces guide the use of attention efficiently. These mechanisms are not necessarily under our conscious control and we do not have explicit access to the underlying context of our visual representations which guide the deployment of visual attention.

1.6 The Low Prevalence Effect

During laboratory visual search experiments, participants typically search for a given target within an array of distractor items, and target prevalence refers to the proportion of trials containing the target stimulus. When target prevalence is relatively low (<0.3 target probability), participants are less likely to detect a target compared to higher target prevalence search conditions (0.5 probability). This robust phenomenon has been referred to as the "Prevalence Effect" or "Low Prevalence Effect" (Godwin, Menneer, Cave, Helman, et al., 2010; Godwin, Menneer, Cave, & Donnelly, 2010; Menneer, Donnelly, Godwin, & Cave, 2010; Van Wert, Horowitz, & Wolfe, 2009; Wolfe et al., 2007; Wolfe, Horowitz, & Kenner, 2005). The effects of prevalence have typically been examined using behavioural measures, including response accuracy and latency. These measures have also been supplemented with signal detection theory (SDT) indices relating to sensitivity (an index of target detectability) and the decision criterion (an index of response bias) (Macmillan & Creelman, 2004).

The prevalence effect is not a new observation. Perhaps the largest body of literature related to the prevalence effect comes from the vigilance research domain in the 1960's. These early psychophysical studies demonstrated how participants changed their pattern of response decisions based on the probability of a target signal occurring (originally referred to as signal frequency) (Broadbent & Gregory, 1965; Colquhoun, 1961; D. M. Green, Swets, & others, 1966; Tanner, Swets, Green, & Macnee, 1956). These studies indicated that for simple detection tasks, the effect of reducing target prevalence would be to make the decision criterion more conservative without altering the sensitivity to the target signal across varying schedules of signal frequency (D. M. Green et al., 1966). In other words, you would be less likely to call something a target if the a priori probability of a target was low. For an extensive summary of the time see Mackworth (1970), who stated that "…one of the most important findings in vigilance research has been the discovery that the probability of a signal being detected is considerably reduced when the number of non-target events is increased over a period of time." (p.60).

However, Wolfe et al. (2007) contend that vigilance tasks differ in subtle ways from search tasks. Specifically, most vigilance tasks require the participant to detect "a faint and infrequent signal at an uncertain time" (Broadbent & Gregory, 1965). In contrast, search stimuli typically remain present until the participant classifies the display as containing a target or not, meaning that a transient lapse in attention would not generate an error. Secondly, vigilance tasks generally present stimuli one at a time, whereas search tasks, by definition, involve searching for a target among simultaneously displayed distractor items. Finally, vigilance decrements appear as time-on-task increases. While there are interesting theoretical similarities between vigilance decrements and prevalence effects in search, Wolfe et al. contend that they are not the same phenomenon, as these similarities may arise from fundamental aspects of human decision processes that play out in different ways in different settings.

1.6.1 Medical Imaging Research

More recently, this phenomenon was brought back to the fore within the medical imaging context by Kundel (1982) who published a paper titled "Disease prevalence and radiological decision making," a retrospective study that centered around the diagnostic utility of a radiological imaging system. Signal detection theory (SDT) and receiver operating characteristic (ROC) analysis were introduced into medical imaging from psychology in 1960 (Lusted & Ledley, 1960) and radiologists started using ROC-derived measures of signal sensitivity, the area under the ROC curve, to analyse the performance of imaging systems. ROC curves are obtained by plotting the fraction of true positives in overall positive cases (hit rate) against the fraction of false positives in overall negative cases (false alarm rate). As a result, the ROC curve is insensitive to the relative number of positive and negative cases, or in a medical context, the prevalence of the disease in question. Kundel emphasised that the diagnostic utility of an imaging system required taking prevalence into account and proposed a measure to achieve this: the positive predictive value (PPV), which is the ratio of hits to total positive responses (hits + false alarms). Kundel (1982) incorporated six retrospective medical performance studies to illustrate his point. By design, all six studies had similar d' scores (another widely used SDT index of sensitivity) but varied dramatically in PPV values due to varying prevalence.

In a follow-up study, Kundel (2000) reported a more comprehensive set of field studies to address the question of whether disease prevalence influenced detection performance. In this data set, sensitivity, as measured by d', appeared to be inversely proportional to prevalence. Unfortunately, the main point of contention with these retrospective radiological studies which Kundel (1982, 2000) used was that each data point obtained represented a different patient population, disease, and medical observer population. The studies incorporated by Kundel were conducted between 1947 (Birkelo et al., 1947) and 1998 (Henschke et al., 1999) at facilities in the United States, Czechoslovakia, and Japan. As a result, there could be many confounding factors besides prevalence varying across these studies which may be influencing performance. Kundel (2000) concluded that: "The fundamental question of what would happen if the same group of readers read a high and a low prevalence test set containing the same cancer cases is still unanswered. The slopes of the ROC curves must be known to make a valid comparison. This issue could be settled by assembling a test set of cancer cases with matched cancer-free cases from a screening study in which confidence levels were reported." (Kundel, 2000, p.142).

Prior to Kundel's (2000) investigation, Egglin and Feinstein (1996) embarked on a fairly small medical imaging study involving six faculty angiographers interpreting arteriograms for pulmonary emboli with either 0.2 or 0.6 prevalence probabilities (40 cases per condition). This study found lower sensitivity for the lower prevalence condition. Following on from this study, Ethell and Manning (2001) asked novice radiographers to look for fractures in wrist radiograph images at 0.22, 0.50, or 0.83 prevalence probabilities (72 cases per condition). They found that sensitivity was highest at 0.5 prevalence and that performance was reduced for both the high and low prevalence conditions. Both these studies suggest small but measurable, decrements in performance under low prevalence conditions. However, both these investigations employed a relatively small number of cases and image readers.

To answer the question posed by Kundel (2000), Gur et al. (2003) embarked upon an ambitious investigation into prevalence within a medical context. This study involved 14 radiologists, comprising eight faculty, two fellows, and four residents. Gur et al. employed 1,632 posteroanterior chest images and measured performance across five prevalence levels ranging from 0.02 to 0.28 probabilities. Instead of reporting whether an abnormality was present or absent, they reported their confidence that there was an abnormality present on a 1-100 scale. These confidence ratings were used to calculate sensitivity using the area under the ROC curve (Az). They showed that while there were differences in detectability across different lesions (interstitial disease was difficult to detect, pneumothorax relatively simple), they did not find any appreciable differences in sensitivity as a function of prevalence. However, they did find differences in confidence ratings, with increased confidence that the target was present in lower prevalence conditions.

On the face of it, the Gur et al. study appeared to provide a conclusive answer to Kundel's (2000) question: When the same group of readers read the same cases at different prevalence levels, sensitivity to the target stimulus does not change. These results suggested that testing imaging systems and training novice radiologists under high prevalence conditions would be reflective of what clinicians would experience in the clinic under low prevalence conditions. Unfortunately, this study did not report an SDT measure for response bias, such as the decision criterion.

1.6.2 The Criterion Shift Hypothesis

In 2005, psychology re-entered the fray when Wolfe, Horowitz, and Kenner conducted a laboratory study investigating the effect of prevalence on visual search. This study examined the question of how low prevalence conditions might affect search performance in complex visual displays within the context of an airport baggage X-ray screening simulation using greyscale, semi-transparent, photorealistic images of common items, superimposed, with various transparency levels, over a background of sinusoidal noise to mimic some of the visual properties of X-ray images; as well as, some of the conceptual properties of a screening task, such as searching for a category of items.

This study found that when the target (a tool) was presented with a trial probability of 0.5, participants made few "miss" errors (7% of target-present trials). In contrast, when the target was presented with a trial probability of 0.01, the miss error rate increased substantially (30% of target-present trials). Moreover, a curious pattern emerged in the response latency data (reaction times). As mentioned in the previous visual search chapter, mean response latencies for target-absent trials are generally longer than latencies for target-present trials, approximating a 2:1 ratio. This ratio suggests that in a serial, self-terminating search task with 0.5 prevalence probability, participants appear to examine twice as many items to respond "no, the target is not present" than to respond "yes, the target is present." At 0.5 prevalence, Wolfe et al. (2005) replicated this established response pattern. However, during the 0.01 prevalence condition, this pattern reversed, so that participants were faster to report the target as absent rather than present. Surprisingly, in the low prevalence condition, mean latencies for correct target-absent trials were less than the mean latency for correct target-present trials, suggesting that there may be some form of execution errors influencing response patterns, namely response priming or a simple speed-accuracy trade-off.

One difficulty with this study was that participants did not produce many false alarm errors, making it difficult to estimate SDT parameters for sensitivity and the decision criterion. These measures are important for both theorising about the effects of prevalence and for making comparisons with applied contexts. Therefore, in most subsequent prevalence experiments, Wolfe and colleagues increased the fidelity of the display simulation. This methodology comprised of X-ray images of empty luggage and various component items which could then be superimposed within the luggage. This realistic simulation produced false alarm errors and allowed for SDT analysis (Wolfe et al., 2007).

Wolfe and Van Wert (2010) used this methodology to investigate the effects of very high prevalence. This study showed that at a very high prevalence (0.98 prevalence probability), miss errors reduced to almost nothing relative to a 0.5 prevalence condition, while false alarms increased from 18% at 0.5 prevalence to 58% at 0.98 prevalence. Target-present latencies were similar across both prevalence conditions. However, target-absent latencies increased substantially in the very high prevalence condition. As both false alarms and miss errors were produced in this task, the SDT indices for sensitivity (*d'*) and the decision criterion (*c*) were able to be computed. These measures showed that sensitivity was similar across prevalence conditions, whereas the decision criterion, which was roughly neutral at 0.5 prevalence, was markedly more liberal (i.e., negative) during the 0.98 prevalence condition.

In a subsequent experiment, Wolfe and Van Wert (2010) generalized these findings across the complete spectrum of prevalence values by varying prevalence sinusoidally over 1,000 trials, starting at a 1.0 prevalence probability, systematically reducing down to a 0.0 prevalence level, before systematically returning up to 1.0 prevalence. Unsurprisingly, this procedure demonstrated that false alarm and miss errors traded off, with high false alarm rates at high prevalence and high miss rates at low prevalence. Target-present latencies were relatively insensitive to prevalence, while target-absent responses varied markedly, with shorter latencies at low prevalence and increasing under high prevalence conditions. The SDT index for sensitivity remained fairly constant throughout the procedure, while the criterion varied systematically with prevalence, reaching a conservative (positive) high at approximately the 0.0 prevalence level and converting to a liberal (negative) zenith at approximately the 1.0 prevalence level. Of interest was the finding that the sinusoidal points of the criterion function lagged behind the underlying prevalence function by approximately 50 trials, suggesting that participants may be interpreting and integrating prevalence information over this timeframe.

Returning to the medical imaging context, Evans, Tambouret, Evered, Wilbur, and Wolfe (2011) asked 10 cytologists in the United States and 12 in the United Kingdom to read 2,100 Papanicolaou (Pap) test images under medium or low prevalence conditions. In the medium condition, prevalence was set at a 0.5 probability, while in the U.S. the low prevalence condition was set at a 0.02 probability, and in the U.K. a 0.05 probability. In the low prevalence condition, miss errors increased and false alarm errors decreased compared to error rates in the medium prevalence condition. This finding equated to a significant conservative shift in the decision criterion under low prevalence conditions, while sensitivity did not change significantly.

In an ambitious second study, Evans, Birdwell, and Wolfe (2013) inserted 100 mammography images, 50 of which were known to be positive cancer cases, into the normal workload of a radiology clinic over 9 months (low prevalence condition). These cases were presented with false header information so that the radiologists did not know that a given case was associated with the study. Following the conclusion of the low prevalence arm of the study, six of the 14 clinic radiologists were asked to read the same 100 cases in the laboratory. As a result, this study was able to compare the same images at a 0.5 probability to the typical prevalence rates found within a clinical setting (approximately 0.003 probability). These radiologists missed more cancer cases in the low-prevalence clinical setting and made more false alarm errors in the medium-prevalence laboratory setting, equating to a conservative decision criterion shift under low prevalence conditions.

Adding further weight to the criterion shift hypothesis was a meta-analysis study conducted by Horowitz (2017) which focused on the SDT parameters of sensitivity (d') and the decision criterion (c). This analysis incorporated five medical imaging studies involving 12 experiments and eight psychology laboratory studies comprising 25 experiments. Low prevalence probabilities ranged between

0.02 and 0.3, with a mean prevalence of 0.065; while the medium prevalence sample ranged between 0.4 and 0.6, with a mean of 0.492.

While some medical studies showed improved sensitivity at low prevalence and some the opposite, overall there was little evidence for an effect of prevalence on sensitivity, consistent with Gur et al. (2003). The corresponding data from the psychology studies showed the same general pattern of results or lack of one. The criterion shift account, however, predicts that the primary effect of prevalence is on the decision criterion, rather than sensitivity. Accordingly, within the medical context, this is what was found, with the decision criterion becoming more liberal as prevalence transitioned from low to medium. The psychology data showed a similar pattern of results. This meta-analysis did not analyse response latency data as these results are rarely recorded within medical imaging studies.

To summarise, this pattern of findings whereby sensitivity remains fairly constant as a function of prevalence, while the decision criterion becomes increasingly conservative as prevalence decreases, constitutes the standard prevalence effect. In the error domain, miss rates increase with a corresponding decrease in false alarm rates under low prevalence conditions, and the inverse occurs under high prevalence conditions. In the response latency domain, there is a strong modulation of target-absent latencies with the respect that “no” responses are faster than “yes” responses under low prevalence conditions, and considerably slower at high prevalence, while target-present latencies are relatively insensitive to prevalence variations.

There has also been an interest in the literature to determine whether the prevalence of a target stimulus changes search behaviour in fundamental ways aside from the decision criterion, in that the excessive increase in miss errors under low prevalence conditions may reflect three basic error types (Kundel, Nodine, & Carmody, 1978). Specifically, motor execution errors, where the participant recognises the target but makes the wrong response; early search termination errors, where the participant decides that no target is present before having seen the target; and perceptual errors, where the participant has seen the target but does not recognise it. In all likelihood, under certain circumstances, all three of these error types may be involved in the increased rates of miss errors found under low prevalence conditions.

1.6.3 The Motor Prepotency Hypothesis

Increased miss errors under low prevalence conditions may result from the artificial environment within laboratory settings. Participants are generally instructed to respond “as quickly and accurately as possible.” If the search task is relatively easy, the participant may rapidly respond “no” repeatedly over long periods of time under low prevalence conditions. Fleck and Mitroff (2007) demonstrated that when participants were given the opportunity to correct their previous response, without the stimuli present, most of the additional low prevalence miss errors were eliminated. This led them to propose the “motor prepotency hypothesis,” which states that miss errors under low prevalence conditions may be attributed to the “no” response being strengthened by repetition, so that even when the participant intended to respond “yes” the “no” response prevails.

Adding weight to this hypothesis, Rich et al. (2008) tested the prevalence effect in an identification task. In this procedure, targets were Ts amongst L distractors, and participants were required to identify whether the Ts were upright or upside-down. In the medium prevalence condition orientations were equally proportioned; while in the low prevalence condition, one target only was presented with a probability of 0.02. When the low prevalence target was an upright T, participants were more likely to report an inverted T compared the medium prevalence condition. Furthermore, Rich et al. observed that correct response latencies for the rare target were longer than correct latencies for the common target. They concluded that this finding was inconsistent with a strategy of quitting the search early and guessing, because with this strategy, the correct rare target responses would only be made when the rare target happened to be identified fairly early in the search. Rich et al. interpreted this finding as likely to reflect these errors as motor perseverance of the type proposed by Fleck and Mitroff (2007).

However, these motor perseverance errors may not account for all of the additional miss errors under low prevalence conditions. In 2008, Rich et al. used eye tracking data to identify what types of errors participants were making. For example, on trials where the participant did not fixate the target, it may be assumed that they terminated searching prematurely. While trials where the participant fixated the target but continued searching may reflect a failure in recognising the target, which would support the criterion shift hypothesis. Finally, on trials where the target was fixated shortly before a “no” response may reflect motor perseverance errors. However, these errors of judgment may not be reflective of the whole story. A participant may not fixate the target, but process it in their peripheral vision and fail to recognise it; similarly, they might respond “no” just after fixating the target, but fail to recognise it, etc. However, if we accept the validity of these miss error assumptions, Rich et al.’s (2008) eye-movement data allows for an estimate of these differing error type distributions as a function of prevalence. This study employed a spatial configuration task, whereby participants searched for rotated Ts amongst rotated Ls at either 0.5 or 0.02 target probabilities. They replicated the standard prevalence effect and observed that the additional miss errors under low prevalence conditions were primarily the result of premature search termination (failing to fixate the target).

Godwin, Menneer, Riggs, Cave, and Donnelly (2015) used a relative prevalence dual-target search procedure with eye tracking to investigate perceptual failures in the selection (failing to fixate the target) and the identification (fixating the target but failing to identify it) of low prevalence targets. In this procedure overall prevalence was held at a probability of 0.5 to minimise motor errors caused by repetitive target-absent responses. Participants searched for two targets, one of which appeared with a probability of 0.45 on target-present trials and the other appeared on 0.05 of these trials. In the first experiment, participants searched for dissimilar targets, and in the second experiment participants searched for two similar targets, to minimise dual-target cost effects. In contrast to Rich et al. (2008), they showed that the primary influence of perceptual errors under low prevalence conditions were errors of identification, rather than selection.

Adding weight to Godwin et al.’s (2015) findings was a study by Hout, Walenchok, Goldinger, and Wolfe (2015) who employed both a relative prevalence task and a rapid serial visual presentation (RSVP) task, in which participants responded following the display sequence ending, to control for both motor errors and early search termination errors. Eye-tracking analysis showed that high prevalence targets evoked greater attentional guidance and faster perceptual decisions, and that when participants directly fixated low prevalence targets they failed to identify it on 12-34% of trials.

The contribution of motor errors may depend on the type of search task employed. Kunar, Rich, and Wolfe (2010) split a spatial configuration search display in half, presenting half the display 1,000 ms before the other half, in an attempt to slow responses and reduce motor errors. However, the prevalence effect persisted for both whole and split displays, and still persisted when participants were given the opportunity to correct their previous response (also see Russel and Kunar, 2012). Wolfe et al. (2007), similarly found that issuing participants with "speeding tickets" when they were responding too quickly, slowed responses but the prevalence effect remained.

It appears plausible that the contribution of motor errors to the prevalence effect may be reflective of the search task; in this respect, when the search task is relatively easy, participants may find themselves rhythmically responding "no" in quick succession. Therefore, while it seems likely that prepotent motor responses may contribute to the prevalence effect in some experimental environments, it should be considered one of several factors by which low prevalence induces miss errors in visual search, rather than the primary cause.

1.6.4 Expectation and Experience

Research has illustrated that participants are quite good at estimating the real prevalence of a search task, and that prevalence expectations are developed over a window of approximately 20-50 trials (Ishibashi, Kita, & Wolfe, 2012; Wolfe et al., 2007; Wolfe & Van Wert, 2010). When prevalence is reduced, the target appears less often within the sequence of trials, which results in a reduction in bottom-up priming for the target item, influencing both the speed and accuracy of target detection (Kristjansson, 2008). Over time, participants may also form a top-down expectation of how likely the target will occur on any given trial. Prevalence studies examining the contributions of top-down and bottom-up information have used various forms of cues and feedback to modulate top-down expectations, while holding bottom-up priming constant. Studies which provide participants feedback after each trial, such as text information regarding their accuracy, have shown a strong effect; so much so, that they are able to swing the decision criterion within 50 trials (Wolfe et al., 2007; Wolfe & Van Wert, 2010). However, studies which have provided pre-trial cues regarding the upcoming target prevalence have found weaker effects from this top-down information (Ishibashi et al., 2012; Lau & Huang, 2010). This Indicates that this pre-trial information may be competing with a participant's internal representation of prevalence which has accumulated over many trials and stored in memory (Kristjansson, 2008).

J. D. Schwark, Macdonald, Sandry, and Dolgov (2013) demonstrated that when perceptual decisions are difficult, participants are inclined to make prevalence-based decisions. In this study, participants searched for letter stimuli in cluttered arrays with as many as 300 distractor letters (increasing to 700 in later studies). In these displays, letters often overlap, and the juxtapositions of letter fragments could create ambiguous targets. Participants were given three response options: Respond by clicking on the actual target, by clicking on a target-absent button, or by clicking on the target-present button. The latter response option would be indicative of a prevalence-based decision, as opposed to a perceptual decision; in that, participants were reasonably confident that the target was present but could not perfectly identify it. They found that participants used the target-present button more frequently when the search was difficult, as indexed by the number of distractor letters, which suggests that they were relying on cognition when they could not rely on perception.

In a similar study, J. Schwark, Sandry, Macdonald, and Dolgov (2012) used false feedback to manipulate the participants' perception of miss errors. When participants were led to believe that they were committing many miss errors, they shifted their decision criterion and began committing more false alarm errors. In both these studies, when participants could not adequately rely on perception, they relied instead on the memory of their search history for prior search outcomes (J. Schwark, Sandry, & Dolgov, 2013)(see also Schwark, Sandry, & Dolgov, 2013). We now turn to a model of search behaviour, which we will discuss in some detail, as it has specific relevance to the prevalence effect and also eye-movement behaviour, the Multiple Decision Model (Wolfe & Van Wert, 2010).

1.6.5 The Multiple Decision Model

The prevalence effect does not appear to be the result of the participants' vigilance declining under low prevalence conditions, as their ability to detect the target (sensitivity) remains fairly constant as a function of prevalence. However, as the participant's decision criterion becomes more conservative under low prevalence conditions, they require more evidence to respond "yes" to the target is present. Wolfe and Van Wert (2010) built on the Guided Search Model to propose a dual-threshold model, termed the Multiple Decision Model (MDM) based on this hypothesis, to explain changes in both error and response latency as a function of prevalence. At the heart of this model is the assumption that the search array constitutes a discrete set of items and the process of visual search may be conceptualised as a series of decisions.

One type of decision involves item identification. In this respect, items are attentionally selected one after the other, and the participant must decide whether the current item under investigation is a target or not. This identification decision is modelled according to SDT parameters for a two-alternative forced-choice decision, and the participant must make this decision based upon some "activation threshold," drawn from either the target or distractor distributions. These distributions overlap to some extent so the decision process is noisy. The participant sets a criterion threshold, so that if the activation exceeds this value, the item will be identified as the target, otherwise the item will be classified as a distractor. This is the first decision within the model.

If the identification decision is "target," the participant then responds "yes" and the trial is terminated. However, if the decision is "distractor" there is a second decision to be made; specifically, continue searching or terminate the search by responding "no." This second decision to terminate or continue is modelled as a drift-diffusion process (Palmer, Huk, & Shadlen, 2005; Ratcliff, 2006). As with the first decision, an activation threshold is generated. Under the drift-diffusion process, this activation threshold begins with a value of zero and "drifts" upwards until a decision threshold limit is reached. This upward drift does not increase incrementally, but over time will eventually reach the threshold limit and the search will be terminated. This limit may be set by the participant at the lower end of the scale, resulting in quick “no” responses, or set high, resulting in slow “no” responses. This termination threshold represents the second decision within the model. This cycle, identify and then terminate, or not, is repeated until either the identification process results in a “yes” response or the termination process reaches its limit and results in a “no” response.

Prevalence acts simultaneously on both thresholds. The identification threshold becomes more liberal under high prevalence conditions, so that it takes less evidence to convince the participant that an item is a target when targets are frequent; and more conservative under low prevalence conditions, so that it requires more evidence when targets are rare. The termination threshold is increased under high prevalence conditions, resulting in participants being more reluctant to terminate the search when targets are frequent and decreases under low prevalence conditions so that participants are more likely to terminate the trial prematurely when targets are rare.

Shifting from a medium or high prevalence condition to a low prevalence condition will result in the identification threshold becoming more conservative and the termination threshold becoming lower. As a result, targets are more likely to be missed because some targets will be mistakenly identified as a distractor, and because the participant will be more likely to terminate the search before the target has been identified. The decreased termination threshold also explains why correct target-absent response latencies are faster under low prevalence conditions, as the participant terminates the search after observing a small proportion of the display items. Most of the time, this strategy pays off because targets are rare. Hit latencies are generally impervious to this mechanism because hits are only generated when the target is identified, and this takes approximately the same amount of time regardless of prevalence.

Shifting from a low or medium prevalence to a high prevalence, miss errors become rarer because participants will spend increased time searching the display due to the higher termination threshold, meaning that they are more likely to identify the target if it is there. This results in longer target-absent response latencies. They are also more likely to make false alarm errors due to their reduced identification threshold. Hit responses may become contaminated with some “lucky guesses,” meaning that participants would incorrectly indicate an item as the target before locating the actual target so that hit latencies are slightly reduced. Most of the time, this "lucky guess" strategy pays off because the target stimuli are relatively frequent.

To summarise, this model illustrates the effects of prevalence on both accuracy and response latencies. Miss errors increase under low prevalence conditions because the identification decision threshold is set at a high level, meaning that only highly recognisable targets are identified, and false alarms are rare. The second threshold regarding the search termination decision, explains why prevalence only affects target-absent response latencies and not target-present latencies. The time taken to find a target is determined by the perceptual difficulty of the task, which does not change as a function of prevalence. However, the decision to terminate the search is made more quickly under low prevalence conditions, and more slowly under high prevalence conditions. As task difficulty increases, the target and distractor distributions become increasingly overlapped which decreases the signal-to-noise ratio, meaning that this model may easily generalise to more complex search tasks, such as those in which the number of potential targets are greater than one, or the task involves a high degree of uncertainty, allowing this model to provide a useful framework for understanding how prevalence affects search behaviour.

1.7 Search Behaviour

Reinforcement has long been known to influence behaviour. As early as 1911, Thorndike proposed the "Law of Effect", which states that behaviours soliciting a reward will become strengthened and more likely to occur again in that situation over time. Conversely, behaviours which solicit a discomforting effect would become less likely to occur again in that situation (Thorndike, 1911, p.244). His seminal work on learning theory led to the development of operant conditioning theory within behaviourism. Following these revelations, Pavlov (1927) observed that stimuli reliably predicting a reward would eventually elicit a behavioural response. This learning process created a conditioned response through associations between an unconditioned stimulus and a neutral stimulus, termed classical or Pavlovian conditioning. In this respect, classical conditioning is learning through the development of associations between events, whereas operant conditioning involves learning from the consequences of behaviours. In the laboratory and the natural environment, we (animals including humans) tend to perform behaviours in a ratio which matches the ratio of available reinforcement for those behaviours according to "The Matching Law" (Herrnstein, 1961).

While the behavioural processes of learning were initially the focus of behaviourism, later researchers were trying to understand what internal states accompanied the strengthening of response patterns due to reinforcement (Simon, 1967; Berridge & Robinson, 1998; Mackintosh, 1975; Pearce & Hall, 1980). In regards to the guidance of visual attention, it has become evident that similar principles as those formulated in classic behavioural learning theories may apply to how reinforcement influences memory and the allocation of visual attention (Boehler, Schevernels, Hopf, Stoppel, & Krebs, 2014; Patai, Doallo, & Nobre, 2012; Braem, Verguts, Roggeman, & Notebaert, 2012). These studies provide evidence for the influence of reinforcement on executive functioning, cognitive control, and the allocation of attention, which has been suggested to influence the control of top-down attentional allocation and have a beneficial effect on perceptual and attentional processes (Kiss, Driver, & Eimer, 2009; Sawaki, Luck, & Raymond, 2015; Seitz, Kim, & Watanabe, 2009; Serences, 2008; Botvinick & Braver, 2015). However, the findings that reinforcement manipulations were congruent with task demands do not necessarily mean that top-down attentional control was the source of attentional selection. Studies employing an inter-trial reinforcement priming procedure strongly suggest that reinforcement history independently influences the perception of stimulus features on the following trial, irrespective of whether that feature was associated with the target or reward payout, which cannot be fully explained in terms of top-down activation processes (Della Libera & Chelazzi, 2009; Hickey, Chelazzi, & Theeuwes, 2011, 2010).

In a conditioning/test procedure, Della Libera and Chelazzi (2009) demonstrated that reinforcement-based attentional capture, as shown by response latencies, could persist for up to five days, even when the follow-up experimental session was no longer associated with reinforcement. Moreover, this attentional capture may even be observed over a considerably longer time frame, up to six months after the initial reinforcement-based contingency session (Anderson & Yantis, 2013).

Anderson and Yantis (2013) argue that the increased response latency was due to attentional capture by the previously associated reinforced stimulus, which cannot be explained in terms of either top-down or bottom-up processes. Hickey et al. (2011) suggest that reinforcement modulations only impacted search when they were tied to the target stimulus, as reinforcement appeared to make it no easier to ignore a distractor when it subsequently reappeared as a distractor. However, Theeuwes and Belopolsky (2012) have demonstrated that a task-irrelevant stimulus which has previously been associated with high monetary reinforcement captures the eyes attention much stronger than that very same stimulus when previously associated with low monetary reinforcement. They contend that reinforcement changes the salience of a stimulus, such that a stimulus which has been associated with a reinforcer becomes more pertinent and therefore captures the eyes' attention above and beyond its physical salience. Furthermore, because the stimulus captured the eyes' attention and disrupted goal-directed behaviour, the authors argue that this effect was automatic and not driven by strategic top-down control. This suggests that targets associated with good outcomes became visually salient.

Attentional capture by Pavlovian reward-signalling distractors has been shown to persist in an unrewarded test session. In this procedure, the colours of distractor stimuli in the conditioning phase signalled the availability of high or low reward. Response latencies showed that performance in the initial conditioning session was impaired by the presence of a distractor signalling the availability of high, relative to low, reward and this attentional capture was not reduced in the subsequent unrewarded follow-up session (P. Watson et al., 2019). This study provides evidence that Pavlovian conditioning can govern attentional selection priority due to reward association, even when this association was entirely task-irrelevant.

Eye-movement behaviour has also been shown to be modulated by inter-trial reinforcement priming, in regards to oculomotor capture and saccadic trajectories (Hickey & van Zoest, 2012, 2013). Furthermore, oculomotor capture due to reinforcement-based selection history has been observed when participants were unconstrained regarding where they looked (Anderson & Yantis, 2012; Bucker, Silvis, Donk, & Theeuwes, 2015), or in an additional singleton task when participants were explicitly instructed to search for a target defined by a different feature than the one previously associated with reinforcement (Theeuwes & Belopolsky, 2012).

Montagnini and Chelazzi (2005) have demonstrated that saccadic latencies and saccadic peak velocities may be controlled by the ability to perform visual discrimination on the saccade target. In this study, the target was briefly displayed such that the shorter the saccadic latency, the longer the target was presented and the easier the discrimination task was. In this situation, saccadic latencies decreased while saccadic peak velocities increased. This suggests that the successful performance of the task reinforces the short latencies and indicates that saccade properties are plastic and may be influenced by their environmental contingencies.

Paeye and Madelain (2014) manipulated the reinforcement contingencies of saccadic amplitude variability. It has been demonstrated that variability may be controlled by making reinforcement contingent on variability itself (Blough, 1966; Machado, 1989; Neuringer, 2002). Variability is an operant dimension and is used to *shape* a behaviour by systematically reinforcing successive approximations towards a desired behavioural outcome (Skinner, 1953). In other words, operant variability is functional, as it allows the shaping of new behaviours by increasing the probability of emitting new approximations to the target response and the probability of detecting changes in environmental contingencies (Neuringer, 2002). This suggests that part of the variability observed in motor responses may be related to environmental contingencies, such that a given range of behavioural instances of an operant response might be reinforced. In Paeye and Madelain's study, one group were reinforced (i.e., with a tone) for the least frequent saccadic amplitudes to increase variability, and then reinforced for the central part of the amplitude distributions to reduce variability, while in a second yoked control group reinforcement was independent of the saccade's amplitude. They observed that the spread of the amplitude distributions was controlled by reinforcement. Furthermore, the reinforcement-induced increases in variability did not affect the mean saccadic amplitude. This suggests that the natural level of variability in saccadic control may result from operant learning in which vision provides functional consequences.

This suggests that reinforcement learning influences, or changes, our search strategies in a similar way to the significant and long-lasting improvements in search performance after repeated training (Ahissar & Hochstein, 1997; Karni & Sagit, 1993; Kellman & Garrigan, 2008). Furthermore, evidence indicates that eye movements made during visual search tasks rely on neural computations similar to those employed when eye movements are planned and executed to obtain explicit reinforcement and suggest that eye movements originate from the processing of, predominantly, visual sensory information, feedback about previous errors or successes based on dopaminergic feedback loops, and expectations regarding environmental factors, such as reward (Trommershäuser, Glimcher, & Gegenfurtner, 2009). Eye movements serve the purpose of gathering and updating information across successive eye movements, and these feedback loops suggest that obtaining our visual goal (i.e., finding the target stimulus) may be an environmental reinforcer. This intrinsic reinforcement regarding the success of a visual decision may induce persistent changes in sensory and motor processes which are represented in shifts in our attentional strategies.

In an ingenious study designed to shed more light on the reinforcing effects of finding the target, Paeye, Schütz, and Gegenfurtner (2016) incorporated a gaze-contingency paradigm. In this procedure, the target (i.e., a *Gabor* patch) was not visible at the beginning of each trial and reinforcement (i.e., the target being made visible) was provided on a saccade-to-saccade basis, according to two concurrent schedules. In one of these schedules, saccades landing in a particular quadrant (Experiment 2) or saccades with a direction located in a specific angle range (Experiment 3) were reinforced with a high probability, and at the same time according to the other reinforcement schedule, saccades in the three other quadrants (Experiment 2) or with other directions (Experiment 3) were reinforced with a low probability. This study demonstrated that the proportion of saccades meeting the reinforcement criteria increased dramatically, and participants matched their search behaviour to the relative reinforcement rates of targets, according to the *Generalized Matching Law* (Baum, 1974). These findings show that finding a search target at a particular location increases the probability that participants make saccades to that region in the future (i.e., the behaviour is strengthened) and, importantly, the act of finding the target, without any form of monetary or other reward associated with it, serves as a reinforcer during visual search. This suggests that people learn to prioritize important locations (i.e., those which are more likely to contain the target) from experience, based upon regularities in the environment. In other words, reinforcement learning may serve as the mechanism for optimizing eye-movement strategies, and this learning may form the basis of intelligent behaviour without requiring prior knowledge or extensive computations. With this in mind, we now turn towards our current study.

1.8 The Current Study

1.8.1 The Research Question and Hypothesises

We know that the prevalence of a target signal influences visual search behaviour dramatically in systematic ways. For example, in visual search tasks, the likelihood that an observer will miss a target varies with the frequency in which it is presented across trials (Ishibashi, Kita, & Wolfe, 2011; Rich et al., 2008; J. Schwark et al., 2012; Wolfe et al., 2005, 2007). When target prevalence is low, target-absent responses become faster and miss errors tend to increase, while false alarm errors are rare under these conditions; as a consequence, the participants' response patterns become increasingly conservative, meaning that there is a bias towards making the target-absent response (Wolfe et al., 2005, 2007). Conversely, when target prevalence is high, target-absent responses are slowed and miss errors are rare while false alarm errors tend to increase; as a consequence, the participants' response patterns become increasingly liberal, meaning that there is a bias towards making the target-present response (Wolfe & Van Wert, 2010).

The power of the schedule of signal presentation to control visual search behaviour has been pondered over for a long time and speculation over this control dates back to Holland (1958) who reported a series of vigilance experiments in which participants, working in a darkened room, would repeatedly press a key to briefly (i.e., 700ms) illuminate a dial. A non-signal was defined as a dial with a needle in the vertical position. A signal occurred when this needle was deflected from the vertical to the horizontal position. Horizontal dial signals remain present until detected. Unbeknown to the participants, the primary measure of interest was the illuminating key press rate across various schedules of signal presentation. Cumulative response curves were generated, and it was observed that when signals were presented on fixed time interval schedules in which the first response after the specified interval was followed by a dial deflection, the dial illumination rate was higher on shorter interval schedules. When signals were withheld, there was a high rate of responding before gradually declining until the dial viewing response was extinguished. Holland contended that this decline was dependent upon the removal of the signal and could not be explained by fatigue, as high rates of behaviour were maintained over similar periods on other schedules.

Interestingly, he also noted that immediately following the detection of a signal needle deflection there was a brief period of no dial illumination key pressing and an increasing rate of responding occurred as the fixed-interval period elapsed. These observed scallop-shaped response patterns were similar to the patterns of responding produced by non-human studies when animals worked for primary reinforcers on fixed-interval schedules of reinforcement (Ferster & Skinner, 1957) and that this brief pause in responding following a signal deflection was reminiscent of a post-reinforcement pause when an animal working on a schedule of reinforcement does not respond for a brief time following the delivery of a reinforcer (Felton & Lyon, 1966). After intensively studying various schedules of signal presentation, Holland concluded that the response patterns seen in his dial observation experiments were consistent with animals working under similar schedules of reinforcement and inferred that the signal presentation schedule was controlling observing behaviour.

Blair (1958) and Baker (1960) expressed scepticism that Holland's (1958) findings would generalise to behaviours with greater face validity, such as head orientation or eye movements. To address this scepticism, Schroeder and Holland (1968) evaluated the signals' reinforcing effects on eye movement patterns and signal detection performance. They demonstrated that eye movement patterns were similarly controlled by signal reinforcement schedules in earlier studies involving key-pressing (Holland, 1958). Importantly, they also observed that eye movements decreased as signal prevalence decreased and that these declines were associated with increased rates of miss errors. They concluded that an eye movement can act as an operant controlled by its consequences and that operant control of eye movements has important implications for human factor analysts concerned with "attention".

Holland's (1958) vigilance tasks involved simple visual signals. However, increased complexity may affect search behaviour. To address this concern, Hogan, Bell, and Olson (2009) applied Holland's procedure to a simulated X-ray baggage screening task, similar to Wolfe and colleagues' (2005) task. Signals were an image of a knife presented according to a variable-interval 6-minute schedule of signal presentation, in which signals were presented following random intervals of time centred around a 6-minute mean and an extinction condition where signals were withheld. A linear regression line was fitted to each cumulative response record, and the slope for the variable-interval schedule was approximately 30% steeper (i.e., observing response rate) than those recorded under extinction conditions. The extinction condition produced greater response variability, and some participants exhibited periods of no responding followed by a response burst. These response patterns supported Holland's conclusion that the schedule of signal presentation was controlling observing behaviour and led Hogan, Bell, and Olson to conclude that the schedule of signal presentation was controlling search behaviour through the operant processes of reinforcement and extinction. They hypothesized that signal detection functions as reinforcement for search behaviour. Unfortunately, Holland's (1958) and Hogan, Bell, and Olson's (2009) procedures were strictly vigilance tasks, which repeatedly presented the signal stimulus in the same location and cannot be considered an active visual search task. Furthermore, as the signal stimulus remained present until a detection response had been made, signal detection measures of performance were unobtainable.

Current evidence suggests that eye movements made during visual search tasks rely on neural computations similar to these employed when eye movements are planned and executed to obtain explicit reinforcement (Trommershäuser et al., 2009) and that variability in saccades may also be affected by reinforcement (Madelain, Champrenaut, & Chauvin, 2007; Madelain, Paeye, & Wallman, 2011; Paeye & Madelain, 2014). Theoretical models of human search behaviour assume that fixation locations are selected to maximise the information gained across successive eye movements (Najemnik & Geisler, 2005) and to minimize uncertainty regarding the location of the target stimulus (Renninger, Verghese, & Coughlan, 2007). This suggests that the visual information gained may be conceptualised as a reinforcing consequence, controlling fixation locations. In other words, we gain visual information which may be indirectly used to act and obtain reinforcement from the environment.

Behaviour analysis views the principle of selection by reinforcement as playing the same role in our understanding of behaviours that natural selection plays in understanding the evolution of organisms (Donahoe & Palmer, 1994; Madelain, Paeye, & Darcheville, 2011). It has been demonstrated that variability may be controlled by making reinforcement contingent upon variability itself. Variability is an operant dimension, and functional because it allows the shaping of new behaviours (Blough, 1966; Machado, 1989; Neuringer, 2002; Pryor, Haag, & O'Reilly, 1969). Therefore, it may be suggested that part of the variability observed in motor responses is related to environmental contingencies, such that a given range of behavioural instances of an operant response may be reinforced. This suggests that oculomotor responses (i.e., search behaviour) may also be reinforced by environmental contingencies and there is a small but growing body of evidence to support this understanding.

For example, Montagnini and Chelazzi (2005) demonstrated that saccadic latencies and saccadic peak velocities may be controlled by the ability to perform visual discriminations. In this study, participants were instructed to perform a discrimination task on the saccade target. The target was briefly displayed such that the shorter the saccadic latency, the longer the target was viewed and the easier the discrimination task was. They observed that saccadic latencies decreased while saccadic peak velocities increased and concluded that the successful performance of the task reinforces the short latencies which suggests that saccade properties are plastic and may be influenced by their environmental contingencies.

Furthermore, Paeye, Darcheville, and Madelain (2007) used a target whose position was continuously adjusted to match the gaze position, such that the target did not move if the eye was fixed. Therefore, pursuit initiation could not be controlled by the physical motion of the target. A shaping procedure was implemented using an auditory tone based on the eye velocity of anticipatory movements. They observed that participants progressively learned to initiate smooth movements in the absence of a motion signal. These results demonstrate that smooth pursuit, even in the absence of a visual target or external motion, is an operant behaviour.

Similarly, Paeye and Madelain (2011) demonstrated that saccadic endpoint variability may result from operant learning mechanisms and not the outcome of neural noise during sensorimotor processing. In this study, the least frequent saccadic amplitudes were arbitrarily reinforced (i.e., with a tone) to increase variability and then the central portion of the amplitude distribution was reinforced to reduce variability. To examine the effects of the contingency, high levels of variability were reinforced in four participants and four other participants were assigned to a yoked control group in a second experiment. On average, saccadic amplitude standard deviations were doubled while the medians remained mostly unchanged in both experiments, and variability returned to baseline levels when low variability was reinforced. No consistent changes in amplitude distributions were observed in the control group. These results demonstrate that variability may be reinforced and challenge the idea that saccadic variability is due to stochastic neural noise. Paeye and Madelain (2011) contend that saccades are operant behaviours and that selection processes constrain saccadic amplitude distributions. In real life these movements are controlled by an enhanced clarity of a visual target, allowing adapted visually guided behaviours which act as reinforcers (Madelain, Paeye, & Darcheville, 2011).

Furthermore, Madelain, Harwood, Herman, and Wallman (2010) investigated saccadic adaptation to either a target stimulus or a distractor, such that the unadapted saccade landed either on the distractor or on the target. Three experiments were conducted, in which the post-saccadic locations of the target and distractor were varied. They showed that decreased amplitude adaptation, increased amplitude adaptation, and recovery from adaptation were controlled by the post-saccadic position of the target rather than the distractor. These results demonstrated that in the presence of two conflicting post-saccadic visual stimuli (i.e., the target or the distractor), the saccadic system selectively adapted its amplitude as though only the target was present. Importantly, Madelain et al. (2010) contend that the target-oriented adaptation observed may be explained by the reinforcing effect of viewing the target stimulus. The saccades landing near the target were selected whereas those landing near the distractor were not. In other words, finding the target is visually reinforcing.

Moreover, Madelain, Paeye, and Wallman (2008) induced changes in saccade amplitude in the absence of retinal error, using arbitrary reinforcement (i.e., either an auditory tone or viewing the target on the fovea). To eliminate retinal error, they either had the target vanish when the saccade started or placed the target on the fovea at the end of the saccade. They computed saccadic amplitudes in real time and a percentile procedure was used to compute the reinforcement criteria on each trial. It was observed that the amplitude of saccades followed the reinforcement contingencies, such that amplitudes decreased when small saccades were reinforced while amplitudes increased when large saccades were reinforced. They concluded that saccadic amplitudes were being controlled by a general learning process relying on the outcome of each movement given a particular state of the environment. In other words, adjusting saccadic amplitude may be explained by a selection process in which those responses that produce a valuable outcome are selected over those that do not. Accordingly, viewing the target may be reinforcing because maximizing visual perception would allow other adaptive interactions within the environment (Madelain et al., 2008).

Importantly, Paeye et al. (2016) have demonstrated that finding the target stimulus at a particular location increases the probability that participants will make saccades to that region of the display in the future, and they suggest that the act of finding the target, without any monetary or other reinforcing consequences associated with it, may serve as a reinforcer for visual search behaviour. This suggests that finding the target stimulus may be intrinsically reinforcing and that visual search, in the form of observed oculomotor responses, is a set of operant responses controlled by environmental contingencies. Simply put, visual search is an operant behaviour.

These studies demonstrate that voluntary eye movements are controlled by operant learning processes relying on the outcome of each movement given a particular state of the environment; behaviours that produce a favourable outcome are selected and strengthened while behaviours that do not are weakened and rejected. Accordingly, eye movements may be controlled by the ability to perform visually guided behaviours that act as reinforcers because of their adaptive values and viewing the target may act as reinforcement because maximizing visual perception would allow other adaptive interactions with the environment. In other words, environmental contingencies are controlling eye-movement behaviour. In the current study, it should come as no surprise that there would be an increase in search behaviour when target signals are abundant due to operant learning processes and the reinforcing effect of finding what you are looking for.

With this in mind, we asked the question, "Is the schedule of signal presentation controlling eye-movement behaviour and if so, how?" We suggest that search behaviour may itself be influenced by the outcomes of search behaviour through the operant processes of reinforcement and extinction. Specifically, in operant conditioning, behaviours are differentially selected by their consequences (Skinner, 1984). According to operant theory, which is primarily concerned with reinforcement learning mechanisms, behaviours reinforced in a particular context are "selected" and more likely to occur under similar conditions in the future. If reinforcement is no longer forthcoming (i.e., extinction is in effect), the likelihood of the behaviour will decrease over time (Skinner, 1984). Therefore, reinforcing consequences will strengthen a behaviour by maintaining or increasing those behaviours that they follow and if this contingent reinforcement is withheld, the dependent behaviour will be weakened and over time this behaviour will be extinguished (Pierce & Cheney, 2017). Following Hogan, Bell and Olson's (2009) premise, if we consider that signal detection functions as reinforcement for visual search behaviour, operant theory will predict more intense searching behaviours and improved detection performance when target signals are plentiful, due to the strengthening effects of reinforcement on search behaviour when we find what we are looking for. Conversely, the absence of target signals for long periods, would weaken search behaviours and increase signal detection errors. This weakening of search behaviour and the corresponding increase in signal detection miss errors under low-prevalence target signal conditions would influence the decision criterion (i.e., a conservative criterion shift), and lead to the commonly observed Low Prevalence Effect. We term this complimentary pair of predictions the *Reinforcement Strengthening Hypothesis*.

In an attempt to answer our research question, we incorporated a repeated-measures signal detection experimental design containing two prevalence conditions, namely a medium-prevalence condition and a low-prevalence condition, and incorporated the *Advance-key response procedure* across three experiments. We used eye-tracking technology to investigate eye movements to ascertain the effects of the schedule of signal presentation on eye-movement behaviour. To clarify the mechanisms underlying the Low Prevalence Effect and provide a richer understanding of visual search behaviour.

1.8.2 The Advance-key Response Procedure

The advance-key response procedure is a variant of the yes-no signal detection task and researchers have predominantly applied this procedure as a laboratory analogue to natural prey-detection and foraging behaviour (Bond & Kamil, 2002; Kamil, Lindstrom, & Peters, 1985; Kamil, Yoerg, & Clements, 1988; Kono, Reid, & Kamil, 1998; Mccarthy, Voss, & Davison, 1994; Pietrewicz & Kamil, 1977; Voss, Mccarthy, & Davison, 1993; Weisman, Gibson, & Rochford, 1984). When a predatory animal searches for cryptic prey, they are faced with a situation in which prey transmit sensory cues that cannot be discriminated with certainty (Curio & Curio, 1976; Edmunds & Brunner, 1999; Endler, 1988). In this respect, foraging animals may be conceptualised as decision-makers attempting to solve a problem (Kamil & Yoerg, 1982). When searching for prey and the quality of a patch is unpredictable, a predator must invest time and energy in discovering the patch's value. The trade-off between time spent assessing and time spent exploiting a patch is similar to the problem faced by laboratory animals in a concurrent reinforcement schedule.

From an ecological perspective, a concurrent reinforcement schedule can be thought of as analogous to having two (or more) patches in which to hunt, each of which pays off differently. The possible outcomes of an encounter between a predator and a potential prey item can be represented in a stimulus-response matrix, similar to the requirements of the yes-no paradigm in a signal detection task (D. M. Green et al., 1966). In this case, the two stimuli are prey and non-prey, and the responses are attack or reject. Each stimulus-response event is associated with some benefit or cost to the predator (Staddon, 1983, Prey-detection Model), and Steddon contends that rejecting a non-prey item (i.e., a correct rejection) has no beneficial consequences, or may result in a time and energy cost in moving from one patch to the next.

Kamil and colleagues' pioneering ecological approach to laboratory animal detection behaviour and Staddon's prey-detection model inspired a variant of the yes/no signal detection procedure, termed the advance-key procedure (Weisman et al., 1984). In this laboratory response procedure, the advance response simulates the behaviour of moving from one stimulus set to another. In this respect, the participant observes a stimulus set and can make a detection response ("attack" in the prey-detection model), much the same as a "yes" response in a yes/no signal detection task. However, the participant must "advance" to the next stimulus set to terminate a trial (move to the next patch in the prey-detection model). If the advance response is made before a detection response, this is considered a rejection or a "no" response in the yes/no procedure. Therefore, in many respects, the advance-key procedure may be considered similar to that of the yes/no procedure. However, there are some important differences between these two response procedures which, we think, help align the advance-key procedure with many aspects of vocational detection tasks in the field.

Firstly, with the advance-key procedure, if feedback is provided, it is only ever provided for correct detection responses. This is in line with the prey-detection model, as a correct rejection (leaving a patch) would never be directly reinforced with prey, and there is a cost involved, concerning the time and energy moving from one prey site to another. Similarly, we would not expect to see security personnel congratulating each other for correctly rejecting a piece of luggage which does not contain illicit contraband.

Secondly, with the advance-key procedure, a conscious decision to advance must be made on each trial, even if a detection response has been made. In this respect, each trial and the overall pace of the search task is self-governed. In the field, an operative moves at their own pace from one stimulus set to another, until they detect a target anomaly. Once a potential target has been identified, the operative is then free to flag this stimulus set for further investigation before moving on to the next set in their own time.

Another potential difference between these two procedures is that, with the yes/no procedure, there is an explicit "no" response, whereas with the advance-key procedure, the "advance" response is not an explicit "no" but the same response required to present the next stimulus set. This also appears to align closely with many applied signal detection tasks, in which the operative advances through X-ray images, pieces of luggage or some other stimulus set and only makes a detection response when a signal has been detected. Therefore, while these two response procedures may appear similar on the surface, there are subtle differences.

In our Advance-key procedure, the participants were provided with a two-key button box with one key labelled “Detect” and the other key labelled ”Advance”. The participant would observe a stimulus set and if a target was observed, they could make an indication response by pressing the detect key (“attack” in the prey-detection model or a “yes” response in the yes-no procedure). However, an “advance” response must be made to terminate any trial and advance to the next stimulus set (move to the next patch in the prey-detection model). If the participant did not detect a target in the stimulus set, they would simply press the advance key to terminate the trial and advance to the next stimulus set.

Method

2.1 Participants

An *a priori* power analysis was performed for sample size estimation based on data from two published studies involving five experiments that were similar to our study, Rich et al. (2008) and Van Wert et al. (2009). The effect sizes (partial eta-squared) for these experiments ranged between 0.54 - 0.81, with a mean of 0.62. Using G*Power (version 3.1) with alpha = 0.05, power = 0.80, and an effect size equivalent to partial eta-squared = 0.50, a required sample size of N = 8 was obtained for within-group comparisons. Thus, we selected a sample size of N = 20 for each Experiment to allow for expected attrition and additional analyses.

A total of sixty participants were recruited from the University of Waikato's School of Psychology student body and received 1.5 course credits for participating. All participants reported normal or corrected-to-normal visual acuity and no deficits in colour discrimination. All recruitment and experimental protocols were subjected to and received ethical approval from the University of Waikato's School of Psychology Human Research and Ethics Committee (# 17:15).

2.2 Apparatus

All stimuli were presented using a Dell OptiPlex 760 Minitower PC, and displayed on a VIEWPixx display (VPixx Technologies) with a 1920 x 1200 pixel resolution and a refresh rate set at 60Hz. Eye movement data were recorded using an EyeLink 1000 Desktop System (Eyelink 1000, SR Research, Ltd., Ontario, Canada), averaging 0.25° - 0.5° accuracy. A chinrest was used to ensure that each participant's head remained fixed for the duration of the experiment, and this was located 57cm away from the monitor screen, producing a field of view (FOV) of 40° x 30° (horizontal x vertical). Participants responded via a two-key button box, with one key labelled “Detected” and the other labelled “Next Trial”. A 100-watt incandescent desk lamp projected light in the opposite direction from the computer monitor onto the rear wall of the laboratory, apart from the illumination of the monitor screen, this was the only light source in the laboratory during the experiment.

2.3 Stimuli

On each trial image, a set size of 15 blue-coloured rain droplet stimuli were presented on random loci within an invisible 5 x 6 stimulus matrix. The visual angle for each droplet subtended 1.2° x 1.7° (horizontal x vertical), each droplet matrix box subtended 2.4° x 2.0° (horizontal x vertical) producing an overall stimulus matrix subtending 12.0° x 12.1° (horizontal x vertical).

The target and distractors were defined by their triple conjunction of colour, ranging from 0 – 255 across the red, green, and blue colour values. Throughout the entire experiment, the red colour value was set at 0 and the blue colour value was set at 255 for both the distractors and target stimuli. The distractors’ green colour value was set at 128. The colour of the target stimulus was either set (Experiment One) or individually titrated for each participant using a two-down/one-up single staircase procedure (Experiments Two and Three: see Design 2.4). Smaller values of the green channel (i.e., < 128) made the target droplet a “darker” blue than the distractor droplets and the target droplet was always a “darker” blue.

These droplet stimuli were presented over a random multi-scale grey textured noise background. Following the procedure of Cheng and Li (2012), this textured noise image was generated by randomly assigning shades of the grey colour scale to each pixel and convolved with a range of 2-D *Gaussian* filters of different sizes (SD=16, 8, 4, 2, 1 pixel: Amplitude=1, 1, 1, 1.5, 2 pixels). The resulting filtered images were summed to produce an overall noise image which simulates the type of noise textures typically found in natural settings (Cheng & Li, 2012).

Ninety target-absent trial images were generated with distractor droplet stimuli located on random loci within the stimulus matrix. Ninety target-present trial images were also produced with the target stimulus placed in each locus within the thirty loci stimulus matrix on three separate occasions and the remaining distractor stimuli were randomly assigned within this matrix for each titrated colour value. These trial images were randomly presented with replacement for each participant throughout the experiment according to the prevalence condition's random-ratio schedule of signal presentation. This meant that for target-present trials, the one target stimulus presented was of a darker blue compared to the distractors.Below the stimulus matrix box was a cartoon garden scene which was present in each trial image. Figure 2.1 presents a screenshot image for a target-present trial.

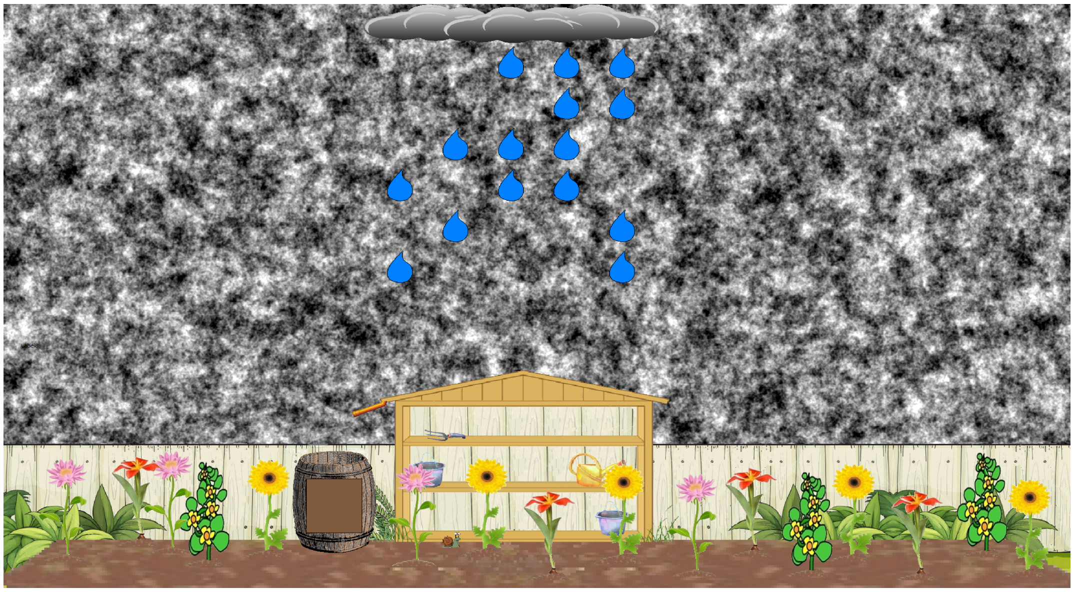

Figure 2.1: A screenshot presenting a trial image and the stimulus set for a target-present trial. The darker blue droplet is the target stimulus.

2.4 Design

For each trial in the advance-key procedure, the participant was presented with a droplet stimulus set and, if a target was not detected, they would press the button box's "advance-key" to instantaneously progress to the next trial (no inter-trial interval). If the participant concluded they had detected a target, they were required to press the "detected key", followed by the "advance-key" to progress.

The first trial block consisted of practice trials. During this block, a large discrepancy between the target stimulus and distractors was set, with the target green colour value of 16, and the target prevalence set at a probability of 0.5. Participants were required to correctly identify 90% of both target-present and target-absent trials within a six-minute period. If they failed to achieve this, the experiment was terminated.

The second trial block involved a single staircase titration procedure which ran for six minutes. A target stimulus was presented on each trial and the procedure began with the same target colour conjunction used in the practice trial block. This procedure was designed to find the target detection threshold for each participant using a two-down/one-up staircase method (Levitt, 1971). During this procedure, the blue colour component of the target stimulus was adjusted by adding or subtracting the green hue within the conjunction. The blue colour of the target stimulus would lighten by one step size following two consecutive correct responses, bringing the target colour closer to the distractors' colour conjunction; and darken in colour by one step size following each incorrect response, distancing the target colour from the distractors. Step sizes halved on the first, third, seventh, and thirteenth reversals and the colour step sizes were 16, 8, 4, 2, and 1 of the 0-255 green colour spectrum or $6.3\%, 3.1\%, 1.6\%, 0.8\%$, and 0.4% of the green colour value respectively. At the end of this procedure, the mean colour value for the preceding eight reversals became the target stimulus colour for the remainder of the experiment. This mean colour value converged on an approximate difficulty probability of 0.71 for each participant (Lu & Dosher, 2013). The inclusion of this staircase procedure controlled for individual variance in visual discrimination across the sample and ensured that the task difficulty was similar for all participants. This procedure also allowed the experimenter to set the level of task difficulty.

Following from previous visual search literature (Fleck & Mitroff, 2007; Rich et al., 2008; Van Wert et al., 2009; Wolfe et al., 2007), the third trial block, termed the Experimental block, consisted of two conditions incorporating random-ratio schedules of signal presentation. This Experimental block started and finished with a fixation cross in the centre of the screen. The medium-prevalence condition consisted of 200 trials with a 0.5 probability of a target signal randomly appearing throughout the condition. The low-prevalence condition consisted of 1,000 trials with a 0.02 probability of a target randomly appearing (Experiment One and Two) or a 0.00 probability (Experiment Three). These schedules equated to the participant receiving 100 target-present trials during the medium-prevalence condition, and 20 target-present trials during the low-prevalence condition (Experiment One and Two). Participants seamlessly transitioned prevalence conditions and there was no stimulus change. As we were interested in the progression of eye-movement behaviour over time, this study was strictly a freeview experimental design.

The order of prevalence conditions was counterbalanced across participants; half the participants experienced the medium-prevalence condition first followed by the low-prevalence condition, and the remainder experienced the low-prevalence condition first followed by the medium-prevalence condition. No feedback was provided during the experiment.

An instruction frame preceded each trial block, which stated: "When you detect a darker coloured raindrop on the screen press the 'detection button' on the left of the button box and then the 'next trial button' on the right to advance to the next trial. If there is no colour difference in the raindrops on the screen, simply press the 'next trial button' on the right to advance to the next trial. Please respond as quickly and as accurately as possible."

2.5 Procedure

Following the consent process, the first block of trials was the practice block. This block began with the instruction frame and the participant read this frame in their own time before pressing the button box's "advance-key" to progress onto the first practice trial. The second block of trials was the staircase titration procedure which also began with the instruction frame.

Following the titration procedure, eye tracking calibration and validation were conducted before the start of the experimental trial block; using a 9-point calibration grid (Eyelink 1000, SR Research, Ltd., Ontario, Canada), if the calibration and validation process exceeded 4° of visual angle, the procedure was repeated until this threshold was not exceeded. Following this process, the instruction frame was again presented, and the participant would read the frame before pressing the "advance-key" to progress onto the next frame, which was a fixation cross. Participants fixated on the crosshairs of the cross before pressing the "advance-key" to progress onto the first raindrop trial image. Following all 1200 experimental trial images, another fixation cross frame was presented. Again, participants fixated on the crosshairs before pressing the "advance-key" to end the experiment. The initial fixation period and the one at the end were used to check for drift in the eye tracker during the experiment. If this drift was two standard deviations above the participant's mean, the data was disregarded. The experiment was concluded with a debriefing session in which the rationale for the experiment was presented and participants could have any questions answered.

2.6 Data Analysis

Data analysis was conducted using custom-built MatLab programmes (R2017b, Mathworks, Natick, Massachusetts, USA). For each trial, the x and y eye movement coordinates, the trial image number, and the button press response were extracted from the EyeLink output file. The trial image was paired with the button press response to produce the signal detection theory measures for hits, misses, false alarms, and correct rejections.

An invisible data catchment box subtending 15.8° x 15.8° was placed over the stimulus matrix and only the eye movement coordinates within this box were analysed. Any missing x and y eye-movement coordinates due to blinking were filled with NaNs and interpolated. A *Gaussian* smoothing filter (SD = 10 pixels, Width = 60 pixels) was then convolved across the x and y eye movement data to smooth these position traces.

The x and y coordinates for the centre of each droplet in the trial image were overlaid across the eye movement data and a circle with a radius subtending 0.85° represented the droplet catchment area for eye-movement behaviour.

Eye movement velocity was then calculated using a one-millisecond sampling period (see Figure 2.2. A velocity less than 5° per second was classified as a fixation, while a velocity greater than 5° per second represented a saccade (Liston, Krukowski, & Stone, 2013). Fixation start and finish times were calculated, and this provided our three measures of eye-movement behaviour: namely, fixation length (ms) per trial, the number of fixations (N) per trial, and the saccadic scanpath (degrees of visual angle) per trial. We had no a priori reason to select one eye-movement measurement over any others and so we are using the "widest possible net" to catch any possible changes in eye-movement behaviour brought about by the different prevalence conditions. These three measures have been used in a wide range of other eye movement studies (Liston et al., 2013; Madelain et al., 2007; Paeye et al., 2016; Tatler, Baddeley, & Gilchrist, 2005; Wienrich et al., 2009) and so have been selected for analysing the data in this thesis.

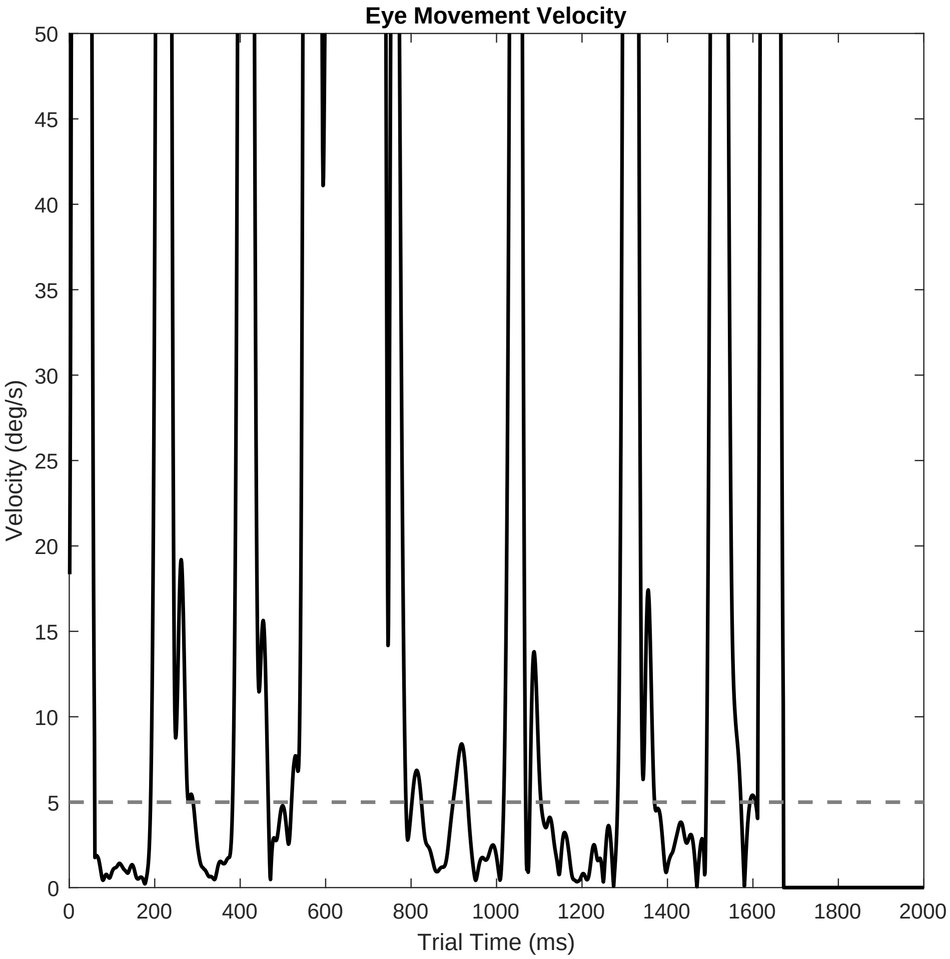

Figure 2.2: The eye-movement velocity across a trial. The dashed horizontal line set at 5 deg/sec represents the boundary between a fixation and a saccade, with a velocity greater than 5 deg/sec representing a saccade and a velocity less than 5 deg/sec representing a fixation.

These per-trial measures of eye-movement behaviour were plotted across prevalence conditions and to estimate the change of eye movement measures as a function of trial number, we fitted an exponential decay curve to each data set. The three decay curve parameters were calculated, which consisted of the magnitude of the curve decay or rise along the y-axis (curve parameter 1), the rate of curve decay or rise (curve parameter 2), and the asymptote of the curve tail on the y-axis (curve parameter 3).

2.7 Stimulus Density Map

When we first examined the target-present scanpath data for Experiment One, we observed that the participants often examined a portion of the stimulus matrix which did not contain the target stimulus which, on occasions, resulted in target misses. On occasions, we observed that the participants would saccade past the target, or even directly over the target stimulus, to reach a search location of their choice (see Figures 2.3 & 2.4).

This pattern of behaviour intrigued us, and we wondered what was drawing their search behaviour away from the target stimulus, in what should be a simple colour singleton-feature search task, or what has been termed a "pop-out" search task (Treisman & Gelade, 1980). Due to this observation, we wondered if the participant's search behaviour was being drawn towards the area of the stimulus matrix with the densest proportion of stimuli, to examine the greatest number of stimuli with the least amount of search behaviour. We decided to investigate this pattern of behaviour by formulating a stimulus density map for each trial image, to examine whether the density of the stimuli within the matrix was interacting with their search behaviour by drawing their behaviour towards this portion of the search image.

The stimulus density map for each trial image was formulated by transforming the trial image into a black background with a 1-pixel white dot representing the centre of each droplet location and a 2-D *Gaussian* filter (Sigma = 50 pixels, Amplitude = 1 pixel) convolved across this image (see Figure 2.5). The resulting image may be represented as a contour plot or a 3-D surface plot (see Figures 2.6 & 2.7).

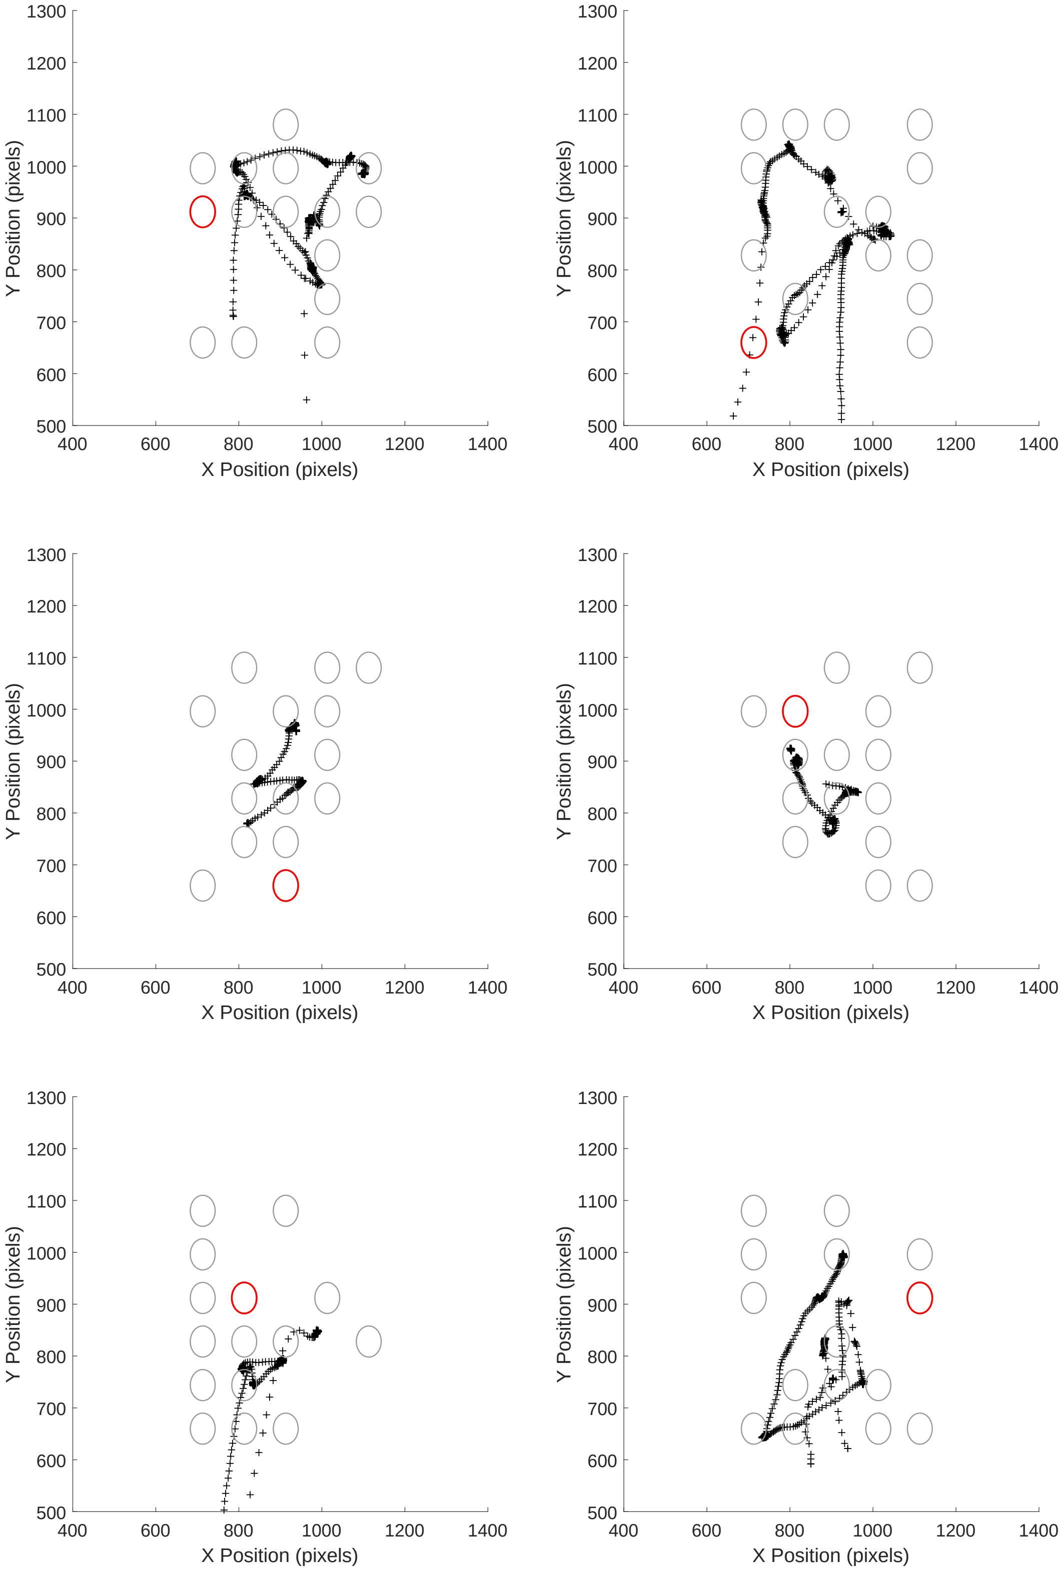

Figure 2.3: A sample of the participant's scanpath trajectories on trials where they missed the target stimulus. The target stimulus is marked in red.

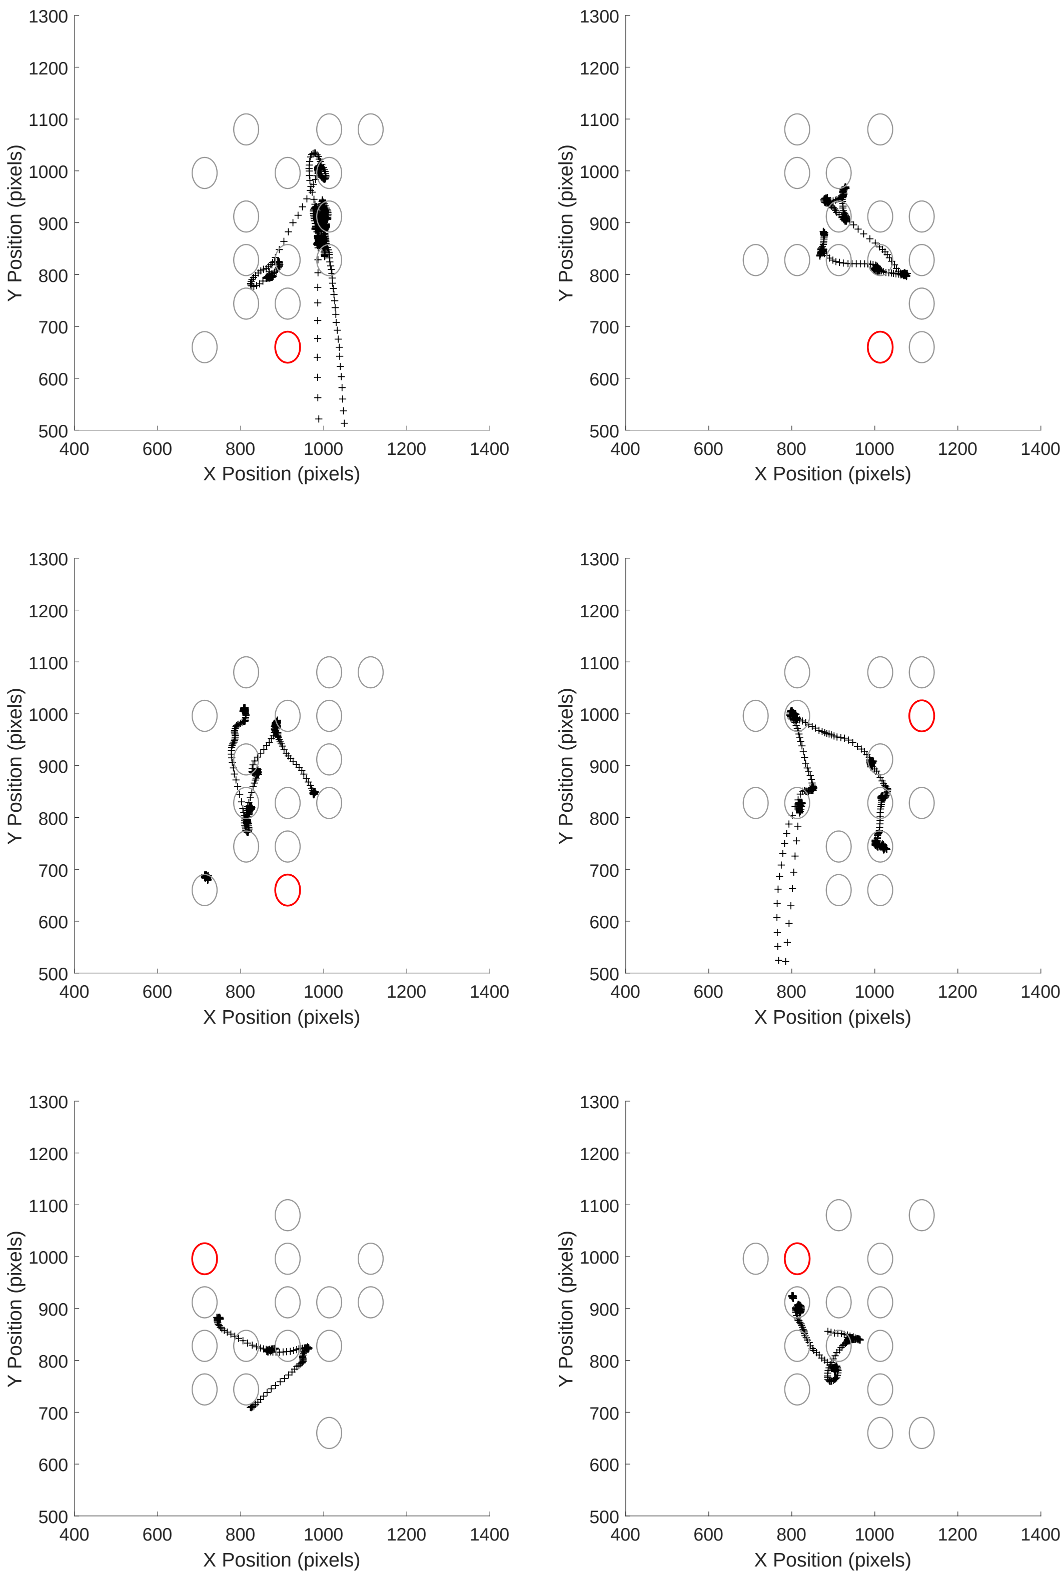

Figure 2.4: This continuation of figures shows a sample of scanpath trajectories across individual participants on trials where they missed the target stimulus. The target stimulus is marked in red.

The data were thresholded at 0.8 and data equal to or greater than this value represented the portion of the image which contained the greatest density of stimulus items, which we termed the ”Centroid” region (see Figure 2.8). Eye-movement behaviour was analysed within and outside this centroid zone.

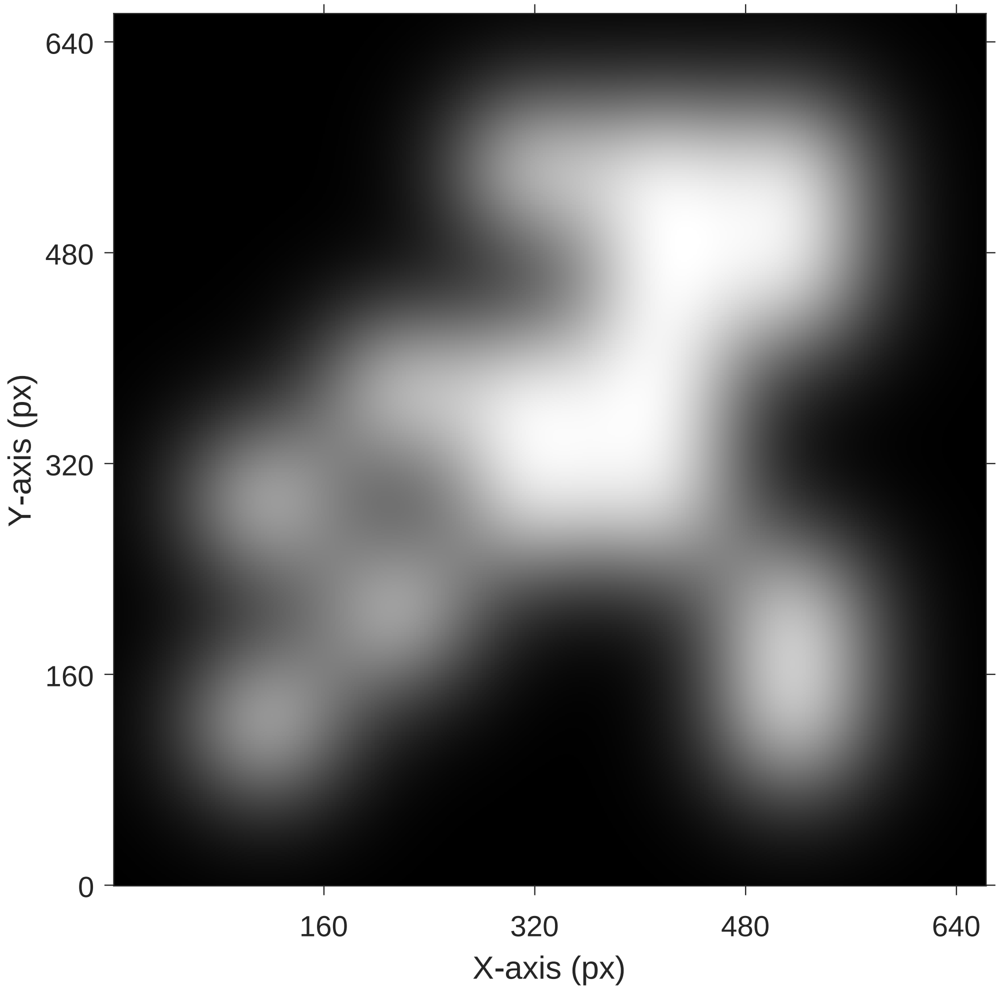

Figure 2.5: This plot shows the droplet “Density Map” created by convolving the drop locations (single white pixel) with a 2-Dimensional *Gaussian*. The lighter areas represent areas with a higher concentration of droplets.

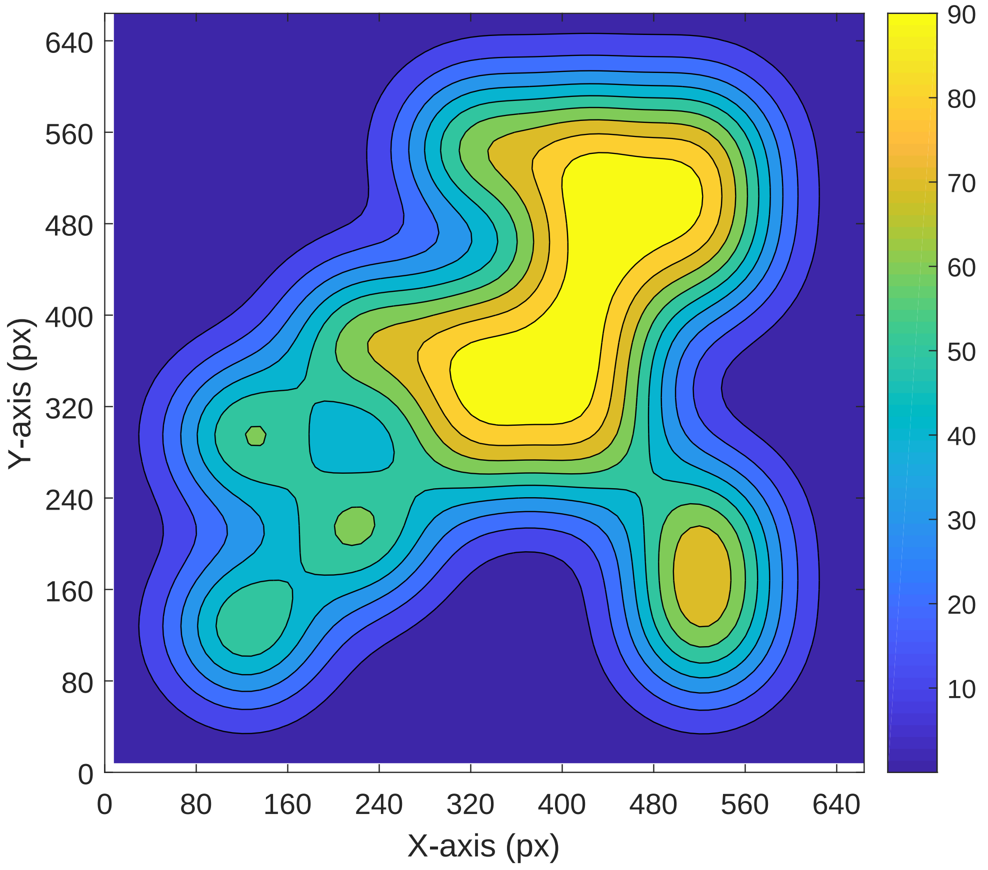

Figure 2.6: This plot shows the droplet "Density Map" in the form of a contour plot. The yellow areas represent areas of higher density.

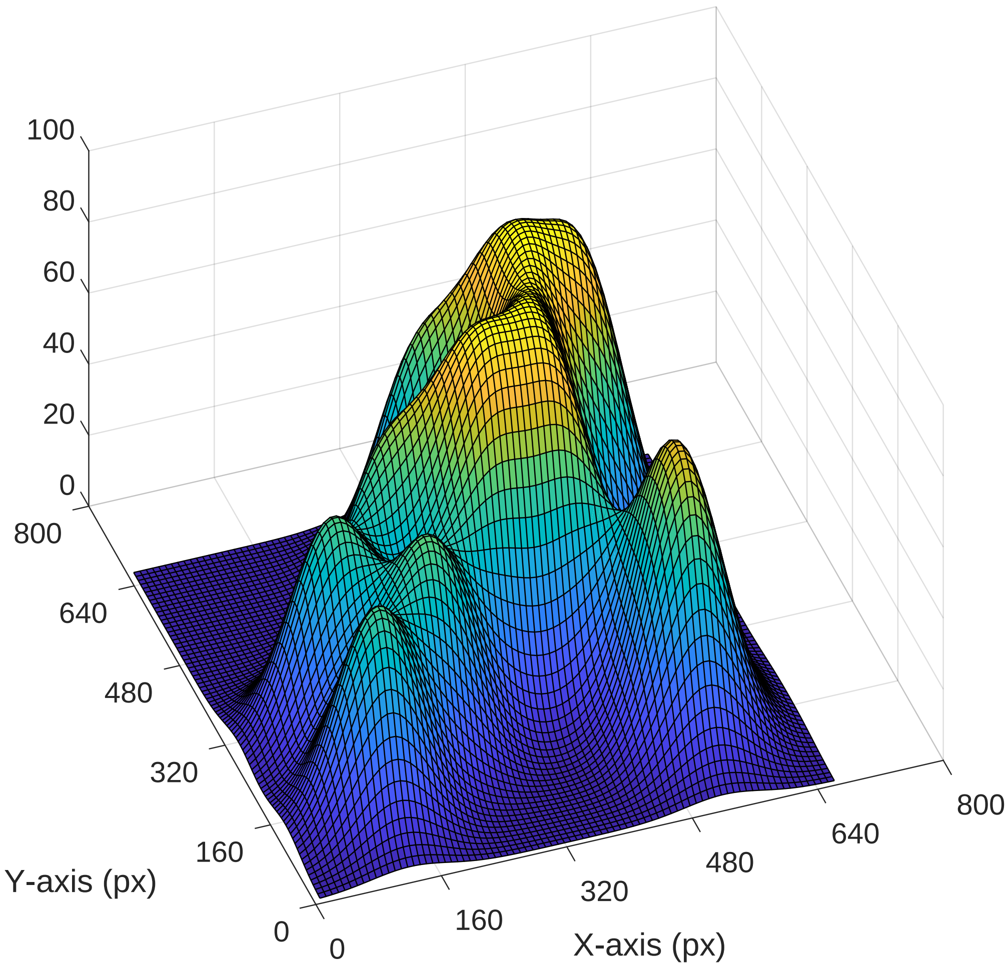

Figure 2.7: The droplet "Density Map" is represented as a 3-D surface plot. The peaks represent higher-density regions.

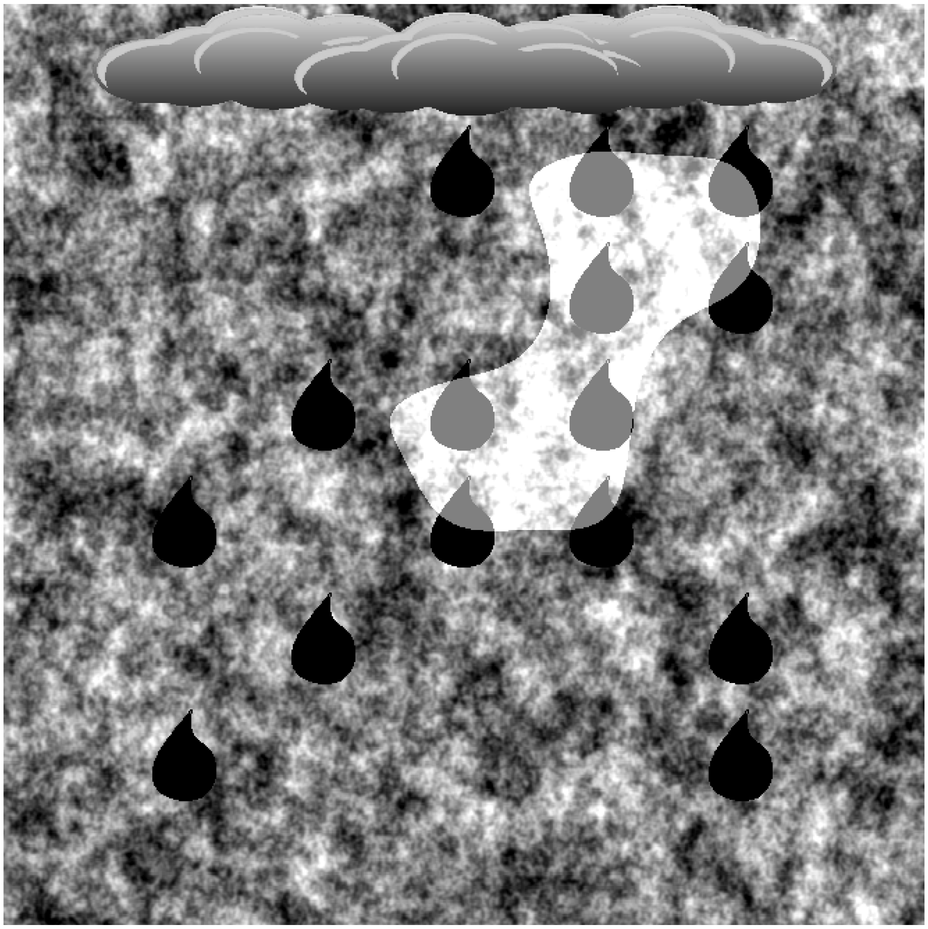

Figure 2.8: A greyscale trial stimulus image overlaid with its centroid density zone.

2.8 Statistical Analysis

Statistical analysis was conducted using IBM SPSS statistical analysis software (Version 24). In the Tables, we report means followed by their standard deviation in parentheses. Repeated measures ANOVAs were used between prevalence conditions and Univariate ANOVAs were used between experiments, with the alpha level set at 0.05. Partial-eta squared (η_P^2) was used as a measure of effect size for these statistical tests with a partial-eta squared value of 0.009 representing a small effect size, 0.0588 representing a medium effect size, and 0.1379 representing a large effect size (Richardson, 2011).

The Preliminary Experiment

The primary purpose of this experiment was to ascertain whether the results obtained from the current study's methodology corresponded to those found in previous low prevalence effect literature. In this respect, participants were titrated using the two-down/one-up single staircase method. However, eye-tracking technology was not incorporated into this experiment. We were interested to see if our results were consistent with the Criterion Shift Hypothesis (Wolfe et al. 2007) and, consequently, the Multiple Decision Model (Wolfe & Van Wert, 2010).

3.1 Participants

Twenty participants (9 male & 11 female) were recruited from the University of Waikato's student body. They ranged in age between 19.4 - 51.6 years ($M =$ 29.7 years, $SD = 9.6$ years). Each participant's threshold value estimates were titrated using the two-down/one-up single staircase method (Levitt, 1971) and the green colour spectrum ranged between 75 – 112 of the colour values, with a mean of 100.74 ($SD = 8.98$).

3.2 Signal Detection Theory Measures

It was found that the low-prevalence condition produced greater miss errors ($M = 49.5\%, SD = 16.8\%$) compared to the high-prevalence condition ($M = 26.0\%, SD = 19.4\%$), demonstrating the emergence of the low prevalence effect using the current study's methodology ($F_{(1,19)} = 32.81, p < .001, \eta^2_P = 0.63$). Conversely, the high-prevalence condition produced more false alarm errors ($M = 2.5\%, SD = 3.3$) than the low-prevalence condition ($M = 1.4\%, SD = 2.5$); however, this difference between the means was not found to be significant.

Participants who reported a probability of either zero or one for their hit rate or false alarm rate were excluded from this analysis. The decision criterion, as measured by c (Stanislaw & Todorov, 1999), shifted from a conservative criterion in the low-prevalence condition ($M = 1.23, SD = 0.49$) to a more liberal criterion ($M = 0.56, SD = 0.30$) in the high-prevalence condition ($F_{(1,10)} = 48.90, p < .001, \eta_P^2 = 0.83$) while sensitivity, as measured by d' , was found to be greater in the high-prevalence condition ($F_{(1,10)} = 5.15, p < .05, \eta_P^2 = 0.34$).

3.3 Response Latency

Response latency decreased in the low-prevalence condition. When the first and last twenty trials in each condition were analysed, a significant decrease in response latency was observed over the duration of the low-prevalence condition ($F_{(1,18)} = 53.9, p < .001, \eta_P^2 = 0.75$). A significant decrease in latency was not observed across the high-prevalence condition. When these start response latencies were subtracted from the finish latencies to provide a difference measure in latency between conditions, the difference was found to be significant ($F_{(1,18)} = 31.53, p < .001, \eta_P^2 = 0.64$). This change in the response pattern was found within 50 trials of the condition change, with increased latency in the high-prevalence condition ($M = 1757$ ms$, SD = 651$ ms) compared to the low-prevalence condition ($M = 1123$ ms$, SD = 334$ ms) ($F_{(1,18)} = 43.46, p < .001, \eta_P^2 = 0.71$).

3.4 Discussion

The familiar low prevalence effect was produced using the current study's methodology with increased miss errors across the low-prevalence condition which was consistent with findings from previous LPE literature (Wolfe, Horowitz & Kenner, 2005; Wolfe et al. 2007; Rich et al. 2008. Horowitz, 2017). Conversely, false alarm errors increased in the medium-prevalence condition which is consistent with a trade-off between miss and false alarm errors as target prevalence increases (Wolfe & Van Wert, 2010). These observations in error rates suggest a conservative shift in the decision criterion under low-prevalence conditions (Green

& Swets, 1966; Horowitz, 2017; Wolfe & Van Wert, 2010).

Adding further weight to the Criterion Shift Hypothesis was the observation that the decision criterion (c) shifted from a conservative criterion under the low-prevalence condition to a more liberal criterion under medium-prevalence conditions. However, the sensitivity index d' was found to be greater in the medium-prevalence condition despite there being no change in the experimental stimuli across prevalence conditions. This observation in sensitivity is not uncommon in prevalence studies with small samples. Horowitz's (2017) meta-analysis study showed that while some studies observed improved stimulus sensitivity under low-prevalence conditions and some the opposite; overall, there was little evidence for an effect of prevalence on sensitivity, in agreement with Gur et al. (2003). Unfortunately, this change in sensitivity across the prevalence conditions suggests that the d' assumption of equal variance may have been violated (Macmillan & Creelman, 2004). This observation questions the validity of both the sensitivity index (d') and the decision criterion (c) in the current study.

It was observed that response latencies decreased throughout the low-prevalence condition. This finding is consistent with Wolfe, Horowitz and Kenner (2005) who observed that participants were faster to report the target as absent rather than present and this was due to target-absent latencies speeding up under low-prevalence conditions, rather than target-present latencies slowing down. This observation suggests that participants may require less evidence to convince themselves that a display did not contain a target under low-prevalence conditions (Horowitz, 2017).

According to the Multiple Decision Model (Wolfe & Van Wert, 2010), when a participant shifts from a medium-prevalence condition to low-prevalence, the identification threshold becomes more conservative, and the termination threshold is lowered. As a result, targets are missed more often, both because some targets will be mistakenly identified as a distractor, and because participants are more likely to terminate their search before the target has been identified. This decrease in the termination threshold explains why target-absent latencies become faster under low-prevalence conditions because the participant quits after searching only a small proportion of the items in the display. Most of the time this search strategy works, simply because targets are rare.

In the current experiment, this change in the response pattern was found within 50 trials of the condition change, with greater response latencies in the medium-prevalence condition relative to the low-prevalence condition. This is consistent with Wolfe and Van Wert's (2010) observation that the inflection points of the criterion function lagged the underlying prevalence function by approximately 50 trials when they varied target prevalence in a sinusoidal fashion, suggesting that participants may integrate prevalence information over this timeframe.

The current study's findings constitute the standard prevalence effect, in that miss errors increased and false alarm errors decreased under low-prevalence conditions. In the response latency domain, a strong effect was observed with responses becoming faster under low-prevalence conditions, which was not observed under medium-prevalence conditions. This demonstrates that we were able to shift the decision criterion and, henceforth, the participant's response patterns simply by changing the schedule of signal presentation using the current study's methodology. While it came as no surprise that a phenomenon as robust as the prevalence effect manifested itself using a simple colour singleton feature task, it does give us confidence that we can observe patterns in eye movements and how these patterns change in other experiments.

Experiment One

In Experiment One the participant's threshold value estimates were individually titrated using the two-down/one-up single staircase method to ensure experimental integrity across all experiments. However, in this experiment, all participants experienced the same target stimulus colour value during the main experimental condition. This stimulus colour component was set at the minimum titrated stimulus value found in the Preliminary Experiment $(0 \text{ red}, 75 \text{ green}, 255 \text{ blue})$. This resulted in a large discrepancy between the target (75 green colour value) and distractor colour values (128 green colour value) and means that subsequent titrated experiments should have greater perceptual difficulty than the current Experiment.

In Experiment One, eye-tracking technology was used to analyse the eye-movement behaviour of participants, and these measures include fixation length, saccade scanpath trajectories, and the number of fixations across each trial.

4.1 Participants

Twenty participants (4 male & 16 female) were recruited from the University of Waikato's student body. They ranged in age between 18.0 – 36.7 years ($M =$ 21.8 years, $SD = 5.2$ years). One participant was excluded from the eye-movement analysis due to inconsistent eye-tracking data. These inconsistencies were the result of the eye-tracker failing to track the participant's eye movements.

4.2 Signal Detection Theory Measures

As with the Preliminary Experiment, participants who produced a probability of either zero or one for their hit or false alarm rates were excluded from this analysis and this resulted in nine participants being excluded. The indices for d' and c were calculated as follows (Stanislaw & Todorov, 1999):

$$d' = z(Hit) - z(FalseAlarm) \tag{4.1}$$

$$c = \frac{z(Hit) + z(FalseAlarm)}{2} \tag{4.2}$$

The decision criterion, as measured by c, shifted from a conservative criterion in the low-prevalence condition ($M = 0.89, SD = 0.60$) to a more liberal criterion ($M = 0.46, SD = 0.35$) in the medium-prevalence condition, and this difference between the criterion means was found to be significant ($F_{(1,9)} = 14.30, p < .05, \eta_P^2 = 0.61$). No difference in sensitivity, as measured by d', was found between prevalence conditions.

These signal detection theory results are what one would expect to see in such an experiment. However, further investigation was required to ascertain whether the underlying assumptions behind the signal detection indices of d' and c had been violated. If the underlying signal and noise distributions are normally distributed with equal variance (i.e., parametric) the slope of the z-ROC (normalized Receiver Operating Characteristic) curve would be unity or one. When the slope is less than one, this indicates that the assumptions of normality and equal variance for the signal and noise distributions have been violated (Macmillan & Creelman, 2004). Therefore, the z-transformed hit rate (i.e., the proportion of trials where a participant correctly responded “present”) was plotted against the z-transformed false alarm rate (i.e., the proportion of trials where a participant responded “present” when no target stimulus was presented) for each participant and a linear regression line was fitted for both prevalence conditions (see Figures 4.2 & 4.1). Unfortunately, the underlying assumptions of normality and equal variance for the signal detection indices of d' and c were found to be violated with the low-prevalence regression line having a slope of 0.005 and an intercept of 0.02, while the medium-prevalence regression line had a slope of 0.188 and an intercept

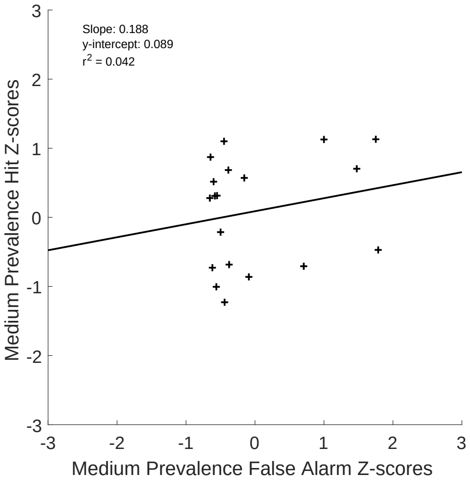

Figure 4.1: The fitted linear regression line to medium-prevalence hit and false alarm z-scores data.

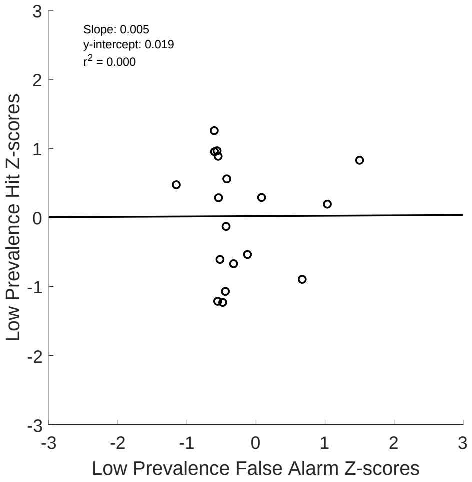

Figure 4.2: The low-prevalence hit and false alarm zscores with a regression line fitted to the data.

of 0.088. As can be seen in Figures 4.1 and 4.2 there was a large amount of participant variance in the sample resulting in a poor fit for the regression lines, in both the low-prevalence condition ($R^2 = 0.00$) and the medium-prevalence condition ($R^2 = 0.04$). As a result, the parametric signal detection indices of d' and c were unable to be used and non-parametric indices were investigated.

Given the non-unity of the z-ROC slopes in Experiment One, we investigated the non-parametric alternative measures of d_a (sensitivity index) and c_2 (decision criterion index), used by Wolfe and Van Wert (2010) and recommended by (Godwin, Menneer, Cave, & Donnelly, 2010) for such cases. The following equations were used to calculate d_a and c_2:

$$d_a = \left(\frac{2}{1+s^2}\right)^{\frac{1}{2}} [z(Hit) - sz(FalseAlarm)] \tag{4.3}$$

$$c_2 = \frac{-s}{1+s} [z(Hit) + z(FalseAlarm)] \tag{4.4}$$

These measures incorporated the slope of the z-ROC curve (s) with the z-transformed hit rate and the z-transformed false alarm rate. However, given the sample variance and the difficult fit of the regression lines, as indicated by the r^2 values, it was decided that these measures would not be appropriate.

We then applied the non-parametric measures of log d (i.e., stimulus discriminability or sensitivity) and log B (i.e., response bias or decision criterion). For these measures, M. C. Davison and Tustin (1978) applied the "Generalized Matching Law" (Baum, 1974) to the standard 2 x 2 matrix of events in the signal detection yes-no paradigm and derived two independent measures of performance: stimulus discriminability or the ability to distinguish between two stimulus classes (i.e., sensitivity):

$$log d = 0.5[log(\frac{Hits}{Misses} * \frac{Correct\ rejections}{False\ Alarms})] \tag{4.5}$$

and response bias (i.e., decision criterion):

$$log B = 0.5[log(\frac{Hits}{Misses} * \frac{False\ Alarms}{Correct\ rejections})] \tag{4.6}$$

Response bias, as measured by log B, shifted from a conservative bias in the low-prevalence condition ($M = -1.90, SD = 1.33$) to a more liberal bias ($M = -0.99, SD = 0.75$) in the medium-prevalence condition, and this difference between the means was found to be significant ($F_{(1,9)} = 13.52, p < .05, \eta_P^2 = 0.60$). No significant difference in stimulus discriminability, as measured by log d, was found between prevalence conditions.

To summarize, the available parametric signal detection measures (i.e., d' and c) had issues with normality and equal variance between the signal and noise distributions, while the non-parametric measures of d_a and c_2 had issues with the fit of the regression lines. Therefore, we applied Davison and Tustin's (1978) non-parametric summary signal detection measures for stimulus discriminability (i.e., log d) and response bias (i.e., log B).

Experiment One produced the low prevalence effect with greater miss errors in the low-prevalence condition ($M = 32.3\%, SD = 21.1\%$) compared to the medium-prevalence condition ($M = 13.2\%, SD = 8.5\%$) and this difference between the means was found to be significant ($F_{(1,19)} = 27.91, p < .001, \eta_P^2 = 0.60$). Conversely, the medium-prevalence condition produced more false alarm errors ($M = 3.2\%, SD = 5.9\%$) compared to the low-prevalence condition ($M = 2.2\%, SD = 4.3\%$). However, this difference was not found to be significant (see Figure 4.3).

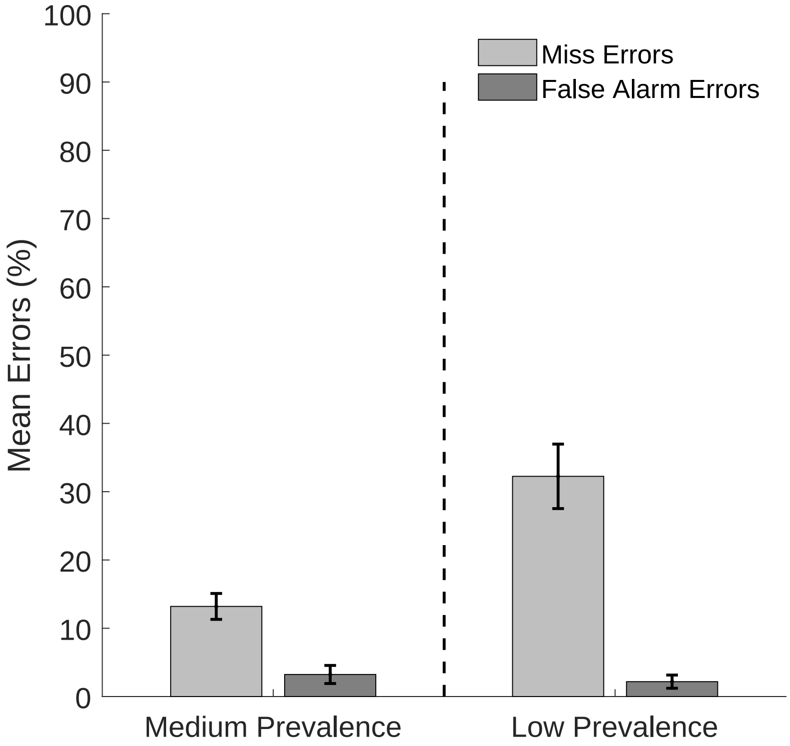

Figure 4.3: Mean signal detection errors found in the medium-prevalence condition and the low-prevalence condition.

Of particular interest was the observation that when compared to the correct rejection mean, there was increased eye-movement behaviour for the first correct rejection immediately following a correct target-present response (see Figures 4.4). This pattern of behaviour was observed for the medium-prevalence eye-movement measures of fixation length ($F_{(1,18)} = 24.81, p < .001, \eta^2_P = 0.58$), saccadic scan-path trajectory ($F_{(1,18)} = 26.11, p < .001, \eta^2_P = 0.59$), and the number of fixations ($F_{(1,18)} = 32.13, p < .001, \eta^2_P = 0.64$), and the low-prevalence measures of scan-path trajectory ($F_{(1,18)} = 16.82, p < .05, \eta^2_P = 0.48$), and the number of fixations ($F_{(1,18)} = 24.93, p < .001, \eta^2_P = 0.58$).

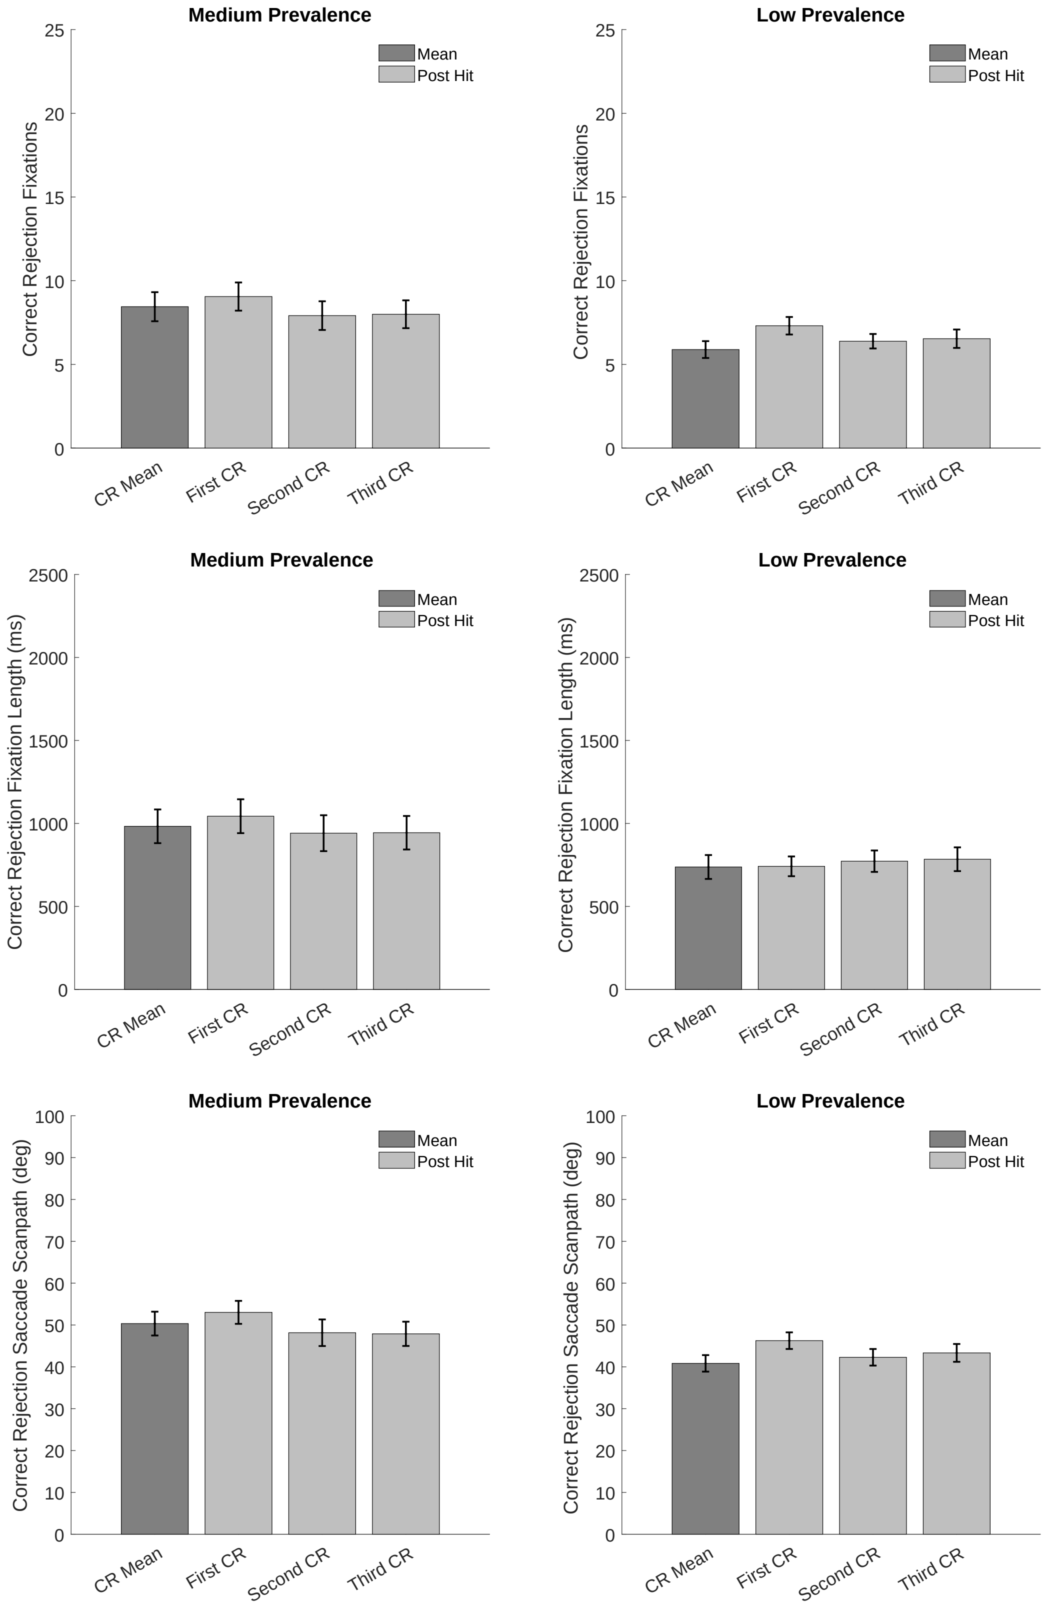

Figure 4.4: Eye-movement behaviour for the correct rejection mean across each condition, the first correct rejection immediately following a hit, the second consecutive correct rejection following a hit, and the third consecutive correct rejection following a hit for the low-prevalence condition and the medium-prevalence condition.

Table 4.1: Eye-movement behaviour for the correct rejection mean and the first correct rejection immediately following a hit across the prevalence conditions. Means are followed by their parenthesised standard deviations. CR is an abbreviation for correct rejection.

	Medium Prevalence		Low Prevalence	
	CR Mean	Post-hit CR	CR Mean	Post-hit CR
Experiment One				
Fixations (*n*)	8.45(3.79)	9.05(3.67)	5.89(2.20)	7.31(2.28)
Length (*ms*)	982.69(442.47)	1043.50(444.77)	737.52(313.76)	741.50(259.36)
Scanpath (*deg*)	50.34(12.40)	53.02(12.00)	40.83(8.61)	46.26(8.67)

correct rejections following the first post-hit correct rejection, nor was it found for misses, false alarms, and subsequent hits immediately following a hit.

4.3 Eye-movement Decay

We know that target-absent response latencies under low-prevalence conditions become faster throughout this condition and that this phenomenon seldom occurs under medium-prevalence conditions (Wolfe et al., 2005) and Horowitz (2017) attributes this pattern of behaviour to a shift in the decision criterion. Specifically, the decision criterion becomes increasingly conservative as target prevalence decreases. This was exactly what we found, a greater amount of decay in eye-movement behaviour throughout the low-prevalence condition compared to the medium-prevalence condition and we were interested in how operant learning theory equated to this pattern of behaviour throughout each prevalence condition. We hypothesize that this observed pattern of behaviour is an operant behaviour controlled by environmental contingencies (Madelain, Paeye, & Darcheville, 2011) rather than the "Criterion Shift Hypothesis" (Horowitz, 2017). Therefore, we fitted an exponential decay curve function to each eye-movement data set for both the medium- and low-prevalence conditions and analysed the curve parameters. Figure 4.5 shows the eye-movement behaviour found in the low- and medium-prevalence conditions with an exponential decay curve fitted to each dataset.

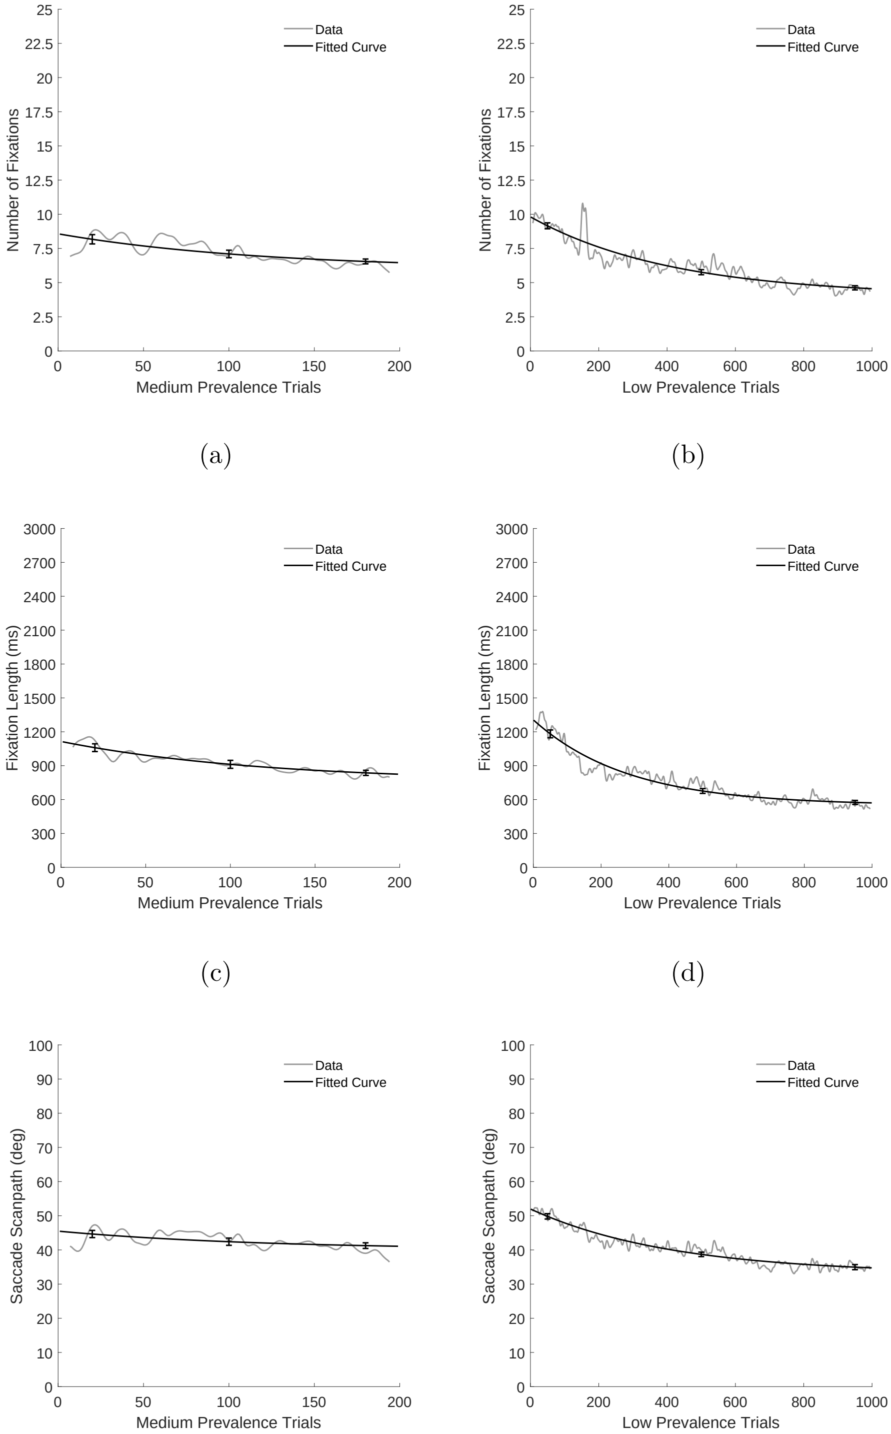

(a) (b)

(c) (d)

(e) (f)

Figure 4.5: The mean eye-movement behaviour across the participants for each prevalence condition with an exponential decay curve fitted to each dataset. The error bars represent the standard error found in a sample consisting of fifteen trials either side of the presented error bar. Note the differing number of trials within each of the two prevalence conditions, with 200 trials in the medium-prevalence condition and 1000 trials in the low-prevalence condition.

Table 4.2: This table shows the per-trial eye-movement behaviour decay curve parameters for both the medium-prevalence and the low-prevalence conditions. Mean percentages are followed by their parenthesised standard deviations, and we remind the reader that the rate of decay before reaching an asymptote is represented in trials.

	Medium Prevalence			Low Prevalence		
	Decay	Rate	Asymptote	Decay	Rate	Asymptote
Experiment One						
Fixations (*n*)	2.63(5.02)	82.56(43.74)	6.44(1.78)	5.11(3.29)	324.03(165.92)	4.67(1.99)
Length (*ms*)	461.61(427.83)	63.44(39.00)	847.36(237.03)	795.07(499.40)	307.84(174.43)	567.80(280.12)
Scanpath (*deg.*)	6.16(8.94)	88.97(39.84)	42.74(5.82)	18.39(11.43)	347.08(193.95)	35.82(7.88)

The fitting of the exponential decay curve to the number of fixation data produced a RMS value of 2.33 fixations ($SD = 0.68\ fixations$) for the low-prevalence curve fit and a RMS value of 3.25 fixations ($SD = 2.49\ fixations$) for the medium-prevalence curve fit. The RMS value (Root-Mean-Squared value) provides a measure of the variance of the dataset around the exponential curve, with smaller values indicating a better curve fit. These values indicate that there was a greater amount of variance in the curve fitted to the medium-prevalence condition data and that the variance was reduced in the low-prevalence condition.

As can be seen in Table 4.2, the magnitude of the curve decay for fixations (as represented on the Y-axis) was greater in the low-prevalence condition, however this difference between the prevalence conditions was not found to be significant. Note the large amount of sample variance in the medium-prevalence sample. A difference was found in the rate of this behavioural decay (the time-constant to reach asymptote as represented on the X-axis) ($F_{(1,18)} = 40.46, p < .001, \eta_P^2 = 0.69$), with participants taking longer to reach a steady state of behaviour in the low-prevalence condition (see Figure 4.5). A difference was also found between prevalence conditions for the curve asymptotes (as represented on the Y-axis) ($F_{(1,18)} = 21.80, p < .001, \eta_P^2 = 0.55$). Specifically, once a steady state of fixational behaviour had been reached, this level of behaviour was lower in the low-prevalence condition.

These curve parameter results inform us that the participants were able to reach a steady state of fixational behaviour more rapidly and maintain a higher level of this behaviour in the medium-prevalence condition. As a result, there were a greater number of fixations found in this condition ($M = 7.54\ fixations, SD = 2.73\ fixations$) compared to the low-prevalence condition ($M = 5.95\ fixations, SD = 2.21\ fixations$), and this difference between the conditions was found to be significant ($F_{(1,18)} = 6.80, p < .05, \eta_P^2 = 0.27$).

This study incorporated the practices of previous Low Prevalence Effect studies (Godwin et al., 2015; Hout et al., 2015; Rich et al., 2008; Wolfe et al., 2005, 2007) and included their prevalence conventions for each condition (i.e., the number of trials). However, from a behavioural perspective, measuring behavioural change across the prevalence conditions when their time constants are different may be considered poor practice. Therefore, to provide a richer understanding of this behavioural change we analysed the data over the first 50 trials in each prevalence condition and observed that the mean number of fixations in the medium-prevalence condition decreased by 0.80 of a fixation, while in the low-prevalence condition, fixations decreased by 0.90 of a fixation. When we pushed this timestamp out to include the full 200 trials in the medium-prevalence condition and the first 200 trials in the low-prevalence condition, we observed that the mean number of fixations in the medium-prevalence condition decreased by 1.91 fixations, and in the low-prevalence condition the number of fixations decreased by 2.72 fixations, and we know that the number of fixations in the low-prevalence condition continued to decay until trial number 324. Note that this trial number was generated from the curve parameter equation.

The fitting of the exponential decay curve to the fixation length data produced a RMS value of 300.34ms ($SD = 124.06ms$) for the low-prevalence curve fit and a RMS value of 406.13ms ($SD = 224.12ms$) for the medium-prevalence curve fit.

When we examined the curve parameters for the fixation length data, we found that the magnitude of behavioural decay was greater in the low-prevalence condition. However, as with the fixation data, there was also a large amount of sample variance within the medium-prevalence sample, and this difference between the means was not found to be significant. However, there was a significant difference found in the rate of this behavioural decay in fixation length ($F_{(1,18)} = 36.09, p < .001, \eta_P^2 = 0.67$), with participants taking longer to reaching a steady state of behaviour in the low-prevalence condition (see Figure 4.5). A difference was also found between prevalence conditions for the curve asymptotes ($F_{(1,18)} = 27.30, p < .001, \eta_P^2 = 0.60$). Specifically, once a steady state of behaviour had been reached, fixation times were shorter in the low-prevalence condition.

These fixation length curve parameter results inform us that, participants were able to reach a steady state of behaviour more rapidly and maintain a higher level of this behaviour in the medium-prevalence condition. As there was a greater number of per-trial fixations found in the medium-prevalence condition, it should come as no surprise that there was a greater amount of time fixating throughout a trial in this condition ($M = 974.87ms, SD = 315.50ms$) compared to the low-prevalence condition ($M = 753.85ms, SD = 318.19ms$), and this difference between the conditions was found to be significant ($F_{(1,18)} = 12.55, p < .05, \eta_P^2 = 0.41$).

When we examined this data over the first 50 trials within each experimental condition, we found that the mean length of fixations in the medium-prevalence condition decreased by 133.12ms while in the low-prevalence condition, fixation lengths decreased by 138.56ms. When we pushed this timestamp out to include the full 200 trials in the medium-prevalence condition and the first 200 trials in the low-prevalence condition, we observed that the fixation lengths in the medium-prevalence condition decreased by 308.75ms and in the low-prevalence fixation lengths decreased by 412.52ms, and we know that fixation lengths continued to decrease in the low-prevalence condition until trial 307. Note, that this trial number was generated from the curve parameter equation.

The fitting of the exponential decay curve to the saccadic scanpath trajectory data produced an RMS value of 10.04° ($SD = 2.1°$) for the low-prevalence condition curve fit and an RMS value of 13.38° ($SD = 6.20°$) for the medium-prevalence condition curve fit.

The curve parameters for the saccadic scanpath trajectory data showed that the magnitude of behavioural decay was greater in the low-prevalence condition. As with the previous eye-movement data, this difference in behavioural decay between the conditions was not found to be significant and again, note the large amount of sample variance within the medium-prevalence (see Table 4.2). However, a difference was found in the rate of this behavioural decay ($F_{(1,18)} = 44.98, p < .001, \eta_P^2 = 0.71$), with participants reaching a steady state of behaviour more rapidly in the medium-prevalence condition. We also found a significant difference in the curve asymptotes ($F_{(1,18)} = 44.70, p < .001, \eta_P^2 = 0.71$). Specifically, once a steady state of behaviour had been reached, search trajectories were shorter in the low-prevalence condition.

These curve parameters results inform us that participants were able to reach a steady state of behaviour more rapidly and maintain a higher level of this behaviour in the medium-prevalence condition, and this was reflected in the mean trial scanpath trajectories, with longer trajectories occurring in this condition ($M = 45.32°, SD = 8.00°$) compared to the low-prevalence condition ($M = 40.81°, SD = 8.27°$), and this difference between the conditions was found to be significant ($F_{(1,18)} = 10.68, p < .05, \eta_P^2 = 0.37$).

When we examined the data over the first 50 trials, we found that the scanpath trajectories in the medium-prevalence condition decreased by 1.59°, while in the low-prevalence condition, this behaviour decreased by 2.84°. When we pushed this timestamp out to 200 trials, we observed that the scanpath trajectories in the medium-prevalence condition decreased by 3.76° and in the low-prevalence condition they decreased by 8.66°, and we know that trajectories continued to decay in the low-prevalence condition until trial 347. .

Overall, this pattern of behaviour mimics that found in previous eye-movement behavioural measures. Again, participants reached a steady state of behaviour more rapidly and maintained a higher rate of this behaviour in the medium-prevalence condition. Whereas, participants in the low-prevalence condition were slower to reach a steady state of behaviour, and the magnitude of their behavioural change was greater than that found in the medium-prevalence condition. These results suggest that eye-movement behaviour is an operant controlled by environmental contingencies and the outcome of these contingencies is resulting in the Low Prevalence Effect.

4.4 Stimulus Density Map

We were also interested in the progression of eye-movement behaviour throughout a trial and when we first examined the scanpath trajectory data, we observed that the participants often examined a portion of the stimulus matrix which did not contain the target stimulus and, on occasions, this resulted in target misses. On occasions, we also observed that the participants would saccade past the target, or even directly over the target stimulus, to reach a search location of their choice (see Figures 2.3 & 2.4). We wondered what was capturing their attention and guiding their eye movements away from the target stimulus in what should be a relatively simple search task. Therefore, we constructed the "Stimulus Density Map" which represents the portion of the stimulus matrix with the densest population of droplet stimulus items. We hypothesize that eye-movement behaviour was being guided towards this centroid zone.

The mean surface area of the centroid zone within the stimulus density map constituted 10.18% (*SD* = 0.11%) of the stimulus matrix box area and represents the proportion of the stimulus matrix with the greatest stimulus density

($\geq$ 0.8 *density probability*). If the centroid zone sliced through a stimulus this stimulus was counted as being within the centroid zone. As a result, the mean number of stimuli within the centroid zone was 6.22 ($SD = 0.78$) of the 15 stimuli presented within the stimulus matrix box on each trial, and on target-present trials the probability of the target stimulus appearing within the centroid zone was 0.146 ($SD = 0.011$).

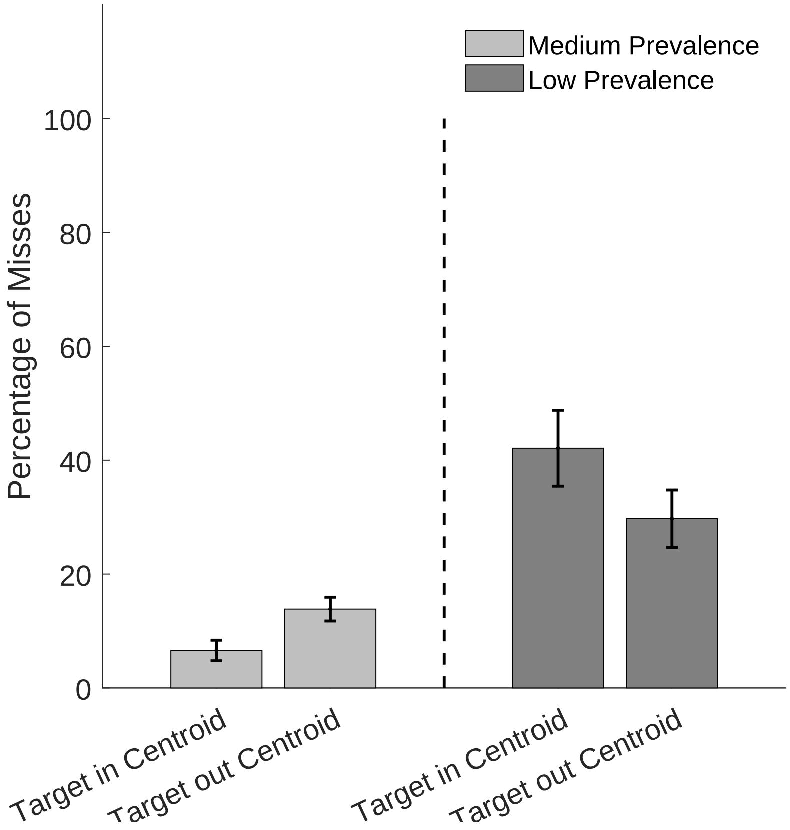

Figure 4.6: The miss errors that occurred when the target stimulus was presented inside the centroid zone and those that occurred when the target stimulus was presented outside the centroid zone across the prevalence conditions.

It was observed that only 6.58% ($SD = 7.90\%$) of miss errors occurred when the target was presented within the centroid zone and 13.85% ($SD = 9.12\%$) occurred when the target was presented outside the centroid zone in the medium-prevalence condition ($F_{(1,18)} = 17.12, p < .001, \eta_P^2 = 0.49$). In the low-prevalence condition, 42.11% ($SD = 29.06\%$) of miss errors occurred when the target was presented within the centroid zone and 29.72% ($SD = 21.96\%$) occurred when the

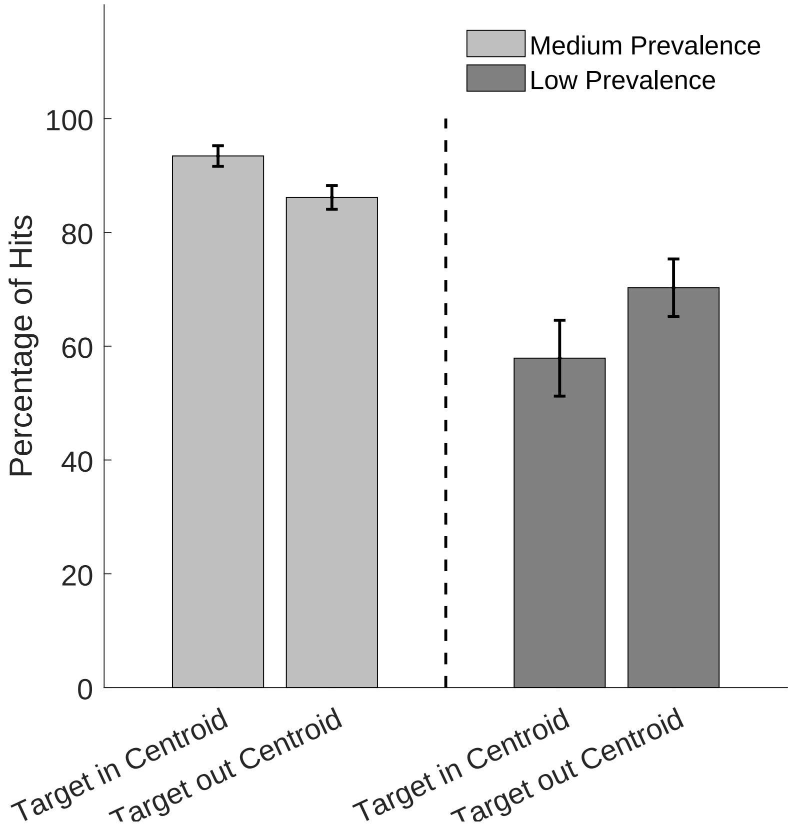

Figure 4.7: The hits that occurred when the target stimulus was presented inside the centroid zone and those that occurred when the target stimulus was presented outside the centroid zone across the prevalence conditions.

target was presented outside the centroid zone ($F_{(1,18)} = 4.92, p < .05, \eta^2_P = 0.22$). These results suggest that the increase in miss errors which occur under low-prevalence conditions (Low Prevalence Effect), are primarily occurring within the centroid zone (see Figure 4.6). Furthermore, significant differences were found for miss errors between the prevalence conditions when the target appeared within the centroid zone ($F_{(1,18)} = 32.94, p < .001, \eta^2_P = 0.65$) and outside this zone($F_{(1,18)} = 16.30, p < .001, \eta^2_P = 0.48$).

Of the hits which occurred outside the centroid zone in the medium-prevalence condition, 30.86% (SD = 8.85) occurred before entering the centroid zone, and 69.14% (SD = 8.85) occurred after leaving this zone. In the low-prevalence condition, 31.53% (SD = 18.49) of hits occurred before entering the centroid zone and 68.47% (SD = 18.49) occurred after leaving this zone. No significant

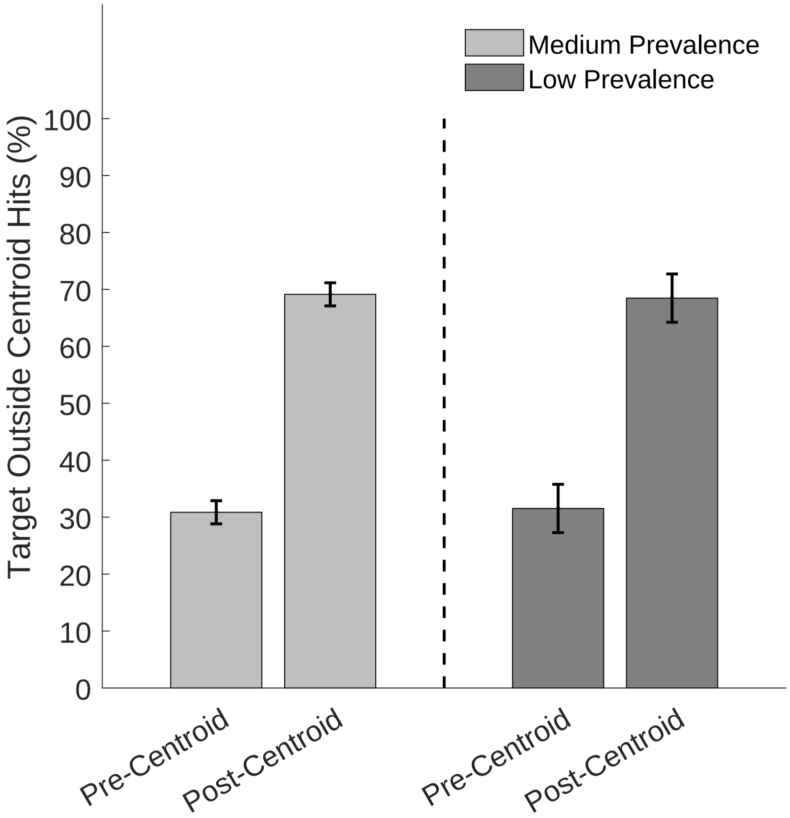

Figure 4.8: The proportion of hits that occurred before entering the centroid zone and the proportion of hits that occurred after leaving the centroid zone across the prevalence conditions.

differences were found between the prevalence conditions for the pre- and post-centroid hits.

When we examined the eye-movement behaviour within the centroid zone, we observed that there was a greater proportion of behaviour within this zone compared to the proportion of the designated centroid area within the stimulus matrix box ($M = 10.18\%, SD = 0.11\%$) for all three eye-movement measures (see Figure 4.9).

Expanding on this, there was a greater proportion of fixations occurring within the centroid zone in the low-prevalence condition ($M = 37.31\%, SD = 6.83\%$) compared to the medium-prevalence condition ($M = 31.68\%, SD = 6.86\%$) and this difference between the conditions was found to be significant ($F_{(1,18)} = 26.14, p < .001, \eta^2_P = 0.59$). These differences between the centroid zone area and the cor-

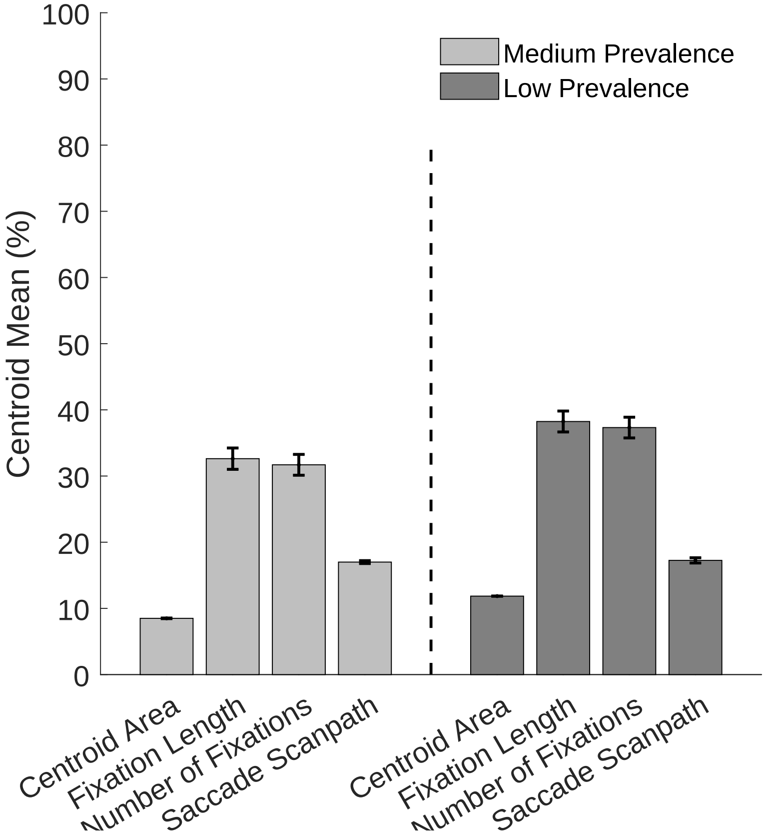

Figure 4.9: The percentage of the centroid area within the stimulus matrix, the percentage of the fixation lengths which occur within the centroid zone, the percentage of fixations which occur within the centroid zone, and the percentage of saccade scanpath trajectories found within the centroid zone across the prevalence conditions.

responding proportions of eye-movement behaviour across prevalence conditions are shown in Figure 4.9.

This was also the case for the length of these fixations, with a greater proportion of fixation time occurring within the centroid zone in the low-prevalence condition ($M = 38.22\%, SD = 6.89\%$) compared to the medium-prevalence condition ($M = 32.60\%, SD = 7.03\%$) ($F_{(1,18)} = 23.51, p < .001, \eta_P^2 = 0.57$). No significant difference was found between the conditions for the proportion of the saccadic scanpath trajectory within the centroid zone with only 0.25° separating the means.

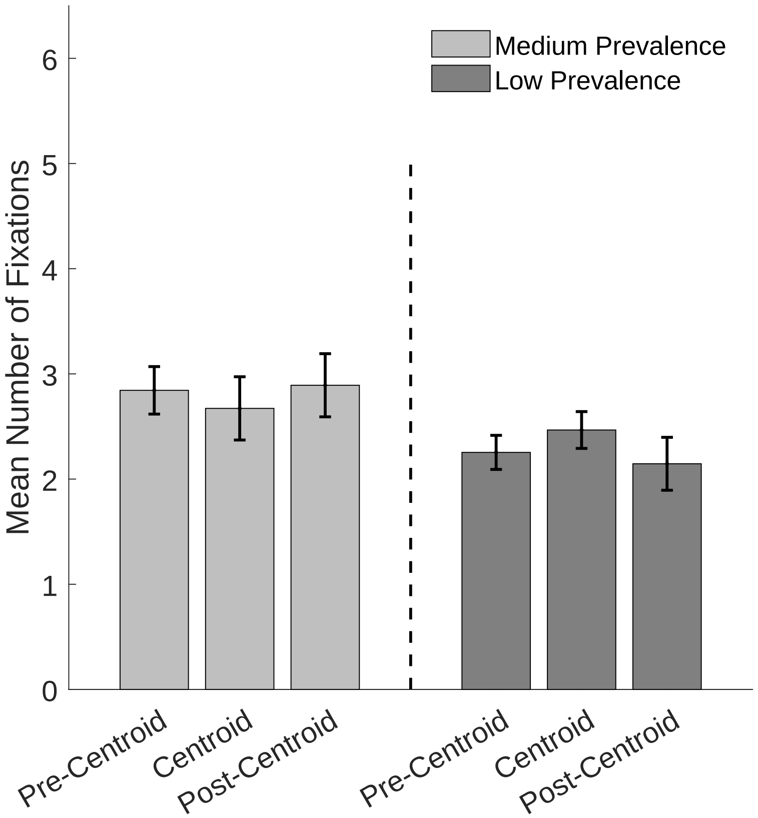

Figure 4.10: The average number of fixations before entering the centroid zone, the number of fixations within the centroid, and the number of fixations after leaving the centroid zone across prevalence conditions.

Table 4.3: This table shows the eye-movement behaviour before entering the centroid zone (Pre-), the eye-movement behaviour within the centroid zone (Centroid), and the eye-movement behaviour after leaving the centroid zone (Post-) for both the medium-prevalence and the low-prevalence conditions. Mean percentages are followed by their parenthesised standard deviations.

	Medium Prevalence			Low Prevalence		
	Pre-	Centroid	Post-	Pre-	Centroid	Post-
Experiment One						
Fixations (*n*)	2.80(0.97)	2.66(1.31)	2.87(1.30)	2.25(0.71)	2.46(0.76)	2.14(1.09)
Length (*ms*)	394.60(140.37)	379.37(161.61)	372.16(148.58)	324.40(111.72)	357.01(125.71)	261.57(152.19)
Scanpath (*deg*)	8.89(0.38)	7.89(1.53)	18.55(4.61)	8.59(0.41)	7.27(1.70)	15.17(4.56)

As can be seen in Table 4.3, the number of fixations in the medium-prevalence condition were similar across the pre-centroid, centroid, and post-centroid data categories, with 33.60% of fixations occurring before entering the centroid zone, 31.95% occurred within the centroid zone, and 34.45% occurred after leaving this zone, no significant differences were found between the means. In the low-prevalence condition, fixations were greater within the centroid zone, with 32.84% of fixations occurring before entering the centroid zone, 35.92% occurring within the centroid zone, and 31.23% occurring after leaving this zone, no significant differences were found.

Table 4.3 shows the length of these fixations throughout a trial for each data category across the medium- and the low-prevalence conditions which provide a measure of the information processing time within each data category. In the medium-prevalence condition, these lengths equate to a mean pre-centroid length of 140.93ms, 142.62ms for each fixation within the centroid zone, and 129.67ms for each fixation post-centroid, no significant differences were found between means. In the low-prevalence condition, this equates to a mean pre-centroid length of 144.18ms for each fixation, 145.13ms within the centroid zone, and 122.23ms for each fixation post-centroid. As can be seen in Table 4.3, the greatest fixation lengths occurred within the centroid zone ($F_{(2,17)} = 5.60, p < .05, \eta_P^2 = 0.24$), with the posthoc analysis showing that the main significant difference occurred between the centroid and post-centroid means.

The saccadic scanpath trajectories provide a measure of the total distance searched within each data category. In the medium-prevalence condition, these trajectories equate to a mean pre-centroid saccadic distance between each fixation of 3.18° , 2.97° within the centroid zone, and 6.46° post-centroid. In the low-prevalence condition, this equates to a mean pre-centroid saccadic distance of 3.82° between each fixation, 2.96° within the centroid zone, and 7.09° post-centroid. These results across prevalence conditions inform us that the saccadic distance travelled between fixations to reach the centroid zone was longer than the distance travelled between fixations within the centroid zone, and the longest saccadic distance travelled occurred between fixations after leaving this zone. These results come as no surprise given that the surface area of the centroid zone only constituted 10.18% ($SD = 0.11\%$) of the stimulus matrix box area.

We observed that these pre-, post- and centroid trajectories in the medium-prevalence condition were shortest within the centroid zone ($F_{(2,17)} = 115.81, p < .001, \eta_P^2 = 0.87$), with the posthoc analysis showing that the main significant difference occurred between the centroid and the post-centroid means ($p < .001$), a larger p-value occurred between the pre- and post-centroid means ($p < .001$), and the largest p-value occurred between the pre-centroid and centroid means ($p = .027$). This trend in trajectory behaviour was also seen in the low-prevalence condition ($F_{(2,17)} = 67.55, p < .001, \eta_P^2 = 0.79$), with the posthoc analysis showing that the main significant difference occurred between the centroid and the post-centroid means ($p < .001$), and the pre- and post-centroid means ($p < .001$), the largest p-value occurred between the pre-centroid and centroid means ($p = .005$).

To further our understanding of the search strategies which may be occurring, we then became interested in what the participants were doing within their pre-centroid fixations and we asked the question, "Is their eye-movement behaviour drifting towards the centre of the display for a scene *preview*, termed the centre bias phenomenon (Tatler & Vincent, 2009) before entering the centroid zone?" we hypothesize that eye movements were going to the centre of the display (centre bias) before saccading to the centroid zone. For this analysis, we placed an invisible circle in the centre of the stimulus matrix box. This circle's radius subtended 1.7° and was twice the diameter of the catchment circle area placed around each stimulus item. Figure 4.11 shows a greyscale trial image with an overlaid circle placed in the centre of the stimulus matrix. If the centre circle overlapped with the centroid zone any fixations within this overlap were counted as being within the centroid. This overlap was calculated because the centroid zone is often centred around the centre of the stimulus matrix. In this respect, the mean percentage of overlap between the centre circle and the centroid zone was 28.15% ($SD = 0.38\%$).

When we examined the number of pre-centroid fixations which occur before entering the centre circle of the stimulus matrix across prevalence conditions, we found that the medium-prevalence condition produced 1.56 ($SD = 0.57$) pre-centre fixations and the low-prevalence condition produced 1.10 ($SD = 0.43$) fixations. This difference between the prevalence conditions was found to be significant ($F_{(1,18)} = 36.84, p < .001, \eta_P^2 = 0.67$). The number of fixations within

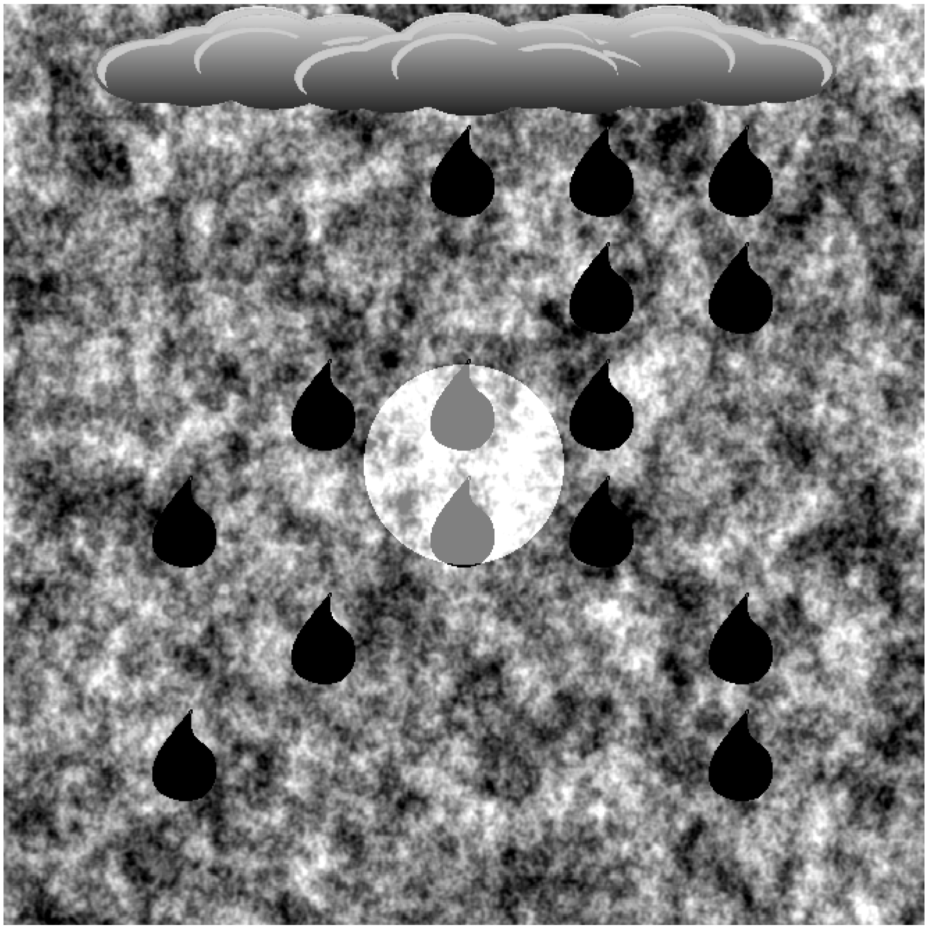

Figure 4.11: A greyscale trial stimulus image overlaid with the centre circle.

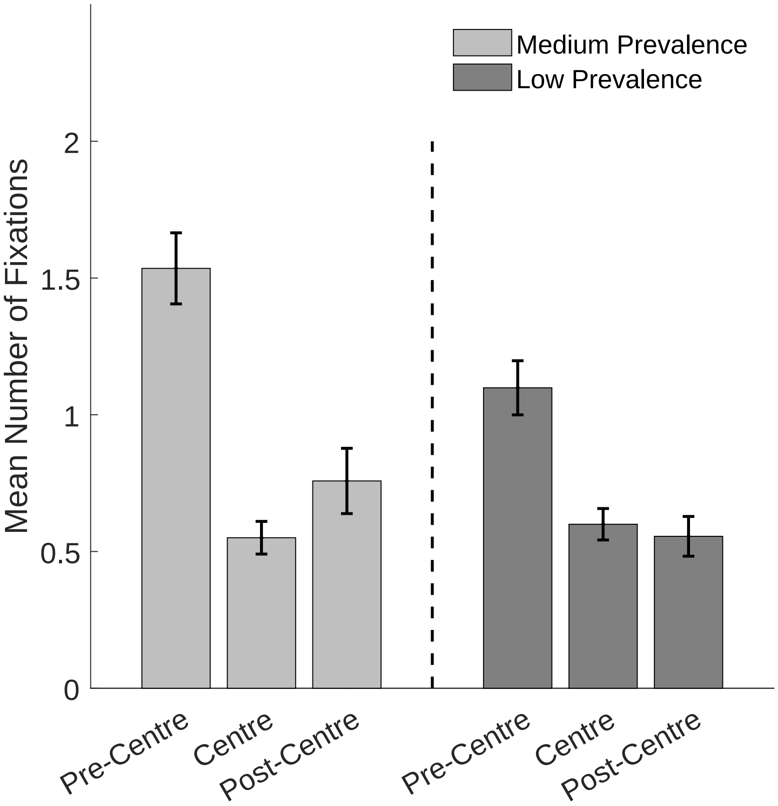

Figure 4.12: Pre-centre, centre, and post-centre fixations before entering the centroid zone.

the centre circle in the medium-prevalence condition was 0.55 ($SD = 0.26$) of a fixation and 0.60 ($SD = 0.25$) of a fixation in the low-prevalence condition ($F_{(1,18)} = 5.15, p < .05, \eta_P^2 = 0.22$). The number of pre-centroid fixations which occurred after leaving the centre circle in the medium-prevalence condition was 0.76 ($SD = 0.52$) of a fixation and 0.56 ($SD = 0.32$) of a fixation in the low-prevalence condition ($F_{(1,18)} = 8.93, p < .05, \eta_P^2 = 0.33$). These results do not support our hypothesis that eye movements were travelling to the centre of the display before saccading towards the centroid zone and suggest that they were being guided towards the centroid zone during their scene preview as a search strategy. These pre-, post-, and centre circle fixations across both prevalence conditions are shown in Figure 4.12.

4.4.1 Centroid Eye-movement Decay

The data suggests that at different stages of their search patterns, the participants change their eye-movement behaviour at different rates. To provide a richer understanding of these patterns, we asked the question, "Are different patterns of eye movements causing the differences observed between the low- and medium-prevalence conditions?" Consistent with the Reinforcement Strengthening Hypothesis, we observed that there was a greater amount of eye-movement behaviour in the medium-prevalence condition compared to the low-prevalence condition (see section Eye-movement Decay 4.3). We were interested in this behavioural change throughout a trial and how this equated to the centroid zone. We hypothesize that the data would be similar across the prevalence conditions before entering the centroid zone and within this zone. However, we hypothesize that there would be a greater amount of decay in eye-movement behaviour after leaving the centroid zone in the low-prevalence condition compared to the medium-prevalence condition. Therefore, we fitted exponential decay functions to the data and analysed the curve parameters for eye-movement behaviour before entering the centroid zone, within the centroid zone, and after leaving this zone in both the medium- and low-prevalence conditions for the number of fixations, the length of these fixations, and their scanpath trajectories to provide a richer understanding of these patterns and the underlying mechanisms and strategies which may be involved. For the reader's convenience, we report the findings for the number of fixations and only present the data for the length of these fixations, and their scanpath trajectories in table format. Figures 4.13 and 4.14 show the medium- and low-prevalence conditions pre-, post-, and centroid eye-movement behaviour with an exponential decay function fitted to each dataset.

The fitting of the exponential decay curves to the medium-prevalence condition data produced an RMS value of 2.99 fixations ($SD = 0.91\ fixations$) for the curve fitted to the pre-centroid data, an RMS value of 2.25 fixations ($SD = 1.24\ fixations$) for the curve fitted to the centroid data, and an RMS value of 2.86 fixations ($SD = 1.31\ fixations$) for the curve fitted to the post-centroid data. The fitting of the decay curves to the low-prevalence data produced a RMS value of 2.47 fixations ($SD = 0.74\ fixations$) for the curve fitted to the pre-centroid data, a RMS value of 1.89 fixations ($SD = 0.44\ fixations$) for the curve fitted to the centroid data, and a RMS value of 2.13 fixations ($SD = 0.69\ fixations$) for the curve fitted to the post-centroid data.

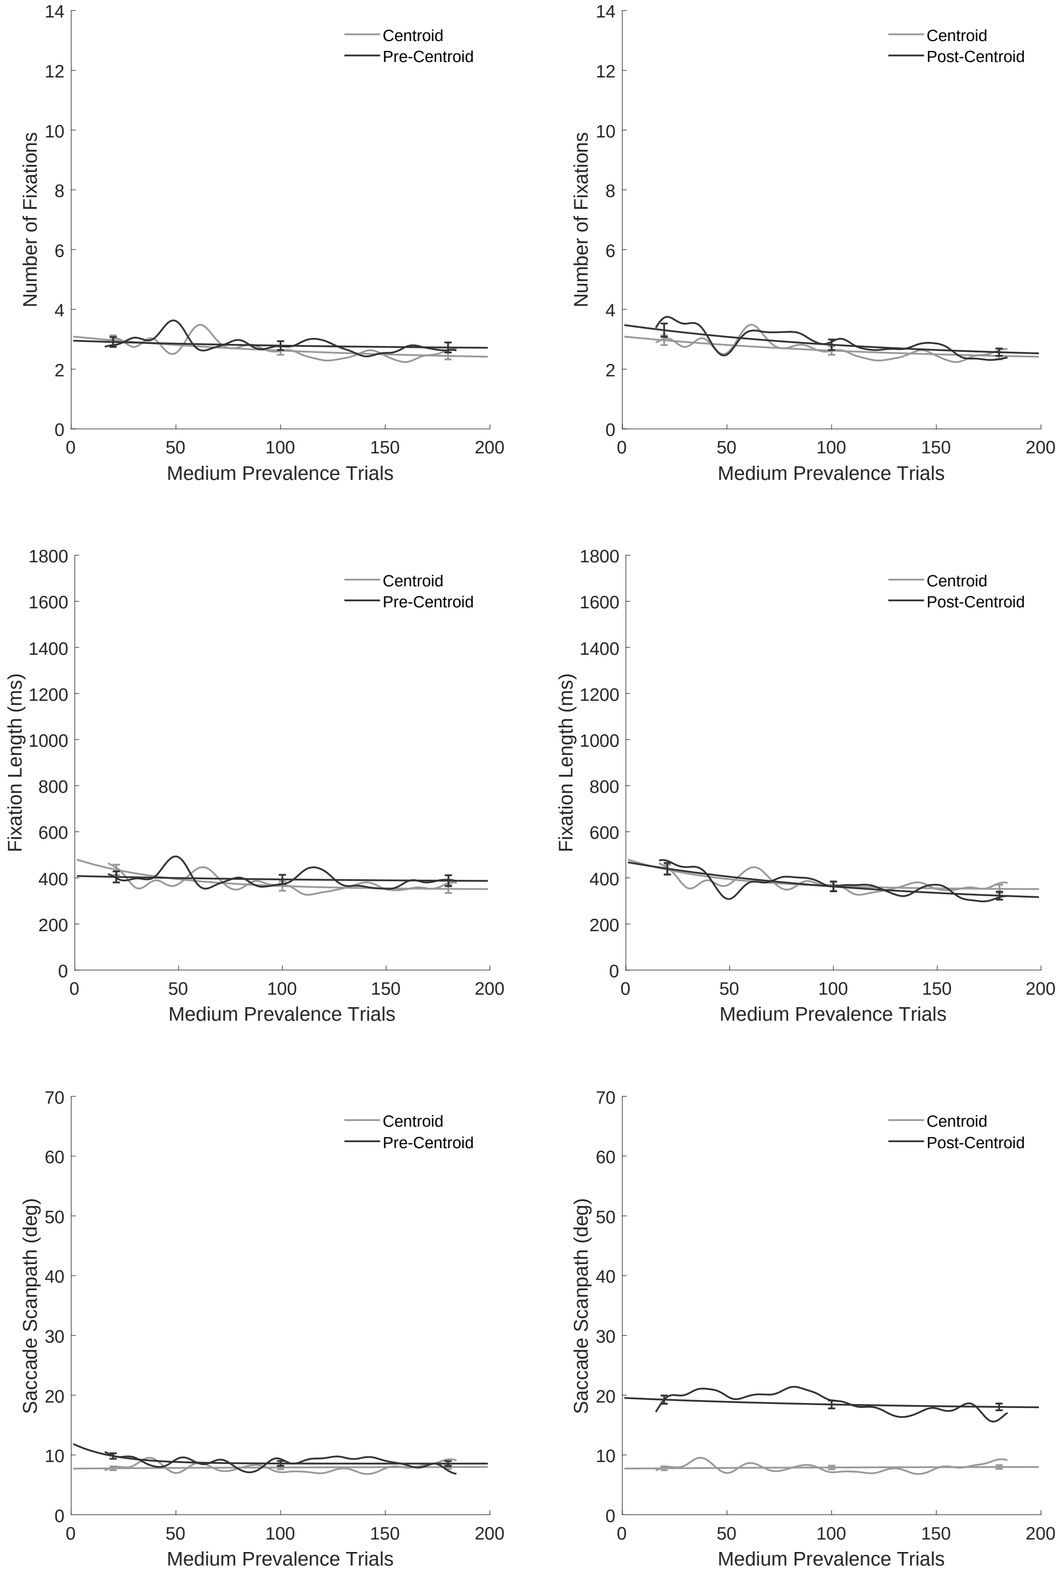

Figure 4.13: The medium-prevalence condition eye-movement behaviour before entering the centroid zone, within the centroid zone, and after leaving the centroid zone with an exponential decay function fitted to each dataset. The error bars represent the standard error found in a sample consisting of fifteen trials either side of the presented error bar.

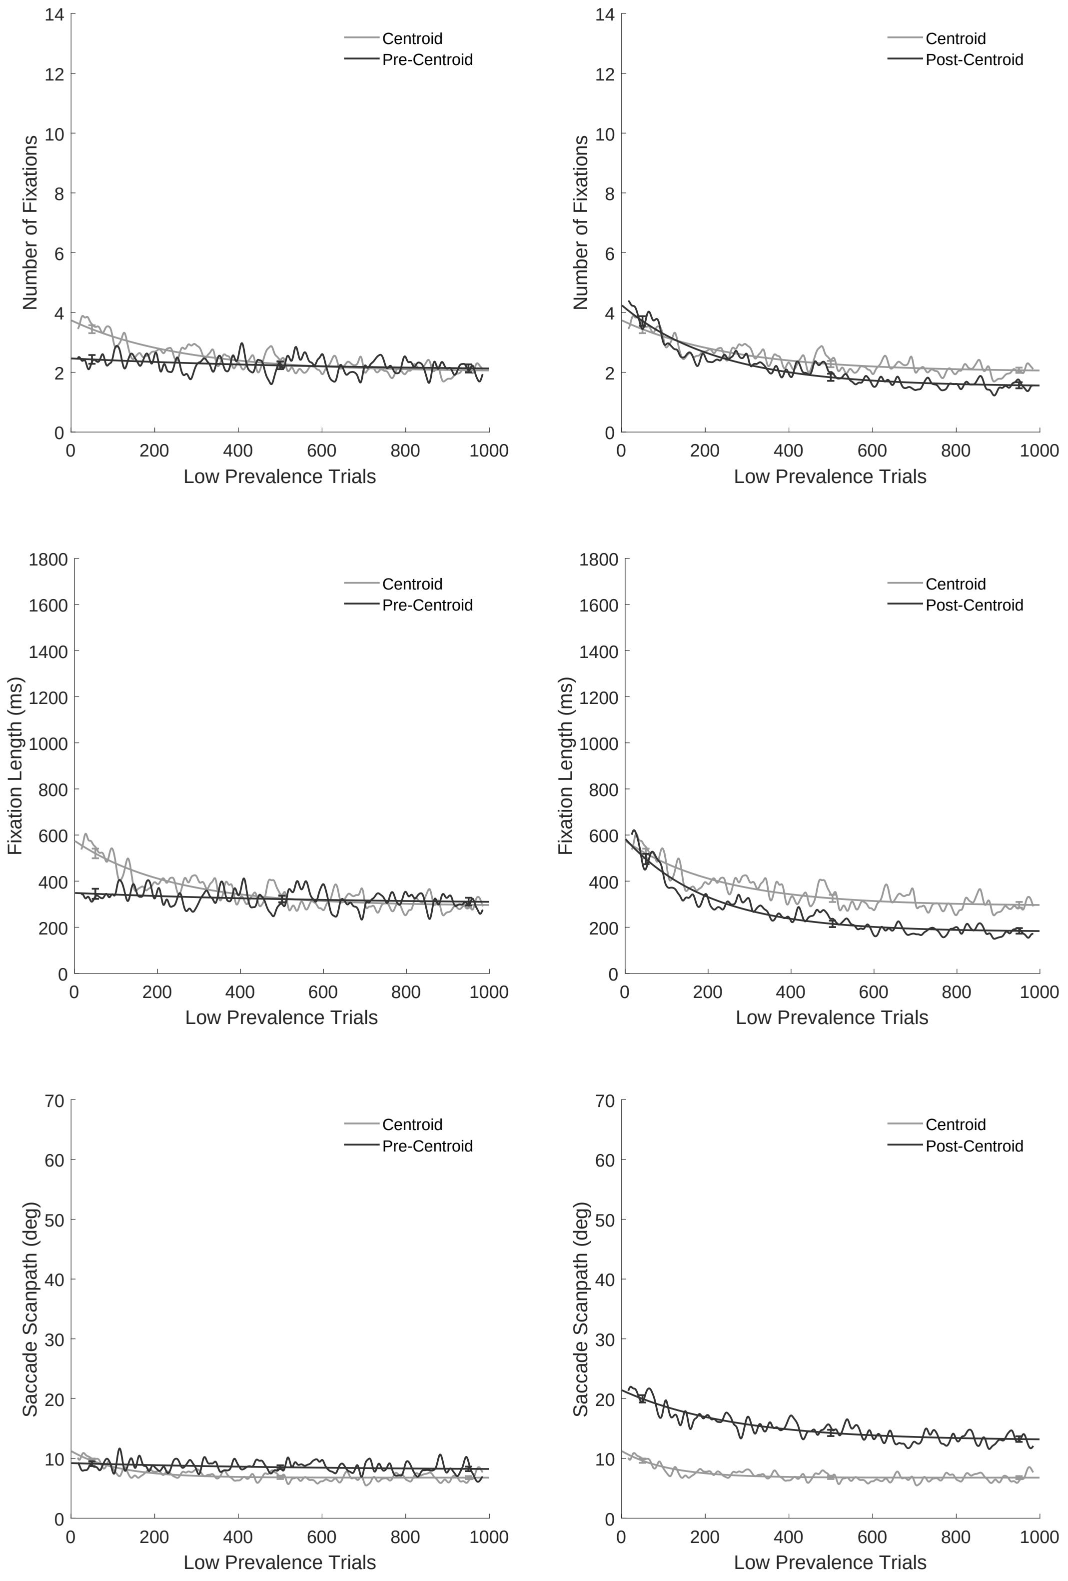

Figure 4.14: The low-prevalence eye-movement behaviour before entering the centroid zone, within the centroid zone, and after leaving the centroid zone with an exponential decay function fitted to each dataset. The error bars represent the standard error found in a sample consisting of fifteen trials either side of the presented error bar.

Table 4.4: Pre-centroid, centroid, and post-centroid eye-movement behaviour across the prevalence conditions. Means are followed by their parenthesised standard deviations.

	Experiment One					
	Medium Prevalence			Low Prevalence		
	Pre-	Centroid	Post-	Pre-	Centroid	Post-
Fixations						
Decay (*n*)	0.22(0.98)	1.01(2.01)	1.26(2.81)	0.37(0.81)	1.90(1.15)	3.01(1.89)
Rate (*trial*)	72.93(49.00)	67.32(47.55)	93.65(40.36)	442.92(209.43)	324.10(196.38)	308.55(172.23)
Asymptote (*n*)	2.72(1.17)	2.37(0.70)	2.37(0.81)	2.13(0.65)	1.97(0.75)	1.40(0.93)
Length						
Decay (*ms*)	91.47(92.50)	200.26(214.48)	236.92(211.36)	66.87(119.17)	317.04(209.86)	427.40(329.39)
Rate (*trial*)	82.05(40.52)	64.21(44.67)	75.49(46.52)	446.15(168.03)	255.71(168.45)	296.44(168.83)
Asymptote (*ms*)	368.15(154.87)	340.88(113.74)	304.32(91.76)	296.60(79.30)	286.93(108.34)	176.04(110.38)
Scanpath						
Decay (*deg*)	3.66(2.83)	0.43(2.07)	3.20(6.04)	1.16(1.03)	4.42(2.21)	9.33(6.66)
Rate (*trial*)	47.89(44.01)	85.11(39.34)	88.06(39.21)	512.57(179.20)	129.89(45.42)	422.74(213.30)
Asymptote (*deg*)	8.25(0.55)	7.76(1.19)	16.92(3.91)	8.06(0.51)	6.71(1.60)	12.29(4.64)

As can be seen in Table 4.4, the magnitude of the curve's decay in fixations was the least before entering the centroid zone and the greatest decay occurred after leaving the centroid zone in the low-prevalence condition ($F_{(2,17)} = 24.30, p < .001, \eta_P^2 = 0.57$), with the posthoc analysis showing that the main significant difference occurred between the pre-centroid and post-centroid means ($p < .001$), a larger p-value occurred between the pre-centroid and centroid means ($p = .002$), and the largest p-value occurred between the centroid and post-centroid means ($p = .025$). In the medium-prevalence condition, there was a similar trend in behaviour; however, no significant differences were found between the pre-, post-, and centroid means. No significant differences in the rate of this pre-, post-, and centroid decay in fixations before reaching asymptote were observed in both the medium- and low-prevalence conditions. Once the participants had reached an asymptote in the low-prevalence condition, no significant differences were found with only 0.72 of a fixation separating the pre-, post-, and centroid means. This was also the case for the medium-prevalence condition with only 0.35 of a fixation separating these data categories.

Between the prevalence conditions, no significant difference was observed for the pre-centroid and centroid fixational decay. However, after leaving the centroid zone, there was a greater amount of decay in the low-prevalence condition ($F_{(1,18)} = 4.96, p < .05, \eta_P^2 = 0.22$). We found that the participants took a longer amount of time to reach a steady state of fixational behaviour in the low-prevalence condition before entering the centroid zone ($F_{(1,18)} = 58.15, p < .001, \eta_P^2 = 0.76$), within the centroid zone ($F_{(1,18)} = 33.20, p < .001, \eta_P^2 = 0.65$), and after leaving this zone ($F_{(1,18)} = 27.30, p < .001, \eta_P^2 = 0.60$). Once the participants had reached asymptote, their level of fixational behaviour was greater in the medium-prevalence condition before entering the centroid zone ($F_{(1,18)} = 8.84, p < .05, \eta_P^2 = 0.33$), within the centroid zone ($F_{(1,18)} = 10.84, p < .05, \eta_P^2 = 0.38$), and after leaving this zone ($F_{(1,18)} = 22.89, p < .001, \eta_P^2 = 0.56$).

These results show that there was very little decay in behaviour pre-centroid and within the centroid in both the prevalence conditions and the greatest amount of decay occurred post-centroid in the low-prevalence condition and this was the same between the prevalence conditions. The participant's rates of decay were significantly longer in the low-prevalence condition and once the asymptote was reached this level of behaviour was greater in the medium-prevalence condition across the pre-, post-, and centroid data categories. These findings confirm our hypothesis that the data would be similar across the prevalence conditions before entering the centroid zone and within this zone; however, there would be a greater amount of decay in eye-movement behaviour after leaving the centroid zone in the low-prevalence condition compared to the medium-prevalence condition. The behavioural trends reported for the pre-, post-, and centroid fixations between the conditions were similar to those observed for fixation lengths and saccadic scanpath trajectories.

4.5 Discussion

4.5.1 Visual Search

When target prevalence in a visual search task is low, the chances of a participant failing to detect a target increases dramatically, with miss rates up to two to five times greater (Ishibashi et al., 2012; Rich et al., 2008; J. D. Schwark et al., 2013; Wolfe et al., 2005, 2007). This phenomenon is referred to as the "Low Prevalence Effect" and is a major concern for many critical real-world visual search tasks, such as airport security and medical screening, where targets can be extremely rare (Gur et al., 2003). Therefore, it is of theoretical and practical significance to have a greater understanding of the underlying mechanisms behind this phenomenon. We asked the question, "Is the schedule of signal presentation controlling eye-movement behaviour and if so, how?" We suggest that search behaviour may itself be influenced by the outcomes of search behaviour through the operant processes of reinforcement and extinction. Specifically, in operant conditioning, behaviours are differentially selected by their consequences (Skinner, 1984). Reinforcing consequences will strengthen a behaviour by maintaining or increasing the behaviours that they follow and if this contingent reinforcement is withheld, the dependent behaviour will be weakened and over time this behaviour will be extinguished (Pierce & Cheney, 2017). When signal detection functions as reinforcement for visual search behaviour, operant theory would predict more intense searching behaviours and improved detection performance when target signals are plentiful (i.e., search behaviour is strengthened). Conversely, the absence of target signals for long periods would weaken search behaviours and increase signal detection errors. We term this complementary pair of predictions the *Reinforcement Strengthening Hypothesis*.

For Experiment One, all participants worked with the same disparity between the target stimulus (i.e., 75 green component value) and non-target stimuli (i.e., 128 green component value) which may be considered a large discrepancy. We observed that the current experiment produced the Low Prevalence Effect with a significantly greater amount of miss errors occurring within the low-prevalence condition compared to the medium-prevalence condition, while no significant differences were observed between the prevalence conditions for the false alarm errors. These findings were in line with that of Wolfe et al. (2007) who in a series of experiments showed that the prevalence effect was an extremely robust phenomenon, and they attribute this change in behaviour to a shift in the decision criterion rather than a change in sensitivity (Wolfe & Van Wert, 2010; Horowitz, 2017). However, as the decision criterion (c) is simply an average of the z-score transformed hit rate and the z-score transformed false alarm rate, an increase in miss errors under low-prevalence conditions would shift the decision criterion to a more conservative position (Stanislaw & Todorov, 1999; Macmillan & Creelman, 2004). An operant interpretation of the Low Prevalence Effect would suggest that when signal detection functions as reinforcement for visual search behaviour, this theory would predict more intense searching behaviours and, consequently, improved target signal detection when target signals are abundant (i.e., search behaviours are strengthened). Conversely, the absence of target signals for long periods would weaken search behaviours and increase signal detection errors (i.e., an extinction effect).

We observed that the participants settled into their search task more rapidly (i.e., eye-movement behaviour metrics stabilized significantly faster) and they maintained a significantly higher rate of behaviour across all three measures of eye-movement behaviour (i.e., the number of fixations, the length of these fixations, and their scanpath trajectories) throughout the medium-prevalence condition compared to the low-prevalence condition (see Figures 4.5). This increase in behaviour under medium-prevalence conditions may be viewed as a manifestation of the strengthening of behaviour (Mccarthy & Davison, 1979; Thorndike, 1911) due to the abundance of target-present trials in this condition and the reinforcing effects of finding what you are looking for. As a result, there would be a greater number of eye movements under these circumstances, resulting in a reduced rate of miss errors.

Furthermore, when compared to the medium-prevalence condition, the participants took a significantly longer amount of time to reach a steady state of behaviour, and once this state had been reached, the level of this behaviour was significantly lower in the low-prevalence condition for all our measures of eye-movement behaviour. Interestingly, a significantly greater amount of behavioural decay was only observed for the saccadic trajectories in the low-prevalence condition and not for the number of fixations or their lengths, which may be attributed to the large amount of sample variance across these measures. Overall, these observations were also in line with the Reinforcement Strengthening Hypothesis, which would predict a weakening of search behaviour under these conditions, resulting in fewer eye movements and an increase in detection errors.

Nonhuman signal detection studies (i.e., pigeon studies) have shown that changes in performance and response bias (i.e., the decision criterion) are a function of variations in the relative reinforcement ratio for choice responses, and not a function of variations in the probability of stimulus presentation (M. C. Davison & Tustin, 1978; Mccarthy & Davison, 1979; J. A. Nevin, Jenkins, Whittaker, & Yarensky, 1982). Therefore, if we consider that signal detection functions as reinforcement for visual search behaviour, it may be suggested that the reinforcing effects of finding the target stimulus are primarily influencing search behaviour. Adding to our understanding of these intrinsic reinforcement effects is the knowledge that eye movements made during visual search tasks rely on neural computations similar to these employed when eye movements are planned and executed to obtain explicit reinforcement, which suggests that eye movements originate from the processing of predominantly visual sensory information, dopaminergic feedback regarding previous errors or successes, and reinforcement expectations (Trommershäuser et al., 2009).

Of interest was the observation that, when compared to the correct rejection mean, there was a significant increase in eye-movement behaviour for the first correct rejection immediately following a correct target-present response. This pattern of behaviour was observed for all eye-movement measures across both the prevalence conditions, except for the length of fixations in the low-prevalence condition (see Figures 4.4). This increase in eye-movement behaviour for the first post-hit correct rejection may be viewed as a manifestation of the local effects of obtained reinforcement (Mccarthy & Davison, 1979; Jia et al., 2021; Williams, 1976). Given the greater number of target-present trials in the medium-prevalence condition, it may be suggested that this increase in search behaviour immediately following a hit is a manifestation of the strengthening of search behaviours due to the reinforcing effect of finding what you are looking for, and this reinforcement strengthening mechanism would maintain eye-movement behaviour at a consistently higher level throughout this condition due to the greater consistency of target-present trials and the reinforcing effects of finding the target stimulus. Furthermore, it is not surprising that target-absent latencies become shorter under low-prevalence conditions (A. M. Treisman & Gelade, 1980; Wolfe, 1998; Wolfe et al., 2005; Wolfe & Van Wert, 2010) due to the lower number of target-present trials, search behaviour would result in finding the target less often and behaviour would be weakened. In other words, the schedule of reinforcement (i.e., signal presentation) appears to be controlling the intensity of eye-movement behaviour and this observable increase in search behaviour immediately following the identification of a target (i.e., a reinforcer) impacts detection performance and influences the decision criterion throughout each prevalence condition.

4.5.2 The Stimulus Density Map

We were also interested in the progression of eye-movement behaviour throughout a trial. To provide an understanding of this progression and the search strategies which may be involved, we constructed the centroid zone, which represented the portion of the display with the greatest density of stimulus items (see Figure 2.8). We were interested in the centroid zone because we had observed that there was increased eye-movement activity in areas of the display which did not contain the target stimulus across both the prevalence conditions (see Figures 2.3 & 2.4), and we hypothesized that visual attention was being guided or drawn towards this centroid zone.

In this respect, we observed that there was a much greater proportion of eye-movement behaviour within the centroid zone compared to the relative surface area of this zone within the stimulus matrix box in both the prevalence conditions. This centroid pattern of eye-movement behaviour would make for a statistically efficient search strategy, in that the centroid zone only makes up a small proportion of the search array and yet incorporates over one-third of the array's stimuli and accounts for approximately one-third of the fixations across a trial. The density of the stimuli within the centroid zone would therefore allow the participant to search this area with both overt and covert visual attention in an energy-efficient manner.

As a result, it was observed that there were very few miss errors occurring within the centroid zone for the medium-prevalence condition while in the low-prevalence condition, there were a significantly greater number of misses within this zone. These results tentatively suggest that the increase in miss errors which occur under low-prevalence conditions (i.e., Low Prevalence Effect), may primarily occur within the centroid zone. Of the small number of hits which occurred outside the centroid zone, approximately one-third occurred before entering this zone in both prevalence conditions. This suggests that if the target appeared close to the point of fixation upon the presentation of a target-present trial or the participant was able to identify the target within their visual periphery they would saccade directly towards the target and terminate the trial thereafter, before entering the centroid zone.

We then asked the question, "Is this pre-centroid eye-movement behaviour being guided towards the centre of the search array before entering the centroid zone?" We understand that when an image of a visual scene is presented on a computer screen, there is a strong tendency for the observer to look at the centre of the scene first (Bindemann, 2010; Buswell, 1935; Mannan, Ruddock, & Wooding, 1995; Parkhurst, Law, & Niebur, 2002; Parkhurst & Niebur, 2003) and this may account for a substantial proportion of behaviour during eye guidance (Tatler et al., 2005; Vincent, Baddeley, Correani, Troscianko, & Leonards, 2009). The basis for this central fixation bias remains unresolved; however, studies have ruled out some potential explanations for this effect. For example, this bias may arise from the fact that studies have confounded the centre of a scene with the location of a fixation cross which is presented before the scene onset, and could inherently bias observers to initially view the scene centre (Mannan et al., 1995; Parkhurst et al., 2002; Parkhurst & Niebur, 2003; Tatler et al., 2005; Vincent et al., 2009). However, the central viewing tendency persists when non-central fixation markers are used, so that observers are not biased in this manner (Bindemann, Scheepers, Ferguson, & Burton, 2010). This tendency was even observed when highly eccentric fixation markers were presented in the periphery of a visual display (Tatler, 2007). This indicates that this effect does not arise from a motor bias to perform small length saccades, which are then scattered near the scene centre when a central or near-central fixation point is employed (Tatler, 2007). Another possible explanation for the central viewing bias is that it reflects a tendency to look straight-ahead, to a location which typically coincides with the scene centre. However, observers are still drawn to the display centre when the screen is shifted left or right of a central viewing position (Vitu, Kapoula, Lancelin, & Lavigne, 2004). This indicates that this effect does not reflect a viewing preference for a straight-ahead position but reflects a systematic bias towards the centre of a visual display. The central fixation bias also appears to be unaffected by the distribution of visual features in a scene. In photographic images, objects of interest often provide a focal point in a central location that could therefore give rise to a central viewing effect (Reinagel & Zador, 1999; Tatler et al., 2005). However, this cannot explain the central fixation bias either, which persists when the distribution of visual features is systematically biased to off-centre locations in a scene (Tatler, 2007).

Furthermore, the human visual system is foveated, that is, outside the central visual field resolution and acuity decline rapidly. Nonetheless, much of a visual scene is perceived after only a few saccadic eye movements, suggesting that saccade targets are not chosen randomly, but according to an effective strategy. It has been known for some time that the local image structure may influence this selection process. However, the question of what the most relevant visual features are is still debated (Kienzle, Franz, Scholkopf, & Wichmann, 2009).

Therefore, we placed a circle in the centre of the stimulus matrix box with a radius subtending 1.7° , which was twice the diameter of the catchment circle placed around each stimulus item (see Figure 4.11). While the centroid zone is often centred around the centre of the display (e.g., 28.15% overlap between the centroid and centre circle) the data did not support a central fixation bias, with very few fixations occurring within the centre circle (i.e., 0.55 fixation mean within the medium-prevalence condition and 0.60 in the low-prevalence condition). Alternatively, it may be suggested that these pre-centroid eye movements were being guided towards the centroid zone as a search strategy. Particularly, if they were unable to identify the target upon stimulus presentation (i.e., a pop-out effect).

Interestingly, it was observed that there was a significantly greater number of fixations, and the length of these fixations was longer within the centroid zone in the low-prevalence condition compared to the medium-prevalence condition and suggests that when reinforcement is lean (i.e., low-prevalence condition) there is an increase in search behaviour within this zone (see Figures 4.9).

To enrich our understanding of the centroid zone and the possible role that the centroid zone could have in search behaviour, we broke the eye-movement curve parameters down into the pre-, post-, and centroid data sub-categories for each of the prevalence conditions, and we observed that in the low-prevalence condition, the least amount of decay occurred before entering the centroid zone (i.e., scene preview) and a significantly greater amount of decay occurred after leaving this zone while in the medium-prevalence condition, no significant differences in decay were observed across these data categories. Once a steady state of eye-movement behaviour had been reached, which took a significantly greater amount of time in the low-prevalence condition, there was a significantly greater number of fixations in the medium-prevalence condition the pre-, post-, and centroid data categories. These findings show that eye-movement behaviour was similar across the prevalence conditions before entering the centroid zone and within this zone; however, there was a greater deterioration in eye-movement behaviour after leaving the centroid zone in the low-prevalence condition due to the weakening of behaviour when reinforcement was scarce.

Interestingly, it was observed that for the absolute means, there was a significantly greater number of fixations, and the length of these fixations was longer within the centroid zone for the low-prevalence condition, suggesting that when reinforcement is lean (i.e., low-prevalence condition) there is an increase in behaviour within this zone. However, when we examined the centroid curve asymptote parameters, we observed that there was a significantly greater number of centroid fixations, and the length of these fixations was significantly longer in the medium-prevalence condition once a steady state of behaviour had been reached. This suggests that there was a spike in behaviour in the low-prevalence condition during the transition period before the asymptote had been reached.

Overall, these observations suggest that eye-movement behaviour may be considered stable across the conditions during their scene preview before entering the centroid zone and within the centroid zone, while the weakening of search behaviour predominantly occurs after leaving this zone in the low-prevalence condition due to the scarcity of obtained reinforcement. Furthermore, the strengthening of behaviour across the pre-, post-, and centroid data categories in the medium-prevalence condition when target signals are abundant is observed in the post-hit increase in correct rejection eye-movement behaviour due to the local effects of reinforcement when we find the target stimulus.

Experiment Two

The dominant theories of visual search (J. Duncan & Humphreys, 1992; A. M. Treisman & Gelade, 1980; Wolfe et al., 1989) propose that search is a process involving comparisons of items against a target's visual properties. Therefore, we were interested in establishing whether the search strategies (i.e., centroid eye-movement behaviour) and the mechanisms (i.e., Reinforcement Strengthening Hypothesis) which were observed in Experiment One would also manifest themselves in Experiment Two when task difficulty is increased. We hypothesise that search behaviour and performance errors would also increase when the task's difficulty is increased.

Experiment Two was a replica of Experiment One except that the participants' discrimination threshold estimates were individually titrated using the two-down/one-up single staircase method (Levitt, 1971) to control for participant variance in visual discriminability across the sample and converge the task demand to a difficulty probability of approximately 0.71 for each participant (Lu & Dosher, 2013). Figure 5.1 shows a sample of the titration procedure using the two-down/one-up staircase method.

To remind the reader, this procedure was designed to find the target detection threshold for each participant using the two-down/one-up staircase method (Levitt, 1971). During this procedure, the blue colour component of the target stimulus was adjusted by adding or subtracting the green colour value within the colour conjunction. The blue colour of the target stimulus would lighten by one step-size following two consecutive correct responses, bringing the target colour closer to the distractors' colour conjunction; and darken in colour by one step-size following each incorrect response, distancing the target colour from the distractors. Step-sizes halved on the first, third, seventh, and thirteenth reversals.

At the end of this procedure, the mean colour value for the preceding eight

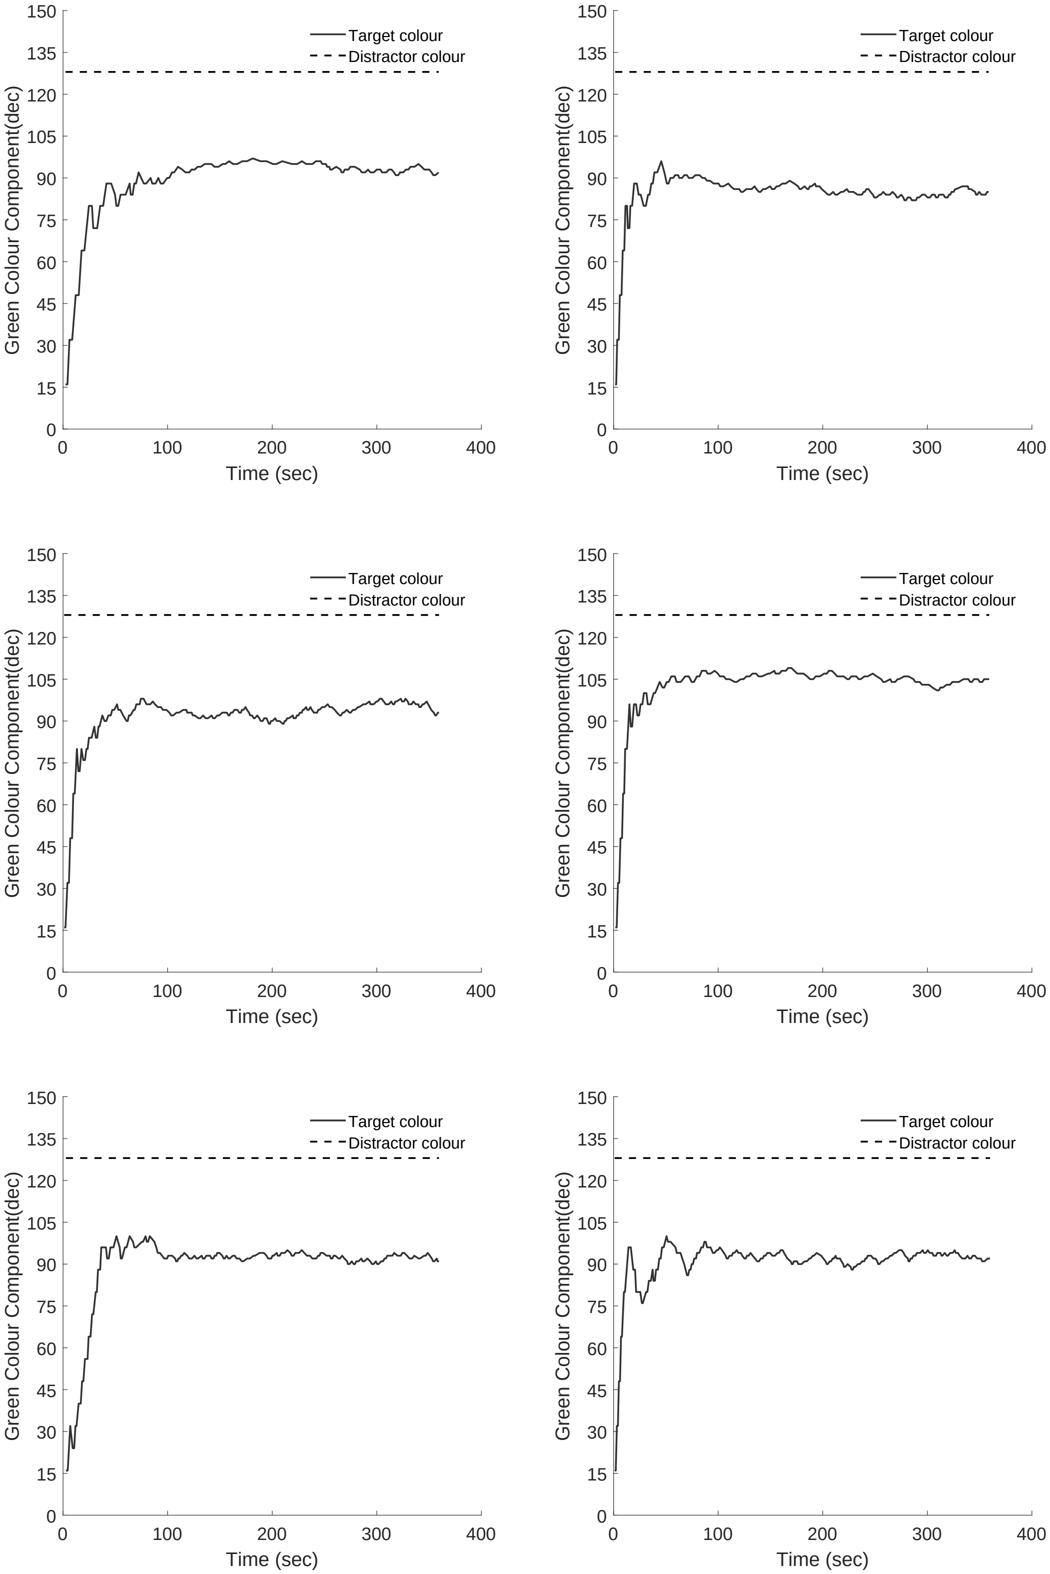

Figure 5.1: Shows a sample of the participants discrimination threshold titration results. The x-axis represents the duration of the titration procedure, and the y-axis represents the green colour component amount (i.e., 0 – 255 green colour units).

reversals became the target stimulus colour for each participant. This mean colour value converged on a difficulty probability of 0.71 (Lu & Dosher, 2013). This titration procedure resulted in a mean green target colour value of 93.2 ($SD =$ 9.13) across participants while the green distractor colour value was set at 128 throughout all the experiments. When we compared these results to Experiment One, the green target colour value assigned to each participant in Experiment One was 75. In the current experiment, this means that there was a 7.14% increase in the green colour component which brings the target stimulus colour conjunction closer to that of the distractor's colour conjunction, resulting in increased difficulty in locating the target stimulus.

5.1 Participants

Twenty participants (8 male & 12 female) were recruited from the University of Waikato's student body. They ranged in age between 17.8 – 45.5 years ($M =$ 21.7 years, $SD = 6.8$ years). One participant's data was excluded from the eye-movement analysis due to head movement during the experiment which was greater than two standard deviations above the mean (exclusion criteria).

5.1.1 Signal Detection Theory Measures

As with Experiment One, the z-transformed hit rate was plotted against the z-transformed false alarm rate for each participant and a linear regression line was fitted for both the prevalence conditions (see Figures 5.2 & 5.3). Unfortunately, the underlying assumptions of normality and equal variance for the signal detection indices of d' and c were found to be violated with the low-prevalence regression line having a slope of 0.54 and an intercept of -0.34, while the medium-prevalence regression line had a slope of 0.35 and an intercept of -0.35. As can be seen in Figures 5.2 and 5.3, there was a large amount of participant variance in the sample resulting in a poor fit for the regression lines in both the low-prevalence condition ($r^2 = 0.22$) and the medium-prevalence condition ($r^2 = 0.10$). As a result, the parametric signal detection indices of d' and c were unable to be used.

Therefore, we then applied Davison and Tustin's (1978) non-parametric sum-

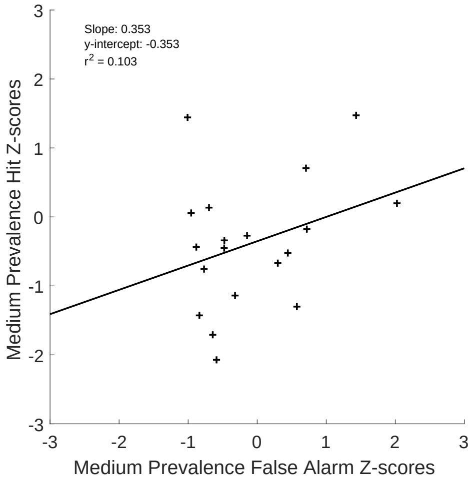

Figure 5.2: The medium-prevalence hit and false alarm z-scores with a regression line fitted to the data.

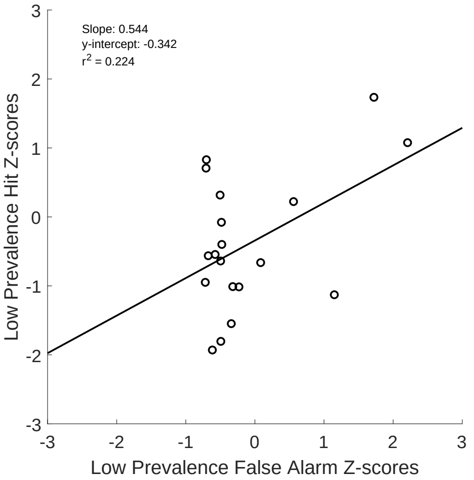

Figure 5.3: The low-prevalence hit and false alarm z-scores with a regression line fitted to the data.

mary measures of log d (i.e., stimulus discriminability or sensitivity) and log B (i.e., response bias or decision criterion) which were used in Experiment One. These measures showed that response bias, as measured by log B (see equation 4.6), shifted from a conservative bias in the low-prevalence condition ($M = -2.22, SD = 0.92$) to a more liberal bias ($M = -1.26, SD = 0.76$) in the medium-prevalence condition ($F_{(1,15)} = 11.98, p < .05, \eta_P^2 = 0.44$). No significant difference in stimulus discriminability, as measured by log d (see equation 4.5), was found between the prevalence conditions.

Between Experiments One and Two, there was a shift in discriminability (Log d), for the medium-prevalence condition indicating that task difficulty was increased in Experiment Two ($F_{(2,28)} = 11.30, p < .05, \eta_P^2 = 0.29$) and this trend was repeated in the low-prevalence condition; however, no significant difference was observed. No significant differences in response bias (log B) were observed for either of the prevalence conditions across Experiments One and Two.

As can be seen in Table 5.1, Experiment Two produced the familiar Low Prevalence Effect with a greater amount of miss errors in the low-prevalence condition compared to the medium-prevalence condition ($F_{(1,19)} = 35.65, p < .001, \eta_P^2 = 0.65$). No significant difference in false alarm errors was observed between the conditions (see Figure 5.4).

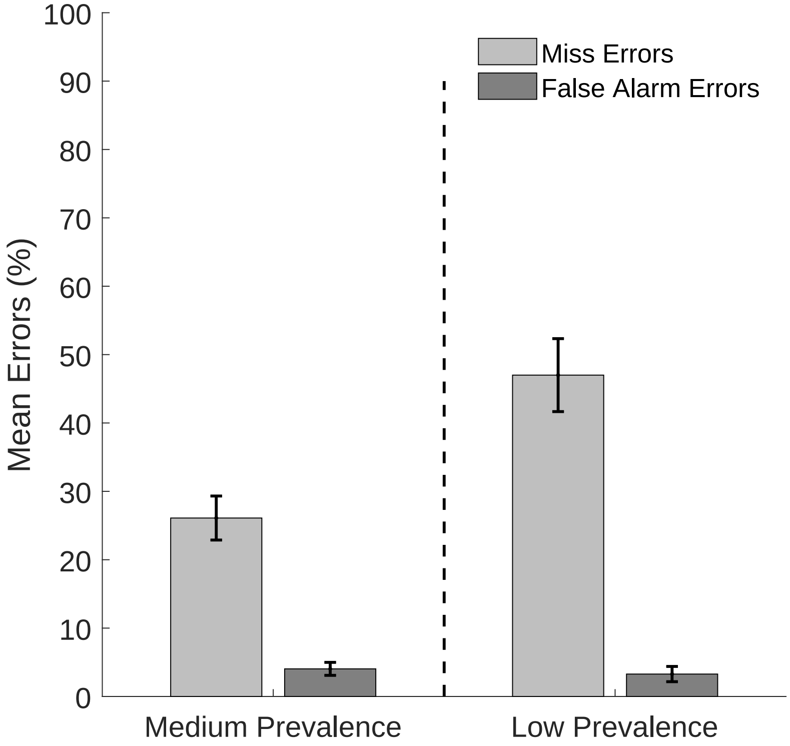

Figure 5.4: Signal detection miss and false alarm errors in the medium-prevalence condition and the low-prevalence condition.

Compared to Experiment One, miss errors increased in Experiment Two by 14.7% in the low-prevalence condition ($F_{(1,39)} = 4.29, p < .05, \eta_P^2 = 0.10$), and by 12.9% in the medium-prevalence condition ($F_{(1,39)} = 11.92, p < .05, \eta_P^2 = 0.24$). No significant differences in false alarm errors for each condition were observed between the experiments. These results support our hypothesis that there would be a greater amount of performance errors when task difficulty is increased.

As with Experiment One, it was observed that when compared to the correct rejection mean, there was an increase in eye-movement behaviour for the first correct rejection immediately following a correct target-present response. This pattern of behaviour was observed in the medium-prevalence eye-movement measures for the number of fixations ($F_{(1,18)} = 6.96, p < .05, \eta_P^2 = 0.28$), and the saccade scanpath trajectories ($F_{(1,18)} = 8.60, p < .05, \eta_P^2 = 0.32$). This significant pattern of behaviour was not observed for fixation lengths, subsequent correct re-

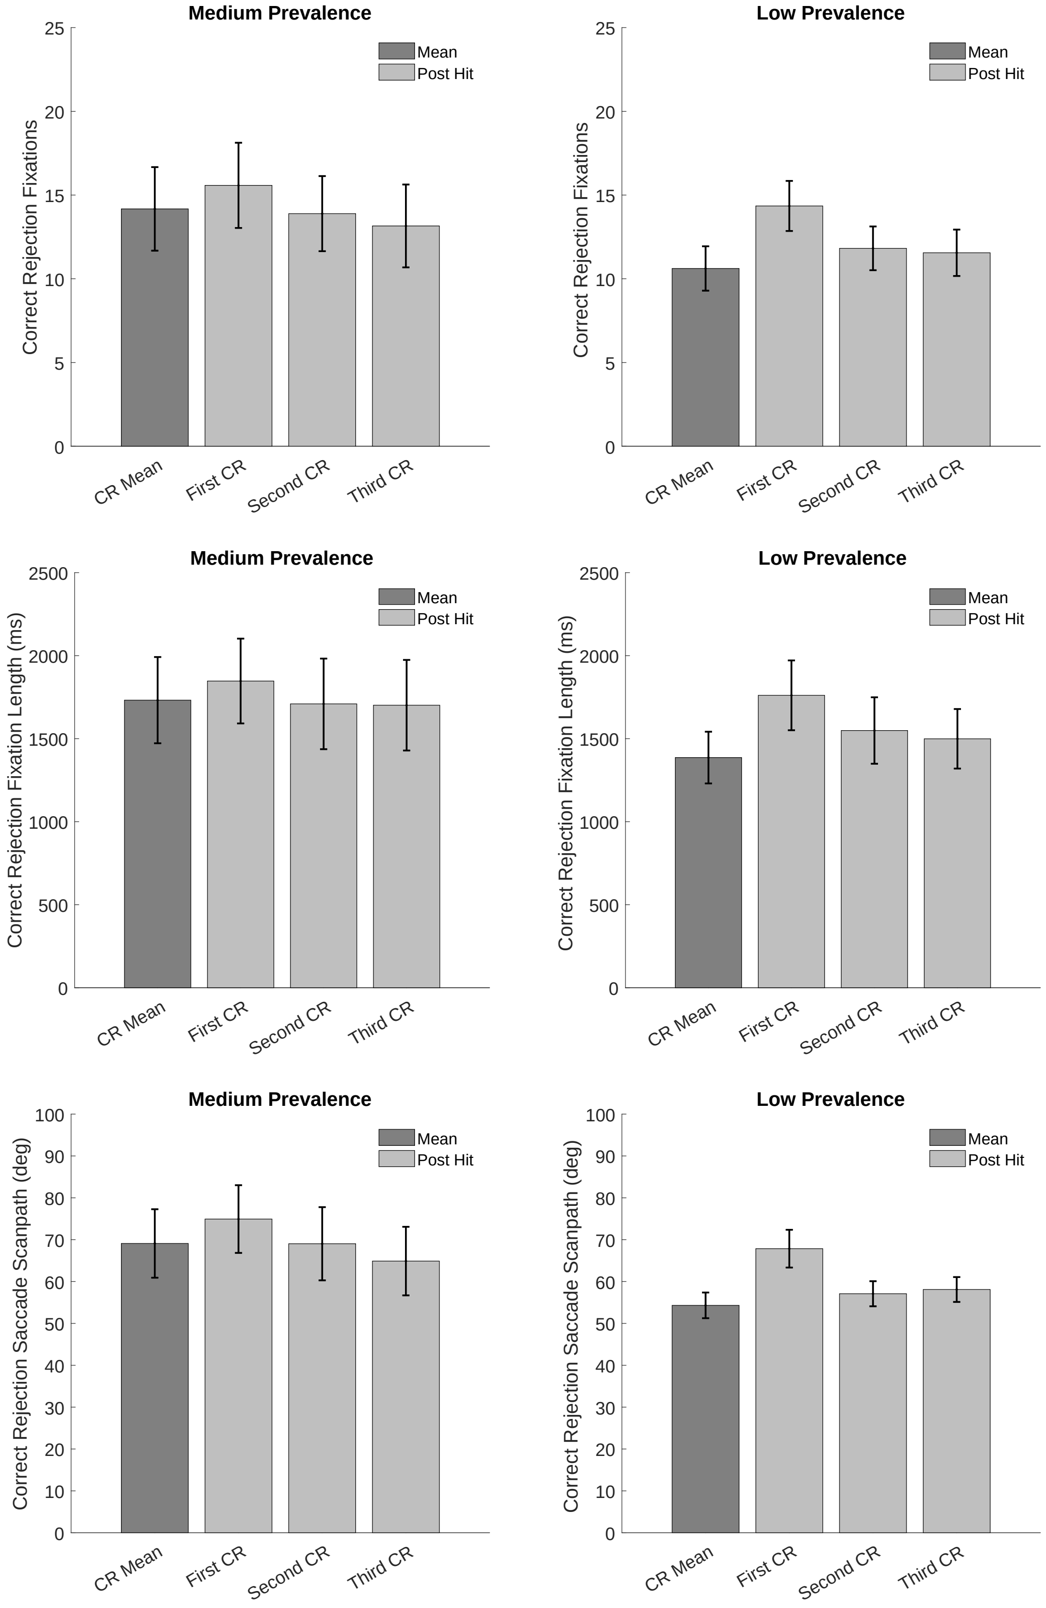

Figure 5.5: The eye-movement behaviour for the correct rejection mean, the first correct rejection immediately following a hit, the second consecutive correct rejection following a hit, and the third consecutive correct rejection following a hit across both prevalence conditions.

Table 5.1: Signal detection miss and false alarm errors for the medium-prevalence and the low-prevalence conditions in Experiments One and Two.

	Medium Prevalence		Low Prevalence	
	Miss (%)	False Alarm (%)	Miss (%)	False Alarm (%)
Experiment One	13.20(8.49)	3.23(5.92)	32.25(21.12)	2.18(4.32)
Experiment Two	26.10(14.40)	4.04(4.25)	47.00(23.86)	3.28(4.99)

jections following the first post-hit correct rejection, and for misses, false alarms, and subsequent hits immediately following a correct target-present response in the medium-prevalence condition (see Figures 5.5).

Table 5.2: Eye-movement behaviour for the correct rejection mean and the first correct rejection immediately following a hit across prevalence conditions for Experiments One and Two. Means are followed by their parenthesised standard deviations. CR is an abbreviation for correct rejection.

	Medium Prevalence		Low Prevalence	
	CR Mean	Post-hit CR	CR Mean	Post-hit CR
Experiment One				
Fixations (*n*)	8.45(3.79)	9.05(3.67)	5.89(2.20)	7.31(2.28)
Length (*ms*)	982.69(442.47)	1043.50(444.77)	737.52(313.76)	741.50(259.36)
Scanpath (*deg*)	50.34(12.40)	53.02(12.00)	40.83(8.61)	46.26(8.67)
Experiment Two				
Fixations (*n*)	14.18(10.88)	15.58(11.09)	10.62(5.78)	14.35(6.52)
Length (*ms*)	1731.96(1132.01)	1847.16(1113.32)	1386.08(679.08)	1760.98(916.32)
Scanpath (*deg*)	69.10(35.66)	74.93(35.25)	54.31(13.36)	67.85(19.67)

In the low-prevalence condition, compared to the correct rejection mean there was increased eye-movement behaviour for the first correct rejection immediately following a hit for all of the eye-movement behavioural measures, starting with the number of fixations ($F_{(1,18)} = 34.58, p < .001, \eta_P^2 = 0.66$), the length of these fixations ($F_{(1,18)} = 8.52, p < .05, \eta_P^2 = 0.32$), and the scanpath trajectories ($F_{(1,18)} = 22.94, p < .001, \eta_P^2 = 0.56$). Again, this pattern of behaviour was not observed for subsequent correct rejections following the first post-hit correct rejection, nor was it found for misses, false alarms, and subsequent hits immediately following a correct target-present response in the low-prevalence condition.

When we compared these results to those found in Experiment One, we found a similar trend in behaviour. In Experiment One, significant increases in behaviour were found for all three eye-movement measures in the medium-prevalence condition. However, in Experiment Two this trend was only observed for fixations and scanpath trajectories. In the low-prevalence condition for Experiment One, there was a post-hit increase in fixations and scanpath trajectories whereas this pattern of behaviour was found for all three eye-movement measures in the low-prevalence condition for Experiment Two.

5.2 Eye-movement Decay

We had observed that over time (trials) some of the measures dropped off in their values, suggesting a particular change in eye-movement behaviour over time. To capture this trend, we fitted the data with an exponential decay function for both the medium- and low-prevalence conditions and the three curve parameters were analysed. These decay curve parameters represent the magnitude of the decay curve fall (represented on the Y-axis), the rate of this fall (represented on the X-axis), and the asymptote (represented on the Y-axis). Figure 5.6 shows the eye-movement behaviour found in the low- and medium-prevalence conditions with an exponential decay function fitted to each dataset.

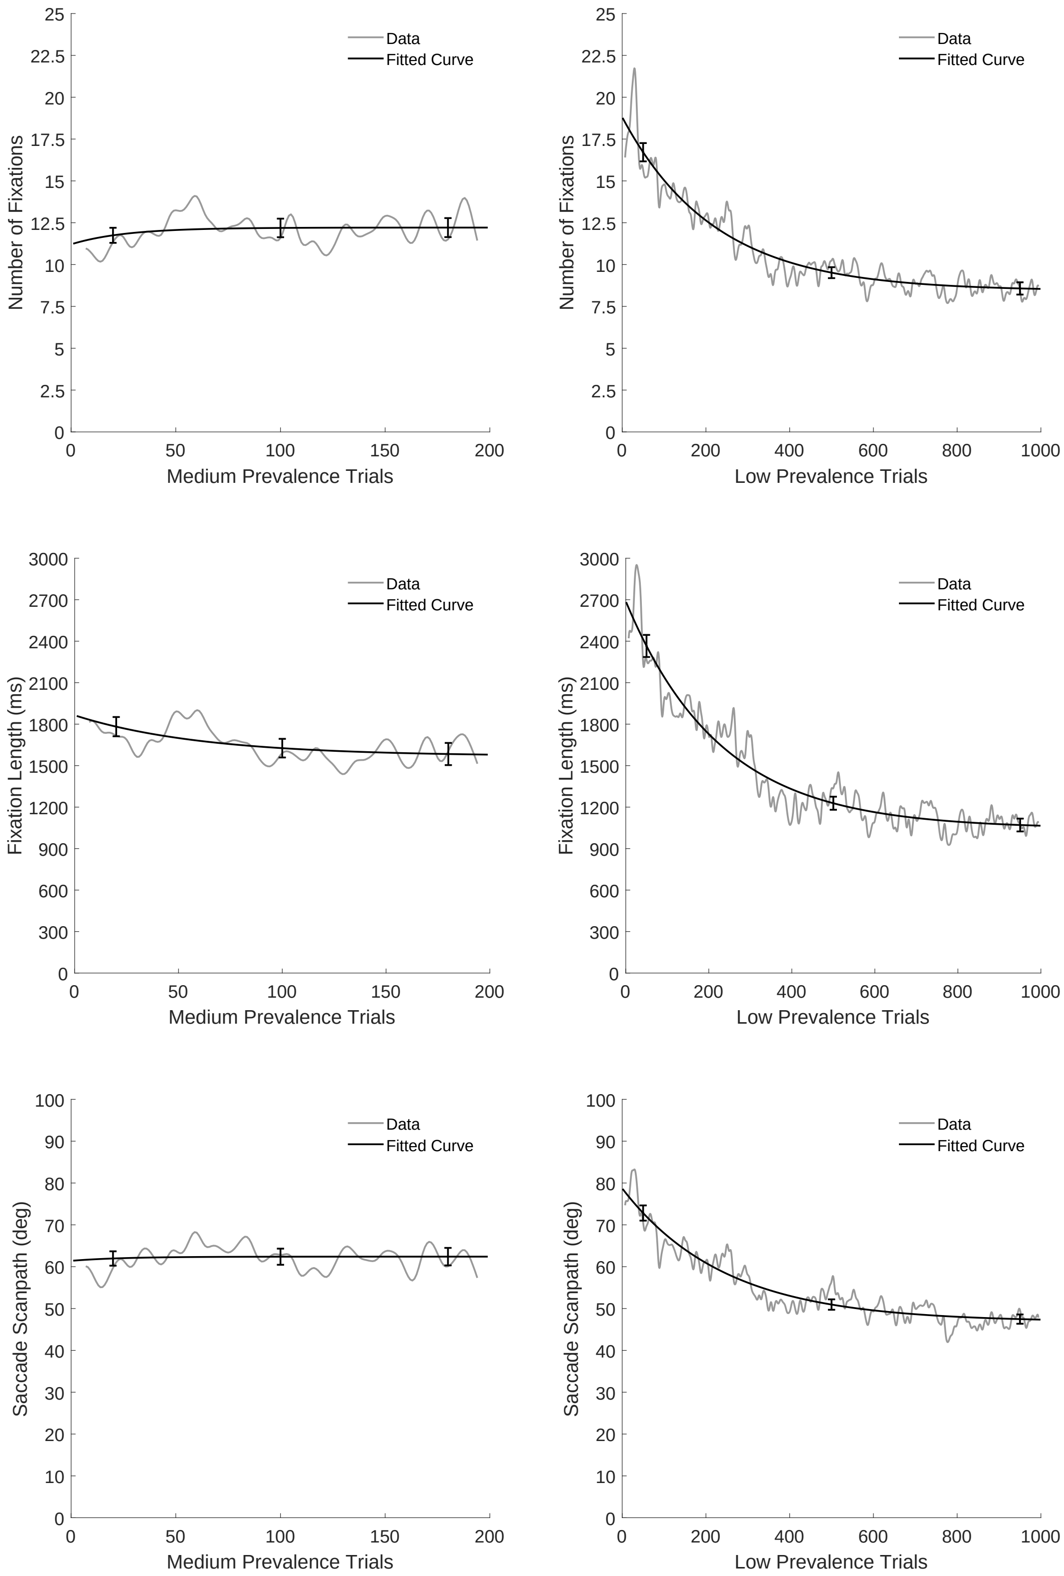

Figure 5.6: The mean eye-movement behaviour across the participants for each prevalence condition with an exponential decay curve fitted to each dataset. The error bars represent the standard error found in a sample consisting of fifteen trials either side of the presented error bar. Note the differing number of trials within each of the two prevalence conditions, with 200 trials in the medium-prevalence condition and 1000 trials in the low-prevalence condition.

Table 5.3: The eye-movement behaviour decay curve parameters for both the medium-prevalence and the low-prevalence conditions for Experiments One and Two. Mean percentages are followed by their parenthesised standard deviations and we remind the reader that the rate of decay before reaching asymptote is represented in trials.

	Medium Prevalence			Low Prevalence		
	Decay	Rate	Asymptote	Decay	Rate	Asymptote
Experiment One						
Fixations (*n*)	2.63(5.02)	82.56(43.74)	6.44(1.78)	5.11(3.29)	324.03(165.92)	4.67(1.99)
Length (*ms*)	461.61(427.83)	63.44(39.00)	847.36(237.03)	795.07(499.40)	307.84(174.43)	567.80(280.12)
Scanpath (*deg.*)	6.16(8.94)	88.97(39.84)	42.74(5.82)	18.39(11.43)	347.08(193.95)	35.82(7.88)
Experiment Two						
Fixations (*n*)	1.05(3.16)	76.15(39.28)	11.76(6.91)	11.89(8.74)	344.39(175.67)	7.47(5.28)
Length (*ms*)	585.97(667.35)	73.16(39.19)	1483(827.04)	1787.77(1174.02)	346.45(196.82)	931.69(569.48)
Scanpath (*deg.*)	7.53(12.86)	77.64(40.11)	59.96(28.46)	36.94(28.62)	340.35(208.87)	45.21(12.57)

The fitting of the exponential decay function to the fixations data produced a RMS value of 5.15 fixations ($SD = 2.76\ fixations$) for the low-prevalence curve fit and a RMS value of 6.19 fixations ($SD = 5.21\ fixations$) for the medium-prevalence curve fit.

As can be seen in Table 5.3, the magnitude of the curve decay for fixations was greater in the low-prevalence condition ($F_{(1,18)} = 26.30, p < .001, \eta_P^2 = 0.59$). This pattern of behaviour was consistent with that found in Experiment One, except that no significant difference was observed between the prevalence conditions. When we compared these results to those found in Experiment One, we found that there was a significantly greater decay in fixations within the current experiment for the low-prevalence condition ($F_{(1,18)} = 10.01, p < .05, \eta_P^2 = 0.22$), but not for the medium-prevalence condition.

A difference was found in the rate of this behavioural decay ($F_{(1,18)} = 40.58, p < .001, \eta_P^2 = 0.69$), with participants taking a longer amount of time to reach a steady state of behaviour in the low-prevalence condition (see Figure 5.6). These results were consistent with those found in Experiment One. A difference was also found between the prevalence conditions for the curve asymptotes ($F_{(1,18)} = 19.19, p < .001, \eta_P^2 = 0.52$). Specifically, once a steady state of fixational behaviour had been reached, this level of behaviour was lower in the low-prevalence condition. When we compared these results to those found in Experiment One, we found that the participants maintained a higher steady state of behaviour in the current experiment for both the medium-prevalence condition ($F_{(1,37)} = 10.57, p < .05, \eta_P^2 = 0.23$) and the low-prevalence condition ($F_{(1,37)} = 4.65, p < .05, \eta_P^2 = 0.11$).

From a behavioural perspective, measuring behavioural change across the prevalence conditions when their time constants are different (i.e., the number of trials) may be considered poor practice. Therefore, we analysed this data over the first 50 trials in each prevalence condition and found that the mean number of fixations increased by 0.86 in the medium-prevalence condition while in the low-prevalence condition, fixations decreased by 2.10. When we pushed this timestamp out to include the full 200 trials in the medium-prevalence condition and the first 200 trials in the low-prevalence condition, we found that the mean number of fixations increased by 1.00 in the medium-prevalence condition and decreased by 6.15 fixations in the low-prevalence condition, and we know that search behaviour continued to decrease in the low-prevalence condition until participants reach a steady state of behaviour on trial number 344. Note, this trial number was generated from the curve parameter equation. This increase in search behaviour found in the medium-prevalence condition was not found in Experiment One.

The fitting of the exponential decay curve to the fixation length data produced a RMS value of 776.13ms ($SD = 363.96ms$) for the low-prevalence curve fit and a RMS value of 878.82ms ($SD = 582.14ms$) for the medium-prevalence curve fit.

When we examined the curve parameters for the length of these fixations, we observed that the magnitude of the decay in fixation length was greater in the low-prevalence condition ($F_{(1,18)} = 19.06, p < .001, \eta^2_P = 0.51$). This trend in fixation length behaviour was like that found in Experiment One, except that no significant difference was found which may have been due to the high levels of sample variance. Between Experiments One and Two, there was a greater level of decay found in the low-prevalence condition for the current experiment ($F_{(1,37)} = 11.50, p < .05, \eta^2_P = 0.24$). Participants taking a longer amount of time (rate of decay) to reach a steady state of behaviour in the low-prevalence condition ($F_{(1,18)} = 29.06, p < .001, \eta^2_P = 0.62$), with was consistent with Experiment One. Once a steady state of behaviour had been reached, fixation times were shorter in the low-prevalence condition ($F_{(1,18)} = 11.82, p < .05, \eta^2_P = 0.40$) and this was also consistent with Experiment One. However, when we compared the current experiment with Experiment One, there were significantly greater fixation times found in the current experiment for both the medium-prevalence condition ($F_{(1,37)} = 10.39, p < .05, \eta^2_P = 0.22$) and the low-prevalence condition ($F_{(1,37)} = 6.25, p < .05, \eta^2_P = 0.15$).

To further our understanding of the differences between the conditions and the mechanisms which may be involved, we then examined this data over the first 50 trials within each prevalence condition and found that the mean length of trial fixations in the medium-prevalence condition decreased by 164.69ms while in the low-prevalence condition, these fixation lengths decreased by 324.19ms. Furthermore, when we push this timestamp out to include the full 200 trials in the medium-prevalence condition and the first 200 trials in the low-prevalence condition, we found that the fixation lengths in the medium-prevalence condition decreased by 284.29ms and in the low-prevalence condition these fixation times decreased by 958.18ms, and we know that these fixation times continued to decrease in the low-prevalence condition until participants reach a steady state of behaviour on trial number 346. Note, this trial number was generated from the curve parameter equation.

To provide greater insight into the differences between the prevalence conditions and the behaviours employed we next examined the saccadic scanpath trajectory data. The fitting of the exponential decay curve to the saccadic scanpath trajectory data produced a RMS value of 17.90° ($SD = 7.04°$) for the curve fitted to the low-prevalence condition data and a RMS value of 21.95° ($SD = 13.58°$) for the medium-prevalence curve fit.

When we investigated the fitted curve parameters for this data, we found that the magnitude of trajectory decay was greater in the low-prevalence condition ($F_{(1,18)} = 25.77, p < .001, \eta_P^2 = 0.59$). This trend in behaviour was consistent with Experiment One, except that no significant difference was found. Furthermore, there was a significantly greater amount of decay in the current experiment for both the medium-prevalence condition ($F_{(1,37)} = 4.79, p < .05, \eta_P^2 = 0.12$) and the low-prevalence condition ($F_{(1,37)} = 6.88, p < .05, \eta_P^2 = 0.16$) compared to Experiment One. The participants reached a steady state of trajectory behaviour more rapidly in the medium-prevalence condition ($F_{(1,18)} = 25.02, p < .001, \eta_P^2 = 0.58$), which was consistent with Experiment One and no significant differences were observed in the rate of decay for both the prevalence conditions across experiments. Once a steady state of behaviour had been reached, this level of trajectory behaviour was greater in the medium-prevalence condition ($F_{(1,18)} = 8.70, p < .05, \eta_P^2 = 0.33$) which was consistent with Experiment One. However, when we compared the current results to those found in Experiment One, significantly longer trajectories were found for both the medium-prevalence condition ($F_{(1,37)} = 6.68, p < .05, \eta_P^2 = 0.16$) and the low-prevalence condition ($F_{(1,37)} = 7.61, p < .05, \eta_P^2 = 0.18$) in the current experiment.

When we examined this data over the first 50 trials within each experimental condition, surprisingly, scanpath trajectories increased in the medium-prevalence condition by 0.87° while in the low-prevalence condition, trajectories decreased by 5.90°. When we pushed this timestamp out to 200 trials, we found that the scanpath trajectories in the medium-prevalence condition increased by 1.00° and in the low-prevalence condition they decreased by 17.76°.

Overall, participants reached a steady state of search behaviour more rapidly and maintained a higher level of this behaviour in the medium-prevalence condition due to the higher rate of reinforcement in this condition (Reinforcement Strengthening Hypothesis). Across both prevalence conditions, participants maintained a significantly higher degree of steady state behaviour in the current experiment which supports our task difficulty hypothesis, despite there also being a greater magnitude of behavioural decay in the low-prevalence condition, this greater amount of decay is consistent with operant theory and a reduction in obtained reinforcement (finding fewer targets). These results and those observed in Experiment One strongly suggest that eye-movement behaviour is an operant controlled by environmental contingencies and the outcome of these contingencies is resulting in the Low Prevalence Effect.

5.3 Stimulus Density Map

We were interested in the stimulus density map because we had observed in Experiment One that eye-movement behaviour was being guided towards the centroid zone during their scene preview and we hypothesised that the density of the stimuli within this zone would make this behaviour a statistically efficient search strategy by allowing the participant to search this area with both overt and covert visual attention in an energy-efficient manner. Therefore, we hypothesise that this search strategy would be repeated in Experiment Two when task difficulty was increased.

The mean surface area of the centroid zone within the stimulus density map constituted 10.10% (SD = 0.08%) of the stimulus matrix box area and represents the proportion of the stimulus matrix with the greatest stimulus density ($\geq$ 0.8 *density probability*). If the centroid zone sliced through a stimulus, this stimulus was counted as being within the centroid zone. As a result, the mean number of stimuli within the centroid zone was 6.28 (SD = 0.23) of the 15 stimuli presented within the stimulus matrix box on each trial. This equated to 0.68 ($SD = 0.36$) fixations per centroid stimulus. These fixations per centroid stimulus are significantly greater than those found in Experiment One ($F_{(1,18)} = 14.22, p < .001, \eta_P^2 = 0.28$), which indicates that there was a greater amount of search behaviour within the centroid zone in the current experiment when the discriminability between the target stimulus and the distractor items was reduced.

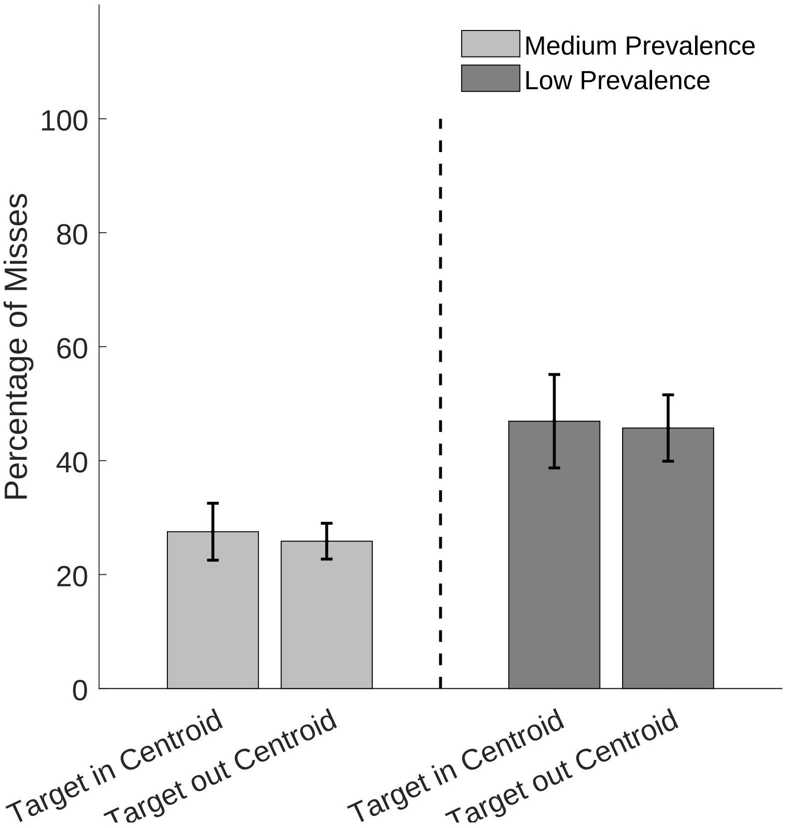

Figure 5.7: This figure shows the miss errors which occurred when the target stimulus was presented inside the centroid zone and those which occurred when the target stimulus was presented outside the centroid zone across the prevalence conditions.

On target-present trials in the medium-prevalence condition, the probability of the target stimulus appearing within the centroid zone was 0.19 ($SD = 0.07$) and it was observed that 27.53% ($SD = 21.79\%$) of miss errors occurred when the target was presented within the centroid zone and 25.86% ($SD = 13.69\%$) occurred when the target was presented outside this zone in the medium-prevalence condition, this difference was not found to be significant. Between Experiments One and Two, there was a significantly greater amount of miss errors within

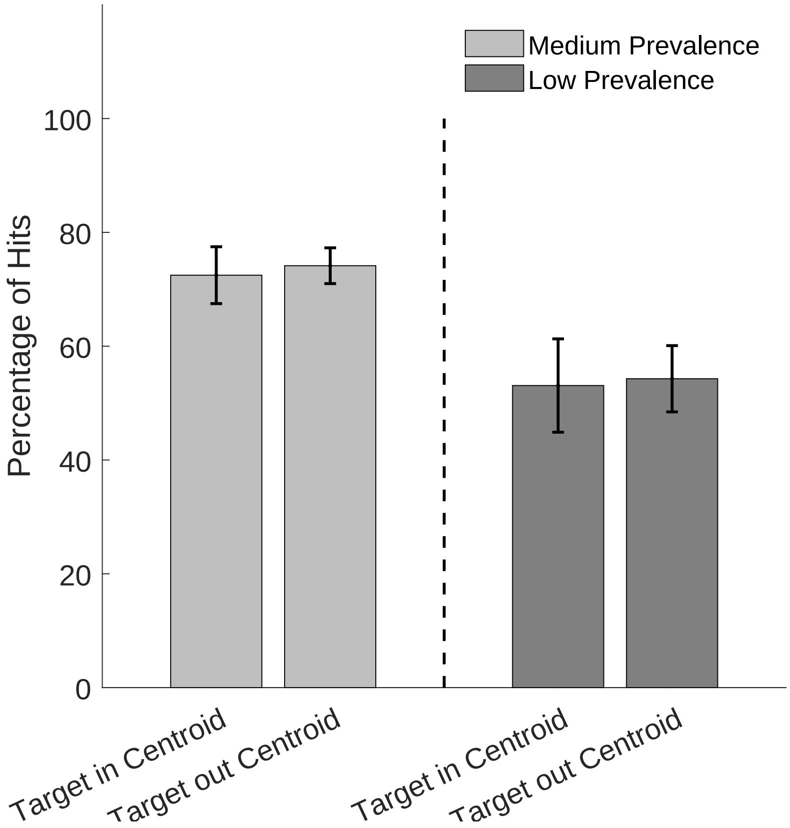

Figure 5.8: The hits which occurred when the target stimulus was presented inside the centroid zone and those which occurred when the target stimulus was presented outside this zone across the prevalence conditions.

the centroid zone ($F_{(1,37)} = 15.52, p < .001, \eta^2_P = 0.30$) and outside this zone ($F_{(1,37)} = 10.13, p < .05, \eta^2_P = 0.22$) in the medium-prevalence condition for the current experiment.

In the low-prevalence condition, on target-present trials the probability of the target stimulus appearing within the centroid zone was 0.17 ($SD = 0.08$) and it was observed that 46.92% ($SD = 35.74\%$) of miss errors occurred within the centroid zone and 44.65% ($SD = 25.67\%$) occurred outside this zone, again, this difference was not found to be significant. As with Experiment One, compared to the medium-prevalence condition there was a greater amount of miss errors occurring in the low-prevalence condition within the centroid zone ($F_{(1,18)} = 4.99, p < .05, \eta^2_P = 0.23$) and outside this zone ($F_{(1,18)} = 27.97, p < .001, \eta^2_P = 0.61$). Between Experiments One and Two, there was a significantly greater amount of

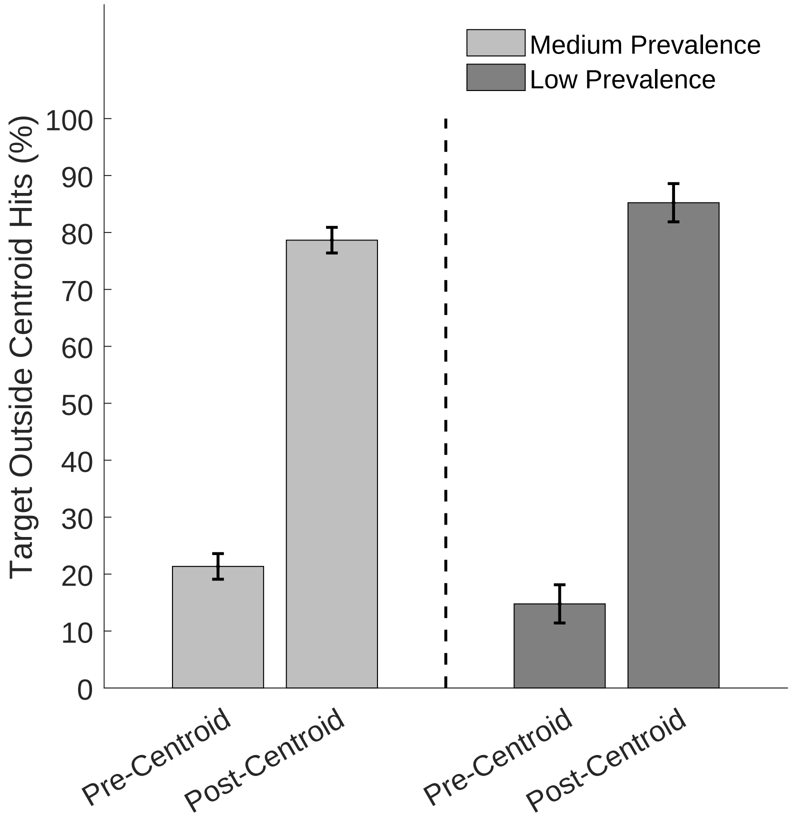

Figure 5.9: The proportion of hits which occurred before entering the centroid zone and the proportion of hits which occurred after leaving this zone across the prevalence conditions.

miss errors in the low-prevalence condition outside the centroid zone ($F_{(1,37)} = 4.32, p < .05, \eta_P^2 = 0.11$) in the current experiment. However, no significant difference was observed within the centroid zone between the experiments. The suggestion from Experiment One that the increase in miss errors which occurred under low-prevalence conditions (Low Prevalence Effect) was primarily occurring within the centroid zone was not supported by the current results. This suggests more intensive search behaviour within the centroid zone as a search strategy when the target stimulus was more difficult to discriminate from the background noise.

Of the hits which occurred outside the centroid zone in the medium-prevalence condition, 21.35% (SD = 9.81) occurred before entering the centroid zone and 78.65% (SD = 9.81) occurred after leaving this zone. When compared to Experiment One, there was a significant reduction in pre-centroid hits ($F_{(1,37)} = 9.85, p < .05, \eta^2_P = 0.22$) and a significant increase in post-centroid hits ($F_{(1,37)} = 9.85, p < .05, \eta^2_P = 0.22$) in the medium-prevalence condition for the current experiment. In the low-prevalence condition, 14.78% ($SD = 14.64$) of hits occurred before entering the centroid zone and 85.22% ($SD = 14.64$) occurred after leaving this zone. When compared to Experiment One, there was a significant reduction in pre-centroid hits ($F_{(1,37)} = 9.59, p < .05, \eta^2_P = 0.21$) and a significant increase in post-centroid hits ($F_{(1,37)} = 9.59, p < .05, \eta^2_P = 0.21$) in the low-prevalence condition for the current experiment. These differences between the experiments suggest that the participants were employing a search strategy in which they were relying on the target stimulus to "pop-out" when the task was less difficult (i.e., Experiment One) and in a more difficult task (i.e., Experiment Two) they search the display more intently.

Table 5.4: The percentage of signal detection miss errors which occurred when the target stimulus was presented inside the centroid zone and the miss errors which occurred when the target stimulus was presented outside the centroid zone for the medium-prevalence condition and the low-prevalence condition in Experiments One and Two. Means are followed by their parenthesised standard deviations. Prev. is an abbreviation for prevalence.

	Medium Prev. Misses		Low Prev. Misses	
	In Centroid (%)	Out Centroid (%)	In Centroid (%)	Out Centroid (%)
Experiment One	6.58(7.90)	13.85(9.12)	42.11(29.06)	29.72(21.96)
Experiment Two	27.52(21.79)	25.86(13.69)	46.92(35.74)	45.72(25.38)

When we analysed the eye-movement behaviour within the centroid zone, we observed that there was a greater proportion of behaviour within this zone compared to the proportion of the centroid area for all three eye-movement measures (see Figure 5.10).

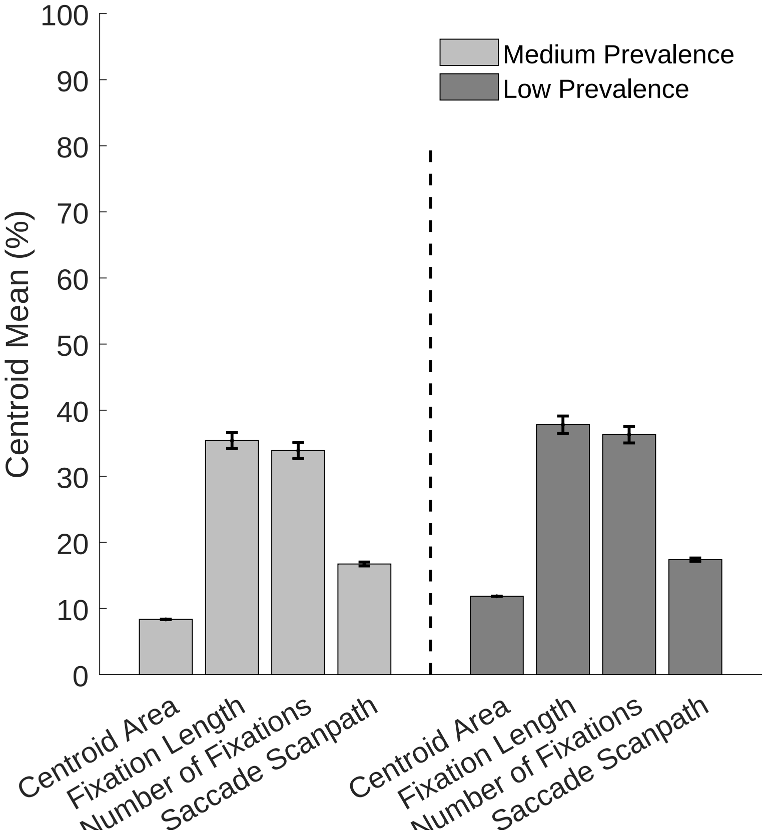

Figure 5.10: Centroid area and eye-movement behaviour within the centroid.

Table 5.5: The trial percentages for the number of fixations, the length of these fixations, and the saccadic scanpath found within the centroid zone for both the medium-prevalence and the low-prevalence conditions for Experiments One and Two. Mean percentages are followed by their parenthesised standard deviations.

	Medium Prevalence (%)			Low Prevalence (%)		
	Fixations	Length	Scanpath	Fixations	Length	Scanpath
Experiment One	31.68(6.86)	32.60(7.03)	17.00(0.95)	37.31(6.83)	38.22(6.89)	17.25(1.72)
Experiment Two	33.86(5.25)	35.35(5.25)	16.74(1.29)	36.33(5.50)	37.85(5.69)	17.39(1.12)

Expanding on this, the fixations found within the centroid zone constituted 35.93% ($SD = 5.12\%$) of the total number of fixations within the stimulus matrix box over a trial. Across the prevalence conditions, there was a greater percentage of fixations found within the centroid zone in the low-prevalence condition ($F_{(1,18)} = 4.47, p < .05, \eta^2_P = 0.20$), which was also seen in Experiment One. Fixation lengths within the centroid zone constituted 37.43% ($SD = 5.25\%$) of the total fixation time within the matrix box over a trial. No significant difference was found between the prevalence conditions; however, in Experiment One fixation lengths were significantly longer in the low-prevalence condition. Scanpath trajectories within the centroid zone constituted only 17.28% ($SD = 1.07\%$) of the total scanpath trajectory found within the matrix box which may be due to the relatively small area of the centroid zone compared to the stimulus matrix. Across prevalence conditions, there was a longer trajectory found in the low-prevalence condition ($F_{(1,18)} = 5.67, p < .05, \eta^2_P = 0.24$). This significantly longer centroid trajectory was not found in Experiment One for the low-prevalence condition.

5.4 Centroid Eye-movement Decay

Experiment One showed that at different stages of their search patterns throughout a trial, the participants changed their eye-movement behaviour at different rates. We were interested in this behavioural change throughout a trial and how this equated to the centroid zone in the current experiment when task difficulty was increased. We hypothesize that the data would be similar across the prevalence conditions before entering the centroid zone and within this zone. However, we hypothesize that there would be a greater amount of decay in eye-movement behaviour after leaving the centroid zone in the low-prevalence condition compared to the medium-prevalence condition. Therefore, we fitted exponential decay functions to the data and analysed the curve parameters for the eye-movement behaviour before entering the centroid zone, within the centroid zone, and after leaving this zone in both the medium- and low-prevalence conditions for the number of fixations, the length of these fixations, and their scanpath trajectories to provide a richer understanding of these patterns and the underlying mechanisms and strategies which may be involved. For the reader's convenience, we report the findings for the number of fixations and only present the data for the length of these fixations, and their scanpath trajectories in table format. Figures 5.11 and 5.12 show the medium- and low-prevalence conditions pre-, post-, and centroid eye-movement behaviour with an exponential decay function fitted to each dataset.

The fitting of the exponential decay function to the medium-prevalence condition fixation data produced a RMS value of 3.31 fixations ($SD = 0.99\ fixations$) for the curve fitted to the pre-centroid data, a RMS value of 3.39 fixations ($SD = 2.05\ fixations$) for the curve fitted to the centroid data, and a RMS value of 4.91 fixations ($SD = 3.58\ fixations$) for the curve fitted to the post-centroid data. The fitting of the decay curves to the low-prevalence data produced a RMS value of 3.30 fixations ($SD = 1.13\ fixations$) for the curve fitted to the pre-centroid data, a RMS value of 3.12 fixations ($SD = 1.23\ fixations$) for the curve fitted to the centroid data, and a RMS value of 4.07 fixations ($SD = 2.10\ fixations$) for the curve fitted to the post-centroid data.

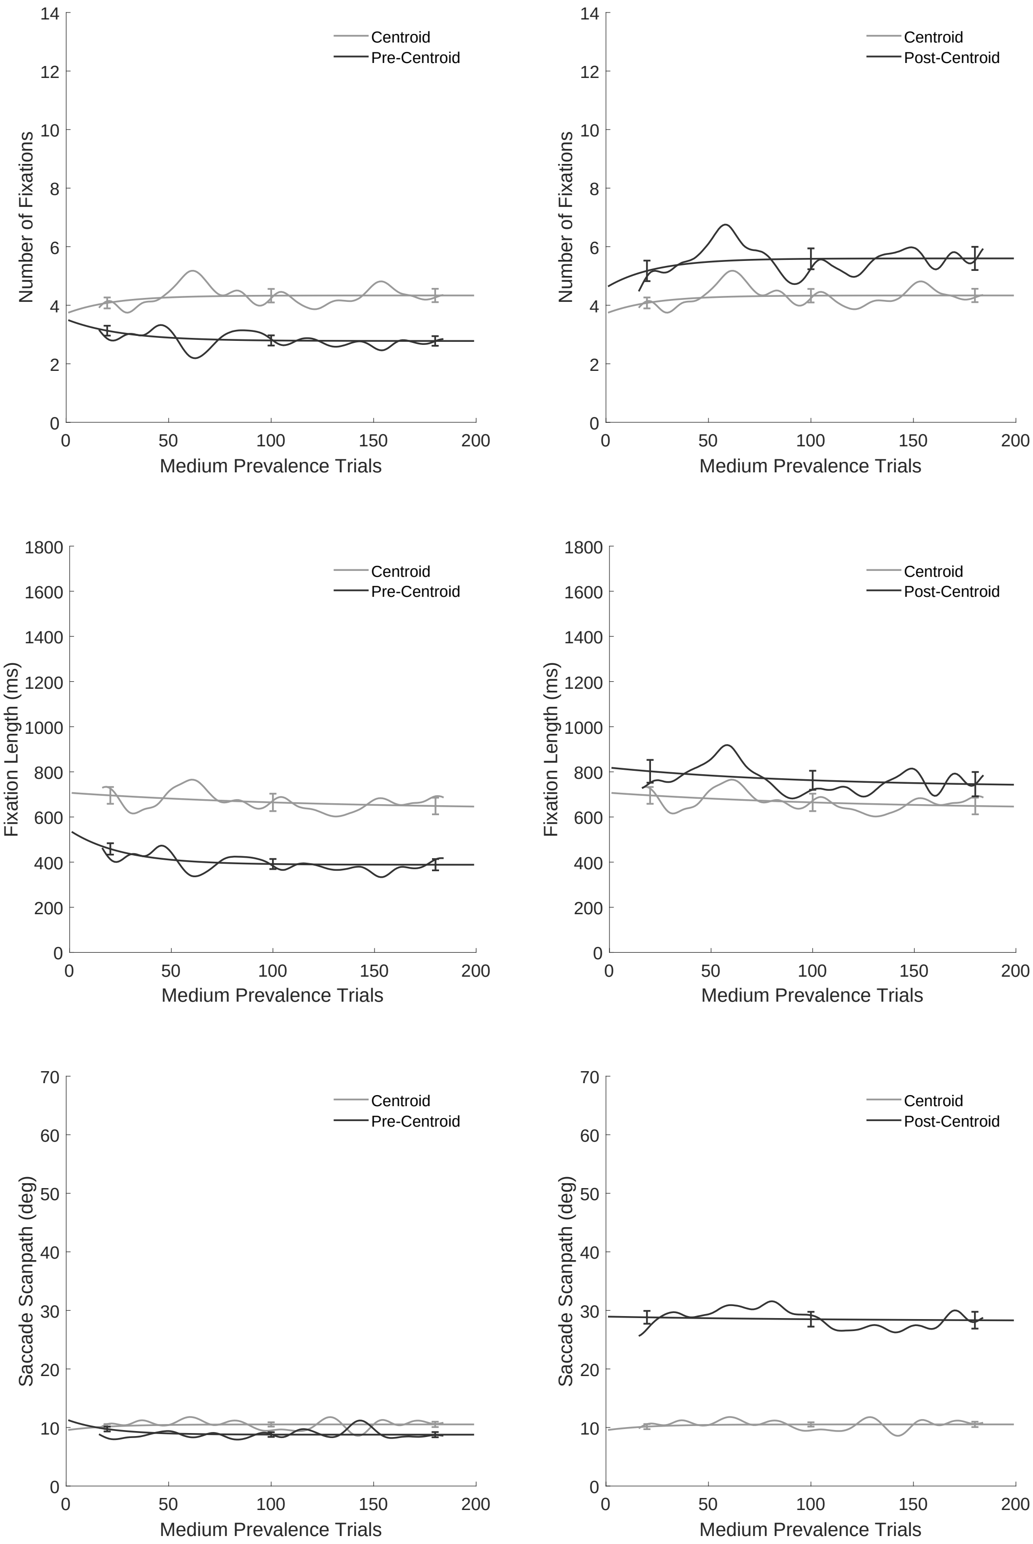

Figure 5.11: The medium-prevalence condition eye-movement behaviour before entering the centroid zone, the eye-movement behaviour within the centroid zone, and the eye-movement behaviour after leaving the centroid zone with an exponential decay curve fitted to each dataset. The error bars represent the standard error found in a sample consisting of fifteen trials either side of the presented error bar.

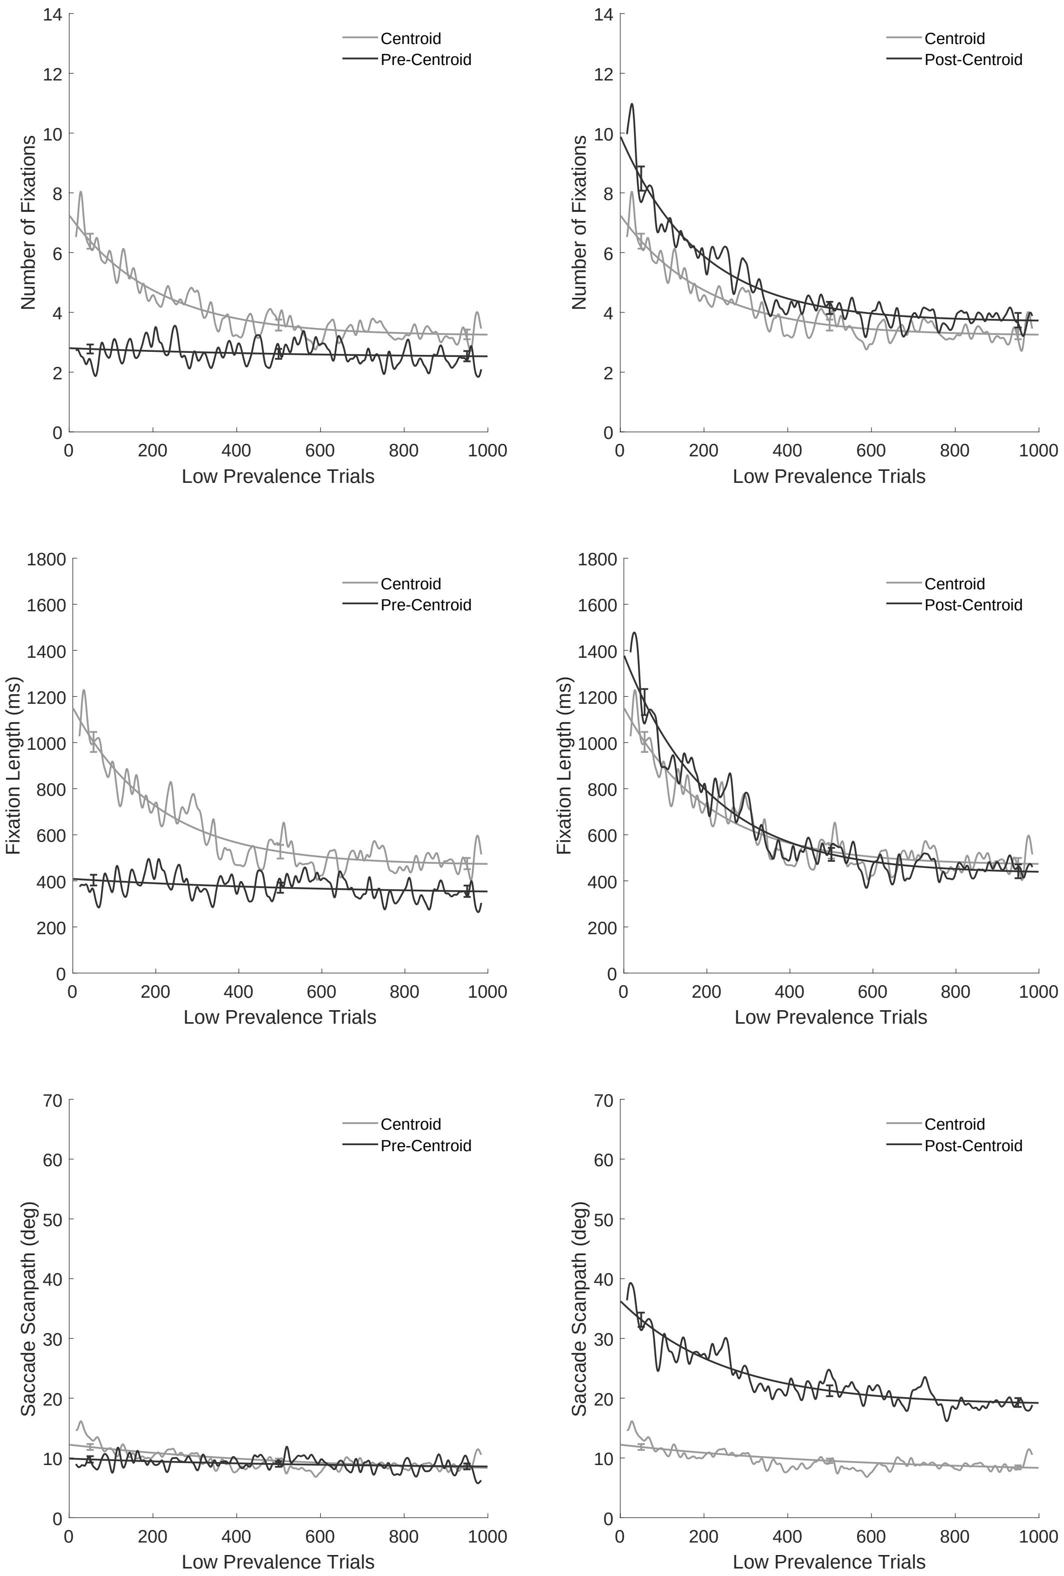

Figure 5.12: These Figures show the low-prevalence eye-movement behaviour before entering the centroid zone, the eye-movement behaviour within the centroid zone, and the eye-movement behaviour after leaving the centroid zone with an exponential decay curve fitted to each dataset. The error bars represent the standard error found in a sample consisting of fifteen trials either side of the presented error bar.

Table 5.6: Pre-centroid, centroid, and post-centroid exponential decay curve parameters across prevalence conditions for Experiment One. Means are followed by their parenthesised standard deviations.

	Experiment One					
	Medium Prevalence			Low Prevalence		
	Pre-	Centroid	Post-	Pre-	Centroid	Post-
Fixations						
Decay (*n*)	0.22(0.98)	1.01(2.01)	1.26(2.81)	0.37(0.81)	1.90(1.15)	3.01(1.89)
Rate (*trial*)	72.93(49.00)	67.32(47.55)	93.65(40.36)	442.92(209.43)	324.10(196.38)	308.55(172.23)
Asymptote (*n*)	2.72(1.17)	2.37(0.70)	2.37(0.81)	2.13(0.65)	1.97(0.75)	1.40(0.93)
Length						
Decay (*ms*)	91.47(92.50)	200.26(214.48)	236.92(211.36)	66.87(119.17)	317.04(209.86)	427.40(329.39)
Rate (*trial*)	82.05(40.52)	64.21(44.67)	75.49(46.52)	446.15(168.03)	255.71(168.45)	296.44(168.83)
Asymptote (*ms*)	368.15(154.87)	340.88(113.74)	304.32(91.76)	296.60(79.30)	286.93(108.34)	176.04(110.38)
Scanpath						
Decay (*deg*)	3.66(2.83)	0.43(2.07)	3.20(6.04)	1.16(1.03)	4.42(2.21)	9.33(6.66)
Rate (*trial*)	47.89(44.01)	85.11(39.34)	88.06(39.21)	512.57(179.20)	129.89(45.42)	422.74(213.30)
Asymptote (*deg*)	8.25(0.55)	7.76(1.19)	16.92(3.91)	8.06(0.51)	6.71(1.60)	12.29(4.64)

Table 5.7: Pre-centroid, centroid, and post-centroid exponential decay curve parameters across prevalence conditions for Experiment Two. Means are followed by their parenthesised standard deviations.

	Experiment Two					
	Medium Prevalence			Low Prevalence		
	Pre-	Centroid	Post-	Pre-	Centroid	Post-
Fixations						
Decay (*n*)	0.81(1.26)	0.19(1.43)	0.26(2.12)	0.33(0.96)	4.49(3.10)	6.77(6.06)
Rate (*trial*)	73.65(45.57)	74.39(43.68)	83.57(38.50)	416.09(223.78)	328.23(193.58)	351.84(175.10)
Asymptote (*n*)	2.79(0.98)	4.21(2.42)	5.50(4.06)	2.48(1.04)	2.96(1.61)	3.13(3.15)
Length						
Decay (*ms*)	198.75(178.14)	208.77(238.70)	294.90(471.31)	58.89(91.72)	732.93(481.11)	957.32(788.20)
Rate (*trial*)	65.83(44.21)	86.14(34.01)	75.34(38.67)	259.17(208.87)	277.29(195.68)	354.33(202.94)
Asymptote (*ms*)	376.98(106.38)	595.41(362.51)	700.35(551.61)	390.27(134.86)	457.49(219.81)	393.71(322.38)
Scanpath						
Decay (*deg*)	2.70(4.61)	0.66(3.15)	4.83(7.42)	1.45(1.69)	7.30(5.59)	20.63(16.23)
Rate (*trial*)	65.81(45.65)	80.89(42.56)	64.84(38.92)	501.31(194.55)	224.08(160.55)	342.95(196.73)
Asymptote (*deg*)	8.81(1.08)	10.15(3.56)	27.17(17.02)	8.58(1.01)	8.35(2.05)	17.88(7.35)

As can be seen in Table 5.7, the magnitude of the curve's decay for the number of fixations within a trial in the medium-prevalence condition was the greatest before entering the centroid zone and the least amount of fixational decay occurred within the centroid zone. However, no significant differences were found between the pre-, post-, and centroid means. Interestingly in the low-prevalence condition, the least amount of decay occurred before entering the centroid zone and the greatest amount occurred after leaving the centroid zone ($F_{(2,17)} = 18.06, p < .001, \eta_P^2 = 0.50$), with the posthoc analysis showing that the main significant difference occurred between the pre-centroid and centroid means ($p < .001$), a larger p-value occurred between the pre-centroid and post-centroid means ($p = .001$), and the largest p-value occurred between the centroid and post-centroid means ($p = .028$). These results were consistent with that found in Experiment One, in that no significant differences were found between the pre-, post-, and centroid means in the medium-prevalence condition and in the low-prevalence condition, the least amount of decay occurred before entering the centroid zone and a significantly greater amount occurred after leaving the centroid zone.

The rate of this pre-, post-, and centroid decay in fixations before reaching the asymptote in the medium-prevalence condition showed that the participants were the quickest to reach a steady state of behaviour before entering the centroid zone and took the longest amount of time to reach this state after leaving the centroid zone. However, no significant differences were found between the means. In the low-prevalence condition, participants were the quickest to reach the asymptote within the centroid zone and took the longest amount of time to reach this state before entering the centroid zone. Again, no significant differences were found between means. These results were, again, consistent with that found in Experiment One, in that no significant differences were observed in the pre-asymptote rate of decay for both the medium- and low-prevalence conditions.

Once the participants had reached a steady state of fixational behaviour in the medium-prevalence condition, we observed that this level of behaviour was the lowest before entering the centroid zone and the highest after leaving the centroid zone ($F_{(2,17)} = 10.66, p < .001, \eta_P^2 = 0.37$), with the posthoc analysis showing that the main significant difference occurred between the pre-centroid and post-centroid means ($p = .010$), a larger p-value occurred between the centroid and post-centroid means ($p = .015$), and the largest p-value occurred between the pre-centroid and centroid means ($p = .023$). In Experiment One, no significant differences were found. In the low-prevalence condition, the lowest level of behaviour occurred before entering the centroid zone and the highest after leaving the centroid zone. However, with only 0.64 of a fixation separating these data categories, no significant differences were found between the asymptote means. This consistency between the pre-, post-, and centroid asymptotes was also seen in Experiment One.

Between the prevalence conditions, we found no significant difference between the conditions for the pre-centroid fixational decay which shows the consistency in behaviour across the conditions during their scene preview. However, there was a greater amount of decay within the centroid zone ($F_{(1,18)} = 24.98, p < .001, \eta_P^2 = 0.58$) and after leaving this zone ($F_{(1,18)} = 25.85, p < .001, \eta_P^2 = 0.59$) in the low-prevalence condition. The centroid results contrasted with those found in Experiment One, which only found a significantly greater amount of decay in the low-prevalence condition after leaving the centroid zone. Once the participants had reached their pre-centroid asymptote in behaviour, no significant difference was observed between the conditions with only 0.31 of a fixation separating the means. However, the asymptote levels of behaviour were greater within the centroid zone ($F_{(1,18)} = 9.75, p < .05, \eta_P^2 = 0.35$) and after leaving this zone ($F_{(1,18)} = 17.70, p < .001, \eta_P^2 = 0.50$) in the medium-prevalence condition due to the abundance of target-present trials in this condition and the reinforcing effects of finding what you are looking for. This was also the case for Experiment One with greater levels of behaviour in the medium-prevalence condition across all three data categories when the task was less difficult.

Across Experiments One and Two, no significant differences were observed for the pre-, post-, and centroid behavioural decay in the medium-prevalence condition across the experiments. In the low-prevalence condition, no significant difference was observed in the decay of the pre-centroid behaviour; however, a significantly greater amount of decay occurred within the centroid zone ($F_{(1,37)} = 11.69, p < .05, \eta_P^2 = 0.25$) and after leaving this zone ($F_{(1,37)} = 6.68, p < .05, \eta_P^2 = 0.16$) in Experiment Two which may be due the reduced rate of obtained reinforcement within this experiment. No significant differences were found between the Experiments for the rate of this pre-, post-, and centroid decay in fixations in both the medium- and low-prevalence conditions. Once the asymptote level of behaviour had been reached in the medium-prevalence condition, this level of behaviour was significantly greater within the centroid zone ($F_{(1,37)} = 10.12, p < .05, \eta_P^2 = 0.22$) and after leaving this zone ($F_{(1,37)} = 10.92, p < .05, \eta_P^2 = 0.23$) in Experiment Two. This was also the case within the centroid zone ($F_{(1,37)} = 5.89, p < .05, \eta_P^2 = 0.14$) and after leaving this zone ($F_{(1,37)} = 5.22, p < .05, \eta_P^2 = 0.13$) in the low-prevalence condition for Experiment Two. These results show the consistency of behaviour across the experiments during the pre-centroid scene preview and that when task demands are increased there is an increase in search behaviour within the centroid zone and after leaving this zone.

Overall, these findings confirm our hypothesises that there would be an increase in behaviour when the difficulty of the task was increased and and that search behaviour is being guided towards the centroid zone during their scene preview. However, the results disprove our hypothesis that the decay in search behaviour predominantly occurs after leaving the centroid zone in the low-prevalence condition when obtained reinforcement is scarce as the results showed that when task demands are increased this weakening of behaviour occurs both within the centroid zone and after leaving this zone in the low-prevalence condition. This weakening of behaviour within the centroid zone and after leaving this zone in Experiment Two which only occurred post-centroid in Experiment One may reflect the reduced rate of reinforcement observed in Experiment Two when task difficulty was increased.

5.5 Discussion

5.5.1 Visual Search

The dominant theories of visual search (J. Duncan & Humphreys, 1992; A. M. Treisman & Gelade, 1980; Wolfe et al., 1989) assume that search is a process involving comparisons of individual items against a target description that is based on the properties of the target in isolation, and an understanding of these parameters which determine our choice of visual search strategy may, in turn, shed light on visual behaviour within everyday situations. Therefore, it is critical to establish an understanding of the underlying strategies and mechanisms behind visual search behaviour. With this in mind, we were interested in establishing whether the search strategies (i.e., centroid eye-movement behaviour) and the mechanisms (i.e., Reinforcement Strengthening Hypothesis) which were observed in Experiment One would also manifest themselves in Experiment Two when the task difficulty was increased.

Therefore, in Experiment Two, we reduced the disparity between the target stimulus and the distractor items by titrating the discrimination threshold estimate for each individual participant using the two-down/one-up staircase method (Levitt, 1971) (see Figures 5.1). All other aspects of Experiment One were repeated in Experiment Two. This staircase procedure standardizes the sample's visual discriminability and sets task difficulty at a probability of 0.71 (Lu & Dosher, 2013). In the signal detection model, this equates to an increase in the noise distribution (D. M. Green et al., 1966) and most theories of visual search would predict response latencies and performance errors to increase when the difficulty of the task is increased (Heinke & Humphreys, 2003; Itti & Koch, 2000; A. M. Treisman & Gelade, 1980; Zelinsky, 2008).

Experiment Two produced the familiar low prevalence effect, with a significantly greater amount of miss errors occurring in the low-prevalence condition compared to the medium-prevalence condition, while no significant difference was observed between the prevalence conditions for the false alarm errors (see Figure 5.4).

We again observed that the eye-movement behaviour metrics (i.e., the number of fixations, their lengths, and their scanpath trajectories) stabilized more rapidly and, once a steady state of behaviour had been reached, there was a significantly greater number of fixations, their lengths were longer, and their trajectories were also longer throughout the medium-prevalence condition compared to the low-prevalence condition (see Figures 5.6). Furthermore, when compared to the medium-prevalence condition, there was a significantly greater amount of decay in eye-movement behaviour, the participants took a significantly longer amount of time to reach a steady state of behaviour and, once this state had been reached, the level of this behaviour was significantly lower in the low-prevalence condition for all three of our measures of eye-movement behaviour. These patterns of eye-movement behaviour observed throughout each prevalence condition add further weight to the Reinforcement Strengthening Hypothesis, which states that: When signal detection functions as reinforcement for visual search behaviour, operant theory would predict more intense searching behaviours and improved detection performance when target signals are plentiful (i.e., search behaviour is strengthened). Conversely, the absence of target signals for long periods of time would weaken search behaviours and increase signal detection errors.

It has been demonstrated that variability may be controlled by making reinforcement contingent upon variability itself. Variability is an operant dimension, and functional because it allows the shaping of new behaviours (Blough, 1966; Machado, 1989; Neuringer, 2002; Pryor et al., 1969). Therefore, it may be suggested that part of the variability observed in motor responses is related to environmental contingencies, such that a given range of behavioural instances of an operant response may be reinforced. This suggests that oculomotor responses (i.e., search behaviour) may also be reinforced by environmental contingencies and there is a growing body of evidence to support this understanding (Montagnini & Chelazzi, 2005; Madelain et al., 2008, 2010; Madelain, Paeye, & Darcheville, 2011; Paeye et al., 2007; Paeye & Madelain, 2011).

In this respect, voluntary eye movements are controlled by operant learning processes relying on the outcome of each movement given a particular state of the environment; behaviours that produce a favourable outcome are selected and strengthened while behaviours that do not are weakened and rejected. Accordingly, eye movements may be controlled by the ability to perform visually guided behaviours that act as reinforcers because of their adaptive values and viewing the target may act as reinforcement because maximizing visual perception would allow other adaptive interactions with the environment. In other words, environmental contingencies are controlling eye-movement behaviour (Madelain, Paeye, & Darcheville, 2011, for a review). In the current study, it should come as no surprise that there is an increase in search behaviour when target signals are abundant due to operant learning processes and the reinforcing effect of finding what you are looking for.

It was again observed that when compared to the correct rejection mean, there was a significant increase in eye-movement behaviour for the first correct rejection immediately following a correct target-present response (see Figures 5.5). This pattern of behaviour was observed for all three of our eye-movement behavioural measures across both the prevalence conditions, except for the length of fixations in the medium-prevalence condition. This increase in search behaviour immediately following a correct target-present response (i.e., identifying the target) may be viewed as the local effects of reinforcement and a manifestation of the strengthening of search behaviour, which is primarily influencing the degree to which the display is searched across a trial (i.e., scanpath trajectory) and the number of fixations within a trial. Furthermore, the abundance of target-present trials throughout the medium-prevalence condition, to which the correct response is a hit, and the post-hit increase in correct rejection eye-movement behaviour would maintain search behaviour at a higher rate across the course of this condition, which is what we are seeing. In other words, the schedule of signal presentation is controlling eye-movement behaviour through the operant mechanisms of reinforcement and extinction.

5.5.2 The Stimulus Density Map

We were also interested in the progression of eye-movement behaviour throughout a trial, to provide an understanding of the search strategies which may be involved. Therefore, we constructed the centroid zone, which represented the portion of the display with the greatest density of stimulus items.

We again observed that there was a much greater proportion of eye-movement behaviour within the centroid zone compared to the relative surface area of this zone within the stimulus matrix box for both the prevalence conditions (see Figure 5.10). Of interest was the observation that there were no significant differences found for miss errors between the prevalence conditions within the centroid zone.

We then blocked the eye-movement curve parameters down into the pre-, post-, and centroid sub-categories for the number of fixations in each of the prevalence conditions and we observed that, between the prevalence conditions, there was a significantly greater amount of decay within the centroid zone and after leaving this zone in the low-prevalence condition. Interestingly, no significant pre-centroid differences in the amount of decay were observed between the prevalence conditions. Once a steady state of eye-movement behaviour had been reached, which took a significantly longer amount of time in the low-prevalence condition, eye-movement behaviour was significantly greater within the centroid zone and after leaving this zone in the medium-prevalence condition while no significant difference was observed between the prevalence conditions for the pre-centroid behaviour (see Figures 5.11 and 5.12).

These observations suggest that the patterns of eye-movement behaviour may be considered consistent across the conditions with the least amount of decay occurring during their "scene preview" before entering the centroid zone. However, the weakening of search behaviour (i.e., a greatest amount of decay) predominantly occurred within the centroid zone and after leaving this zone in the low-prevalence condition which may be due to the scarcity of obtained reinforcement weakening search behaviour within this condition. This weakening of search behaviour would increase the rate of miss errors and result in a shift in the decision criterion throughout this condition. Once asymptote had been reached in the medium-prevalence condition, there was a significantly greater amount of post-centroid behaviour for all of our three measures of eye-movement behaviour, and this was also the case within the centroid zone for the number of fixations and their scanpath trajectories. These observations suggest that the strengthening of behaviour predominantly occurs within the centroid zone and after leaving this zone, which may be due to the abundance of target-present trials in this condition and the reinforcing effects of finding what you are looking for, which predominantly manifests itself as an increase in eye-movement behaviour for the first post-hit correct rejection and over the course of this condition would increase the eye-movement behaviour for target-absent trials.

These patterns of behaviour observed within each prevalence condition may be considered consistent with the understanding that if the target appeared close to the point of fixation upon the presentation of a target-present trial or the participant was able to identify the target within their visual periphery, they would saccade directly towards the target and terminate the trial thereafter, before entering the centroid zone. If this was not the case, they would saccade towards the centroid zone and if the target was within this zone, there was a high probability that they would find the target. If the target was not within the centroid zone, they would begin to search outside this zone and often terminate the trial before they had identified the target in the low-prevalence condition due to the scarcity of obtained reinforcement throughout this condition and the weakening of search behaviour.

Interestingly, it was observed that for the absolude means, there was a significantly greater number of fixations, and the scanpath trajectories were longer within the centroid zone in the low-prevalence condition suggesting that when reinforcement is lean (i.e., low-prevalence condition) there is an increase in search behaviour within this zone. However, when we examined the centroid curve asymptote parameters, we discovered that there was a significantly greater number of centroid fixations in the medium-prevalence condition once a steady state of behaviour had been reached and this pattern of behaviour was also consistent for the centroid scanpaths. In other words, there was an initial increase in eye-movement behaviour within the centroid zone in the low-prevalence condition compared to the medium-prevalence condition. However, over time this behaviour decays to the point where there is significantly less eye-movement behaviour, in terms of fixations and their scanpath trajectories, in the low-prevalence condition. This pattern of behaviour is reminiscent of an *extinction burst*, in which there is an initial increase in response frequency (i.e., a behaviour) when reinforcement is no longer forthcoming, or much less frequently, and then the dependent behaviour subsides (Fisher, Greer, Shahan, & Norris, 2023; Harris, Pentel, & Lesage, 2007; Lerman, Iwata, & Wallace, 1999; Pushparaj, Pryslawsky, Forget, Yan, & Le Foll, 2012; Shahan, 2022).

5.5.3 Experimental Comparisons

Importantly, it should be noted that the cohorts for each experiment were recruited over different timeframes and the sample was not randomly assigned to each experimental group. Therefore, firm conclusions cannot be drawn from the comparisons between the experiments. With this in mind, we tentatively present some of the key findings across Experiments One and Two.

The current findings in Experiment Two may be considered consistent with those observed in Experiment One; however, there were some subtle differences. For example, compared to Experiment One, there was a significantly greater amount of eye-movement behaviour across both the prevalence conditions in the current experiment, this increase in eye-movement behaviour is consistent with visual search reaction time studies, which have shown that reaction times become slower (i.e., search becomes more inefficient) when the task difficulty is increased (A. M. Treisman & Gelade, 1980; Wolfe, 1998). Furthermore, the physiology of the retina limits the number of items that may be processed within a single fixation (Gegenfurtner, 2016). However, it has also been shown that the size of the functional visual field (FVF: the number of items processed during a fixation) does not only depend on visual acuity, but also the difficulty of the task (Geisler & Chou, 1995; Hulleman & Olivers, 2017; Hulleman, Lund, & Skarratt, 2019; Motter & Simoni, 2008). For example, Geisler and Chou (1995) have shown that the FVF was smaller for a conjunction of colour and orientation task than for an orientation feature task. In other words, the FVF shrinks when the task becomes more difficult, and this implies that when a search task is easy, multiple items may be processed within a single fixation (i.e., a parallel process) whereas when a search is difficult, only a single item is processed with each fixation (i.e., a serial process), which may explain why we see an increase in eye-movement behaviour when the task difficulty was increased.

With this in mind, it was also observed that a significantly greater amount of miss errors occurred across both the prevalence conditions in the current experiment compared to Experiment One, which may be considered consistent with an increase in the signal detection noise distribution (D. M. Green et al., 1966; J. Nevin, 1969; Mccarthy & Davison, 1979) and current theories of visual search (Heinke & Humphreys, 2003; Itti & Koch, 2000; A. M. Treisman & Gelade, 1980; Wolfe, 1994; Zelinsky, 2008) and we hypothesized that increasing task difficulty would increase search behaviour and performance errors which is what we found.

Of interest was the observation that there were no significant differences found for miss errors between the prevalence conditions within the centroid zone in the current experiment. These results contrasted with those observed in Experiment One, which found a significantly higher rate of miss errors within the centroid zone for the low-prevalence condition and suggested that this increased amount of miss errors under low-prevalence conditions (Low Prevalence Effect) was primarily occurring within the centroid zone in Experiment One. However, the current findings in Experiment Two may simply reflect an increase in search behaviour within the centroid zone when task difficulty is increased.

In summary, in the current experiment, the disparity between the target stimulus and the distractor items was reduced (i.e., task difficulty increased) and standardized across the sample. It was observed that there was a significantly greater amount of eye-movement behaviour across both the prevalence conditions, despite there also being a significantly greater amount of decay in search behaviour within the low-prevalence condition, miss errors also significantly increased, and the sample variance was reduced. In line with the Reinforcement Strengthening Hypothesis, it may be suggested that the greater amount of decay observed within the low-prevalence condition represents the weakening of search behaviour when obtained reinforcement was reduced (i.e., fewer targets found) and this weakening of search behaviour (i.e., eye-movement decay) primarily occurred within the centroid zone and after leaving this zone in the low-prevalence condition. Furthermore, there was, again, a post-hit correct rejection increase in search behaviour which may be viewed as the local effects of reinforcement (Jia et al., 2021; Williams, 1976) and a manifestation of the strengthening of search behaviour.

As with Experiment One, there was a much greater proportion of eye-movement behaviour within the centroid zone compared to the relative surface area of this zone within the stimulus matrix box for both the prevalence conditions and our data suggests that eye-movement behaviour is drawn or guided towards this zone as an efficient search strategy.

Experiment Three

We know that the frequency of a target signal influences search behaviour dramatically in systematic ways (D. M. Green et al., 1966). For example, in visual search tasks, the likelihood that an observer will miss a target signal varies with the frequency in which the target stimulus is presented across trials (Ishibashi et al., 2012; Rich et al., 2008; J. D. Schwark et al., 2013; Wolfe et al., 2005, 2007). And we asked the question, "Is the schedule of signal presentation controlling eye-movement behaviour and if so, how?" We suggest that search behaviour may itself be influenced by the outcomes of search behaviour through the operant processes of reinforcement and extinction. Therefore, in Experiment Three we removed all target-present trials in the low-prevalence condition. Specifically, the participants experienced a target probability of 0.50 in the medium-prevalence condition and in the low-prevalence condition they experienced a target probability of 0.00. Experiment Three replicated Experiment Two in all other aspects of experimental design. In line with the Reinforcement Strengthening Hypothesis, we predict that this schedule manipulation would result in fewer eye movements and a greater decay in eye-movement behaviour across the low-prevalence condition in Experiment Three compared to Experiment Two, while the findings in the medium-prevalence condition in Experiment Three would be like those found in Experiment Two. As with Experiment Two, in Experiment Three the participant's discrimination threshold estimate values were individually titrated using the two-down/one-up single staircase method to normalize participant variance in visual discriminability across the sample and ensured that the task difficulty was similarly set across Experiments Two and Three for all participants.

6.1 Participants

Twenty participants (6 male & 14 female) were recruited from the University of Waikato's student body. They ranged in age between 17.9 – 31.3 years ($M = 22.2$ years, $SD = 4.2$ years). Two participants were excluded from the eye-movement analysis due to head movement during the experiment which was two standard deviations above the sample mean (exclusion criteria).

6.2 Signal Detection Theory Measures

As with the previous Experiments, the z-transformed hit rate was plotted against the z-transformed false alarm rate for each participant in the medium-prevalence condition and a linear regression line was fitted (see Figure 6.1). As with the previous experiments, the underlying assumptions of normality and equal variance for the signal detection indices of d' and c were found to be violated with the medium-prevalence regression line showing a slope of 0.19 and an intercept of -0.02. As can be seen in Figure 6.1 there was a large amount of participant variance in the small sample resulting in a poor fit for the regression line ($R^2 = 0.04$). As a result, the parametric signal detection indices of d' and c were unable to be reported.

We then applied the non-parametric summary measures of log d (i.e., stimulus discriminability or sensitivity) and log B (i.e., response bias or decision criterion). As there were no target-present trials in the low-prevalence condition for Experiment Three, response bias, as measured by log B (see equation 4.6) and discriminability, as measured by log d (see equation 4.5) were unable to be calculated between the prevalence conditions. However, no significant differences were observed in the medium-prevalence condition between Experiments Two and Three for both log d and log B (M. C. Davison & Tustin, 1978).

As can be seen in Table 6.1, the medium-prevalence condition produced an average of 23.04% ($SD = 11.66\%$) for miss errors. However, as no target stimuli were presented within the low-prevalence condition, a comparison in miss errors between prevalence conditions was unable to be calculated. False alarm errors were found to be similar across prevalence conditions, with an average of 9.22%

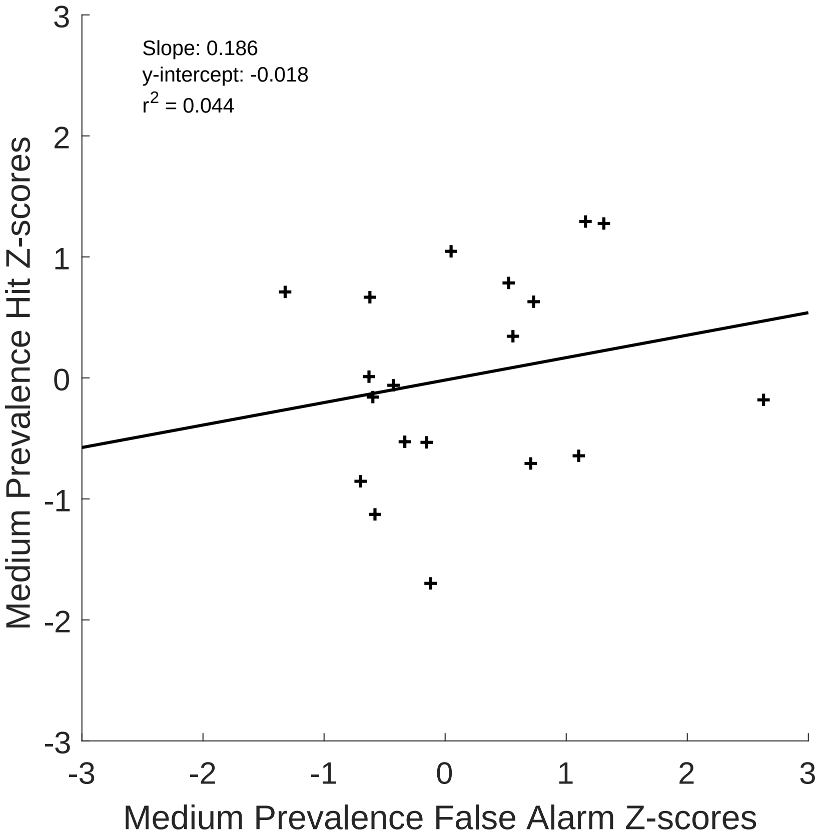

Figure 6.1: The fitted regression line to medium-prevalence hit and false alarm z-scores for each participant.

($SD = 13.22\%$) false alarm errors occurring within the medium-prevalence condition and 9.82% ($SD = 18.90\%$) occurring in the low-prevalence condition. Figure 6.2 shows the signal detection errors across prevalence conditions for Experiment Three.

Across the experiments in the medium-prevalence condition (see Table 6.1), miss errors increased considerably in Experiment Two and were similar between Experiments Two and Three ($F_{(2,58)} = 6.57, p < .05, \eta_P^2 = 0.19$), with the posthoc analysis showing a main significant difference occurring between Experiments One and Two ($p = .003$) and a larger p-value occurring between Experiments One and Three ($p = .032$). The false alarm errors within the medium-prevalence condition increased slightly in Experiment Two and increased again in Experiment Three, no significant differences were found between the experimental means.

Table 6.1: The signal detection theory miss and false alarm errors for the medium-prevalence condition and the low-prevalence condition across the Experiments.

	Medium Prevalence		Low Prevalence	
	Miss (%)	False Alarm (%)	Miss (%)	False Alarm (%)
Experiment One	13.20(8.49)	3.23(5.92)	32.25(21.12)	2.18(4.32)
Experiment Two	26.10(14.40)	4.04(4.25)	47.00(23.86)	3.28(4.99)
Experiment Three	23.04(11.66)	9.22(13.22)	NA	9.82(18.90)

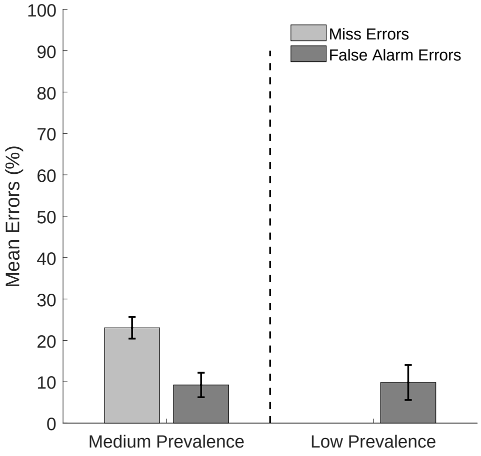

Figure 6.2: Signal detection performance errors across the prevalence conditions for Experiment Three.

Note that miss errors could not occur in the low-prevalence condition for Experiment Three. The false alarm errors in the low-prevalence conditions increased slightly in Experiment Two and again in Experiment Three, no significant differences were found between the means. These elevated rates of false alarm errors in both prevalence conditions for Experiment Three may be attributed to one participant in the medium-prevalence condition (55.00%) and two participants in the low-prevalence condition (53.95% & 60.36%) reporting false alarm rates that were greater than two standard deviations above the mean.

As with the previous Experiments, it was observed that when compared to the correct rejection mean there was increased eye-movement behaviour for the first correct rejection immediately following a correct target-present response in the medium-prevalence condition (see Figures 6.3). This pattern of behaviour was observed for the medium-prevalence eye-movement measures for the number

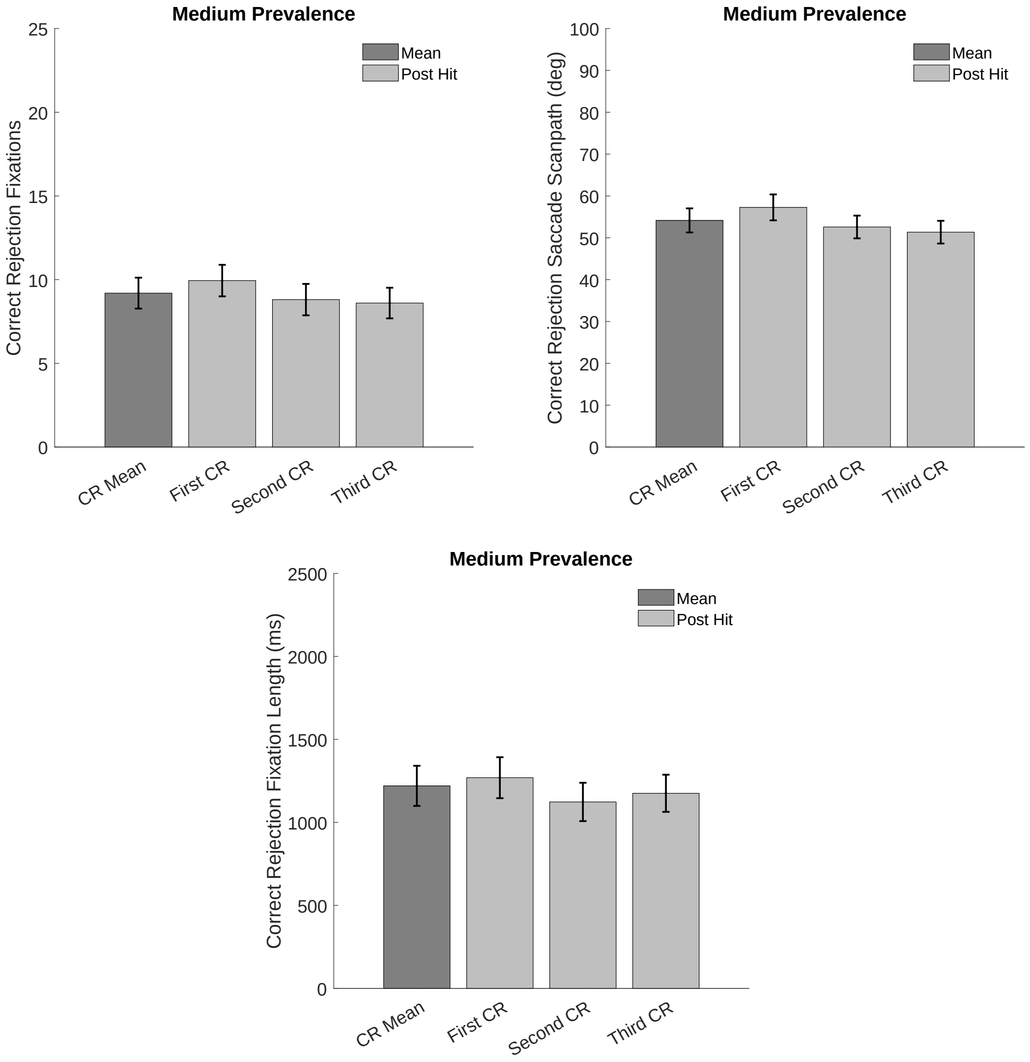

Figure 6.3: Eye-movement behaviour for the correct rejection mean in the medium-prevalence condition, the first correct rejection immediately following a hit, the second consecutive correct rejection following a hit, and the third consecutive correct rejection following a hit.

of fixations ($F_{(1,17)} = 19.66, p < .001, \eta_P^2 = 0.54$), and the saccadic scanpath ($F_{(1,17)} = 32.51, p < .001, \eta_P^2 = 0.66$), no difference was observed between the fixation length correct rejection mean and the first post-hit correct rejection. This pattern of behaviour has been consistently observed across all experiments and is in keeping with the strengthening of search behaviour through the local effects of reinforcement when we find what we are looking for.

Table 6.2: Eye-movement behaviour for the correct rejection mean and the first correct rejection immediately following a correct target-present response across prevalence conditions and Experiments. Means are followed by their parenthesised standard deviations. CR is an abbreviation for correct rejection.

	Medium Prevalence		Low Prevalence	
	CR Mean	Post-hit CR	CR Mean	Post-hit CR
Experiment One				
Fixations (*n*)	8.45(3.79)	9.05(3.67)	5.89(2.20)	7.31(2.28)
Length (*ms*)	982.69(442.47)	1043.50(444.77)	737.52(313.76)	741.50(259.36)
Scanpath (*deg*)	50.34(12.40)	53.02(12.00)	40.83(8.61)	46.26(8.67)
Experiment Two				
Fixations (*n*)	14.18(10.88)	15.58(11.09)	10.62(5.78)	14.35(6.52)
Length (*ms*)	1731.96(1132.01)	1847.16(1113.32)	1386.08(679.08)	1760.98(916.32)
Scanpath (*deg*)	69.10(35.66)	74.93(35.25)	54.31(13.36)	67.85(19.67)
Experiment Three				
Fixations (*n*)	9.19(3.92)	9.94(4.00)	NA	NA
Length (*ms*)	1220.11(512.49)	1269.06(523.85)	NA	NA
Scanpath (*deg*)	54.15(12.21)	57.27(13.14)	NA	NA

6.3 Eye-movement Decay

We had observed that over time (trials) some of the measures dropped off in their values, suggesting a particular change in eye-movement behaviour over time and we hypothesize that this pattern of behaviour is the result of operant learning due to the processes of reinforcement and extinction across each condition, we predict that there will be a greater extinction of behaviour throughout the low-prevalence condition when all reinforcement is removed in the current experiment compared to Experiment Two. To capture this trend, we fitted the data with an exponential decay function for both the medium- and low-prevalence conditions and the three curve parameters were analysed. These parameters represent the magnitude of the curve fall (represented on the Y-axis), the rate of this fall (represented on the X-axis), and the asymptote (represented on the Y-axis). Figure 5.6 shows the eye-movement behaviour found in the low- and medium-prevalence conditions with an exponential decay function fitted to each dataset.

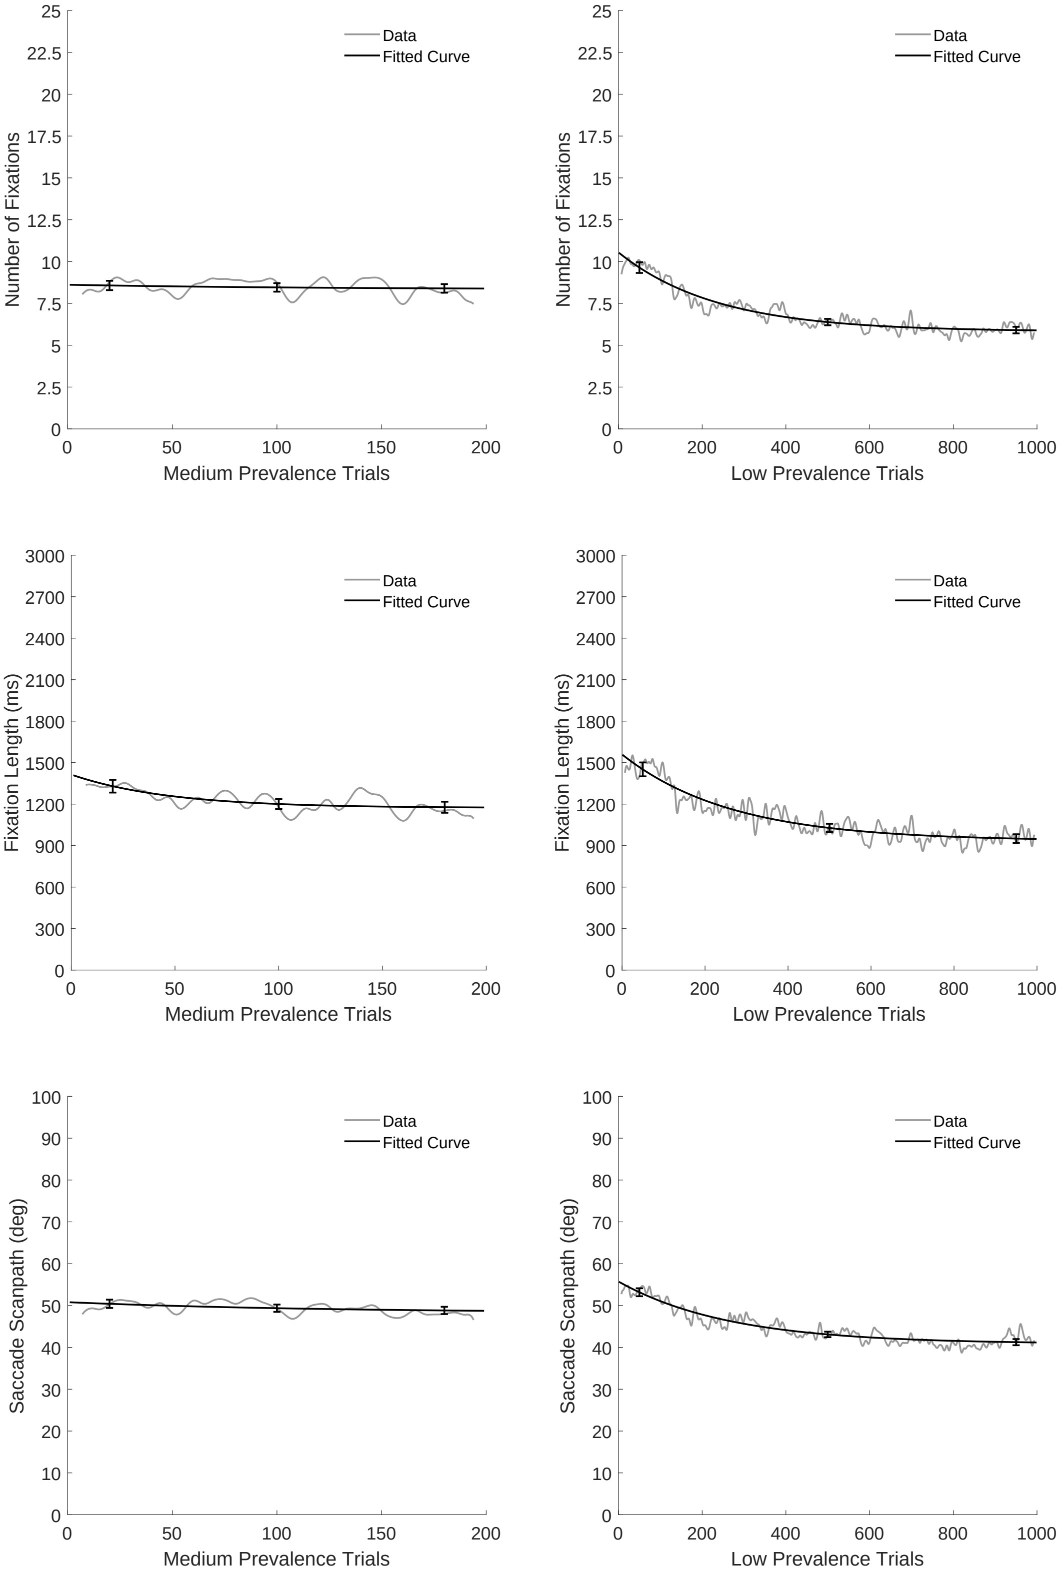

Figure 6.4: The mean eye-movement behaviour across the participants for each prevalence condition with an exponential decay curve fitted to each dataset. The error bars represent the standard error found in a sample consisting of fifteen trials either side of the presented error bar. Note the differing number of trials within each of the two prevalence conditions, with 200 trials in the medium-prevalence condition and 1000 trials in the low-prevalence condition.

Table 6.3: The eye-movement behaviour decay curve parameters for both the medium-prevalence and the low-prevalence conditions across Experiments. Mean percentages are followed by their parenthesised standard deviations and we remind the reader that the rate of the curves decay is represented as the number of trials to reach asymptote.

	Medium Prevalence			Low Prevalence		
	Decay	Rate	Asymptote	Decay	Rate	Asymptote
Experiment One						
Fixations (*n*)	2.63(5.02)	82.56(43.74)	6.44(1.78)	5.11(3.29)	324.03(165.92)	4.67(1.99)
Length (*ms*)	461.61(427.83)	63.44(39.00)	847.36(237.03)	795.07(499.40)	307.84(174.43)	567.80(280.12)
Scanpath (*deg*)	6.16(8.94)	88.97(39.84)	42.74(5.82)	18.39(11.43)	347.08(193.95)	35.82(7.88)
Experiment Two						
Fixations (*n*)	1.05(3.16)	76.15(39.28)	11.76(6.91)	11.89(8.74)	344.39(175.67)	7.47(5.28)
Length (*ms*)	585.97(667.35)	73.16(39.19)	1483(827.04)	1787.77(1174.02)	346.45(196.82)	931.69(569.48)
Scanpath (*deg*)	7.53(12.86)	77.64(40.11)	59.96(28.46)	36.94(28.62)	340.35(208.87)	45.21(12.57)
Experiment Three						
Fixations (*n*)	0.70(1.46)	83.51(42.99)	8.09(3.48)	4.90(4.47)	368.75(208.44)	5.40(2.24)
Length (*ms*)	367.44(312.61)	69.63(41.49)	1127.39(464.21)	744.89(709.22)	316.29(195.16)	856.68(334.66)
Scanpath (*deg*)	4.57(7.10)	76.47(42.83)	47.42(9.56)	17.56(13.73)	354.15(204.59)	39.47(8.74)

The fitting of the exponential decay function to the fixation data produced a RMS value of 2.99 fixations ($SD = 1.19\ fixations$) for the low-prevalence curve fit and a RMS value of 3.61 fixations ($SD = 1.88\ fixations$) for the medium-prevalence curve fit. As can be seen in Table 6.3 across the prevalence conditions, the magnitude of the curve decay for fixations was greater in the low-prevalence condition ($F_{(1,17)} = 14.05, p < .05, \eta_P^2 = 0.45$), the participants taking longer to reach a steady state of behaviour ($F_{(1,17)} = 31.91, p < .001, \eta_P^2 = 0.65$), and once the asymptote had been reached the level of behaviour was lower in the low-prevalence condition ($F_{(1,17)} = 20.41, p < .001, \eta_P^2 = 0.55$).

From a behavioural perspective, measuring behavioural change across the prevalence conditions when their time constants are different (i.e., the number of trials) may be considered poor practice. Therefore, we analysed the data over the first 50 trials and found that the mean number of fixations decreased by 0.10 of a fixation in the medium-prevalence condition and 0.91 of a fixation in the low-prevalence condition. When we pushed this timestamp out to include the full 200 trials in the medium-prevalence condition and the first 200 trials in the low-prevalence condition, we found that the mean number of fixations decreased by 0.23 of a fixation in the medium-prevalence condition and 2.70 fixations in the low-prevalence condition, and we know from the parameter equation that fixations continued to decay in the low-prevalence condition until trial 368. This analysis demonstrates the nature of operant processes and how these mechanisms influence behaviour over similar time constraints for each condition.

Across the Experiments, no significant differences were found between the experiments in the medium-prevalence condition for the magnitude of decay. In the low-prevalence condition, the magnitude of decay was the greatest for Experiment Two ($F_{(2,54)} = 8.27, p < .001, \eta_P^2 = 0.24$), with the posthoc analysis showing that the main significant difference occurred between Experiments Two and Three ($p = .002$) and a larger p-value occurred between Experiments One and Two ($p = .003$). No significant differences were found between the Experiments for the rate of this decay in both the medium- and low-prevalence conditions. Once the asymptote had been reached the levels of behaviour were the highest in Experiment Two for both the medium-prevalence condition ($F_{(2,54)} = 6.65, p < .05, \eta_P^2 = 0.20$) and the low-prevalence condition ($F_{(2,54)} = 3.20, p < .05, \eta_P^2 = 0.11$), with posthoc analyses showing that the main significant difference occurred between Experiments One and Two, and a larger p-value occurred between Experiments Two and Three. We hypothesized a reduction in search behaviour in the low-prevalence condition in Experiment Three due to the absence of target-present trials. However, it appears there is also a reduction in behaviour in the medium-prevalence condition in Experiment Three.

The fitting of the exponential decay curve to the fixation length data, produced a RMS value of 500.59ms ($SD = 214.00ms$) for the low-prevalence curve fit and a RMS value of 553.18ms ($SD = 252.53ms$) for the medium-prevalence curve fit.

Between the prevalence conditions, we observed that the behavioural decay was greater in the low-prevalence condition ($F_{(1,17)} = 5.78, p < .05, \eta_P^2 = 0.25$) and the participants took longer to reach a steady state of behaviour ($F_{(1,17)} = 27.11, p < .001, \eta_P^2 = 0.62$). Once the asymptote had been reached, the level of behaviour was lower in the low-prevalence condition ($F_{(1,17)} = 8.95, p < .05, \eta_P^2 = 0.35$) (see Figure 6.4).

When we analysed this data over the first 50 trials, surprisingly, we found that the mean length of fixations decreased by 158.56ms in the medium-prevalence condition and only 108.11ms in the low-prevalence condition. However, when we pushed this timestamp out to include the full 200 trials in the medium-prevalence condition and the first 200 trials in the low-prevalence condition, we found that the mean trial fixation length decreased by 237.90ms in the medium-prevalence condition and 331.55ms in the low-prevalence condition, and we know from the parameter equation that fixation lengths continued to decay in the low-prevalence condition until trial 316.

Across the experiments, no significant differences were found for the magnitude of decay in the medium-prevalence condition. In the low-prevalence condition, the greatest magnitude of decay occurring in Experiment Two ($F_{(2,54)} = 9.11, p < .001, \eta_P^2 = 0.26$), with the posthoc analysis showing that the main significant difference occurred between Experiments Two and Three ($p = .001$) and a larger p-value occurred between Experiments One and Two ($p = .002$). No significant differences were found between the Experiments for the rate of this decay in both prevalence conditions.

Once the participants had reached a steady state of behaviour, the amount of time fixating was the longest in Experiment Two for both the medium-prevalence condition ($F_{(2,54)} = 6.03, p < .05, \eta^2_P = 0.19$) and the low-prevalence condition ($F_{(2,54)} = 4.05, p < .05, \eta^2_P = 0.13$), with the posthoc analyses showing that the main significant difference occurred between Experiments One and Two and a larger p-value occurring between Experiments Two and Three.

The fitting of the exponential decay function to the saccadic scanpath trajectory data produced a RMS value of 12.09° ($SD = 3.27°$) for the low-prevalence curve fit and a RMS value of 14.00° ($SD = 4.97°$) for the medium-prevalence curve fit. Across the prevalence conditions, the magnitude of decay was greater in the low-prevalence condition ($F_{(1,17)} = 13.87, p < .05, \eta^2_P = 0.45$) and the participants taking longer to reach a steady state of behaviour ($F_{(1,17)} = 30.83, p < .001, \eta^2_P = 0.65$). Once the asymptote had been reached, the level of behaviour was lower in the low-prevalence condition ($F_{(1,17)} = 9.32, p < .05, \eta^2_P = 0.35$).

When we analysed this data over the first 50 trials, we found that the mean trial trajectory decreased by 0.86° in the medium-prevalence condition and 2.56° in the low-prevalence condition. When we pushed this timestamp out to include the full 200 trials in the medium-prevalence condition and the first 200 trials in the low-prevalence condition, we found that the mean trial trajectory decreased by 2.06° in the medium-prevalence condition and 7.87° in the low-prevalence condition, and we know from the parameter equation that scanpath trajectories continued to decay in the low-prevalence condition until trial 354.

Across the experiments, no significant differences were observed for the magnitude of the scanpath decay in the medium-prevalence condition. In the low-prevalence condition, the greatest behavioural decay occurring in Experiment Two ($F_{(2,54)} = 5.89, p < .05, \eta^2_P = 0.18$), with the posthoc analysis showing that the main significant difference occurred between Experiments Two and Three ($p = .012$) and a larger p-value occurred between Experiments One and Two ($p = .015$). No significant differences were found between the Experiments in the rate of this decay for both the medium- and low-prevalence conditions. Once the asymptote had been reached, the trajectories were the longest in Experiment Two for both the medium-prevalence condition ($F_{(2,54)} = 4.76, p < .05, \eta^2_P = 0.15$) and the low-prevalence condition ($F_{(2,54)} = 4.29, p < .05, \eta^2_P = 0.14$), with the posthoc analyses showing that the main significant difference occurred between

Experiments One and Two and a larger p-value occurring between Experiments Two and Three.

Note, a familiar pattern of eye-movement behaviour is emerging across the Experiments, in that there were no significant differences in the rate of decay, whereas the magnitude of decay increased significantly in Experiment Two and then decreased again in Experiment Three for the low-prevalence condition, and this pattern of behaviour was also observed for the curve asymptotes in both the medium- and low-prevalence conditions.

We hypothesized a reduction in search behaviour in the low-prevalence condition in Experiment Three due to the absence of target-present trials (i.e., extinction condition). However, the data shows a reduction in behaviour was also occurring in the medium-prevalence condition in Experiment Three when compared to Experiment Two despite there being no change in this condition across these two experiments. This finding may be a confounding consequence related to the order in which the participants experienced the conditions and demonstrated the power of operant conditioning in controlling search behaviour. Specifically, half of the participants experienced the low-prevalence condition first followed by the medium-prevalence condition and the remaining half experienced the medium-prevalence condition first followed by the low-prevalence condition. It may be suggested that the participants who experienced 1,000 trials without the presentation of a target-present trial in the low-prevalence condition (i.e., extinction condition) had a carry-over effect in the following medium-prevalence condition (see Figures 6.5. These figures show a carry-over effect from the condition the participant experienced first and disprove our hypothesis that there would be no change in eye-movement behaviour between Experiments Two and Three for the medium-prevalence conditions.

6.4 Stimulus Density Map

We were interested in the stimulus density map because we had observed in Experiments One and Two that eye-movement behaviour was being guided towards the centroid zone during their scene preview and we hypothesised that the density of the stimuli within this zone would make this behaviour a statistically efficient search strategy by allowing the participant to search this area with both overt and covert visual attention in an energy-efficient manner. Therefore, we hypothesise that this search strategy would be repeated in both conditions for Experiment Three all target-present trials had been removed from the low-prevalence condition.

The mean surface area of the centroid zone within the stimulus density map constituted 11.18% ($SD = 0.10\%$) of the stimulus matrix box area and represents the proportion of the stimulus matrix with the greatest stimulus density ($\geq$ 0.8 *density probability*). If the centroid zone sliced through a stimulus, this stimulus was counted as being within the centroid zone.

As a result, the mean number of stimuli within the centroid zone was 6.17 ($SD = 0.06$) of the 15 stimuli presented within the stimulus matrix box on each trial. This equated to 0.51 ($SD = 0.23$) fixations per centroid stimulus. When we examined the number of fixations per centroid stimulus across the Experiments we found that for Experiment One there were 0.44 ($SD = 0.13$) fixations per centroid stimulus, while in Experiment Two there was a greater amount of search behaviour within the centroid zone with 0.71 ($SD = 0.30$) fixations per centroid stimulus, and in Experiment Three this search behaviour within the centroid zone declined with 0.51 ($SD = 0.23$) fixations per centroid stimulus ($F_{(2,54)} = 7.33, p < .05, \eta_P^2 = 0.22$), with the posthoc analysis showing a main significant difference occurring between Experiments One and Two ($p = .002$) and a larger p-value occurring between Experiments Two and Three ($p = .028$). Consistent with our hypotheses, these observations show that the number of centroid fixations increases in Experiment Two when task difficulty was increased and then decreased in Experiment Three when target-present trials (i.e., rate of reinforcement) was reduced (extinction effect).

When comparing the miss errors in the medium-prevalence condition across the experiments, there was significantly less miss errors occurring when the target was presented within the centroid zone in Experiment One ($F_{(2,54)} = 8.28, p < .001, \eta_P^2 = 0.24$), with the posthoc analysis showing a main significant difference occurring between Experiments One and Two ($p = .001$) and a larger p-value occurring between Experiments One and Three ($p = .010$). This was also the case for the miss errors which occurred when the target was presented outside the centroid zone in the medium-prevalence condition ($F_{(2,54)} = 5.88, p < .05, \eta_P^2 = 0.18$), with the posthoc analysis showing a main significant difference occurring between Experiments One and Two ($p = .007$) and a larger p-value occurring between Experiments One and Three ($p = .029$). These results are consistent with an increase in task difficulty for Experiments Two and Three.

Table 6.4: The percentage of signal detection miss errors which occurred when the target stimulus was presented inside the centroid zone and the miss errors which occurred when the target stimulus was presented outside the centroid zone for the medium-prevalence condition and the low-prevalence condition in Experiments One, Two, and Three. Means are followed by their parenthesised standard deviations. Prev. is an abbreviation for prevalence.

	Medium Prev. Misses		Low Prev. Misses	
	In Centroid (%)	Out Centroid (%)	In Centroid (%)	Out Centroid (%)
Experiment One	6.58(7.90)	13.85(9.12)	42.11(29.06)	29.72(21.96)
Experiment Two	27.52(21.79)	25.86(13.69)	46.92(35.74)	45.72(25.38)
Experiment Three	23.66(17.78)	24.10(11.63)	NA	NA

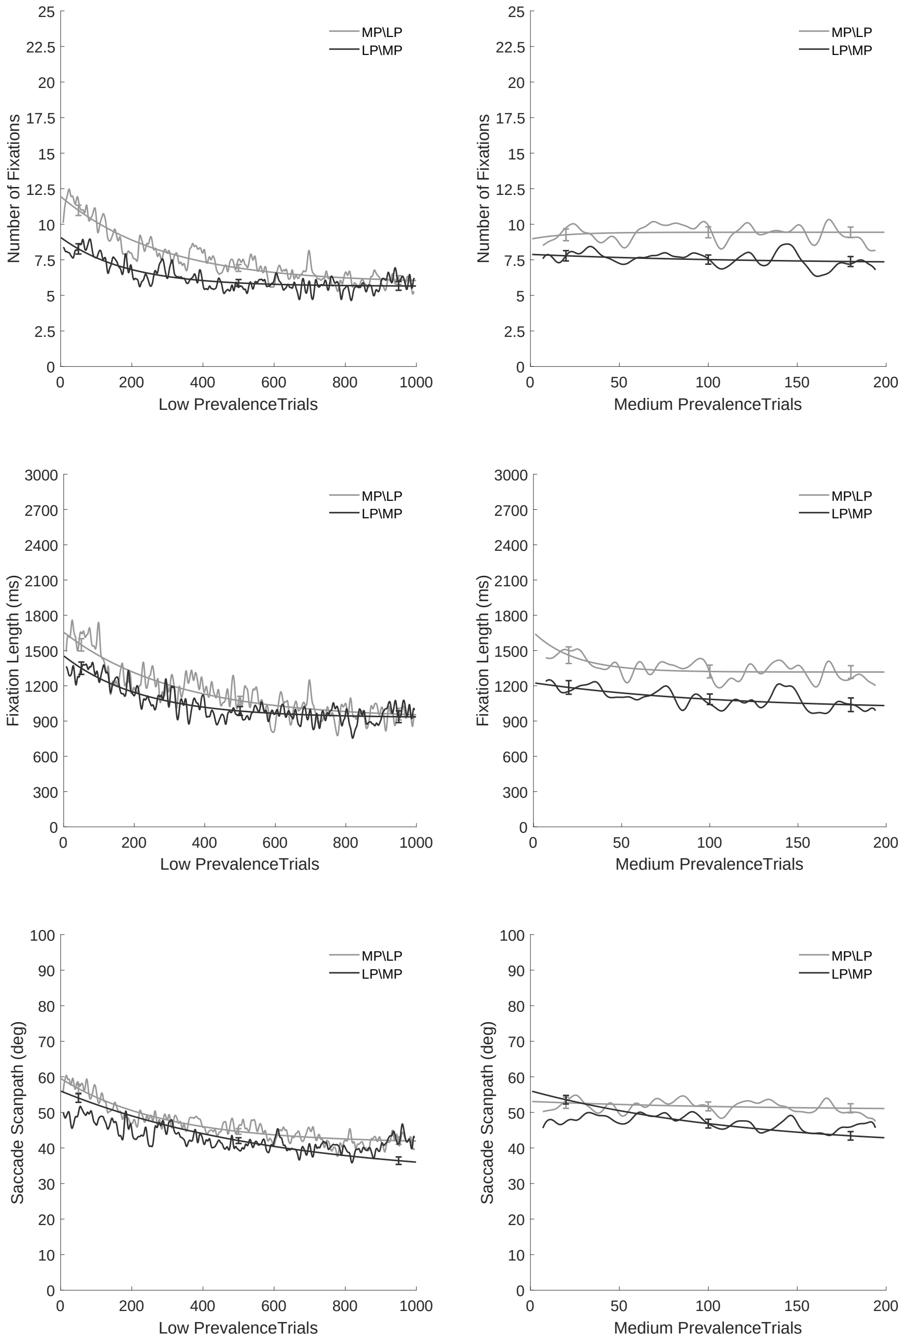

Figure 6.5: The eye-movement behaviour when the participants experienced the low-prevalence condition first before experiencing the medium-prevalence condition and the eye-movement behaviour when the participants experienced the medium-prevalence condition first before experiencing the low-prevalence condition. MP is an abbreviation for the medium-prevalence condition and LP is an abbreviation for the low-prevalence condition. MP/LP indicates that the participants experienced the medium-prevalence condition first before experiencing the low-prevalence condition and LP/MP indicates that the participants experienced the low-prevalence condition first before experiencing the medium-prevalence condition.

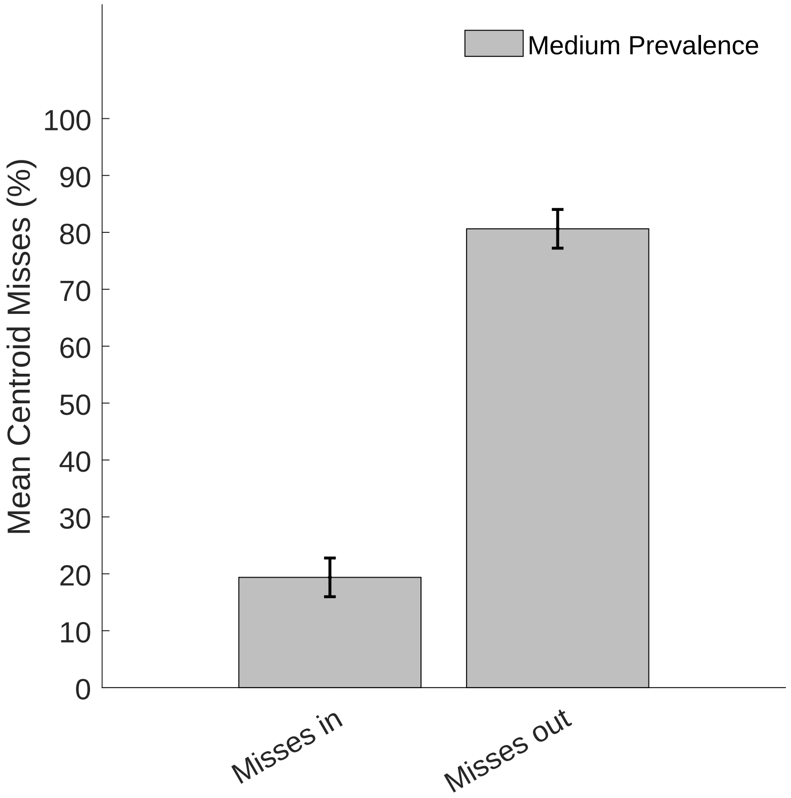

Figure 6.6: The miss errors which occurred inside the centroid zone and those which occurred outside the centroid zone in the medium-prevalence condition.

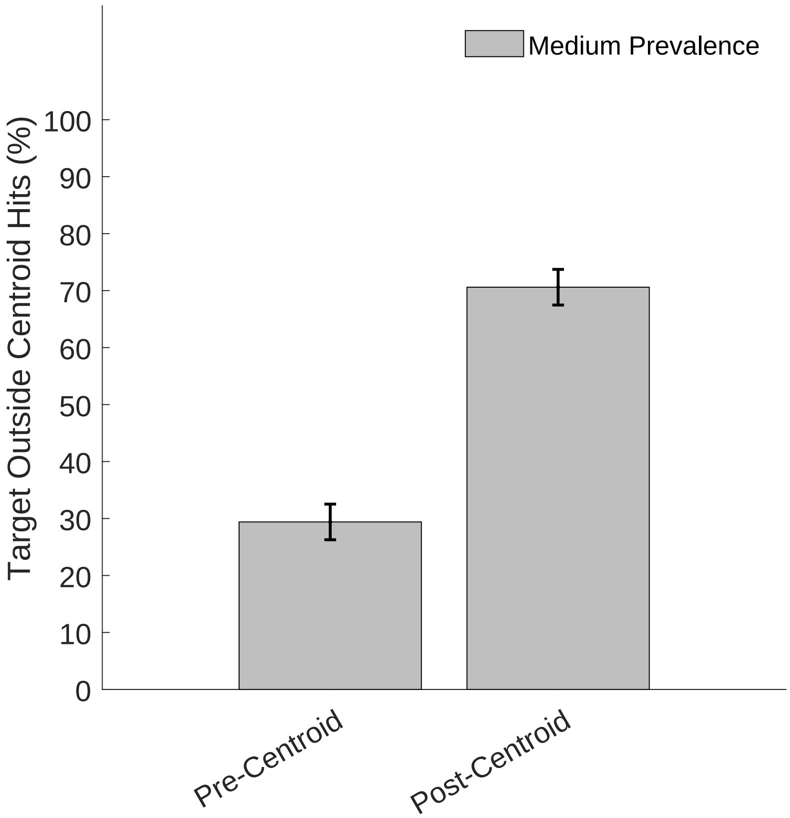

Figure 6.7: The proportion of hits which occurred before entering the centroid zone and the proportion of hits which occurred after leaving the centroid zone in the medium-prevalence condition.

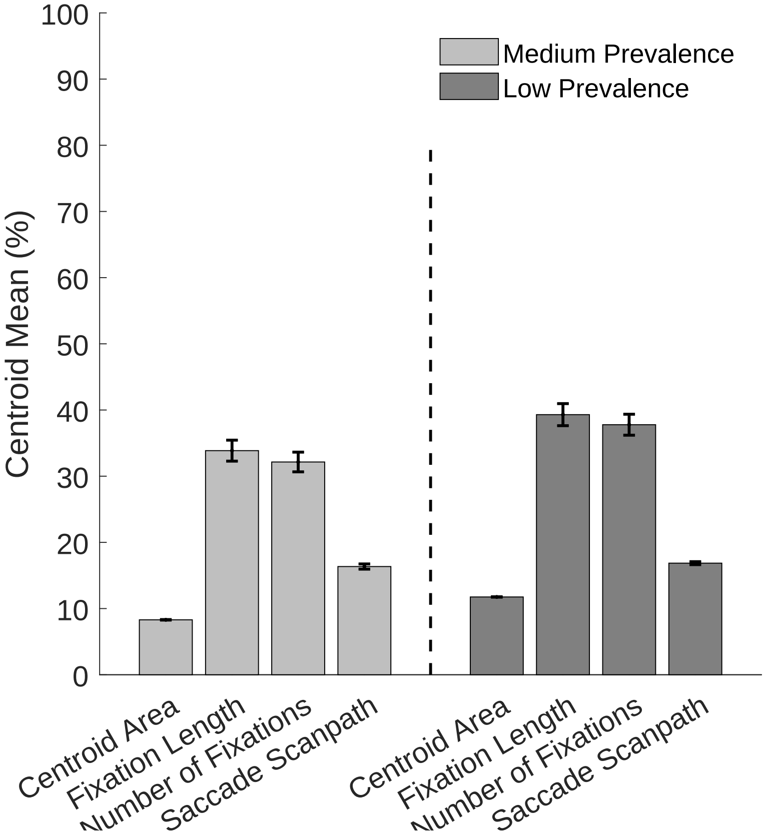

Figure 6.8: Centroid area and eye-movement behaviour within the centroid.

Table 6.5: This table shows the trial percentages for the number of fixations, the length of these fixations, and the saccadic scanpath found within the centroid zone for both the medium-prevalence and the low-prevalence conditions across Experiments. Mean percentages are followed by their parenthesised standard deviations.

	Medium Prevalence (%)			Low Prevalence (%)		
	Fixations	Length	Scanpath	Fixations	Length	Scanpath
Experiment One	31.68(6.86)	32.60(7.03)	17.00(0.95)	37.31(6.83)	38.22(6.89)	17.25(1.72)
Experiment Two	33.86(5.25)	35.35(5.25)	16.74(1.29)	36.33(5.50)	37.85(5.69)	17.39(1.12)
Experiment Three	32.13(6.32)	33.84(6.73)	16.35(1.67)	37.77(6.72)	39.29(7.07)	16.86(0.95)

As can be seen in Table 6.5, there was a greater percentage of eye-movement behaviour within the centroid zone compared to the surface area of the centroid zone within the stimulus matrix box ($M = 11.18\%, SD = 0.10\%$) for Experiment Three and this was also the case in the two previous Experiments (see Figure 6.8). There was a significantly greater percentage of fixations within the centroid zone in the low-prevalence condition compared to the medium-prevalence condition ($F_{(1,17)} = 20.20, p < .001, \eta_P^2 = 0.54$) in the current experiment which was also found in Experiment One and suggest that when reinforcement is lean (i.e., the low-prevalence condition) there is an increased amount of search behaviour within the centroid zone.

6.4.1 Centroid Eye-movement Decay

Experiments One and Two showed that the participants changed their pattern of eye-movement behaviour at different stages throughout a trial. Therefore, we analysed eye-movement behaviour before entering the centroid zone, within the centroid zone, and after leaving this zone in both the medium- and low-prevalence conditions. We fitted an exponential decay function to the data and report the curve parameters to provide a richer understanding of these patterns and the underlying mechanisms and strategies which may be involved when all target-present trials were removed from the low-prevalence condition. For the reader's convenience, we report the findings for the number of fixations and only present the data for the length of these fixations, and their scanpath trajectories in table format. Figures 6.9 and 6.10 show the medium- and low-prevalence conditions pre-, post-, and centroid eye-movement behaviour with an exponential decay function fitted to each dataset.

The fitting of the exponential decay function to the medium-prevalence condition data produced a RMS value of 3.09 fixations ($SD = 0.75\ fixations$) for the curve fitted to the pre-centroid data, a RMS value of 2.44 fixations ($SD = 1.03\ fixations$) for the curve fitted to the centroid data, and a RMS value of 3.15 fixations ($SD = 1.35\ fixations$) for the curve fitted to the post-centroid data. The fitting of the decay curves to the low-prevalence data produced a RMS value of 2.51 fixations ($SD = 0.57\ fixations$) for the curve fitted to the pre-centroid data, a RMS value of 2.20 fixations ($SD = 0.76\ fixations$) for the curve fitted to the centroid data, and a RMS value of 2.43 fixations ($SD = 0.92\ fixations$) for the curve fitted to the post-centroid data.

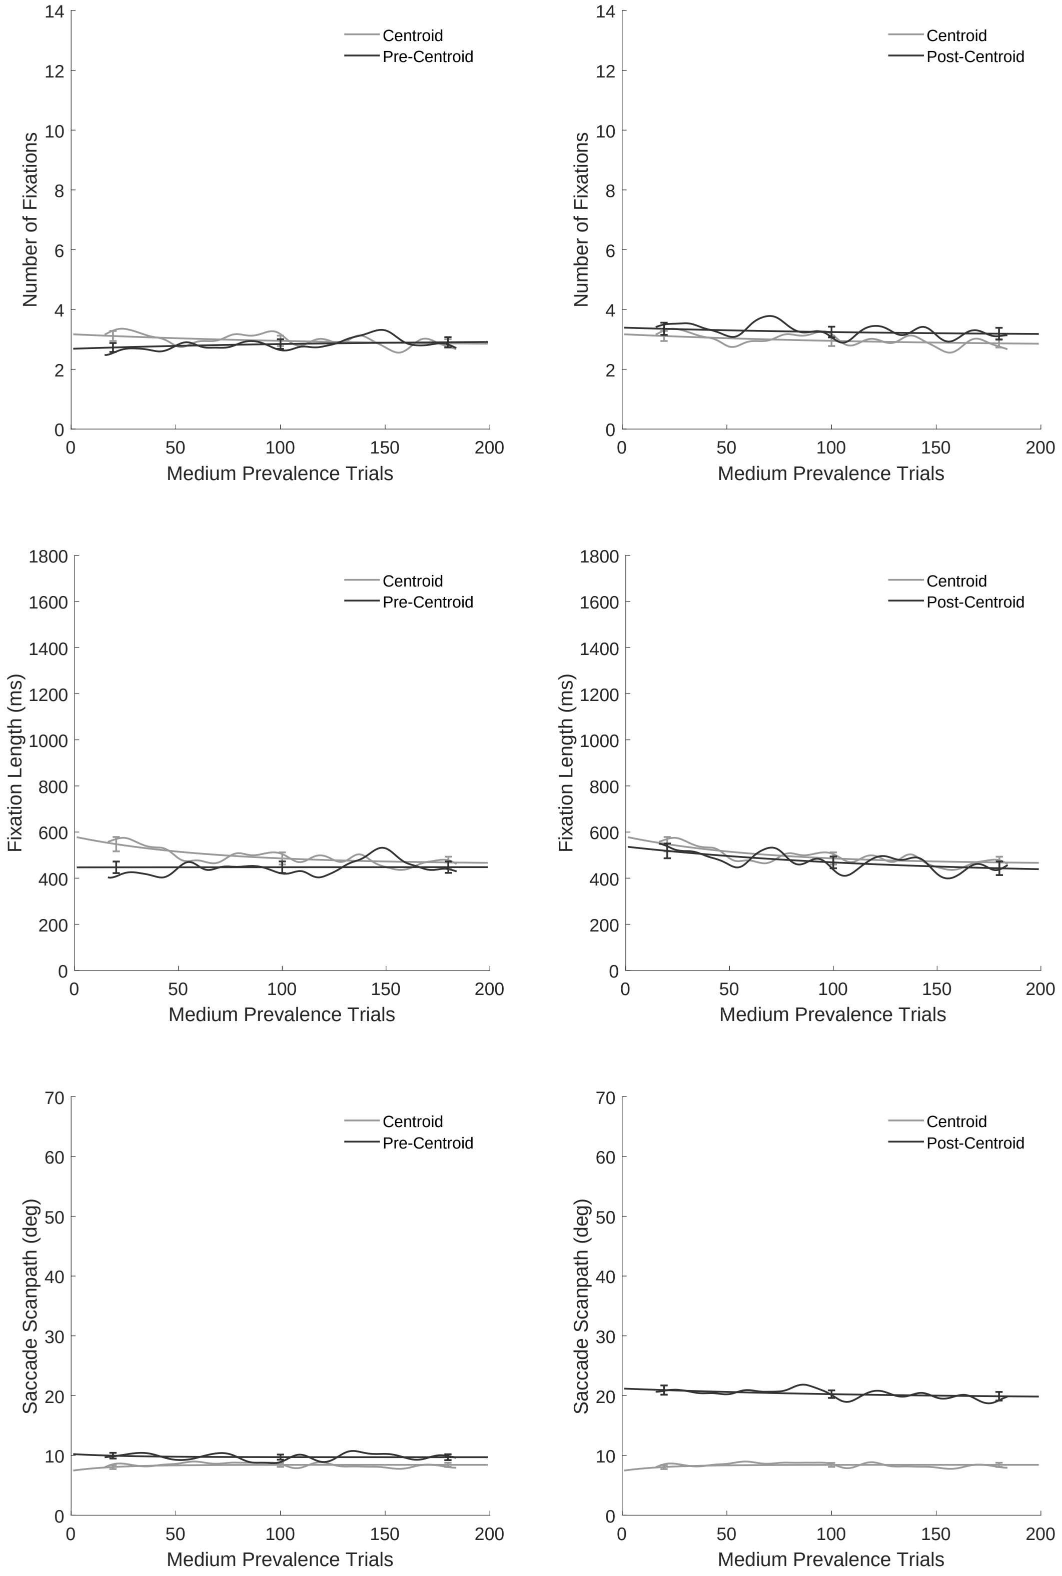

Figure 6.9: The medium-prevalence condition eye-movement behaviour before entering the centroid zone, the eye-movement behaviour within the centroid zone, and the eye-movement behaviour after leaving the centroid zone with an exponential decay curve fitted to each dataset in Experiment Three. The error bars represent the standard error found in a sample consisting of fifteen trials either side of the presented error bar.

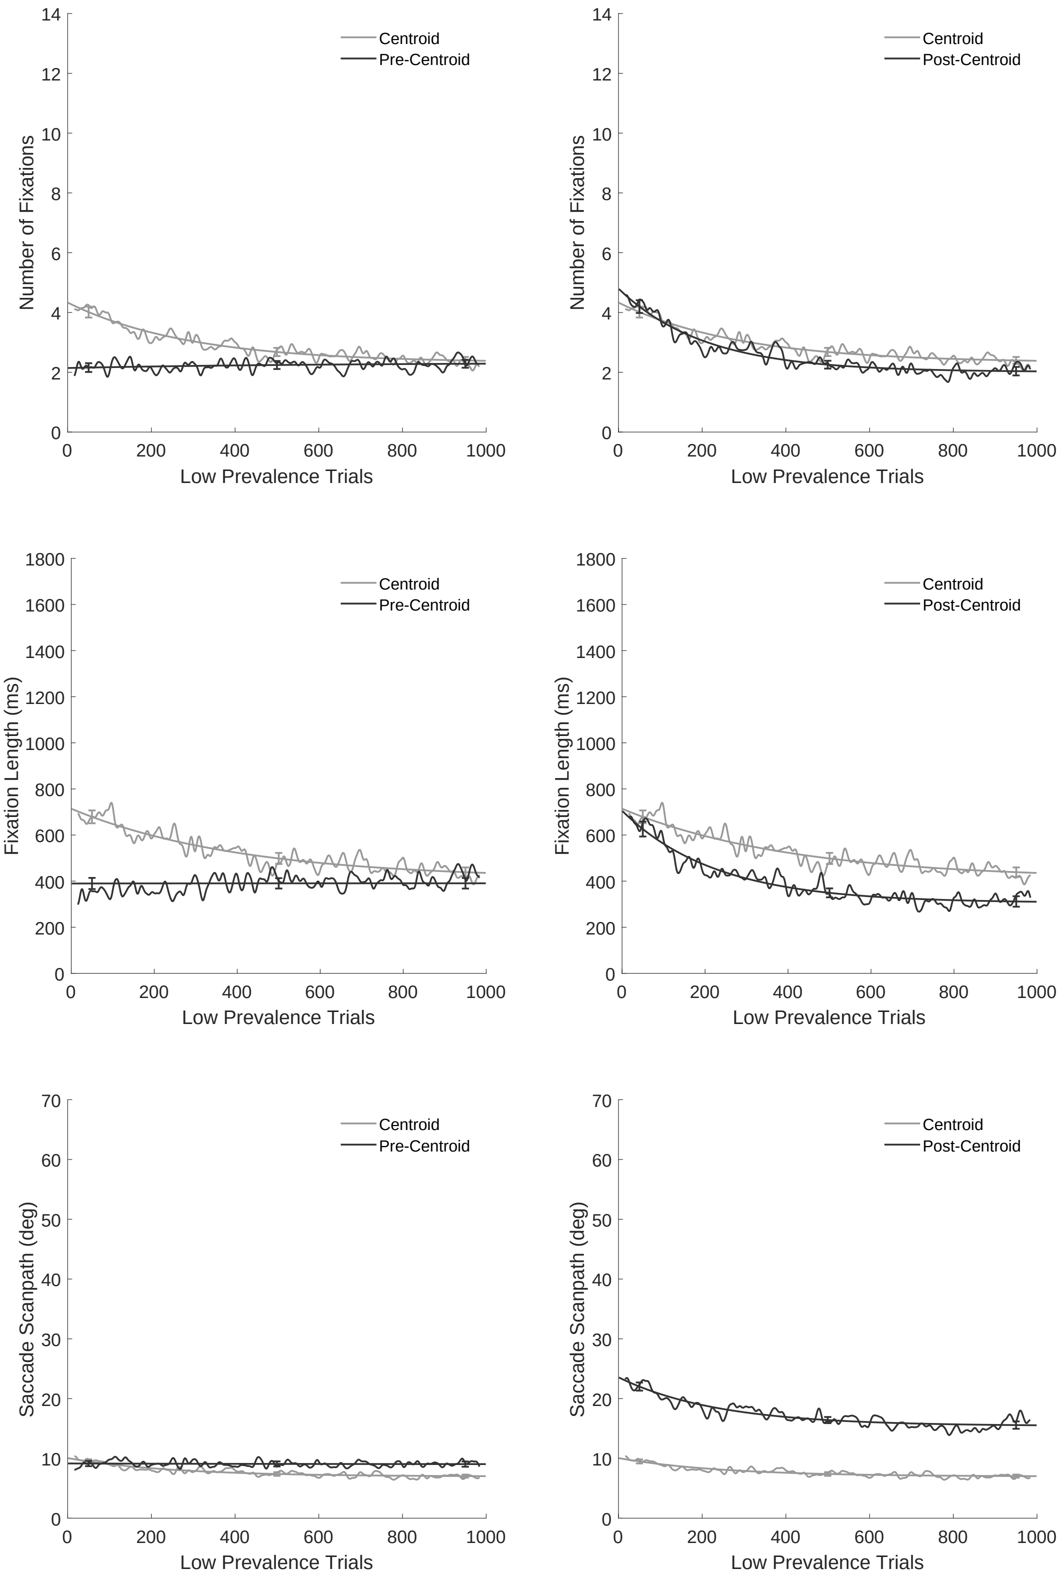

Figure 6.10: The low-prevalence eye-movement behaviour before entering the centroid zone, the eye-movement behaviour within the centroid zone, and the eye-movement behaviour after leaving the centroid zone with an exponential decay curve fitted to each dataset in Experiment Three. The error bars represent the standard error found in a sample consisting of fifteen trials either side of the presented error bar.

Table 6.6: Pre-centroid, centroid, and post-centroid exponential decay curve parameters across prevalence conditions for Experiment One. Means are followed by their parenthesised standard deviations.

	Experiment One					
	Medium Prevalence			Low Prevalence		
	Pre-	Centroid	Post-	Pre-	Centroid	Post-
Fixations						
Decay (*n*)	0.22(0.98)	1.01(2.01)	1.26(2.81)	0.37(0.81)	1.90(1.15)	3.01(1.89)
Rate (*trial*)	72.93(49.00)	67.32(47.55)	93.65(40.36)	442.92(209.43)	324.10(196.38)	308.55(172.23)
Asymptote (*n*)	2.72(1.17)	2.37(0.70)	2.37(0.81)	2.13(0.65)	1.97(0.75)	1.40(0.93)
Length						
Decay (*ms*)	91.47(92.50)	200.26(214.48)	236.92(211.36)	66.87(119.17)	317.04(209.86)	427.40(329.39)
Rate (*trial*)	82.05(40.52)	64.21(44.67)	75.49(46.52)	446.15(168.03)	255.71(168.45)	296.44(168.83)
Asymptote (*ms*)	368.15(154.87)	340.88(113.74)	304.32(91.76)	296.60(79.30)	286.93(108.34)	176.04(110.38)
Scanpath						
Decay (*deg*)	3.66(2.83)	0.43(2.07)	3.20(6.04)	1.16(1.03)	4.42(2.21)	9.33(6.66)
Rate (*trial*)	47.89(44.01)	85.11(39.34)	88.06(39.21)	512.57(179.20)	129.89(45.42)	422.74(213.30)
Asymptote (*deg*)	8.25(0.55)	7.76(1.19)	16.92(3.91)	8.06(0.51)	6.71(1.60)	12.29(4.64)

Table 6.7: Pre-centroid, centroid, and post-centroid exponential decay curve parameters across prevalence conditions for Experiment Two. Means are followed by their parenthesised standard deviations.

	Experiment Two					
	Medium Prevalence			Low Prevalence		
	Pre-	Centroid	Post-	Pre-	Centroid	Post-
Fixations						
Decay (*n*)	0.81(1.26)	0.19(1.43)	0.26(2.12)	0.33(0.96)	4.49(3.10)	6.77(6.06)
Rate (*trial*)	73.65(45.57)	74.39(43.68)	83.57(38.50)	416.09(223.78)	328.23(193.58)	351.84(175.10)
Asymptote (*n*)	2.79(0.98)	4.21(2.42)	5.50(4.06)	2.48(1.04)	2.96(1.61)	3.13(3.15)
Length						
Decay (*ms*)	198.75(178.14)	208.77(238.70)	294.90(471.31)	58.89(91.72)	732.93(481.11)	957.32(788.20)
Rate (*trial*)	65.83(44.21)	86.14(34.01)	75.34(38.67)	259.17(208.87)	277.29(195.68)	354.33(202.94)
Asymptote (*ms*)	376.98(106.38)	595.41(362.51)	700.35(551.61)	390.27(134.86)	457.49(219.81)	393.71(322.38)
Scanpath						
Decay (*deg*)	2.70(4.61)	0.66(3.15)	4.83(7.42)	1.45(1.69)	7.30(5.59)	20.63(16.23)
Rate (*trial*)	65.81(45.65)	80.89(42.56)	64.84(38.92)	501.31(194.55)	224.08(160.55)	342.95(196.73)
Asymptote (*deg*)	8.81(1.08)	10.15(3.56)	27.17(17.02)	8.58(1.01)	8.35(2.05)	17.88(7.35)

Table 6.8: Pre-centroid, centroid, and post-centroid exponential decay curve parameters across prevalence conditions for Experiment Three. Means are followed by their parenthesised standard deviations.

	Experiment Three					
	Medium Prevalence			Low Prevalence		
	Pre-	Centroid	Post-	Pre-	Centroid	Post-
Fixations						
Decay (*n*)	0.02(0.97)	0.30(1.18)	0.17(1.06)	-0.12(0.57)	1.99(1.80)	2.77(2.77)
Rate (*trial*)	80.82(42.74)	53.74(40.39)	71.34(45.22)	393.31(221.25)	355.50(184.80)	323.11(194.26)
Asymptote (*n*)	2.92(0.90)	2.86(1.51)	3.15(1.89)	2.32(0.69)	2.29(1.16)	1.81(1.05)
Length						
Decay (*ms*)	94.38(120.36)	158.76(197.83)	168.44(192.24)	9.19(21.68)	350.00(295.36)	463.40(439.72)
Rate (*trial*)	75.81(40.28)	60.57(39.38)	95.54(36.03)	366.91(176.98)	339.77(196.80)	323.00(195.76)
Asymptote (*ms*)	431.87(131.90)	462.78(254.33)	418.96(254.67)	394.64(137.83)	405.14(192.27)	261.42(127.71)
Scanpath						
Decay (*deg*)	1.37(2.77)	0.08(1.42)	3.08(5.52)	0.34(1.05)	3.65(2.69)	9.33(7.07)
Rate (*trial*)	92.06(41.04)	92.03(39.47)	87.13(37.86)	380.63(212.39)	305.82(207.63)	325.03(280.55)
Asymptote (*deg*)	9.48(1.55)	8.20(2.43)	19.16(6.03)	8.94(0.62)	6.66(1.60)	14.84(5.07)

As can be seen in Table 6.6, the magnitude of the curve decay for the number of trial fixations in the medium-prevalence condition was negligible before entering the centroid zone and the greatest fixational decay occurred within the centroid zone, no significant differences were found between the pre-, post-, and centroid means. In the low-prevalence condition, behavioural performance improved (negative decay) before entering the centroid zone and a significantly greater amount of decay occurred after leaving this zone ($F_{(2,16)} = 19.34, p < .001, \eta_P^2 = 0.53$), with the posthoc analysis showing that the main significant difference occurred between the pre-centroid and post-centroid means ($p < .001$) and a larger p-value occurred between the pre-centroid and centroid means ($p < .001$). These observations show the consistency of search behaviour throughout a trial in the medium-prevalence condition and also in the pre-centroid scene preview in both conditions while in the low-prevalence condition, the greatest effect of extinction occurred post-centroid.

Between the prevalence conditions, no significant difference was found for the pre-centroid fixational decay, which shows the consistency of search behaviour during their scene preview across the conditions. There was a significantly greater amount of decay in the low-prevalence within the centroid zone ($F_{(1,17)} = 8.43, p < .05, \eta_P^2 = 0.33$) and after leaving this zone ($F_{(1,17)} = 14.34, p < .05, \eta_P^2 = 0.46$), which shows that the decay of behaviour when reinforcement was withheld primary occurs within the centroid and after leaving this zone. Once the asymptote had been reached, the level of behaviour was significantly greater in the medium-prevalence condition pre-centroid ($F_{(1,17)} = 17.00, p < .05, \eta_P^2 = 0.36$), within the centroid ($F_{(1,17)} = 8.15, p < .05, \eta_P^2 = 0.32$), and post-centroid ($F_{(1,17)} = 14.21, p < .05, \eta_P^2 = 0.46$) which demonstrates how the high rate of reinforcement in this condition maintains search behaviour at a consistent level throughout a trial.

Across the experiments, once the asymptote had been reached, no significant pre-centroid differences were observed in both the medium- and low-prevalence conditions which shows the consistency of behaviour during the scene preview. Within the centroid zone, the greatest amount of behaviour occurred in Experiment Two when task difficulty was increased for both the medium-prevalence ($F_{(2,54)} = 5.95, p < .05, \eta_P^2 = 0.18$) and low-prevalence conditions ($F_{(2,54)} =$

$3.21, p < .05, \eta_P^2 = 0.11$), with the posthoc analyses showing that the main significant difference occurred between Experiments One and Two. After leaving the centroid zone, the greatest amount of behaviour again occurred in Experiment Two for both the medium-prevalence ($F_{(2,54)} = 7.26, p < .05, \eta_P^2 = 0.22$) and low-prevalence conditions ($F_{(2,54)} = 3.82, p < .05, \eta_P^2 = 0.13$), with the posthoc analyses showing that the main significant difference occurred between Experiments One and Two. Of interest was the observation that there was a reduction in centroid and post-centroid behaviour in the medium-prevalence condition for Experiment Three. These results show the weakening of behaviour when all target-present trials were removed in the low-prevalence condition in the current experiment and that this weakening effect also transposes over to the medium-prevalence post-centroid behaviour when the extinction condition is followed by the medium-prevalence condition.

Across the experiments, these results show that search behaviour is being consistently guided towards the centroid zone during their scene preview as a search strategy in both conditions. This suggests that examining the centroid zone is an efficient search strategy due to the number of stimulus items within this relatively small area.

It may also be suggested that the centroid zone is acting as a discriminative stimulus, signalling the availability of reinforcement, due to operant learning and this is guiding search behaviour towards the centroid zone. Specifically, during the transition period before the asymptote has been reached within each condition, the participant learns that if the target is not in the centroid zone, there is a high probability that the target will be outside this zone in the medium-prevalence condition due to the high rate of reinforcement (i.e., finding the target) in this condition and this reinforcement strengthens and maintains search behaviour at a consistent level throughout a trial.

Whereas, in the low-prevalence condition, search behaviour is being guided towards the centroid zone and if the target is not in this zone, there is a low probability that the target will be outside this zone, this weakens search behaviour, and we see a behavioural decay both within the centroid zone and after leaving this zone in the low-prevalence condition throughout a trial.

6.5 Discussion

6.5.1 Visual Search

We know that the prevalence of a target signal influences visual search behaviour dramatically in systematic ways (Ishibashi et al., 2012; Rich et al., 2008; J. D. Schwark et al., 2013; Wolfe et al., 2005, 2007). And we asked the question, “Is the schedule of signal presentation controlling eye-movement behaviour and if so, how?” We suggest that search behaviour may itself be influenced by the outcomes of search behaviour through the operant learning processes of reinforcement and extinction. Therefore, to try to exaggerate the eye-movement behaviour even more than in Experiments One and Two and to enrich our understanding of the operant control of search behaviour, we removed all target-present trials in the low-prevalence condition in Experiment Three. Specifically, the participants experienced a target probability of 0.50 in the medium-prevalence condition and in the low-prevalence condition they experienced a target probability of 0.00. All other aspects of Experiment Three were the same as Experiment Two. In line with the Reinforcement Strengthening Hypothesis, we predict that this schedule manipulation would result in fewer eye movements across the low-prevalence condition.

This schedule of signal presentation manipulation prevented miss errors occurring within the low-prevalence condition and, as such, we were unable to ascertain whether the Low Prevalence Effect occurred within Experiment Three. However, we again observed that the eye-movement behaviour metrics stabilized more rapidly and were significantly greater throughout the medium-prevalence condition compared to the low-prevalence condition. Furthermore, when compared to the medium-prevalence condition, there was a significantly greater amount of decay in eye-movement behaviour, the participants took a significantly longer amount of time to reach a steady state of behaviour and, once this state had been reached, the level of this behaviour was significantly lower in the low-prevalence condition for all three measures of eye-movement behaviour which supports the Reinforcement Strengthening Hypothesis (see Figures 6.4).

Furthermore, it was again observed that when compared to the correct rejection mean there was a significant increase in the number of fixations and their scanpath trajectories for the first correct rejection immediately following a correct target-present response in the medium-prevalence condition. As there were no target-present trials in the low-prevalence condition, it was not possible to observe this pattern of behaviour. These findings add further weight to the Reinforcement Strengthening Hypothesis, in that, the increase in eye-movement behaviour for the first post-hit correct rejection may be viewed as a manifestation of the local effects of obtained reinforcement and the strengthening of search behaviour when we find what we are looking for (see Figures 6.3).

Moreover, the abundance of target-present trials throughout the medium-prevalence condition, to which the correct response is a hit, and the subsequent post-hit increase in eye-movement behaviour for correct rejections would increase search behaviour for target-absent trials across the course of this condition and maintain search behaviour at a higher rate. Moreover, it is not surprising that target-absent responses become faster under low-prevalence conditions when there are very few, or in this case no, target-present trials to which identifying the target is reinforcing, resulting in a greater amount of decay in search behaviour and a lower rate of steady state behaviour under these circumstances. The greater decay and the lower rate of steady state behaviour in the low-prevalence condition may be considered reminiscent of an extinction curve when reinforcement has been withheld. In other words, the schedule of signal presentation is controlling eye-movement behaviour through the operant processes of reinforcement and extinction. The abundance of target-present trials in the medium-prevalence schedule and the reinforcing effect of finding the target would strengthen search behaviour within this condition resulting in an increased level of behaviour. Conversely, the absence of target-present trials in the low-prevalence schedule would weaken search behaviour due to the absence of obtained reinforcement (i.e., extinction) and we observed a markedly reduction in behaviour.

We know that explicit reinforcement, in the form of monetary compensation, profoundly increases our perceptual abilities, regarding the allocation of visual attention to stimulus features, oculomotor capture, and saccadic trajectories (Boehler et al., 2014; Braem et al., 2012; Hickey & van Zoest, 2012, 2013; Patai et al., 2012). In other words, reinforcement makes the target stimulus more salient (Theeuwes & Belopolsky, 2012), and this effect has been shown to persist for up to six months after the initial reinforcement-based contingency session (Anderson &

Yantis, 2013). Furthermore, studies employing inter-trial reinforcement priming procedures strongly suggest that reinforcement history independently influences the perception of stimulus features in the following trial, which cannot be fully explained in terms of top-down or bottom-up attentional processes (Della Libera & Chelazzi, 2009; Hickey et al., 2011, 2010), and these contingencies are questioning the dichotomy between top-down and bottom-up attentional control which prominent models of visual search assert (Awh, Belopolsky, & Theeuwes, 2012).

We also understand that reinforcer magnitude has a strong effect on behaviour, in that larger reinforcers affect behaviour more dramatically in much the same way as higher reinforcer frequency (M. Davison & Baum, 2003; Landon, Davison, & Elliffe, 2003; Mccarthy & Voss, 1995). For example, M. Davison and Baum (2003) showed that larger reinforcers produced larger and longer post-reinforcement preference pulses (i.e., a measure of the local effects of reinforcement) than did smaller reinforcers, sensitivity to reinforcer magnitude was considerably higher and was greater for magnitudes that differed more from one another, and visit durations following reinforcers, measured either as number of responses emitted or time spent responding before a changeover, were longer following larger reinforcers than following smaller reinforcers. These results are consistent with our understanding that the variables controlling choice behaviour, such as reinforcer magnitude, have both short- and long-term effects. Therefore, it may be suggested that, if finding the target stimulus is intrinsically reinforcing (Paeye et al., 2016), increasing the magnitude of this reinforcement by adding a secondary explicit reinforcer (i.e., monetary compensation) would strengthen visual behaviour to a greater extent, and this may be why we see such a profound impact of explicit reinforcement on visual search behaviour.

6.5.2 The Stimulus Density Map

We again observed that there was a significantly greater number of fixations, and the length of these fixations was longer within the centroid zone in the low-prevalence condition compared to the medium-prevalence condition. These results suggest that when reinforcement is lean (i.e., low-prevalence condition) there is an increase in search behaviour within the centroid zone. Furthermore, between the prevalence conditions, eye-movement behaviour was similar across the conditions during the scene preview before entering the centroid zone and the weakening of search behaviour (i.e., a greater amount of decay) predominantly occurred within the centroid zone and after leaving this zone during extinction conditions and this was reflected in the asymptote levels of behaviour.

We again observed an extinction burst, a temporary increase in the rate of an operant behaviour following the transition to extinction, within the centroid zone for fixations and their lengths in the low-prevalence condition. There was an initial increase in eye-movement behaviour when reinforcement was no longer forthcoming and then the dependent behaviour subsided (Fisher et al., 2023; Harris et al., 2007). These bursts of behaviour were first described by Sidman (1960) and subsequent research (Shahan, 2022) has shown that these bursts appear to be most common following schedules with a high rate of reinforcement, such as the medium-prevalence condition in our study. These bursts are not ubiquitous across all individuals, they occur for only a short time following the implementation of extinction (i.e., when reinforcement is no longer forthcoming), and they may be eliminated with the availability of alternative reinforcement. This strongly indicates that it is reinforcement, which is producing these bursts of behaviour, or rather, the transition from a high rate of reinforcement to a much lower rate of reinforcement and adds further weight to the Reinforcement Strengthening Hypothesis.

6.5.3 Experimental Comparisons

Importantly, it should be noted that the cohorts for each experiment were recruited over different timeframes and the sample was not randomly assigned to each experimental group. Therefore, firm conclusions cannot be drawn from the comparisons between the experiments. With this in mind, we now cautiously present some of the key findings between the experiments.

It was observed across all the experiments that, when compared to the medium-prevalence condition, there was a greater amount of behavioural decay, the participants took a longer amount of time to reach a steady state of search behaviour, and once this state had been reached, the level of this behaviour was lower in the low-prevalence condition for all three measures of eye-movement behaviour; however, there were some subtle differences between the experiments.

Once asymptote had been reached in the low-prevalence condition, for example, the pre-centroid fixation lengths and scanpath trajectories were significantly longer for Experiment Three when no target-present trials were presented indicating that there was an increase in information processing in the scene preview on the way to the centroid zone and that they were travelling to this zone more rapidly. Furthermore, fixations were fewer, the length of these fixations was shorter, and the scanpath trajectories were also shorter within the centroid zone and after leaving this zone in Experiment Three under extinction conditions compared to Experiment Two in the low-prevalence condition. Moreover, the mean number of fixations per centroid stimulus was significantly greater in the cohorts for Experiment Two and then significantly decreased in the cohorts for Experiment Three. No significant differences were found for the miss errors within the centroid zone and outside this zone in the medium-prevalence condition across Experiments Two and Three. These observations indicate that, once again, the schedule of signal presentation was controlling search behaviour.

Our results show that when there was a large disparity between the target stimulus and the distractor items in the control experiment (i.e., a simple singleton-feature task in Experiment One) search was efficient and required less effort as represented in the amounts of eye-movement behaviour and performance errors (see Figures 4.5). When this disparity was reduced in Experiment Two, visual search became more inefficient, requiring a significantly greater amount of eye-movement behaviour and an increase in performance errors (i.e., miss rate increased) and we observed a greater decay in behaviour in the low-prevalence condition due to the reduction in obtained reinforcement (see Figures 5.6). When we introduced an extinction condition in Experiment Three, search behaviour significantly decreased in the extinction condition (i.e., less visual effort) to a level that was similar to, or slightly above, that observed in Experiment One when less effort was required to find the target and the absence of reinforcement in this condition also had a carry-over effect in the medium-prevalence condition when this condition followed the extinction condition with a reduction in eye-movement behaviour (see Figures 6.4).

Interestingly, while it was observed that there were differences in the rate of behavioural decay between the prevalence conditions, with participants in the low-prevalence condition taking a significantly longer amount of time to reach the asymptote, no significant differences were observed between the experiments for the rate of this behavioural decay within each of the prevalence conditions. This consistency in the rate of decay in search behaviour within each prevalence condition demonstrates how the schedule of signal presentation (i.e., the rate of reinforcement) controls search behaviour and is represented in the robustness of the Low Prevalence Effect when the amount of obtained reinforcement is reduced.

Of particular interest was the finding that, when compared to the cohorts in the previous experiments, false alarm error rates were elevated, although not significantly, in both prevalence conditions for Experiment Three. This elevation may be attributed to three participants making idiosyncratic target-present decisions and reporting false alarm rates that were greater than two standard deviations above the mean. This finding was unexpected as miss and false alarm errors tend to trade-off as prevalence increases (D. M. Green et al., 1966). Wolfe and Van Wert (2010) investigated the effects of very high-prevalence conditions and found that at a 0.98 prevalence probability, miss errors became very rare, relative to a 0.5 prevalence probability, and false alarm errors increased dramatically from 18% at a 0.5 prevalence to 58% at a 0.98 prevalence. They also observed that target-absent latencies were subsequently longer under high-prevalence conditions, which would produce a negative decay curve under these circumstances.

This trade-off in errors under very high or very low prevalence conditions may reflect the history of signal presentation (i.e., history of reinforcement) experienced by the participant and their expectation that the target will, or will not, be presented under differing conditions. In other words, their behaviour may be driven by a “bottom-up” response to the experimental environment, and/or a “top-down” result of a change in strategy due to the participant’s idiosyncratic beliefs. In the context of the Multiple Decision Model (Wolfe & Van Wert, 2010), it has been suggested that the item identification threshold may be driven by the actual “bottom-up” experience of prevalence in the environment and the quitting threshold could be influenced by the “top-down” expectations (Horowitz, 2017). In an attempt to separate out expectations and experience, probability cueing studies (Ishibashi & Kita, 2014; Ishibashi et al., 2011) have shown that miss rates were not affected by the prevalence cues; however, the cues did affect the latencies on target-absent trials, with participants taking longer to respond "target-absent" following a high-prevalence cue. However, when Ishibashi and Kita (2014) reduced their low-prevalence condition down to a probability of 0.03 (compared to 0.06 in their previous studies) they observed significant effects on both accuracy and latencies, meaning that both thresholds in the Multiple Decision Model (Wolfe & Van Wert, 2010) may be influenced by both the history of signal presentation (i.e., rate of reinforcement) and the idiosyncratic beliefs that this history produces. On the other hand, our findings may simply reflect that search behaviour could become increasingly unstable in some participants when the schedule of signal presentation becomes too lean. For example, when they experienced 1,000 trials without the presentation of a target-present image, they were unable to update their visual target template during this time (Hout et al., 2015). In other words, the data became idiosyncratic for some participants when the schedule of reinforcement became too lean.

General Discussion

7.1 Visual Search

In visual search tasks, the likelihood that an observer will miss a target varies with the frequency in which it is presented across trials, with rare targets being missed more often than common ones. This phenomenon is referred to as the Low Prevalence Effect and can have serious consequences in real-world scenarios, such as airport security and medical screening, where targets can be extremely rare. Despite a large body of literature surrounding this phenomenon, we simply do not have a complete understanding of the underlying mechanisms behind this robust occurrence. ? (?) contend that the rate or intensity of search behaviour throughout an experiment is a measure which may be incorporated within most visual search studies, but is rarely reported, and is necessary for providing a richer understanding of visual search.

With this in mind, we suggest that the rate or intensity of search behaviour may itself be influenced by the outcomes of search behaviour through the processes of reinforcement and extinction. Specifically, if we consider that signal detection functions as reinforcement for visual search behaviour, we would predict that there would be more intense search behaviour when target signals are abundant, due to the strengthening effects of reinforcement when we find what we are looking for, and conversely, when target signals are rare, we would predict a reduction in search behaviour, which may result in an increase in miss errors and lead to the commonly observed Low Prevalence Effect. To answer this question, we incorporated a repeated-measures signal detection experimental design containing two prevalence conditions, namely a medium-prevalence condition and a low-prevalence condition, across three experiments and used eye-tracking technology to investigate eye-movements to ascertain the effects of the schedule of signal presentation on eye-movement behaviour, in an attempt to clarify the mechanisms underlying the Low Prevalence Effect and provide a richer understanding of visual search behaviour.

Unsurprisingly, we replicated the Low Prevalence Effect with a significantly greater number of miss errors occurring under the low-prevalence condition compared to the medium-prevalence condition, while no significant differences were observed between the prevalence conditions for the false alarm errors across Experiments One and Two. Furthermore, there was a significantly greater number of miss errors across both the prevalence conditions when the task difficulty was increased in Experiment Two compared to Experiment One, while no difference was found for miss errors in the medium-prevalence condition across Experiments Two and Three (note that no target-present trials were presented in the low-prevalence condition for Experiment Three). These findings were in line with those of ? (?), who in a series of experiments demonstrated that the prevalence effect was an extremely robust occurrence, and this effect has been attributed to a phenomenon whereby the target-absent response latencies become shorter under low-prevalence conditions (?, ?, ?).

Specifically, in visual search a well-established response pattern emerges with a prevalence probability of 0.5 or above, latencies for target-absent trials are typically longer than latencies for target-present trials. In a serial, self-terminating search task, the time to find a target may vary considerably (?, ?, ?). For example, the time to find a target may be very short, if by chance the first item the participant attends to is the target, or it may be much longer if the participant looks at all the wrong items before finally attending the target stimulus. However, the time to finally convince oneself that the target was not present is unlikely to be very short (?, ?). However, under low-prevalence conditions, this familiar pattern of behaviour reverses, with participants reporting that the target was absent faster than reporting that the target was present (?, ?, ?), which we also observed.

We observed that the participants settled into their search task more rapidly (i.e., search behaviour metrics stabilized) and maintained a consistently higher rate of eye-movement behaviour throughout the medium-prevalence condition compared to the low-prevalence condition. We remind the reader that eye-movement behaviour is an umbrella term incorporating all three of our eye-movement measures, namely the number of fixations, the length of these fixations (i.e., fixation amplitude), and their scanpath trajectories and that there were no significant differences observed across these measures unless otherwise noted.

Furthermore, when compared to the medium-prevalence condition, there was a greater amount of decay in eye-movement behaviour, the participants took a longer amount of time to reach a steady state of behaviour and, once this state had been reached, the level of this behaviour was lower in the low-prevalence condition for all three measures of eye-movement behaviour. This pattern of behaviour was observed across all three experiments; however, there were some subtle differences. For example, compared to Experiment One, there was a significant increase in eye-movement behaviour across both the prevalence conditions in Experiment Two when task difficulty was increased, despite this experiment also showing the greatest amount of decay in the low-prevalence condition. Furthermore, compared to Experiment Two, there was a significant decrease in eye-movement behaviour for Experiment Three in the low-prevalence condition when no target-present trials were presented.

From a cognitive perspective, this change in the response patterns under low-prevalence conditions has been attributed to a change in the decision criterion, in that the decision criterion becomes increasingly conservative as target prevalence decreases while sensitivity remains fairly constant across the prevalence manipulations (?, ?, ?). This mechanism was demonstrated in a study by ? (?) who varied the prevalence of signal presentation continuously in blocks of 50 trials, in a sinusoidal fashion across 1,000 trials, starting at 100% target prevalence, going all the way down to 0%, before coming back up to 100% again. They observed that sensitivity remained constant across the prevalence manipulation, while the decision criterion varied smoothly with prevalence, reaching a conservative peak at 0% prevalence and a liberal peak at approximately 100% prevalence. As one would expect, false alarm errors and miss errors traded off (?, ?), with high false alarm rates at high-prevalence and high miss rates at low-prevalence, while target-absent responses speeded up under low-prevalence conditions and slowed down under high-prevalence conditions.

These observations led ? (?) to formulate their Multiple Decision Model to explain the Low Prevalence Effect. According to this model, changes in the decision criterion manifest themselves across two distinct thresholds. When we move from a medium-prevalence condition to a low-prevalence condition, our internal identification threshold becomes more conservative, and the termination threshold is lowered. This results in targets becoming more likely to be missed, both because some targets will be mistakenly identified as a distractor, and because the participant will be more likely to terminate their search before the target has been identified. When we move from a low-prevalence condition to a medium-prevalence condition, miss errors become rarer because there would be more time spent searching the display items due to an increased termination threshold, resulting in an increased chance that the target will be found. They are also more likely to produce false alarm errors because of the reduced identification threshold, meaning that they require less evidence to identify a stimulus item as the target under these circumstances (?, ?). However, we believe that our thesis provides a more detailed explanation (i.e., mechanism) for the increased threshold concept behind the Low Prevalence Effect in the Multiple Decision Model.

A major difference between Wolfe and colleagues' interpretation of the Low Prevalence Effect and a behavioural interpretation is that a cognitive interpretation treats this shift in the decision criterion as an internal causal mechanism, whereas a behavioural interpretation treats this criterion shift as a behavioural outcome or effect.

From a behavioural standpoint, it may be suggested that finding an item that we are looking for may function as reinforcement for visual search behaviour due to a long and complex history of relevant conditioning. In this respect, the item itself may function as a reinforcer; for example, when we find our favourite coffee cup that we are looking for, we are then able to enjoy a rewarding cup of coffee, or there may be social reinforcement for finding the item which we are looking for; for example, when a young child finds an item of interest in a picture book and is then praised by a parent or teacher. In regards to eye-movement behaviour, Paeye et al. (2016) suggests that finding a target item is intrinsically reinforcing due to neurological dopaminergic feedback loops (?, ?), and that oculomotor responses are a set of operant responses controlled by environmental contingencies. Therefore, in the laboratory, a visual search yes-no signal detection task with differing prevalence conditions (e.g., a medium-prevalence condition and a low-prevalence condition) may be essentially viewed as two different schedules of reinforcement for search behaviour.

With this in mind, a behavioural interpretation of the mechanisms behind the Low Prevalence Effect would view this increase in eye-movement behaviour under medium-prevalence conditions as supportive of the Reinforcement Strengthening Hypothesis. This increase in behaviour may be viewed as a manifestation of the strengthening of behaviour due to the abundance of target-present trials in the medium-prevalence condition and the reinforcing effects when the participants find what they are looking for. As a result, this strengthening of search behaviour would result in a greater number of eye movements under these circumstances. Furthermore, the greater amount of decay and the lower rate of steady state behaviour in the low-prevalence condition may be considered reminiscent of an extinction curve when reinforcement has been withheld or a transition from a high rate of behaviour to a lower rate of behaviour as a function of a leaner schedule of reinforcement. Therefore, miss errors would be expected to increase under these circumstances. Given that the decision criterion (c) is simply a statistical measure of response bias calculated by averaging the z-score which corresponds to the hit rate and the z-score which corresponds to the false alarm rate (?, ?, ?), an increase in miss errors due to the weakening of search behaviour under low-prevalence conditions when reinforcement is lean would shift the decision criterion to a more conservative outcome.

Adding further weight to the reinforcement strengthening hypothesis was the observation that, when compared to the correct rejection means, there was an increase in eye-movement behaviour for the first correct rejection immediately following a correct target-present response, and this post-hit correct rejection increase in eye-movement behaviour was observed for all three measures of eye-movement behaviour in both the prevalence conditions across all three experiments. From a cognitive perspective, this post-hit correct rejection increase in search behaviour may be viewed as a manifestation of the extended termination threshold, due to a shift in the decision criterion. However, this increase in behaviour occurred in both the prevalence conditions independent of a shift in the decision criterion and, as such, these momentary shifts in the decision criterion on a trial-by-trial basis have not been accounted for within existing theory. From a behavioural perspective, this post-hit correct rejection increase in behaviour may be viewed as the momentary strengthening of search behaviour due to the reinforcing effects of finding what you are looking for and is reminiscent of a preference pulse (?, ?).

Specifically, preference pulses are a brief period of increased responding on the just-productive-alternative in a choice scenario. In our experiment, if searching was just reinforced with finding the target, a pulse of more search behaviour immediately after this trial would be expected and if the target was not present on this trial, there is an increased allocation of attention to the stimuli which have been associated with reinforcement (?, ?, ?, ?, ?, ?, ?, ?). This suggests that these "pulses" in search behaviour following a hit may be a result of the same underlying reinforcement mechanisms that underly similar preference pulses. In this respect, it has been suggested that these "pulses" represent the immediate effects of reinforcement on choice behaviour (?, ?, ?, ?).

Therefore, this increase in search behaviour immediately following a correct target-present response may reflect the local effects of reinforcement on search behaviour. Furthermore, the global impact of these "pulses" would influence search behaviour across each scheduled condition. Specifically, search behaviour would be maintained at a high rate throughout the medium-prevalence condition (i.e., search behaviour is strengthened), by altering response patterns and, as such, shifts the decision criterion due to the greater number of target-present trials in this condition and the local reinforcing effects of finding the target. Moreover, given that the correct response for a target-absent trial is a correct rejection and that there is an increase in search behaviour for correct rejections immediately following a hit, it is not surprising that target-absent responses become faster under low-prevalence conditions (i.e., search behaviour is weakened) when there are very few target-present trials to which the correct response is a hit. It should be noted, however, that these "pulses" and their associated increase in behaviour may not be solely responsible for the schedule-induced effects influencing overall search behaviour in each prevalence condition.

In the same manner, it has become abundantly clear that similar principles as those formulated in classic operant learning theory also apply to how reinforcement, in the form of monetary rewards, influences not only visual search behaviour but also executive functioning, cognitive control, and the allocation of visual attention resulting in a beneficial effect on perceptual processes through, some would say, the control of top-down attentional allocation (?, ?, ?, ?, ?, ?, ?, ?, ?). Furthermore, eye-movement behaviour, including oculomotor capture and saccadic trajectories, is beneficially influenced by inter-trial reinforcement priming (?, ?). Moreover, the reinforcing effects of inter-trial feedback have also been shown to influence search behaviour by rapidly shifting the decision criterion (?, ?).

If we consider that signal detection functions as reinforcement for visual search, we believe that operant theory accounts for the phenomenon of the Low Prevalence Effect and meshes with the two parameters within Wolfe and Van Wert's (2010) Multiple Decision Model, but rather than treating these parameters as hypothetical causes, it treats them as outcomes and provides an explanation for these behavioural shifts. Simply put, under a lean schedule of signal presentation (i.e., low reinforcement rate), search behaviour weakens resulting in decreased search times (i.e., termination threshold goes down) and an increase in miss errors (i.e., conservative decision criterion). As we enrich the schedule, search behaviour is strengthened resulting in longer search times (i.e., termination threshold goes up) and reduces miss errors (i.e., liberal decision criterion). In other words, the schedule of signal presentation controls eye-movement behaviour and influences search performance due to the relative rate of reinforcement and the reinforcing effects of finding what you are looking for.

Adding further weight to visual search and the theory of reinforcement learning, is that our findings align with Davison's discrimination model and a theory of attending (Nevin, Davison, & Shahan, 2005)which posits that attending is positively related to the rate of reinforcement correlated with stimuli. Higher reinforcement rates lead to increased probabilities of attending, impacting response accuracy and allocation, and a greater resistance to change in behaviour.Specifically, higher reinforcement rates (e.g., medium-prevalence condition) increases search behaviour (e.g., attending) increasing search accuracy (e.g., fewer miss errors) and produces a greater resistance to change in behaviour (e.g., a temporary increase in search behaviour following a medium- to low-prevalence change in prevalence conditions. In this respect, reinforcement learning may function as a mechanism to optimise visual search behaviour concerning eye-movement behaviour.

7.2 The Stimulus Density Map

We were also interested in the progression of eye-movement behaviour throughout a trial. To provide an understanding of this progression and the search strategies which may be involved, we introduced the concept of a 'centroid zone', which represented the portion of the display with the greatest density of stimulus items. We were interested in this centroid zone because we had observed, early in the study, that there was increased eye-movement behaviour in areas of the display which did not contain the target, and we wondered what was drawing visual attention away from the target stimulus. Furthermore, dot-cloud studies have demonstrated that human vision is highly efficient in estimating the centroid location of spatially scattered display items, and it has been suggested that this may be accomplished by neurons within the primary visual cortex and that each neural response within this receptive field is an increasing function of the number of items that fall within this field (?, ?, ?).

We observed that there was a much greater proportion of eye-movement behaviour within the centroid zone compared to the relative surface area of this zone within the stimulus matrix box in both the prevalence conditions across the three experiments. As a search strategy, the greater proportion of eye-movement behaviour within the centroid zone would make for an efficient statistical search strategy. The centroid zone only makes up an extremely small proportion of the search array and yet incorporates over one third of the array's stimuli and accounts for approximately one third of the fixations across a trial. The density of the stimuli within the centroid zone would therefore allow the participant to search this area with both overt and covert visual attention in an extremely efficient manner. As a result, if the target appeared within the centroid zone there was a high probability that the target would be identified; however, the only exception was the elevated miss rate within the centroid zone in the low-prevalence condition for Experiment One, which may imply that the participants were relying on this search strategy to a greater extent when the task became more difficult. This search strategy is consistent with the law of least effort (Zipf, 2016) which states that we seek out paths to goals that involve the least amount of effort as an energy preserving strategy (?, ?).

The brief exposure to a search array provides an initial understanding of this array (scene guidance) and is enough to guide the deployment of eye-movements once a search has begun (?, ?, ?). Importantly, this preview need not contain the target to be effective (?, ?), as knowledge about the target provides an independent source for the guidance of eye-movement behaviour (?, ?, ?). This "preview" is limited to the initial couple of fixations (?, ?) and once a search has begun, the guidance of eye-movement behaviour is updated, in that the content of the current fixation in an array or scene influences the location of the next fixation (?, ?). This initial "preview" often guides eye-movement behaviour towards the centre of an array or scene, termed the central-tendency phenomenon (?, ?). However, our data suggests that this was not the case in the current study. Instead, our data shows that eye movements were being guided or drawn towards the centroid zone (based on stimulus droplet density) rather than the centre of the display, regardless of whether this zone contained the target stimulus.

This pattern of search behaviour fits well with Zelinsky and Bisley's (2015) definition of a priority map. They contend that the function of a priority map is largely one of a filter for reducing the amount of incoming sensory information for processing, and define a priority map as a neural representation of topographic space for the selection of goal-relevant information to control and coordinate eye-movement behaviour, regarding the prioritisation of locations within this space, for the achievement of some behavioural goal, irrespective of the "bottom-up" or "top-down" factors. Furthermore, it may be said that the history of reinforcement (i.e., locating the target stimulus) may influence the prioritisation of these locations within this map (Paeye et al., 2016).

With this in mind, across the pre-, post-, and centroid data categories within each prevalence condition, the least amount of decay occurred before entering the centroid zone and a significantly greater amount of decay occurred after leaving this zone in the low-prevalence condition across all three experiments, which suggests that the weakening of search behaviour within the low-prevalence condition was primarily occurring after leaving the centroid zone. Once a steady state of eye-movement behaviour had been reached, which took a significantly longer amount of time in the low-prevalence condition across all the experiments, there was a significantly greater amount of post-centroid eye-movement behaviour in the medium-prevalence condition for all of our measures of eye-movement behaviour across all the experiments and this was also the case for the number of fixations and their scanpath trajectories within the centroid zone. Before entering the centroid zone, Experiments One and Three also showed a greater number of fixations in the medium-prevalence condition.

From a behavioural perspective, these observations suggest that the greater amount of eye-movement behaviour and the consistency of this behaviour across a trial in the medium-prevalence condition may be due to the strengthening of search behaviour, as a result of the abundance of target-present trials throughout this condition and the intrinsic reinforcing effects of finding what you are looking for. Conversely, the greater post-centroid behavioural decay in the low-prevalence condition may be representative of a weakening in search behaviour, resulting in a decrease in the termination threshold throughout this condition due to the lean schedule of signal presentation and the scarcity of obtained reinforcement, resulting in an increased rate of miss errors under these circumstances and over time this behaviour would shift the decision criterion. In this respect, a certain number of correct target-present responses (hits) outside the centroid zone may be needed to maintain post-centroid eye-movement behaviour and this number of correct responses may simply be too low in the low-prevalence condition to maintain this behaviour, so we see something which looks like an extinction effect.

Conceivably, the observation that eye-movement behaviour was being guided towards the centroid zone and the increased amount of behaviour within this zone compared to the relative surface area of the display suggests that the centroid zone may be working as an antecedent stimulus correlated with the availability of reinforcement (i.e., discriminative stimulus) due to the history of reinforcement. Specifically, under low-prevalence conditions, if the target stimulus was not observed within the centroid zone, the participant may learn through experience that reinforcement may be unavailable and terminate their search shortly after a thorough investigation of the centroid zone. Conversely, under medium-prevalence conditions, if the target stimulus was not observed within the centroid zone, the participant may learn that reinforcement may still be available outside this zone due to the abundance of target-present trials within this condition. This pattern of behaviour aligns with Wyckoff's (1952) definition of an observing response, which states that an observing response (OR) is an operant behaviour that produces discriminative stimuli (i.e., the centroid zone), which function as conditioned reinforcers that signal the availability of primary reinforcement, but do not affect the delivery of the primary reinforcer (i.e., finding the target stimulus). As such, these responses are a necessary condition for the establishment of stimulus control.

Our results suggest that if the target stimulus appeared close to the point of fixation upon presentation (i.e., functional visual field) and the participant was able to identify the target within their visual periphery during their visual preview (i.e., pre-centroid behaviour), they would saccade towards the target and terminate their search thereafter, before entering the centroid zone. If this was not the case, they would saccade towards the centroid zone and if the target was within this zone, there was a high probability that they would find the target. If the target was not within the centroid zone, they would begin to search outside this zone and in the low-prevalence condition, they would often terminate their search before they had identified the target due to the weakening of search behaviour as a result of the scarcity of obtained reinforcement, which would bring about an increase in miss errors within this condition. Importantly, the increase in eye-movement behaviour for a post-hit correct rejection would result in an increase in post-centroid search behaviour in the medium-prevalence condition. As a result, miss errors would be reduced in this condition which would produce a shift in the decision criterion over time.

Situating visual search in the ecological context of foraging may provide further insight. Ecologists have long known that foraging animals do not behave randomly. For example, predators feeding on cryptic prey tend to eat disproportionately more of the most common types of prey, a phenomenon termed frequency-dependent predation (?, ?). In the laboratory, when different targets are rewarded with uneven frequency, animals tend to match their ratio of behaviour proportionally to the frequency of reinforcement according to Herrnstein's (1961) matching law. Therefore, if searching the centroid zone had been reinforced with finding a target in the past, there would be an increased chance that they would search this area again in the future (Paeye et al., 2016). In other words, their search behaviour would be strengthened. This amounts to learning regularities or departures from randomness that may lead to expectations on the next trial, such as with uneven target probabilities. The priors for assuming that prediction may be possible must be high because completely random phenomena are rare within the natural world (?, ?). We suggest that similar operant learning mechanisms operate in visual search: prediction is obtained by weighting attention to a particular feature in an array proportionally to its previous encounters and the reinforcing effects that this feature has been associated with (?, ?, ?).

7.3 Limitations

One limitation which prevented conclusive cross-experimental comparisons was due to the experimental design, in that the cohorts for each experiment were recruited over different timeframes and the sample was not randomly assigned to each experimental group. Therefore, firm conclusions cannot be drawn from the comparisons between experiments. In hindsight this may have been an option; however, the sample would have to be larger to provide adequate statistical power for group comparisons. It may have been possible to include the same participants in all experiments, but there would have been greater attrition, and we may have experienced the confound of a practice effect. Furthermore, the specific methods of each experiment were determined by the findings from the previous experiment, which was the reason that a group design approach across all experimental conditions was not considered.

This study investigated eye-movement behaviour across three experimental conditions, namely two low-prevalence conditions (i.e., 0.02 & 0.00 target prevalence) and a medium-prevalence condition (i.e., 0.50 target prevalence); however, a high-prevalence condition was not incorporated. It may be said that including a greater range of conditions across the prevalence spectrum to incorporate a more parametric experimental design may provide valuable insights into the mechanisms and strategies behind visual search behaviour and add further weight to the Reinforcement Strengthening Hypothesis. This approach may be an exciting avenue for research in the future. As it may be suggested that, compared to a medium-prevalence condition, high-prevalence conditions may further strengthen search behaviour to a greater extent, and this may manifest itself as an increase in eye-movement behaviour for post-hit correct rejections which would explain why target-absent reaction times become even longer under high-prevalence conditions (Wolfe & Van Wert, 2010). We understand that every reinforcer counts in the strengthening of a behaviour (?, ?) and that a high rate of reinforcement in the past may influence decision processes in the present (?, ?). An understanding of these processes may shed light on the trade-off between miss and false alarm errors at either end of the prevalence spectrum. For example, we may rely more heavily on our history of reinforcement in situations of uncertainty and, as such, if a behaviour (i.e., a response) has been repeatedly reinforced in the past, this response would be more likely to occur again in the future in times of uncertainty, and this would help to explain the high rates of false alarm errors under high-prevalence conditions. Conversely, the same may be said for miss errors under low-prevalence conditions due to the scarcity of past reinforcement. In the same manner, operant theory may also shed light on the phenomenon of identification errors (?, ?) (i.e., fixated misses under low-prevalence conditions and non-fixated hits under high-prevalence conditions) under situations of uncertainty as our history of reinforcement would strengthen or weaken our response behaviours and may influence response decisions in times of uncertainty.

7.4 Conclusion

In summary, this study has made two main contributions towards a richer understanding of visual search behaviour. The first was the observation that there was an increase in eye-movement behaviour for the first correct rejection immediately following a correct target-present response. We suggest that this increase in search behaviour is due to the local reinforcing effects of finding what we are looking for which, in turn, may be due to a long and complex history of conditioning, and we suggest that this reinforcement mechanism influences the outcome of search behaviour across the schedule. In other words, the schedule of signal presentation is controlling eye-movement behaviour through the processes of reinforcement and extinction. The second contribution was the observation that search behaviour was being drawn or guided towards the densest cluster of stimulus items in the display as an energy-efficient search strategy, regardless of whether this centroid zone contained the target stimulus, and we propose that the centroid zone may act as a discriminative stimulus signalling the availability of reinforcement, which may provide valuable insight into our natural search behaviours. For example, when I am looking for my favourite coffee cup, I am often drawn towards the cluster of dirty coffee cups on the workstation kitchen bench.

It may be said that, for a stimulus to control behaviour, we must first attend to it (Dinsmoor, 1985). Unfortunately, visual attention is a limited resource (?, ?, ?). Therefore, when multiple stimuli are present, as is the case in the natural environment, we must divide our visual attention between the stimuli, attending to some and ignoring others. Understanding the processes which control such divided attention in visual space has implications for comprehending just how stimuli exert control over our visual behaviour. ? (?) contended that attention is the "controlling relation" (p.123) between stimuli and behaviour; that is, a stimulus will only exert control over our behaviour if it has been attended to.

Operant theory, which is primarily concerned with reinforcement learning mechanisms, may provide valuable insight into the processes behind visual search behaviour within real-world scenarios where performance errors can have serious consequences and may shed light on practices which may optimise visual search performance within the field. Eye-tracking technology and the study of eye-movement behaviour are unlocking the cognitive realms of visual search to behaviour analysis and is an exciting frontier for behavioural research in the future.

References

Adams, W. J. (2008). Frames of reference for the light-from-above prior in visual search and shape judgements. *Cognition*, *107*(1), 137–150.

Ahissar, M., & Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. *Nature*, *387*(6631), 401–406.

Anderson, B. A., Laurent, P. A., Yantis, S., & Smith, E. E. (2011, May). Value-Driven Attentional Capture.
(Publication Title: Proceedings Of The National Academy Of Sciences)

Anderson, B. A., & Yantis, S. (2012). Value-driven attentional and oculomotor capture during goal-directed, unconstrained viewing. *Attention, Perception, & Psychophysics*, *74*(8), 1644–1653. doi: 10.3758/s13414-012-0348-2

Anderson, B. A., & Yantis, S. (2013). Persistence of value-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(1), 6–9. doi: 10.1037/a0030860

Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in cognitive sciences*, *16*(8), 437–443.

Baker, C. (1960). Observing behavior in a vigilance task. *Science*, *132*(3428), 674–675.

Baum, W. M. (1974). On two types of deviation from the matching law: bias and undermatching 1. *Journal of the experimental analysis of behavior*, *22*(1), 231–242.

Beck, V. M., Hollingworth, A., & Luck, S. J. (2012). Simultaneous Control of Attention by Multiple Working Memory Representations. *Psychological Science*, *23*(8), 887–898. doi: 10.1177/0956797612439068

Berridge, K. C., & Robinson, T. E. (1998). What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? *Brain Research Reviews*, *28*(3), 309–369. doi: 10.1016/S0165-0173(98)00019-8

Biederman, I., Glass, A. L., & Stacy, E. W. (1973). Searching for objects in real-world scenes. *Journal of Experimental Psychology*, *97*(1), 22–27. doi: 10.1037/h0033776

Bindemann, M. (2010). Scene and screen center bias early eye movements in scene viewing. *Vision research*, *50*(23), 2577–2587.

Bindemann, M., Scheepers, C., Ferguson, H. J., & Burton, A. M. (2010). Face, body, and center of gravity mediate person detection in natural scenes. *Journal of Experimental Psychology: Human Perception and Performance*, *36*(6), 1477.

Birkelo, C. C., Chamberlain, W. E., Phelps, P. S., Schools, P. E., Zacks, D., & Yerushalmy, J. (1947). Tuberculosis case finding: A comparison of the effectiveness of various roentgenographic and photofluorographic methods. *Journal of the American Medical Association*, *133*(6), 359–366. (Publisher: American Medical Association)

Blair, W. C. (1958). Measurement of observing responses in human monitoring. *Science*, *128*(3318), 255–256.

Blough, D. S. (1966). The reinforcement of least-frequent interresponse times 1. *Journal of the Experimental Analysis of behavior*, *9*(5), 581–591.

Boehler, C. N., Schevernels, H., Hopf, J.-M., Stoppel, C. M., & Krebs, R. M. (2014). Reward prospect rapidly speeds up response inhibition via reactive control. *Cognitive, Affective, & Behavioral Neuroscience*, *14*(2), 593–609. doi: 10.3758/s13415-014-0251-5

Bond, A. B., & Kamil, A. C. (2002, February). Visual Predators Select For Crypticity And Polymorphism In Virtual Prey. *Nature*, *415*(6872), 609–613. Retrieved from `Https://Doi.Org/10.1038/415609a` doi: 10.1038/415609a

Botvinick, M., & Braver, T. (2015). Motivation and Cognitive Control: From Behavior to Neural Mechanism. *Annual Review of Psychology*, *66*(1), 83–113. doi: 10.1146/annurev-psych-010814-015044

Braem, S., Verguts, T., Roggeman, C., & Notebaert, W. (2012). Reward modulates adaptations to conflict. *Cognition*, *125*(2), 324–332. doi: 10.1016/j.cognition.2012.07.015

Breitmeyer, B. G., & Ganz, L. (1976). Implications of sustained and transient channels for theories of visual pattern masking, saccadic suppression, and information processing. *Psychological review*, *83*(1), 1. (Publisher: American Psychological Association)

Broadbent, D. E. (1957). A mechanical model for human attention and immediate memory. *Psychological Review*, *64*(3), 205–215. doi: 10.1037/h0047313

Broadbent, D. E., & Gregory, M. (1965). Effects of Noise and of Signal Rate upon Vigilance Analysed by Means of Decision Theory. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *7*(2), 155–162. doi: 10.1177/001872086500700207

Brockmole, J. R., Castelhano, M. S., & Henderson, J. M. (2006). Contextual cueing in naturalistic scenes: Global and local contexts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(4), 699–706. doi: 10.1037/0278-7393.32.4.699

Brockmole, J. R., & Henderson, J. M. (2006). Using real-world scenes as contextual cues for search. *Visual Cognition*, *13*(1), 99–108. doi: 10.1080/13506280500165188

Brown, A. C., & Grant, D. (1895). Abstract of a Lecture upon the relation between the Movements of the Eyes and the Movements of the Head. *The Journal of Laryngology & Otology*, *9*(9), 625–626. (Publisher: Cambridge University Press)

Bucker, B., Silvis, J. D., Donk, M., & Theeuwes, J. (2015). Reward modulates oculomotor competition between differently valued stimuli. *Vision Research*, *108*, 103–112. (Publisher: Elsevier)

Burr, D. C., Morrone, M. C., & Ross, J. (1994). Selective suppression of the magnocellular visual pathway during saccadic eye movements. *Nature*, *371*(6497), 511–513. doi: 10.1038/371511a0

Buswell, G. (1935). The place of the psychological laboratory in educational diagnosis. *Teachers College Record*, *36*(9), 155–167.

Castelhano, M., & Henderson, J. (2005). Incidental visual memory for objects in scenes. *Visual Cognition*, *12*(6), 1017–1040. doi: 10.1080/13506280444000634

Castelhano, M. S., & Heaven, C. (2011). Scene context influences without scene gist: Eye movements guided by spatial associations in visual search. *Psychonomic Bulletin & Review*, *18*(5), 890–896. doi: 10.3758/s13423-011-0107-8

Castelhano, M. S., & Krzyś, K. (2020). Rethinking space: A review of perception, attention, and memory in scene processing. *Annual Review of Vision Science*, *6*, 563–586. (Publisher: Annual Reviews)

Chelazzi, L., Duncan, J., Miller, E. K., & Desimone, R. (1998). Responses of Neurons in Inferior Temporal Cortex During Memory-Guided Visual Search. *Journal of Neurophysiology*, *80*(6), 2918–2940. doi: 10.1152/jn.1998.80.6.2918

Cheng, J. C. K., & Li, L. (2012). Effects of reference objects and extra-retinal information about pursuit eye movements on curvilinear path perception from retinal flow. *Journal of Vision*, *12*(3), 12–12. doi: 10.1167/12.3.12

Chun, M. M., & Jiang, Y. (1998). Contextual Cueing: Implicit Learning and Memory of Visual Context Guides Spatial Attention. *Cognitive Psychology*, *36*(1), 28–71. doi: 10.1006/cogp.1998.0681

Chun, M. M., & Nakayama, K. (2000). On the Functional Role of Implicit Visual Memory for the Adaptive Deployment of Attention Across Scenes. *Visual Cognition*, *7*(1), 65–81. doi: 10.1080/135062800394685

Colquhoun, W. (1961). The effect of ‘unwanted’signals on performance in a vigilance task. *Ergonomics*, *4*(1), 41–51. (Publisher: Taylor & Francis)

Curio, E., & Curio, E. (1976). Prey Recognition. *The Ethology of Predation*, 85–112. (Publisher: Springer)

Daniel, P. M., & Whitteridge, D. (1961). The representation of the visual field on the cerebral cortex in monkeys. *The Journal of Physiology*, *159*(2), 203–221. doi: 10.1113/jphysiol.1961.sp006803

Davison, M., & Baum, W. M. (2003). Every Reinforcer Counts: Reinforcer Magnitude And Local Preference. , *80*. (Publication Title: Journal Of The Experimental Analysis Of Behavior)

Davison, M. C., & Tustin, R. D. (1978, February). The Relation Between The Generalized Matching Law And Signal-Detection Theory. *Journal Of The Experimental Analysis Of Behavior*, *29*, 331–336.

Della Libera, C., & Chelazzi, L. (2009). Learning to Attend and to Ignore Is a Matter of Gains and Losses. *Psychological Science*, *20*(6), 778–784. doi:

10.1111/j.1467-9280.2009.02360.x

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual review of neuroscience*, *18*(1), 193–222. (Publisher: Annual Reviews 4139 El Camino Way, PO Box 10139, Palo Alto, CA 94303-0139, USA)

Diamond, M. R., Ross, J., & Morrone, M. C. (2000). Extraretinal Control of Saccadic Suppression. *The Journal of Neuroscience*, *20*(9), 3449–3455. doi: 10.1523/JNEUROSCI.20-09-03449.2000

Dinsmoor, J. A. (1985). The role of observing and attention in establishing stimulus control. *Journal of the Experimental Analysis of Behavior*, *43*(3), 365–381.

Donahoe, J. W., & Palmer, D. C. (1994). *Learning and complex behavior.* Allyn & Bacon.

Donk, M., & Theeuwes, J. (2003). Prioritizing selection of new elements: Bottom-up versus top-down control. *Perception & Psychophysics*, *65*(8), 1231–1242. doi: 10.3758/BF03194848

Dosher, B. A., Han, S., & Lu, Z.-L. (2010). Information-limited parallel processing in difficult heterogeneous covert visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *36*(5), 1128–1144. doi: 10.1037/a0020366

Dukewich, K. R., & Klein, R. M. (2009). Finding the target in search tasks using detection, localization, and identification responses. *Canadian Journal of Experimental Psychology / Revue canadienne de psychologie expérimentale*, *63*(1), 1–7. doi: 10.1037/a0012780

Duncan, J., & Humphreys, G. (1992). Beyond the search surface: visual search and attentional engagement.
(Publisher: American Psychological Association)

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological review*, *96*(3), 433. (Publisher: American Psychological Association)

Duncan, R. O., & Boynton, G. M. (2003). Cortical Magnification within Human Primary Visual Cortex Correlates with Acuity Thresholds. *Neuron*, *38*(4), 659–671. doi: 10.1016/S0896-6273(03)00265-4

Eckstein, M. P. (2011). Visual search: A retrospective. *Journal of Vision*, *11*(5), 14–14. doi: 10.1167/11.5.14

Edmunds, M., & Brunner, D. (1999). Ethology of defenses against predators. *The praying Mantids. The Johns Hopkins University Press, Baltimore, Maryland and London*, 276–299.

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

Egglin, T. K., & Feinstein, A. R. (1996). Context bias: a problem in diagnostic radiology. *Jama*, *276*(21), 1752–1755. (Publisher: American Medical Association)

Ehinger, K. A., Hidalgo-Sotelo, B., Torralba, A., & Oliva, A. (2009). Modelling search for people in 900 scenes: A combined source model of eye guidance. *Visual Cognition*, *17*(6), 945–978. doi: 10.1080/13506280902834720

Endler, J. A. (1988). Frequency-Dependent Predation, Crypsis And Aposematic Coloration. , *319*. Retrieved from `Https://Royalsocietypublishing.Org/` (Publication Title: Philosophical Transactions Of The Royal Society Of London. Series B, Biological Sciences)

Estes, W., & Taylor, H. (1964). A detection method and probabilistic models for assessing information processing from brief visual displays. *Proceedings of the National Academy of Sciences*, *52*(2), 446–454. (Publisher: National Acad Sciences)

Ethell, S. C., & Manning, D. (2001). Effects of prevalence on visual search and decision making in fracture detection. In *Medical Imaging 2001: Image Perception and Performance* (Vol. 4324, pp. 249–257). SPIE.

Evans, K. K., Birdwell, R. L., & Wolfe, J. M. (2013). If You Don't Find It Often, You Often Don't Find It: Why Some Cancers Are Missed in Breast Cancer Screening. *PLoS ONE*, *8*(5), e64366. doi: 10.1371/journal.pone.0064366

Evans, K. K., Tambouret, R. H., Evered, A., Wilbur, D. C., & Wolfe, J. M. (2011). Prevalence of Abnormalities Influences Cytologists' Error Rates in Screening for Cervical Cancer. *Archives of Pathology & Laboratory Medicine*, *135*(12), 1557–1560. doi: 10.5858/arpa.2010-0739-OA

Felton, M., & Lyon, D. O. (1966). The post-reinforcement pause 1. *Journal of* *the experimental analysis of behavior*, *9*(2), 131–134.

Ferster, C. B., & Skinner, B. F. (1957). Schedules of reinforcement. (Publisher: Appleton-Century-Crofts)

Findlay, J. M., & Gilchrist, I. D. (2003). *Active Vision.* Oxford University Press. doi: 10.1093/acprof:oso/9780198524793.001.0001

Fisher, W. W., Greer, B. D., Shahan, T. A., & Norris, H. M. (2023). Basic And Applied Research On Extinction Bursts. *Journal Of Applied Behavior Analysis*, 4–28. doi: 10.1002/Jaba.954

Fleck, M. S., & Mitroff, S. R. (2007). Rare Targets Are Rarely Missed in Correctable Search. *Psychological Science*, *18*(11), 943–947. doi: 10.1111/j.1467-9280.2007.02006.x

Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human perception and performance*, *18*(4), 1030. (Publisher: American Psychological Association)

Folk, C. L., Remington, R. W., & Wright, J. H. (1994). The structure of attentional control: contingent attentional capture by apparent motion, abrupt onset, and color. *Journal of Experimental Psychology: Human perception and performance*, *20*(2), 317. (Publisher: American Psychological Association)

Foulsham, T., Kingstone, A., & Underwood, G. (2008). Turning the world around: Patterns in saccade direction vary with picture orientation. *Vision Research*, *48*(17), 1777–1790. doi: 10.1016/j.visres.2008.05.018

Friedman-Hill, S., & Wolfe, J. M. (1995). Second-order parallel processing: visual search for the odd item in a subset. *Journal of Experimental Psychology: Human Perception and Performance*, *21*(3), 531. (Publisher: American Psychological Association)

Gegenfurtner, K. R. (2016). The Interaction Between Vision and Eye Movements. *Perception*, *45*(12), 1333–1357. doi: 10.1177/0301006616657097

Gegenfurtner, K. R., & Sperling, G. (1993). Information transfer in iconic memory experiments. *Journal of Experimental Psychology: Human Perception and Performance*, *19*(4), 845. (Publisher: American Psychological Association)

Geisler, W. S., & Chou, K.-L. (1995, January). Separation Of Low-Level And

High-Level Factors In Complex Tasks: Visual Search. *Psychological Review*, *102*(2), 356–378. Retrieved from `Https://Doi.Org/10.1037//0033-295x.102.2.356` doi: 10.1037/0033-295x.102.2.356

Godwin, H. J., Menneer, T., Cave, K. R., & Donnelly, N. (2010). Dual-target search for high and low prevalence X-ray threat targets. *Visual Cognition*, *18*(10), 1439–1463. doi: 10.1080/13506285.2010.500605

Godwin, H. J., Menneer, T., Cave, K. R., Helman, S., Way, R. L., & Donnelly, N. (2010). The impact of Relative Prevalence on dual-target search for threat items from airport X-ray screening. *Acta Psychologica*, *134*(1), 79–84. doi: 10.1016/j.actpsy.2009.12.009

Godwin, H. J., Menneer, T., Riggs, C. A., Cave, K. R., & Donnelly, N. (2015). Perceptual failures in the selection and identification of low-prevalence targets in relative prevalence visual search. *Attention, Perception, & Psychophysics*, *77*(1), 150–159. doi: 10.3758/s13414-014-0762-8

Godwin, H. J., Reichle, E. D., & Menneer, T. (2017). Modeling lag-2 revisits to understand trade-offs in mixed control of fixation termination during visual search. *Cognitive Science*, *41*(4), 996–1019. (Publisher: Wiley Online Library)

Gray, J., & Wedderburn, A. (1960). Shorter articles and notes grouping strategies with simultaneous stimuli. *Quarterly Journal of Experimental Psychology*, *12*(3), 180–184. (Publisher: SAGE Publications Sage UK: London, England)

Green, B. F., & Anderson, L. K. (1956). Color coding in a visual search task. *Journal of experimental psychology*, *51*(1), 19. (Publisher: American Psychological Association)

Green, D. M., Swets, J. A., & others. (1966). *Signal detection theory and psychophysics* (Vol. 1). Wiley New York.

Greene, M. R., & Oliva, A. (2009). The Briefest of Glances: The Time Course of Natural Scene Understanding. *Psychological Science*, *20*(4), 464–472. doi: 10.1111/j.1467-9280.2009.02316.x

Gur, D., Rockette, H. E., Armfield, D. R., Blachar, A., Bogan, J. K., Brancatelli, G., … others (2003). Prevalence effect in a laboratory environment. *Radiology*, *228*(1), 10–14. (Publisher: Radiological Society of North America)

Harris, A. C., Pentel, P. R., & Lesage, M. G. (2007, July). Prevalence, Magnitude, and Correlates of an Extinction Burst in Drug-Seeking Behavior in Rats Trained to Self-Administer Nicotine During Unlimited Access (23 H/Day) Sessions. *Psychopharmacology*, *194*(3), 395–402. doi: 10.1007/S00213-007-0848-2

He, P., & Kowler, E. (1989). The role of location probability in the programming of saccades: Implications for "center-of-gravity" tendencies. *Vision research*, *29*(9), 1165–1181. (Publisher: Elsevier)

Heinke, D., & Humphreys, G. W. (2003). Attention, spatial representation, and visual neglect: simulating emergent attention and spatial memory in the selective attention for identification model (saim). *Psychological review*, *110*(1), 29.

Heitz, R. P. (2014). The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Frontiers in Neuroscience*, *8*. doi: 10.3389/fnins.2014.00150

Henmon, V. A. C. (1911). The relation of the time of a judgment to its accuracy. *Psychological review*, *18*(3), 186. (Publisher: The Review Publishing Company)

Henschke, C. I., McCauley, D. I., Yankelevitz, D. F., Naidich, D. P., McGuinness, G., Miettinen, O. S., … others (1999). Early Lung Cancer Action Project: overall design and findings from baseline screening. *The Lancet*, *354*(9173), 99–105. (Publisher: Elsevier)

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the experimental analysis of behavior*, *4*(3), 267. (Publisher: Society for the Experimental Analysis of Behavior)

Hickey, C., Chelazzi, L., & Theeuwes, J. (2010). Reward Changes Salience in Human Vision via the Anterior Cingulate. *The Journal of Neuroscience*, *30*(33), 11096–11103. doi: 10.1523/JNEUROSCI.1026-10.2010

Hickey, C., Chelazzi, L., & Theeuwes, J. (2011). Reward has a residual impact on target selection in visual search, but not on the suppression of distractors. *Visual Cognition*, *19*(1), 117–128. doi: 10.1080/13506285.2010.503946

Hickey, C., Kaiser, D., & Peelen, M. V. (2015, January). Reward Guides Attention To Object Categories In Real-World Scenes. *Journal Of Experimental Psychology General*, *144*(2), 264–273. Retrieved from `Https://Doi.Org/10.1037/A0038627` doi: 10.1037/A0038627

Hickey, C., & van Zoest, W. (2012). Reward creates oculomotor salience. *Current Biology*, *22*(7), R219–R220. doi: 10.1016/j.cub.2012.02.007

Hickey, C., & van Zoest, W. (2013). Reward-associated stimuli capture the eyes in spite of strategic attentional set. *Vision Research*, *92*, 67–74. doi: 10.1016/j.visres.2013.09.008

Hillstrom, A. P. (2000). Repetition effects in visual search. *Perception & Psychophysics*, *62*(4), 800–817. doi: 10.3758/BF03206924

Hillstrom, A. P., Scholey, H., Liversedge, S. P., & Benson, V. (2012, January). The Effect of the First Glimpse at a Scene on Eye Movements During Search. *Psychonomic Bulletin & Review*, *19*(2), 204–210. doi: 10.3758/S13423-011-0205-7

Hogan, L. C., Bell, M., & Olson, R. (2009). A Preliminary Investigation of the Reinforcement Function of Signal Detections in Simulated Baggage Screening: Further Support for the Vigilance Reinforcement Hypothesis. *Journal of Organizational Behavior Management*, *29*(1), 6–18. doi: 10.1080/01608060802660116

Holland, J. G. (1958). Human Vigilance. *Science*, *128*(3315), 61–67. Retrieved from `Https://Www.Science.Org` (Publication Title: Science)

Hollingworth, A. (2009). Two forms of scene memory guide visual search: Memory for scene context and memory for the binding of target object to scene location. *Visual Cognition*, *17*(1), 273–291. doi: 10.1080/13506280802193367

Hollingworth, A., & Henderson, J. M. (2002). Accurate visual memory for previously attended objects in natural scenes. *Journal of Experimental Psychology: Human Perception and Performance*, *28*(1), 113. (Publisher: American Psychological Association)

Holt, E. B. (1903). Eye-movement and central anaesthesia. *The Psychological Review: Monograph Supplements*. (Publisher: The Macmillan Company)

Horowitz, T. S. (2017). Prevalence in visual search: From the clinic to the lab and back again. *Japanese Psychological Research*, *59*(2), 65–108.

Horowitz, T. S., & Wolfe, J. M. (1998). Visual search has no memory. *Nature*, *394*(6693), 575–577. doi: 10.1038/29068

Hout, M. C., Walenchok, S. C., Goldinger, S. D., & Wolfe, J. M. (2015). Failures of perception in the low-prevalence effect: Evidence from active and passive visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(4), 977–994. doi: 10.1037/xhp0000053

Huebner, G. M., & Gegenfurtner, K. R. (2010). Effects of Viewing Time, Fixations, and Viewing Strategies on Visual Memory for Briefly Presented Natural Objects. *Quarterly Journal of Experimental Psychology*, *63*(7), 1398–1413. doi: 10.1080/17470210903398139

Hulleman, J. (2010). Inhibitory tagging in visual search: Only in difficult search are items tagged individually. *Vision Research*, *50*(20), 2069–2079. doi: 10.1016/j.visres.2010.07.017

Hulleman, J., Lund, K., & Skarratt, P. A. (2019, July). Medium versus Difficult Visual Search: How a Quantitative Change in the Functional Visual Field Leads to a Qualitative Difference in Performance. *Attention Perception & Psychophysics*, *82*(1), 118–139. doi: 10.3758/S13414-019-01787-4

Hulleman, J., & Olivers, C. N. L. (2017). The Impending Demise Of The Item In Visual Search. Retrieved from `Https://Doi.Org/10.1017/S0140525x15002794` (Publication Title: Behavioral And Brain Sciences) doi: 10.1017/S0140525x15002794

Hwang, A. D., Wang, H.-C., & Pomplun, M. (2011). Semantic guidance of eye movements in real-world scenes. *Vision Research*, *51*(10), 1192–1205. doi: 10.1016/j.visres.2011.03.010

Ishibashi, K., & Kita, S. (2014, January). Probability Cueing Influences Miss Rate And Decision Criterion In Visual Searches. *I-Perception*, *5*(3), 170–175. doi: 10.1068/I0649rep

Ishibashi, K., Kita, S., & Wolfe, J. M. (2011, October). The Effects of Local Prevalence and Explicit Expectations on Search Termination Times. *Attention Perception & Psychophysics*, *74*(1), 115–123. doi: 10.3758/S13414-011-0225-4

Ishibashi, K., Kita, S., & Wolfe, J. M. (2012). The effects of local prevalence and explicit expectations on search termination times. *Attention, Perception, & Psychophysics*, *74*(1), 115–123. doi: 10.3758/s13414-011-0225-4

Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision research*, *40*(10-12), 1489–1506.

Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience*, *2*(3), 194–203. doi: 10.1038/35058500

Jia, D., Guo, H., Song, Z., Shi, L., Deng, X., Perc, M., & Wang, Z. (2021, August). Local And Global Stimuli In Reinforcement Learning. *New Journal Of Physics*, *23*(8), 083020. doi: 10.1088/1367-2630/Ac170a

Jiang, Y., Chun, M. M., & Olson, I. R. (2004). Perceptual grouping in change detection. *Perception & Psychophysics*, *66*, 446–453. (Publisher: Springer)

Jiang, Y., & Wagner, L. C. (2004). What is learned in spatial contextual cuing—configuration or individual locations? *Perception & Psychophysics*, *66*(3), 454–463. doi: 10.3758/BF03194893

Jiang, Y. V., Won, B.-Y., & Swallow, K. M. (2014). First saccadic eye movement reveals persistent attentional guidance by implicit learning. *Journal of Experimental Psychology: Human Perception and Performance*, *40*(3), 1161–1173. doi: 10.1037/a0035961

Kamil, A. C., Lindstrom, F., & Peters, J. (1985). The detection of cryptic prey by blue jays (Cyanocitta cristata) I: The effects of travel time. *Animal Behaviour*, *33*(4), 1068–1079. (Publisher: Elsevier)

Kamil, A. C., & Yoerg, S. I. (1982). Learning and foraging behavior. In *Ontogeny* (pp. 325–364). Springer.

Kamil, A. C., Yoerg, S. I., & Clements, K. C. (1988). Rules to leave by: Patch departure in foraging blue jays. *Animal Behaviour*, *36*(3), 843–853. (Publisher: Elsevier)

Karni, A., & Sagit, B. D. (1993). The Time Course Of Learning A Visual Skill. (Publication Title: Nature)

Kellman, P. J., & Garrigan, P. (2008, December). Perceptual Learning And Human Expertise. , *6*. Retrieved from `Https://Www.Elsevier.Com/Locate/Plrev` (Publication Title: Physics Of Life Reviews) doi: 10.1016/J.Plrev.2008.12.001

Kienzle, W., Franz, M. O., Scholkopf, B., & Wichmann, F. A. (2009). Center-surround patterns emerge as optimal predictors for human saccade targets. *Journal of Vision*, *9*(5), 7–7. doi: 10.1167/9.5.7

Kingsley, H. L. (1932). An Experimental Study of 'Search'. *The American Journal of Psychology*, *44*(2), 314. doi: 10.2307/1414831

Kiss, M., Driver, J., & Eimer, M. (2009). Reward Priority of Visual Target Singletons Modulates Event-Related Potential Signatures of Attentional Selection. *Psychological Science*, *20*(2), 245–251. doi: 10.1111/j.1467-9280.2009.02281.x

Klein, R. (1988). Inhibitory tagging system facilitates visual search. *Nature*, *334*(6181), 430–431. (Publisher: Nature Publishing Group UK London)

Klein, R. M., & MacInnes, W. J. (1999). Inhibition of Return is a Foraging Facilitator in Visual Search. *Psychological Science*, *10*(4), 346–352. doi: 10.1111/1467-9280.00166

Kono, H., Reid, P. J., & Kamil, A. C. (1998). The effect of background cuing on prey detection. *Animal Behaviour*, *56*(4), 963–972.

Kristjansson, A. (2008). "I know what you did on the last trial" - a selective review of research on priming in visual search. *Frontiers in Bioscience*, *13*(13), 1171. doi: 10.2741/2753

Kristjánsson, A. (2000). In Search of Remembrance: Evidence for Memory in Visual Search. *Psychological Science*, *11*(4), 328–332. doi: 10.1111/1467-9280.00265

Kruijne, W., & Meeter, M. (2015). The long and the short of priming in visual search. *Attention, Perception, & Psychophysics*, *77*(5), 1558–1573. doi: 10.3758/s13414-015-0860-2

Kunar, M. A., Rich, A. N., & Wolfe, J. M. (2010). Spatial and temporal separation fails to counteract the effects of low prevalence in visual search. *Visual cognition*, *18*(6), 881–897. (Publisher: Taylor & Francis)

Kundel, H. L. (1982). Disease prevalence and radiological decision making. *Investigative Radiology*, *17*(1), 107–109. (Publisher: LWW)

Kundel, H. L. (2000). Disease prevalence and the index of detectability: a survey of studies of lung cancer detection by chest radiography. In *Medical imaging 2000: Image perception and performance* (Vol. 3981, pp. 135–144). SPIE.

Kundel, H. L., Nodine, C. F., & Carmody, D. (1978). Visual scanning, pattern recognition and decision-making in pulmonary nodule detection. *Investigative radiology*, *13*(3), 175–181. (Publisher: LWW)

Lamy, D., Yashar, A., & Ruderman, L. (2013). Orientation search is mediated by distractor suppression: Evidence from priming of pop-out. *Vision Research*, *81*, 29–35. doi: 10.1016/j.visres.2013.01.008

Lamy, D., Zivony, A., & Yashar, A. (2011). The role of search difficulty in intertrial feature priming. *Vision Research*, *51*(19), 2099–2109. doi: 10.1016/j.visres.2011.07.010

Landon, J., Davison, M., & Elliffe, D. (2003). Concurrent Schedules: Reinforcer Magnitude Effects. , *79*. (Publication Title: Journal Of The Experimental Analysis Of Behavior)

Lau, J. S.-H., & Huang, L. (2010). The prevalence effect is determined by past experience, not future prospects. *Vision Research*, *50*(15), 1469–1474. doi: 10.1016/j.visres.2010.04.020

Lee, J., & Shomstein, S. (2013, June). The Differential Effects Of Reward On Space- And Object-Based Attentional Allocation. *Journal Of Neuroscience*, *33*(26), 10625–10633. Retrieved from `Https://Doi.Org/10.1523/Jneurosci.5575-12.2013` doi: 10.1523/Jneurosci.5575-12.2013

Le-Hoa Võ, M., & Wolfe, J. M. (2015). The role of memory for visual search in scenes. *Annals of the New York Academy of Sciences*, *1339*(1), 72–81. (Publisher: Wiley Online Library)

Lerman, D. C., Iwata, B. A., & Wallace, M. D. (1999). Side Effects Of Extinction: Prevalence Of Bursting And Aggression During The Treatment Of Self-Injurious Behavior. , *32*. (Publication Title: Journal Of Applied Behavior Analysis)

Levi, D. M., Klein, S. A., & Aitsebaomo, A. (1985). Vernier acuity, crowding and cortical magnification. *Vision Research*, *25*(7), 963–977. doi: 10.1016/0042-6989(85)90207-X

Levitt, H. (1971). Transformed up-down methods in psychoacoustics. *The Journal of the Acoustical society of America*, *49*(2B), 467–477.

Liston, D. B., Krukowski, A. E., & Stone, L. S. (2013). Saccade detection during smooth tracking. *Displays*, *34*(2), 171–176.

Lu, Z.-L., & Dosher, B. (2013). *Visual Psychophysics: From Laboratory to Theory.* Cambridge: MIT Press. (Publication Title: Visual Psychophysics)

Lusted, L. B., & Ledley, R. S. (1960). Mathematical models in medical diagnosis.

Academic Medicine, *35*(3), 214–222. (Publisher: LWW)

Machado, A. (1989). Operant conditioning of behavioral variability using a percentile reinforcement schedule. *Journal of the Experimental Analysis of behavior*, *52*(2), 155–166.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*(4), 276–298. doi: 10.1037/h0076778

Mackworth, J. F. (1970). *Vigilance and attention: A signal detection approach.* Penguin Books.

Macmillan, N. A., & Creelman, C. D. (2004). *Detection theory: A user's guide.* Psychology press.

Madelain, L., Champrenaut, L., & Chauvin, A. (2007). Control of Sensorimotor Variability by Consequences. *Journal of Neurophysiology*, *98*(4), 2255–2265. (Publication Title: J Neurophysiol) doi: 10.1152/Jn.01286.2006

Madelain, L., Harwood, M. R., Herman, J. P., & Wallman, J. (2010). Saccade adaptation is unhampered by distractors. *Journal of vision*, *10*(12), 29–29.

Madelain, L., Paeye, C., & Darcheville, J.-C. (2011). Operant control of human eye movements. *Behavioural processes*, *87*(1), 142–148.

Madelain, L., Paeye, C., & Wallman, J. (2008). Saccadic adaptation: reinforcement can drive motor adaptation. *Journal of Vision*, *8*(6), 919–919.

Madelain, L., Paeye, C., & Wallman, J. (2011, April). Modification of Saccadic Gain by Reinforcement. *Journal of neurophysiology*, *106*(1), 219–232. (Publication Title: J Neurophysiol) doi: 10.1152/Jn.01094.2009

Malcolm, G. L., & Henderson, J. M. (2010, January). Combining Top-down Processes to Guide Eye Movements During Real-World Scene Search. *Journal Of Vision*, *10*(2), 1–11. doi: 10.1167/10.2.4

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

Mannan, S., Ruddock, K. H., & Wooding, D. S. (1995). Automatic control of saccadic eye movements made in visual inspection of briefly presented 2-d images. *Spatial vision*, *9*(3), 363–386.

Mccarthy, D., & Davison, M. (1979, November). Signal Probability, Reinforcement And Signal Detection.

(Publication Title: Journal Of The Experimental Analysis Of Behavior)

Mccarthy, D., & Voss, P. (1995). Delayed matching-to-sample performance: Effects of relative reinforcer frequency and of signaled versus unsignaled reinforcer magnitudes. *Journal Of The Experimental Analysis Of Behavior*, *63*, 33–51.

Mccarthy, D., Voss, P., & Davison, M. (1994, September). Leaving Patches: Effects of Travel Requirements. , *62*(1), 89–108. (Publication Title: Journal Of The Experimental Analysis Of Behavior)

Menneer, T., Donnelly, N., Godwin, H. J., & Cave, K. R. (2010). High or low target prevalence increases the dual-target cost in visual search. *Journal of Experimental Psychology: Applied*, *16*(2), 133–144. doi: 10.1037/a0019569

Montagnini, A., & Chelazzi, L. (2005, December). The Urgency To Look: Prompt Saccades To The Benefit Of Perception. *Vision Research*, *45*(27), 3391–3401. doi: 10.1016/J.Visres.2005.07.013

Motter, B. C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *Journal of Neurophysiology*, *70*(3), 909–919. doi: 10.1152/jn.1993.70.3.909

Motter, B. C., & Simoni, D. A. (2008, October). Changes In The Functional Visual Field During Search With And Without Eye Movements. *Vision Research*, *48*(22), 2382–2393. Retrieved from `Https://Doi.Org/10.1016/J.Visres.2008.07.020` doi: 10.1016/J.Visres.2008.07.020

Müller, H. J., & Mühlenen, A. v. (2000). Probing distractor inhibition in visual search: inhibition of return. *Journal of Experimental Psychology: Human Perception and Performance*, *26*(5), 1591. (Publisher: American Psychological Association)

Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, *434*(7031), 387–391. (Publisher: Nature Publishing Group UK London)

Nakayama, K., & Martini, P. (2011). Situating visual search. *Vision Research*, *51*(13), 1526–1537. doi: 10.1016/j.visres.2010.09.003

Nakayama, K., & Silverman, G. H. (1986). Serial and parallel processing of visual feature conjunctions. *Nature*, *320*(6059), 264–265. doi: 10.1038/320264a0

Navalpakkam, V., & Itti, L. (2005). Modeling the influence of task on attention.

Vision Research, *45*(2), 205–231. doi: 10.1016/j.visres.2004.07.042

Neider, M. B., & Zelinsky, G. J. (2008). Exploring set size effects in scenes: Identifying the objects of search. *Visual Cognition*, *16*(1), 1–10. doi: 10.1080/13506280701381691

Neisser, U., & Beller, H. K. (1965). Searching through word lists. *British Journal of Psychology*, *56*(4), 349–358. (Publisher: Wiley Online Library)

Neuringer, A. (2002). Operant Variability: Evidence, Functions, And Theory. *Psychonomic Bulletin & Review*, *9–9*(4), 672–705. (Publisher: Psychonomic Society, Inc.)

Nevin, J. (1969, May). Signal Detection Theory and Operant Behavior. , *12*(3), 475. (Publication Title: Journal Of The Experimental Analysis Of Behavior)

Nevin, J. A., Jenkins, P., Whittaker, S., & Yarensky, P. (1982). Reinforcement contingencies and signal detection. *Journal of the Experimental Analysis of Behavior*, *37*(1), 65–79.

Nordfang, M., & Wolfe, J. M. (2014). Guided search for triple conjunctions. *Attention, Perception, & Psychophysics*, *76*(6), 1535–1559. doi: 10.3758/s13414-014-0715-2

Nothdurft, H.-C. (1999). Focal attention in visual search. *Vision Research*, *39*(14), 2305–2310. doi: 10.1016/S0042-6989(99)00006-1

Oh, S.-H., & Kim, M.-S. (2004). The role of spatial working memory in visual search efficiency. *Psychonomic Bulletin & Review*, *11*(2), 275–281. doi: 10.3758/BF03196570

Oliva, A., & Torralba, A. (2001). Modeling the shape of the scene: A holistic representation of the spatial envelope. *International journal of computer vision*, *42*, 145–175. (Publisher: Springer)

Paeye, C., Darcheville, J.-C., & Madelain, L. (2007). Shaping induces long-term enhancements in pursuit-like-smooth movements under image-stabilization condition. In *European conference on visual perception.*

Paeye, C., & Madelain, L. (2011). Reinforcing saccadic amplitude variability. *Journal of the Experimental Analysis of Behavior*, *95*(2), 149–162.

Paeye, C., & Madelain, L. (2014). Reinforcing Saccadic Amplitude Variability In A Visual Search Task.

doi: 10.1167/14.13.20

Paeye, C., Schütz, A. C., & Gegenfurtner, K. R. (2016). Visual reinforcement shapes eye movements in visual search. *Journal of vision*, *16*(10), 15–15.

Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, *5*(5), 1. doi: 10.1167/5.5.1

Parkhurst, D., Law, K., & Niebur, E. (2002). Modeling the role of salience in the allocation of overt visual attention. *Vision research*, *42*(1), 107–123.

Parkhurst, D., & Niebur, E. (2003). Scene content selected by active vision. *Spatial vision*, *16*(2), 125–154.

Pashler, H. (1987). Detecting conjunctions of color and form: Reassessing the serial search hypothesis. *Perception & Psychophysics*, *41*(3), 191–201. doi: 10.3758/BF03208218

Pashler, H., Johnston, J. C., & Ruthruff, E. (2001). Attention and performance. *Annual review of psychology*, *52*(1), 629–651. (Publisher: Annual Reviews 4139 El Camino Way, PO Box 10139, Palo Alto, CA 94303-0139, USA)

Patai, E. Z., Doallo, S., & Nobre, A. C. (2012, December). Long-Term Memories Bias Sensitivity And Target Selection In Complex Scenes. *Journal Of Cognitive Neuroscience*, *24*(12), 2281–2291. Retrieved from `Https://Doi.Org/10.1162/JocnA00294` doi: 10.1162/JocnA00294

Pavlov, I. P. (1927). *Conditioned reflexes: An investigation of the physiological activities of the cerebral cortex.*

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological review*, *87*(6), 532. (Publisher: American Psychological Association)

Pierce, W. D., & Cheney, C. D. (2017). *Behavior analysis and learning: A biobehavioral approach*. Routledge.

Pietrewicz, A. T., & Kamil, A. C. (1977). Visual detection of cryptic prey by blue jays (cyanocitta cristata). *Science*, *195*(4278), 580–582.

Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of experimental psychology: General*, *109*(2), 160. (Publisher: American Psychological Association)

Pryor, K. W., Haag, R., & O'Reilly, J. (1969). The creative porpoise: Training for novel behavior 1. *Journal of the Experimental Analysis of behavior*, *12*(4), 653–661.

Pushparaj, A., Pryslawsky, Y., Forget, B., Yan, Y., & Le Foll, B. (2012, October). Extinction Bursts In Rats Trained To Self-Administer Nicotine Or Food In 1-H Daily Sessions. *Am J Transl Res*, *4–4*, 422–431. Retrieved from `Https://Www.Ajtr.Org`

Ratcliff, R. (2006). Modeling response signal and response time data. *Cognitive Psychology*, *53*(3), 195–237. doi: 10.1016/j.cogpsych.2005.10.002

Rayner, K. (1998). Eye Movements in Reading and Information Processing: 20 Years of Research. *EYE MOVEMENTS IN READING*.

Reinagel, P., & Zador, A. M. (1999). Natural scene statistics at the centre of gaze. *Network: Computation in Neural Systems*, *10*(4), 341.

Renninger, L. W., Verghese, P., & Coughlan, J. (2007, February). Where to Look Next? Eye Movements Reduce Local Uncertainty. *Journal Of Vision*, *7–3*, 1–17. doi: 10.1167/7.3.6

Rensink, R. A. (2000). Seeing, sensing, and scrutinizing. *Vision Research*, *40*(10), 1469–1487. doi: 10.1016/S0042-6989(00)00003-1

Rich, A. N., Kunar, M. A., Van Wert, M. J., Hidalgo-Sotelo, B., Horowitz, T. S., & Wolfe, J. M. (2008). Why do we miss rare targets? Exploring the boundaries of the low prevalence effect. *Journal of Vision*, *8*(15), 15–15. doi: 10.1167/8.15.15

Richardson, J. T. (2011). Eta squared and partial eta squared as measures of effect size in educational research. *Educational Research Review*, *6*(2), 135–147. doi: 10.1016/j.edurev.2010.12.001

Rovamo, J., Leinonen, L., Laurinen, P., & Virsu, V. (1984). Temporal Integration and Contrast Sensitivity in Foveal and Peripheral Vision. *Perception*, *13*(6), 665–674. doi: 10.1068/p130665

Sawaki, R., Luck, S. J., & Raymond, J. E. (2015). How Attention Changes in Response to Incentives. *Journal of Cognitive Neuroscience*, *27*(11), 2229–2239. doi: 10.1162/jocna00847

Schill, H. M., Cain, M. S., Josephs, E. L., & Wolfe, J. M. (2020). Axis of rotation as a basic feature in visual search. *Attention, Perception, & Psychophysics*, *82*(1), 31–43. doi: 10.3758/s13414-019-01834-0

Schroeder, S. R., & Holland, J. G. (1968). Operant control of eye movements during human vigilance. *Science*, *161*(3838), 292–293.

Schwark, J., Sandry, J., & Dolgov, I. (2013). Evidence for a Positive Relationship between Working-Memory Capacity and Detection of Low-Prevalence Targets in Visual Search. *Perception*, *42*(1), 112–114. doi: 10.1068/p7386

Schwark, J., Sandry, J., Macdonald, J., & Dolgov, I. (2012, August). False Feedback Increases Detection of Low-Prevalence Targets in Visual Search. *Attention Perception & Psychophysics*, *74*(8), 1583–1589. doi: 10.3758/S13414-012-0354-4

Schwark, J. D., Macdonald, J., Sandry, J., & Dolgov, I. (2013, July). Prevalence-Based Decisions Undermine Visual Search. *Visual Cognition*, *21–21*, 541–568. Retrieved from `Http://Dx.Doi.Org/10.1080/13506285.2013.811135` doi: 10.1080/13506285.2013.811135

Seitz, A. R., Kim, D., & Watanabe, T. (2009, March). Rewards Evoke Learning of Unconsciously Processed Visual Stimuli in Adult Humans. *Neuron*, *61*(5), 700–707. doi: 10.1016/J.Neuron.2009.01.016

Serences, J. T. (2008). Value-Based Modulations in Human Visual Cortex. *Neuron*, *60*(6), 1169–1181. doi: 10.1016/j.neuron.2008.10.051

Shahan, T. A. (2022, June). A Theory of the Extinction Burst. , *45–519*. (Publication Title: Perspectives On Behavior Science) doi: 10.1007/S40614-022-00340-3

Shiffrin, R. M., & Gardner, G. T. (1972). Visual processing capacity and attentional control. *Journal of experimental psychology*, *93*(1), 72. (Publisher: American Psychological Association)

Sidman, M. (1960). *Tactics of scientific research.* Basic Books, Incorporated, Pub.

Simon, H. A. (1967). Motivational and emotional controls of cognition. *Psychological Review*, *74*(1), 29–39. doi: 10.1037/h0024127

Skinner, B. F. (1953). Some contributions of an experimental analysis of behavior to psychology as a whole. *American Psychologist*, *8*(2), 69. (Publisher: American Psychological Association)

Skinner, B. F. (1984). Selection by consequences. *Behavioral and brain sciences*, *7*(4), 477–481.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological monographs: General and applied*, *74*(11), 1. (Publisher: American Psychological Association)

Staddon, J. E. R. (1983). *Adaptive behavior and learning.* Cambridge [Cambridgeshire] ; New York: Cambridge University Press.

Stanislaw, H., & Todorov, N. (1999). Calculation Of Signal Detection Theory Measures.
(Publication Title: Behavior Research Methods, Instruments, & Computers)

Tanner, W. P., Swets, J. A., Green, D. M., & Macnee, A. B. (1956). *Some general properties of the hearing mechanism* (Tech. Rep.). University of Michigan.

Tatler, B. W. (2007). The central fixation bias in scene viewing: Selecting an optimal viewing position independently of motor biases and image feature distributions. *Journal of vision*, *7*(14), 4–4.

Tatler, B. W., Baddeley, R. J., & Gilchrist, I. D. (2005). Visual correlates of fixation selection: Effects of scale and time. *Vision research*, *45*(5), 643–659.

Tatler, B. W., & Vincent, B. T. (2009). The prominence of behavioural biases in eye guidance. *Visual Cognition*, *17*(6-7), 1029–1054. (Publisher: Taylor & Francis)

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

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

Theeuwes, J. (1994). Endogenous and exogenous control of visual selection. *Perception*, *23*(4), 429–440. (Publisher: SAGE Publications Sage UK: London, England)

Theeuwes, J. (2010). Top–down and bottom–up control of visual selection. *Acta Psychologica*, *135*(2), 77–99. doi: 10.1016/j.actpsy.2010.02.006

Theeuwes, J., & Belopolsky, A. V. (2012). Reward grabs the eye: Oculomotor capture by rewarding stimuli. *Vision Research*, *74*, 80–85. doi: 10.1016/j.visres.2012.07.024

Thorndike, E. L. (1911). *Animal intelligence: Experimental studies.* Transaction

Publishers.

Torralba, A., Oliva, A., Castelhano, M. S., & Henderson, J. M. (2006). Contextual guidance of eye movements and attention in real-world scenes: The role of global features in object search. *Psychological Review*, *113*(4), 766–786. doi: 10.1037/0033-295X.113.4.766

Townsend, J. T. (1976). Serial and within-stage independent parallel model equivalence on the minimum completion time. *Journal of Mathematical Psychology*, *14*(3), 219–238. doi: 10.1016/0022-2496(76)90003-1

Townsend, J. T., & Wenger, M. J. (2004). A Theory of Interactive Parallel Processing: New Capacity Measures and Predictions for a Response Time Inequality Series. *Psychological Review*, *111*(4), 1003–1035. doi: 10.1037/0033-295X.111.4.1003

Treisman, A. (1991). Search, similarity, and integration of features between and within dimensions. *Journal of Experimental Psychology: Human Perception and Performance*, *17*(3), 652. (Publisher: American Psychological Association)

Treisman, A., & Sato, S. (1990). Conjunction search revisited. *Journal of experimental psychology: human perception and performance*, *16*(3), 459. (Publisher: American Psychological Association)

Treisman, A. M. (1960). Contextual cues in selective listening. *Quarterly Journal of Experimental Psychology*, *12*(4), 242–248. (Publisher: Taylor & Francis)

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

Trommershäuser, J., Glimcher, P. W., & Gegenfurtner, K. R. (2009, November). Visual Processing, Learning And Feedback In The Primate Eye Movement System. *Trends In Neurosciences*, *32*(11), 583–590. Retrieved from `Https://Doi.Org/10.1016/J.Tins.2009.07.004` doi: 10.1016/J.Tins.2009.07.004

Van Wert, M. J., Horowitz, T. S., & Wolfe, J. M. (2009). Even in correctable search, some types of rare targets are frequently missed. *Attention, Perception, & Psychophysics*, *71*(3), 541–553. doi: 10.3758/APP.71.3.541

Vincent, B. T., Baddeley, R., Correani, A., Troscianko, T., & Leonards, U. (2009). Do we look at lights? using mixture modelling to distinguish between low-and high-level factors in natural image viewing. *Visual Cognition*, *17*(6-7), 856–879.

Vitu, F., Kapoula, Z., Lancelin, D., & Lavigne, F. (2004). Eye movements in reading isolated words: Evidence for strong biases towards the center of the screen. *Vision research*, *44*(3), 321–338.

Volkmann, F. C. (1962). Vision during voluntary saccadic eye movements. *JOSA*, *52*(5), 571–578. (Publisher: Optica Publishing Group)

Voss, P., Mccarthy, D., & Davison, M. (1993, September). Stimulus Control And Response Bias In An Analogue Prey-Detection Procedure. , *60*. (Publication Title: Journal Of The Experimental Analysis Of Behavior)

Watson, D. G., & Humphreys, G. W. (1997). Visual marking: prioritizing selection for new objects by top-down attentional inhibition of old objects. *Psychological review*, *104*(1), 90. (Publisher: American Psychological Association)

Watson, P., Pearson, D., Most, S. B., Theeuwes, J., Wiers, R. W., & Le Pelley, M. E. (2019). Attentional capture by Pavlovian reward-signalling distractors in visual search persists when rewards are removed. *PLOS ONE*, *14*(12), e0226284. doi: 10.1371/journal.pone.0226284

Weisman, R., Gibson, M., & Rochford, J. (1984). Testing models of delayed sequence discrimination in pigeons: The advance key procedure. *Canadian Journal of Psychology/Revue canadienne de psychologie*, *38*(2), 256. (Publisher: Canadian Psychological Association)

Wertheim, T. (1894). Uber die indirekte Sehscharfe. *Zeitschrift fur Psychologie*, *7*, 172–187.

Wiecek, E., Pasquale, L. R., Fiser, J., Dakin, S., & Bex, P. J. (2012). Effects of Peripheral Visual Field Loss on Eye Movements During Visual Search. *Frontiers in Psychology*, *3*. doi: 10.3389/fpsyg.2012.00472

Wienrich, C., Heße, U., & Müller-Plath, G. (2009). Eye movements and attention in visual feature search with graded target-distractor-similarity. *Journal of Eye Movement Research*, *3*(1). doi: 10.16910/jemr.3.1.4

Williams, B. A. (1976, September). Elicited Responding to Signals for Reinforcement: The Effects of Overall versus Local Changes in Reinforcement Probability. *Journal Of The Experimental Analysis Of Behavior*, *26*, 213–220.

Wolfe, J. M. (1994). Guided Search 2.0: A Revised Model Of Visual Search. *Psychonomic Bulletin & Review*, *1*(2), 202–238. (Publisher: Psychonomic Society, Inc.)

Wolfe, J. M. (1998). What Can 1 Million Trials Tell Us About Visual Search? *Psychological Science*, *9*(1), 33–39. doi: 10.1111/1467-9280.00006

Wolfe, J. M. (2021). Guided Search 6.0: An updated model of visual search. *Psychonomic Bulletin & Review*, *28*(4), 1060–1092. doi: 10.3758/s13423-020-01859-9

Wolfe, J. M., Alvarez, G. A., Rosenholtz, R., Kuzmova, Y. I., & Sherman, A. M. (2011). Visual search for arbitrary objects in real scenes. *Attention, Perception, & Psychophysics*, *73*(6), 1650–1671. doi: 10.3758/s13414-011-0153-3

Wolfe, J. M., Butcher, S. J., Lee, C., & Hyle, M. (2003). Changing your mind: On the contributions of top-down and bottom-up guidance in visual search for feature singletons. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(2), 483–502. doi: 10.1037/0096-1523.29.2.483

Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: an alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human perception and performance*, *15*(3), 419.

Wolfe, J. M., & Gray, W. (2007). Guided search 4.0. *Integrated models of cognitive systems*, 99–119.

Wolfe, J. M., & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, *5*(6), 495–501. doi: 10.1038/nrn1411

Wolfe, J. M., & Horowitz, T. S. (2017). Five factors that guide attention in visual search. *Nature Human Behaviour*, *1*(3), 0058. doi: 10.1038/s41562-017-0058

Wolfe, J. M., Horowitz, T. S., Kenner, N., Hyle, M., & Vasan, N. (2004). How fast can you change your mind? The speed of top-down guidance in visual search. *Vision Research*, *44*(12), 1411–1426. doi: 10.1016/j.visres.2003.11.024

Wolfe, J. M., Horowitz, T. S., & Kenner, N. M. (2005). Rare items often missed in visual searches. *Nature*, *435*(7041), 439–440. doi: 10.1038/435439a

Wolfe, J. M., Horowitz, T. S., Van Wert, M. J., Kenner, N. M., Place, S. S., & Kibbi, N. (2007). Low target prevalence is a stubborn source of errors in visual search tasks. *Journal of Experimental Psychology: General*, *136*(4), 623–638. doi: 10.1037/0096-3445.136.4.623

Wolfe, J. M., & Van Wert, M. J. (2010). Varying Target Prevalence Reveals Two Dissociable Decision Criteria in Visual Search. *Current Biology*, *20*(2), 121–124. doi: 10.1016/j.cub.2009.11.066

Woodman, G. F., & Luck, S. J. (2004). Visual search is slowed when visuospatial working memory is occupied. *Psychonomic Bulletin & Review*, *11*(2), 269–274. doi: 10.3758/BF03196569

Wurtz, R. H. (2008). Neuronal mechanisms of visual stability. *Vision Research*, *48*(20), 2070–2089. doi: 10.1016/j.visres.2008.03.021

Yantis, S., & Egeth, H. E. (1999). On the distinction between visual salience and stimulus-driven attentional capture. *Journal of experimental psychology: Human perception and performance*, *25*(3), 661–676. (Publisher: American Psychological Association)

Yantis, S., & Jonides, J. (1984). Abrupt visual onsets and selective attention: evidence from visual search. *Journal of Experimental Psychology: Human perception and performance*, *10*(5), 601. (Publisher: American Psychological Association)

Yu, X., & Geng, J. J. (2019). The attentional template is shifted and asymmetrically sharpened by distractor context. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(3), 336–353. doi: 10.1037/xhp0000609

Zelinsky, G. J. (2008). A theory of eye movements during target acquisition. *Psychological Review*, *115*(4), 787–835. doi: 10.1037/a0013118

Zelinsky, G. J., & Bisley, J. W. (2015, January). The what, where, and why of priority maps and their interactions with visual working memory. *Annals Of The New York Academy Of Sciences*, *1339*(1), 154–164. doi: 10.1111/Nyas.12606

Zipf, G. K. (2016). *Human behavior and the principle of least effort: An introduction to human ecology.* Ravenio books.