UNIVERSITY OF BIRMINGHAM

Understanding the Effects of Vision Impairment on Visual Search Through Computational Models

by

DALAL KHALID ALJASEM

A thesis submitted to the University of Birmingham for the degree of

DOCTOR OF PHILOSOPHY

School of Computer Science

College of Engineering and Physical Sciences

University of Birmingham

September 2023

Abstract

This PhD thesis explores the application of computational models to help better understand the effects of vision loss on visual search behaviours. It analyses how people adapt to cognitive limitations when interacting with a visual search task. Visual search is a cognitive function that plays a crucial role in everyday activities, and gaining insights into how it differs amongst individuals with vision impairment holds significant practical implications for designing more accessible user interfaces. The research is exploratory and interpretative. It examines the implementation of a cognitive information processing system, which serves as the thinking brain to a computational model (i.e. the mechanism). Specifically, it explores how the mechanism affects visual search strategies in vision impairment. Researchers have used computational-rational agents to analyse and predict eye movement strategies. However, most studies have only focused on the strategies of normal vision without considering vision impairment. Therefore, in this research, we are taking a new approach by attempting to construct a model of eye movements in individuals with different types of vision impairment. To achieve optimal adaptation strategies in the model, we formulate different visual search tasks as a Partially Observable Markov Decision Process (POMDP). This approach would allow the implementation of the constraints imposed by the environment and the human visual system, including the limitations resulting from vision impairment. Solving real-life problems as POMDP is complex and requires significant computational resources. Therefore, we use different Reinforcement Learning (RL) techniques to achieve optimal strategies that we can compare against human data. The thesis considers three types of vision impairments (Glaucoma, Cataracts, and Age-related Macular Degeneration (AMD)). We use the Q-learning method to find the optimal policy for the Glaucoma task. Whereas

for the other tasks, we use Deep Reinforcement Learning (DRL), specifically Proximal Policy Optimisation (PPO). The model finds optimal strategies in each experiment through adaptation to an impaired visual field. The results show that the theoretical approach to optimal rationality can give insight into how people with various types of vision impairment would interact with the screen. Overall, the model achieved the following: (1) adapted strategies to peripheral vision and environment constraints, (2) it showed a decline in performance when given a crowded setting, and (3) it predicted the accuracy of people with Cataracts in two different tasks. The research findings suggest that people with vision impairments adapt to their types of vision loss and the nature of their visual field defect. The findings have significant implications for Human-Computer Interaction (HCI). They can provide insights into the interactions of people with a specific type of vision impairment.

Dedicated to my father, in loving memory.

Acknowledgements

I would like to express my deepest appreciation and gratitude to all those who have supported and guided me throughout the journey of completing this research.

First, I would like to thank my supervisor, Professor Andrew Howes, for giving me the opportunity to work on such an interesting project. Thank you for assessing and directing me throughout the PhD. You have been an inspiration to me, and I am truly lucky to have had the opportunity to work under your guidance. I extend my thanks to the members of the RSMG group, Dr Hyung Jin Chang and Dr Mohan Sridharan. Your feedback and comments were always appreciated, and they definitely helped improve my skills and enhanced the quality of the research.

Thank you to the researchers Muriel Boucart, Miguel Thibaut and Yu Wan, for providing the data from their experiments. I really appreciated their collaboration and their quick responses to the inquiries I had regarding their experiments.

I want to express my deepest appreciation to the Government of the Kingdom of Saudi Arabia for sponsoring my research and to the staff of the Saudi Arabian Cultural Bureau in London for resolving all my sponsorship requests promptly.

I would also like to acknowledge the University of Birmingham faculty and staff, especially the staff in IT, in BlueBear and Baskerville. Thank you to Thomas Kappas, for all the technical help in setting up my projects in BlueBear and for helping me understand how to run my experiments using these resources. Thank you Gavin Yearwood for setting up my Baskerville projects and responding to my requests promptly. This helped a lot in finishing my experiments in a good time.

Thank you to my colleague Dr Haiyang Chen for all the help throughout this PhD. Another special thank you to Dr Xuili Chen and Dr Aditya Acharya for all the help they provided during the early months of my research, for explaining their previous work and for setting up meetings when needed. This, without a doubt, helped me understand the concept of their research, which shaped my approach to my own research.

Thank you to my friend Kha Yiu for all the encouragement, especially during the pandemic, for always reminding me to put my mental health first and for sharing her experience finishing her own PhD research, which gave me a lot of motivation when I most needed it.

Thank you to my friend and yoga teacher, Tina. Learning yoga during PhD was one of the best decisions I've ever made, and I am lucky that I had the chance to learn in her classes. It played a big part in my PhD journey, and I am forever grateful for it. Thank you to my friend Paul for the continuous support and for always giving me valuable advice. Thank you to my parents Amal and Khalid, my brother Mohammed and my sister Mariam for supporting me throughout my academic journey from day one and for always believing in my abilities. I am forever grateful.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The process of designing a new computer interface requires many decisions to consider. For example, during the initial stages of the design, the designers should consider how users are supposed to interact with the given interface and what the best representation of the elements appearing on the computer screen is (Dudley and Kristensson, 2022). Visual communication is the main element when it comes to building computer interfaces, this, in turn, negatively reduces the impact and uptake of Human-Computer Interaction (HCI) techniques for people with vision impairment (Edwards, 1995). Based on World Health Organization (2022) findings, worldwide, there is an estimation of 2.2 billion people who are either blind or visually impaired (BVI). It also shows that the changes in lifestyle (including children spending more time indoors using devices such as computers and tablets) is a main factor of the increasing number of vision impairments. Many diseases cause vision impairment around the world, such as uncorrected refractive errors, cataracts, age-related macular degeneration, and glaucoma, where the causes vary between countries, based on their average income. World Health Organization (2022) also stated that most people with vision impairment are above the age of 50 and as the population grows, the number of elderly people will increase which in turn impacts the number of visually impaired people. Visually impaired people interact with technology and use computer devices on a daily basis, and despite the significant effort at designing accessibility, interaction for the visually impaired remains challenging (Griffin-Shirley et al., 2017; Pal et al., 2017; Vatavu, 2017).

Although there is a fast-growing trend in designing user-friendly interfaces to enhance user experience, people with vision impairment still face difficulties interacting with even the most recent interfaces, including touch screens (Kane et al., 2011), mobile screen readers for blind people (Kuber, Hastings, and Tretter, 2020) or mobile devices provided with VoiceOver (e.g. in Apple devices) (Leporini, Buzzi, and Buzzi, 2012). Assistive technology supporting people with different types of vision impairments has been addressed through different research, and varying solutions were proposed using diverse technologies such as image detection and voice recognition (Barakat et al., 2020), Augmented Reality (AR) (Zhao et al., 2016) and navigation system using Radio Frequency Identification (RFID) (Kulyukin et al., 2004). However, the impact of what are sometimes generic technological solutions on people with very specific impairments and very particular needs is often unclear. One response, amongst many, to this problem must be a deeper understanding of how people with specific impairments interact with computers. Effective redesign may be urgent, but in order for it to be successful, it has to be informed by a deep scientific understanding of impaired vision in interactive contexts.

In this thesis, we propose a computational model that explains the visual behaviours of people with different types of vision impairments and explore *why* and *how* they adapt to their environment the way they do. The theoretical framework that informed building the proposed model uses cognitive science and rationality as an empirical tool to explain *why* people adapt to their environment, while the implementation of decision theory (i.e Partially Observable Markov Decision Processes, or POMDPs) and policy optimisation in Reinforcement Learning model (RL) methods explains *how* people with vision impairment would interact with their environment. Despite the fact that Machine Learning (ML) methods have been widely used to make predictions, in this thesis, our aim is to model HCI, this approach has been discussed in a review by Oulasvirta, Jokinen, and Howes (2022). The goal of the presented model is to provide analysis of how BVI people adapt to their impairment while interacting with a visual search task which have the potential to help improving the design of more accessible computer User Interfaces (UI). This in turn can support the decision-making process for visually impaired people and enhance their overall experience when interacting with UI. Before discussing the

model further throughout the thesis, a brief introduction to the human visual system and the causes of vision impairment will be explained in the next section.

1.1 Vision and Vision Impairment

The way the human eye works is by perceiving light rays reflected or emitted from objects; a healthy eye would perceive the light rays and use optic nerves to send electric signals to the brain. Before reaching the retina, the light hits the eye and passes through different stages starting from the *cornea*, going to the *aqueous humor*, the *iris*, the *lens*, and then the *vitreous humor*. The lights are then detected by the retina and converted into electrical signals by *photoreceptors*. There are two types of photoreceptors, which are the rods and cones, where rods are responsible for dark vision (i.e. *scotopic vision*) as they do not identify colours, and cones are responsible for daylight vision (i.e. *photopic vision*) to distinguish colours (Baxes, 1994). A human vision is considered "normal" if it meets certain visual acuity measurements (i.e. 6/6 in Snellen Chart) (Sue, 2007); this means that a person can read from 6 meters distance the small letters placed on the chart that are supposed to be clear in 6 meters distance in a normal acuity. Visual acuity is the ability of the eye to distinguish small details, which defines the clarity of the image. It is affected by both optical and neural factors. A good visual acuity indicates that the retina is functioning well, and the brain is able to recognise the signals received from the cones in the retina, so it can create the clear desired image (Sandywell, 2016; Snowden et al., 2012). Additionally, the human visual system has a limited visual field. This visual field is a portion of space where objects can be visible while fixating the gaze in one direction. The monocular visual field consists of both central and peripheral vision, where the central vision covers the inner vision of 30 degrees and the central fixation, where the peripheral visual field is where the vision can extend to 100 degrees laterally, 60 degrees medially, 60 degrees upward, and 75 degrees downward (Spector, 1990). Figure 1.1 shows an image of the visual field of the right eye as presented in the book titled *An introduction to clinical perimetry* by Traquair (1931).

Whereas in vision impairment, the ability of the eye to see decreases to a certain level that

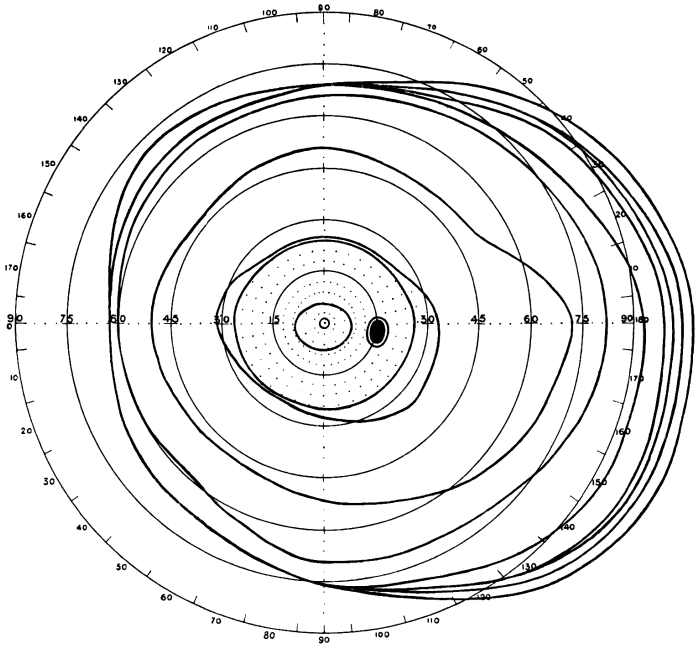

Figure 1.1: Visual field measures as presented by Traquair (1931)

cannot be treated medically or corrected by wearing glasses or contact lenses (World Health Organization, 2022). There has been variation in the exact definition of vision impairments (World Health Organization, 2007, 2008). However, Sight (2023) highlights the main two factors that eye specialists will look for to decide if a person has a type of vision impairment, these two factors are visual acuity and visual field. Vision impairments is categorised based on the visual acuity (World Health Organization, 2008, 2022) as follows:

- Category 0: Equal to or better than 6/18 - 3/10: Mild or no vision impairment
- Category 1: Less than 6/18 or greater than or equal to 6/60: Moderate visual impairment
- Category 2: Less than 6/60 or greater than or equal to 3/60: Severe visual impairment
- Category 3: Less than 3/60 1/20 or greater than or equal to 1/60* (Or counts fingers (CF) at 1 metre) 1/50: Blindness
- Category 4: Less than 1/60* 1/50 with light perception: Blindness
- Category 5: No light perception: Blindness

There are different types of vision impairment, but the focus in this thesis will be on three different types, which are also considered one of the leading causes of blindness worldwide based on (World Health Organization, 2022). These types are glaucoma, Age-related macular degeneration (AMD), and cataracts. A brief explanation of these types will be given in subsection 1.1.1 in order to have an understanding of these types and their effects on people's lives. These types will be used in the upcoming experiments in this thesis.

1.1.1 Types of Vision Impairment

Glaucoma: Glaucoma is a group of eye diseases that is a result of neuropathies leading to progressive loss of ganglion cells (National Eye Institute, 2022; Vingrys, 2000). The results in a study done by Crabb et al. (2013) showed that people with Glaucoma do not observe black patches in the visual field, but their observation is mostly described as blurred vision; similarly, a study done by Hu et al. (2014) reported that people with glaucoma usually need more lights due to blurred vision, they find it difficult to see objects to one or both sides, this is described as looking through dirty glasses, they also had difficulties differentiating boundaries as well as colours. People with glaucoma face challenges in everyday life activities (Hochberg et al., 2012), such as driving (Crabb et al., 2010; Ramulu et al., 2009), reading (Burton, Smith, and Crabb, 2014; Ramulu et al., 2013) as well as negative impacts on physical activities such as fear of falling (Ramulu et al., 2012) and postural stability (Kotecha et al., 2012).

Cataracts: A cataract is the case of cloudiness of the crystalline lens which results when the refractive index of the lens varies significantly over distances approximating the wavelength of the transmitted light (Benedek, 1971; National Eye Institute, 2023; Shiels and Hejtmancik, 2007; Shiels and Hejtmancik, 2013). A person with cataracts would experience blurred vision, and light glare would increase. The quality of vision for people with cataracts can be improved with cataract surgery as mentioned by Skiadaresi et al. (2012) and Raza (2018).

Age-related Macular Degeneration: Age-related macular degeneration (AMD) is an eye condition causing blurred central vision, it happens when ageing causes damage in the macula which is a light-sensitive tissue located at the back of the eye that controls the sharpness of the central vision and lights detection (Jager, Mieler, and Miller, 2008; Macular Society, 2022; National Eye Institute, 2021). There are two types of AMD, which are 'wet' and 'dry'. In the early stages of AMD, the chance of the person having visual loss is moderate, and it is asymptomatic, however, a person might experience some symptoms such as blurred vision and visual scotomas (Jager, Mieler, and Miller, 2008). Dry AMD can develop central visual loss in the course of months to years, whereas patients with the wet type might experience sudden visual loss within days to weeks (Alexander, 2015; Jager, Mieler, and Miller, 2008). Quality of life can be negatively impacted in terms of daily activities as well as psychological well-being as a result of having AMD (Cimarolli et al., 2012; Hochberg et al., 2012; Mathew et al., 2011; Taylor et al., 2016) even in those with mild vision loss cases can be negatively affected due to being diagnosed with AMD (Hassell, Lamoureux, and Keeffe, 2006).

After briefly discussing the biological side of the human visual system, an introduction to the field of active vision and the relevant topic of visual search and its relations to vision impairment is given in the next section, which will be part of this research.

1.2 Towards Computational Models of Vision Impairment in HCI

Visual search involves the human eye and cognitive perception to try and find target objects located within distractors in a certain environment (Treisman and Gelade, 1980; Wolfe, 2001). When humans are presented with a pattern involving texts, shapes and colours, a complex set of eye movements (i.e saccades) is required to allow the eyes to collect information about the environment through direct fixations (Hayhoe and Ballard, 2014; Rayner and Castelhano, 2007; Sprague, Ballard, and Robinson, 2007; Trommershäuser, Glimcher, and Gegenfurtner, 2009). Furthermore, the distance between a cue and the fixation location increases when the

eye fixates on a specific location. This leads to a fall in the density of the cells, and the outcome is increasing blurriness (Rayner and Castelhano, 2007). Eye movements play an important part in psychological processes involved in different tasks such as reading, visual search, and scene perception (Rayner and Castelhano, 2007). *Foveated vision* covers 1-2 degrees of the visual angle (Bringmann and Wiedemann, 2022; Duchowski, 2018), and it is the area with the highest density as well as colour vision receptor cells. The information outside the foveated vision cannot be fully processed without eye movements (Rayner, 1998). On the other hand, and despite reduced acuity, as stated by Geisler (2011), the *periphery* covers a larger area than the fovea and therefore allows information on colour, text and size to be gathered during eye movements' engagement of visual search. Having a distinct acuity function, the impact of colour, shape and size can be studied independently or collectively (Kieras and Hornof, 2014). Previous research findings showed that eye movement performance is optimised through visual feedback (Ono, Das, and Mustari, 2012; Robinson, 1975; Schubert and Zee, 2010). Eye movements in vision impairment have been intensively studied in psychology. Those studies are mainly done through experiments involving human participants (Boucart et al., 2020; Glen, Smith, and Crabb, 2013; Loughman, Davison, and Flitcroft, 2007; Thibaut, Boucart, and Tran, 2020; Van der Stigchel et al., 2013; Wan et al., 2020). Using this approach. Researchers have been able to measure eye movement behaviours and gaze strategies. The experiments designed for visually impaired people usually include a stimulus for those who can observe (or see) the stimulus. However, if participants cannot see the stimulus due to their vision impairment, different alternative methods can be used, such as using the magnetic field and search coil method (Robinson, 1963; Schneider et al., 2013). The research presented in this thesis will focus on conducting computer science experiments based on human experiments that include stimuli presented to human participants.

This thesis discusses the hypothesis of having a probabilistic model of adaptive interaction in vision impairment to help understand the visual behaviours of individuals with vision loss and, consequently, build better and more accessible user interfaces. The model was built based on the theoretical framework, *computational rationality* (Lewis, Howes, and Singh, 2014), which is an application of *bounded optimality* (Russell and Subramanian, 1994) assuming

that human behaviours are the results of the adaptation to the bounds of the human brain as well as the structure of the environment. Several authors have considered the importance of building computational models using both cognitive science and ML where they also discussed its potential to help to understand the human perceptual strategies to support HCI (Acharya, 2019; Acharya et al., 2017; Chen, 2015; Chen, Acharya, and Oulasvirta, 2021; Chen et al., 2015, 2017; Howes et al., 2018; Li et al., 2023; Olson and Olson, 1995; Oulasvirta, Jokinen, and Howes, 2022; Tseng and Howes, 2015). However, the majority of those studies have only focused on normal vision and did not consider those models' importance in supporting people with vision impairment in HCI. In the next chapter, we will explore computational rationality and the related studies in more depth, along with discussing current computational approaches supporting vision impairments.

1.3 Contributions

The present thesis makes four important contributions to computational rationality further to understand visual search behaviours in people with vision impairment.

1. The thesis represents a computational model that explains the behaviours of different types of vision impairment by implementing the cognitive bounds into the ML model.

2. The model represented in the thesis finds strategies for people with vision impairments through the implementation of probabilistic environments, information processing and utility maximisation.

3. The thesis shows the application of different methods of RL used to solve visual search problems formulated as sequential decision process (POMDP).

4. The model supports HCI as it presents an analysis of computational rationality which can explain why people with different types of vision impairment choose their strategies when they interact with a given visual search task on the screen.

1.4 Thesis Outline

- **Chapter 2:** Discusses the literature review, including computational rationality and different models of visual search. It also discusses different computational approaches for vision impairment, and computational models supporting HCI.

- **Chapter 3:** Provides an overview of the model implemented in the thesis, including the background of POMDP mathematical formulations and Reinforcement Learning methods. It also illustrates the technical implementation of the model with workflow explaining the iterative nature of the experiments conducted in the research.

- **Chapter 4:** This chapter discusses the first experiment implemented in this research which involves a visual search task specified for glaucoma patients.

- **Chapter 5:** In this chapter, we investigate the crowding effect in a task that involves patients with AMD and test the crowding effect in the visual search model by giving different crowding settings.

- **Chapter 6:** This chapter considers cataracts patients and implements the model into two visual search tasks. The first task in subsection 6.4.1 is a an object search task, and the second task is a face recognition task (subsection 6.4.2).

- **Chapter 7:** This is the final chapter in the thesis which discusses the summary of the model's results. It also discusses how this research contributes to HCI. Lastly, it discusses possible improvements and future work that can extend the research.

Chapter 2

Literature Review

2.1 Introduction

There are many aspects of the model presented in this thesis. As mentioned in Chapter 1, the model was built based on the *computational rationality* framework (Lewis, Howes, and Singh, 2014). Therefore, a systematical review of the framework will be discussed in this chapter, in addition to exploring different approaches to visual search modelling in the literature. Another theoretical approach is *Adaptive Interaction* framework (Payne and Howes, 2013), which was also used in this research and is closely related to computational rationality. The framework explains *interaction* as adaptive behaviours in the human brain through the combination of a utility function and ecology (as observed by individuals) given the bounds of the human brain, which in return can explain the strategy an individual follows in their decision-making process. In other words, the framework illustrates that humans interact by adapting to their environment based on their knowledge of the environment within the limitations of their cognitive system. The three main elements of the model are utility, ecology, and mechanism, which can help define the strategy. In this research, we implement the bounds of both normal vision systems alongside vision impairments where we can show how each (i.e. people with normal vision and people with vision impairment) would interact with the same visual search task given the differences in how they observe the environment. The four components of *Adaptive Interaction* framework

will be reviewed in more depth in this chapter. We will explain how each piece can play a part in directing the eyes to the target in visual search tasks. Moreover, studies of assistive technology for people with vision impairment will be highlighted, including different approaches supporting visual search and decision-making strategies in vision impairment. Lastly, we will delve into the literature on the applications of gaze computational models in HCI, including Augmented Reality (AR), Virtual Reality (VR), or display screen equipment.

2.2 Computational Rationality

The questions of *why* and *how* people adapt to their environment were discussed by Anderson (1991) in a review on the *rational analysis* framework and its applications. It was explained that the cognitive processes of humans' decision-making are usually optimisation processes which maximise expected gain in contrast with the cost of the psychological efforts in the mind, for example, maximising the trade-off between the chance of finding the relevant information in the memory and the cost associated with the attempts finding this information. Using this approach, Anderson (1991) explained that the results of these processes could be predicted based on the environment's statistical nature rather than the mind's structure. Rational theory was applied as a 6-step iterative process that was grouped into two categories of a 3-step process: the first category (consisting of steps 1-3) is considering the process of creating the theory, where steps 4-6 in the second category consider finding the optimal behaviour and predicting the human behaviours. The steps are as follows:

1. Specifying a goal for the cognitive system to achieve.
2. Structuring the environment as a formal model that the system can adapt to.
3. Giving assumptions about the computational constraints in the system that might prevent it from achieving global optimum.
4. The combination of the three previous steps to set an optimal behaviour.
5. Comparing the optimal behaviours in the model against human data for validation.

6. Iteratively refining the theory that was created in the previous steps and making adjustments to improve the model.

Anderson's strategy assumes that the human will always behave optimally towards the given structure of the environment without considering the mechanisms of the mind. This suggests that his approach can only explain *why* people adapt to their environment, leaving the question *how* they adapt unanswered. On the other hand, Russell and Subramanian (1994) introduced the *bounded-optimality* framework; in this framework, an intelligent agent can be created by the combination of a given formal structure of the environment and a suitable device that can run an optimisation program. It stated that the feasible behaviours of the bounded-optimal agent are generated by the given task's environment and a constrained optimisation problem which describes the agent's mechanism. In other words, a bounded-optimal agent uses both rational analysis as well as a pre-defined environment to choose its behaviours which can explain the answer to *"why"* an intelligent agent is choosing certain behaviours.

Computational Rationality is a framework proposed by Lewis, Howes, and Singh (2014) where the term computational rationality stands for combining both rationality's definition and computational mechanisms. The framework was developed based on the idea that the chosen behaviours by cognitive systems are not only adapted by the environment's structure but also by the mechanism the mind is following. The main goal of the computational rationality framework is to address the theories of developing the mind's mechanisms and the adapted behaviours by maximising utility values in psychological models. The framework introduces processing bounds to expand the standard rational approaches question (i.e. *what should a rational agent do in this environment?*) to be (*What should a (utility-maximising) rational agent, with its own available information-processing mechanisms, do in this environment?*). In order to build a meaningful strategy in this framework, three components are needed: (1) utility: which describes what a person would prefer to do based on a gained value, (2) ecology: which represents the constraints in a person's mind imposed by interaction with the environment and the experience they gained over time through this interaction, (3) mechanism: which defines the human brain and the processes associated with making decisions (Payne and Howes, 2013), the

three components are illustrated in Figure 2.1.

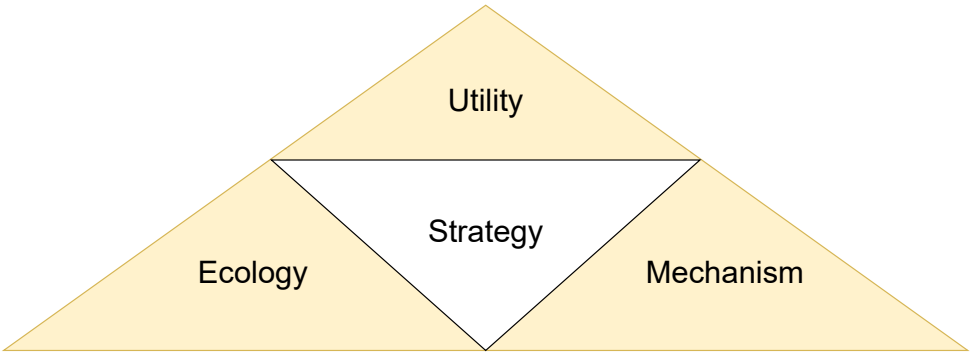

Figure 2.1: Computational Rationality framework components as represented in Payne and Howes (2013)

The three components are explained in more detail below, in addition to a further explanation of the strategy, which is formalised based on these three components.

Utility

The term *utility* refers to the human preference for gaining pleasure over experiencing pain (Etzioni, 1986). Different concepts of utility have been discussed by (Kahneman and Tversky, 2013; Von Neumann, Morgenstern, and Kuhn, 2007). In *expected utility theory* (Von Neumann, Morgenstern, and Kuhn, 2007), it stated that people under uncertainty would compare expected utility values between different options to make a final decision; their expected value is the statistical expected value of the final decision. Wheres, in *prospect theory* (Kahneman and Tversky, 2013), it was explained that people in risky situations would prefer to go with immediate choices that give them the specific value of gain or loss (i.e. gambling their decisions) over making a final decision based on the end outcome.

In Computational Rationality, utility concerns what people desire and consider valuable, and it shapes peoples' preferences and choices by defining how people trade quantities that hold physiological, cognitive, or social importance. This concept is closely connected to *intrinsic motivation and reward* concept in Singh et al. (2010), where decision-making strategies of humans are guided by the optimised utility of future states, which includes internal mental events. In Reinforcement Learning (RL), the utility function can be computed as a reward

function (Singh et al., 2010).

Ecology

In human interaction, *ecology* refers to the constraints imposed on an individual's perception of the environment, resulting from immediate and lifelong interactions with it. The framework describes the statistical distribution of the environment the agent interacts with, given that it has been interacting with it throughout its lifetime while considering the agent to be computationally rational. As discussed in *rational analysis* (Anderson, 1991), a statistical environment can lead to an optimal decision. The same concept was also addressed in (Schooler and Anderson, 1997), where the goal was to predict the odds of encountering an item in a speech to children and word usage in the front page headlines of the New York Times using the theory that human memory can optimally adapt to a given task environment. Also, given a visual search task, it was shown that humans adapt visual search strategies to the task ecology statistic (Chen et al., 2015; Tseng and Howes, 2015). An example of how ecology plays a part in the strategies of human visual search is shown in (Vlaskamp, Over, and Hooge, 2005); in their experiments, they showed that changing the space between elements in a display would affect the search time, where the search time is the number of eye movements multiplied by the average of fixation duration. Their findings showed that if the spacing distance is in the range between 3.4° - 7.1°, the search time will increase, whereas decreasing the distance less than 3.5° did not show any effect, and only if the distance between elements was less than 1.5°, then this would lead to increase in the search time. These findings show how the statistical representation of the environment would impact humans' decision-making strategies.

Mechanism

The study of cognitive mechanism investigates how the human brain recognises the various mental activities, including perception, learning, memory, thinking, emotion, and consciousness (Shi, 2021). The cognitive mechanism involves recognising the mechanism responsible for a specific phenomenon and breaking it down into components and operations (Bechtel, 2008).

In computational rationality, the term *mechanism* refers to the fundamental architecture of the human mind. This architecture maps sensory inputs, such as sight, hearing, touch, and smell, to be stored, processed, and eventually transformed into interactive behaviours like eye movements (Payne and Howes, 2013). Computational models have been developed, including cognitive architectures with production rule (Anderson, 1996; Kieras and Meyer, 1997), image-based control of attention deployment (Itti and Koch, 2001), and artificial neural networks (ANN) as reported in (McClelland and Cleeremans, 2009). These models offer insights into how the human mind processes information. Theoretical assumptions about the underlying mechanism enable the construction of inferences and decision-making processes under specific cognitive constraints.

Strategy

Strategies are shaped by the previous three components (i.e. utility, ecology and mechanism) as shown in Fig. 2.1; a strategy is defined as a discretionary method for achieving good behaviour, and in order to find the optimal strategy (i.e. the optimal policy), the utility value must be maximised. Strategies themselves cannot be bounded; however, when the three components are combined, they result in a bounded strategy space. If any of these components is missing or not adequately considered, it can lead to an unbounded strategy space. Thus, it is the integration of these three components that imposes the necessary boundaries and constraints on strategies, enabling them to be effectively analysed and optimised (Payne and Howes, 2013). A strategy is a computational program involving a sequence of actions taken by an agent to interact with the environment to ultimately accomplish the main task goal. Achieving the task's goal may require running multiple programs with different combinations of actions. Therefore, the framework assumes that the human mind can be seen as a bounded computational unit capable of executing various programs with different sequences of actions (Lewis, Howes, and Singh, 2014). Optimal actions taken by individuals are not solely a result of the adaptation to the environment but also a result of the running bounded computational unit (i.e. the human brain). If a strategy is found in a developed computational model that matches the original human data, then the model can be

used for further predictions and analysis. However, failure to match the human data may indicate that the chosen utility function is not suitable or the way the human information is processed, or incorrect implementation of human information processing constraints in the model.

In the following section, we delve into the phenomena of adaptation within the field of Human-Computer Interaction (HCI). We will go in more depth exploring the answer as to *why* humans choose and adapt their strategies to the limitations of their brains when engaging with computers.

2.3 Adaptive Interaction Phenomena

The theory of interaction in computational rationality suggests that people adapt to various types of bounds when they interact with their environment. In HCI, adaptation is also shaped through regular interaction with computers and dealing with problems that arise through those interactions (e.g. navigating icons). Human strategies in HCI are shaped by different phenomena that are mainly formed based on an individual's cognitive system, capabilities and preferences that are driven by differences in personalities (Ackermans et al., 2020; Ashktorab et al., 2019; Egan, 1988; Mackay, 1988; Oulasvirta, Jokinen, and Howes, 2022; Oviatt, Lunsford, and Coulston, 2005; Vorvoreanu et al., 2019). The purpose of this section is to demonstrate the various phenomena in HCI and how they help in understanding the reasons behind the diverse ways in which humans adapt. The phenomena below are listed in the same order as (Oulasvirta, Jokinen, and Howes, 2022).

2.3.1 Limitations in Memory

The bounds imposed by human memory (e.g. working memory, semantic memory, episodic memory etc.) play a significant role in how people adapt to interactive technology. Many aspects of UI design are influenced by the way people remember and process information, such as breaking long numbers into blocks that can be easily memorised by users (Johnson, 2020) or introducing additional features to enhance users' experience, for example, adding

additional menu bar to help developers navigate through complex source code (DeLine et al., 2006). Humans tend to develop bounded optimal strategies when adapting to the limitations of their memory (Gray et al., 2006; Howes et al., 2016). Optimal strategies are usually achieved through experience, considering cognitive constraints and maximised utility. For example, in the experiments done by (Howes et al., 2016), participants were asked to copy information from email into a menu through a no-choice/choice paradigm. Most participants showed bounded optimal adaptation and chose which items to remember based on the given utility aspects of the experiments. Therefore participants chose to remember to maximise utility given speed and accuracy constraints as well as personal preferences (Howes et al., 2016). The experiments in the previously mentioned study demonstrate how individuals can adapt their behaviours and adjust their memory strategies under specific constraints to reach optimal performance.

2.3.2 Bounds in Perceptual Systems

Visual limits imposed by the human visual system play a crucial role in HCI (Byrne, 2001; Dimara and Perin, 2019; Ware, 2019). For instance, due to the fall in acuity within the peripheral vision, visual search is guided based on the information a person can actively observe and collect through eye movements, allowing information that cannot be observed through the peripheral vision to be within the fovea (Kieras and Hornof, 2014; Lavie, 2010; Wolfe and Horowitz, 2017). People tend to adapt their eye movement strategies based on the different cues available within their field of view. For example, within peripheral vision, shape information is more challenging to perceive compared to colour or size, leading people to be less likely to rely on it to guide their search (Findlay and Gilchrist, 2003; Kieras and Hornof, 2014; Williams, 1967). Research also suggests that individuals adapt their eye movement strategies based on each person's unique pattern of retinal cones (Findlay and Gilchrist, 2003; Geisler, 2011; Hofer, Singer, and Williams, 2005). Also, people tend to find strategies to maximise the information gathered that can lead to the correct target location (Najemnik and Geisler, 2008). Furthermore, individuals with vision impairments often develop adapted strategies to cope with their condition. For example, those with AMD may identify a particular area in their visual field known as a preferred retinal locus

(PRL) which can act as an alternative for the fovea (Fletcher and Schuchard, 1997). We will delve deeper into vision impairment adaptation in subsequent chapters.

2.3.3 Motor Bounds

Aimed movement is a fundamental aspect of HCI. For example, in a pointing task, people tend to prioritise object selection based on its size rather than its distance, which results in minimising the time cost of locating the object starting from the first point of movement (Fitts, 1954; MacKenzie, 2018). However, people usually miss their intended target due to uncertainty caused by motor noise. This noise can lead to the users making corrective movements or unintentionally selecting adjacent objects resulting in potential errors and costs. While Fitts's law is commonly used in HCI to predict the time users take in a pointing task (Fitts, 1954), it does not fully consider the adaptation to motor noise in movement time. More precisely, as the target size decreases, movement time increases, but due to the perceptual noise, users tend to adapt differently in how they perform the task. For instance, in an experiment measuring reaction time (RT) latencies of the saccadic eye and goal-directed hand movements, participants adapted their RT latency of hand and eye movements based on the characteristics and conditions of each task (Bekkering et al., 1994). Adapting to perceptual/motor bounds goes beyond the target size or width. It is crucial in HCI, allowing UI designers to improve the overall user experience by considering motor noise and how users would adapt to it. Which, in return, can broaden the scope of design considerations and ensure that interface design is optimised based on users' specific needs.

2.3.4 Environment Structure

As previously mentioned in Section 2.2, people adapt to the environment's *ecology*. Previous studies have shown the importance the environment plays in adaptation to technology (Hollan, Hutchins, and Kirsh, 2000; Sarcar et al., 2018). Sarcar et al. (2018) investigated how improving touchscreen design to match individuals with different types of cognitive disabilities, such as dyslexia, can positively affect the way they interact with a touchscreen keyboard. Another study

by Binetti et al. (2021) examines how visual and auditory cues impact users' ability to locate objects within a head-mounted augmented reality (AR) task. The previously mentioned studies exemplify the human active nature of adaptation to environmental changes, as people modify their task performance to compensate for increased uncertainty or error rates. The empirical analyses by Pirolli (2007) showed that people tune their decision-making-trade-off based on the environmental probabilities and conditions, as users would more likely exit a website when the gained utility no longer outweigh the benefits of exploring other options. The theory in adaptive information interaction introduced by Pirolli (2007) emphasises the need to understand the human psychological mechanisms and strategies adapted to a given environment to enhance user experience in HCI.

2.4 Models of Visual Search

This section summarises various existing models of visual search, including guided search, Bayesian, control and deep learning models, signal detection theory, and cognitive architecture.

2.4.1 Guided search model

Guided Search model is a model of human search behaviour by Wolfe (1994) and Wolfe and Gray (2007), and it was based on the feature-integration theory of attention proposed by Treisman and Gelade (1980). Initially, visual search was discussed by two main types of search, which are serial and parallel search (Wolfe, 1994), and in order to guide attention to the target, independent feature maps were created to work in parallel. Each feature is represented in a feature map such as "colour", "orientation" or "shape". The presence or absence of a feature was distributed in each feature map to direct attention to the target (i.e. activation of the feature).

A weighted sum map describes the activations. In a bottom-up process, low-level features extracted from the display are presented in different channels. Following this process, a top-down weighting process is initiated, where higher weights are allocated to target features and other features are "ignored" by being allocated lower weights. This facilitates higher activation

of targets from feature channels, while distractors are only 'eligible' for activation from a single target channel, which explains ignoring a specific set of items through the selectivity demonstrated by people.

The Guided Search model lacks explicit *utility* function, and *ecology* is defined in the local task environment. Heuristic control (i.e. MAX-rule), where attention is attracted to the highest activation, defines the *strategy*. The decline in acuity is not considered in Guided search models, which constitute one of the key weakness of the models, which focuses less on the overt deployment of attention and more on the covert attention, therefore, not considering the behaviour of eye movement.

2.4.2 Signal detection theory

Visual search models based on Signal Detection Theory (SDT) are search models that focus on a single parallel processing stage (Eckstein et al., 2000; Verghese, 2001). This is contrary to the two-stage information processing architecture presented by Treisman and Gelade, 1980.

The goal in SDT models is to identify the target stimulus against a background (e.g. distrators or surrounding noise). The correct identification of signal (or hit) as opposed to a false alarm defines SDT models' *utility*, and as in Guided Search models. The description of *ecology* is encoded in the local task environment, and a heuristic control (MAX-rule) describes the *strategy*. SDT framework, in addition to presenting itself as rejecting the two-stage information processing architecture, aims to explain the visual search process as an adaptation to the constraints in the visual system (noise) and the local task environment (Acharya, 2019). While SDT models present some success at explaining effects due to distractor set size and similarity, they exhibit an inability to identify the behaviours that form before reaching the end goal.

2.4.3 Bayesian Models

Modelling visual search tasks using Bayesian methods aims to explain the search process by utilising the Bayesian probabilistic framework (Butko and Movellan, 2008; Elazary and Itti, 2010; Myers, Lewis, and Howes, 2013; Najemnik and Geisler, 2005, 2008; Nunez-Varela and

Wyatt, 2013). In this model, Najemnik and Geisler (2008) argues that the *utility* can be described as a function that "minimises the number of fixations" to locate the target while at the same time maintaining the error rate below a certain threshold, and *ecology* is described in the model of the external environment. Furthermore, *strategy* is defined as a heuristic control deploying either a Maximum a Posterior (MAP) or an ideal search strategy to guide attention (Acharya, 2019).

2.4.4 Cognitive Architectures

Several models are inspired by cognitive architecture, all based on the idea that there is a time-persistent set of hypotheses on human information processing systems independent of the task. Among these models are Adaptive Control of Thought—Rational or ACT-R proposed by Anderson et al. (1998), Executive Process-Interactive Control or EPIC (Kieras and Meyer, 1997) and Eye Movements and Movement Attention or EMMA (Salvucci, 2001). The information processing architectures form the basis of these hypotheses, guiding the different (search) behaviours. No explicit *utility* functions are used in the models of this class of search approaches. In line with several other classes, the description of *ecology* is provided in the local task environment. In terms of *strategy*, the model uses heuristic control (i.e. if-else conditions) based on task objectives and is introduced by the modeller, making this manually-coded strategy approach one of the weaknesses of this class of search models.

2.4.5 Control Models

The basis of the control model approach is formed based on the assumptions that people are, in principle, rational and able to find optimal strategies to reach the objectives of each task (Butko and Movellan, 2008; Chen, 2015; Hayhoe and Ballard, 2014; Nunez-Varela and Wyatt, 2013; Rao, 2010; Sprague, Ballard, and Robinson, 2007). The reward function in control models defines the *utility*, with the understanding that in the case of the objective of finding a target, the goal is to maximize reward (acknowledging that the scope of the information processing system impacts maximum reward). While the *strategy* space is defined in control models, it is formed as an optimal adaptation to the *ecology* and *utility*. One associated issue with control models is

their ability to scale up.

2.4.6 Deep Learning Models

Aspects of human vision are used in deep learning models for visual search to primarily classify real-world images into often (but not always) well-defined but hard-to-distinguish classes and categories (Leibo et al., 2018; Mnih et al., 2014; Xu et al., 2015). This class of visual search can be seen as an extension to control models, but their use is particularly prominent when other approaches fail to deal with scalability issues.

In the deep learning model, the *strategy* space is created by designing and training a neural network to act as a classifier, with the output being locations in the display. The description of *ecology* is in the local task environment, and the *utility* is defined as the error function used in the model. The constraints in the inputs presented to the neural network form the *mechanism*.

Despite the recent success of the deep learning models and their ability to deal with several real-world problems (including visual task problems), the model relies on the inputs and their accuracy, without which the output quality cannot be reliably trusted. Furthermore, the difficulty of interpreting the model remains an issue.

2.5 Computational Approaches for Visual Impairment

Various computational approaches in the literature aid visually impaired individuals while interacting with technology. Most of these aids use magnifiers and contrast enhancement features. One example is closed-circuit televisions (CCTVs) in digital devices, allowing people with vision impairment to see the world in a magnified digital display (Wolffsohn and Peterson, 2003). Also, computers and smartphones have built-in applications and settings to deliver users with vision impairment a better user experience while interacting with these devices, such as Android and iOS (Holton, 2014; Zhao et al., 2019). Another approach is the development of smart glasses for the visually impaired to support them in daily life activities; it has become increasingly popular in recent years, such as helping with the clarity of text while reading, guiding

the users through voice and giving information about the surrounding and the obstacles in the user's way (Ali Hassan and Tang, 2016; Mukhiddinov and Cho, 2021; Rajendran, Krishnan, and Aravindhar, 2020). Another way to help the visually impaired interact with their surroundings is to widen the visual field using peripheral prism glasses, which can benefit patients with Homonymous hemianopsia (HH) (Bowers, Keeney, and Peli, 2014). The same concept was implemented in Augmented Reality (AR) glasses (Luo and Peli, 2006; Peli, 2001; Vargas-Martín and Peli, 2001) using what is called *vision multiplexing*, which helps bring a minified contour images of the wider visual field into the patients' functional field of view, this technology showed improvement in visual search behaviours. However, the users needed more training and adaptation to benefit from its potential. Additional functionality tailored to a specific Graphical User Interface (GUI) to improve accessibility was also considered by some researchers, for instance, visual cues used to guide users' attention while interacting with a given visual search task (Zhao et al., 2016). Researchers also considered automated design-time of web pages design by combining design-time generation of Web-pages with responsive design, resulting in generating more accessible web pages and more flexibility in adding features allowing a more comprehensive range of people with different time of vision impairments to be able to access web pages easily (Rathfux et al., 2018). Improving the accessibility of modern technology, such as Virtual Reality (VR) and touchscreens, has also been widely explored. In one such work by Zhao et al. (2019), they developed a system called *SeeingVR* which consists of 14 tools plugged into a VR system to help the visually impaired people have a better and more enjoyable experience while interacting with VR. Some of the tools included in SeeingVR are *peripheral remapping* which have the same concept as *vision multiplexing*, *text augmentation* to help improve text contrast, and *text to speech* delivering audio augmentation to the user, the system showed better interaction between the visually impaired participants and the VR; the results showed improved accuracy, and participants spent less to complete the visual tasks. Given that most devices in the modern world use screen touch, improving the accessibility of such technology is essential. Most tech companies such as Android and Apple use screen readers and optional additional devices (e.g. external buttons, Bluetooth keyboard or trackpad) to make it easier for the visually

impaired users to navigate the information on the touchscreen (Apple Inc., 2023; Google, 2023). Researchers, on the other hand, conducted research to improve the interaction between the visually impaired and touch screen, for example, providing haptic feedback to the user using different methods, such as vibration, piezoelectric actuators and magnetic fluids (Awada et al., 2013; Jansen, Karrer, and Borchers, 2010; Poupyrev and Maruyama, 2003; Vidal-Verdú and Hafez, 2007). There is also research on using 3D printers to add tactile feedback to visually impaired users (Kane, Morris, and Wobbrock, 2013; Stangl, Kim, and Yeh, 2014). Although the previously mentioned approaches improve the interaction for the visually impaired, there are also limitations such as cost and efficiency (Mi et al., 2014; Yang et al., 2021), as well as limited social interactions, for example, difficulties using emojis within text (Tigwell, Gorman, and Menzies, 2020; Zhang et al., 2021). Another limitation is the need for the users to learn how to use the technology and find ways to adapt to it (Peli, 2001). Furthermore, partially sighted people often prefer interacting with computers tailored to their visual ability rather than relying on assistive tools built for people with complete blindness, such as tactile devices or screen readers. One of the main solutions proposed for partially sighted people is enlargement (Huang, 2018; Khan and Khusro, 2019). However, this solution is not suited for certain vision impairments, especially the ones involving tunnel vision (Edwards, 1995).

2.6 Computational Models in HCI

Researchers have explored the use of computational models in HCI and proposed recommendations to designers to deliver better UI. This section summarises some of the computational models and their potential applications in HCI. In a recent work, Leiva, Shiripour, and Oulasvirta (2022) trained a deep learning model to predict users' preferences in website aesthetics. The model had high accuracy and could predict the suitability of an interface design for different age groups. Web page designers can then use the results of this model to have an idea of the design preference of the targeted users. Halverson and Hornof (2011) built a computational model of theories of active vision. They used the cognitive architecture EPIC as their foundation for the

new model keeping all its perceptual features the same, and developed the model using eye-tracking data from two different experiments; one experiment concerns mixed density search task, and the second is a consonant-vowel-consonant (CVC) search task. The model predicted that objects closer to the point of gaze are more likely to be perceived, where the features can be identified, that the eyes move after processing a fixated visual stimulus, the eyes tend to move to closer objects, and lastly, the model showed that the information that remained in the memory between fixation is only coarse spatial information. The model developed by Halverson and Hornof (2011) can offer insights into visual interaction in HCI, which can positively aid in user interface design, helping to identify usability issues in the early stages of design and minimising the need for human participants in testing newly developed interfaces.

HCI scientists also used POMDP and ML methods to model human interactions behaviours by understanding and implementing cognitive strategies, which in return can support the design of UI with a better understanding of people's psychology and with less need for experiments including human participants. Knox et al. (2012) illustrate how humans and artificial agents use reinforcement learning to make decisions. Attempting to predict human behaviour in menu search based on search time and eye movements, Chen et al. (2015) demonstrate that human decision-making is an emergent consequence of the chosen environment and the relevant psychological limitations. Furthermore, Chen et al. (2017) modelled the visual behaviours of individuals when interacting with a visual search task formulated as POMDP and solved using RL techniques involving different types of information visualisations. More recent work proposed a reinforcement learning model for a gaze-based target acquisition task (Chen, Acharya, and Oulasvirta, 2021). The task was formulated as POMDP and solved using deep learning methods. The model predicted eye movement strategies as adaptive to some constraints in the environment and in the human visual system. Those constraints are target eccentricity, target size, ocular-motor saccade noise, fixation jitter, and location estimation noise in peripheral vision. The model could predict selection times and the average number of saccades given target size and eccentricity. Another recent work, Li et al. (2023), implemented a computational model that behaves similarly to blind users in a menu selection task. Different interaction methods

were simulated in the model, including swiping, gliding, and direct touch. The model was developed based on the theory of bounded optimal similar to Chen, Acharya, and Oulasvirta (2021). However, unlike Chen, Acharya, and Oulasvirta (2021) model, which focuses on gaze selection, this model was built based on the assumption that users' selection behaviours result from adaption to the user's memory of where the items are located, which is gained through experience and feedback while interacting with the screen reader. They also assumed that visually impaired users aim to minimise their selection time due to their vision limitations. A menu selection task was formulated as POMDP, and then the model was trained using Deep Q-Network (DQN). Their results showed that the model could learn the strategies of those with vision impairment by predicting the effect of the menu length and object locations on selection time for people with vision impairment. The latter model has similarities to the model we are proposing in this thesis. However, the scientists in the latter research did not focus on a specific type of vision impairment, but their model is more general for people with different levels of vision impairments. In this research, we implement our model for a specific type of vision impairment, focusing on how each type observes the environment based on its specific limitations.

Chapter 3

Framework: Computational Model of Vision Impairment

3.1 Introduction

This thesis aims to construct a computational model of visual search behaviours among individuals with vision impairment while interacting with visual search tasks. The model is built based on the theoretical assumptions that human visual behaviours and decision strategies result from adaptation to the statistical environment, the constraints imposed by the human visual system and the reward associated with actions and decisions. Due to the limitations of the human visual system and the partial vision of the visually impaired, we utilise Partially Observable Markov Decision Processes (POMDP) to formulate the tasks presented in this thesis.

This chapter presents a general overview of the model's framework. Subsequently, it explains Partially Observable Markov Decision Processes, our primary implementation method of the visual search tasks. After that, we talk about the general implementation of the external environment, which differs in each experiment, yet have the same concept. Then we talk about how the agent interacts with the given environment and use the Bayesian update to update its internal belief state. Lastly, we explain our main learning methods (i.e. Reinforcement Learning and Deep Reinforcement Learning), which allow the model to learn the visual search strategies

based on the action space and the reward function, which we will explain in more detail in the upcoming chapters.

3.2 Theoretical Framework

An overview of the framework is presented in Figure 3.1. The model consists of a learning agent with internal processes representing the cognitive theory and an external environment representing the task the agent interacts with. The model was constructed based on the model in Chen, Acharya, and Oulasvirta (2021). The external display includes the objects on the screen which remain unchanged. The agent should have an initial probabilistic belief about the environment. Then, it interacts with the environment through action (i.e. fixations) and starts collecting information about the environment, using both the fovea (fixating on the centre of an object) and peripheral vision (getting information about the surrounding objects). The agent's input is a noisy observation of the environment. Our theory in this thesis is that more noise in the observation would cause partial vision in the model, which can represent vision impairment (the noise will be explained in more detail in the upcoming chapters). At every time step, the agent updates its belief about the environment based on the observations it receives. The model attempts to find the target with a stopping searching strategy that depends on either a final reward associated with the final decision (correctly identifying the target) or a pre-defined time cap to minimise the time spent on each trial (the time cap differs in each experiment). Additionally, there is an internal reward, a time cost for collecting information through fixations. The internal reward is essential in guiding eye movements in visual search tasks (Hayhoe and Ballard, 2005). The reward function determines the agent's best policy or strategy. The model utilises Reinforcement and Deep Reinforcement Learning techniques to approximate an optimal strategy that will be learned throughout learning trials. The strategies found by the model are bounded-optimal, meaning that the learning agent finds the optimal policy by adapting to the architecture of the task, the perceptual processes and the reward function. In the next section, we will discuss the theory of interaction implemented in our model in more detail.

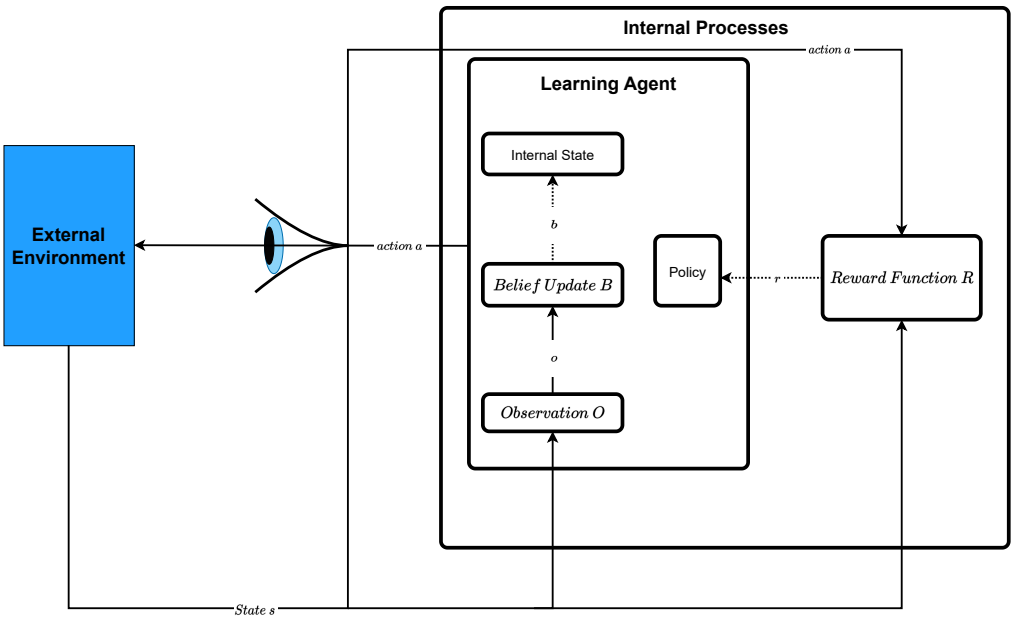

Figure 3.1: Overview of the model's structure. The agent starts by taking an action a in the given environment based on its belief state $b(s)$. It receives observation o based on the state s and updates the belief state b accordingly. It also receives a reward r based on the state s and action a, which updates the policy. The eye icon is only a representation of the agent-environment interaction as the model is a visual search model.

3.3 Agent-Environment Interaction

The model is formulated as a sequential decision process bounded by the limitations of the human visual system. As we are modelling the visual behaviours in people with vision impairment, we assume that the agent interacts with its environment through vision, which attempts to gather information about the environment, send it to the cognitive processes and analyse it to find the target. The model's cognitive system carries out internal procedures and stores the information in a probabilistic belief state, including the agent's beliefs of the target's location and distractors. Based on the belief state, the model chooses to either collect more information about the environment or terminate the search and make a final decision about the target's location. The model might also stop the search due to the pre-defined limited search time. There is an

immediate cost associated with gathering environmental information. However, this might lead to better accuracy in the long term, and therefore, the model has to learn an optimal strategy by finding a trade-off between the cost of gathering more information and its accuracy. In other words, the model needs to learn the best action to take in a given belief state (i.e. when to gather more information and when to gather less information and make a final decision). The model continuously adapts its decision-making strategies by iterating through these steps, gradually improving its ability to locate the target efficiently and accurately.

In this thesis, we use ML methods to solve this problem and find optimal visual search strategies. However, due to the complexity of the model, the solutions are not feasible for human behaviours; therefore, we aim to find approximate strategies that best fit the human data.

As previously mentioned, the visual search tasks addressed in this thesis are formulated as a POMDP framework. In the following section, we will discuss Markov Decision Process (MDP), a foundational element in decision-making and reinforcement learning, from which the concept of POMDP extends.

3.3.1 Markov Decision Processes

Markov Decision Processes (MDP), also known as controlled Markov chains (Altman, 1999; Howard, 1960), provide a mathematical formulation of sequential decision-making problems with stochastic outcomes. MDP is defined as a tuple $\{S, T, A, R\}$, where each element of the tuple is defined by the following:

- $S = \{s_1, s_2, \ldots, s_{|S|}\}$, a finite set of states
- $T(s_i, a, s_j) = p(s_j|s_i, a)$, a finite set of transition probabilities
- $A = \{a_1, a_2, \ldots, a_{|A|}\}$, a finite set of actions
- $R : S \times A \rightarrow \mathbb{R}$, the expected reward of taking action in a given state
- $\gamma \in [0, 1]$, the discount factor

In MDP, the agent interacts with the environment in discrete time steps. The agent observes the current state s_t at any time step t. By taking action a_t, in state s_t, the agent will transition into a new state s_{t+1} in its environment by probability $T(s_t, a, s_{t+1})$. The reward r will, therefore, be received as an indication of how good or bad the action taken in the state. MDPs can be described as a generalisation of (non-controlled) Markov chains with Markovian property indicating that the current state s_t, and the action a_t taken in that state are independent of all previous sets of states. All future interactions will depend only on the current state (Altman, 1999). The Markovian property can be formally defined as follows:

$$P(s_{t+1} = s'|s_t, a_t) = P(s_{t+1} = s'|s_t, a_t; s_{t-1}, a_{t-1}; ....; s_0, a_0) \tag{3.1}$$

For the agent to reach the task's goal, it must take a sequence of independent actions in each given state until it terminates. Upon termination, the discounted cumulative reward should be optimised (maximised or minimised) over a potentially infinite horizon. The agent uses two functions to reach its goal: (1) the control policy, which maps action a_t to take in a state s_t, and (2) the value function, which describes the discounted cumulative reward as a given value of a specific state s, this value is generated by following the policy π. The Bellman equation is used as the value function which represents the expected reward of being in a specific state following a given policy π (see Equation 3.2)

$$V^{\pi}(s) = R(s, \pi(s)) + \gamma \sum T(s, \pi(s), s')V^{\pi}(s') \tag{3.2}$$

where γ is a discount factor $0 \leq \gamma < 1$ that is used to determine the future reward, $V^{\pi}(s')$ is the value function for state s' and $\pi(s)$ is the policy the agent follows when in state s. The stochastic policy is calculated as follows:

$$\begin{aligned} \pi(a|s) &= P(A_t = a|S = s) \\ V^{\pi}(s') &= \sum_a \pi(a|s) \sum_{s \in S} T(s, a, s')[R(s, a) + \gamma V^{\pi}(s')] \end{aligned} \tag{3.3}$$

Usually, the agent will search for the optimal policy π^* by finding the maximum value for a state under all policies. This can be defined as follows:

$$\begin{aligned} V^*(s) &= \max_{\pi} V^{\pi}(s) \\ V^*(s) &= \max_{a \in A} [R(s, \pi(s)) + \gamma \sum_{s \in S} T(s, \pi, (s), s') V^{\pi}(s')] \end{aligned} \tag{3.4}$$

The MDP framework was used in various tasks (Altman, 1999; Littman, 1996). In MDPs, the agent has full knowledge of its environment. Nevertheless, this cannot represent real-world scenarios accurately, because when humans interact with their environment, there are always limitations due to the nature of their mind and senses. This makes the environment not fully but partially accessible. This thesis focuses on modelling visual behaviours, particularly those affected by vision impairment, which are constrained by various factors, making the environment partially accessible. For this reason, in this research, we implement a Partially Observable Markov Decision Process (POMDP), which extends from MDP (Altman, 1999; Lovejoy, 1991; Monahan, 1982).

3.3.2 Partially Observable Markov Decision Processes

A Partially Observable Markov Decision Process (POMDP) is a generalisation of the MDP framework. It is a model for deciding how to act in "an accessible, stochastic environment with a known transition model" (Russell and Norvig, 2016). In POMDP, the agent receives probabilistic observations instead of observing the current state (Kaelbling, Littman, and Cassandra, 1998). POMDPs can be defined by the aforementioned 7-tuple $(S, A, T, R, \Omega, O, \gamma)$, where S, A, T, R and γ have the same definition as in MDP, and the rest (i.e. Ω and O) are defined as follows:

- $\Omega = \{o_1, o_2, \ldots, o_{|O|}\}$, represents a finite set of observations that can exist in the agent's environment
- $O : S \times A \rightarrow \prod(\Omega)$, represents the observation function, which is the probability distribution over possible observations given an action and a resulting state. It is written

as $O(s', a, o)$, which is the probability of making an observation o, given that the agent took action a and transitioned into state s'.

Therefore, rather than perceiving the current state, the agent will receive observations about possible states instead. Consequently, the agent chooses an action under this uncertainty about its current state. The state, in this case, is called a *Belief State*, which is represented as $b(s)$ (Littman, 1996). A method was proposed by Littman (1996), explaining if an agent takes action in its environment, it will receive an observation about the environment. Therefore, it will update the probability distribution of all the observations, and as a result, its belief of the current state will be updated. POMDPs can better explain human interactions in a given environment in several scenarios, including visual search, interactive search and sense-making (Oulasvirta, Bi, and Howes, 2018), which make them a better fit for the proposed research.

3.4 Internal Processes

This section will discuss the internal processes in the model presented in Figure 3.1. Mainly, we will delve into more details about our belief update method, the reward function, eye movements representation in the model, and the spatial uncertainty.

3.4.1 Belief State

As the state is not fully observable, the model maintains a belief state of the current state through taking action, a_t, and receiving an observation, o_t. The belief state differs in each experiment. Therefore, it will be explained in more detail in the upcoming chapters. In order to update the belief state in the model, we used two methods of belief update: the Bayesian update and the Kalman filter, which will be explained in the next subsection.

Bayesian Update

At every time step, the model takes an action, receives an observation and updates its belief about the environment. The process of updating the belief state is called *state estimator* (Kaelbling,

Littman, and Cassandra, 1998), and it works by computing a new belief state, b', based on an old belief state b, an action a, and the observation o that the agent received as a result of action a. The belief of the agent is a probability of the current state. Therefore the agent won't need any information from previous states, this means that the expected reward and the action in the next state will depend only on the previous belief state, which indicates that the state has the Markov property. As the belief is a probabilistic state, it can be updated using the basic Bayes' rule (Kaelbling, Littman, and Cassandra, 1998) as shown in Equation 3.5:

$$P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)} \tag{3.5}$$

where $P(A \mid B)$ is the posterior probability of A given B, $P(B \mid A)$ is the likelihood, $P(A)$ is the prior and $P(B)$ is the marginal distribution factor. The Bayes' rule is a rigorous method that analyses evidence within the prior state and updates its existence belief based on probabilistic observations (Stone, 2013). The following equation shows the process of Bayes' rule updating process (Kaelbling, Littman, and Cassandra, 1998):

$$\begin{aligned} b'(s') &= P(s' \mid o, a, b) \\ P(s' \mid o, a, b) &= \frac{P(o \mid s', a, b)P(s' \mid a, b)}{P(o \mid a, b)} \\ P(s' \mid o, a, b) &= \frac{P(o \mid s', a) \sum_{s \in S} P(s' \mid a, b, s)P(s \mid a, b)}{P(o \mid a, b)} \\ b'(s') &= \frac{O(s', a, o) \sum_{s \in S} T(s, a, s')b(s)}{P(o \mid a, b)} \end{aligned} \tag{3.6}$$

where $P(o \mid a, b)$ is the marginal distribution factor, $b(s)$ represents the previous state belief, $O(s', a, o)$ is the observation function, $T(s, a, s')$ is the transition function and finally $b'(s')$ is the new belief state.

Kalman Filter

The *Kalman filter* is an optimal recursive estimator that derives parameters from uncertain, indirect or inaccurate observations. It is based on Bayesian principles, which involve updating the belief state based on noisy observations. The Kalman filter represents the posterior probability as Gaussian distribution with mean μ and variance σ (Faragher, 2012). The initial belief states b_0 is a Gaussian distribution with mean μ_0 and variance σ_0^2. After taking action a, the model gets a noisy observation o as Gaussian distribution with mean μ_1 and variance σ_1, with the assumption that the transition and observation probability is linear at each iteration, the model updates the mean (see Equation 3.7), and the variance (see Equation 3.8), which then updates the belief to a new estimation of the current belief state (Acharya, 2019; Faragher, 2012).

$$
\begin{aligned}
k &= \frac{\sigma^2}{\sigma^2 + \sigma_0^2} \\
\acute{\mu} &= \mu + k(\mu_0 - \mu) && (3.7) \\
\acute{\sigma}^2 &= \sigma^2 - k\sigma^2 && (3.8)
\end{aligned}
$$

where k is the *Kalman gain*, representing a weight parameter used in the recursively updated process in the equations. It is used to give a weight that indicates the importance of the current observation to be used to estimate the current state.

3.4.2 Reward Function

In the model, the reward function consists of two factors: (1) cost associated with actions collecting more information about the environment, and (2) immediate reward for making a decision and resulting in a transition to another belief state. The reward function differs based on the nature of each experiment; therefore, it will be explained in more detail in the upcoming chapters.

3.5 General POMDP Tasks Formulation

This section will provide a general POMDP formulation for all the visual search tasks used in the research. Due to the similarities in the POMDP formulation in each task, this formulation can be described as a base for all the tasks, allowing us to define and clarify each element of POMDP is implemented in the model. Subsequently, in each experiment chapter, we will explain the differences in more detail.

General POMDP formulations for all tasks are as follows:

- State S

  At each time step t, the environment is in a true state s_t. A state s_t in the model differs based on the task in the human experiment. Generally, a state is a list of multiple tuples, each corresponding to an object. Within each tuple, there is information about the object, such as the object's coordinates and features. However, it is important to note that the features always determine the object's potential to be a target regardless of the task. This feature is referred to as the object's *targetedness*, which is a value representing how close the object is to being a target, where a target value is 1 in all the experiments.

- Action A

  The agent takes action a_t at every time step t to either collect information from the environment or to make a final decision. Each human experiment has different set-up and instructions for the human participants. For this reason, this will be clarified in more detail in each of the experiment chapters.

- Reward function $R(S, A)$

  The subsection 3.4.2 explained a general definition of the reward function in the model. After taking action a_t at time step t. The agent should receive a reward that depends on the state and the agent's action, $r(s_t, a_t)$. The agent should maximise accumulated reward to optimise the learning and find the target.

- Transition function $T(S_{t+1}|S_t, A_t)$

After taking action a_t at state s_t, the environment should transition to a new state s_{t+1}. However, in this task, given the nature of the environment, the state (i.e. the position of the target and distractors) remains unchanged throughout the time steps in each trial.Hence, $T(S_{t+1}|S_t, A_t)$ equals to 1 only when $S_{t+1} = S_t.T(S_{t+1}|S_t, A_t)$, or 0 otherwise.

- Observation O and observation function $O = f(S, A)$

  After every fixation action, the agent receives a new observation about the target location. The observation is a function of state and action, $o_t = f(s_t, a_t)$. The observation in Experiment 1 (Chapter 4), concerning glaucoma patients, is only dependent on the distance between an object location in the peripheral visual field and the fixation location (i.e. eccentricity). The object location has a spatial uncertainty of peripheral vision, represented as noise (i.e. standard deviation), which is linearly dependent on the distance of the observed object from the centre fixation point (Chen, Acharya, and Oulasvirta, 2021; Michel and Geisler, 2011). An additional factor to spatial uncertainty in the observation o_t is the crowding effect, which was added as additional noise in the observation. The crowding noise depends on the spatial distance between objects (i.e. flanker distance). We conducted experiments using two different noise types, Gaussian and what we refer to as *Blending*. The crowding noise was only implemented in the experiments related to patients with AMD and cataracts (Chapter 5 and 6) as it is more relevant to these experiments' setup. More details about both the distance and crowding noise will be explained in the corresponding chapters. In all the experiments, we added an additional hyper-parameters to find the best fit for the human data. We refer to this hyper-parameter as *noise level*, which is used in a similar way as $P_{spatial}$ in (Chen, Acharya, and Oulasvirta, 2021).

- Discount rate γ

  The discount rate $\gamma \in [0, 1]$ determines the future reward. The model receives a reward of $r(s_t, a_t)$ at each time step t, and the agent should learn an optimal strategy that maximises the long-term expected reward of $R = E\ [\ \sum \gamma^t r(s_t, a_t)\ ]$.

- Belief update

As the state is not fully observable, the model maintains a belief state $b(s)$ of the current state through taking action, a_t, and receiving an observation, o_t. Initially, the initial belief is assumed to be a uniform distribution representing the unknown target location. After taking action a_t, the agent receives a noisy observation o_t, which can provide an estimate of the target location. The belief consists of a variance σ and a mean μ, which consequently can be updated using the Kalman filter explained in subsection 3.4.1. The belief state differs for each task and will be explained in more detail within each chapter.

After creating the environment and formulating the tasks as POMDP, we input the environment into the RL agent. RL methods are explained in the following sections.

3.6 Reinforcement Learning

Reinforcement Learning (RL) (Sutton and Barto, 2018) is an ML approach where the agent is not explicitly informed about what action to take; instead, it should learn to choose actions depending on the state it is in by interacting with the environment to maximise a total future reward. One of the challenges attributed to RL is that the agent should proceed with the learning while balancing exploration and exploitation. By only exploiting, the agent could maximise a future reward. However, it should still explore the environment to find better actions that would otherwise be missed or ignored. If the environment is stochastic, the agent should try to take action multiple times to have a reliable estimation of the expected reward. RL has four elements: *policy, reward, value function*, and a *model* of the environment, which can sometimes be created. The problems that can be solved using RL are defined as the Markov Decision Process (MDP), where the states can be formally identified and given to the agent as input information of the environment. Based on (Sutton and Barto, 2018), RL is considered a learning problem rather than a learning method. Therefore if a method can solve it, it is considered a reinforcement learning method.

In this thesis, specifically in the experiment in chapter 4, we use Q-Learning, which we will discuss in more detail in the following subsection.

3.6.1 Q-Learning

Q-Learning (Watkins, 1989) is an off-policy temporal difference (TD) control algorithm. It is a model-free form of reinforcement learning, which allows the agent to act optimally in MDP by recursively finding the optimal state-action value $Q^*(s,a)$ using the following equation:

$$Q^*(s,a) = Q^*(s,a) + \alpha(r + \gamma \max_{a \in A} Q^*(s',a') - Q^*(s,a)) \tag{3.9}$$

where $Q^*(s',a')$ represents the optimal state-action value for the next state s' and action a', r is the reward that the agent receives when taking action a in state s, α is a learning rate where $0 < \alpha \leq 1$, and γ is a discount factor where $0 < \gamma \leq 1$.

Q-learning algorithm consists of a set of states S and possible actions A associated with those states. The agent learns through the state-action values (i.e. $Q(s,a)$ or Q values) created in the Q table based on actions in the given states. First, all Q values are set to an arbitrary value such as 0. Then, at every time step t, the agent takes an action a_t, receives a reward r_t and transitions into a new state s_{t+1}. The Q values get updated throughout exploring state-action space until convergence, meaning the agent learned the policy. The greedy algorithm $\epsilon\text{-}greedy$ can be used to determine the actions taken in the Q table, where actions are selected greedily based on the highest $Q(s,a)$ value with a probability of $1-\epsilon$, otherwise, the agent will select random action. The actions in the states are then evaluated in the model based on their next state. Thus, if an action leads to a non-targeted state, the Q value is reduced for that action in the given state, leading to other actions having higher values and, therefore, a higher probability of being chosen. Whereas if the action leads to a targeted state, then a higher value would be given to this action. It is also important to note that when taking action and updating the Q value, the previous state-action value has been updated in the Q table; this is known as one-step *lookahead*, which results in policy convergence (Sutton and Barto, 2018). The Q-learning algorithm faces two main challenges: exploration and exploitation and state Space Explosion.

- **Exploration and Exploitation**

  Exploration and Exploitation is one of the challenges in RL because of the trade-off

between the two. In the Q-learning algorithm, Q values are used to indicate the rewards for given actions in a given state, and the policy function is used to indicate the best action to choose based on the reward, so the agent is *exploiting* by choosing the best action. However, this indication does not mean that the best action is always truly the best as by only *exploiting*, there is always a possibility that there could be a better action that the agent has not yet discovered. Therefore, the agent has to *explore* other options in the environment by taking random actions. Nevertheless, also by only *exploring* and not *exploiting*, there is a potential that the policy will not reach convergence, and the agent will keep choosing actions randomly. Balancing between the two is challenging; methods such as ϵ-$greedy$ can support the balance and, therefore, the convergence process (Sutton and Barto, 2018).

- **State Space Explosion**

  In the Q-learning algorithm, there is a process of mapping values of all states encountered to available actions in the Q table. A state space explosion could happen when the Q table has a vast state space (i.e. a massive number of rows in the Q table). An example was given by Acharya (2019) when using a Q-learning algorithm to learn how to play the Atari game, and the state is the current screen with an image size of 84×84 pixels, which, if converted to a grey scale it will be represented with 256 values. Therefore, for the Q-learner to learn through the Q table, it would need to create 256^{7056} states (i.e. rows in the Q table), which would cause the Q-learner to take a 'lifetime' to reach convergence. Furthermore, the process would have created rows that would never be visited, rendering them useless in the learning process.

3.7 Deep Reinforcement Learning

Deep Reinforcement Learning (DRL) combines reinforcement learning methods such as Q-learning and proximal policy optimization (PPO) algorithms with deep learning techniques such as convolutional neural network or CNN (François-Lavet et al., 2018).

DRL is powerful and capable of solving problems that used to be unsolvable; the use of these approaches has, for instance, allowed researchers to build decision-making models that play Atari games achieving human behaviour by only using images (i.e. pixels) as input. DRL demonstrated learning the end-to-end process of robot control policies from signals to motors based on deep neural networks search policy methods which would have been otherwise hard to achieve solely with RL methods (Egorov, 2015).

Arulkumaran et al. (2017) demonstrated that DRL technique can scale in high-dimensional decision-making problems. Furthermore, DRL is expected to improve human vision modelling by training end-to-end systems and combining CNN with recurrent neural networks (RNN) that can decide where to look by applying RL methods (LeCun, Bengio, and Hinton, 2015). In this research, specifically in the experiments reported in Chapters 5 and 6, we trained the model using DRL algorithm Proximal Policy Optimization (PPO).

3.7.1 Proximal Policy Optimization

Proximal policy optimization algorithms (PPO) is a recently built policy gradient method for RL by OpenAI team (Schulman et al., 2017). Unlike the standard policy gradient methods performing single gradient updates per data sample, PPO was built as a powerful method that can select minibatch updates for multiple epochs. The main benefit of PPO is that it provides a simpler implementation of trust region policy optimization (TRPO), which is more general with improved complexity (empirically) based on empirical evaluation. Upon testing the PPO method, Schulman et al. (2017) have demonstrated its outperformance over other policy gradient methods in various tasks such as simulated robotic locomotion and playing the Atari game. Our model is quite complex and needs a powerful method to train, especially for the experiments in Chapters 5 and 6; therefore, we make use of PPO to perform the visual search tasks of vision impairments.

3.8 Workflow

The process of generating the results of the model is presented in Figure 3.2. First, we collected suitable research papers in the literature that conducted visual search experiments on human participants with vision impairment (see Subsection 3.8.1 for a more in-depth explanation on how those studies were selected). Second, we developed the model by generating the environment, which is the formulation of the selected tasks as POMDP and building all the relevant functionalities. In parallel, we implemented a grid search to find the best noise level that fits the human data. In all the experiments, we try to fit noise levels to the results of those with normal vision as well as those with vision impairment. Third, we train and test the ML model given the POMDP and the hyper-parameters. Lastly, we validate our results against the human data and record the results. This process is iterative and can require adjusting the model's development until finding the best fit for the human data. As the case in the model generated by Chen, Acharya, and Oulasvirta (2021), the model is not trained using human data, but it was trained based on an artificial environment that represented the task in the chosen human experiment.

3.8.1 Selection of Human Experiments

The process of selecting suitable human experiments which can be implemented in our model considered a number of factors. These factors are listed below:

1. **Year of publication:** The year of publication was considered as a main factor in the selection process of the human experiments. We have prioritised more recent papers as there is a higher chance of their human data being stored and thus a higher availability to researchers; this is crucial for our validation step (see Figure 3.2). Additionally, with recent publications, there is a higher chance that the researchers are still active in their respective fields. This increases the likelihood of a potential collaboration, thus enabling us to discuss the research findings with the vision impairments experts. Furthermore, human experiments conducted more recently may reflect the latest understanding in the

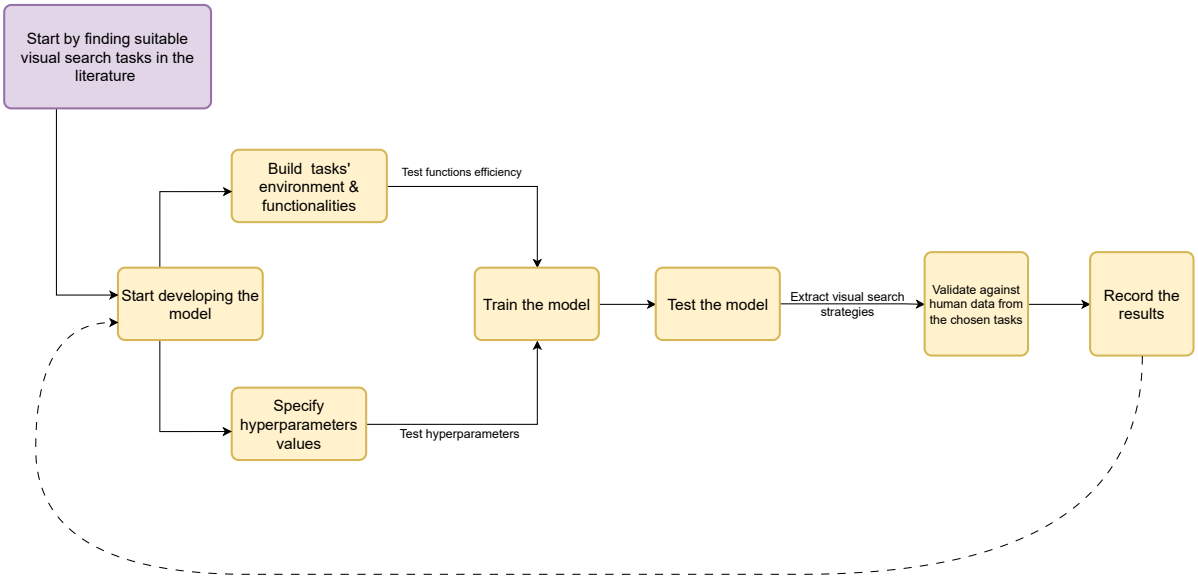

Figure 3.2: This figure shows a simplified illustration of the iterative process of developing the model, extracting visual search strategies, comparing them against human data, and modifying the model again until achieving good results and strategies that are approximately close to the human participants

field of vision impairments, making their data more useful and reliable for the research. The older papers have a higher chance of their data being destroyed and the researchers are no longer active in their respective research areas.

2. **The nature of the visual search task:** The choice of the experiments considered the nature of the visual search task the participants were asked to interact with. This is important due to different factors: (1) if the task is suitable for our methodology and, by extension, implementation; in other words, we have to ask if the task can be formulated as a POMDP. For example, some studies considered finding a target through free viewing (e.g. (Asfaw et al., 2018)), which would be more challenging to formulate as a POMDP than those that had a definite number of objects on the screen. (2) Does the task consider the specific limitation of the type of the vision impairment? For example, those with glaucoma experience limitations in their peripheral vision, therefore, the focus would be on recognising objects located within higher eccentricity whereas people with AMD experience loss in the centre of the visual field which will cause increased uncertainty when objects are close together, particularly within the foveal visual field. In the case of those with cataracts, the uncertainty would be in the whole visual field due to the blurriness they experience.

3. **The outcome measured:** Lastly and most importantly, is the outcome measured in the human experiments, such as the accuracy, the number of fixations, the saccade latency and the search time which can reflect gaze interaction. Our ultimate aim is to build a model that can give insights and predictions of these measures which can reflect the adaptation of humans with different types of vision impairments. Finding the suitable research papers and being able to predict these measures in our models can give the model the potential to help design more accessible user interfaces in a future work (Majaranta et al., 2019; Plopski et al., 2022; Wolf and Ueda, 2021).

3.8.2 Implementation

In this thesis, there are two types of implementations to build the model of vision search in vision impairment:

1. Q-learning implementation: For the visual search task in Chapter 4, the Q-learning method was hard-coded in the model. The chosen hyper-parameters will be explained in the corresponding chapter.

2. Deep Reinforcement Learning implementation: We used OpenAI Gym (Brockman et al., 2016) and built a suitable environment to train the model and make use of the built-in libraries of deep reinforcement, PPO algorithm provided by Stable Baselines[1]. This method was used in Chapter 5 and Chapter 6; the choice of hyper-parameters will be clarified in the corresponding chapters.

Next, in Chapter 4, we will discuss the implementation of a visual search task for glaucoma patients.

[1] https://github.com/DLR-RM/stable-baselines3

Chapter 4

Experiment 1: A Model of Saccadic Choice Task with Glaucoma

4.1 Introduction

In this chapter, the model was applied to the first experimental visual search in this thesis. The human experiment involved participants with glaucoma, and it was done by Boucart et al. (2020). The main goal of the task was to check signal processing behaviours in glaucoma at large eccentricity. The scientists used the two-alternative forced choice (2AFC) method, introducing time limitations and uncertainty constraints to the human participants. Using the 2AFC method to achieve both measurable and comparable results has shown that the method can optimise the humans' mechanisms in decision-making throughout evolution (Bogacz et al., 2006). The eccentricities in the human experiments vary from small (i.e. 10°) to very large (i.e. 80°). The task required participants with normal vision and glaucoma to find a target (i.e. human face) located at different eccentricities. In this research, we assume that human strategies are emerging due to the statistical aspects of the environment, the limitations of the human visual system, the actions involved and the rewards associated with the actions. As glaucoma patients usually experience loss in peripheral vision, therefore we represent the uncertainty in the peripheral vision as Gaussian noise that increases with eccentricity. The task then was

formulated as a sequential decision-making task bound by the human visual system limitations and the environment. Therefore, it was formulated as a POMDP and solved using Q-learning.

The model predicts similar accuracies and saccade latencies to the human participants for both normally sighted people and people with glaucoma for the given problem.

4.2 Background

Glaucoma is a result of neuropathy, which leads to progressive loss of ganglion cells and visual field loss, as was stated by The National Eye Institute (2022) and Vingrys (2000). Glaucoma is considered one of the main reasons for blindness that mostly affects older people globally (Ramrattan et al., 2001; World Health Organization, 2022). Boucart et al. (2020) stated that the visual field loss in glaucoma causes retinotopic loss in contrast sensitivity in the visual field. Contrast sensitivity loss starts by affecting the peripheral vision, but studies have shown that it can also result in loss in the central vision (Hood, 2017; Jung et al., 2019; Lenoble, Lek, and McKendrick, 2016). Furthermore, Boucart et al. (2020) emphasises that in glaucoma, vision quality worsens first in the signal passed through the part where the peripheral visual field is represented in the cortex, and this is due to the peripheral nerve fibres being the first affected. The effects of glaucoma on the visual field and eye movements have been studied through different experiments using eye-tracking tools to detect the visual behaviour and decision-making strategies and how they are affected by the visual field loss caused by glaucoma (Kasneci, Black, and Wood, 2017). Everyday tasks were used in many visual search tasks experiments, such as driving (Crabb et al., 2010; Haymes et al., 2008; Johnson and Keltner, 1983), reading, walking, shopping (Nelson, Aspinall, and O'Brien, 1999) and face recognition (Glen, Smith, and Crabb, 2013). Peripheral vision perception was also studied on people with glaucoma using driving simulations on laboratory-based experiments (Kasneci, Black, and Wood, 2017).

Despite the low spatial resolution of peripheral vision (Westheimer, 1982), peripheral perception plays an essential role in providing important information about the surrounding environment (Thorpe et al., 2001; Velisavljević and Elder, 2008; Wiecek et al., 2012); it can also

direct the saccadic eye movements towards the target, bringing it to foveal resolution through fixation (Torralba et al., 2006; Wiecek et al., 2012). The results of experimental studies on normally sighted people showed that the information with low resolution in peripheral vision can guide the eye to targets in the visual search tasks (Torralba et al., 2006). Detecting objects in a wide visual field is crucial in everyday activities such as detecting crossing pedestrians, moving cars or facial expressions (Bayle et al., 2011); it is also essential in visual search tasks such as spatial navigation and spatial memory, which involve walking or driving as it helps to recognise danger and acting accordingly beside giving better knowledge and information about the environment (Boucart et al., 2020). Additionally, studies on normal vision showed that the low resolution of peripheral vision is efficient for scene gist recognition at eccentricity higher than 60° (Boucart et al., 2013; Loschky et al., 2019). Some experimental studies have examined the peripheral perception in people with glaucoma. They included using a driving simulator to test the safety of driving in patients with glaucoma, where they placed the hazard object within 30° eccentricity (Kübler et al., 2015). The results showed that three out of six patients were considered safe drivers. The unsafe drivers experienced shorter saccade amplitudes, gaze bias to the right, and their focus was only one point ahead, which was believed to be due to tunnel vision reported by Engström, Johansson, and Östlund, 2005. Another experiment was conducted by Szlyk et al. (2005) using a driving simulator with a wider visual field of 160° horizontal viewing. The results showed that people with glaucoma experienced a higher chance of accidents in the given simulator than in the control group. In another work, Boucart et al., 2020 tested the visual behaviour of people with glaucoma when given a target within a very large horizontal viewing up to 180° in order to provide further clinical assessment for patients with glaucoma. As part of this work, a model was built based on the latter study, to predict the visual behaviours and strategies of patients with glaucoma.

4.2.1 Visual Search Behaviours in Glaucoma

The effect of glaucoma on the visual field and eye movements has been studied in different experiments using eye-tracking tools to detect the visual behaviour and decision-making strategies

and how they are affected by the visual field loss caused by glaucoma (Mao et al., 2021). In one such study by Smith, Glen, and Crabb (2012), eye movements were studied during visual search in patients with glaucoma. Participants were given photos of a real-world scene and asked to find an object. Behaviours of eye movements were compared between people with glaucoma and people with a similar age range with normal vision. The results showed that people with glaucoma made fewer saccades, and there was no difference in average saccade amplitude, whereas patients with severe visual field defects made less eye movements during the task. In another study by the same group Smith et al. (2012), glaucoma eye movements were investigated where participants were asked to view real-world images freely. The findings showed that patients with glaucoma made fewer saccades with longer periods of fixations; it was suggested that the reason for this behaviour is related to the loss of the centre vision, as people with normal vision would make more saccades and fixations using their central vision as it is the area of highest quality of vision. This aligns with the results reported by Henderson (2003). In another study, Burton, Smith, and Crabb (2014) tried to find the relationship between reading speed and eye movements in people with glaucomatous visual field (VF) defects. The authors compared eye movements in different experiments, such as providing the participants with short paragraphs or a single word to spot a fault in the word. The results showed that people with glaucoma reported more problems while reading, especially those with advanced VF defects. Furthermore, the results demonstrated that people with glaucoma did what the authors referred to as *text saturation*, where a person fixates on every letter when trying to read the text. This was not observed as much in the control group from the same age range. In this particular research, there was no difference in reading speed between the two groups. The reasons are thought to be linked to the control group being from the same age range and that both groups were given a large text. In another experiment, Kanjee et al. (2012) tested the saccade reaction time among patients with glaucoma and a control group from the same age group (i.e. 40-80) when given a pro-saccade step task. Their findings showed that patients with glaucoma of all levels (early, moderate or advanced) had a delayed reaction time to initiate the search, which did not happen with any of the participants in the control group. Glen, Smith, and Crabb (2013) investigated

saccadic eye movements and face recognition performance in patients with central glaucomatous visual field defects. In their work, face recognition task was tested on three different groups: the first group consisted of 28 people with glaucoma having significant 10° defects in both eyes (10-2 MD p < 1%) where saccadic amplitude was associated with face recognition task (larger saccades), the second group was formed of 23 who did not have significant 10° defects (10-2 MD p > 1%) where the association was less between saccadic amplitude and face recognition, and finally, in the control group, there was no association. The results showed that, in terms of the ability to identify faces in the given task, patients with 10° defects performed worse than those with no significant 10° defects, whereas the control group performed with the highest accuracy. In another visual task and, this time, focusing on the driving scene, Crabb et al. (2010) investigated eye movements in patients with glaucoma when viewing a driving scene. It was shown that patients with glaucoma needed more fixations and more saccades and that some hazards were missed.

In the next section, we delve into more detail of the task we chose for this research, followed by the formulation of the task as a sequential decision-making task.

4.3 Experimental task

This section explains the human experiment done by Boucart et al. (2020) in more detail, including the settings and stimulus provided to the human participants. Next, it illustrates the formal POMDP formulation, implementation, and results.

4.3.1 Visual search task at large visual eccentricities in patients with glaucoma

Measuring visual search in glaucoma requires an experimental design that allows testing of the entire visual field. Accordingly, Boucart et al. (2020) designed an experiment that permitted large eccentricity testing using a panoramic screen of 180°. Participants were asked to find a target (i.e. human face) through a two-alternative forced choice (2AFC) task by saccade towards

the location containing the target. The target is located on either the right or left side of a fixation point in the centre, within eccentricities varying between 10 - 80°. Studies have shown that using a human face as a target for normal vision caused faster saccadic response than other stimuli, such as animals at small eccentricities 8° (Crouzet, Kirchner, and Thorpe, 2010). This was extended to the whole visual field (i.e. 80°) in young participants (Boucart et al., 2016). The experiment included 14 glaucoma participants with 0.1 LogMar visual acuity or better in each eye, and a score over 26/30 on the Mini-Mental State Examination (Folstein, Folstein, and McHugh, 1975) for people over the age of 60 to indicate that there are no major cognitive deficits. It also included 14 normally sighted participants with visual acuity of 8/10 or better in each eye on the Monoyer scale (i.e. 0.1 to 0.0 LogMar), and a score above 26/30 on the Mini-Mental State Examination for people over the age of 60 who had no history of neurological or psychiatric disease. The stimuli included 275 pictures of human faces in grey level and 400 pictures of objects such as animals, vehicles, plants and buildings (see Figure 4.1 for an example of the stimuli). All pictures were represented in full contrast given that the contrast decreases in higher eccentricity (Cannon, 1985). Images were presented at the viewing distance of 2.04m with resolutions of 512×512 pixels, covering $18° \times 18°$ of the visual angle. Participants were initially asked to fixate on a black cross, defined as the first fixation point in the experiment. This was presented for $1s$ and was located in the centre between two images. Subsequently, a blank screen was shown for $2s$, and the images appeared on both sides. The participants were asked to only saccade towards the target, located within eccentricities varied between $10°, 20°, 40°, 60°$ and 80°. The eccentricities and location of the target were chosen randomly at each trial and were equally represented. Only saccade latency and accuracy were recorded, but not the saccade amplitude, given that the scene camera used in the experiment only covered up to 40°. The authors expected that patients with glaucoma would have longer latencies and less accuracy as the eccentricity increases. The results in Figure 4.2b showed that patients with glaucoma had much slower saccade latency than the control group. For both patients and control, the saccade latency slowed at higher eccentricity (with patients being generally slower). Also, both patients and control exhibited better accuracy at lower eccentricity. However, for the patients' group,

the accuracy dropped significantly after eccentricity 40° compared with the control group (see Figure 4.2a); two patients had a drop in performance at eccentricity 40° (see Figure 4.3a). It was stated that the drop in performance at larger eccentricity is due to the loss of optic nerve fibres responsible for transferring information from the peripheral vision. Also, people with glaucoma experience a reduced ability for more exploration in the visual field (Asfaw et al., 2018; Smith, Glen, and Crabb, 2012). As mentioned in a free viewing task done by Smith, Glen, and Crabb (2012), people with glaucoma have a more restricted elliptical region (by 23%) than the control participants. Additionally, degeneration of retinal ganglion cells in glaucoma can cause cortical remapping after reduced neural input to the visual cortex. This is reflected in the reduction of accuracy at 40° eccentricity in the experiment conducted by Boucart et al. (2020). (Figures 4.3 and 4.4 are provided for further insight into each individual's performance)

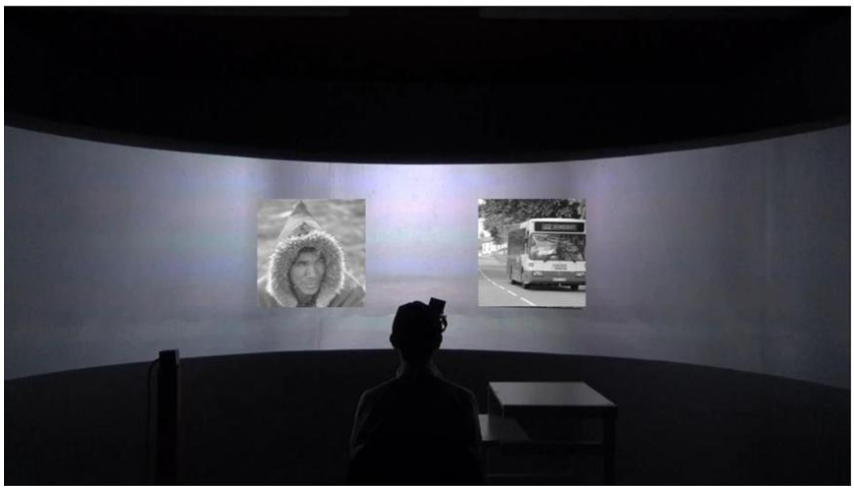

Figure 4.1: An example of the display used in the saccadic choice task experiment by Boucart et al. (2020) as shown in the original paper. (image used with permission)

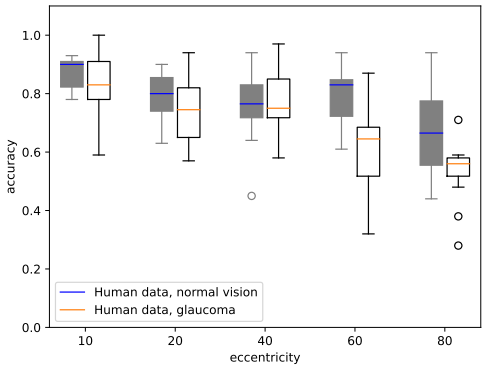

(a) Human data distribution, showing the difference between controls and glaucoma participants (accuracy).

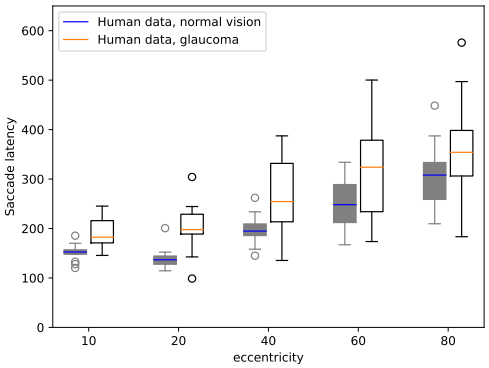

(b) Human data distribution, showing the difference between controls and glaucoma participants (saccade latency).

Figure 4.2: A replicated version of the results in (Boucart et al., 2020), showing the distribution of the human data including the median and outliers (results provided by the author Boucart et al. (2020)).

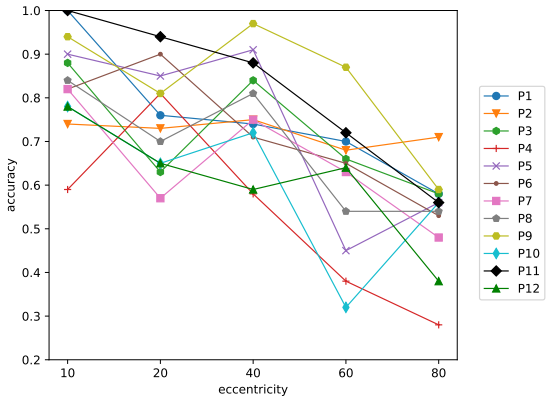

(a) Human data that shows the accuracy for each participant with glaucoma.

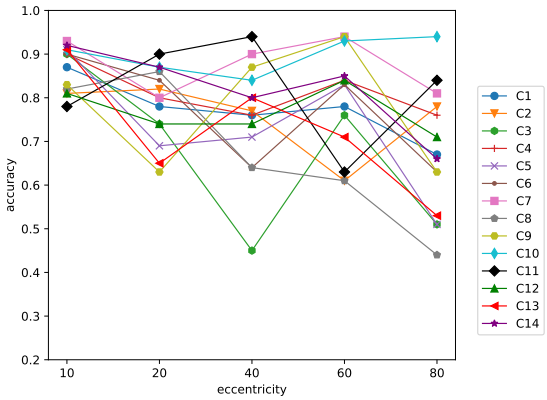

(b) Human data that shows the accuracy for each participant with normal vision.

Figure 4.3: Comparison between accuracies in the human data for each participant in the human experiment as provided by Boucart et al. (2020) where (P) refers to Patient and (C) refers to Control. It can be seen that the relationship between the accuracy and the eccentricity is not always linear. For example, P11 in glaucoma participants had a near-linear performance. whereas P5 and P10 had a better performance at eccentricity 40° and then a drastic drop at eccentricity 60° with their accuracy improved at eccentricity 80°. Furthermore, it is noticeable that the accuracy for P6 and P4 dropped drastically at eccentricity 40°.

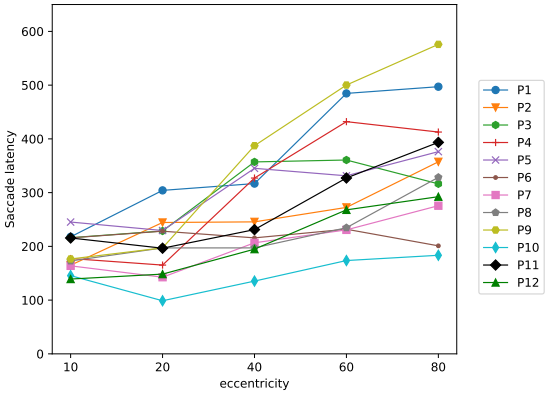

(a) Human data that shows the saccade latency for each participant with glaucoma.

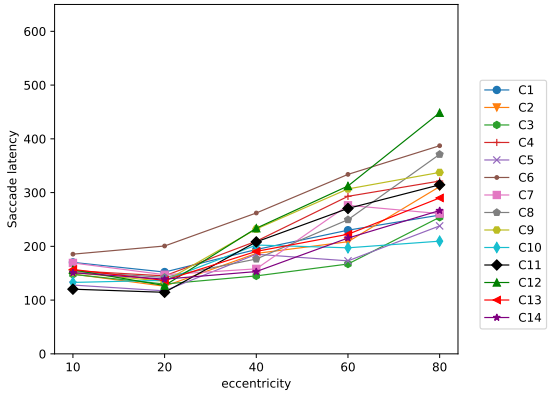

(b) Human data that shows the saccade latency for each participant with normal vision.

Figure 4.4: Comparison between saccade latencies in the human data for each individual participant in the human experiment as provided by Boucart et al. (2020) where (P) refers to Patient and (C) refers to Control. The relationship was mostly linear in the saccade latency for the human participants, especially for those with normal vision. However, some participants showed less saccade latencies at an eccentricity 80°, for example, the case of P3 and P4 in glaucoma patients.

4.4 Theory

In this chapter, we implemented a model of peripheral perception. As explained in Chapter 3, in our theory, we assume that human visual behaviours and decision strategies are the results of adaptation to the statistical environment (i.e. target distance), the constraints imposed by the human visual system (spatial uncertainty due to the peripheral vision) and the reward associated with actions and decisions. The decision-making task based on peripheral observations (Boucart et al., 2020) is considered a sequential decision-making problem that can be formulated as POMDP and solved using Reinforcement Learning (RL) methods to find optimal search and decision strategies. The stopping-search strategy in the model was based on the assumption that people will receive a positive reward if the correct target is found, and a negative reward otherwise. In this model, the internal reward of fixations is the time cost associated with the time a person would need to collect information from the peripheral vision. As in lower eccentricity, fewer fixations would be needed when the target is within 10°. In this particular task, the fixation point is fixed in the middle between the target and the distractor, where correctly identifying a

target was measured by the saccade directions to either left or right and not by directly fixating on the target. Therefore, a fixation hypothetically represents fixations needed to collect information from the peripheral vision. Consequently, the number of fixations is expected to increase as the eccentricity gets higher as more effort would be needed to collect information from the peripheral vision. We assume that people with glaucoma would experience more noise than normally sighted people because of the loss in the visual field, including the peripheral field, as mentioned in Section 4.2. Furthermore, some studies showed that people with glaucoma experience increased equivalent input noise (Liu and Kwon, 2020). It is important to note that the model assumes that the loss of vision equally affects all areas of the visual field. However, this does not represent the full extent of the condition as individual glaucoma patients might experience different levels of visual impairments within different areas of the visual field (refer to Figure 1 in Boucart et al. (2020)).

4.5 Task formulation as POMDP

This section explains how the task by Boucart et al. (2020) is formulated as POMDP. The following is the formal formulation, with similarities to the formulation in Chapter 3.

- **State S:** A state in this task is represented as $s_t = (s_{\text{left}}, s_{\text{right}}) \in S$, which is a representation of the possibility of the target being on the left or on the right, where s_{left} and s_{right} are represented with either 1 or 0, where 1 is considered a target, and 0 is a distractor. There are two states in the model: (0,1) and (1,0). A state is picked randomly at every trial, and it remains the same throughout the steps in each trial.

- **Action A:** At each time step t, the agent takes action a_t to either collect information from the environment or to make a decision. In this model, the agent either fixates to receive a noisy observation (noise depends on the eccentricity) about the target's location or decides that the target is on the left or right.

- **Reward function $R(S, A)$:** After taking action a_t, the agent should receive a reward that

depends on the true state and the action taken, $r(s_t, a_t)$. A reward considers an empirically found cost of -3 for fixation actions because of the time cost of collecting information. As for decision actions, the reward is 0 for correctly identifying the target location, and -100 otherwise.

- **Transition function $T(S_{t+1}|S_t, A_t)$:** As described in Section 3.5 in Chapter 3.

- **Observation O and observation function $O = f(S, A)$:** After every fixation action, the agent receives a new observation o_t. The observation is a function of state and action, $o_t = f(s_t, a_t)$. As the target distances vary (based on the given eccentricity), the observation of the target being on the left or right is calculated based on the target location and the current distance between the target and the fixation point. The target location has a spatial uncertainty of peripheral vision, represented as noise (i.e. standard deviation), which is linearly dependent on the eccentricity. Therefore, $o_t = (o_{t_{\text{left}}}, o_{t_{\text{right}}})$, where $o_{t_{\text{left}}} \sim \mathcal{N}(s_{\text{left}}, \sigma_o)$, $o_{t_{\text{right}}} \sim \mathcal{N}(s_{\text{right}}, \sigma_o)$, and σ_o is the variance of the observation which does not depend on time and is linear to the eccentricity. Meaning that $\sigma_o = noise\,level \times eccentricity$, where $noise\,level$ is a hyper-parameter that we use in the observation to represent the different levels of uncertainty and find the best fit to the human data of both categories (i.e. normal vision and visually impaired).

- **Discount rate γ:** As explained in Section 3.5 in Chapter 3.

- **Belief update:** In this task, the belief is the pair consisting of the Gaussian distribution of the left value and that of the right value, stored as the mean μ and the variance σ^2 for each (four values in total: μ_{left}, μ_{right}, σ^2_{left}, σ^2_{right}). The initial belief is chosen as (0.5, 0.5, 100, 100), which is not the real distribution of the left and right values, as the real distribution is not Gaussian but exactly 0 with probability 0.5 and 1 with probability 0.5. We choose the initial belief with the correct means and very large variances to reflect the very high uncertainty of this initial value (i.e. 0.5). After taking action a, the model will receive a noisy observation o, about the value of the left and the value of the right, which are represented as a Gaussian distribution with mean μ_o and variance σ^2_o, where μ_o and σ^2_o

are the mean and variance of the received observation respectively. Then, the estimate of the belief for each left and right pair of mean and variance will be updated independently using the Kalman filter formula (see subsection 3.4.1 in Chapter 3).

4.5.1 Implementation

As in Section 3.8.2 in Chapter 3 in this experiment, we implemented the Q-learning method. Our choice of hyper-parameters for the method consists of the following:

- $\gamma = 1$, for the discount rate
- $\epsilon = 0.9$, for ϵ-$greedy$
- There is a starting learning rate, $\alpha = 1$, which decreases by 0.99 at every episode.

The model was trained with a random set of eccentricities that varies between 10 and 80 for 500,000 training episodes. Then, it was evaluated against each eccentricity from the experiment (i.e. $[10, 20, 40, 60, 80]$), with 1,000 testing episodes. Our goal is to fit the model results to the human data. Therefore, in order to find the best $noise\ level$ for the observations, we conducted a grid search with different noise levels. Additionally, we added a constant value to the noise in the observation to prevent the model from over-performing the human data. An example of the noise would be, $noise_{\text{left}} = constant + \sim \mathcal{N}(s_{\text{left}}, \sigma_o)$. The constant value was found empirically as 0.25. The human results consist of accuracy and latency values for each eccentricity, and for us to find the best $noise\ level$ that fits these values, we calculated the Euclidean Distance between our results and the human data to find a local minimum that represents the best fit.

In the next section, we provide the results of this experiment.

4.6 Results

In this section, we first present the learning performance of the Q-learning model. After this, we show the results of testing the model against different eccentricities and evaluate the results with the human data. Then, we find the best fit for the human data by using the Euclidean

distance. It is important to note that the human results in Boucart et al. (2020) only showed the performance through accuracy and saccade latency. However, in the model, we tested the eccentricities against the following:

1. **Accuracy:** This was determined by calculating the average number of times the model correctly identified the location of the target (left or right).

2. **Latency:** The latency was measured by multiplying the number of fixations by a constant value representing the time it takes for one saccade or eye movement to happen. We based the constant value on the average saccade time in the literature (i.e. time in ms = 40 ms) (Abrams, Meyer, and Kornblum, 1989; Rayner, 2009; Wang, Epps, and Chen, 2021).

3. **The number of fixations (or saccades):** We evaluated the model's number of fixations (or saccades) to collect more information about the location of the target.

4. **The rewards:** We finally analysed the rewards at each eccentricity.

Our choice of evaluations provides a comprehensive understanding of the model's performance, especially those we can compare to human data. Also, considering other relevant values can give more insight and a better understanding of how the model chooses its optimal strategies.

4.6.1 Model Performance

The model was trained with different values of the hyper-parameter $noiselevel$. The figures in this subsection illustrate the model's training performance. The model was trained for 500,000 learning episodes, 100 episodes per epoch. The results were moving averaged by 50 with a bin size of 5000. The figures above show the learning performance when the model was trained with different noise levels. The figures illustrate the following:

- In Figures 4.5a and 4.5b, the rewards in the learning curve and accuracy drop as the noise increases. This suggest that the agent's learning behaviour is influenced by the noise levels. Higher noise levels lead to lower rewards and less accuracy due to increased exploration (as illustrated in the size of Q-table in Figure 4.5d). In contrast, lower noise levels lead to

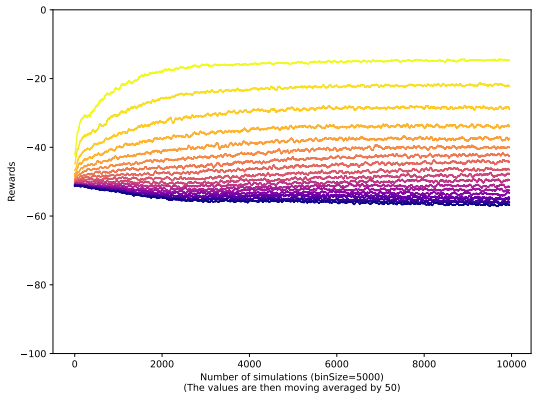

(a) The rewards during learning

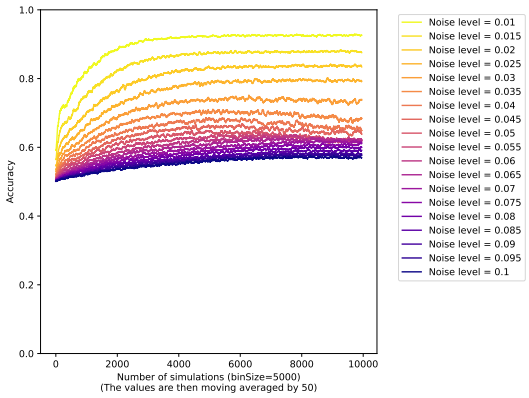

(b) Accuracy during learning

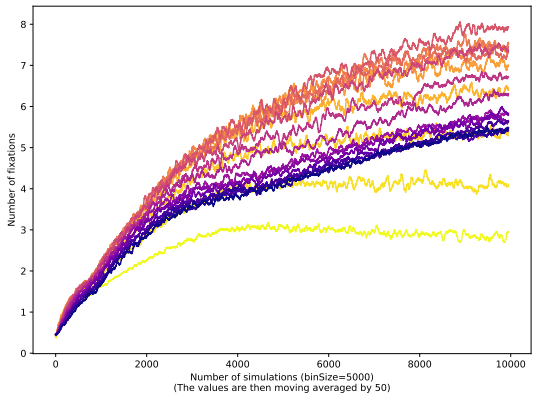

(c) Number of fixations during learning

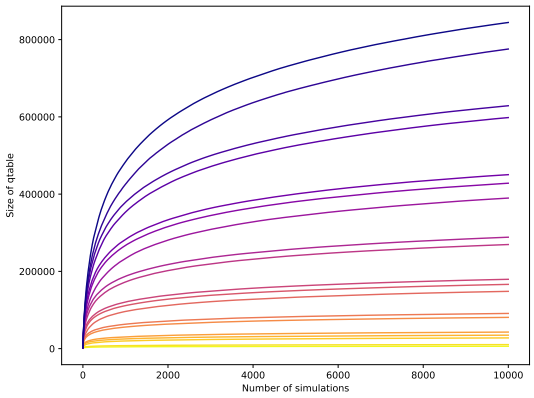

(d) The size of Q-table during learning

Figure 4.5: Model performance illustrated in the rewards, accuracy, number of fixations and the Q-table size

higher rewards, higher accuracy and smaller Q-tables due to more efficient exploitation. This demonstrates the trade-off between exploration and exploitation in reinforcement learning and how it impacts learning progress and rewards.

- In the fixations plot (Figure 4.5c), it shows that the number of actions taken for exploration is closely related to the accuracy, rewards (learning curve) and the size of the Q-table. Higher noise levels require more fixations. However, it is also noticeable that the number of fixations decreases again with higher noise levels (> 0.055). The non-linear relationship between noise levels and the number of fixations highlights the complex nature of learning at higher noise levels in the model. Whereas, at lower noise levels (i.e. 0.01 and 0.015), the agent shows more consistent fixation patterns reflecting efficient exploration, which can also be seen in the learning curve, accuracy and the size of the Q-table.

- Lastly, the clustering that can be observed in the Q-table size curves in Figure 4.5d is likely representing the points at which the agent transitions from more exploitation to more exploration, as it can be seen at higher noise levels (e.g. 0.095 and 0.1) the agent has to continuously adapt and explore a wider range of state-action pairs making it difficult to achieve convergence.

4.6.2 Testing Against Eccentricities

This part presents the results of testing the model against the set of eccentricities [10, 20, 40, 60, 80], and compares it with the human data provided by Boucart et al. (2020).

First, we tested the results from the grid search against the eccentricities and compared them with the human data. This can be illustrated in Figures 4.6 and 4.7.

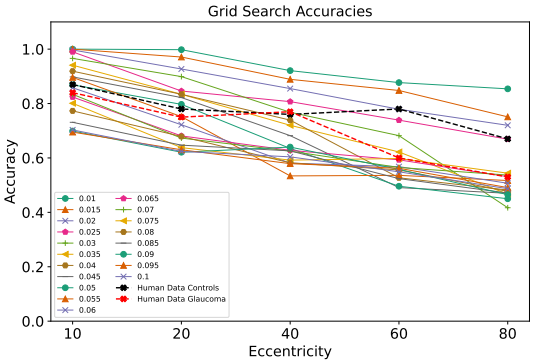

(a) Grid search results (accuracy)

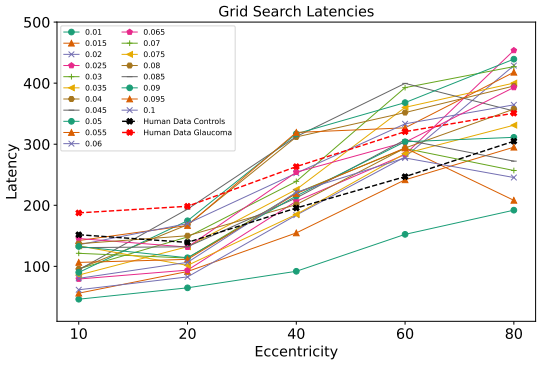

(b) Grid search results (latency)

Figure 4.6: The Figures 4.6a and 4.6b show the grid search results when tested against five different eccentricities, compared with the human data provided by Boucart et al. (2020). It can be seen that there is some decrease in accuracy and an increase in saccade latency as the eccentricity increases.

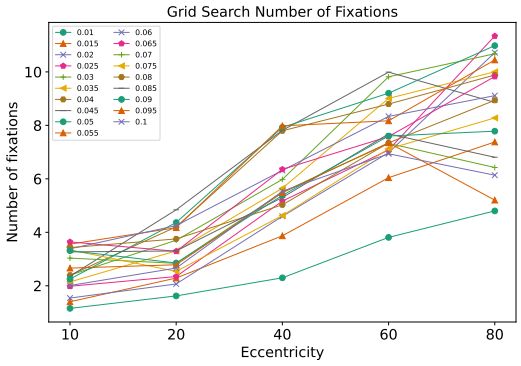

(a) Grid search results (number of fixations)

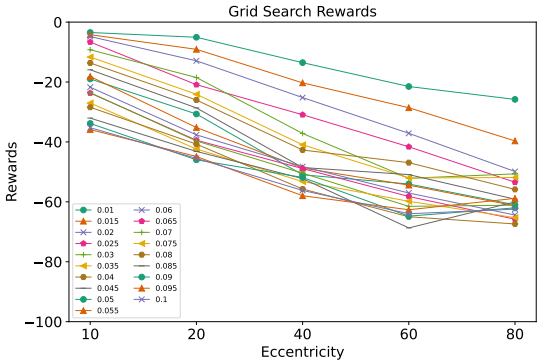

(b) Grid search results (rewards)

Figure 4.7: The Figures 4.7a and 4.7b show the number of fixations and the rewards tested against five different eccentricities. Figure 4.7a illustrates the increase in the number of fixations at higher eccentricities. However, in some cases, when the noise is very high, such as in a noise level of 0.095, the model fails to fixate due to the higher uncertainty. Whereas Figure 4.7b shows the decline in reward as the eccentricity increases, the relationship between the eccentricity and reward is not always linear, specifically for the cases of higher noise such as noise level 0.085 in the provided figure.

Second, in order to find the noise that best fits the human data, we calculated the Euclidean distance between the human data points and the model for the accuracy and latency separately. The results for the Euclidean distance between the accuracies are shown in Figure 4.8.

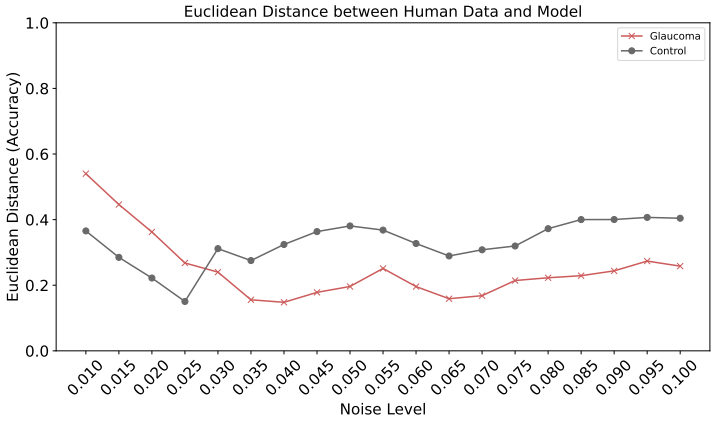

Figure 4.8: Eucludian Distance between model results and human data for accuracy

As can be seen in Figure 4.8, the minimum distance is at noise level 0.025 for controls (normal vision) and 0.040 for those with glaucoma. Figures 4.9a and 4.9b visualise the difference between the values at each eccentricity. Although the best fit in the model does not exactly match the human data, some similarities in the performance can be observed, for example, the drastic drop

in accuracy after eccentricity 40° for Glaucoma patients as shown in Figure 4.9b.

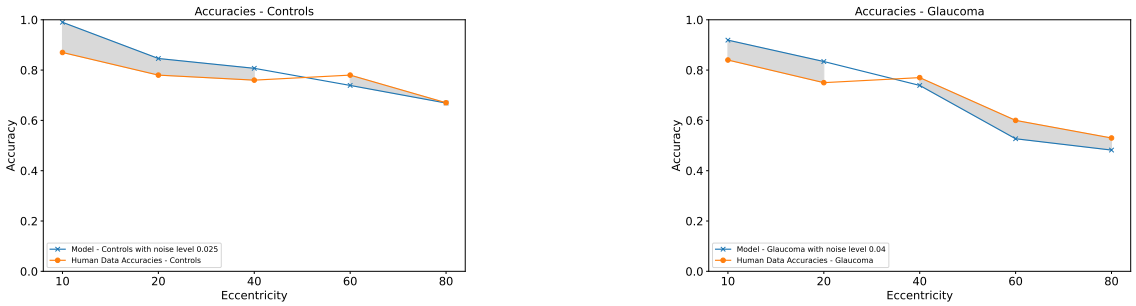

(a) Average accuracies for control participants vs. model

(b) Average accuracies for glaucoma participants vs. model

Figure 4.9: The difference between the best fit from the model results and the human data is highlighted in grey.

Third, the same steps were followed to get the best fit for the latencies values. The Euclidean distance results are presented in Figure 4.10. As shown in Figure 4.10, the best fit for control

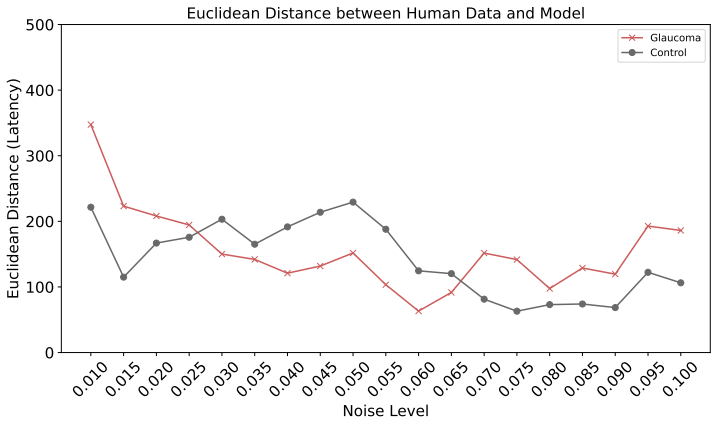

Figure 4.10: Euclidean Distance between model results and human data for latency

participants is at a noise level of 0.070 (with noise level between 0.070-0.090 illustrating comparatively good fit), whereas, for glaucoma participants, the best fit is at a noise level of 0.060. It can demonstrate the differences between each noise level and the human data performance, as lower noise levels had a bigger distance from the human data, while the distance at higher noise levels reaches a local minimum, then gradually increases as the noise increases. For latencies, unlike accuracies, the best fit for the controls is at a higher noise level than glaucoma. This is because the noise level of 0.06 had a higher saccade latency, especially

at eccentricity 60 and 80, which best fits the human data (see Figure 4.11b).

Figures 4.11a and 4.11b illustrate the difference between the model and the human values. The best-fit noise level for glaucoma participants is at noise level 0.06, as shown in Figure 4.11b, where the similarities between the model and patients with glaucoma can be seen especially at eccentricity 60 and 80.

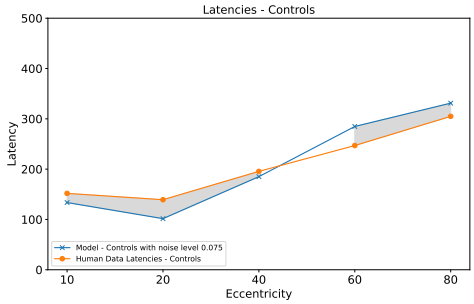

(a) Average latencies for control participants vs. model

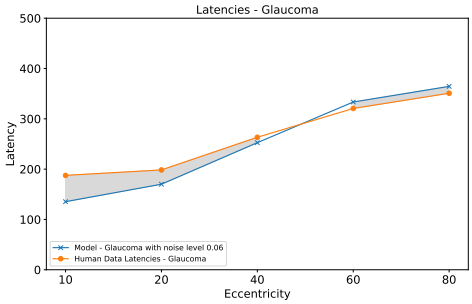

(b) Average latencies for glaucoma participants vs. model

Figure 4.11: The difference between the best fit from the model results and the human data is highlighted in grey.

Lastly, we compared our results with the data set in a box plot to demonstrate the overall results and performance of the model (See Figure 4.12).

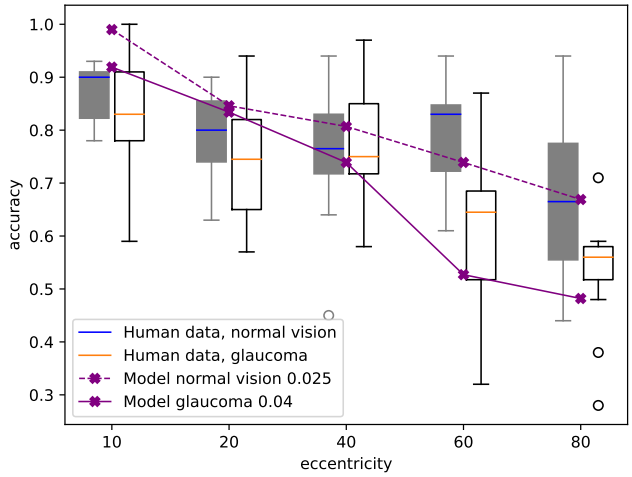

(a) Best fit of the model for accuracy compared against human data

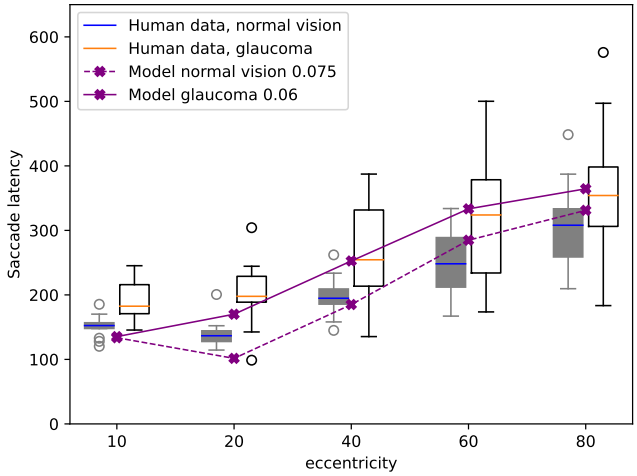

(b) Best fit of the model for latency compared against human data

Figure 4.12: In Figures 4.12a and 4.12b, the model was compared to the human accuracy and latency. It demonstrated the linear relationship between the accuracy of correctly locating a target in the peripheral vision and the eccentricity when a suitable proportion of noise is chosen. It also shows the increase in latency as the eccentricity increases with the right proportion of noise levels and sufficient training.

Overall, the results further highlight the model's ability to learn the visual behaviours of normally sighted people and people with glaucoma, given the methods and parameters tuning. It showed increased saccade latency at higher eccentricities and simultaneously decreased accuracies. The model predicted that less noise is needed to match the normal vision's accuracy,

and a higher noise level is needed to match the glaucoma vision. It also has been shown that less fixations were needed when the target was too close, and more fixations (collecting information actions) were needed when the target was further away from the fixation point. Additionally, the drop in the number of fixations at higher noise levels shows that the increased noise could potentially bias the model to fixate less.

4.7 Discussion

In this chapter, we implemented a computational model of visual search behaviours in people with glaucoma when they try to find a target within wide eccentricity. The goal of our model is to predict and give insight into how people with glaucoma would interact with a visual search task that involves wide eccentricities. We implemented the uncertainty in the peripheral vision in humans as noisy observations. The model was trained based on the assumption that human strategies emerge due to the environment's limitations and the bounds of the human visual system. The model learned optimal strategies under varying noise levels in the observations, shedding light on the relationship between the noise level in the peripheral vision and performance strategies.

Our findings indicate that the $noise\,level$ hyper-parameter introduced in the model can significantly impact the strategies from different aspects. The results presented in Section 4.6 can indicate that lower noise values outperform human participants. In contrast, excessively high noise levels led to a decline in performance, making it challenging to determine an optimal policy.

The results showed that the model predicted that people with glaucoma would experience more noise in the visual field, which can lead to a drop in accuracy, particularly at higher eccentricities. This can reflect on the challenges faced by individuals with glaucoma while interacting with a visual search task. It is important to note that the relationship between the noise levels and accuracy in our model is not strictly linear, especially as the noise gets higher, which can reflect the higher uncertainty at those noise levels. We believe that this can also reflect

the nature of human behaviours, as the accuracies of individuals were not exactly linear to the eccentricity. Each participant in the human experiment had different accuracies that reflected their own abilities, whether in the case of normal vision or glaucoma, where the levels of loss in the visual field differ between each individual (see Figure 4.3 in Section 4.3).

Furthermore, the model predicted an increase in latencies at higher eccentricities, which can give insight into the behaviours of people with glaucoma. These results suggest that people with glaucoma might take a longer reaction time to saccade when they try to find a target that is located within their peripheral visual field. The model predicted this behaviour when the observations were at a higher noise level (i.e. 0.06); see Figure 4.11b, where latency at noise level 0.06 drastically increased at eccentricities 60 and 80, which can best fit the glaucoma patients' behaviours. However, the saccade latency results in the model predicted slightly higher noise for people with normal vision (i.e. noise level 0.075). We believe this is because of the non-linear relationship at higher noise levels, which is caused due to the higher uncertainty in the model. We believe this behaviour can also align with the differences in individual behaviours in the human data especially for those with glaucoma (see Figure 4.4 in Section 4.3).

Furthermore, the model has shown that less fixations were needed when the target was too close, and more fixations (collecting information actions) were needed when the target was further away, which is a behaviour that can also be observed in patients with glaucoma (Crabb et al., 2010).

It is important to acknowledge the limitations of our study. One main limitation is that we built the model presented in this chapter based on the experiment done by Boucart et al. (2020), and one important limitation in this experiment is that the number of human participants was small. This means the human data does not provide a general representation of the visual behaviours in people with glaucoma, making it difficult to bring vigorous conclusions. One way to improve the accuracy of our findings is to extend the model by implementing more visual search tasks tested on people with glaucoma. This can lead to more general findings that can further our understanding of interactive behaviours in individuals with glaucoma. Another limitation of our model is that we only compared the values in the model with the average values

of human data. One way to improve our findings is by fitting the model to each participant's specific visual field defects by specifying areas of greater or lesser vision loss so the vision loss is not uniform across the visual field. Additionally, we might need to introduce more hyper-parameters to fit the model to each individual's vision.

To summarise, we assumed that human behaviours result from adaptation to the limitations of the environment and the bounds of the visual system. This influenced formulating the task as POMDP and solving it using the Q-learning method. The extensive range of phenomena associated with the role of glaucoma in HCI-related tasks, such as visual search and driving, offer an intimidating picture for design. One evident generalisation is that people with glaucoma adapt to their condition (e.g., using different patterns of fixations and saccades). If the design for people with visual impairment is to advance, a deeper understanding of how these adaptations are emergent consequences of the physiological deficit is required. The model showed similarities in performance to the human participants, which shows that the use of these models could predict similar visual behaviours of those with glaucoma. The proposed model could be a starting point to be used as a tool for assisting UI design specifically to be used in interactive technology which includes wide visual fields such as Augmented Reality (AR) and Virtual Reality (VR) (Hosseini, 2017; Renner and Pfeiffer, 2017).

Chapter 5

Experiment 2: A Model of the Crowding Effect with Age-related Macular Degeneration

5.1 Introduction

This chapter demonstrates the implementation of the computational model to a visual search task that investigates the visual search behaviours and strategies amongst individuals with neovascular Age-related macular degeneration (AMD). The human experiment used in this chapter was done by Thibaut, Boucart, and Tran (2020). The main goal of the human experiment was to test the crowding effect and its impact on visual search strategies in those with AMD. It required participants from three categories (i.e. young, age-matched controls, and AMD) to find a target surrounded by distractors within different crowding settings. We built an ML model to predict the visual search behaviours in a crowded setting for the three categories. In order to train the model, the task was formulated as POMDP and solved using the PPO method (see Chapter 3 section 3.8.2). The crowding effect was implemented as a noise in the observation that depends on the flankers' distance (i.e. the spatial distance between a target and distractors) in addition to the spatial noise, which depends on the distance between the fixation point and the object.

The crowding noise increases as the distance between an object and flankers decreases. The spacial noise increases as the distance between the fixation point and an object increases (i.e. the eccentricity). To find the best fit for the human data, we implemented two different types of noise: Gaussian noise and what is referred to as *Blending* noise. In the *Blending* noise, the mean is an average value between an object and its flankers. Therefore, when a flanker is very close to an object, the model might observe "blended" features and mistakenly select a distractor as a target. This feature is applied in the model based on how a crowded setting affects the human vision (Levi, 2011; Whitney and Levi, 2011). The results showed a decline when the crowding effect was present and predicted the human data behaviour, especially in the blended noise setting. However, our model performed better than the humans in the uncrowded setting.

5.2 Background

As explained in Chapter 1, AMD is an eye condition that causes blurriness in the central visual field. It is considered one of the main reasons for blindness that mostly affects older people globally (Lim et al., 2012; World Health Organization, 2022). A systematic review by Wong et al. (2014), showed that 8.7% of the population worldwide had age-related macular degeneration. They also projected that the number of people with the disease will increase to 288 million between 2020 and 2040. Due to the loss in the centre of the visual field, people with AMD adapt their visual search strategies when interacting with a visual search task. One main adaptive strategy is the preferred retinal locus (PRL), and it is when patients use a discrete retinal area for fixation as an alternative to the fovea (i.e. the centre of the visual field) (Crossland et al., 2005; Guez et al., 1993; von Noorden and Mackensen, 1962; Watson et al., 2006; White and Bedell, 1990; Whittaker, Budd, and Cummings, 1988). However, studies showed that PRL is not always the best area for good visual acuity, fixation stability or high retinal sensitivity (Altınbay and İdil, 2022; Bernard and Chung, 2018; Shima, Markowitz, and Reyes, 2010). Therefore, to find the best visual function for centre scotoma, PRL training was introduced, and it showed some improvements in different visual search tasks, such as improving reading strategies, face

recognition and walking (Deruaz et al., 2002; Hassan et al., 2019; Oflaz et al., 2022; Walsh and Liu, 2014). Furthermore, computer-based visual behaviours were tested in people with AMD using a variety of visual search tasks and in daily visual activities such as driving (Higgins et al., 2020), face recognition (Taylor et al., 2018), and reading (Rubin and Feely, 2009).

Crowding is the phenomenon in which objects are rendered in clutter, which results in them being unrecognizable. The crowding effect can cause limitations in the peripheral visual field that can lead to difficulties in daily activities such as reading, driving and object recognition (Whitney and Levi, 2011). It is found to be limiting object recognition, and negatively affecting reading speed in people with healthy vision (Pelli, 2008; Strasburger, Harvey, and Rentschler, 1991). Additionally, Pelli (2008) mentioned that small spacing between objects can make them even harder to recognise than when located further in the peripheral vision. Due to the importance of the centre of the visual field, people with AMD find it more challenging to recognise objects, especially if they are located within a cluttered environment (i.e. a crowded environment) (Wallace, Chung, and Tjan, 2017). The challenges of the crowding effect and people with AMD have been addressed in different visual search experiments. For example, Chung (2014) tested the effects of spacing and letter size in people with AMD in letter recognition tasks. Their findings showed that both crowding and letter size can have a negative impact on visual tasks for those with AMD. In another experiment done by Wallace, Chung, and Tjan (2017), they addressed critical spacing, contrast sensitivity and threshold elevation in crowded conditions in people with AMD. In their experiment, they chose AMD participants with stable and well-developed PRL for fixation. Their findings showed that people with AMD had worse threshold elevation as well as contrast sensitivity than those young and age-matched participants, which made the object recognition task more challenging.

5.3 Visual Search Behaviours in AMD

The effect of AMD on the visual field and eye movements has been studied by different experiments using eye-tracking tools to detect visual behaviour and decision-making strategies. In

such work by Van der Stigchel et al. (2013), they explored the effects of Macular Degeneration (MD) on eye movement through two types of visual search tasks (i.e. serial and pop-out). Participants in this experiment consist of three groups: (1) people with MD, (2) a healthy control group to test the same tasks as the patients and (3) a healthy control group to test their visual search on an MD simulation task. The research reports on two main sets of findings: (1) search latency, which is the time required to find the target, was longer in the MD group; also, for the display with no target, there was no difference between the search latency between MD and control group which the reduced sensitivity in the task might cause; (2) Saccade amplitude: MD group and both control groups are all similar in the serial condition; difference were observed in the pop-out condition, where control group made larger saccadic movements than MD group. Furthermore, there was a trend level difference between the MD and control simulation group, where, on average, the MD group made smaller saccadic movements. In terms of the number of saccades to target, in both conditions, the MD group needed more saccades to find the target, whereas the other groups needed less number of saccades. Another experiment done by Shanidze et al. (2016) focused on testing smooth pursuit in people with MD with the hypothesis that the pursuit of people with MD would depend on the position of the scotoma and the target trajectory. Participants with MD were asked to track a target moving in eight radial directions at speeds between 5°/s to 6°/s. Their findings showed a drop in pursuit gains with the increase of the scotoma extent in the target's heading direction. It was also reported an increase in latencies when scotoma was within the target trajectory. These findings indicate that people with MD are able to make smooth pursuit eye movements, which are limited to the target trajectory and the characteristics of their scotoma. Object identification through visual exploration was also examined in people with AMD by Thibaut et al. (2016). The main purpose of the latter study is to test if people with centre scotoma can recognise an object presented to them. Participants were presented with photos of nature and man-made objects, isolated or in their natural scenes. Then, they were asked to name what they could see, and their visual behaviours were then recorded using an eye tracker. The results showed that people with AMD had lower accuracy by 30% than young and control groups. They also made more saccades, less fixation duration and

longer scan path than the other groups. This is believed to be a result of using the peripheral visual field during the search along unstable fixations. In a more recent study done by the same scientists (Thibaut, Boucart, and Tran, 2020), they examined the visual search behaviours of people with neovascular AMD in an object recognition task, which included the crowding effect. Participants were required to find a target in either crowded or uncrowded conditions. The findings showed that people with neovascular AMD were less accurate and needed longer search times than participants with healthy vision in crowded and uncrowded conditions. Additionally, people with AMD needed to do more fixations to find the target, especially in crowded conditions. The latter task was used for the purpose of our experiment, and it will be explained in more detail in the next section.

5.4 Experimental Task

This section explains the AMD visual search task that we chose from the literature.

5.4.1 Object Search in Neovascular Age-related Macular Degeneration: The Crowding Effect

As mentioned in Section 5.2, crowding is the extent to which an object that a person is looking for (e.g. a button on a screen) is surrounded by other objects (sometimes called flankers). In an experiment designed by Thibaut, Boucart, and Tran (2020), visual behaviours were tested on people with neovascular AMD in an object search task with two different crowding conditions. The task was tested on three different groups of participants, which are AMD, age-matched controls (AMC) and young. Participants were asked to find an object surrounded by distractors. Objects were randomly presented on a display that measured 31.6° horizontally × 25.4° vertically at a viewing distance of 60 cm. There were two spacing types, which are "crowded" with 1.5 degrees spacing, and "uncrowded" with 6 degrees spacing. The object size varied between 5° and 6° horizontally and vertically. Three different displays were created for this experiment: (1) a display with nine objects (1 target, eight distractors), (2) a display with six

objects (1 target, five distractors), and (3) a display with four objects (1 target, three distractors). In order to reduce the 'pop out' effect in the stimuli, the target and distractors had the same colour as well as the same semantic category (either objects from nature or man-made objects). An example of the stimuli is shown in Figure 5.1. There were 96 trials; each trial lasted for a maximum duration of 30 seconds. The order of images was semi-randomised (i.e. a block of four images was first selected by the computer and then presented in a random order). At the start of each trial, participants were asked to start by fixating on a fixation cross ($2\times2^{\circ}$) displayed on the centre of the screen for one second. A target was verbally specified by the experimenter; the target's location was random in each trial, and participants were asked to find it as quickly as possible. After participants found the target, they were asked to first fixate on the target, and then press on the space bar to report their response. The stimuli then disappeared. Accuracy and target search time were measured in all categories and considered the main outcome measures of this study. The results (Figures 5.2) showed that participants with AMD had the lowest accuracy and the highest search time compared to the other two groups of participants. AMD participants were less accurate by 30% than participants with the same age group and the search time increased by 61% (see Figure 5.3 for the results of each individual in the experiment). Other measures were also taken into account, which is the number of fixations and scan path ratio, where the latter measures how directly the eyes move to the target (Brockmole and Henderson, 2006). The results showed that AMD participants did more fixations than both age-matched controls and young participants. The scan path ratio was less efficient in the group with AMD, confirming the low levels of explorations in people with AMD. The young participants, on the other hand, made a lower number of fixations than other categories and had a more efficient scan path ratio. Increasing the number of distractors in the display resulted in increased search (Figure 5.2b) time, fixations, and a less efficient scan path ratio in all three categories, especially when the number of distractors is 8.

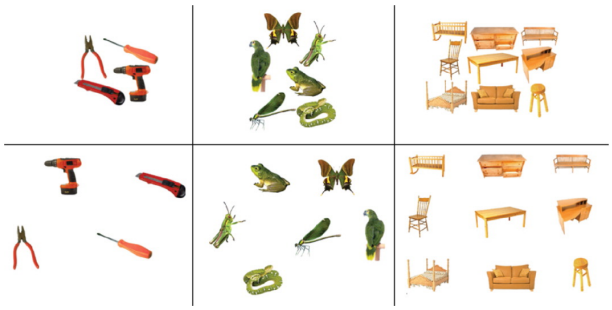

Figure 5.1: An example of the stimuli showing both crowded and uncrowded conditions as presented by Thibaut, Boucart, and Tran (2020). There are three conditions, nine objects: one target and eight distractors(right), six objects: one target and five distractors (middle), and four objects: one target and three distractors (left). (image used with permission)

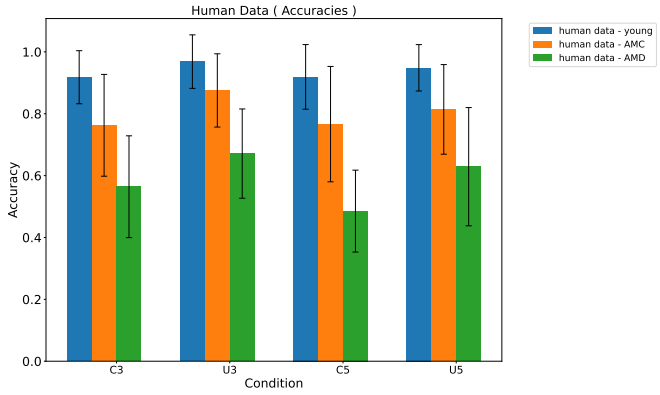

(a) Human data (accuracy).

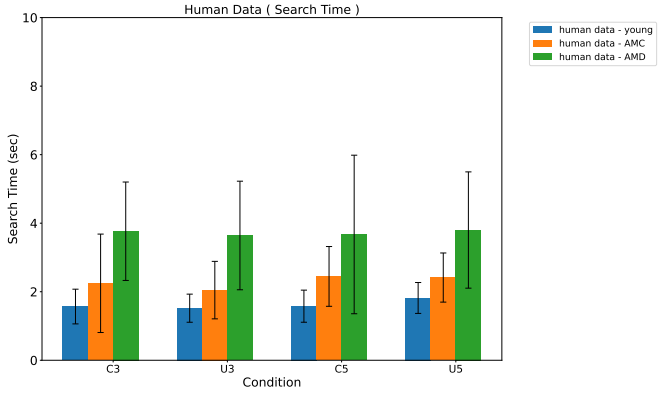

(b) Human data (search time).

Figure 5.2: The figure shows the human data provided by Thibaut, Boucart, and Tran, 2020 showing the results of the experiments with 4 and 6 objects (The figures show the mean and the standard deviation, where C3 refers to (Crowded 3 distractors), U3 (Uncrowded 3 distractors), C5 (Crowded 5 distractors), and U5 (Uncrowded 5 distractors)). We only included the data for four and six objects as these settings are relevant for the comparison with our findings.

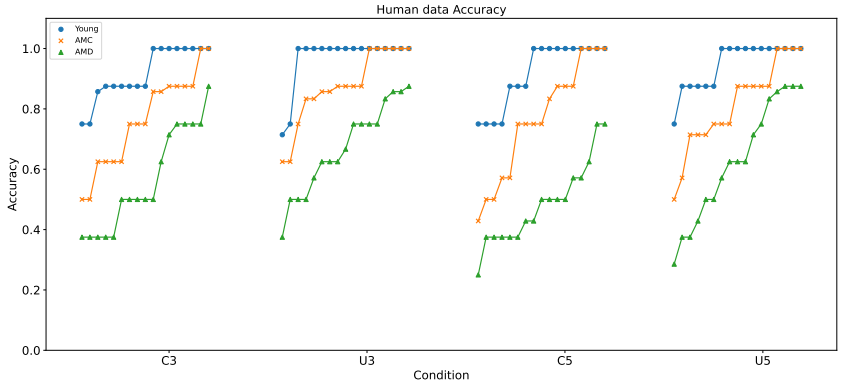

(a) Accuracy

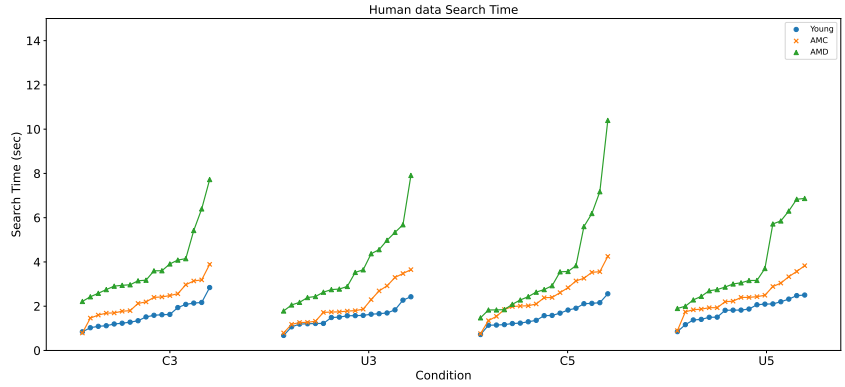

(b) Search time

Figure 5.3: Human data of each participant from the human experiment (Thibaut, Boucart, and Tran, 2020)

5.5 Theory

In this chapter, the model of visual perception was built based on theoretical assumptions that human visual behaviours and decision strategies are the results of adaptation to the statistical environment (target's size, the distance from fixation and the distance between the objects), the constraints imposed by the human visual system in addition to those imposed by having AMD, and the reward associated with actions and decisions. The task mentioned in Section 5.4 was formulated as POMDP and solved using DRL as explained in Chapter 3. Similar to that task in Chapter 4, the internal reward of fixations is the time cost associated with the time a person would need to collect information. However, in this study, the number of fixations was

measured, and the model was required to fixate on the target directly. This was also followed by additional action to confirm the final decision (i.e. pressing on the space bar in the keyboard), which was also implemented in the action space, which will be discussed further in Section 5.7. We assume that people with AMD would experience more noise than normally sighted people due to the loss in the centre of the visual field. Therefore, crowding noise was added to the model to represent the crowding conditions. We expect a high number of fixations and more search time to be needed in higher noise levels, which represents those with AMD. The main theoretical assumptions in this experiment are as follows:

- **Spacial Uncertainty:** As previously stated in Chapter 1, the humans' fovea has the highest density covering covers 1-2 degrees of visual angle (Bringmann and Wiedemann, 2022; Duchowski, 2018). Whereas the perception levels in the peripheral vision drop drastically, as explained in Chapter 4. In this experiment, we present the peripheral vision as a linear standard deviation to the eccentricity (Michel and Geisler, 2011; Todrov, 1998).

- **The Crowding Effect:** In addition to the decline due to the increase in eccentricity in the peripheral vision, the spatial separation between objects located in the peripheral vision plays an essential role in visual search strategies (Levi, 2011; Strasburger, 2020; Tadin et al., 2012; Whitney and Levi, 2011). When the crowding effect is present in a visual search task, it can disrupt the ability to distinguish the target, leading people to confuse a target with a flanker (Levi, 2011; Whitney and Levi, 2011). As the eccentricity increases, greater spatial separation between the targets is necessary to prevent the crowding effect from occurring (Tadin et al., 2012). Two factors can result in confusing flankers with a target, which are the flanker distance and the eccentricity (Strasburger, 2020). In the model, the crowding effect is presented as additional noise that increases as the flanker distance between objects decreases, causing less probability of seeing the objects, especially when there is an increase in eccentricity. More details about how we calculated the crowding in the model will be explained in Section 5.6.

- **Search duration:** The search duration is the duration of the entire episode (i.e. from the

agent's first fixation until making a final decision), which is the sum of the duration of all fixation actions (i.e. the sum of $Action\,Time$). Each fixation action has a duration, which is the sum of the saccade duration (i.e. the time for the fovea to move from the previous location of fixation to the current location of fixation (Baloh et al., 1975)) and the fixation duration (see Equation 5.1). There is a pre-defined maximum number of fixation actions that the model cannot exceed. This adds another constraint that limits the information saved in the model (Acharya, 2019). The search time is calculated as follows:

$$Action\,Time = Saccade\,Duration + Fixation\,Duration \tag{5.1}$$

$$Saccade\,Duration = 2.7 \times Amplitude + 37 \tag{5.2}$$

Where the $Fixation\,Duration$ is a constant value of 0.230, which is the average fixation duration in the distractor ratio task when the target is present (Shen, Reingold, and Pomplun, 2003), and the $Amplitude$ is the distance between the current fixation location and the previous fixation location.

It is important to note that in this experiment, we do not specify an area within the visual field which could represent the loss in the foveal vision people with AMD experience, but we assume that people with AMD are looking at an object and are perceiving information within the foveal visual vision, for example, through finding an alternative foveal vision. In our model, we considered the loss of vision as unifom across the peripheral visual field without specifying an area with visual loss (e.g. top right). We also do not consider features of the objects presented in the stimuli such as all objects having the similar colour and size but different shapes and belonging to the same category.

The next section explains how we computed the crowding effect in this task. It is essential to explain this computation as it will be used when we explain the task's formulation as POMDP.

5.6 How the Crowding Effect is Calculated

The crowding probability is computed from the literature and based on Bouma's law (Bouma, 1970), which states that the critical distance (minimum flankers' distance to the target where crowding starts to occur) is about half the target's eccentricity. Strasburger (2020) derived how the crowding changes with increased eccentricity, emphasising the assumption that the average distance between objects is independent of the visual eccentricity. To develop the crowding probability equation, Strasburger (2020) used both Bouma's law (critical distance vs. eccentricity) and psychometric function that considers the performance vs. flanker distance, which was used from the literature (Yeshurun and Rashal, 2010). Figure 5.4 demonstrates the crowding calculation in relation to eccentricity and flanker distance. The function of crowding probability c used in our model is as follows:

$$c = 1 - \lambda - \Phi(flanker\_distance - \beta \times eccentricity) \tag{5.3}$$

For this problem, lapse rate, λ is considered 0, Φ is used for the proportional correct flanker distance vs. the eccentricity, and β is a scaling factor around 0.5 as given in the original equation. Then, from the crowding probability c, we compute the observation noise explained in Section 5.7.

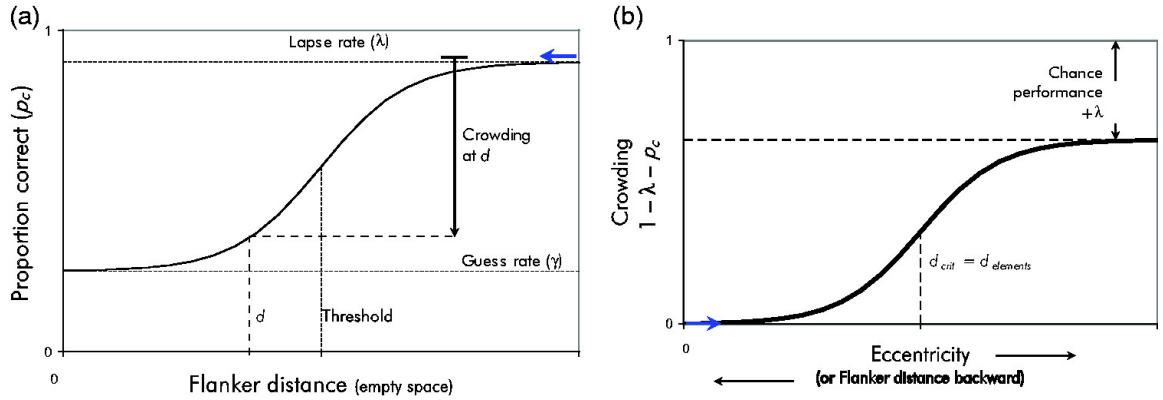

Figure 5.4: This diagram from (Strasburger, 2020) illustrates the conventional concept of crowding, which refers to the decline in the ability to recognize objects when surrounded by other objects (flankers). (a) It presents a psychometric function graph that shows the proportion of correct responses in a crowding task, similar to the original graph presented in (Yeshurun and Rashal, 2010). The impact of crowding is represented by a downward arrow on the right side of the graph, occurring at a specific distance d from the flankers, indicating a decrease in performance from the highest level ($1–\lambda$). (b) Illustrates crowding as shown in Figure (a), but in this case, it is represented as a function of eccentricity instead of a particular flanker distance. Part (b) is derived from Part (a) through graphical adjustments, which involve horizontally and vertically mirroring the psychometric function graph and suitably scaling the Y-axis. The blue arrow helps to demonstrate these transformations visually. (Figure used with permission)

5.7 Task Formulation as POMDP

This section explains the formal POMDP formulation of the visual search task we discussed in section 5.4

- **State S:** A state s_t is a vector of $(x, y, targetedness) \times number\,of\,objects$, which represents the objects' locations. The $targetedness = 1$ if an object is a target, or -1 if it is a distractor. The locations of the objects are randomly distributed in each trial and bounded by the size of the screen ($31.6°$ horizontally $\times$ $25.4°$ vertically) as given in the human experiment, and the object's width ($object\,width = 5$). The state remains unchanged in each trial until a final action is taken.

- **Action A:** The action space consists of fixation to get observations about the environment or to press the space bar to confirm the final decision. At each time step, the agent takes an action a_t where the fixation action $a = [x, y]$ is bounded by the number of objects and the

screen size. The press space bar action = $[0, 1]$, where 1 means press the space bar and end the trial, and 0 means do not press the space bar and continue collecting information about the environment.

- **Reward function $R(S, A)$:** After taking action a_t, the agent should receive a reward that depends on the state and the action taken in that state, $r(s_t, a_t)$. For the given problem, a cost of -1 was empirically chosen for fixation actions. As for the decision action (i.e. when the space bar is pressed), $reward = 100$, for correctly identifying the target location (when the fixation location is within the object, which is a circle of radius 5 around its centre). When the agent does not correctly identify the target, then $reward = -5 \times distance$, where $distance$ is the distance between the fixation location and the centre of the target.

- **Transition function $T(S_{t+1}|S_t, A_t)$:** This is the same as explained in Chapter 3, Section 3.5.

- **Observation O and observation function $O = f(S, A)$:** After every fixation action a_t, the agent receives a new observation o_t, a function of state and action, where $o_t = f(s_t, a_t)$. In our model, the observation is a tuple of the observation $o_{t,obj}$ for each object obj, and it depends on the action the agent takes, the true target location in the state s_t, and the spatial distance between objects (flanker distance). Meaning that (as previously explained in Chapter 3) the distance noise ($d_{t,obj}$) and crowding noise ($c_{t,obj}$) were added to the actual value(s) in the observation $o_{t,obj}$. The distance noise $d_{t,obj}$ is a Gaussian distribution, $d_{t,obj} \sim \mathcal{N}(\mu = 0, \sigma_{o,d})$, where $\sigma_{o,d}(t, obj) = distance\ factor \times eccentricity(t)$ which explains the spatial uncertainty (Todrov, 1998). The distance factor is, $distance\ factor = 0.9 \times Noise\ Level$, in our experiment, the hyper-parameter 0.9 is referred to as $base\ distance\ factor$, and this value showed a better fit to the human data in this experiment than 0.09, which was found in the literature (Chen, Acharya, and Oulasvirta, 2021; Michel and Geisler, 2011). On the other hand, for the crowding noise $c_{t,obj}$, we experimented with two types of noises explained as follows:

1. Gaussian noise $c^{(G)}_{t,obj} \sim \mathcal{N}(\mu = 0, \sigma^{(G)}_{o,c})$:

   Gaussian noise ($c^{(G)}_{t,obj}$) represents the noise with a mean of 0 and standard deviation $\sigma^{(G)}_{o,c} = crowding\ factor \times crowding\ probability$, where $crowding\ factor = 10 \times Noise\ Level$, and 10 is the $base\ crowding\ factor$ that was found to be best at producing the crowding effect in the results. It is important to note that the crowding noise considers the spatial distance between the fixation point and all the objects in the state. The $crowding\ probability$ is explained in the Section 5.6.

2. Blending noise: $c^{(B)}_{t,obj} \sim \mathcal{N}(\mu = average\ crowding, \sigma^{(B)}_{o,c})$:

   The mean $average\ crowding$ is a representation of how close an object is to its flanker. For example, if the flanker distance is 0 (two objects very close together), the average of the observed value will be the average between the object value and the flanker value. The observed value is a weighted average between the observed object feature value and the flanker feature value using a blending crowding factor of 0.5. The blending crowding factor accounts for the balance between these feature values. It is then multiplied by the proximity, which measures the spatial relationship between the objects, indicating how close one object is to its flanker (Kennedy and Whitaker, 2010; Treisman and Gelade, 1980)). This proximity is computed as shown in Equation 5.4:

   $$proximity = 1/(1 + flanker\ distance) \tag{5.4}$$

   Consequently, the calculation of the average crowding is as follows:

   $$average\ crowding = (flanker\ features - object\ feature) \tag{5.5}$$

   $$\times\ blending\ crowding\ factor \times proximity \tag{5.6}$$

   The standard deviation $\sigma^{(G)}_{o,c}$ is computed in the same way as $\sigma^{(G)}_{o,c}$.

The observation o_t is a list of pairs. Each pair includes the noisy value of the feature ($targetedness$) and its confidence for each element in the object observed (i.e.

$(x, y, targetedness)$. The observation is calculated as follows:

$$targetedness = f_{obj} + d_{t,obj} + c_{t,obj} \tag{5.7}$$

Where f_{obj} is the true feature value for the object, and the feature confidence is computed as follows:

$$feature\ confidence = \sqrt{\sigma_{o,c}^2 + \sigma_{o,d}^2} \tag{5.8}$$

Therefore, the observation is a tuple of the $(targetedness, feature\ confidence)$ for each objects.

After computing the observation, the agent updates the belief state. This will be explained in more detail as part of this section.

- **Discount rate γ:** As explained in Chapter 3, Section 3.5, the discount rate for this study was left unchanged as a parameter in the implemented library (see Section 3.8.2 in Chapter 3).

- **Belief update:** As the state is not fully observable, the model maintains a belief state of the current state through taking action, a_t, and receiving an observation, o_t. The belief in this study is a vector of tuples, where each tuple represents one object with two values (targetedness and confidence). The initial belief consists of the initial targetedness value (unknown value = 0), and initial confidence, representing the very high uncertainty (i.e. 1000000). At every fixation action, the agent gets a new noisy observation (including both distance and crowding noise), and then updates its belief state using the Kalman filter formula in Chapter 3, Section 3.4.1.

5.7.1 Implementation

As mentioned in Chapter 3 Section 3.8.2, we used the PPO method from OpenAI Baselines within OpenAI Gym (Brockman et al., 2016). All the default hyper-parameters for PPO remained the same except for the learning rate, which was updated to 1×10^{-4}. In order to choose the values of the hyper-parameters that we introduced in the model, we conducted an iterative process of experiments to examine the performance of the model with different *Noise Level* values. Our final choices of hyper-parameters, based on empirical observations, are $[0.1, 0.15, \ldots, 0.9]$ for the *Noise Level*. As mentioned in section 5.7, the *Noise Level* is multiplied by the base crowding factor 10, and the base distance factor 0.9. Therefore to clarify the actual values for the *Crowding Factor* and the *Distance Factor* , we created the following table:

Noise Level	*Crowding Factor*	*Distance Factor*
0.1	1	0.09
0.15	1.5	0.14
0.2	2	0.18
0.25	2.5	0.23
0.3	3	0.27
0.35	3.5	0.32
0.4	4	0.36
0.45	4.5	0.41
0.5	5	0.45
0.55	5.5	0.49
0.6	6	0.54
0.65	6.5	0.58
0.7	7	0.63
0.75	7.5	0.67
0.8	5	0.72
0.85	8.5	0.76
0.9	9	0.81

Table 5.1: Actual values for *Crowding Factor* and *Distance Factor* as implemented in the observation

We only considered four and six objects in the model due to the high complexity of the problem. The model was trained for 6×10^6 iterations for four objects and 15×10^6 iterations for six objects through a grid search to find the best fit to the human data. We used a 5000 moving

average to produce smooth curves. The model was trained for each noise level on a random set of spacing S, where $S = [1.5, 6.5]$ (representing the spacing between objects within the state). Then, it was tested for 2×10^3 iterations for each experiment (four and six objects). In the testing stage, the model was given the same spacing values as those used in the human experiment (i.e. 6 for the uncrowded condition and 1.5 for the crowded condition). We separated the training into more than one group to organise the noise levels and present them in a clear way (e.g. we trained the noise levels $[0.1, 0.15 \ldots 0.45]$ separately and then trained $[0.5, 0.55 \ldots 0.9]$. There are cases where the model does not learn to explore the environment during learning with some high noise levels. This will be discussed further in Section 5.8 where the results are presented. After testing the model, the results were evaluated against the human data provided by Thibaut, Boucart, and Tran (2020). The human data provided to us only included the accuracy and search time (in seconds), which was considered the main outcome of the experiment. This allowed us to compare the model's accuracy and search time against the data. We also measured the number of fixations and scan path ratio in the model, and compared them to the analysis of the results in the human experiment, as these measures were considered secondary outcomes for the experiment. Lastly, to find the best noise level that fits the data, we computed the difference between the model's results and the data to find a local minimum that best fits the human data.

In the next section, we demonstrate the model's results.

5.8 Results

In this section, we present the results of our experiments. As the experiments considered different noise types (Gaussian and Blending), we have divided the results into subsections, each representing a specific combination of noise types and number of objects.

5.8.1 Four Objects: Model Performance (Gaussian)

This subsection illustrates the learning performance of the models when trained with four objects using the Gaussian noise $c^{(G)}_{t,obj} \sim \mathcal{N}(\mu = 0, \sigma^{(G)}_{o,c})$. First, we present the model performance

while training. We show how the agent interacts with the environment we created and learns to make sequential decisions, which maximises a cumulative reward and optimises the policy. In this experiment, we separated the learning into two groups (i.e. $[0.1, 0.15 \ldots 0.45]$ and $[0.5, 0.55 \ldots 0.9]$). To give a visual representation of the Gaussian noise, we added the histogram in the Appendices (see Appendix A Section A.1.1). After presenting the training performance, we present the results of testing the model with the corresponding spacing values for the crowded and the uncrowded conditions. We evaluate our results against the human data for accuracy and search time.

Grid search was conducted using different noise levels. The model performance for four objects with Gaussian noise is illustrated in Figures 5.5 and 5.6. As can be seen in the figures, the learning within these noise levels is consistent, and the effect of increasing the noise levels can be observed clearly through all ten plots. Additionally, our analysis reveals interesting patterns in the model's learning behaviour. For example, if we compare the accuracy (Figures 5.5b and 5.6b), there is an apparent reduction in accuracy performance when the noise levels increase. Another interesting observation is that the model started learning quicker at lower noise levels (see Figure 5.5a). However, at higher noise levels (Figure 5.6), the model needed more episodes to start learning. It indicates that it chooses a random action at each episode without learning, which is believed to be a result of the increased complexity and the nature of the stochastic environment. Thus, this behaviour can also explain the lower number of fixations, the inefficient scan path ratio and the short search time (when initiating the learning) (see Figures 5.6d, 5.6e, and 5.6c).

In contrast, at lower noise levels (e.g. 0.1 and 0.15), after convergence and learning the optimal strategies, the model used more, yet short, episodes throughout the iterations. Hence, learning the appropriate number of fixations leads to a more efficient scan path ratio, less search time and high accuracy (see Figure 5.5). It is important to note that the number of learning iterations was constant and chosen based on our experiments and observations. The results showed convergence in all the noise levels in this experiment and showed a decline in performance upon increasing the noise levels in the observations. This reduced random

behaviour in the model. In the next section, we test the model's performance on the specified distancing values and evaluate it against the human data.

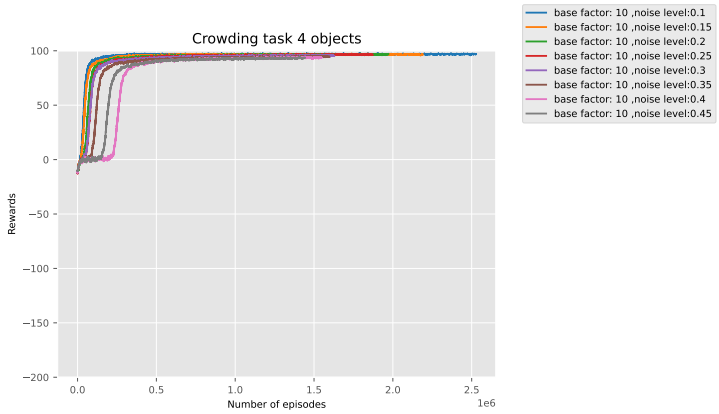

(a) Rewards during learning

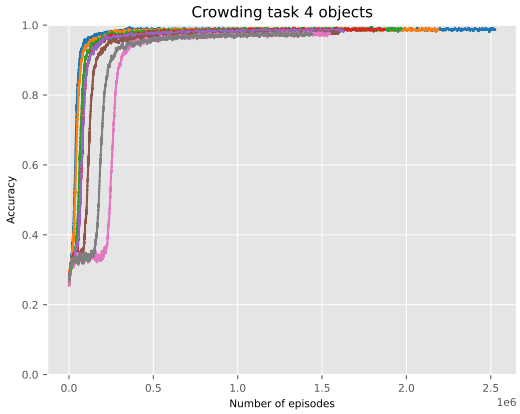

(b) Accuracy during learning

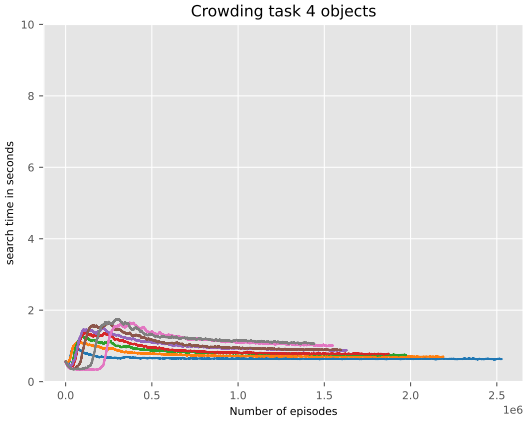

(c) Search time during learning

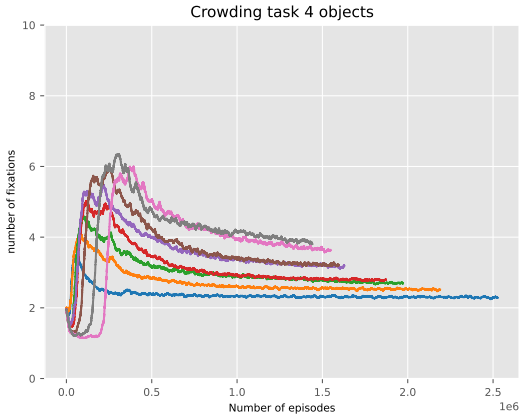

(d) Number of fixations during learning

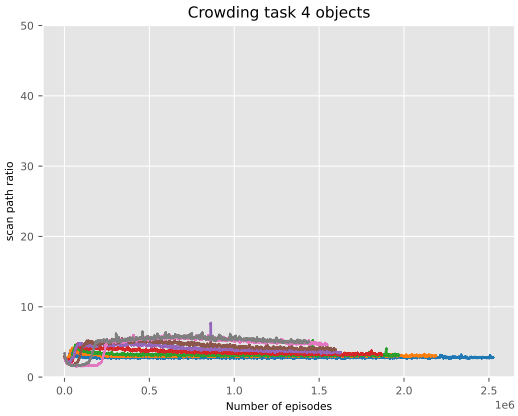

(e) Scan path ratio during learning

Figure 5.5: The figures demonstrate the learning performance (with Gaussian noise) for four objects. As can be seen, using this group of noise levels, the model converged and learned the optimal policy smoothly. It learned the appropriate number of actions (fixations) in a given state with the distribution of noise provided in the observation. It can also clearly provide an illustration of learning an officiant scan path ratio with appropriate search time (Figures 5.5e and 5.5c).

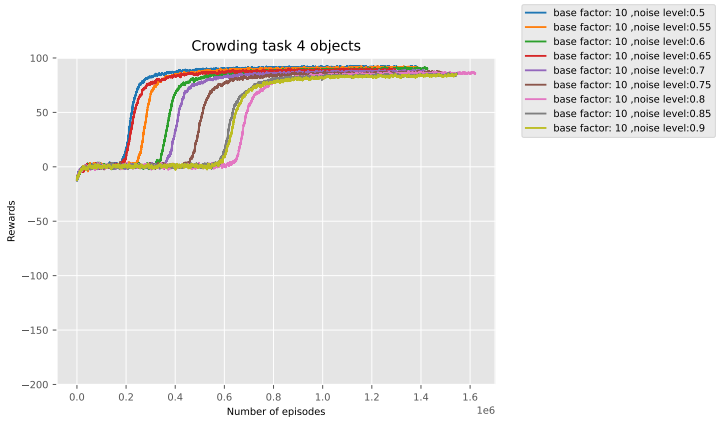

(a) Rewards during learning

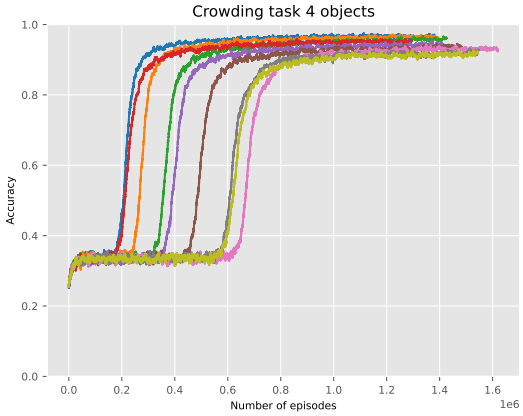

(b) Accuracy during learning

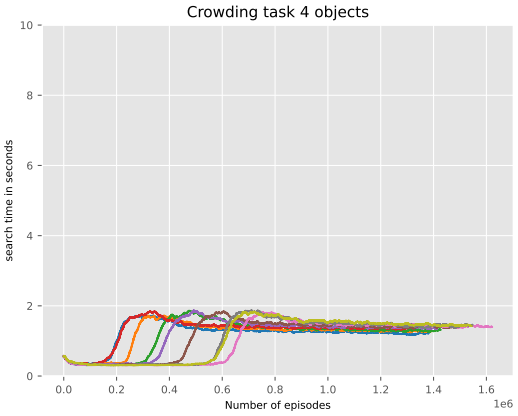

(c) Search time during learning

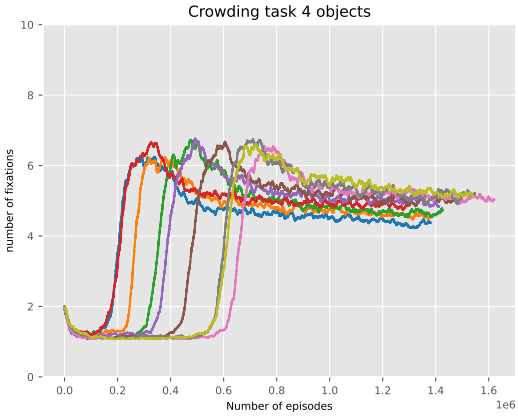

(d) Accuracy over time.

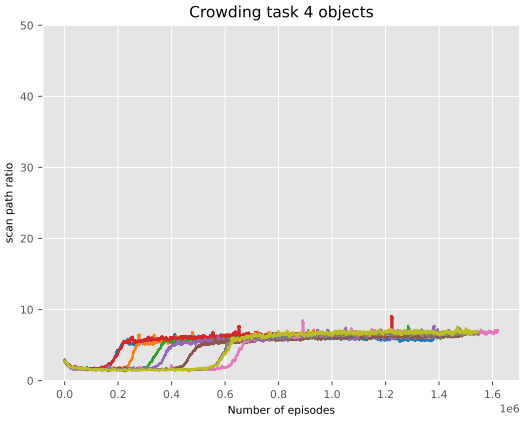

(e) Scan path ratio during learning

Figure 5.6: The figures show the rest of the noise levels for experimenting with Four objects using Gaussian noise. As can be seen, these figures show the learning with higher noise levels than in Figure 5.5. Increasing the noise leads to lower accuracies, increased fixations, longer search time and insufficient scan path ratio. Particularly, the scan path ratio shows that the model's path to find the target increased in comparison to Figure 5.5e.

5.8.2 Four Objects: Evaluation of the Crowding Effect (Gaussian)

This subsection presents the results after testing the model on crowded and uncrowded conditions.

It can be seen in Figure 5.7a that the model showed a reduction in accuracy with increased noise levels for the crowded condition, which is similar to the categories in the human data. However, the model outperformed the human data and produced high accuracies despite increased noise levels for the uncrowded condition. This is believed to be due to the high level of optimisation in the reinforcement learning model and that the model was able to adapt perfectly in uncrowded conditions despite higher noise in the stochastic environment. However, it can be observed in Figures 5.7b, 5.7c and 5.7d that the model chose appropriate learning behaviours in the uncrowded condition, which led to very good accuracy. Figure 5.7b also illustrates the evaluation in search time against the human data. The model showed a similar increase pattern in search time as the human categories with increased noise levels, especially in crowded conditions. However, the model outperformed the human data in crowded and uncrowded conditions.

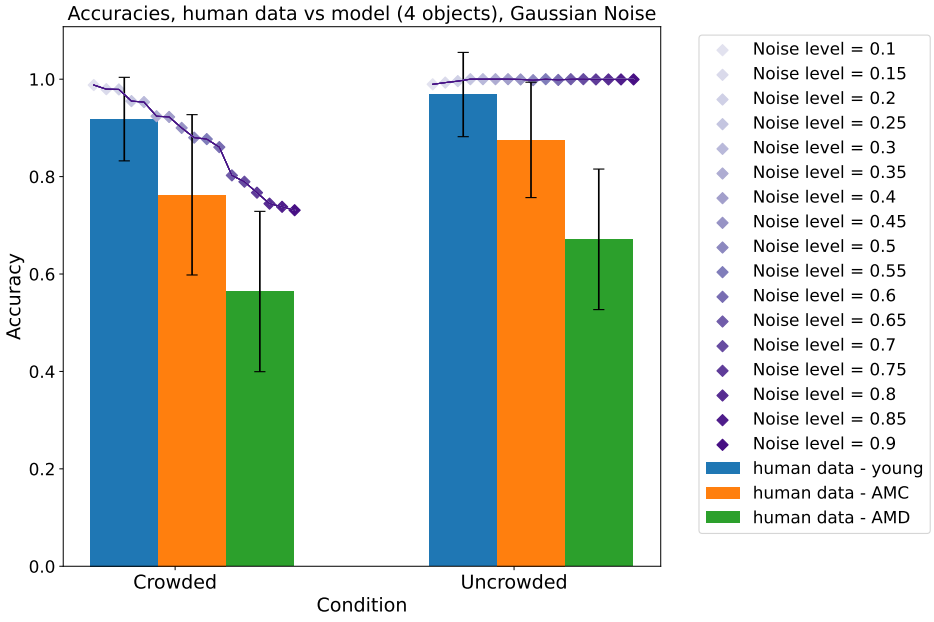

(a) Accuracy during testing

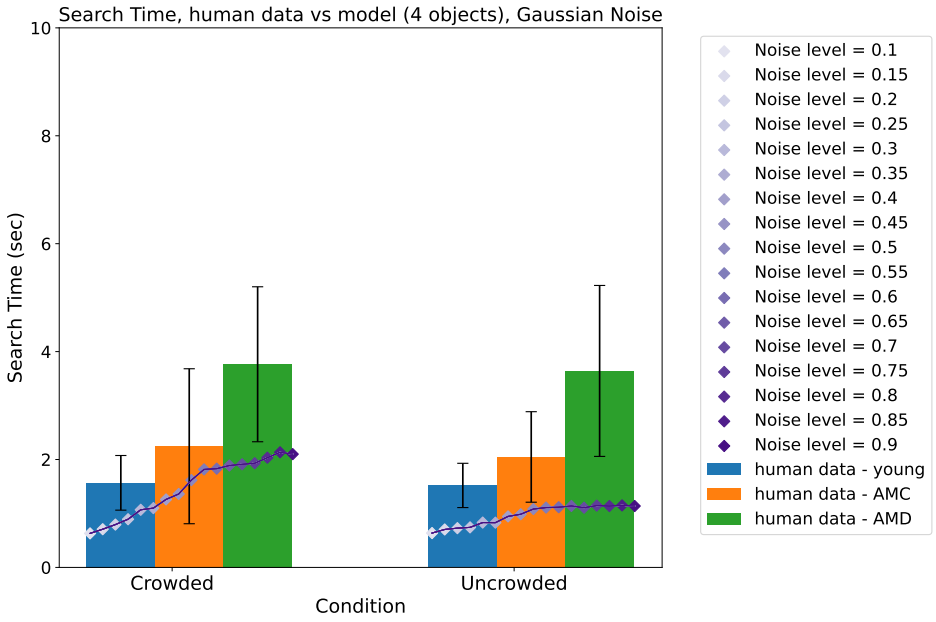

(b) Search time during testing

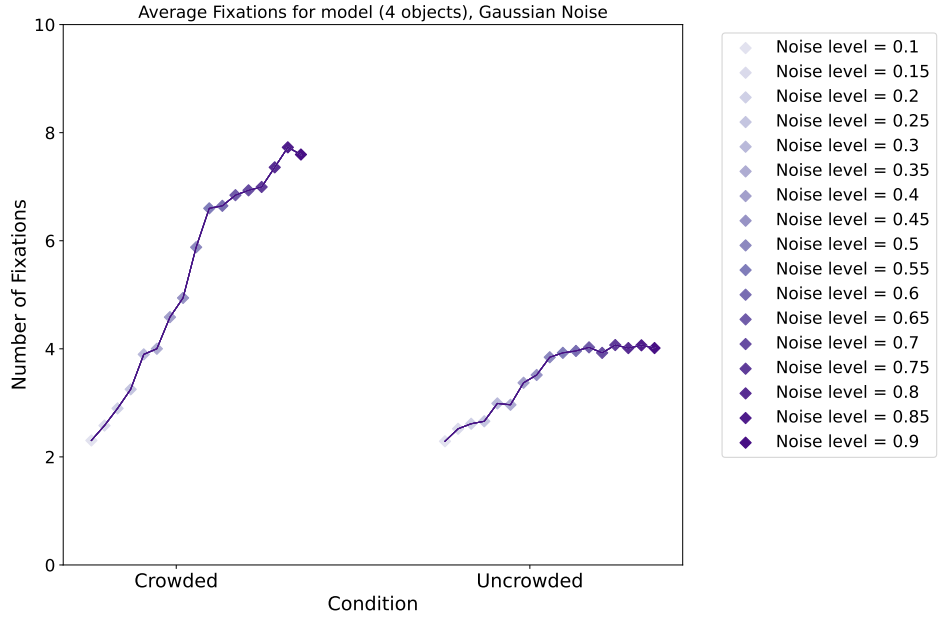

(c) Number of fixations during testing

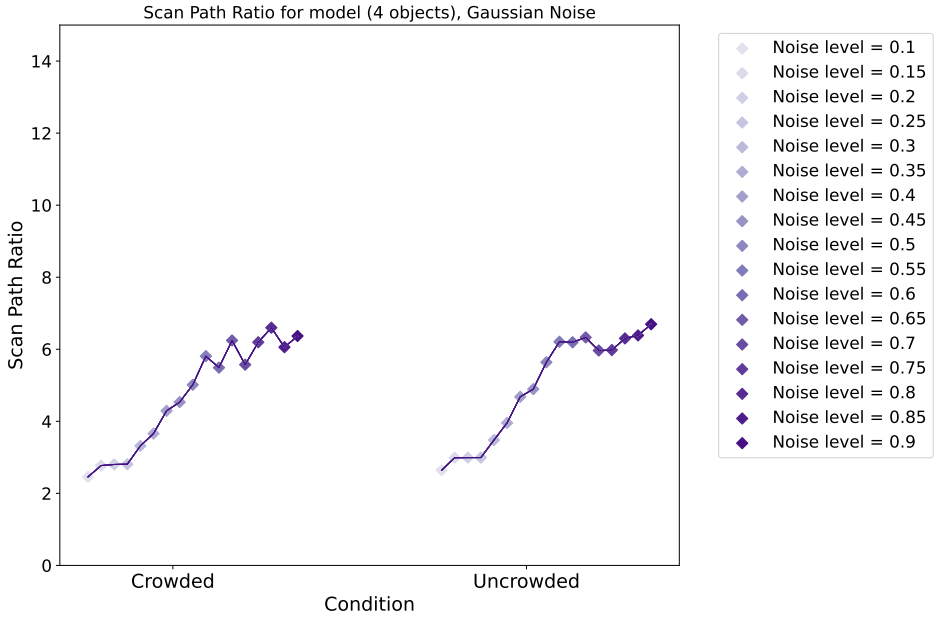

(d) Scan path ratio during testing

Figure 5.7: This figure illustrates the testing results for the four chosen measures in crowded ($S = 1.5$) and uncrowded ($S = 6$) conditions . Figures 5.7a and 5.7b were evaluated against human data.

The results from this experiment can give some insights into the visual behaviours of people with AMD and match our hypothesis to a certain extent. The results suggest that AMD participants experienced more noise during the visual search task than the young and AMC categories for the crowded conditions. It also suggests that people with AMD would make more fixations, and have a less efficient scan path ratio. As the model produced a high accuracy for the uncrowded condition in all the experiments, we will only discuss it in this section.

To find the best noise levels that fit the human data, we created a plot that illustrates the squared difference between the model's results and the human data for both accuracy and search time (see Figure 5.8).

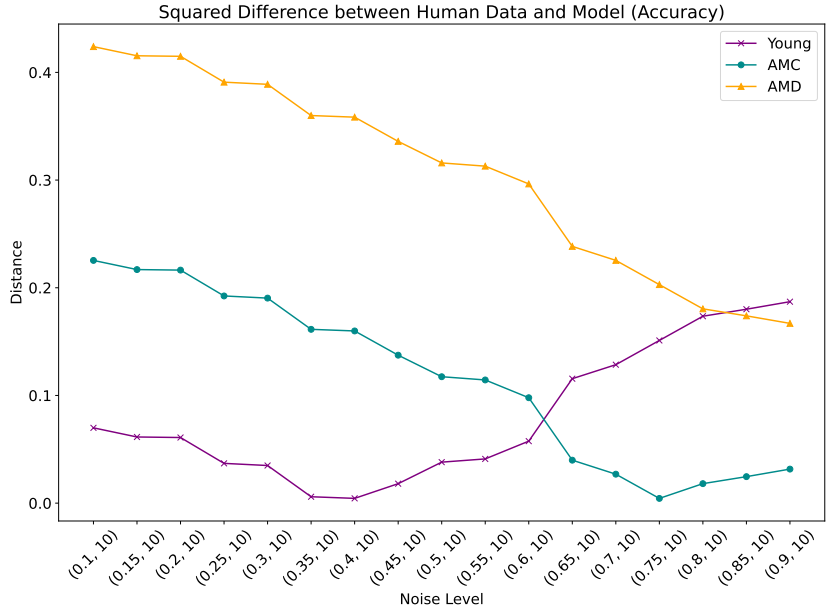

(a) This diagram illustrates the results of the squared difference between accuracies in the model and in the human data. The local minimum can be observed for the young category (located at (0.4,10) with (0.35,10) being at a similar level). The local minimum for AMC is at noise level (0.75,10). Whereas for the AMD participant, the lower point is at the highest noise (0.9,10), where 10 is the *base crowding factor* that we add for illustrate the true value of the *Crowding Factor* (See table 5.1)

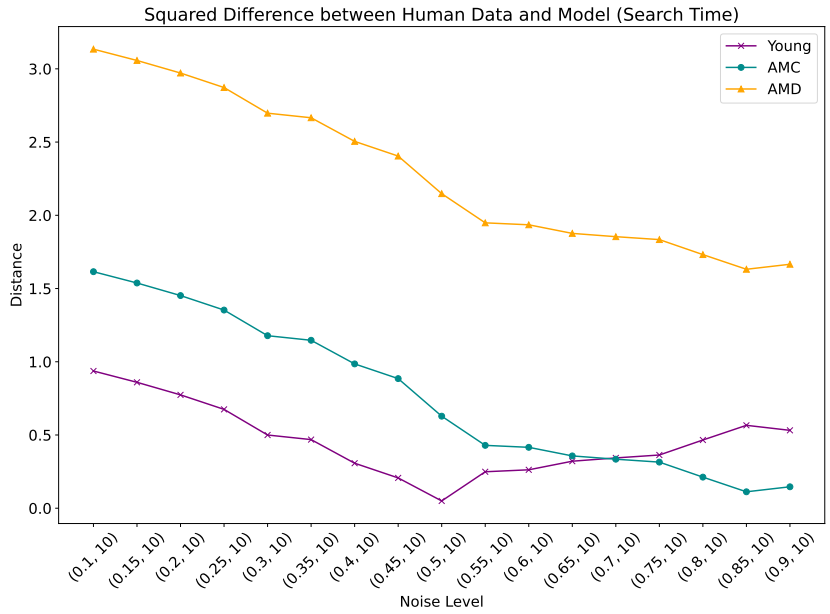

(b) This plot illustrates the local minimum for search time, and it can be seen that the search time in the model is closer to young and AMC groups than those with AMD

Figure 5.8: Squared difference between model results and human data (four objects Blending)

In the next subsection, we illustrate the results of the Blending noise experiment.

5.8.3 Four Objects: Model Performance (Blending)

In this subsection, we illustrate the model's performance with the Blending noise $c^{(B)}_{t,obj} \sim \mathcal{N}(\mu = average\ crowding, \sigma^{(B)}_{o,c})$. The main reason for implementing this noise in the model is to imitate the observation behaviour resulting from the crowding effect when perceived by the human eye. The model's learning behaviours have similarities to the ones presented in Section 5.8.1. However, due to more complex observations in this experiment, the accuracies appear lower than those in the previous subsection. The figures also show that learning using the blending noise required longer episodes throughout the learning iterations (this can be observed by comparing the x-axis length with the Figures 5.5 and 5.6). This indicates that the agent needed to explore the environment more due to increased complexity. The figures also show that the model converged to a policy with higher fixations than the Gaussian noise and a less efficient scan path ratio, which can describe imperfection in the strategies, leading to closer behaviours to the human. The histogram for the blending crowding noise for the experiment is provided in the Appendices (see Appendix A Section A.1.2).

Training the model using the Blending noise led to increased complexity in the noisy observations. For this reason, the model sometimes does not explore the environment and decides immediately on the final decision action rather than learning to fixate in a given belief state. When this happens, a second attempt to re-train the model with the same noise levels was essential to see if the model managed to learn in the second attempt. The effect of increased noise (e.g. motor noise) was addressed in a variety of studies. For example, adding high levels of motor noise to human participants in a RL task (Therrien, Wolpert, and Bastian, 2018) caused impaired RL in the participants. In a different study that considered an artificial environment (similar to this model) increased motor noise showed increased decision-making difficulties (Chen, Mohr, and Galea, 2017). Increase in noise was also tested to study its effects on the stereotypy of saccadic eye movements (van Beers, 2008), which showed that increased noise could decrease an individual's accuracy and increase the chance of missing saccades' desired

destination, whereas another study suggested that different types of noise can influence the human saccades when searching for a target (Van Beers, 2007).

The noise levels that the model did not learn from the first attempt in this experiment are 0.6 and 0.75. However, the model learned with higher noises from the first attempt, which indicates random behaviours in choosing the actions caused by the noisy observations. For example, the agent never learns to fixate and collect more information because it kept choosing the decision action (to press the space bar) without making any exploration in the environment due to the noise. The rest of the figures related to this performance are added to the appendices (see Appendix A, section A.2).

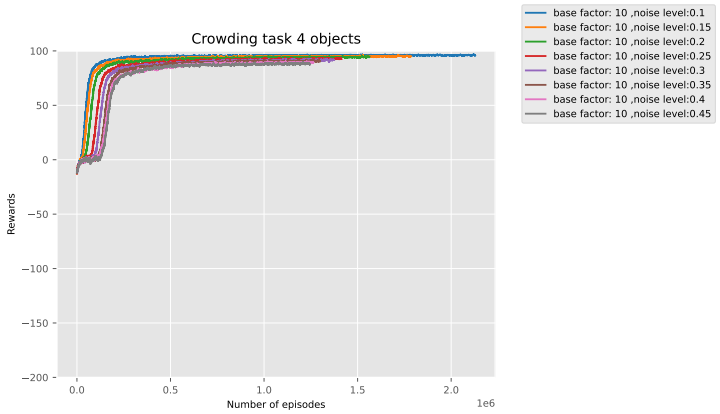

(a) Rewards during learning

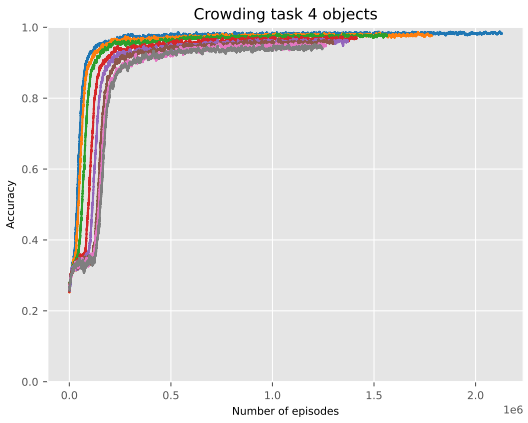

(b) Accuracy during learning

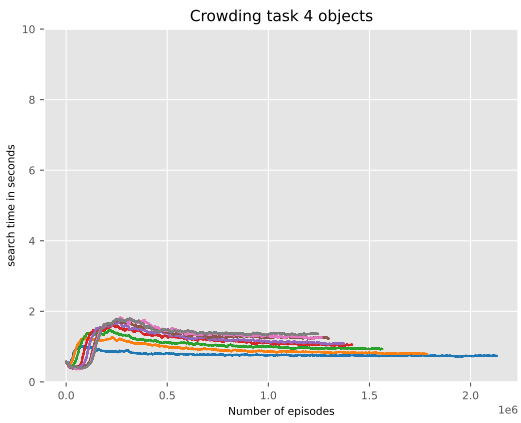

(c) Search time during learning

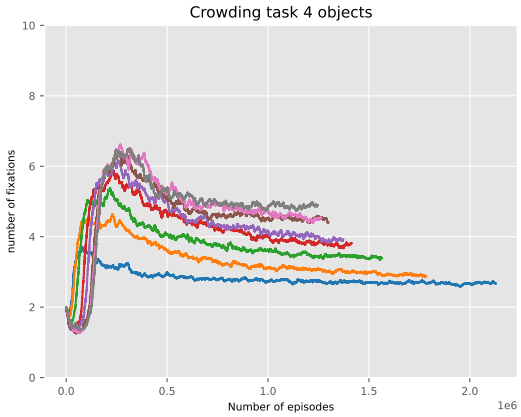

(d) Number of fixations during learning

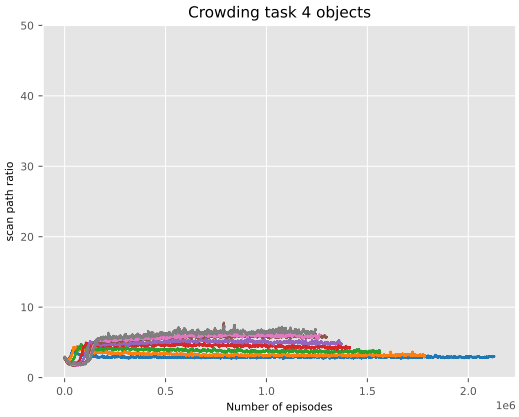

(e) Scan path ratio during learning

Figure 5.9: Learning performance with Blending noise (four objects)

5.8.4 Four Objects: Evaluation of the Crowding effect (Blending)

Figure 5.10a indicates that the model showed closer accuracy to the human data, especially for the AMD. However, the model did not show increase in the search time with higher noise levels, and the search time increased to levels that match the AMC group. The number of fixations, on the other hand, increases with noise levels, but it can be noticed that this increase in not is not always consistent (see Figure 5.10c). This inconsistent behaviour can also be observed in the scan path ratio (Figure 5.10d), and search time in Figure 5.10b. We believe that this is mainly due to decreased uncertainty in the model and the blending of features in the observation (in addition to the constraints introduced in the environment) leading to inconsistent behaviours, which are more similar to the nature of the human behaviours (Son, Mack, and Walther, 2023).

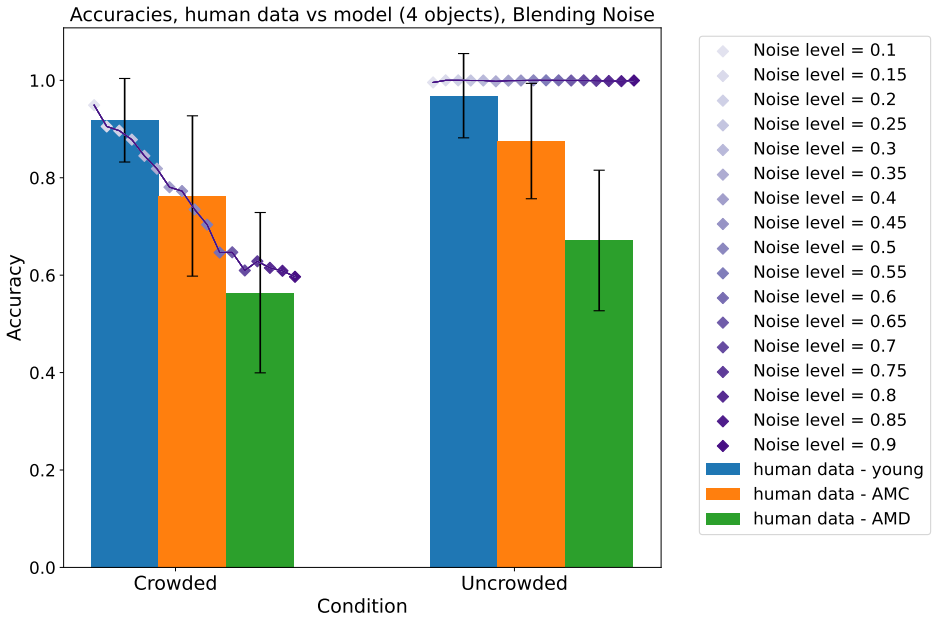

(a) Accuracy during testing

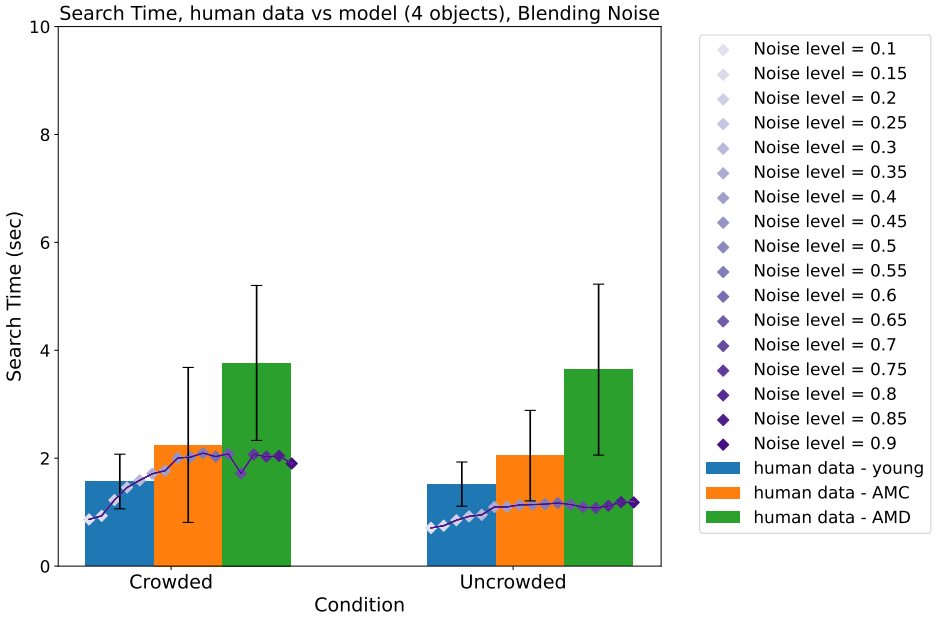

(b) Search time during testing

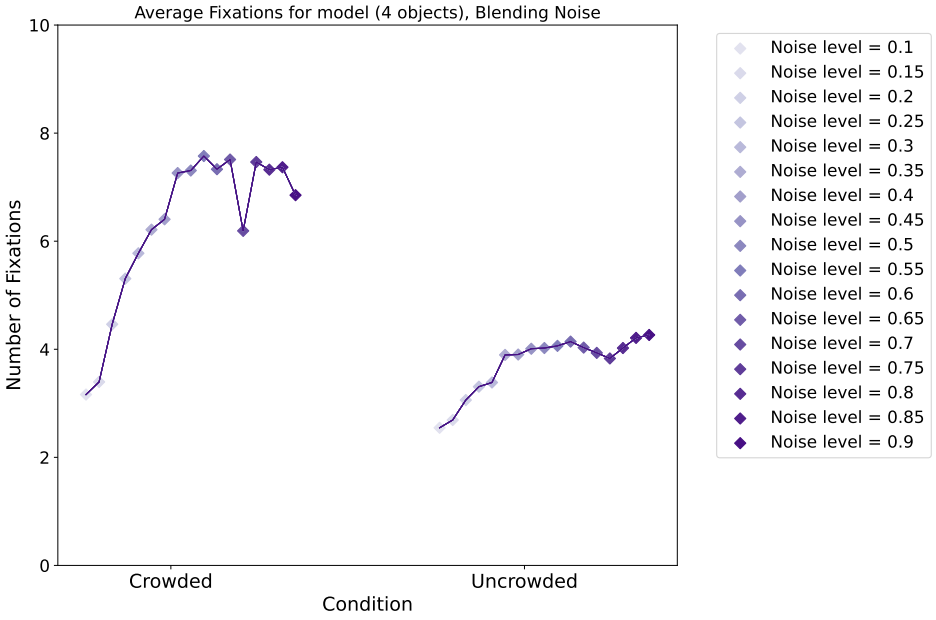

(c) Number of fixations during testing

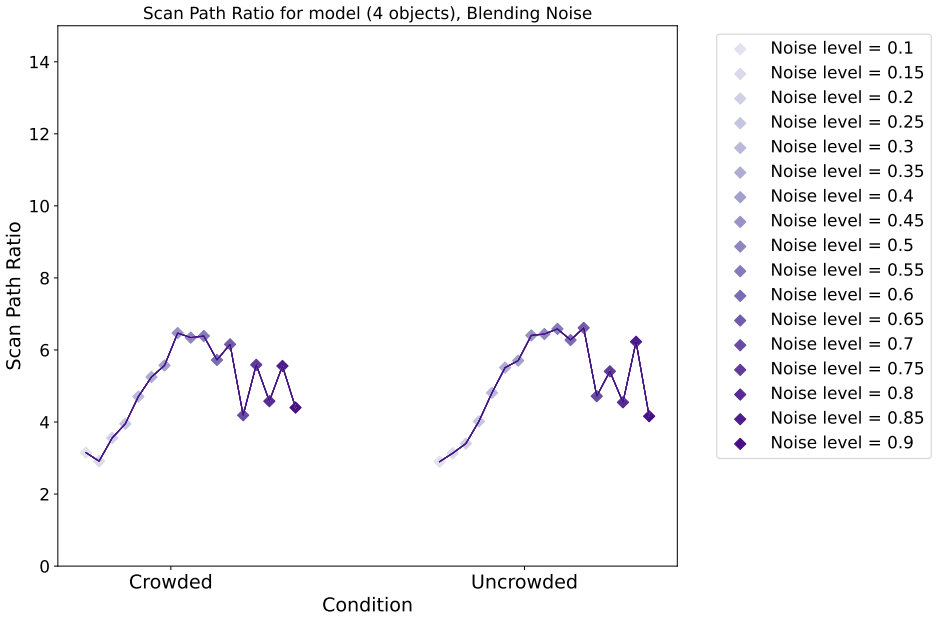

(d) Scan path ratio during testing

Figure 5.10: Testing results for crowded and uncrowded conditions for accuracy, search time, number of fixations and scan path ratio.

The best noise fit to the human data is illustrated in Figure 5.11.

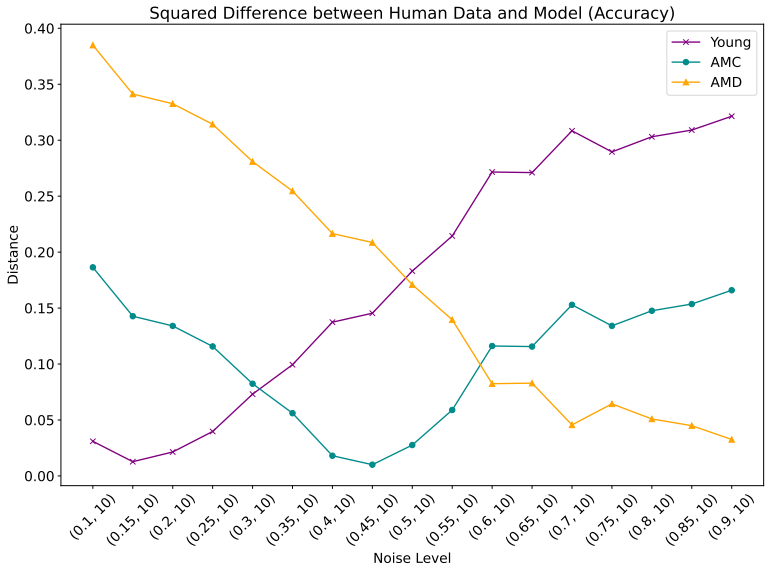

(a) In this diagram the local minimum for young and AMC is very clear: (0.15,10) for young and (0.45,10) for AMC. The effect of the blending noise decreases the difference between the model results and human data for AMD.

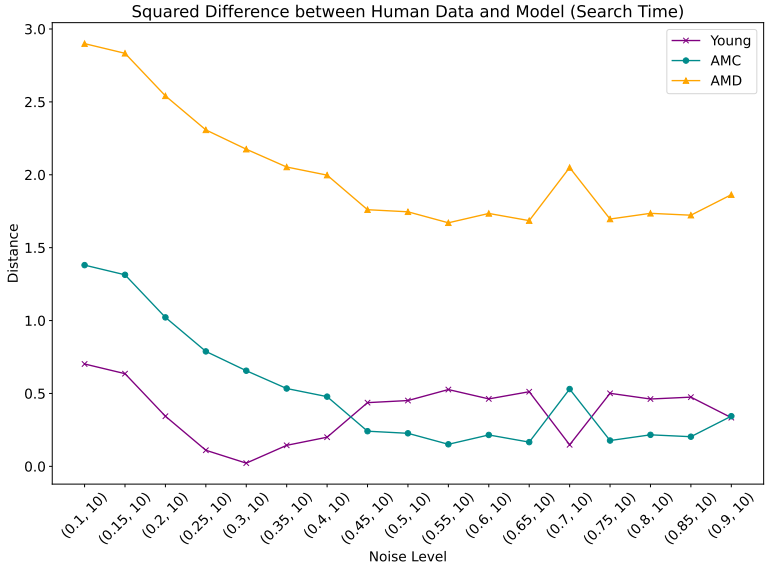

(b) This diagram shows the squared difference in search time with a clear indication of the best for young and AMC participants. However, the distance to the AMD participants is high

Figure 5.11: Squared difference between model results and human data (four objects Blending)

In the next section, we show the experiment with six objects in the display.

5.8.5 Six Objects: Model Performance (Gaussian)

In this subsection, we present the model's learning performance when trained with six objects using Gaussian noise. Due to increased complexity (increased noise levels and number of objects), we expected that the cases where the model does not learn to fixate (collect information) and only learn to press the space bar (like some cases we discussed in section 5.8.3) would also occur in this experiment. For this reason, we separated the learning into smaller groups to minimise the curves with no optimal strategies in the plot. The first noise group contained the following noise levels: $[0.1, 0.15 \ldots 0.45]$. The learning performance is illustrated in Figure 5.12. The model did not learn the strategies when given the noise levels 0.75, 0.85 and 0.9 and re-training a second time did not change the outcome for these conditions. Learning performance for the rest of the noise levels in this experiment is provided in Appendix A, Section A.3.

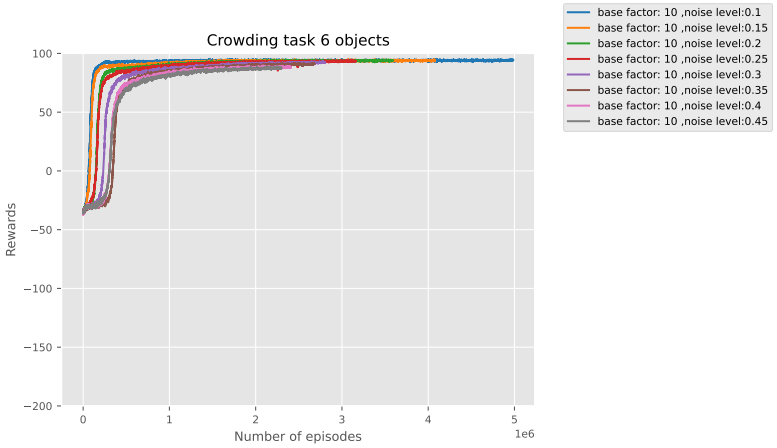

(a) Rewards during learning

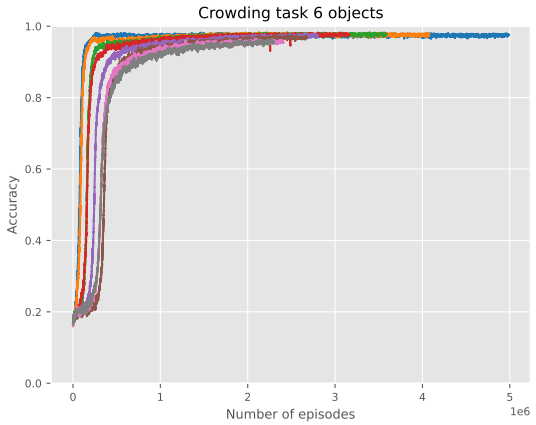

(b) Accuracy during learning

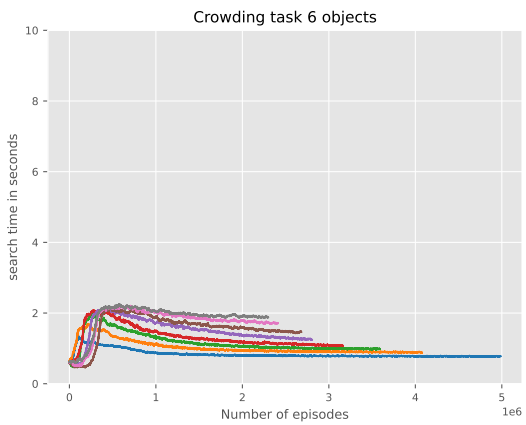

(c) Search time during learning

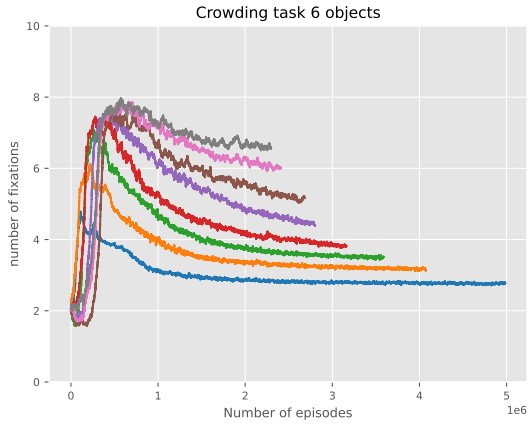

(d) Number of fixations during learning

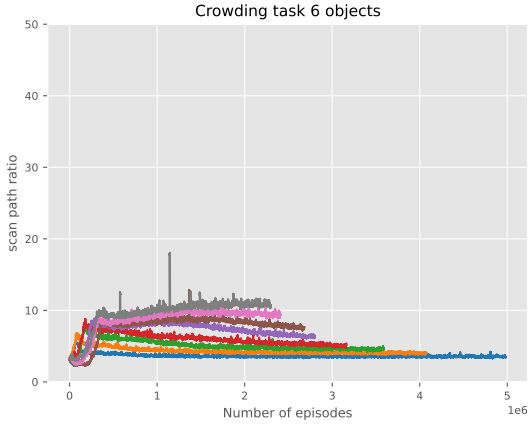

(e) Scan path ratio during learning

Figure 5.12: The figures demonstrate the learning performance (with Gaussian noise) for six objects.

Increasing the number of objects affects the learning performance very clearly as it can be observed when comparing Figures 5.5 and 5.12. One interesting observation is that the model had longer learning episodes while learning for six objects (each episode included more actions than in four objects). One way to observe this visually is by comparing the learning for noise level 0.45. The shorter curve in the six objects experiment indicates the need for the model to explore the environment more by taking more actions in every state. This can also be observed when comparing the number of fixations the agent needed after convergence in each experiment (four fixations for four objects and around seven fixations in the case of six objects) (see Figures 5.5d). Furthermore, this can explain the inconsistent scan path ratio in Figure 5.12e and the increase in the search time in Figure 5.12d.

In the next subsection, we evaluate the results by testing with the crowded and uncrowded conditions and compare the model's strategies against the human data.

5.8.6 Six Objects: Evaluation of the Crowding effect (Gaussian)

Evaluating the results for the crowded and uncrowded conditions (Figure 5.13) showed similar accuracy to the three human categories. Additionally, increasing the number of objects in the model, led to increased complexity, causing more similar behaviours to the human data. For example, there was an increased number of fixations with the increased uncertainty in the

observations and inefficient scan path ratios (Figures 5.13c and 5.13d). It is also noticeable that the scan path ratio does not always worsen with increased noise levels. This is due to the high noise in the observations, causing inconsistent patterns. The search time, on the other hand (Figure 5.13b) shows a consistent increase with the increase in the noise levels. However, there is a slight increase in the noise level of 0.8. Nonetheless, the model has a faster search time than those with AMD regarding the search duration.

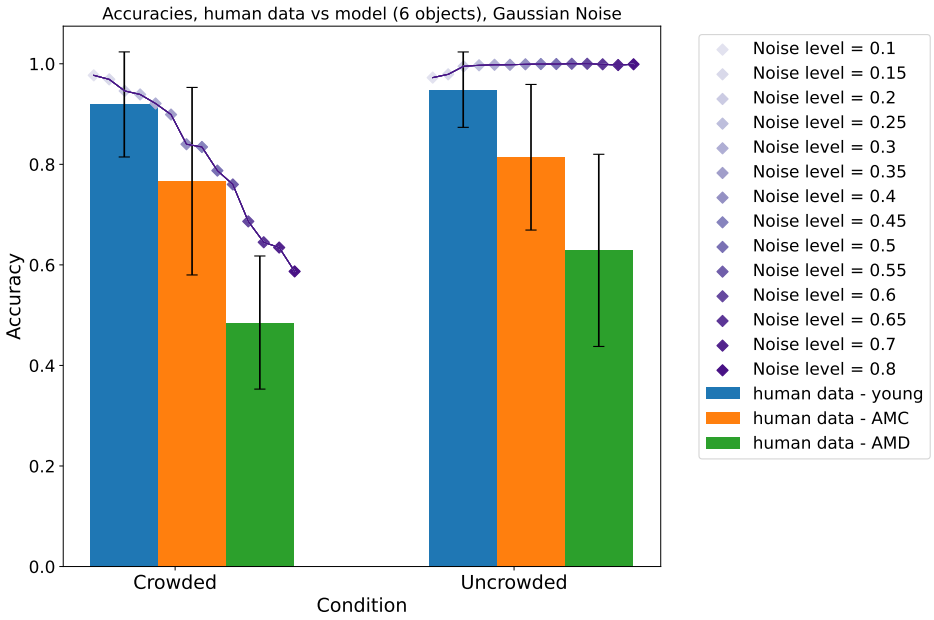

(a) Accuracy during testing compared to human data

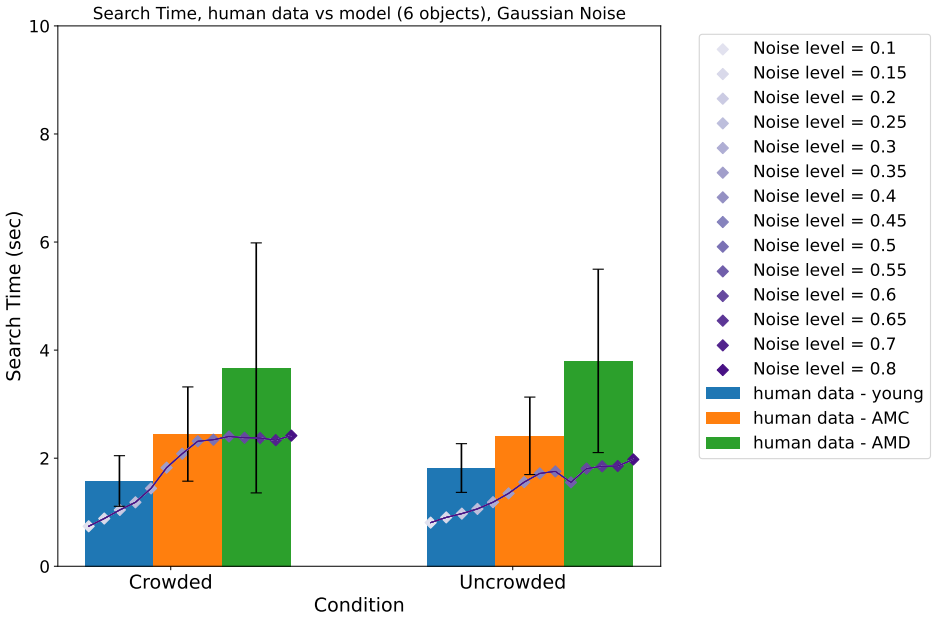

(b) Search time during testing compared to human data

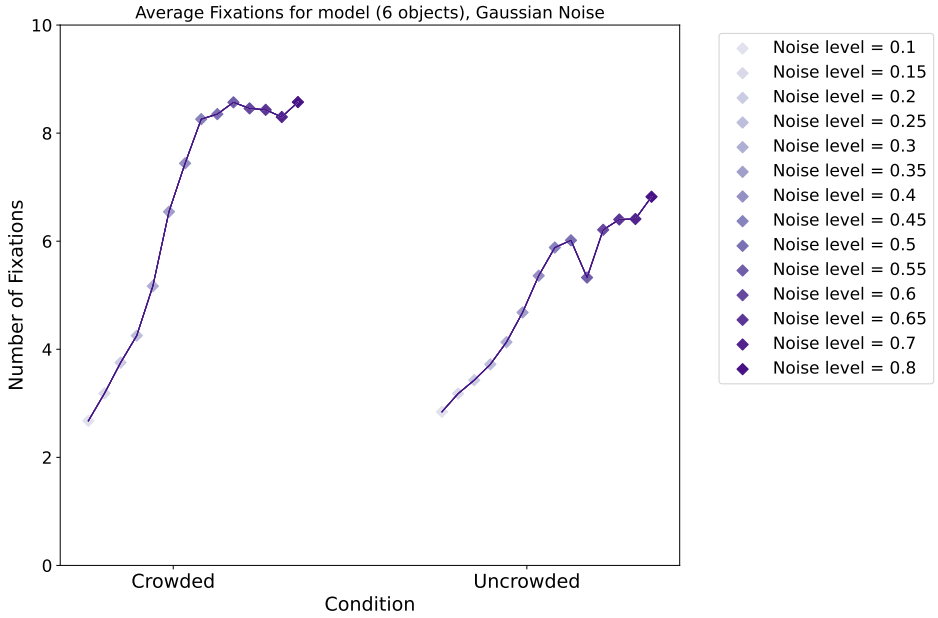

(c) Number of fixations during testing

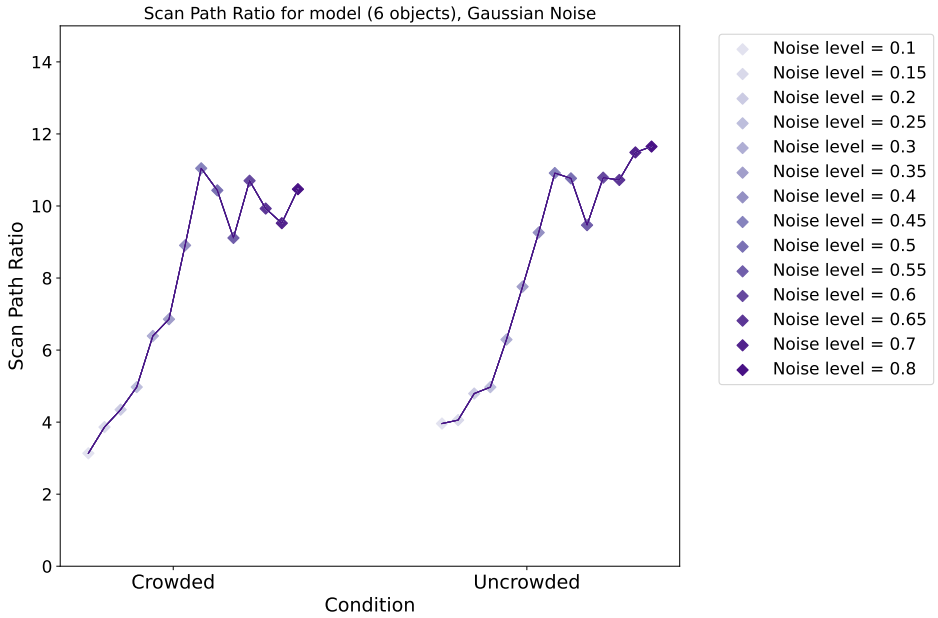

(d) Scan path ratio during testing

Figure 5.13: This figure illustrates the testing results for the crowded and uncrowded conditions (six objects, and using Gaussian noise)

.

The best noise fit to the human data can be illustrated in Figure 5.14.

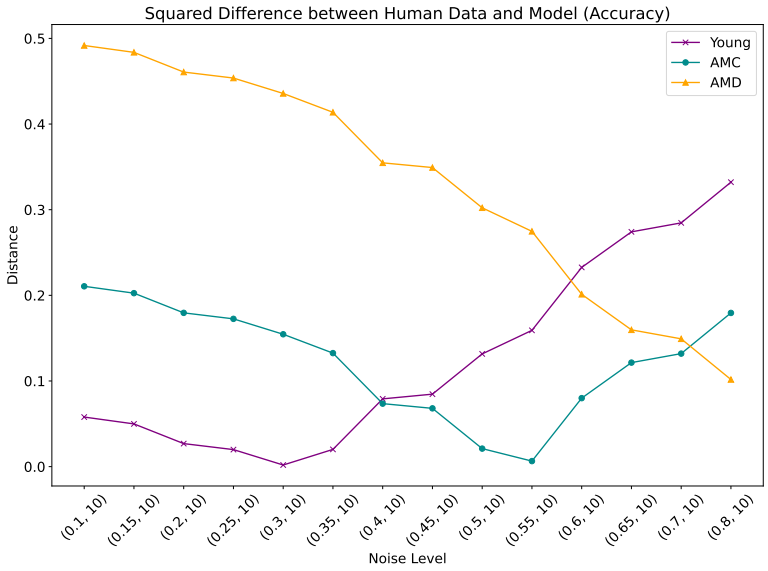

(a) The figures show the distance between accuracy in the model and human data. As it can be seen, the young had a low noise for the best fit (0.3,10), the AMC group a slightly higher noise (0.55,10), and for the AMD the best fit is at the highest noise (08,10).

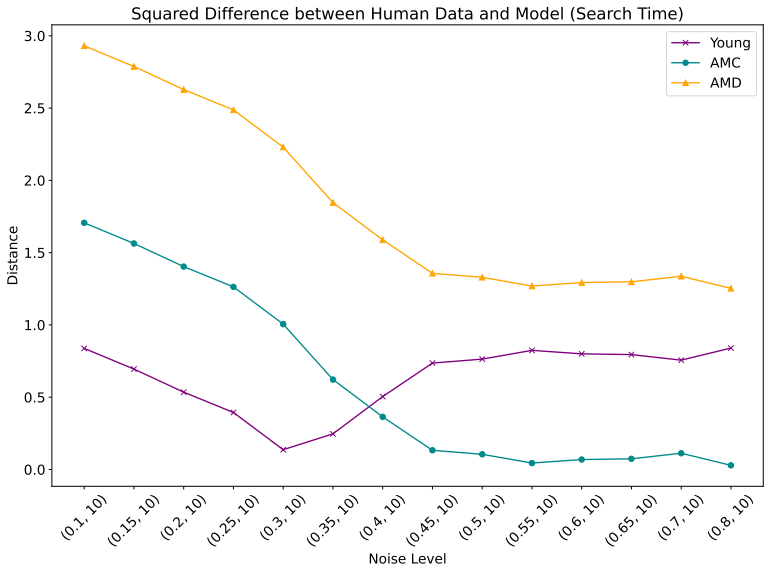

(b) The figures show the distance between search time in the model and human data. The only clear fit is for young participants (0.3,10). The AMC group had a very similar search time across different noises over (0.45,10), whereas, the distance remained big for the AMD group

Figure 5.14: Squared difference between model results and human data (six objects, using Gaussian noise).

In the next subsection, we present the model performance with the Blending noise and six

objects in the display.

5.8.7 Six Objects: Model Performance (Blending)

Due to the increased complexity of the Blending noise, we expected more inaccurate strategies (by the model always choosing to press the space bar). Therefore, we separated the learning into smaller groups of noise levels. The first noise group contained the noise levels $[0.1, 0.15, 0.2, 0.25]$, and the learning performance for those noise levels is presented in Figure 5.15. There are some noise levels where the model only learned the optimal strategies after the second (and sometimes the third) running attempt to the same reason we explained in section 5.8.3. The re-run of those noise levels is provided in the appendices (see Appendix A, Section A.4). However, like in the Gaussian noise, in this experiment, there are some noise levels that the model that did not produce the desired strategies; and for this reason, those noise levels were excluded from the experiment.

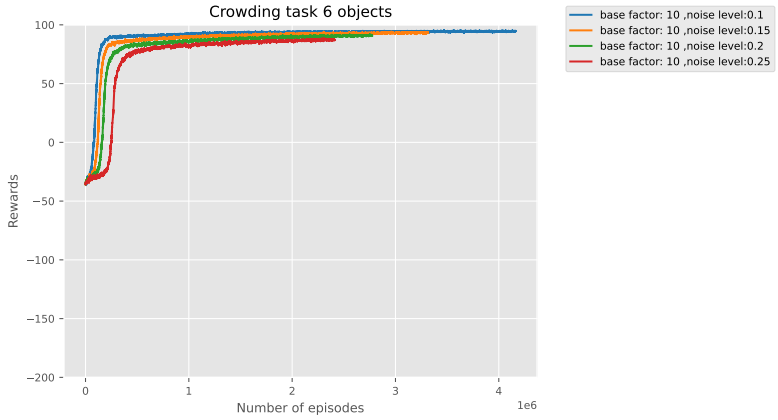

(a) Rewards during learning

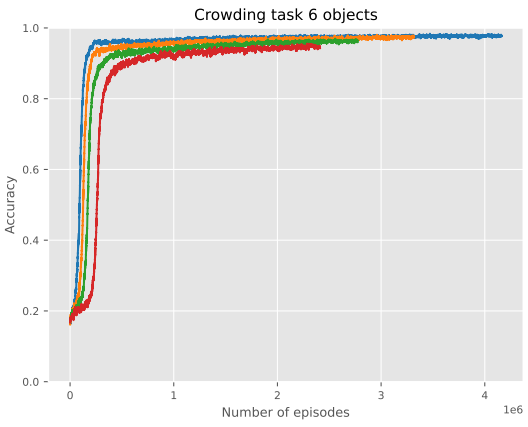

(b) Accuracy during learning

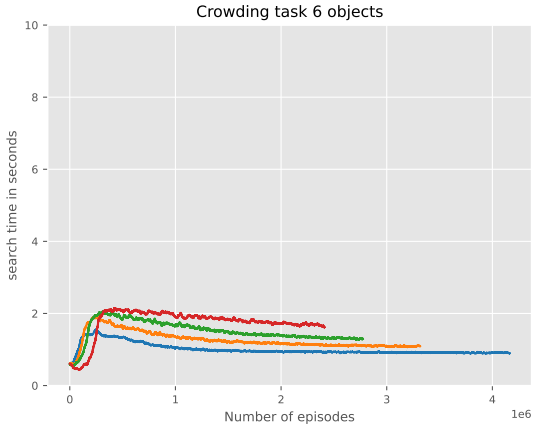

(c) Search time during learning

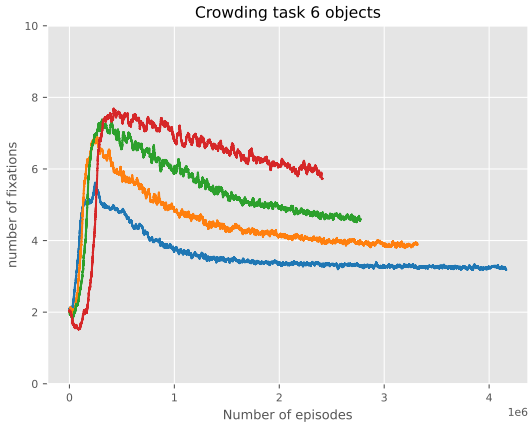

(d) Number of fixations during learning

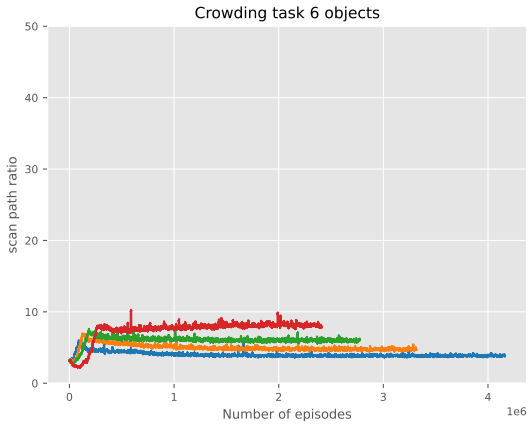

(e) Scan path ratio during learning

Figure 5.15: The figures demonstrate the learning performance (with Blending noise) for six objects.

The effect of the Blending noise in six objects is very similar to the four objects in Section 5.8.3. There is a noticeable increase in the duration of each episode (hence a shorter x-axis). The increase in the fixation actions at lower noise levels can also be observed. For example, for noise level 0.25, the model converged for four fixations as the best strategy when the noise type was set to Gaussian (Figure 5.12d), but for the same noise level using the Blending noise, the number of fixations increased to 6 (see Figure 5.15d). This also explains the increased search time and the lower accuracy for this set of noise levels (Figures 5.15c and 5.15b).

In the next subsection, we evaluate the model's strategies in crowded and uncrowded conditions and compare them against the human data.

5.8.8 Six Objects: Evaluation of the Crowding Effect (Blending)

The testing results with the Blending noise (see Figure 5.16) showed an additional decline in the accuracies in the crowded condition than the results with Gaussian noise if we compare Figure 5.13a with Figure 5.16a. Additionally, the increase in the fixations (Figure 5.16c) is prominent in the crowded condition. It is important to note that the decrease in scan path ratio at higher noise levels (Figure 5.16d) does not indicate an optimal scan path at those higher noise levels (e.g. 0.6 and 0.7), but rather due to higher complexity, and the result from an unstable learning curve (see Figure 5.15e). In this experiment, the model produced a very good fit to the humans'

accuracy, yet the search time did not increase with the increased noise levels.

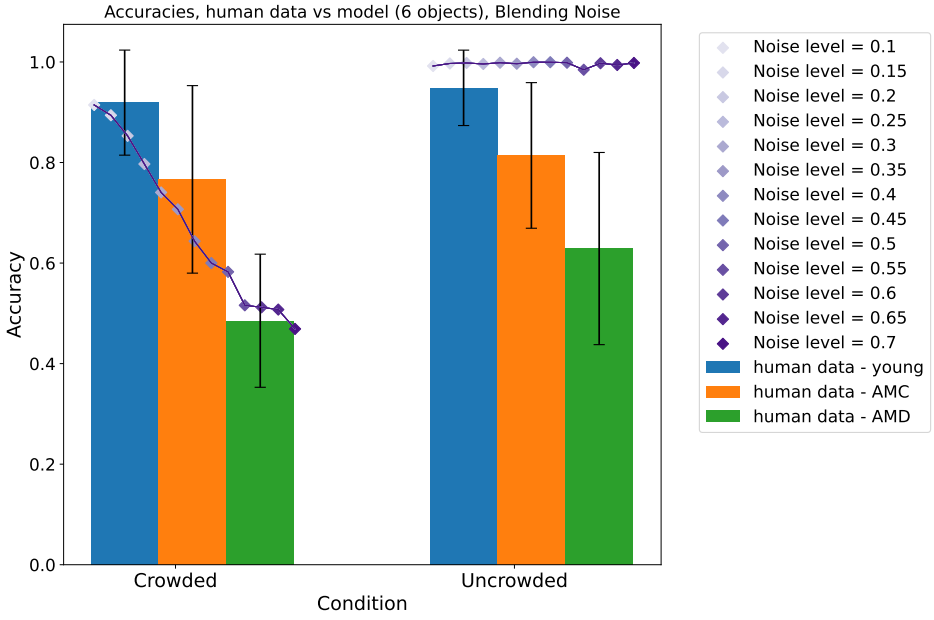

(a) Accuracy during testing compared to human data

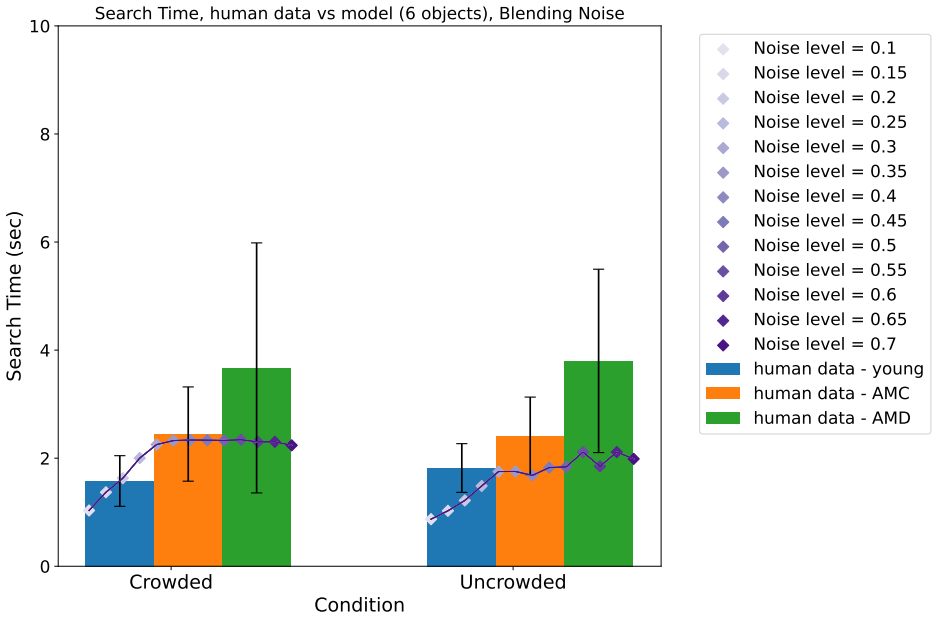

(b) Search time during testing compared to human data

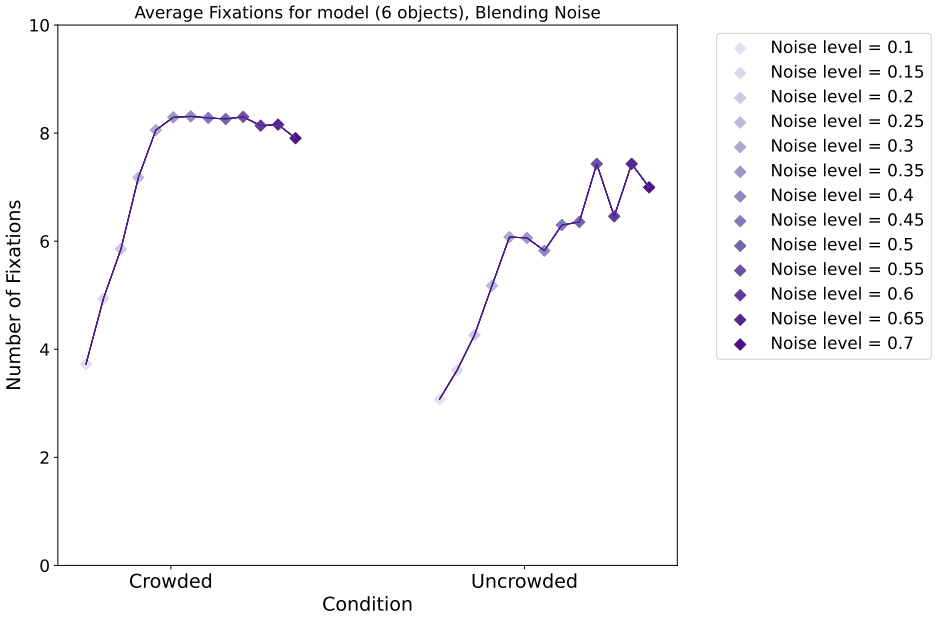

(c) Number of fixations during testing

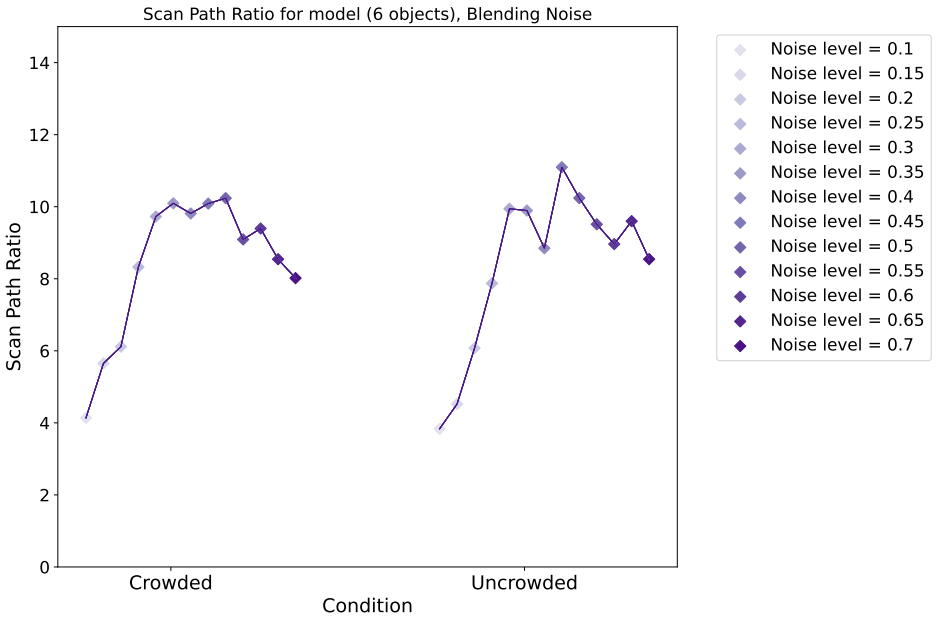

(d) Scan path ratio during testing

Figure 5.16: This figure illustrates the testing results for the crowded and uncrowded conditions (six objects, using Blending noise)

.

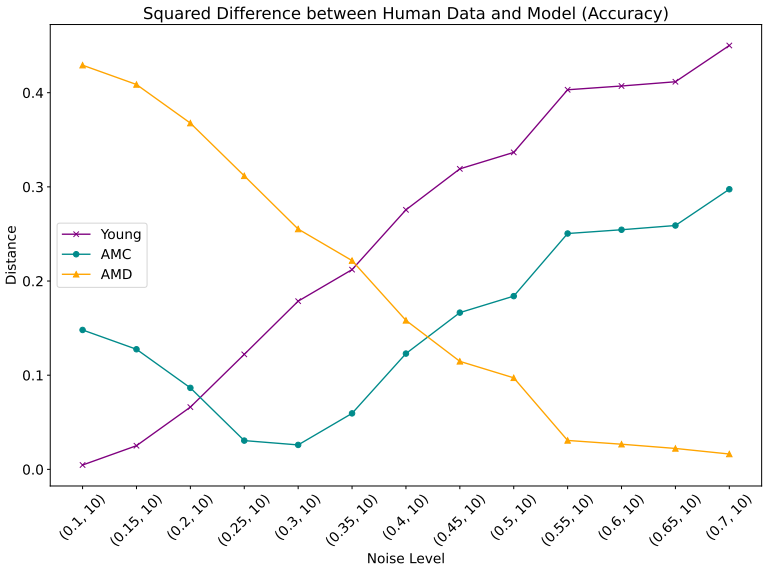

(a) In these results, the young group had the lowest noise level for the best fit (0.1,10), while the AMD group had the highest noise level (0.7,10)

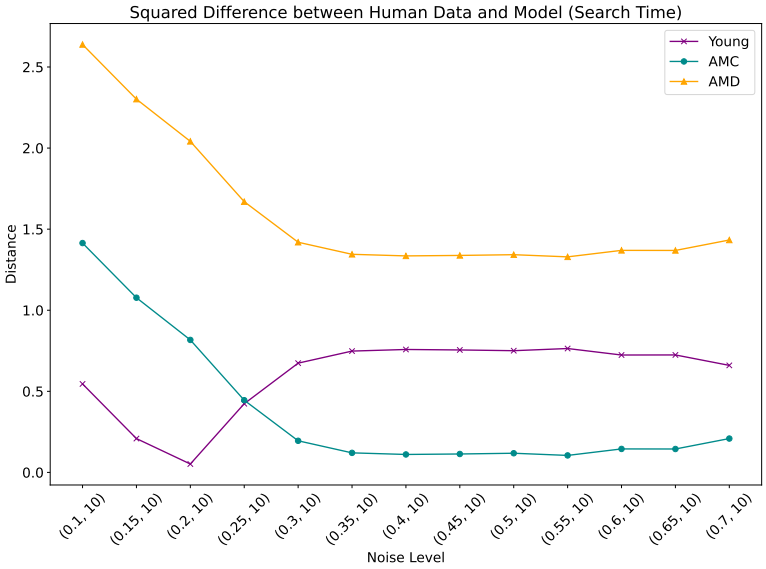

(b) The search time only shows the best fit for the young participants. For the AMD, it can be seen that a group of noise can fit the data $[(0.35, 10) \ldots (0.6, 10)]$. While the difference between the model and the AMD group remains high, yet plateauing similar to the other participants'.

Figure 5.17: Squared difference between model results and human data (six objects Blending)

In the next subsection, we compare all the testing results to provide a more in-depth comparison of the model's strategies and compare it with the analysis of the result in Thibaut, Boucart,

and Tran (2020).

5.8.9 Overall Results

In this subsection, we present all the results specifically to compare the following visually:

1. The effect of increasing the number of objects on accuracy, search time, number of fixations and the scan path ratio

2. The difference between the two types of noises (Gaussian and Blending). Mainly, we compare how the blending noise showed approximately, similar strategies to the people with AMD as well as the other categories with healthy visual fields.

Accuracy

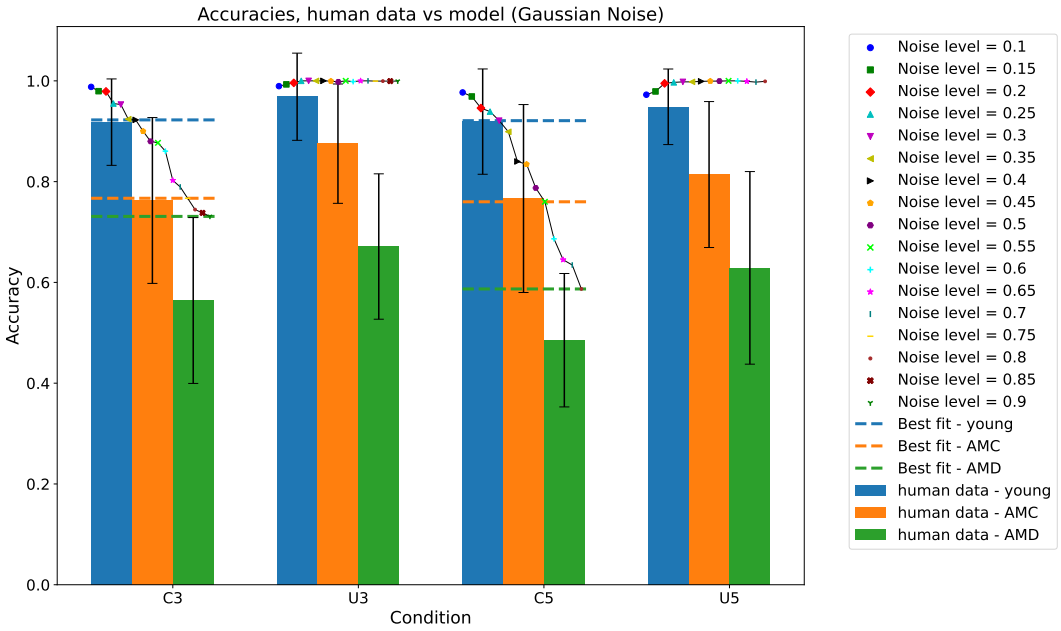

(a) Accuracies, using Gaussian noise. The effect of increasing the number of objects on accuracy is noticeable in this figure. However, the model slightly outperformed the human data, especially at C3.

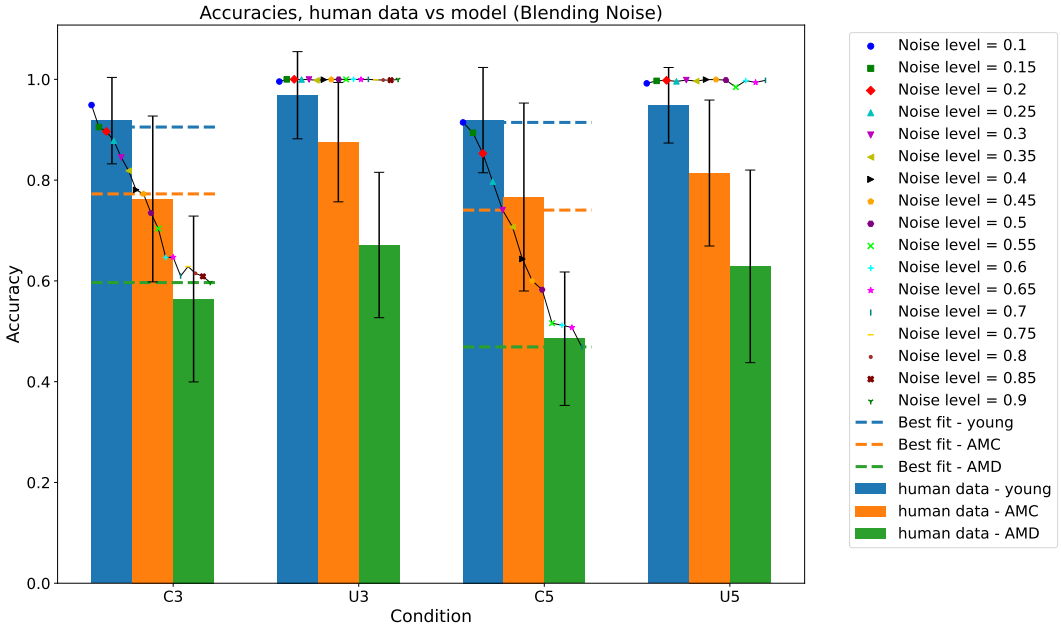

(b) Accuracies, using Blending noise. In this figure the effect of increasing the number of objects is also observable. Additionally, the results show a better fit to the human data, especially in the case of C5.

Figure 5.18: A comparison between the average accuracies and the effect of increasing the number of objects. The figure also compares the overall results with the different noise types.

Search Time

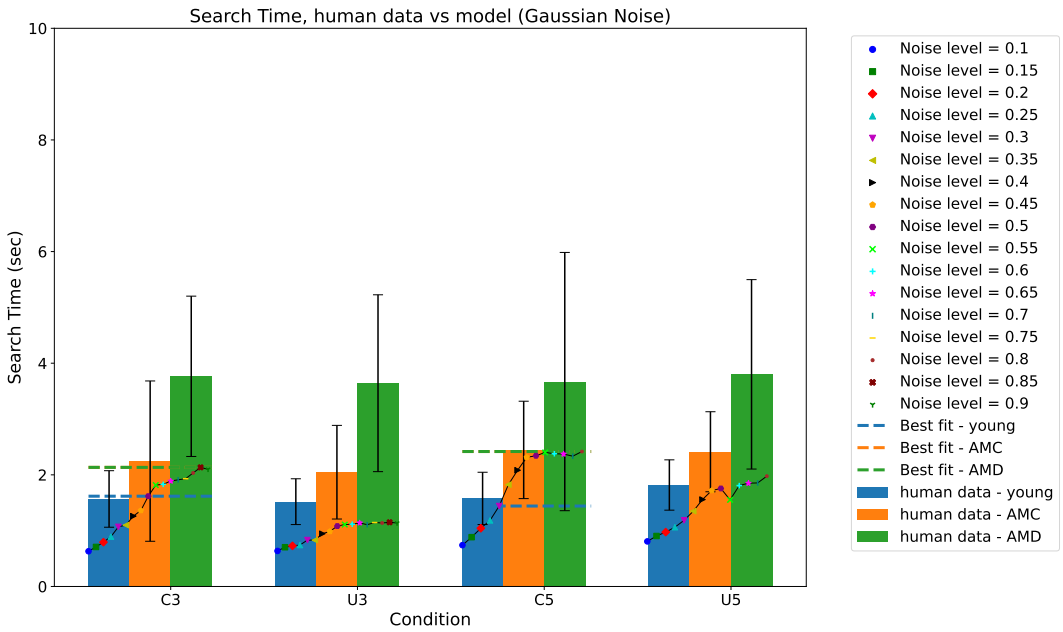

(a) Search time, using Gaussian noise. There is increased search time with the increased number of objects. The increase pattern is similar to the human data. However, the model was faster than the AMD group despite the higher noise levels.

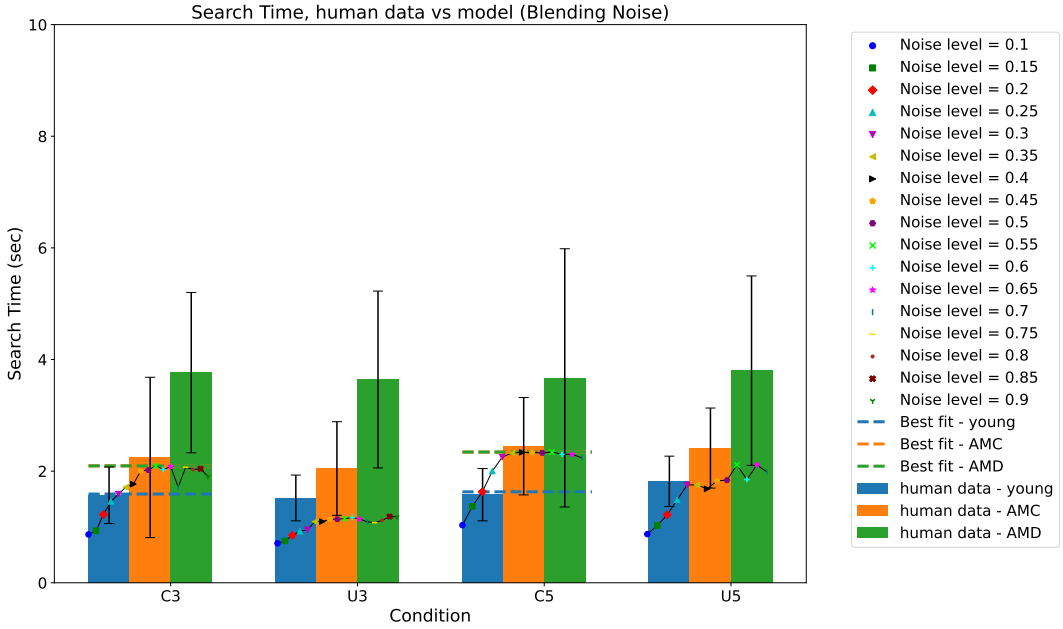

(b) Search time, using Blending noise. Similar to Figure 5.19a. Nonetheless, it appears that the search time was shorter in all four cases when using the Blending noise than the Gaussian noise.

Figure 5.19: A comparison between the average search time and the effect of increasing the number of objects, and comparing the overall results between the different noise types.

Number of Fixations

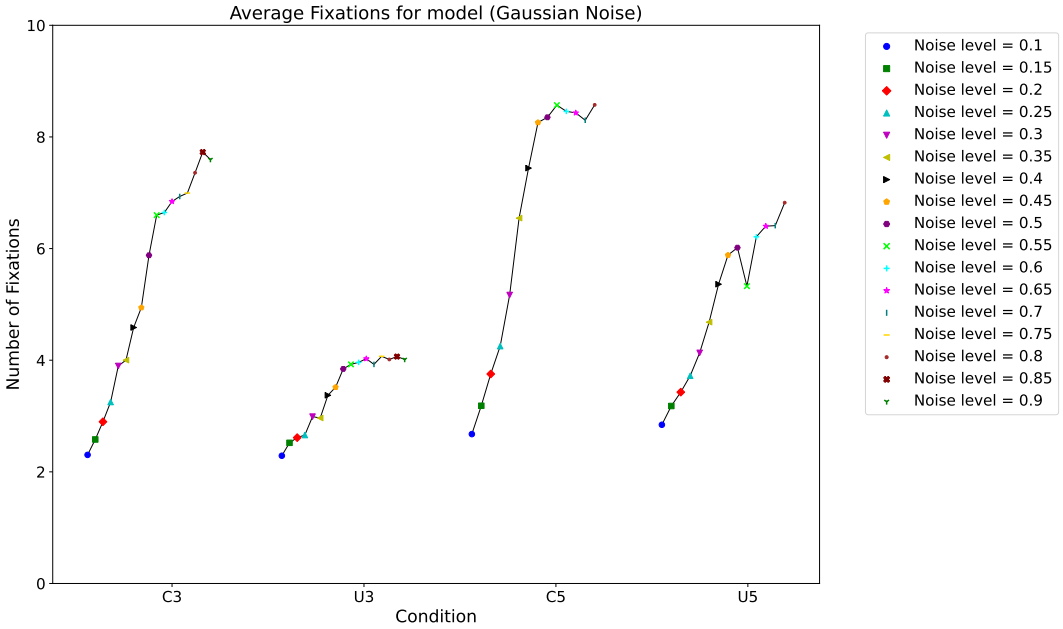

(a) Fixation, using Gaussian noise. The figure shows that with crowded conditions the model needed more fixations than the uncrowded conditions. The effect of increasing the number of objects is also evident.

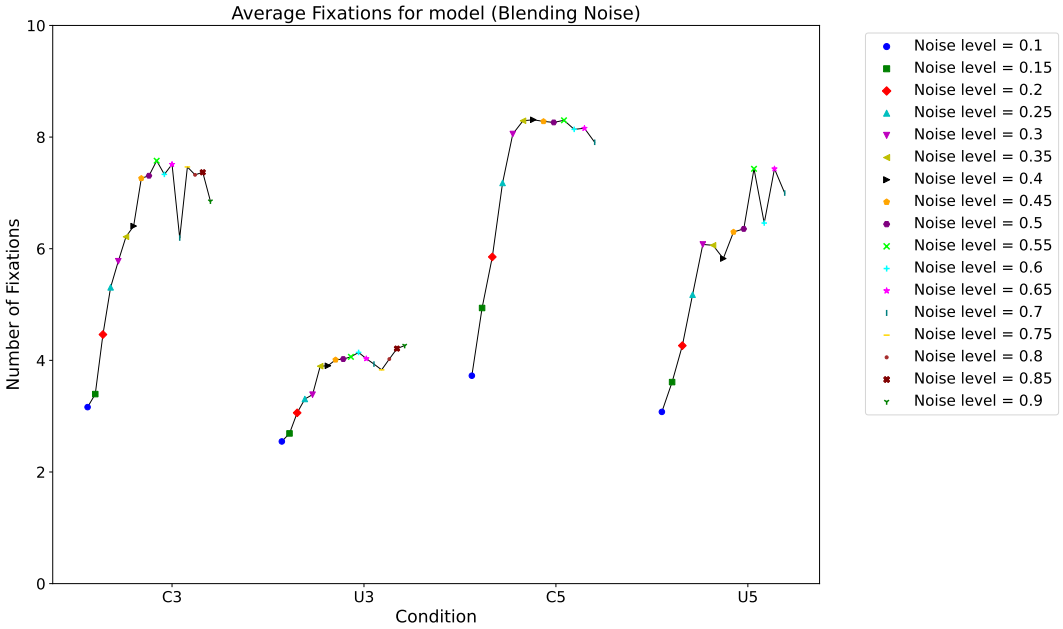

(b) Fixation, using Blending noise. The same as in Figure 5.20a. However, in this figure, there is a clear effect of the noise type when the noise level increases, as some number of fixations drop at higher noise levels making it a less realistic model under these scenarios.

Figure 5.20: A comparison between the number of fixations between the crowded and uncrowded conditions. The figure also compares the effect of increasing the number of objects, and compares the overall results between the different noise types.

Scan Path ration

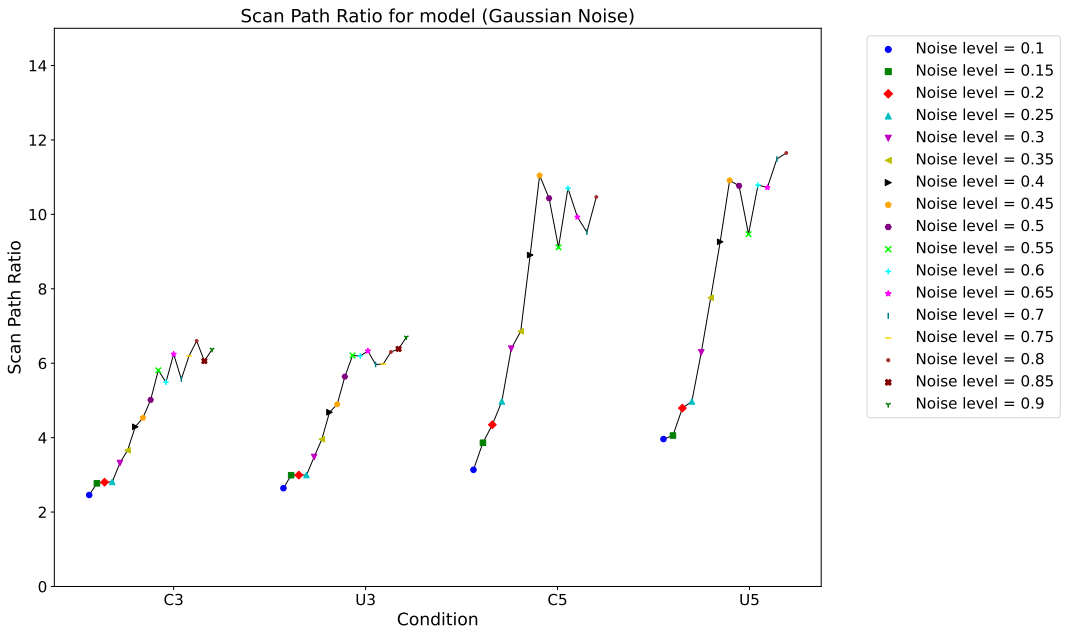

(a) Scan path ratio, using Gaussian noise. The figure shows that the crowded effect did not significantly affect the scan path ratio, yet the effect of the increase in the number of objects is more significant.

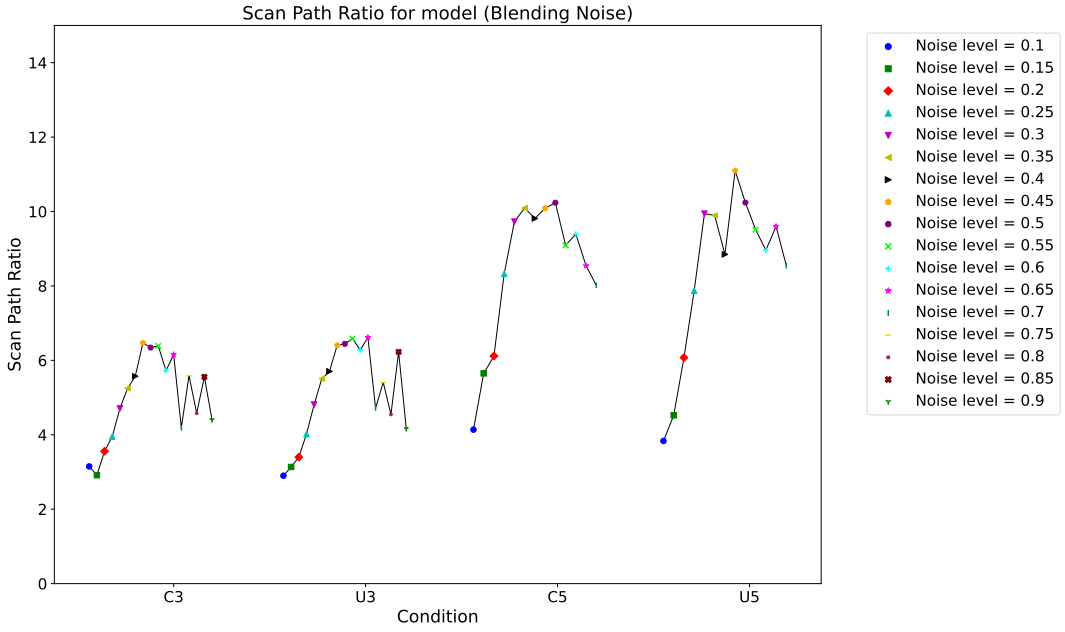

(b) Scan path ratio, using Blending noise. The figure shows that with the Blending noise, the scan path ratio was less consistent, especially at higher noise levels.

Figure 5.21: A comparison between the scan path ratios between the crowded and uncrowded conditions, increasing the number of objects, and comparing the overall effects of the different noise types.

The results presented in this section show clear bounded-optimal adaptation performance in response to the constraints introduced to the stochastic environment and the parameters that were introduced in the model to give some representation of the human vision limitations. It also showed that the model can give an estimation of the visual behaviours of those with AMD when interacting with a task involving crowded conditions, specifically with accuracy and number of fixations (Figures 5.18 and 5.20). Although the accuracy remained high for all the uncrowded conditions, it can be observed that the model is still learning the optimal strategies through the number of fixations, search time and scan path ratio (see Figures 5.20, 5.19 and 5.21). The high accuracy only indicates that the model learned the best policy which always produces high accuracy despite the noise levels in the observations. Despite the model failing to learn in some noise levels, especially for the six objects experiment, the noise levels with which it did manage to optimise behaviours still showed good strategies and were similar to the human data. Although there were no data provided by Thibaut, Boucart, and Tran (2020) for the number of fixations and scan path ratio, it stated in the paper that participants with AMD did more fixations and that the number of fixations increased with adding more distractors in the display. The paper also stated that increasing the number of distractors caused a less efficient scan ratio which is an observation matching our findings.

In the next section, we will discuss these results in more depth in terms of how they can help understand the interactive strategies for those with AMD.

5.9 Discussion

In this chapter, we implemented the computational model to investigate the crowding effect in individuals with age-related macular degeneration (AMD). The main goal was to develop a model that learns the visual search strategies and behaviours in individuals with AMD and to understand how the crowding effect influences their performance. We formulated the task as a Partially Observable Markov Decision Process (POMDP) and used Deep Reinforcement Learning (DRL) to learn the optimal policy. Our theory suggests that human behaviour in people

with AMD emerged as a consequence of the limitations in the environment (the crowding effect) and the bounds in the visual system, in addition to the constraints introduced due to AMD (loss of centre vision). By formulating the task as POMDP, we developed a statistical belief state for the agent that can be optimised through actions, observations and rewards. We introduced two types of noise in the observations, one that depends on the eccentricity and the other on the spatial spacing between objects (i.e. the crowding effect). We experimented with Gaussian noise and a Blending noise where features get blended in the observations if the objects were very close to each other, which imitates the crowding effect when observed through human eyes. Following this approach, the model learned optimal strategies and produced visual behaviours that can be compared and evaluated against human data. The model can provide an understanding of how people with AMD interact with a visual task involving the crowding effect.

Our results in Section 5.8 demonstrate several significant findings. Firstly, we observed a consistent reduction in accuracy in crowded conditions as the noise levels increased. Although the reduction in accuracy was noticeable while implementing the Gaussian noise, the blending noise showed accuracies that better fit the human data. The accuracy also decreased as the number of objects in the visual field increased, confirming the presence of the crowding effect in the experiment. The effect of crowding showed a significant impact on search time and the number of fixations; increased search duration and more fixations are prominent in our results. These behaviours show a similarity between the model's strategies and the human visual behaviours that were discussed in the results of the human experiment (Thibaut, Boucart, and Tran, 2020). Lastly, the scan path ratio was another measure to test the model's strategies, and similar to the human results, showed a decline in performance especially when increasing the number of objects.

These findings align with prior research following the same methodology conducted in individuals with normal vision (Chen, Acharya, and Oulasvirta, 2021; Chen et al., 2015, 2017), and people with different types of vision impairments (Li et al., 2023). We noticed subtle differences in how the crowding effect influences the visual search behaviours in people with AMD, which emphasises the need for specialised models that are built based on a deep

understanding of the reasons that influence the decision-making strategies and which are tailored to this specific type of vision impairment. Furthermore, we believe that our study has two implications. On one hand, our findings provide a deeper understanding of the challenges faced by individuals with AMD in visual search tasks especially when crowding is present. On the other hand, our model has the potential to be used as a tool for researchers and clinicians who are interested in assessing visual search abilities in people with this type of vision impairment. This includes reducing the need for human experiments, which can be costly, as involving few participants, is sometimes inefficient.

One important limitation of this study is the relatively small sample size of the human data as it was recruited from a single study (i.e. the study done by Thibaut, Boucart, and Tran (2020)), which can limit our findings to this specific set-up rather than giving more general findings. Additionally, an important limitation lies in the fact that the model is high in complexity, which limits our results to only four and six objects. Producing those results needed many attempts to improve the development of our model to reduce complexity and enhance learning as illustrated in Figure 3.2 (e.g. finding the set of noise levels that would suit the observation values and lead to efficient learning, re-running the experiments for longer period of time, changing some functionalities in the models dynamic). Achieving results for nine objects is surely something that can be done in the near future. A major limitation in our model is the lack to produce a search time duration that fits the AMD humans data. One reason could be because we chose constant value for the fixation duration from experiment on people with normal vision, so for the future, we could experiment with specifying the fixation duration to those with AMD. Another reason could be the tradeoff in the model between the cost of fixation, the observed value and the accuracy beside the very high noise. All these could be possible reasons for this behaviour. It is important to note that the aim of building this model is to present the ability of computational models to produce bounded-optimal strategies for those with vision impairment by adding more uncertainty in the visual field and observations. An alternative approach could be done by specifying additional limitations for those with AMD. For example, we can further investigate how different scotoma levels could impact the model's observations. Another

important limitation is the model's inability to learn to explore the environment in higher noise levels. More experiments can be done by experimenting with different $base\ crowding\ factors$ and the find noise levels that fir the human data and at the same time show better learning strategies in the model.

Our findings contribute to the theoretical framework of the crowding effect by demonstrating its impact on individuals with AMD. Furthermore, this study contributes to the application of POMDP to formulate a visual search ask for people with AMD and uses DRL techniques to simulate and understand the visual search strategies and behaviours of individuals with AMD. This model represents a significant contribution to the field by providing a framework to analyse and predict the behaviour of individuals with AMD under different crowding conditions. The POMDP framework sets a rigorous foundation for investigating how humans, in a computationally rational way, adapt to constraints in information processing (Howes, Lewis, and Vera, 2009; Lewis, Howes, and Singh, 2014), which connects rationality and cognitive mechanisms.

The experiment could be extended in future work by considering other features than the targetedness in the belief state such as the colour and shape, and implementing a similar grouping to the human experiment in the state (grouping objects with the same features in one display as in Thibaut, Boucart, and Tran (2020). Additionally, we could add additional experiment and implement a more accurate representation of the foveal visual field loss people with AMD experience, for example, by distributing the noise so there is more noise in the centre of the fixation and less noise in the area outside the fixation point while keeping the properties of the peripheral vision. This experiment could add more complexity to the model, but the strategy the agent finds might be closer to the people with AMD strategy. Another extension to the experiment is to consider other cognitive factors that can be affected by the crowding effect such as working memory attentional allocation (Bacigalupo and Luck, 2015). Lastly, we believe that the experiment can be extended to find the best-fit for each person's level of vision loss rather than only comparing the average. As the participants in the human experiments had slightly different vision loss than one another. Figure 5.3 provided in Section 5.4 shows that removing the outliers from the participants, specially the AMD ones, could be contributing to the biggest

difference between the model of human model which would show more similarities between the model's search time and the AMD participants (excluding the outliers).

In conclusion, our study identified the complexities of visual search behaviours in individuals with AMD and highlighted the significance of the crowding effect in the interactive behaviours. These findings can support HCI as they provide some insights on how people with AMD adapt to their limitations when interacting with a screen with crowded objects.

Chapter 6

Experiment 3: A Model of Visual Search with Cataract

6.1 Introduction

This chapter discusses the implementation of a visual search task that is specified for cataract patients. The task aim was to test eye movements and visual performance in people with cataracts when they interact with three visual search tasks (object search, face recognition, and reading). These tasks were detailed by Wan et al. (2020). Only the object search and the face recognition tasks were implemented in our model due to the higher complexity of the reading task. The human experiment aimed to examine the visual search performance of patients who plan to undergo cataract extraction surgery and compare their visual behaviours before and after surgery. Researchers investigated different eye movement measures, such as the accuracy of correctly identifying a target, mean fixation duration, and fixation count. In our model, we investigated the accuracy of correctly finding the target, the scan path ratio (not provided in the human experiment), and we investigated eye movements through visualisations to investigate the model's behaviours while trying to find the target. To the best of our knowledge, research that has been conducted to investigate eye movements in patients with cataracts is minimal. In our experiment, we formulate the two visual search tasks (i.e. object search and

face recognition) as POMDP. Similar to the experiments in Chapters 4 and 5, the experiments were implemented based on the assumption that visual behaviours emerge from the theory of bounded optimal. Visual uncertainty was implemented as a Gaussian distribution with a linear relationship with eccentricity. The crowding noise was also implemented in this experiment. We set up the Crowding settings to be the same as in the human experiment, meaning that the distance between the spacial objects stayed the same in all the trials. Our model predicted the accuracies of cataract patients pre-surgery and post-surgery. The visualisations showed that the model spent more time looking for the object in the post-operation condition. It also showed some differences in the scan path ratio for each category.

6.2 Background

As previously defined in Chapter 1, cataracts is a condition that causes cloudiness in the visual field, causing the patients to experience blurred vision, light sensitivity and difficulties seeing at night (Benedek, 1971; National Eye Institute, 2023; Shiels and Hejtmancik, 2007; Shiels and Hejtmancik, 2013). The symptoms of cataracts can be improved by cataract extraction surgery (National Eye Institute, 2023; Skiadaresi et al., 2012). Cataracts have three forms: nuclear, cortical, and posterior subcapsular. Technology has been used to understand the effects of each condition; for example, Krösl et al. (2020) created an augmented reality (AR) simulator that imitates how a person with cataracts would observe the surroundings. The patients reported that the simulation in the AR created worse visual fields than what they usually experience. Their findings emphasised the importance of lighting conditions in 3D architecture scene models on individuals with cataracts and how the perception differs based on their own condition and abilities. Another study investigated visually guided eye movements of individuals who recovered their sight, including individuals with different types of cataracts, such as those starting to develop sight after they underwent cataract-removal surgery. The study showed that the individuals performed well despite the effects of their previous blinding (Zerr et al., 2020). Most of the research that is done on cataracts focuses on how the surgery can improve the quality

of life for patients with cataracts, which showed that surgery improves the quality of life for those with cataracts in many aspects, such as socialising, mobility, reading and driving (Cillino et al., 2014; Elliott et al., 2000; Lamoureux et al., 2011).

6.3 Visual Search Behaviours in Cataracts

The experiment in (Wan et al., 2020) showed that the visual search behaviours in cataracts improved after the patients underwent surgery. In the first experiment, participants were asked to search for an object on the screen surrounded by 23 distractors. The results showed significant improvement in finding the correct objects, with a noticeable decrease in the search time. It also showed that post-surgery, people made more fixations on the target, and visited the target more frequently than before surgery, which is believed to be the result of improved vision after surgery as participants can pay more attention to the target's location. In the second experiment conducted in the same research, they tried to compare the visual search behaviours in a face recognition task, where participants were asked to memorise a face and then search for it when it was located within two other faces. The results showed similarities to those reported in the previous experiment. Accurately identifying the face increased, and the search time decreased after the surgery. The results also showed that participants fixated more on the area of interest (AOF) after surgery, increasing the total fixation duration and visits to the target. Lastly, the scientists conducted a reading experiment in the same study. The results showed that reading speed significantly increased post-operation compared to pre-operation, where the proportion of regressive saccades (i.e. when the eyes move to parts of the text that was read previously (Vitu and McConkie, 2000)) noticeably decreased. A previous study was done by Thepass et al. (2015), which examined saccade latency in patients with different levels of cataract severity. The study reported an increase in saccadic reaction time in the more severe type of cataract.

In the next section, we will explain the visual search and face recognition experiments done by Wan et al. (2020) in more detail.

6.4 Experimental Tasks

The setup for all the human experiments was the same. Participants were instructed to sit 65 cm from a 21.5 inch computer monitor with a resolution of 1920 × 1200. An eye tracker was placed on the monitor to record eye movement behaviours throughout the experiments. For a trial to start, participants were asked to fixate at the centre on a fixation target that appears on the screen for 3 seconds at the start of every trial. This was done to ensure a fixed fixation location.

6.4.1 Task 1: Object Search

At the beginning of the experiment, participants were provided with objects presented on the computer screen so they could familiarise themselves with the objects included in the task. Afterwards, participants were asked to find a specific target at every trial (see Figure 6.1 for an example of the display). Participants were restricted to finding the target within 6 seconds. After the 6-second period, the screen will disappear. The participants were asked to keep a still fixation on the target once they found it within the 6 seconds period. There were 12 trials, and participants were asked to try and find a target during each trial.

Figure 6.1: The object search displayed for participants in the visual search as presented in (Wan et al., 2020). In this example, the object search behaviours of one participant are shown when they were asked to find a target (a clock in this example). The yellow circles represent the fixation positions, and the yellow line represents the scan path made by the participant. The circle size corresponds to the fixation duration. Whereas the numbers inside the circles show the rank of each fixation. The green colour surrounding the target corresponds to the area of interest (AOI). (image used with permission from the author)

6.4.2 Task 2: Face Recognition

In the face recognition task, there were two stages of each trial. In the first stage (i.e. viewing stage), participants were presented with a target face for 5 seconds, which they were asked to memorise. In the second stage (i.e. recognition stage), participants were presented with three faces (see Figure 6.2 for an example of the display), two of the faces the participants had not seen before, and the third face was the target face they memorised in the previous stage. The participants were then asked to find the target face within 5 seconds. Once they it, they should keep a still fixation on it until the image disappears.

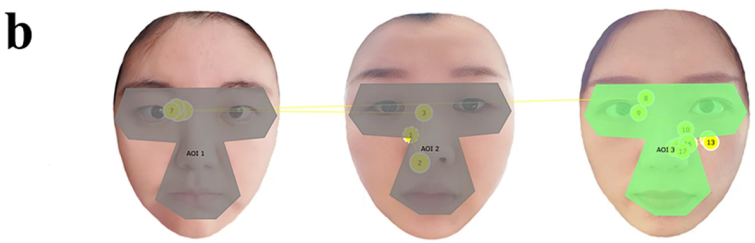

Figure 6.2: An example of the display in the face recognition task as presented in (Wan et al., 2020) including photos of the researches for illustration purpose. The yellow colour represents fixations and saccades. Whereas the AOI of the target (face features) is presented with green highlight. The AOI coloured with gray highlight represents non-target AOI. (image used with permission from the author)

6.5 Theory

The theoretical assumptions for both cataract tasks are the same as the crowding task in Chapter 5. The only difference is that the search time is restricted to 6 seconds in the object search task and 5 seconds in the face recognition task without a final action. Therefore, the search time in this task is considered to be the duration from the first fixation action until the cut-off time

is reached (to match the constraints in the human experiment). The calculations of the search duration remain the same as in Section 5.7. There are some differences in formulating the tasks as POMDP in this chapter due to the differences in each stimulus.

The formulation for each task is in the following sections.

6.6 Task 1: Formulation as POMDP

The section explains the formulation of the object search task (i.e. Task 1 6.4.1) as POMDP

- **State S:** As discussed in Chapter 5 Section 5.7. In this experiment, the objects stay in the same locations at each trial. Whereas the target's location is chosen randomly at every trial. The locations of the objects are distributed within the screen size ($40° horizontally \times 23.5° vertically$) , and the object's width (= 4). . The distractor value in this experiment is the average between a $minmum\ distractor\ value = 0$, and $maximum\ distractor\ value = -1$, which is -0.5, while the target value is 1. The reason for creating this distractor value rather than -1 is to create similarities between the target and the distractor features. The state remains unchanged in each trial until it finishes.

- **Action A:** The action space is the same as in the experiment in Chapter 5 except that in this experiment, there is no action to confirm the decision. Therefore, the agent would remain fixated on the target (once it finds it) until the time cap (6 seconds) is over.

- **Reward function $R(S, A)$:** As explained in Chapter 5, Section 5.7

- **Transition function $T(S_{t+1}|S_t, A_t)$:** As explained in Chapter 5, Section 5.7

- **Observation O and observation function $O = f(S, A)$:** The observation is as discussed in Chapter 5, Section 5.7. Although we do take the crowding noise into consideration, we reduced the $base\ crowding\ factor$ to 5. The reason for this choice is that there are more objects in this experiment, and having a high crowding noise can increase the complexity of the observations. We add the crowding noise in the observations to produce the crowding effect, but we do not test with different spacing.

- **Discount rate γ:** As explained in Chapter 5, Section 5.7

- **Belief update:** The belief in this task is the same as the belief in the crowding task in Chapter 5. However, based on the State S, there are a few differences in the belief state for this task. The unknown value is considered as the average between the distractor value and the target value of 1, (0.25) and the initial confidence 1000000 for each object. After taking action a_t and receiving observation o_t, the belief state would be updated as usual (using the Kalman filter formula in Chapter 3, Section 3.4.1)

6.7 Task 1: Implementation

The implementation of this task is very similar to the crowding task (Chapter 5, Subsection 5.7.1). However, in this task, we considered a different set of noise levels. Because we only measured the accuracy of the model (the average of correctly identifying the object) and the scan path. The reason for not including all the measures examined by Wan et al. (2020) will be discussed in more depth in the results, in Section 6.8. It is also important to note that in this experiment, the spacing between objects remains the same during all the trials, to match the human experiment. The spacing between objects is $S = 1.5$. Due to the high accuracy of the humans, we expect that the best fit would be within the low noise levels. Therefore, for the grid search, we chose the noise levels $[0.1, 0.2, \ldots, 1]$. The base crowding factor is 5 in this experiment, while the base distance factor remains the same as reported in Subsection 5.7.1 (0.9). Table 6.1 illustrates the actual noise values for this experiment.

The number of objects for this experiment is 12, which is half the number of objects in the human experiments. Although our model was able to learn with more objects (up to 18 objects), the experiment would potentially be time-consuming due to the nature of the workflow (Figure 3.2). The model was trained for 8×10^6 iterations with a 5000 moving average to produce smooth curves. Then, we ran testing for 2×10^3 to evaluate the behaviour of the model against the human data. The results will be discussed in the next section.

Noise Level	*Crowding Factor*	*Distance Factor*
0.1	0.5	0.09
0.2	1	0.18
0.3	1.5	0.27
0.4	2	0.36
0.5	2.5	0.45
0.6	3	0.54
0.7	3.5	0.63
0.8	4	0.72
0.9	4.5	0.81
1	5	0.9

Table 6.1: Actual values for *Crowding Factor* and *Distance Factor* implemented in the observation

6.8 Task 1: Results

In this section, we present the results of the model, including the training and the testing results, which we compare against the human data. The results of the Gaussian and Blending noise did not seem to have big differences in this experiment, and we will discuss both of them in this section. After that, we evaluate both results against the human data. Lastly, we provide a visualisation of the model's behaviour to clarify some points in the results.

6.8.1 Model Performance (Gaussian)

The learning performance shows consistent learning behaviour (high rewards) at the lower noise levels. Whereas at higher noise levels, the reward values decreased, leading to less accuracy, and less efficient scan path ratio (Figure 6.3). The visual representation for the histograms of the Gaussian and Blending noise is provided in Appendix B, Sections B.1.1 and B.1.2.

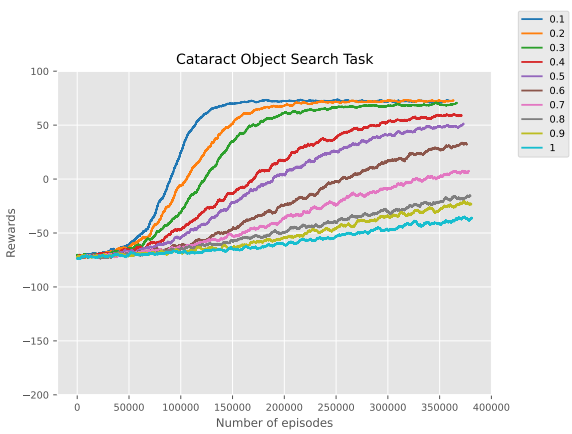

(a) Reward during learning

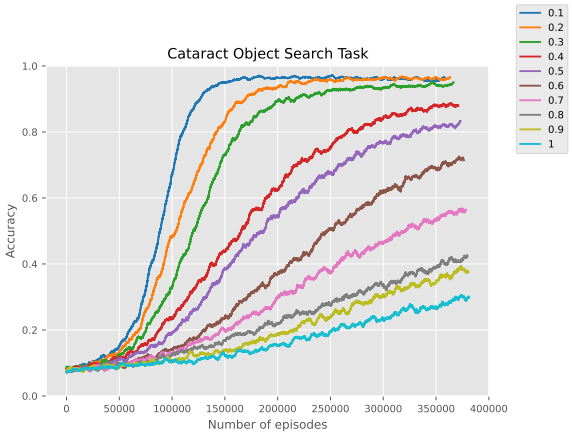

(b) Accuracy during learning

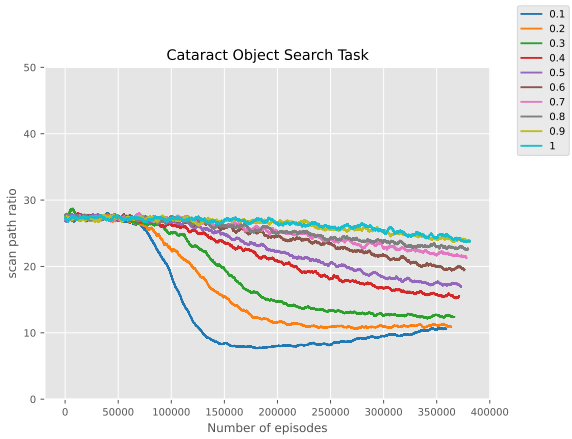

(c) Scan path ratio during learning

Figure 6.3: The learning performance of the model with the object search task (12 objects - Gaussian noise)

The plots in Figure 6.3 can be self-explanatory. It shows the drop in learning as the noise levels increase in the observations. It also shows how the scan path ratio was efficient at lower noise levels, indicating that the model took a shorter path to find the target. The learning performance with Blending noise was very similar to the Gaussian noise (see Appendix B, Section B.1.3). Due to the fixed search duration in the model, the number of fixations at each episode remains the same (the count of the number of actions is almost the same). However, when we visualise the fixations behaviours, the model showed less fixations on distractors until it finds the target during searching with lower noise levels; contrary to that, the model fixates more on distractors and away from the AOI when noise levels are higher. Once it finds the target, it keeps fixating on it until the episode time is over. The visualisations of the search behaviours

are provided in the next section to compare the fixations actions in the model when presented with different noise levels. Another important observation is that even though the scan path ratio behaviours align with the noise levels, this does not always reflect the efficiency of the scan path, because the fixations count is the same in all the trials and there is no action that can end the episode. This means that the agent is possibly fixating near the target (with very short scan path) without correctly identifying it (which would show an efficient scan path ratio but low accuracy). This behaviour is not very observable in these results, but it is more apparent in the results of the second task (i.e. face recognition, in Section 6.11.1).

6.8.2 Evaluation of the Model

In this section, we first compare the mode's accuracy against the human's accuracy and compare the accuracies between the Gaussian and the Blending noise (Figure 6.4), we also present the scan path ratio in the Gaussian and the Blending noise (Figure 6.5). Then, we present visualisations of the fixation actions during the visual search when given different noise levels and report our observations on the results (Figure 6.6).

Accuracy

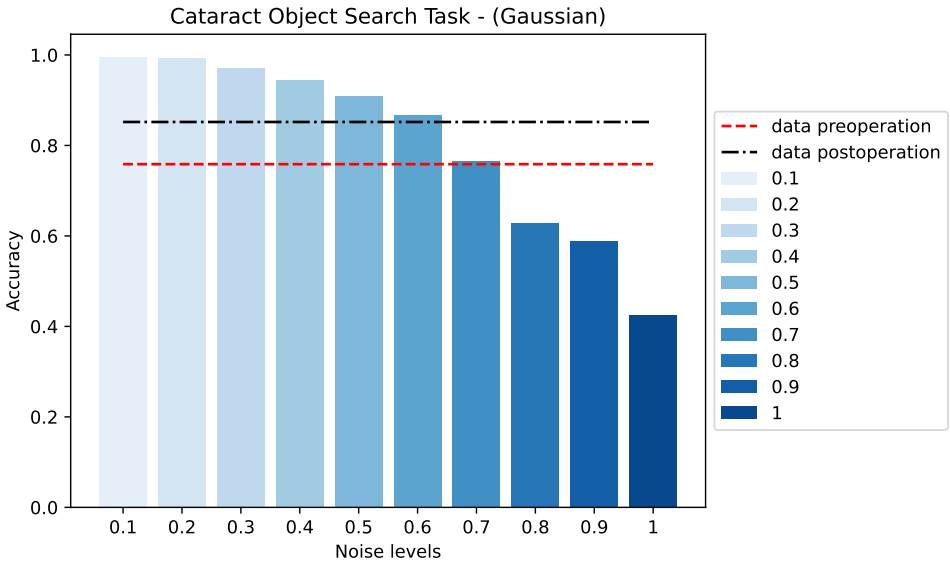

(a) Accuracy of the model against the human's accuracy (Gaussian)

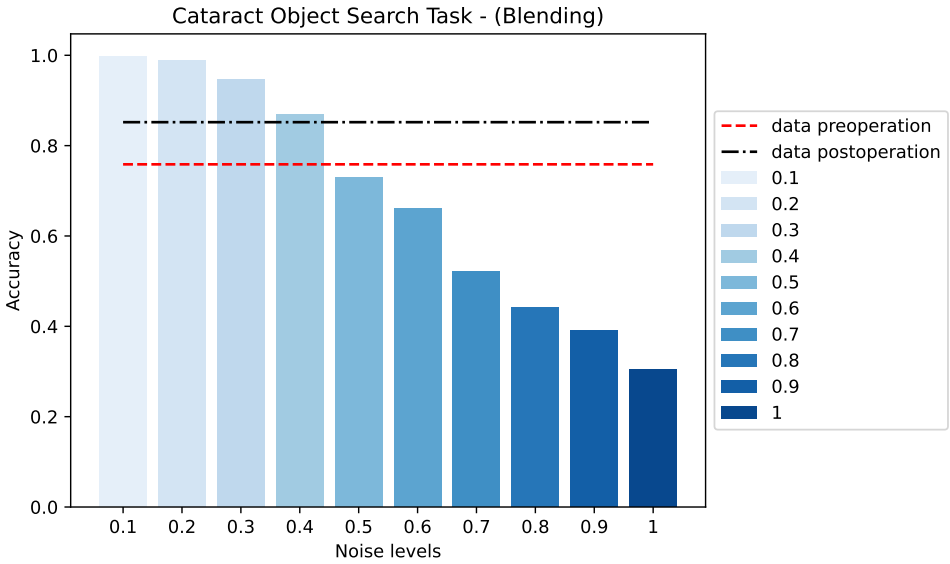

(b) Accuracy of the model against the human's accuracy (Blending)

Figure 6.4: Figures 6.4a represent the accuracy during testing (Gaussian noise). The noise level of 0.7 fits the human data pre-operation, and the noise level of 0.6 fits the post-operation data from (Wan et al., 2020). Whereas the figures 6.4b shows the accuracy during testing with the Blending noise, and it can be observed that the best fit is at smaller noise levels (0.5 for pre-operations and 0.4 for post-operation).

Scan Path Ratio

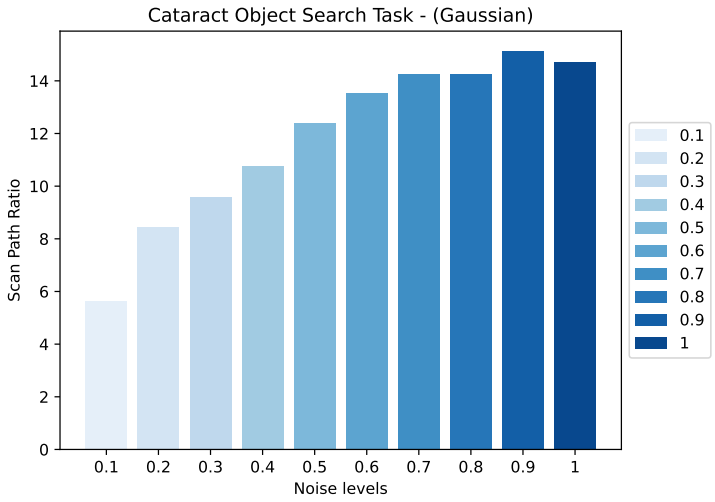

(a) The effect of noise levels on the scan path

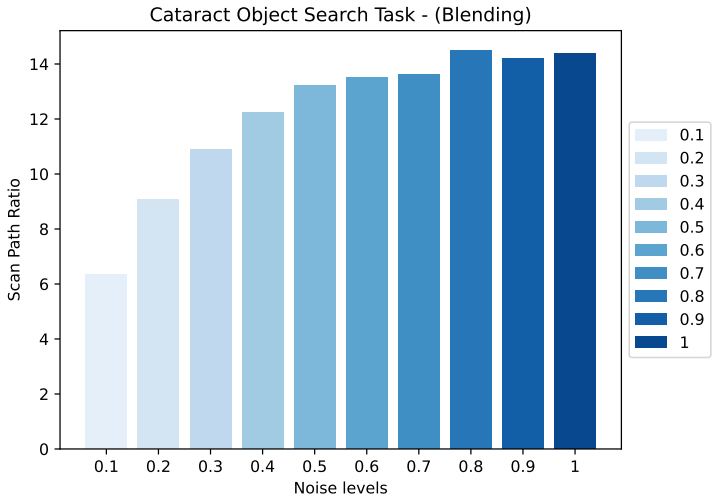

(b) The effect of noise levels on the scan path

Figure 6.5: Figure 6.5a represent the scan path ratio during testing (Gaussian noise) and Figure 6.5b show the scan path ratio during testing with the Blending noise.

Figures 6.4 and 6.5 show that the model's results met our hypothesis as higher noise levels matched the accuracy of people with cataracts before the operation. Although the human data was from a visual search task with 24 objects, we believe that these results can indicate the outcome of the model if it was set to be trained with 24 objects. The only difference would be finding a suitable set of noise levels to experiment with, possibly smaller values. The results also

showed the decrease in scan path ratio efficiently as the noise levels increased. The difference between using the Gaussian noise and the Blending noise can also be observed, as the Blending noise showed less accuracy at lower noise levels (e.g. if we compare the accuracy at noise level 0.4 with Blending noise, it is lower than the accuracy at noise level 0.4 with Gaussian noise). Hence, the best-fit values were at 0.5 for pre-operation rather than 0.7 (see Figures 6.4a and 6.4b). This clearly indicates that the crowding effect is more apparent in the Blending noise. Additionally, as the search duration was fixed in the human experiment, Wan et al. (2020) measured the search duration by measuring the time of search before participant start to make continuous fixations within the AOI. In our model, the agent can only fixate on the feature of the object, which is located at the centre of the object. It was not possible to measure the search duration as the authors Wan et al. (2020) did. However, the visualisations we provide can give an indication of the search time through the number of fixations done until the model starts to make continuous fixations on the target.

Furthermore, the measures of total fixations in the human experiment are a measure of only counting the number of fixations within the AOI. This can be interpreted as the number of fixations actions on the centre of the object in our model, which is also shown in the visualisations. The fixation actions in the model can give an insight into the effect of the noise in the peripheral vision on fixations behaviours. Figure 6.6 compares fixations in the model, which can reflect on the search time between the first fixation until finding the target and then continuing to fixate on it (within the AOI). We chose noise level 0.1 to compare the effect of the Blending noise at even the lowest level. We also compare the visualisation of the best-fit values (0.6, 0.7 for Gaussian and 0.4, 0.5 for Blending). Lastly, another noticeable observation is that at higher noise levels, especially in the Blending noise, there are cases where the model chooses a distractor as a target despite the fact that sometimes it has fixated on the target during the search. Figure 6.7 shows an example of this behaviour.

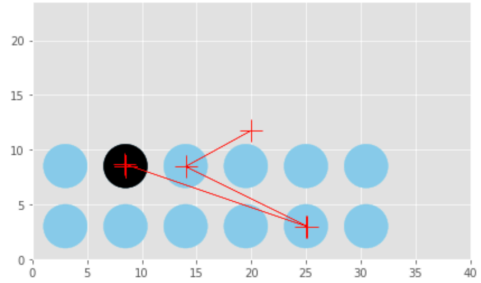

(a) Visualisation of fixation actions at noise level 0.1 (Gaussian). There are 3 fixations before finding the target (including two on the same object).

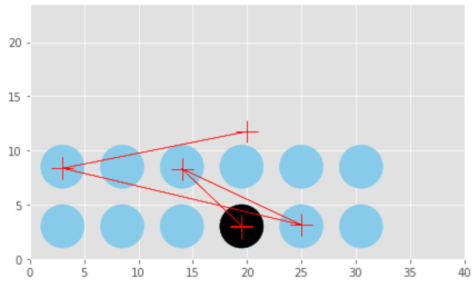

(b) Visualisation of fixation actions at noise level 0.1 (Blending). There are 3 fixations at different locations before finding the target.

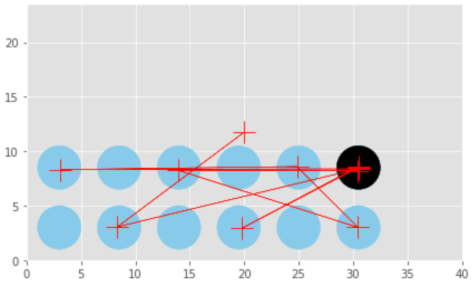

(c) Visualisation of fixation actions at noise level 0.6 (Gaussian), (6 fixations before finding the target)

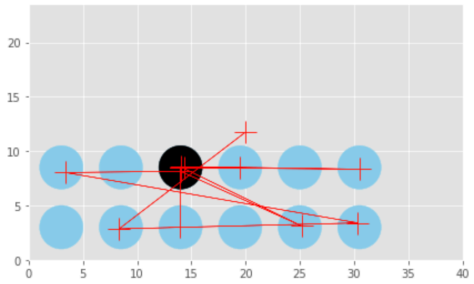

(d) Visualisation of fixation actions at noise level 0.4 (Blending), (10 fixations before finding the target)

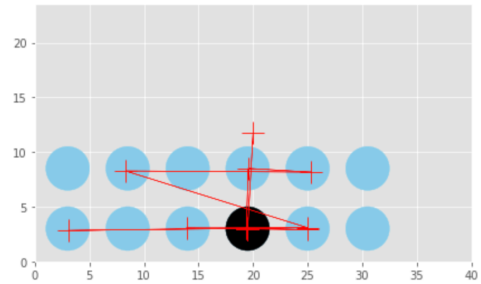

(e) Visualisation of fixation actions at noise level 0.7 (Gaussian), (8 fixations before finding the target, including fixations on the same object).

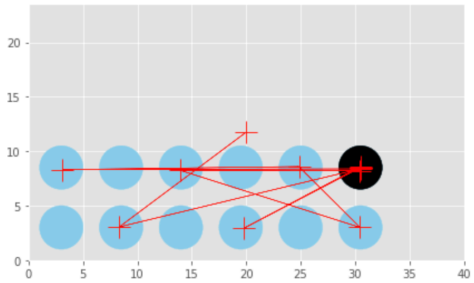

(f) Visualisation of fixation actions at noise level 0.5 (Blending), (8 fixations before finding the target, including fixations on the same object).

Figure 6.6: This figure provides an illustration of how the fixation actions behave in the model, where the blue circles represent distractors, the black circle represents the target, and the red crosses represent the fixation locations, which is at the centre of the start of each trial. Despite the fixed search time, the model tries to find the target within the limitations introduced in the observations. As it can be seen in Figure 6.6a, the model made three fixations (indicating less search time) and then kept fixating on the target. Whereas in Figure 6.6b, the model made more fixations, which indicate the effect of the Blending noise. It can observed that the fixations until finding the target increase with the increase of noise, and with the Blending noise, sometimes having a higher number of fixations.

Figure 6.7 shows examples of cases when the model does not find the target. By testing the simulations, this behaviour is mostly noticeable with the Blending noise where the model confuses targets with flankers, as discussed in Chapter 5.

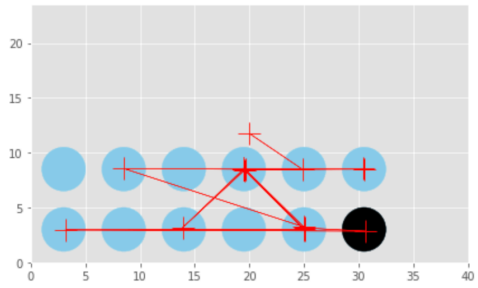

(a) Visualisation of fixation actions at noise level 0.5 (Blending) where the model fixated on the target but chose a different object as a target (the chosen object is the top fourth object (from the left) with several fixations at the same location)

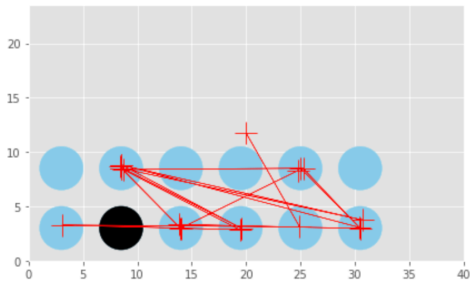

(b) Visualisation of fixation actions at noise level 0.7 (Blending) where the model chose a different object as the target (second object at the top left) without fixating on the target

Figure 6.7: This figure illustrates the model behaviours when it fixates on distractors (sometimes on the target) while incorrectly identifying the target

The visualisations do not represent the humans visual behaviours but they reflect the behaviours in our model and how it is impacted by the noise. The results will be discussed in more depth in Section 6.12.

In the next section, we will discuss the face recognition task.

6.9 Task 2: Formulation as POMDP

The task formulation for the face recognition task (i.e. Task 2 6.4.2) is almost the same as the object search. The only difference lies in the State and Action space, which will be described as follows:

- **State S:** A state s_t is a vector of $(x, y, targetedness) \times 4 \times 3$, which represents three faces with four features (two eyes, mouth and nose). The target value is 1 and the distractor value is the average between a $minmum\ distractor\ value = 0.5$ and $maximum\ distractor\ value = 0.9$, which is 0.7. The reason for creating this distractor value is to create similarities be-

tween the target and distractor features, which can represent similarities in faces. The three faces are equally distributed within the screen size ($40°horizontally \times 23.5°vertically$), where the features were distributed within each face object.

- **Action A:** There are 12 actions a_t in this task, which correspond to the AOIs of each face. We arbitrarily map the numbers $[0, 1 \ldots 11]$, representing the twelve actions, to one of the twelve features. For example, if the action number is 5, we calculate the object number (the face) by using floor division over the number of features (i.e. $\lfloor 5/4 \rfloor = 1$). Then, we get the value of the feature by computing the feature's location within the object, which is the difference between the action number and the number of objects multiplied by the number of features in each object (4) (i.e. $5 - 1 * 4 = 1$). This means that when the action number is 5, it indicates that the model is fixating on the second face (right eye). The agent then will fixate on the perceived (x,y) coordinates of that AOI, which it gets from the repeated Bayesian update (in Section 3.4.1) it had based on each observation about this object (the face).

- **Reward function $R(S, A)$:** As explained in Chapter 5, Section 5.7.

- **Transition function $T(S_{t+1}|S_t, A_t)$:** As explained in Chapter 5, Section 5.7.

- **Observation O and observation function $O = f(S, A)$:** The observation o_t can be described the same as in Section 6.6. However, in this task, the distance noise $d_{t,obj} \sim \mathcal{N}(\mu = 0, \sigma_{o,d})$, where $\sigma_{o,d}(t) = distance\ factor \times eccentricity(t)$ considers the eccentricity between a fixation point and other features in the observed face. Whereas the crowding noise $c_{t,obj}$ considers the crowding between the objects (faces in this task).

- **Discount rate γ:** As explained in Chapter 5, Section 5.7.

- **Belief update:** The belief can be described in the same way it was presented in Section 5.7. However, there are a few differences in the belief state for this task. The unknown value is considered as the average between the distractor value (0.7) and the target value (1), which is 0.85. Therefore, the initial belief consists of the unknown value (0.85) and

the initial confidence 1000000 for each feature in each face. Meaning that the belief is an array consisting of 12 values (i.e. $[0.85, 1000000] \times 4 \times 3$. After taking action a_t and receiving observation o_t, the belief state would be updated as usual (using the Kalman filter formula in Chapter 3, Section 3.4.1).

6.10 Task 2: Implementation

The implementation for this task is very similar to the implementation of the object search task in Section 6.7. There are a few differences, which will be discussed in this section.

In this task, we considered the object width to be 8, and the spacing between objects is $S = 4$. Because the crowding effect is not apparent in this task (due to the big spacing), we only consider the Gaussian noise $c^{(G)}_{t,obj} \sim \mathcal{N}(\mu = 0, \sigma^{(G)}_{o,c})$ in the experiment. Grid search was implemented with the same noise levels as in Table 6.1. We trained the model with 8×10^6 iterations with a 5000 moving average to produce smooth curves and 2×10^3 testing iterations. The results will be discussed in the next section.

6.11 Task 2: Results

As mentioned previously, the results in this section were implemented using the Gaussian noise (for the crowding noise). First, we will present the model performance through reward, accuracy and scan path ratio. Second, we will illustrate the results of the testing iterations.

6.11.1 Model Performance

The learning performance shows that there is consistent learning behaviour (high rewards) at the lower noise levels. Whereas in this task, the model showed a significant drop in performance at higher noise levels. The reward values dropped significantly, leading to much less accuracy (not correctly identifying the target face) and a less efficient scan path ratio (Figure 6.8).

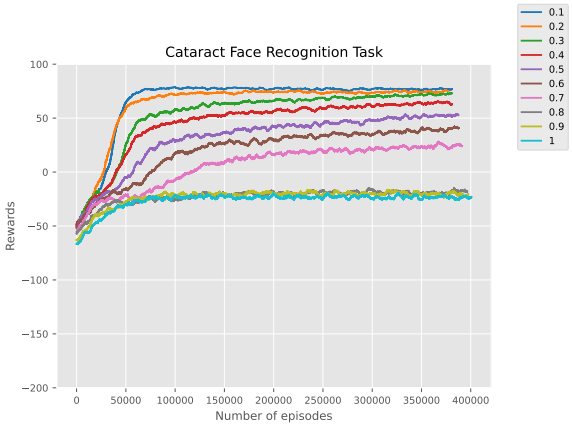

(a) Reward during learning

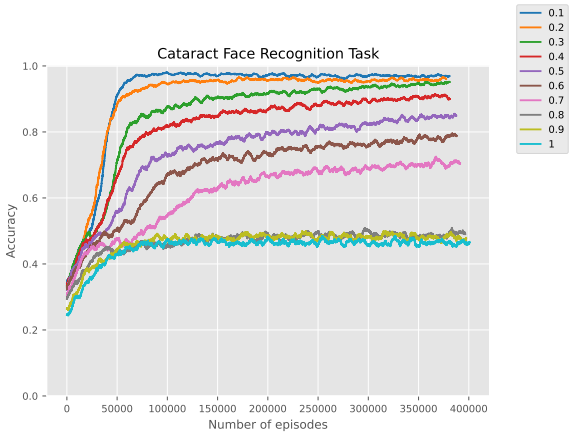

(b) Accuracy during learning

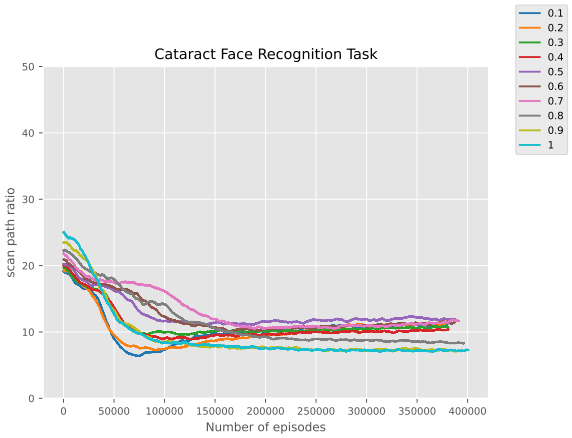

(c) Scan path ratio during learning

Figure 6.8: The learning performance of the model with the face recognition task - Gaussian noise.

The plots in Figure 6.8 show that the model learned the optimal policy at the noise levels $[0.1, 0.2, \ldots 0.7]$. Whereas the performance drastically dropped at noise levels 0.8, 09, and 1. It is also noticeable how the scan path ratio does not always reflect the scan efficiency, as we mentioned in Section 6.8. In the case of these noise levels, the model was fixating on a location very close to the target with very short paths, and it kept fixating until the duration ended. This can be clarified with visualisations that we will provide in the next section.

6.11.2 Testing Results

This subsection represents the behaviours of the model during testing in the face recognition task (Figure 6.9). The accuracy that fits the human data can be observed in Figure 6.9a, where the model predicted the accuracy for people with cataracts pre-operation at noise levels 0.5 and 0.3 for the post-operation. Again, the model's accuracy can meet the hypothesis of this thesis. It can also be seen that the results for the scan path (Figure 6.9b) are not exactly linear with the noise levels for the reason mentioned previously in Section 6.8. We believe that this behaviour is more apparent in this task than in the object search task due to the locations of the features and the noise associated with it. Another observation in the results is that at higher noise levels, the model would either fixate on an area close to the target without finding the target (Figure 6.10c), or it fixates on areas on the edges of the objects or outside the screen limits (Figure 6.10b). Although the same noise was included in the previous tasks, in this task, the locations of the features are distributed within the face, which is also close to the edge of the screen. Whereas in the previous tasks, the locations of the features were usually in the centre, and the distance between objects was smaller. Therefore, even with the added noise, this behaviour did not appear in the object search task and only appeared in this task. Visualisations of the fixations behaviours are provided in Figure 6.10.

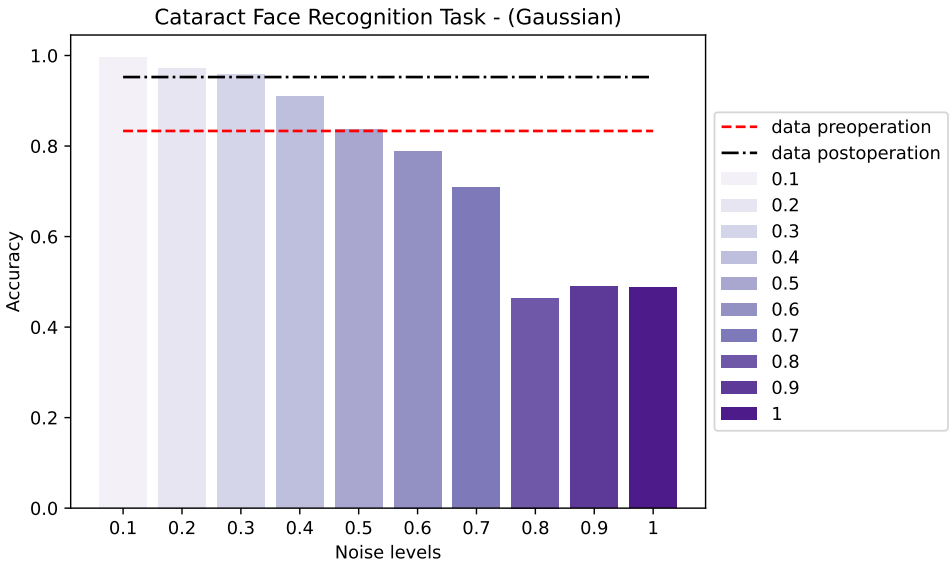

(a) The accuracy of finding the correct target in the model compared with the human data

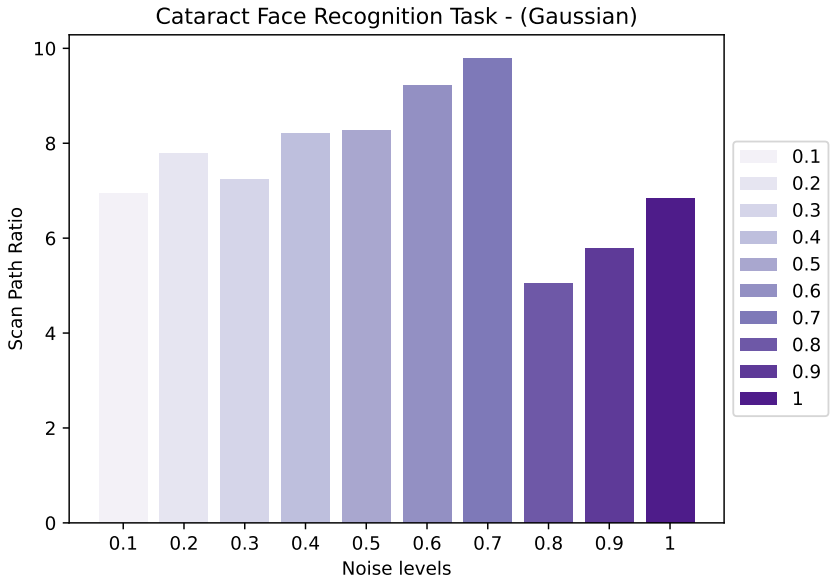

(b) The scan path during the face recognition

Figure 6.9: The figures illustrate the model's accuracy and scan path during the face recognition task.

Figure 6.10 represents examples of visualisations of the noise levels that matched the human's accuracy. We also provide an example visualisation from the results of noise level 0.9, where the scan path ratio is very short. Yet, the model fails to find the target.

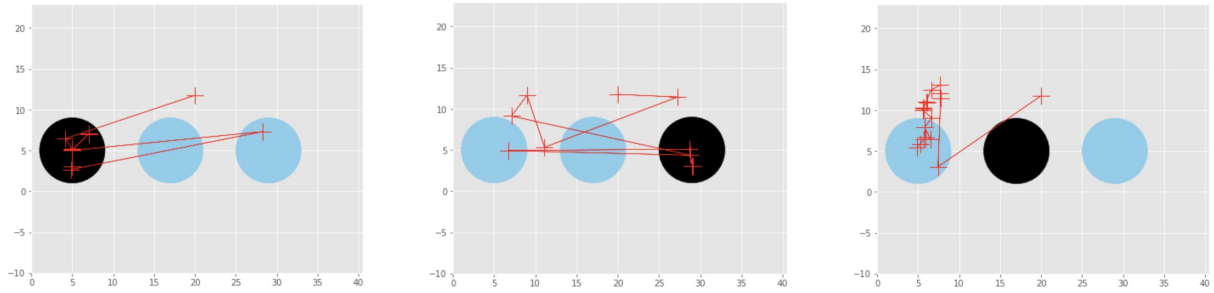

(a) Example visualisation of the visual search behaviours with noise level 0.3

(b) Example visualisation of the visual search behaviours with noise level 0.5

(c) Example visualisation of the visual search behaviours with noise level 0.9

Figure 6.10: The figures show examples of the visual search behaviours in the model. Figure 6.10a shows an example of the behaviours of the model at noise level 0.3, which matched the data for post-operation. The visualisation can clearly show that the model recognised the target by fixating on the features. Figure 6.10b shows an example with a noise level of 0.5 (matched the human data pre-operation, and the effect of the increase of the noise can be observed by looking at the fixations locations. The model seems to fixate on areas next to the objects rather than directly fixating on the features. Lastly, Figure 6.10c is a good example of how a short scan path can look like when there is high noise, which results in low accuracy.

It is important to report that Figure 6.10 does not exactly reflect the humans' visual behaviours, and it is only provided to illustrate how the model is searching for the target given the bounds proposed in the model.

In the next section, we will provide a more in-depth discussion about the results, the limitations of this study and what can be improved for future work.

6.12 Discussion

This chapter discusses implementing the model to two visual search tasks tested on people with cataracts. The goal of the task is to test some visual search measures in people with cataracts. We implemented the task to the same model as the crowding task, providing some flexibility to our model and the ability to plug in different tasks (Chapter 5). The two tasks were formulated as POMDP with some similarities and few differences to the crowding task. The differences were related to the state S, action space A, different hyper-parameters in the $Observation\ O$ and slightly different belief state in the face recognition task. These differences were implemented to simulate the task in the POMDP formulation based on the set-up of the human experiment and the estimated levels of vision impairment in the participants. The model followed that same methodology and implementation as in Chapter 5. It learned the optimal policy and showed accuracy similar to the humans. Due to the nature of our model, it was impossible to compare other measures to all the measures in the human data. However, the model's visual search strategies can give valuable insights into how people with blurred visual fields might interact with a visual search task.

The experiment done by Wan et al. (2020) showed that visual search behaviours drastically improve in people with cataracts after surgery. The participants in the experiment showed a lower number of fixations within the AOI, higher search time and lower accuracy before the surgery. Whereas, after surgery, the fixations within the AOI increased, the search time decreased, and the accuracy increased. Studying these visual search behaviours and trying to simulate them from a computer science perspective can give some insight and understanding of how people

with cataracts would interact with a computer screen with a specific design, which can benefit the HCI and enhance the user experience for those with cataracts.

Our model showed that computational rational models can predict (or give insight into) the visual behaviours of those with vision impairment by formulating the task as POMDP and giving suitable hyper-parameters which reflect the observations. The model showed lower accuracy at higher noise levels, which predicted the average accuracy of participants before undergoing surgery. Additionally, the visualisations we provided can reflect the search time of the model before finding the target. Our model currently does not distinguish areas outside the target's AOI (or within). This would give room for improvement to make a more sophisticated model that can distinguish between the different areas. Another limitation of this study is the lack of visual search experiments on people with cataracts, which was also mentioned by Wan et al. (2020). This made it difficult to find much comparison between our results and the visual search behaviours in people with cataracts. For example, in the face recognition task, at the noise level of 0.5, the model was fixating outside the area of interest and sometimes outside the screen limits (Appendix B, Section B.2). We know that this behaviour resulted from noise added to the coordinates to make the task more challenging for the model. This behaviour was only observed in the face recognition task. We believe that the reason is that the locations of the AOI in the face are not in the centre of the objects. They are distributed within the face (close to the edges of the object and the screen), resulting in fixations being located away from the AOI coordinates due to additional noise. This behaviour was not reported in the human experiment by Wan et al. (2020). However, an experiment done on people with glaucoma (Smith, Glen, and Crabb, 2012) showed that when a participant had a severe visual field defect, they had a long scan path and fixations away from the target. However, our model might still not reflect the behaviours of humans as simulating the visual behaviours using computational models is a complex approach due to the complexity of the human vision (Andersson et al., 2017; Eckstein, 2011; Fuhl et al., 2018).

This experiment has a lot of room for improvement. One main improvement is to update the model and make implementing the other measures possible. Another improvement is finding

a balance between the noise levels and the desired results by tuning the hyper-parameters. For example, removing the noise from the coordinates might make the search very easy for the model, and it would outperform the humans, yet it could prevent fixations outside the screen limits. Therefore, finding a good trade-off between the amount of noise and the model's behaviours could show some improvement in the results. We could also represent cataract vision impairment better by adding an additive noise to represent the blurriness in all areas including the foveal visual field.

In summary, this experiment showed that our approach of formulating the task as POMDP and using deep learning makes it possible to predict the behaviours in those with cataracts. However, given it is a new approach, more collaboration with vision researchers could benefit our implementations and theoretical approach.

In the next chapter, we will discuss the three experiments conducted in this thesis in more depth and give suggestions on how to improve this research in future work.

Chapter 7

General Discussion

This chapter discusses our findings, possible applications of the model, limitations and future work.

7.1 Insights and Findings from the Computational Model

The thesis presents a computational model of visual search behaviours in people with different types of vision impairments. The work offers an insight into eye movement strategies for those with vision impairment. Our theory suggests that those strategies emerged as a consequence of the visual system's bounds and the limitations in the environment. We tested the theory against three visual search tasks from the literature by formulating the tasks as a sequential decision process (i.e. POMDP) and solving them using RL and DRL methods.

The results of the experiments conducted in this research simulate the cognitive processes that lead to visual behaviour. It can provide valuable insights into computational modelling of vision impairment in HCI. Our search suggests that if interaction for people with visual impairment is to be improved, then HCI research must understand the nature of the problems faced by visually impaired users. The model that we have reported in this thesis explains visual search under impairment due to glaucoma, AMD and cataracts as a policy that is optimally adapted to the selective deficits imposed by the conditions. As a consequence of this adaptation, the model offers quantitative predictions of why people with glaucoma and AMD require more

fixations to find a target on computer screens, why this number increased when the crowding effect was introduced to those with AMD and why people with cataracts, before surgery, were less accurate in finding a target placed within distractors and needed more search time.

In Chapter 4, in the saccadic choice task with glaucoma patients, our results showed that patients with glaucoma require longer latencies and more fixations as eccentricities increased in the visual field with a decline in accuracy. Our results also showed that the values we reported were not always linear to the eccentricity, which could match some individual cases of those with glaucoma. This suggests that our model can be extended to model individual participants. Given the nature of glaucoma and how the visual field defect can differ from one person to another, the spatial noise in the observation function can be implemented for each participant based on their visual field.

In Chapter 5, which focused on the crowding effect in individuals with AMD, our results discussed the crowding effects in participants with AMD compared to groups of people with healthy vision. The results showed a decline in accuracy when the crowding effect is present with an increased number of fixations as an inefficient scan path, indicating the difficulties those with AMD would face when interacting with a screen with crowded objects. In this experiment, we tested two types of crowding noise that were added to the observations (i.e. Gaussian and Blending noise). The reason for experimenting with two types of crowding noise is to test which type would best fit the human data. The Gaussian noise produced more consistent results (e.g., decreased accuracy with higher noise and increased number of objects, longer search time, higher number of fixations and less efficient scan path ratio with increased noise). On the other hand, the Blending noise blends the flankers' features with the object features, adding a more realistic crowding effect. However, the downside is increased uncertainty, which in some cases produced a decrease in the number of fixations and inefficient scan paths at higher noise levels. At the same time, it showed a better fit to the human's accuracy in four and six objects. Furthermore, in this experiment, the model produced a faster search time than human data on average, but if we compare the search time in the model with the individual search time, it shows that the model matches the majority of AMD participants' search time. Overall, the experiment

in this chapter shows an estimation of the cognitive processes leading to some visual search behaviours in those with AMD.

Chapter 6 discusses implementing the model to a visual search that tested the behaviours in individuals with cataracts. The finding showed the ability of the computational model to predict how accurate people with cataracts can be when they interact with a visual search task that involves many objects on the screen or recognising facial features. The model predicted the accuracy of people with cataracts before and after surgery. We also provided eye movement simulations, which can reflect the search time of individuals with cataracts in the object search and face recognition tasks.

In summary, the explanation of those three experiments suggests that people with glaucoma, AMD and cataracts adapt to their condition by searching differently, and while their performance is still impaired, it is not as bad as it would have been had no adaptation occurred.

7.2 Computational Models and Machine Learning for Vision Impairment

Our research emphasises the importance of understanding the visual search behaviours of those with vision impairment from a psychological perspective by understanding the cognitive processes that lead to certain decisions and strategies in visual search. In this thesis, our aim is to analyse how someone with a specific type of vision impairment would interact with a given interface.

Computational models of interactions have been proven to better understand human interaction activities. Machine learning (ML) methods, such as deep learning, can produce an analysis of human behaviours and the mechanisms behind them (Li et al., 2020).

Our model is implemented in order to find what actions to take under uncertain observations, and to identify the optimal strategies through this uncertainty, which is a proven way to reflect on human strategies (Geisler, 2011).

Our proposed model employs a model-free RL and DRL algorithms to tackle the visual search

problems in vision impairment within a framework of human cognition theories. We chose this approach for its proven ability to provide a usable framework for modeling decision-making processes under uncertainty, which can represent human behaviours including visual search processes (Dayan and Daw, 2008). Therefore the aim of our model is to capture the optimal adaptation to the humans visual constraints. The RL learning agent could find the optimal visual search policy by formulating visual search as a sequential decision process (POMDP) which could present the statistical nature of the task. It also allows the representation of the internal processes in the learning agent and the agent-environment interaction, allowing the agent to find the optimal strategies for people with vision impairments through adaptation to the task and the cognitive limitations. The agent learns through reward maximisation, observations and state estimation. The use of a noisy observation representing an impaired visual field and introducing the hyper-parameter , *noise-level,* in the observation showed that using the method proposed can lead to an estimation of human behaviours without the need to train the ML agent with a big set of human data. Furthermore, the use of DRL methods in the experiments in Chapter 5 and 6 shows the methods abilities in finding optimal policies in complex problems in the human mind such as the problem of visual search in vision impairment. The methods proposed in this thesis can bridge the gap between cognitive theories and computational models.

7.3 Contributions to HCI

Based on previous research, utilising the methods used in this thesis, demonstrated their reliability in producing simulations of eye movements (Acharya, 2019; Acharya et al., 2017; Chen, 2015; Chen, Acharya, and Oulasvirta, 2021; Chen et al., 2015, 2017; Howes et al., 2018; Li et al., 2023; Olson and Olson, 1995; Oulasvirta, Jokinen, and Howes, 2022; Tseng and Howes, 2015). However, to the best of our knowledge, most of the previous studies only focused on understanding the behaviours of healthy vision rather than vision impairment. Li et al., 2023 implemented a model which could simulate the effect of menu length on visual search strategies in blind users without focusing on a specific type of vision impairment. The empirical findings

in this thesis provide a new understanding of interaction in an impaired visual field, which can be extended or combined with similar research in HCI to expand the outcome, producing significantly insightful results for HCI. For example, in our model, we consider the effect of the spatial location of an object, the spacing between objects and the number of objects and how the effect can differ from one vision impairment to another. These findings can help us understand the reason behind each interaction when it comes to a type of vision impairment. It can give an insight into how people with glaucoma would observe the objects located within the peripheral visual field, and what the eccentricity is when peripheral vision starts declining in a person with glaucoma. It can give an insight into how spatial spacing between objects effects the observation and perceiving of information in those with AMD and how someone with cataract might find it difficult to observe objects given on the screen. These insights can reflect the contributions of this study to HCI.

It is important to note that the model reported in this thesis is very much an initial exploration, and it has a big room for expansion and further development. In the next section, we will discuss the limitations of this research and the possible future work.

7.4 Limitations and Future Work

It is essential to recognise the limitations of this study. One limitation is the small number of tasks used in the study. Implementing the model using more tasks and measuring different visual search behaviours can provide better support to the results the model produced and extract more generalisable conclusions, for instance, testing the impact of the crowding effect using different tasks related to AMD or the wide visual field for those with glaucoma. It would also be useful to implement different spacings to our AMD model in Chapter 5 and examine when the model starts to have high accuracy (and stops to match the human behaviours). Another important limitation in the model is that we only implemented one feature to the objects in the visual search tasks (i.e. the $targetedness$). Future work needs to consider other features such as shape, colour and size, which can give a better insight into how people with vision impairments perceive the

differences in these features. Another major limitation of this study is the inability of the model to learn with some of the high noise levels in the observation in the experiment in Chapter 5. We believe this is caused by the randomness associated with the actions, which is caused by the decision action (i.e. 'press space bar'), which could cause the agent to always exploit under very high uncertainty rather than explore. Although this limitation did not affect the overall outcome of the results, as the model predicted human behaviours with different noise levels, our understanding could be improved by further experiments. We believe further investigating the reward function and the environment dynamic can improve the outcome (Steinparz et al., 2022). We also believe that it is worth experimenting with different RL methods to check if we get the same outcome. It is also worth investigating different elements in the task's specifications, such as the spacing between objects and the objects' size.

Another challenge lies in predicting all the measures of the human experiments. For example, in the experiment in Chapter 6, it was not possible to compare the measures in the human experiment to ours. However, we believe it to be achievable in future study by changing the implementation of the model. Another limitation is the duration of fixation we implemented in the model, which was based on a human experiment that was done on people with normal vision (Shen, Reingold, and Pomplun, 2003). Although the search time of our model matched the majority of the human participants, further experiments with the fixation duration of those with AMD might be more efficient. The way we represented the foveal and peripheral vision in the studies is another limitation that could have been improved by differentiating the areas of noise for each experiment; for example, in the AMD experiment in Chapter 5, we could have added noise in the central of fixation rather than assuming the person with AMD found an alternative fixation point. This improvement to the model could add more accuracy on how the model is implemented and the strategies found by the agent. Additionally, using grid search to find the best-fit hyper-parameters can be time-consuming and computationally expensive. We believe the first step to improve this limitation is to estimate the parameters using the Approximate Bayesian Computation (ABC) inference model. This method showed that it can obtain fixation duration from human experiments, predict the model's parameters and find the best fit for individuals

(Kangasrääsiö et al., 2017). We believe this is the next step forward to improve the model.

7.4.1 Improving User Interface Design and Accessibility

The findings in this study can lay the groundwork for future research into providing better design for UI in computers, phones and tablets using computational modelling. Tailoring the design of the interface with the visually impaired in mind can help better interactions for the visually impaired users, which can improve confidence in using those devices. Furthermore, the experiments we conducted were based on the visual loss experience of each type. For example, the experiment in Chapter 4 would be useful in designing interfaces specified for people with glaucoma in systems which include wider visual fields such as Virtual Reality (VR), Augmented Reality (AR) and widescreen simulators for gaming as the difficulties in glaucoma lay in the peripheral visual field. In Chapter 5, we considered the crowding effect and how spatial spacing can influence the decisions of those with vision AMD, as their difficulties lay in the centre of the visual fields, and crowded objects can cause more difficult interaction. Therefore, considering the spacing between objects that optimise their strategies is essential. In Chapter 6, we used two tasks to test the effects of the blurred visual field on those with cataracts, and the model predicted that if people with cataracts undergo surgery, they might have better strategies than before surgery. The different focus of each experiment can give an example of how perception differs from one vision impairment to another and what to take into consideration when designing user interfaces.

The results of this research support the idea that future interactive systems will more actively engage with users (Dafoe et al., 2021), anticipating their needs and adjusting accordingly. For example, an interactive system might automatically adjust the display layout, including button sizes, eccentricities and latencies, to an individual's visual impairment. Nevertheless, if such systems are to be successful, then they must do more than merely classify users into normal versus visually impaired. They must also be able to explain why users behave as they do. Currently, model-free systems for classification offer a limited capacity for explanation. The kind of model described in this thesis offers a better approach, in which behaviours are explained

by simulating models under varying conditions, e.g. glaucoma versus AMD. Such explanations could be used to adapt interaction design to individual users and test them (in silico) before they are tested on people.

7.5 Conclusion

In conclusion, the thesis reported a computational model that learns to perform visual search tasks given bounds that model psychological limits imposed by glaucoma, AMD and cataracts. The model behaviour corresponds well to previously reported human behaviour in three tasks. The thesis demonstrates that computational rationality can predict the visual search behaviours in those with vision impairments.

The project proposed in this thesis could be a used as a starting project, to be extended in the future, to support the design of more accessible interfaces that are built based on a better understanding of the psychology of vision impairment and the factors influencing the cognitive strategies of the decision-making process, such as the type of visual impairment, target size, shape and distance. It might also support the design of virtual reality and augmented reality environments, which might include functionality that is fitted to an individual user's visual 'profile'; there is no benefit in augmenting reality with labels that cannot be seen, but perhaps labels could be provided in areas of the peripheral visual field with greater acuity and less crowded environment. Achieving this goal would require models to be fitted to individual users.

The results in Chapters 5 and 6 were produced using a combination of the following services:

1. BluBear HPC service [1] at the University of Birmingham, which gives access to High-Performance Computing service.

2. Baskerville Tier 2 HPC service [2] which is funded by the EPSRC and UKRI under the World Class Labs scheme (EP/T022221/1) and the Digital Research Infrastructure program (EP/W032244/1). It is operated by the Advanced Research Computing team at

[1] http://www.birmingham.ac.uk/bear
[2] https://www.baskerville.ac.uk

the University of Birmingham

References

Abrams, R. A., Meyer, D. E., and Kornblum, S. (1989). 'Speed and accuracy of saccadic eye movements: characteristics of impulse variability in the oculomotor system.' *Journal of Experimental Psychology: Human Perception and Performance* 15 (3), p. 529.

Acharya, A. (2019). 'APPROXIMATE OPTIMAL CONTROL MODEL FOR VISUAL SEARCH TASKS'. PhD thesis.

Acharya, A., Chen, X., Myers, C. W., Lewis, R. L., and Howes, A. (2017). 'Human Visual Search as a Deep Reinforcement Learning Solution to a POMDP.' *CogSci*, pp. 51–56.

Ackermans, S., Dey, D., Ruijten, P., Cuijpers, R. H., and Pfleging, B. (2020). 'The effects of explicit intention communication, conspicuous sensors, and pedestrian attitude in interactions with automated vehicles'. *Proceedings of the 2020 chi conference on human factors in computing systems*, pp. 1–14.

Alexander, L. J. (2015). 'Age-Related Macular Degeneration, 2nd ed. D. Virgil Alfaro, III, John B. Kerrison, Kenneth A. Sharpe, and Monica Rodriguez-Fontal'. *Optometry and Vision Science* 92 (1). URL: https://journals.lww.com/optvissci/fulltext/2015/01000/age_related_macular_degeneration,_2nd_ed__d_.28.aspx.

Ali Hassan, E. and Tang, T. B. (2016). 'Smart glasses for the visually impaired people'. *Computers Helping People with Special Needs: 15th International Conference, ICCHP 2016, Linz, Austria, July 13-15, 2016, Proceedings, Part II 15*. Springer, pp. 579–582.

Altınbay, D. and İdil, Ş. A. (2022). 'Fixation stability and preferred retinal locus in advanced age-related macular degeneration'. *Turkish journal of ophthalmology* 52 (1), p. 23.

Altman, E. (1999). *Constrained Markov decision processes*. Vol. 7. CRC Press.

Anderson, J. R. (1991). 'Is human cognition adaptive?' *Behavioral and Brain Sciences* 14 (3), pp. 471–485.

Anderson, J. R. (1996). 'ACT: A simple theory of complex cognition.' *American psychologist* 51 (4), p. 355.

Anderson, J. R., Lebiere, C., Lovett, M., and Reder, L. (1998). 'ACT-R: A higher-level account of processing capacity'. *Behavioral and Brain Sciences* 21 (6), pp. 831–832.

Andersson, R., Larsson, L., Holmqvist, K., Stridh, M., and Nyström, M. (2017). 'One algorithm to rule them all? An evaluation and discussion of ten eye movement event-detection algorithms'. *Behavior research methods* 49, pp. 616–637.

Apple Inc. (2023). *Vision For every point of view.* Available at: https://www.apple.com/uk/accessibility/vision/#:~:text=VoiceOver%20%2B%20braille&text=VoiceOver%20describes%20exactly%20what's%20on,touchscreen%20using%20Braille%20Screen%20Input. (Accessed: 9 Aug. 2023).

Arulkumaran, K., Deisenroth, M. P., Brundage, M., and Bharath, A. A. (2017). 'Deep reinforcement learning: A brief survey'. *IEEE Signal Processing Magazine* 34 (6), pp. 26–38.

Asfaw, D. S., Jones, P. R., Mönter, V. M., Smith, N. D., and Crabb, D. P. (2018). 'Does glaucoma alter eye movements when viewing images of natural scenes? A between-eye study'. *Investigative ophthalmology & visual science* 59 (8), pp. 3189–3198.

Ashktorab, Z., Jain, M., Liao, Q. V., and Weisz, J. D. (2019). 'Resilient chatbots: Repair strategy preferences for conversational breakdowns'. *Proceedings of the 2019 CHI conference on human factors in computing systems*, pp. 1–12.

Awada, A., Issa, Y. B., Tekli, J., and Chbeir, R. (2013). 'Evaluation of touch screen vibration accessibility for blind users'. *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*, pp. 1–2.

Bacigalupo, F. and Luck, S. J. (2015). 'The allocation of attention and working memory in visual crowding'. *Journal of cognitive neuroscience* 27 (6), pp. 1180–1193.

Baloh, R. W., Sills, A. W., Kumley, W. E., and Honrubia, V. (1975). 'Quantitative measurement of saccade amplitude, duration, and velocity'. *Neurology* 25 (11), pp. 1065–1065.

Barakat, B., Steponenaite, A., Lall, G. S., Arshad, K., Wassell, I. J., and Keates, S. (2020). 'Assistive technology for the visually impaired: Optimizing frame rate (freshness) to improve the performance of real-time objects detection application'. *International Conference on Human-Computer Interaction*. Springer, pp. 479–492.

Baxes, G. A. (1994). *Digital image processing: principles and applications*. Wiley New York.

Bayle, D. J., Schoendorff, B., Hénaff, M.-A., and Krolak-Salmon, P. (2011). 'Emotional facial expression detection in the peripheral visual field'. *PloS one* 6 (6), e21584.

Bechtel, W. (2008). 'Mechanisms in cognitive psychology: What are the operations?' *Philosophy of Science* 75 (5), pp. 983–994.

Bekkering, H., Adam, J. J., Kingma, H., Huson, A., and Whiting, H. (1994). 'Reaction time latencies of eye and hand movements in single-and dual-task conditions'. *Experimental brain research* 97, pp. 471–476.

Benedek, G. (1971). 'Theory of transparency of the eye'. *Applied optics* 10 (3), pp. 459–473.

Bernard, J.-B. and Chung, S. T. (2018). 'Visual acuity is not the best at the preferred retinal locus in people with macular disease'. *Optometry and vision science: official publication of the American Academy of Optometry* 95 (9), p. 829.

Binetti, N., Wu, L., Chen, S., Kruijff, E., Julier, S., and Brumby, D. P. (2021). 'Using visual and auditory cues to locate out-of-view objects in head-mounted augmented reality'. *Displays* 69, p. 102032. ISSN: 0141-9382. DOI: https://doi.org/10.1016/j.displa.2021.102032. URL: https://www.sciencedirect.com/science/article/pii/S0141938221000433.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). 'The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks.' *Psychological review* 113 (4), p. 700.

Boucart, M., Bachet, V., Corveleyn, X., Bacchetti, P., and Rouland, J. F. (2020). 'A Saccadic Choice Task for Target Face Detection at Large Visual Eccentricities in Patients with Glaucoma'. *Optometry and Vision Science* 97 (10). URL: https://journals.lww.com/optvissci/fulltext/2020/10000/a_saccadic_choice_task_for_target_face_detection.7.aspx.

Boucart, M., Lenoble, Q., Quettelart, J., Szaffarczyk, S., Despretz, P., and Thorpe, S. J. (2016). 'Finding faces, animals, and vehicles in far peripheral vision'. *Journal of vision* 16 (2), pp. 10–10.

Boucart, M., Moroni, C., Thibaut, M., Szaffarczyk, S., and Greene, M. (2013). 'Scene categorization at large visual eccentricities'. *Vision Research* 86, pp. 35–42.

Bouma, H. (1970). 'Interaction effects in parafoveal letter recognition'. *Nature* 226 (5241), pp. 177–178.

Bowers, A. R., Keeney, K., and Peli, E. (2014). 'Randomized crossover clinical trial of real and sham peripheral prism glasses for hemianopia'. *JAMA ophthalmology* 132 (2), pp. 214–222.

Bringmann, A. and Wiedemann, P. (2022). 'Chapter 5 - Primate fovea'. *The Fovea*. Ed. by A. Bringmann and P. Wiedemann. Academic Press, pp. 83–117. ISBN: 978-0-323-90467-4. DOI: https://doi.org/10.1016/B978-0-323-90467-4.00001-2. URL: https://www.sciencedirect.com/science/article/pii/B9780323904674000012.

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. (2016). 'Openai gym'. *arXiv preprint arXiv:1606.01540*.

Brockmole, J. R. and Henderson, J. M. (2006). 'Short article: Recognition and attention guidance during contextual cueing in real-world scenes: Evidence from eye movements'. *Quarterly journal of experimental psychology* 59 (7), pp. 1177–1187.

Burton, R., Smith, N. D., and Crabb, D. P. (2014). 'Eye movements and reading in glaucoma: observations on patients with advanced visual field loss'. *Graefe's Archive for Clinical and Experimental Ophthalmology* 252 (10), pp. 1621–1630.

Butko, N. J. and Movellan, J. R. (2008). 'I-POMDP: An infomax model of eye movement'. *2008 7th IEEE International Conference on Development and Learning*. IEEE, pp. 139–144.

Byrne, M. D. (2001). 'ACT-R/PM and menu selection: Applying a cognitive architecture to HCI'. *International Journal of Human-Computer Studies* 55 (1), pp. 41–84.

Cannon, M. W. (1985). 'Perceived contrast in the fovea and periphery'. *JOSA A* 2 (10), pp. 1760–1768.

Chen, X. (2015). 'An optimal control approach to testing theories of human information processing constraints'. PhD thesis. University of Birmingham.

Chen, X., Acharya, A., and Oulasvirta, A. (2021). 'An adaptive model of gaze-based selection'. *CHI Conference on Human Factors in Computing Systems (CHI'21)*. Association for Computing Machinery.

Chen, X., Bailly, G., Brumby, D. P., Oulasvirta, A., and Howes, A. (2015). 'The emergence of interactive behavior: A model of rational menu search'. *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. ACM, pp. 4217–4226.

Chen, X., Mohr, K., and Galea, J. M. (2017). 'Predicting explorative motor learning using decision-making and motor noise'. *PLoS computational biology* 13 (4), e1005503.

Chen, X., Starke, S. D., Baber, C., and Howes, A. (2017). 'A cognitive model of how people make decisions through interaction with visual displays'. *Proceedings of the 2017 CHI conference on human factors in computing systems*. ACM, pp. 1205–1216.

Chung, S. T. (2014). 'Size or spacing: Which limits letter recognition in people with age-related macular degeneration?' *Vision research* 101, pp. 167–176.

Cillino, G., Casuccio, A., Pasti, M., Bono, V., Mencucci, R., and Cillino, S. (2014). 'Working-age cataract patients: visual results, reading performance, and quality of life with three diffractive multifocal intraocular lenses'. *Ophthalmology* 121 (1), pp. 34–44.

Cimarolli, V. R., Boerner, K., Brennan-Ing, M., Reinhardt, J. P., and Horowitz, A. (2012). 'Challenges faced by older adults with vision loss: a qualitative study with implications for rehabilitation'. *Clinical rehabilitation* 26 (8), pp. 748–757.

Crabb, D. P., Smith, N. D., Glen, F. C., Burton, R., and Garway-Heath, D. F. (2013). 'How does glaucoma look?: patient perception of visual field loss'. *Ophthalmology* 120 (6), pp. 1120–1126.

Crabb, D. P., Smith, N. D., Rauscher, F. G., Chisholm, C. M., Barbur, J. L., Edgar, D. F., and Garway-Heath, D. F. (2010). 'Exploring eye movements in patients with glaucoma when viewing a driving scene'. *PloS one* 5 (3), e9710.

Crossland, M. D., Culham, L. E., Kabanarou, S. A., and Rubin, G. S. (2005). 'Preferred retinal locus development in patients with macular disease'. *Ophthalmology* 112 (9), pp. 1579–1585.

Crouzet, S. M., Kirchner, H., and Thorpe, S. J. (2010). 'Fast saccades toward faces: face detection in just 100 ms'. *Journal of vision* 10 (4), pp. 16–16.

Dafoe, A., Bachrach, Y., Hadfield, G., Horvitz, E., Larson, K., and Graepel, T. (2021). *Cooperative AI: machines must learn to find common ground.*

Dayan, P. and Daw, N. D. (2008). 'Decision theory, reinforcement learning, and the brain'. *Cognitive, Affective, & Behavioral Neuroscience* 8 (4), pp. 429–453.

DeLine, R., Czerwinski, M., Meyers, B., Venolia, G., Drucker, S., and Robertson, G. (2006). 'Code thumbnails: Using spatial memory to navigate source code'. *Visual Languages and Human-Centric Computing (VL/HCC'06)*. IEEE, pp. 11–18.

Deruaz, A., Whatham, A., Mermoud, C., and Safran, A. (2002). 'Reading with multiple preferred retinal loci: implications for training a more efficient reading strategy'. *Vision research* 42 (27), pp. 2947–2957.

Dimara, E. and Perin, C. (2019). 'What is interaction for data visualization?' *IEEE transactions on visualization and computer graphics* 26 (1), pp. 119–129.

Duchowski, A. T. (2018). 'Gaze-based interaction: A 30 year retrospective'. *Computers & Graphics* 73, pp. 59–69.

Dudley, J. and Kristensson, P. O. (2022). 'Bayesian Optimisation of Interface Features'. *Bayesian Methods for Interaction and Design*, p. 259.

Eckstein, M. P. (2011). 'Visual search: A retrospective'. *Journal of vision* 11 (5), pp. 14–14.

Eckstein, M. P., Thomas, J. P., Palmer, J., and Shimozaki, S. S. (2000). 'A signal detection model predicts the effects of set size on visual search accuracy for feature, conjunction, triple conjunction, and disjunction displays'. *Perception & psychophysics* 62 (3), pp. 425–451.

Edwards, A. (1995). *Extraordinary Human-Computer Interaction: Interfaces for Users with Disabilities*. Vol. 7. CUP Archive.

Egan, D. E. (1988). 'Individual differences in human-computer interaction'. *Handbook of human-computer interaction*. Elsevier, pp. 543–568.

Egorov, M. (2015). 'Deep reinforcement learning with pomdps'. *Tech. Rep.(Technical Report, Stanford University, 2015), Tech. Rep.*

Elazary, L. and Itti, L. (2010). 'A Bayesian model for efficient visual search and recognition'. *Vision research* 50 (14), pp. 1338–1352.

Elliott, D. B., Patla, A. E., Furniss, M., and Adkin, A. (2000). 'Improvements in clinical and functional vision and quality of life after second eye cataract surgery'. *Optometry and vision science* 77 (1), pp. 13–24.

Engström, J., Johansson, E., and Östlund, J. (2005). 'Effects of visual and cognitive load in real and simulated motorway driving'. *Transportation research part F: traffic psychology and behaviour* 8 (2), pp. 97–120.

Etzioni, A. (1986). 'The case for a multiple-utility conception'. *Economics & Philosophy* 2 (2), pp. 159–184.

Faragher, R. (2012). 'Understanding the Basis of the Kalman Filter Via a Simple and Intuitive Derivation [Lecture Notes]'. *IEEE Signal Processing Magazine* 29 (5), pp. 128–132. DOI: 10.1109/MSP.2012.2203621.

Findlay, J. M. and Gilchrist, I. D. (2003). *Active vision: The psychology of looking and seeing*. 37. Oxford University Press.

Fitts, P. M. (1954). 'The information capacity of the human motor system in controlling the amplitude of movement.' *Journal of experimental psychology* 47 (6), p. 381.

Fletcher, D. C. and Schuchard, R. A. (1997). 'Preferred retinal loci relationship to macular scotomas in a low-vision population'. *Ophthalmology* 104 (4), pp. 632–638.

Folstein, M. F., Folstein, S. E., and McHugh, P. R. (1975). '"Mini-mental state": a practical method for grading the cognitive state of patients for the clinician'. *Journal of psychiatric research* 12 (3), pp. 189–198.

François-Lavet, V., Henderson, P., Islam, R., Bellemare, M. G., and Pineau, J. (2018). 'An Introduction to Deep Reinforcement Learning'. *Foundations and Trends® in Machine Learning*

11 (3-4), pp. 219–354. ISSN: 1935-8237. DOI: 10.1561/2200000071. URL: http://dx.doi.org/10.1561/2200000071.

Fuhl, W., Santini, T., Kuebler, T., Castner, N., Rosenstiel, W., and Kasneci, E. (2018). 'Eye movement simulation and detector creation to reduce laborious parameter adjustments'. *arXiv preprint arXiv:1804.00970*.

Geisler, W. S. (2011). 'Contributions of ideal observer theory to vision research'. *Vision research* 51 (7), pp. 771–781.

Glen, F. C., Smith, N. D., and Crabb, D. P. (2013). 'Saccadic eye movements and face recognition performance in patients with central glaucomatous visual field defects'. *Vision Research* 82, pp. 42–51.

Google (2023). *Android accessibility overview*. Available at: https://support.google.com/accessibility/android/answer/6006564?hl=en-GB (Accessed: 18 Aug. 2023).

Gray, W. D., Sims, C. R., Fu, W.-T., and Schoelles, M. J. (2006). 'The soft constraints hypothesis: a rational analysis approach to resource allocation for interactive behavior.' *Psychological review* 113 (3), p. 461.

Griffin-Shirley, N., Banda, D. R., Ajuwon, P. M., Cheon, J., Lee, J., Park, H. R., and Lyngdoh, S. N. (2017). 'A survey on the use of mobile applications for people who are visually impaired'. *Journal of Visual Impairment & Blindness* 111 (4), pp. 307–323.

Guez, J.-E., Le Gargasson, J.-F., Rigaudiere, F., and O'Regan, J. K. (1993). 'Is there a systematic location for the pseudo-fovea in patients with central scotoma?' *Vision research* 33 (9), pp. 1271–1279.

Halverson, T. and Hornof, A. J. (2011). 'A computational model of "active vision" for visual search in human–computer interaction'. *Human–Computer Interaction* 26 (4), pp. 285–314.

Hassan, S. E., Ross, N. C., Massof, R. W., and Stelmack, J. (2019). 'Changes in the properties of the preferred retinal locus with eccentric viewing training'. *Optometry and Vision Science* 96 (2), pp. 79–86.

Hassell, J. B., Lamoureux, E. L., and Keeffe, J. E. (2006). 'Impact of age related macular degeneration on quality of life'. *British Journal of Ophthalmology* 90 (5), pp. 593–596.

Hayhoe, M. and Ballard, D. (2005). 'Eye movements in natural behavior'. *Trends in cognitive sciences* 9 (4), pp. 188–194.

Hayhoe, M. and Ballard, D. (2014). 'Modeling task control of eye movements'. *Current Biology* 24 (13), R622–R628.

Haymes, S. A., LeBlanc, R. P., Nicolela, M. T., Chiasson, L. A., and Chauhan, B. C. (2008). 'Glaucoma and on-road driving performance'. *Investigative Ophthalmology & Visual Science* 49 (7), pp. 3035–3041.

Henderson, J. M. (2003). 'Human gaze control during real-world scene perception'. *Trends in cognitive sciences* 7 (11), pp. 498–504.

Higgins, B. E., Taylor, D. J., Bi, W., Binns, A. M., and Crabb, D. P. (2020). 'Novel computer-based assessments of everyday visual function in people with age-related macular degeneration'. *Plos one* 15 (12), e0243578.

Hochberg, C., Maul, E., Chan, E. S., Van Landingham, S., Ferrucci, L., Friedman, D. S., and Ramulu, P. Y. (2012). 'Association of vision loss in glaucoma and age-related macular degeneration with IADL disability'. *Investigative ophthalmology & visual science* 53 (6), pp. 3201–3206.

Hofer, H., Singer, B., and Williams, D. R. (2005). 'Different sensations from cones with the same photopigment'. *Journal of Vision* 5 (5), pp. 5–5.

Hollan, J., Hutchins, E., and Kirsh, D. (2000). 'Distributed cognition: toward a new foundation for human-computer interaction research'. *ACM Transactions on Computer-Human Interaction (TOCHI)* 7 (2), pp. 174–196.

Holton, B. (2014). 'A review of iOS access for all: Your comprehensive guide to accessibility for iPad, iPhone, and iPod touch, by Shelly Brisbin'. *AccessWorld Magazine* 15 (7).

Hood, D. C. (2017). 'Improving our understanding, and detection, of glaucomatous damage: an approach based upon optical coherence tomography (OCT)'. *Progress in retinal and eye research* 57, pp. 46–75.

Hosseini, M. (2017). 'View-aware tile-based adaptations in 360 virtual reality video streaming'. *2017 IEEE Virtual Reality (VR)*. IEEE, pp. 423–424.

Howard, R. A. (1960). 'Dynamic programming and markov processes.'

Howes, A., Chen, X., Acharya, A., and Lewis, R. L. (2018). 'Interaction as an emergent property of a partially observable Markov decision process'. *Computational interaction*, pp. 287–310.

Howes, A., Duggan, G. B., Kalidindi, K., Tseng, Y.-C., and Lewis, R. L. (2016). 'Predicting short-term remembering as boundedly optimal strategy choice'. *Cognitive Science* 40 (5), pp. 1192–1223.

Howes, A., Lewis, R. L., and Vera, A. (2009). 'Rational adaptation under task and processing constraints: implications for testing theories of cognition and action.' *Psychological review* 116 (4), p. 717.

Hu, C. X., Zangalli, C., Hsieh, M., Gupta, L., Williams, A. L., Richman, J., and Spaeth, G. L. (2014). 'What do patients with glaucoma see? Visual symptoms reported by patients with glaucoma'. *The American journal of the medical sciences* 348 (5), pp. 403–409.

Huang, H. (2018). 'Blind users' expectations of touch interfaces: factors affecting interface accessibility of touchscreen-based smartphones for people with moderate visual impairment'. *Universal Access in the Information Society* 17, pp. 291–304.

Itti, L. and Koch, C. (2001). 'Computational modelling of visual attention'. *Nature reviews neuroscience* 2 (3), pp. 194–203.

Jager, R. D., Mieler, W. F., and Miller, J. W. (2008). 'Age-related macular degeneration'. *New England Journal of Medicine* 358 (24), pp. 2606–2617.

Jansen, Y., Karrer, T., and Borchers, J. (2010). 'MudPad: tactile feedback and haptic texture overlay for touch surfaces'. *ACM International Conference on Interactive Tabletops and Surfaces*, pp. 11–14.

Johnson, C. A. and Keltner, J. L. (1983). 'Incidence of visual field loss in 20,000 eyes and its relationship to driving performance'. *Archives of ophthalmology* 101 (3), pp. 371–375.

Johnson, J. (2020). *Designing with the mind in mind: simple guide to understanding user interface design guidelines*. Morgan Kaufmann.

Jung, K. I., Jeon, S., Kim, Y. C., and Park, C. K. (2019). 'Comparison of pattern electroretinograms of glaucoma patients with parafoveal scotoma versus peripheral nasal step'. *Scientific reports* 9 (1), pp. 1–8.

Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. (1998). 'Planning and acting in partially observable stochastic domains'. *Artificial intelligence* 101 (1-2), pp. 99–134.

Kahneman, D. and Tversky, A. (2013). 'Prospect theory: An analysis of decision under risk'. *Handbook of the fundamentals of financial decision making: Part I*. World Scientific, pp. 99–127.

Kane, S. K., Morris, M. R., Perkins, A. Z., Wigdor, D., Ladner, R. E., and Wobbrock, J. O. (2011). 'Access overlays: improving non-visual access to large touch screens for blind users'. *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pp. 273–282.

Kane, S. K., Morris, M. R., and Wobbrock, J. O. (2013). 'Touchplates: low-cost tactile overlays for visually impaired touch screen users'. *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*, pp. 1–8.

Kangasrääsiö, A., Athukorala, K., Howes, A., Corander, J., Kaski, S., and Oulasvirta, A. (2017). 'Inferring cognitive models from data using approximate Bayesian computation'. *Proceedings of the 2017 CHI conference on human factors in computing systems*. ACM, pp. 1295–1306.

Kanjee, R., Yücel, Y. H., Steinbach, M. J., González, E. G., and Gupta, N. (2012). 'Delayed saccadic eye movements in glaucoma'. *Eye and Brain* 4, p. 63.

Kasneci, E., Black, A. A., and Wood, J. M. (2017). 'Eye-tracking as a tool to evaluate functional ability in everyday tasks in glaucoma'. *Journal of ophthalmology* 2017.

Kennedy, G. J. and Whitaker, D. (2010). 'The chromatic selectivity of visual crowding'. *Journal of Vision* 10 (6), pp. 15–15.

Khan, A. and Khusro, S. (2019). 'Blind-friendly user interfaces–a pilot study on improving the accessibility of touchscreen interfaces'. *Multimedia Tools and Applications* 78, pp. 17495–17519.

Kieras, D. E. and Hornof, A. J. (2014). 'Towards accurate and practical predictive models of active-vision-based visual search'. *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, pp. 3875–3884.

Kieras, D. E. and Meyer, D. E. (1997). 'An overview of the EPIC architecture for cognition and performance with application to human-computer interaction'. *Human–Computer Interaction* 12 (4), pp. 391–438.

Knox, W. B., Otto, A. R., Stone, P., and Love, B. (2012). 'The nature of belief-directed exploratory choice in human decision-making'. *Frontiers in psychology* 2, p. 398.

Kotecha, A., Richardson, G., Chopra, R., Fahy, R. T., Garway-Heath, D. F., and Rubin, G. S. (2012). 'Balance control in glaucoma'. *Investigative ophthalmology & visual science* 53 (12), pp. 7795–7801.

Krösl, K., Elvezio, C., Luidolt, L. R., Hürbe, M., Karst, S., Feiner, S., and Wimmer, M. (2020). 'CatARact: Simulating cataracts in augmented reality'. *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, pp. 682–693.

Kuber, R., Hastings, A., and Tretter, M. (2020). 'Determining the accessibility of mobile screen readers for blind users'. *UMBC Faculty Collection*.

Kübler, T. C., Kasneci, E., Rosenstiel, W., Heister, M., Aehling, K., Nagel, K., Schiefer, U., and Papageorgiou, E. (2015). 'Driving with glaucoma: task performance and gaze movements'. *Optometry and Vision Science* 92 (11), pp. 1037–1046.

Kulyukin, V., Gharpure, C., Nicholson, J., and Pavithran, S. (2004). 'RFID in robot-assisted indoor navigation for the visually impaired'. *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*. Vol. 2. IEEE, pp. 1979–1984.

Lamoureux, E. L., Fenwick, E., Pesudovs, K., and Tan, D. (2011). 'The impact of cataract surgery on quality of life'. *Current opinion in ophthalmology* 22 (1), pp. 19–27.

Lavie, N. (2010). 'Attention, distraction, and cognitive control under load'. *Current directions in psychological science* 19 (3), pp. 143–148.

LeCun, Y., Bengio, Y., and Hinton, G. (2015). 'Deep learning'. *nature* 521 (7553), p. 436.

Leibo, J. Z., de Masson d'Autume, C., Zoran, D., Amos, D., Beattie, C., Anderson, K., Castañeda, A. G., Sanchez, M., Green, S., Gruslys, A., Legg, S., Hassabis, D., and Botvinick, M. M. (2018). 'Psychlab: A Psychology Laboratory for Deep Reinforcement Learning Agents'. *CoRR* abs/1801.08116. arXiv: 1801.08116. URL: http://arxiv.org/abs/1801.08116.

Leiva, L. A., Shiripour, M., and Oulasvirta, A. (2022). 'Modeling how different user groups perceive webpage aesthetics'. *Universal Access in the Information Society*, pp. 1–8.

Lenoble, Q., Lek, J. J., and McKendrick, A. M. (2016). 'Visual object categorisation in people with glaucoma'. *British Journal of Ophthalmology* 100 (11), pp. 1585–1590.

Leporini, B., Buzzi, M. C., and Buzzi, M. (2012). 'Interacting with mobile devices via VoiceOver: usability and accessibility issues'. *Proceedings of the 24th Australian Computer-Human Interaction Conference*, pp. 339–348.

Levi, D. M. (2011). 'Visual crowding'. *Current Biology* 21 (18), R678–R679.

Lewis, R. L., Howes, A., and Singh, S. (2014). 'Computational rationality: Linking mechanism and behavior through bounded utility maximization'. *Topics in cognitive science* 6 (2), pp. 279–311.

Li, Y., Kumar, R., Lasecki, W. S., and Hilliges, O. (2020). 'Artificial intelligence for HCI: a modern approach'. *Extended Abstracts of the 2020 CHI conference on human factors in computing systems*, pp. 1–8.

Li, Z., Ko, Y.-J., Putkonen, A., Feiz, S., Ashok, V., Ramakrishnan, I., Oulasvirta, A., and Bi, X. (2023). 'Modeling Touch-based Menu Selection Performance of Blind Users via Reinforcement Learning'. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–18.

Lim, L. S., Mitchell, P., Seddon, J. M., Holz, F. G., and Wong, T. Y. (2012). 'Age-related macular degeneration'. *The Lancet* 379 (9827), pp. 1728–1738. ISSN: 0140-6736. DOI: https://doi.org/10.1016/S0140-6736(12)60282-7. URL: https://www.sciencedirect.com/science/article/pii/S0140673612602827.

Littman, M. L. (1996). *Algorithms for sequential decision making*. Brown University Providence, RI.

Liu, R. and Kwon, M. (2020). 'Increased Equivalent Input Noise in Glaucomatous Central Vision: Is it Due to Undersampling of Retinal Ganglion Cells?' *Investigative Ophthalmology & Visual Science* 61 (8), pp. 10–10.

Loschky, L. C., Szaffarczyk, S., Beugnet, C., Young, M. E., and Boucart, M. (2019). 'The contributions of central and peripheral vision to scene-gist recognition with a 180 visual field'. *Journal of vision* 19 (5), pp. 15–15.

Loughman, J., Davison, P., and Flitcroft, I. (2007). 'Open angle glaucoma effects on preattentive visual search efficiency for flicker, motion displacement and orientation pop-out tasks'. *British Journal of Ophthalmology* 91 (11), pp. 1493–1498.

Lovejoy, W. S. (1991). 'A survey of algorithmic methods for partially observed Markov decision processes'. *Annals of Operations Research* 28 (1), pp. 47–65.

Luo, G. and Peli, E. (2006). 'Use of an augmented-vision device for visual search by patients with tunnel vision'. *Investigative ophthalmology & visual science* 47 (9), pp. 4152–4159.

Mackay, W. E. (1988). 'Diversity in the use of electronic mail: A preliminary inquiry'. *ACM Transactions on Information Systems (TOIS)* 6 (4), pp. 380–397.

MacKenzie, I. S. (2018). 'Fitts' law'. *The wiley handbook of human computer interaction* 1, pp. 347–370.

Macular Society (2022). *What is the macula?* Available at: https://www.macularsociety.org/macular-disease/macula/ (Accessed: 1 Sept. 2023).

Majaranta, P., Räihä, K.-J., Hyrskykari, A., and Špakov, O. (2019). 'Eye movements and human-computer interaction'. *Eye movement research: An introduction to its scientific foundations and applications*, pp. 971–1015.

Mao, C., Go, K., Kinoshita, Y., Kashiwagi, K., Toyoura, M., Fujishiro, I., Li, J., and Mao, X. (2021). 'Different Eye Movement Behaviors Related to Artificial Visual Field Defects–A Pilot Study of Video-Based Perimetry'. *IEEE Access* 9, pp. 77649–77660.

Mathew, R. S., Delbaere, K., Lord, S. R., Beaumont, P., and Madigan, M. C. (2011). 'Depressive symptoms and quality of life in people with age-related macular degeneration'. *Ophthalmic and Physiological Optics* 31 (4), pp. 375–380.

McClelland, J. L. and Cleeremans, A. (2009). *Consciousness and Connectionist Models*. Oxford University Press.

Mi, N., Cavuoto, L. A., Benson, K., Smith-Jackson, T., and Nussbaum, M. A. (2014). 'A heuristic checklist for an accessible smartphone interface design'. *Universal access in the information society* 13, pp. 351–365.

Michel, M. and Geisler, W. S. (2011). 'Intrinsic position uncertainty explains detection and localization performance in peripheral vision'. *Journal of Vision* 11 (1), pp. 18–18.

Mnih, V., Heess, N., Graves, A., and kavukcuoglu koray, k. (2014). 'Recurrent Models of Visual Attention'. *Advances in Neural Information Processing Systems 27*. Ed. by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger. Curran Associates, Inc., pp. 2204–2212. URL: http://papers.nips.cc/paper/5542-recurrent-models-of-visual-attention.pdf.

Monahan, G. E. (1982). 'State of the art—a survey of partially observable Markov decision processes: theory, models, and algorithms'. *Management science* 28 (1), pp. 1–16.

Mukhiddinov, M. and Cho, J. (2021). 'Smart glass system using deep learning for the blind and visually impaired'. *Electronics* 10 (22), p. 2756.

Myers, C., Lewis, R., and Howes, A. (2013). 'Bounded optimal state estimation and control in visual search: Explaining distractor ratio effects'. *Proceedings of the Annual Meeting of the Cognitive Science Society*. Vol. 35. 35.

Najemnik, J. and Geisler, W. S. (2005). 'Optimal eye movement strategies in visual search'. *Nature* 434 (7031), pp. 387–391.

Najemnik, J. and Geisler, W. S. (2008). 'Eye movement statistics in humans are consistent with an optimal search strategy'. *Journal of Vision* 8 (3), pp. 4–4.

National Eye Institute (2021). *At a glance: AMD*. Available at: https://www.nei.nih.gov/learn-about-eye-health/eye-conditions-and-diseases/age-related-macular-degeneration (Accessed: 1 Sept. 2023).

National Eye Institute (2022). *At a glance: Glaucoma*. Available at: https://www.nei.nih.gov/learn-about-eye-health/eye-conditions-and-diseases/glaucoma (Accessed: 17 July 2023).

National Eye Institute (2023). *At a glance: Cataracts*. Available at: https://www.nei.nih.gov/learn-about-eye-health/eye-conditions-and-diseases/cataracts (Accessed: 25 Aug. 2023).

Nelson, P., Aspinall, P., and O'Brien, C. (1999). 'Patients' perception of visual impairment in glaucoma: a pilot study'. *British Journal of Ophthalmology* 83 (5), pp. 546–552. ISSN: 0007-1161. DOI: 10.1136/bjo.83.5.546. eprint: https://bjo.bmj.com/content/83/5/546.full.pdf. URL: https://bjo.bmj.com/content/83/5/546.

Nunez-Varela, J. and Wyatt, J. L. (2013). 'Models of gaze control for manipulation tasks'. *ACM Transactions on Applied Perception (TAP)* 10 (4), pp. 1–22.

Oflaz, A. B., Öztürk, B. T., Gönül, Ş., Bakbak, B., Gedik, Ş., and Okudan, S. (2022). 'Short-Term Clinical Results of Preferred Retinal Locus Training'. *Turkish Journal of Ophthalmology* 52 (1), p. 14.

Olson, J. R. and Olson, G. M. (1995). 'The growth of cognitive modeling in human-computer interaction since GOMS'. *Readings in Human–Computer Interaction*. Elsevier, pp. 603–625.

Ono, S., Das, V. E., and Mustari, M. J. (2012). 'Conjugate adaptation of smooth pursuit during monocular viewing in strabismic monkeys with exotropia'. *Investigative Ophthalmology & Visual Science* 53 (4), pp. 2038–2045.

Oulasvirta, A., Bi, X., and Howes, A. (2018). *Computational interaction*. Oxford University Press.

Oulasvirta, A., Jokinen, J. P., and Howes, A. (2022). 'Computational rationality as a theory of interaction'. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–14.

Oviatt, S., Lunsford, R., and Coulston, R. (2005). 'Individual differences in multimodal integration patterns: What are they and why do they exist?' *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 241–249.

Pal, J., Viswanathan, A., Chandra, P., Nazareth, A., Kameswaran, V., Subramonyam, H., Johri, A., Ackerman, M. S., and O'Modhrain, S. (2017). 'Agency in assistive technology adoption: visual impairment and smartphone use in Bangalore'. *Proceedings of the 2017 CHI conference on human factors in computing systems*, pp. 5929–5940.

Payne, S. J. and Howes, A. (2013). 'Adaptive interaction: A utility maximization approach to understanding human interaction with technology'. *Synthesis Lectures on Human-Centered Informatics* 6 (1), pp. 1–111.

Peli, E. (2001). 'Vision multiplexing: an engineering approach to vision rehabilitation device development'. *Optometry and Vision Science* 78 (5), pp. 304–315.

Pelli, D. G. (2008). 'Crowding: A cortical constraint on object recognition'. *Current opinion in neurobiology* 18 (4), pp. 445–451.

Pirolli, P. L. T. (2007). *Information Foraging Theory: Adaptive Interaction with Information*. 1st ed. USA: Oxford University Press, Inc. ISBN: 0195173325.

Plopski, A., Hirzle, T., Norouzi, N., Qian, L., Bruder, G., and Langlotz, T. (2022). 'The eye in extended reality: A survey on gaze interaction and eye tracking in head-worn extended reality'. *ACM Computing Surveys (CSUR)* 55 (3), pp. 1–39.

Poupyrev, I. and Maruyama, S. (2003). 'Tactile interfaces for small touch screens'. *Proceedings of the 16th annual ACM symposium on User interface software and technology*, pp. 217–220.

Rajendran, P. S., Krishnan, P., and Aravindhar, D. J. (2020). 'Design and Implementation of Voice Assisted Smart Glasses for Visually Impaired People Using Google Vision API'. *2020 4th International Conference on Electronics, Communication and Aerospace Technology (ICECA)*. IEEE, pp. 1221–1224.

Ramrattan, R. S., Wolfs, R. C., Panda-Jonas, S., Jonas, J. B., Bakker, D., Pols, H. A., Hofman, A., and de Jong, P. T. (2001). 'Prevalence and causes of visual field loss in the elderly and associations with impairment in daily functioning: the Rotterdam Study'. *Archives of Ophthalmology* 119 (12), pp. 1788–1794.

Ramulu, P. Y., Swenor, B. K., Jefferys, J. L., Friedman, D. S., and Rubin, G. S. (2013). 'Difficulty with out-loud and silent reading in glaucoma'. *Investigative ophthalmology & visual science* 54 (1), pp. 666–672.

Ramulu, P. Y., Van Landingham, S. W., Massof, R. W., Chan, E. S., Ferrucci, L., and Friedman, D. S. (2012). 'Fear of falling and visual field loss from glaucoma'. *Ophthalmology* 119 (7), pp. 1352–1358.

Ramulu, P. Y., West, S. K., Munoz, B., Jampel, H. D., and Friedman, D. S. (2009). 'Driving cessation and driving limitation in glaucoma: the Salisbury Eye Evaluation Project'. *Ophthalmology* 116 (10), pp. 1846–1853.

Rao, R. P. (2010). 'Decision making under uncertainty: a neural model based on partially observable markov decision processes'. *Frontiers in computational neuroscience* 4, p. 146.

Rathfux, T., Thöner, J., Kaindl, H., and Popp, R. (2018). 'Combining Design-Time Generation of Web-Pages with Responsive Design for Improving Low-Vision Accessibility'. *Proceedings of the ACM SIGCHI Symposium on Engineering Interactive Computing Systems*. EICS '18. Paris, France: Association for Computing Machinery. ISBN: 9781450358972. DOI: 10.1145/3220134.3220141. URL: https://doi.org/10.1145/3220134.3220141.

Rayner, K. and Castelhano, M. (2007). 'Eye movements'. *Scholarpedia* 2 (10). revision #126973, p. 3649. DOI: 10.4249/scholarpedia.3649.

Rayner, K. (1998). 'Eye movements in reading and information processing: 20 years of research.' *Psychological bulletin* 124 (3), p. 372.

Rayner, K. (2009). 'The 35th Sir Frederick Bartlett Lecture: Eye movements and attention in reading, scene perception, and visual search'. *Quarterly journal of experimental psychology* 62 (8), pp. 1457–1506.

Raza, M. I. J. A. (2018). 'Subjective Quality of Vision before and after Cataract Surgery at Holy Family Hospital, Rawalpindi'. *Pakistan Journal of Ophthalmology* 34 (2).

Renner, P. and Pfeiffer, T. (2017). '[POSTER] Augmented Reality Assistance in the Central Field-of-View Outperforms Peripheral Displays for Order Picking: Results from a Virtual Reality Simulation Study'. *2017 IEEE international symposium on mixed and augmented reality (ISMAR-Adjunct)*. IEEE, pp. 176–181.

Robinson, D. (1975). 'How the oculomotor system repairs itself.' *Investigative ophthalmology* 14 (6), pp. 413–415.

Robinson, D. A. (1963). 'A method of measuring eye movemnent using a scieral search coil in a magnetic field'. *IEEE Transactions on bio-medical electronics* 10 (4), pp. 137–145.

Rubin, G. S. and Feely, M. (2009). 'The role of eye movements during reading in patients with age-related macular degeneration (AMD)'. *Neuro-Ophthalmology* 33 (3), pp. 120–126.

Russell, S. J. and Norvig, P. (2016). *Artificial intelligence: a modern approach*. Malaysia; Pearson Education Limited,

Russell, S. J. and Subramanian, D. (1994). 'Provably bounded-optimal agents'. *Journal of Artificial Intelligence Research* 2, pp. 575–609.

Salvucci, D. D. (2001). 'An integrated model of eye movements and visual encoding'. *Cognitive Systems Research* 1 (4), pp. 201–220.

Sandywell, B. (2016). *Dictionary of visual discourse: a dialectical lexicon of terms*. Routledge.

Sarcar, S., Jokinen, J. P., Oulasvirta, A., Wang, Z., Silpasuwanchai, C., and Ren, X. (2018). 'Ability-based optimization of touchscreen interactions'. *IEEE Pervasive Computing* 17 (1), pp. 15–26.

Schneider, R. M., Thurtell, M. J., Eisele, S., Lincoff, N., Bala, E., and Leigh, R. J. (2013). 'Neurological basis for eye movements of the blind'. *PloS one* 8 (2), e56556.

Schooler, L. J. and Anderson, J. R. (1997). 'The role of process in the rational analysis of memory'. *Cognitive Psychology* 32 (3), pp. 219–250.

Schubert, M. C. and Zee, D. S. (2010). 'Saccade and vestibular ocular motor adaptation'. *Restorative neurology and neuroscience* 28 (1), pp. 9–18.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). 'Proximal policy optimization algorithms'. *arXiv preprint arXiv:1707.06347*.

Shanidze, N., Fusco, G., Potapchuk, E., Heinen, S., and Verghese, P. (2016). 'Smooth pursuit eye movements in patients with macular degeneration'. *Journal of vision* 16 (3), pp. 1–1.

Shen, J., Reingold, E. M., and Pomplun, M. (2003). 'Guidance of eye movements during conjunctive visual search: the distractor-ratio effect.' *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale* 57 (2), p. 76.

Shi, Z. (2021). 'Chapter 1 - Introduction'. *Intelligence Science*. Ed. by Z. Shi. Elsevier, pp. 1–31. ISBN: 978-0-323-85380-4. DOI: https://doi.org/10.1016/B978-0-323-85380-4.00001-4. URL: https://www.sciencedirect.com/science/article/pii/B9780323853804000014.

Shiels, A. and Hejtmancik, J. F. (2007). 'Genetic origins of cataract'. *Archives of ophthalmology* 125 (2), pp. 165–173.

Shiels, A. and Hejtmancik, J. (2013). 'Genetics of human cataract'. *Clinical genetics* 84 (2), pp. 120–127.

Shima, N., Markowitz, S. N., and Reyes, S. V. (2010). 'Concept of a functional retinal locus in age-related macular degeneration'. *Canadian Journal of Ophthalmology* 45 (1), pp. 62–66.

Sight, S. 4. (2023). *How bad does your visual impairment need to be to be registered as a blind or partially sighted person?* Available at: https://www.support4sight.org.uk/how-we-can-help-you/eye-conditions/registered-as-blind-or-partially-sighted/ (Accessed: 2 July 2023).

Singh, S., Lewis, R. L., Barto, A. G., and Sorg, J. (2010). 'Intrinsically motivated reinforcement learning: An evolutionary perspective'. *IEEE Transactions on Autonomous Mental Development* 2 (2), pp. 70–82.

Skiadaresi, E., McAlinden, C., Pesudovs, K., Polizzi, S., Khadka, J., and Ravalico, G. (2012). 'Subjective quality of vision before and after cataract surgery'. *Archives of ophthalmology* 130 (11), pp. 1377–1382.

Smith, N. D., Crabb, D. P., Glen, F. C., Burton, R., and Garway-Heath, D. F. (2012). 'Eye movements in patients with glaucoma when viewing images of everyday scenes'. *Seeing and perceiving* 25 (5), pp. 471–492.

Smith, N. D., Glen, F. C., and Crabb, D. P. (2012). 'Eye movements during visual search in patients with glaucoma'. *BMC ophthalmology* 12 (1), p. 45.

Snowden, R., Snowden, R. J., Thompson, P., and Troscianko, T. (2012). *Basic vision: an introduction to visual perception*. Oxford University Press.

Son, G., Mack, M. L., and Walther, D. B. (2023). 'Feature integration in visual search for real-world scenes'. *Journal of Vision* 23 (9), pp. 4832–4832.

Spector, R. H. (1990). *Clinical Methods: The History, Physical, and Laboratory Examinations. 3rd edition.* Boston: Butterworths;

Sprague, N., Ballard, D., and Robinson, A. (2007). 'Modeling embodied visual behaviors'. *ACM Transactions on Applied Perception (TAP)* 4 (2), p. 11.

Stangl, A., Kim, J., and Yeh, T. (2014). '3D printed tactile picture books for children with visual impairments: a design probe'. *Proceedings of the 2014 conference on Interaction design and children*, pp. 321–324.

Steinparz, C. A., Schmied, T., Paischer, F., Dinu, M.-C., Patil, V. P., Bitto-Nemling, A., Eghbal-zadeh, H., and Hochreiter, S. (2022). 'Reactive exploration to cope with non-stationarity in lifelong reinforcement learning'. *Conference on Lifelong Learning Agents*. PMLR, pp. 441–469.

Stone, J. V. (2013). 'Bayes' rule: a tutorial introduction to Bayesian analysis'.

Strasburger, H. (2020). 'Seven myths on crowding and peripheral vision'. *i-Perception* 11 (3), p. 2041669520913052.

Strasburger, H., Harvey, L. O., and Rentschler, I. (1991). 'Contrast thresholds for identification of numeric characters in direct and eccentric view'. *Perception & psychophysics* 49 (6), pp. 495–508.

Sue, S. (2007). 'Test distance vision using a Snellen chart.' *Community Eye Health* 20 (63), p. 52.

Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.

Szlyk, J. P., Mahler, C. L., Seiple, W., Edward, D. P., and Wilensky, J. T. (2005). 'Driving performance of glaucoma patients correlates with peripheral visual field loss'. *Journal of glaucoma* 14 (2), pp. 145–150.

Tadin, D., Nyquist, J. B., Lusk, K. E., Corn, A. L., and Lappin, J. S. (2012). 'Peripheral vision of youths with low vision: motion perception, crowding, and visual search'. *Investigative ophthalmology & visual science* 53 (9), pp. 5860–5868.

Taylor, D. J., Hobby, A. E., Binns, A. M., and Crabb, D. P. (2016). 'How does age-related macular degeneration affect real-world visual ability and quality of life? A systematic review'. *BMJ open* 6 (12), e011504.

Taylor, D. J., Smith, N. D., Binns, A. M., and Crabb, D. P. (2018). 'The effect of non-neovascular age-related macular degeneration on face recognition performance'. *Graefe's Archive for Clinical and Experimental Ophthalmology* 256 (4), pp. 815–821.

Thepass, G., Pel, J., Vermeer, K., Creten, O., Bryan, S., Lemij, H., and van der Steen, J. (2015). 'The effect of cataract on eye movement perimetry'. *Journal of ophthalmology* 2015.

Therrien, A. S., Wolpert, D. M., and Bastian, A. J. (2018). 'Increasing motor noise impairs reinforcement learning in healthy individuals'. *eneuro* 5 (3).

Thibaut, M., Boucart, M., and Tran, T. H. C. (2020). 'Object search in neovascular age-related macular degeneration: the crowding effect'. *Clinical and Experimental Optometry* 103 (5), pp. 648–655.

Thibaut, M., Delerue, C., Boucart, M., and Tran, T. H. C. (2016). 'Visual exploration of objects and scenes in patients with age-related macular degeneration'. *Journal Francais d'Ophtalmologie* 39 (1), pp. 82–89.

Thorpe, S. J., Gegenfurtner, K. R., Fabre-Thorpe, M., and BuÈlthoff, H. H. (2001). 'Detection of animals in natural images using far peripheral vision'. *European Journal of Neuroscience* 14 (5), pp. 869–876.

Tigwell, G. W., Gorman, B. M., and Menzies, R. (2020). 'Emoji accessibility for visually impaired people'. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–14.

Todrov, E. V. (1998). 'Studies of goal directed movements'. PhD thesis. Massachusetts Institute of Technology.

Torralba, A., Oliva, A., Castelhano, M. S., and 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), p. 766.

Traquair, H. M. (M. (1931). *An introduction to clinical perimetry / by H. M. Traquair.* eng. London.

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

Trommershäuser, J., Glimcher, P. W., and Gegenfurtner, K. R. (2009). 'Visual processing, learning and feedback in the primate eye movement system'. *Trends in neurosciences* 32 (11), pp. 583–590.

Tseng, Y.-C. and Howes, A. (2015). 'The adaptation of visual search to utility, ecology and design'. *International Journal of Human-Computer Studies* 80, pp. 45–55.

Van Beers, R. J. (2007). 'The sources of variability in saccadic eye movements'. *Journal of Neuroscience* 27 (33), pp. 8757–8770.

Van der Stigchel, S., Bethlehem, R. A., Klein, B. P., Berendschot, T. T., Nijboer, T., and Dumoulin, S. O. (2013). 'Macular degeneration affects eye movement behavior during visual search'. *Frontiers in psychology* 4, p. 579.

Van Beers, R. J. (2008). 'Saccadic eye movements minimize the consequences of motor noise'. *PloS one* 3 (4), e2070.

Vargas-Martín, F. and Peli, E. (2001). 'P-16: Augmented View for Tunnel Vision: Device Testing by Patients in Real Environments'. *SID Symposium Digest of Technical Papers*. Vol. 32. 1. Wiley Online Library, pp. 602–605.

Vatavu, R.-D. (2017). 'Visual impairments and mobile touchscreen interaction: state-of-the-art, causes of visual impairment, and design guidelines'. *International Journal of Human–Computer Interaction* 33 (6), pp. 486–509.

Velisavljević, L. and Elder, J. H. (2008). 'Visual short-term memory for natural scenes: Effects of eccentricity'. *Journal of Vision* 8 (4), pp. 28–28.

Verghese, P. (2001). 'Visual search and attention: A signal detection theory approach'. *Neuron* 31 (4), pp. 523–535.

Vidal-Verdú, F. and Hafez, M. (2007). 'Graphical tactile displays for visually-impaired people'. *IEEE Transactions on neural systems and rehabilitation engineering* 15 (1), pp. 119–130.

Vingrys, A. J. (2000). 'The many faces of glaucomatous optic neuropathy'. *Clinical and Experimental Optometry* 83 (3), pp. 145–160.

Vitu, F. and McConkie, G. W. (2000). 'Chapter 12 - Regressive Saccades and Word Perception in Adult Reading'. *Reading as a Perceptual Process*. Ed. by A. Kennedy, R. Radach, D. Heller, and J. Pynte. Oxford: North-Holland, pp. 301–326. ISBN: 978-0-08-043642-5. DOI: https://doi.org/10.1016/B978-008043642-5/50015-2. URL: https://www.sciencedirect.com/science/article/pii/B9780080436425500152.

Vlaskamp, B. N., Over, E. A., and Hooge, I. T. C. (2005). 'Saccadic search performance: the effect of element spacing'. *Experimental brain research* 167 (2), pp. 246–259.

Von Neumann, J., Morgenstern, O., and Kuhn, H. W. (2007). *Theory of games and economic behavior (commemorative edition)*. Princeton university press.

Von Noorden, G. K. and Mackensen, G. (1962). 'Phenomenology of eccentric fixation'. *American Journal of Ophthalmology* 53 (4), pp. 642–661.

Vorvoreanu, M., Zhang, L., Huang, Y.-H., Hilderbrand, C., Steine-Hanson, Z., and Burnett, M. (2019). 'From gender biases to gender-inclusive design: An empirical investigation'. *Proceedings of the 2019 CHI Conference on human factors in computing systems*, pp. 1–14.

Wallace, J. M., Chung, S. T., and Tjan, B. S. (2017). 'Object crowding in age-related macular degeneration'. *Journal of vision* 17 (1), pp. 33–33.

Walsh, D. V. and Liu, L. (2014). 'Adaptation to a simulated central scotoma during visual search training'. *Vision Research* 96, pp. 75–86.

Wan, Y., Yang, J., Ren, X., Yu, Z., Zhang, R., and Li, X. (2020). 'Evaluation of eye movements and visual performance in patients with cataract'. *Scientific Reports* 10 (1), p. 9875.

Wang, Z., Epps, J., and Chen, S. (2021). 'An investigation of automatic saccade and fixation detection from wearable infrared cameras'. *2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, pp. 2250–2257.

Ware, C. (2019). *Information visualization: perception for design*. Morgan Kaufmann.

Watkins, C. J. C. H. (1989). 'Learning from delayed rewards'.

Watson, G. R., Schuchard, R. A., De l'Aune, W. R., and Watkins, E. (2006). 'Effects of preferred retinal locus placement on text navigation and development of advantageous trained retinal locus'. *Journal of Rehabilitation Research and Development* 43 (6), p. 761.

Westheimer, G. (1982). 'The spatial grain of the perifoveal visual field'. *Vision Research* 22 (1), pp. 157–162. ISSN: 0042-6989. DOI: https://doi.org/10.1016/0042-6989(82)90177-8. URL: https://www.sciencedirect.com/science/article/pii/0042698982901778.

White, J. M. and Bedell, H. E. (1990). 'The oculomotor reference in humans with bilateral macular disease.' *Investigative ophthalmology & visual science* 31 (6), pp. 1149–1161.

Whitney, D. and Levi, D. M. (2011). 'Visual crowding: A fundamental limit on conscious perception and object recognition'. *Trends in cognitive sciences* 15 (4), pp. 160–168.

Whittaker, S. G., Budd, J., and Cummings, R. (1988). 'Eccentric fixation with macular scotoma.' *Investigative ophthalmology & visual science* 29 (2), pp. 268–278.

Wiecek, E. W., Pasquale, L. R., Fiser, J., Dakin, S., and Bex, P. J. (2012). 'Effects of peripheral visual field loss on eye movements during visual search'. *Frontiers in psychology* 3, p. 472.

Williams, L. G. (1967). 'The effects of target specification on objects fixated during visual search'. *Acta psychologica* 27, pp. 355–360.

Wolf, A. and Ueda, K. (2021). 'Contribution of eye-tracking to study cognitive impairments among clinical populations'. *Frontiers in Psychology* 12, p. 590986.

Wolfe, J. M. (1994). 'Guided search 2.0 a revised model of visual search'. *Psychonomic bulletin & review* 1 (2), pp. 202–238.

Wolfe, J. M. (2001). 'Guided Search 4.0: A guided search model that does not require memory for rejected distractors'. *Journal of Vision* 1 (3), pp. 349–349.

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

Wolfe, J. M. and Horowitz, T. S. (2017). 'Five factors that guide attention in visual search'. *Nature Human Behaviour* 1 (3), p. 0058.

Wolffsohn, J. S. and Peterson, R. C. (2003). 'A review of current knowledge on Electronic Vision Enhancement Systems for the visually impaired'. *Ophthalmic and Physiological Optics* 23 (1), pp. 35–42.

Wong, W. L., Su, X., Li, X., Cheung, C. M. G., Klein, R., Cheng, C.-Y., and Wong, T. Y. (2014). 'Global prevalence of age-related macular degeneration and disease burden projection for 2020 and 2040: a systematic review and meta-analysis'. *The Lancet Global Health* 2 (2), e106–e116.

World Health Organization (2007). 'Global Initiative for the Elimination of Avoidable Blindness: action plan 2006-2011'.

World Health Organization (2008). 'Change the definition of blindness'. *retrieved October* 31, p. 2008.

World Health Organization (2022). *Blindness and vision impairment*. Available at: https://www.who.int/en/news-room/fact-sheets/detail/blindness-and-visual-impairment (Accessed: 7 June 2023).

Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., and Bengio, Y. (July 2015). 'Show, Attend and Tell: Neural Image Caption Generation with Visual Attention'. *Proceedings of the 32nd International Conference on Machine Learning*. Ed. by F. Bach and D. Blei. Vol. 37. Proceedings of Machine Learning Research. Lille, France: PMLR, pp. 2048–2057. URL: https://proceedings.mlr.press/v37/xuc15.html.

Yang, W., Huang, J., Wang, R., Zhang, W., Liu, H., and Xiao, J. (2021). 'A survey on tactile displays for visually impaired people'. *IEEE Transactions on Haptics* 14 (4), pp. 712–721.

Yeshurun, Y. and Rashal, E. (2010). 'Precueing attention to the target location diminishes crowding and reduces the critical distance'. *Journal of vision* 10 (10), pp. 16–16.

Zerr, P., Ossandón, J. P., Shareef, I., Van der Stigchel, S., Kekunnaya, R., and Röder, B. (July 2020). 'Successful visually guided eye movements following sight restoration after congenital cataracts'. *Journal of Vision* 20 (7), pp. 3–3. ISSN: 1534-7362. DOI: 10.1167/jov.20.7.3. eprint: https://arvojournals.org/arvo/content_public/journal/jov/938480/i0035-8711-660-1-07038_1613022345.96972.pdf. URL: https://doi.org/10.1167/jov.20.7.3.

Zhang, M. R., Wang, R., Xu, X., Li, Q., Sharif, A., and Wobbrock, J. O. (2021). 'Voicemoji: Emoji Entry Using Voice for Visually Impaired People'. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. Yokohama, Japan: Association for Computing Machinery. ISBN: 9781450380966. DOI: 10.1145/3411764.3445338. URL: https://doi.org/10.1145/3411764.3445338.

Zhao, Y., Cutrell, E., Holz, C., Morris, M. R., Ofek, E., and Wilson, A. D. (2019). 'SeeingVR: A set of tools to make virtual reality more accessible to people with low vision'. *Proceedings of the 2019 CHI conference on human factors in computing systems*, pp. 1–14.

Zhao, Y., Szpiro, S., Knighten, J., and Azenkot, S. (2016). 'CueSee: exploring visual cues for people with low vision to facilitate a visual search task'. *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pp. 73–84.

Appendix A

A Model of The Crowding Effect with Age-related Macular Degeneration

A.1 Crowding Noise Histograms (Four Objects)

A.1.1 Gaussian

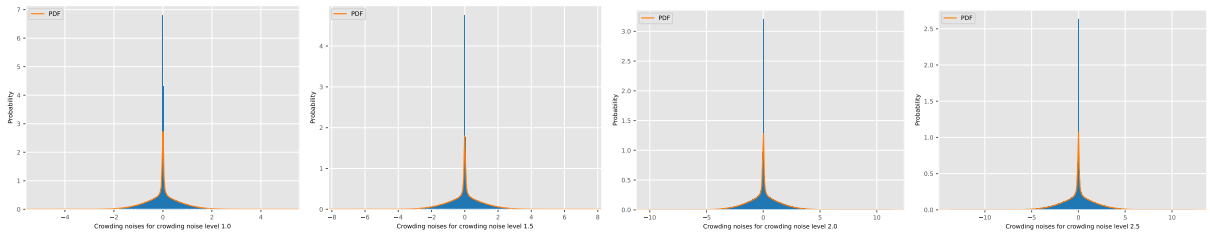

Crowding noise 1.0 Crowding noise 1.5 Crowding noise 2.0 Crowding noise 2.5

Crowding noise 3.0 Crowding noise 3.5

Figure A.1: Example of the crowding noise (Gaussian), we believe the high peak probability at 0 represents the object the agent is fixating on where the eccentricity is near 0, resulting the crowding probability "c" to be almost 0, meaning the noise is 0.

A.1.2 Blending

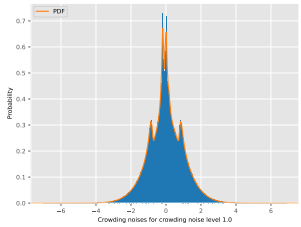

Crowding noise 1.0

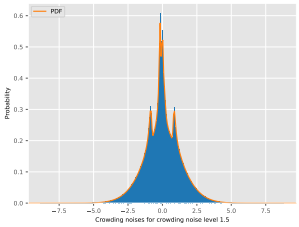

Crowding noise 1.5

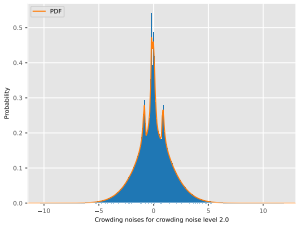

Crowding noise 2.0

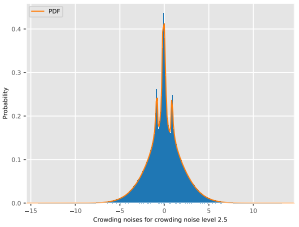

Crowding noise 2.5

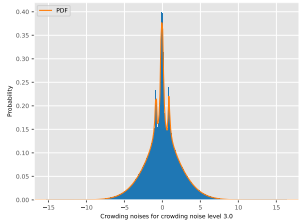

Crowding noise 3.0

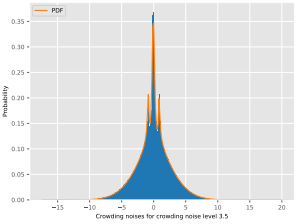

Crowding noise 3.5

Figure A.2: Example of the crowding noise (Blending). The multiple peaks represent the noise of the difference between the object the model is fixating on and the flanker.

A.2 Four Objects: Model Performance (Blending)

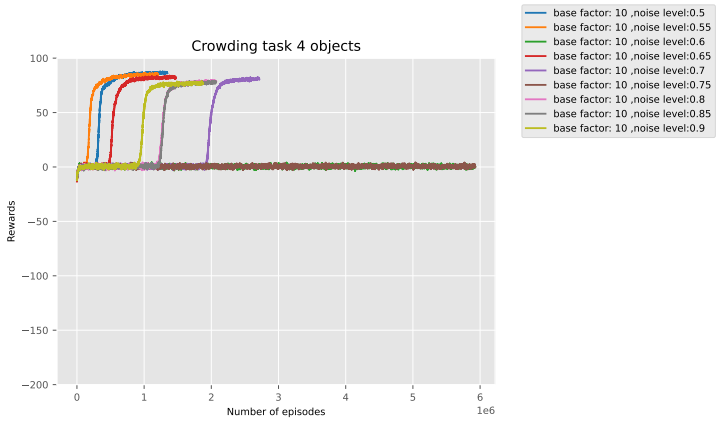

(a) Rewards during learning

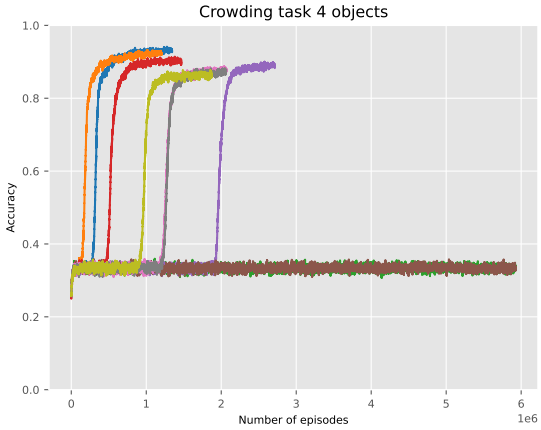

(b) Accuracy during learning

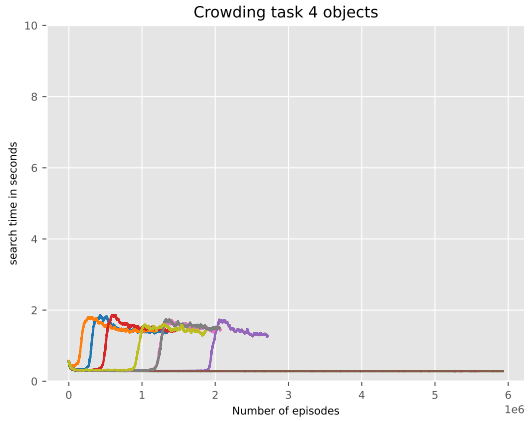

(c) Search time during learning

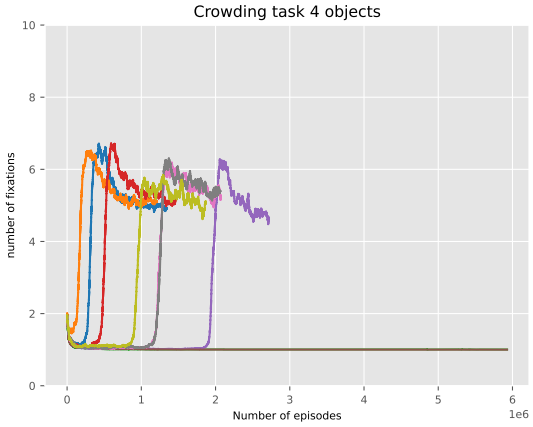

(d) Number of fixations during learning

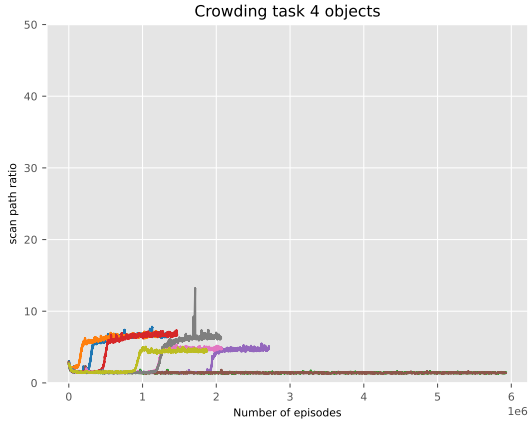

(e) Scan path ratio during learning

Figure A.3: Learning performance with Blending noise (four objects)

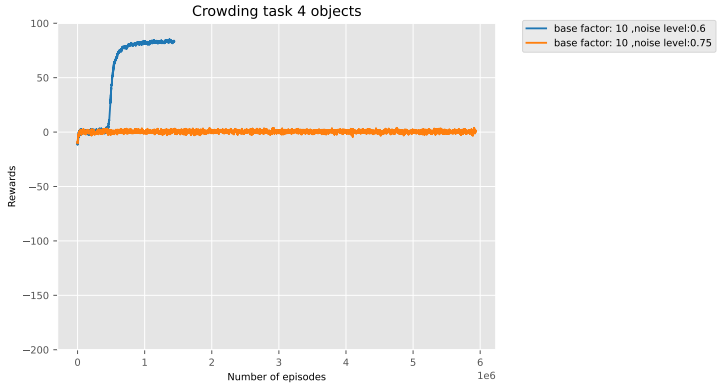

(a) Rewards during learning

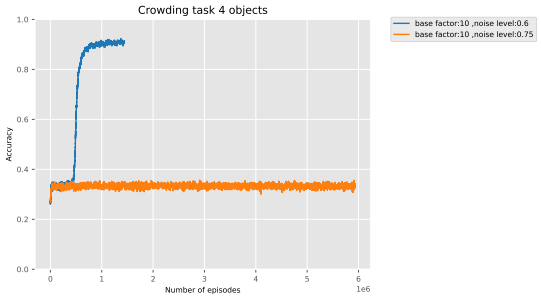

(b) Accuracy during learning

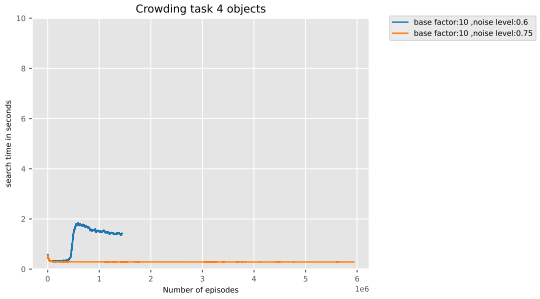

(c) Search time during learning

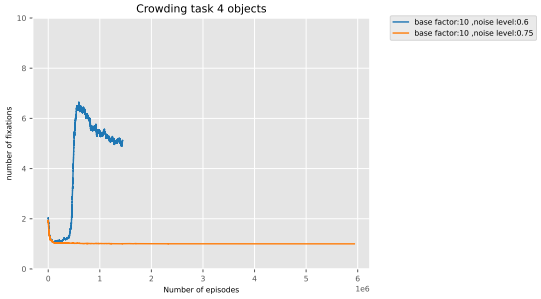

(d) Number of fixations during learning

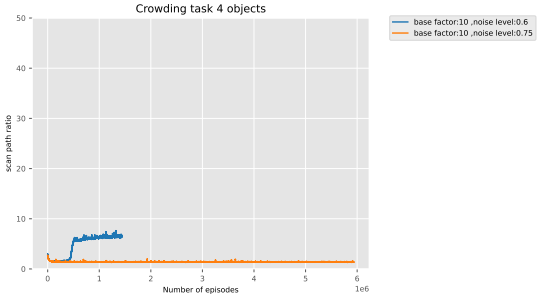

(e) Scan path ratio during learning

Figure A.4: Learning performance with Blending noise (four objects)

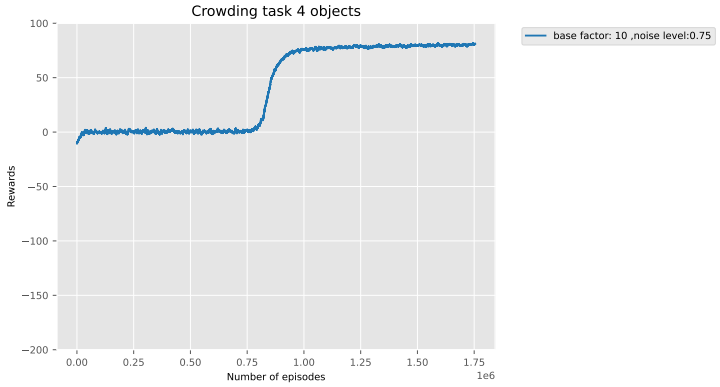

(a) Rewards during learning

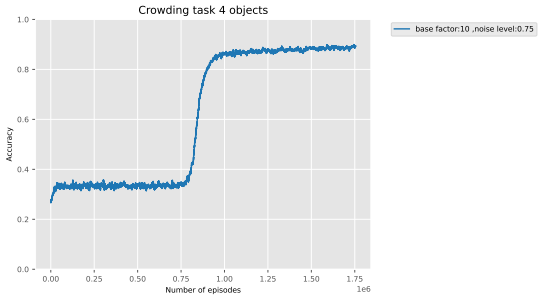

(b) Accuracy during learning

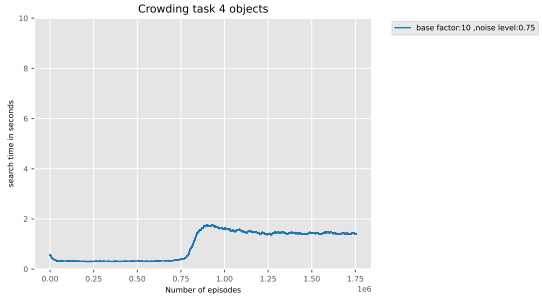

(c) Search time during learning

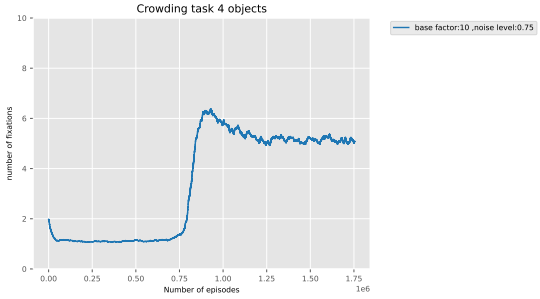

(d) Number of fixations during learning

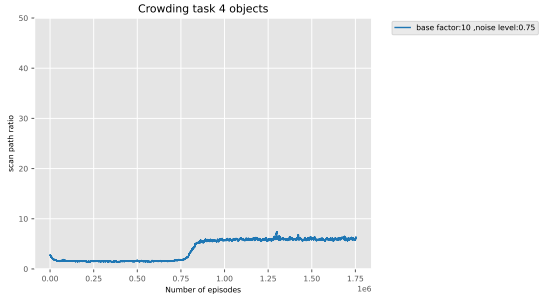

(e) Scan path ratio during learning

Figure A.5: Learning performance with Blending noise (four objects)

A.3 Six Objects: Model Performance (Gaussian)

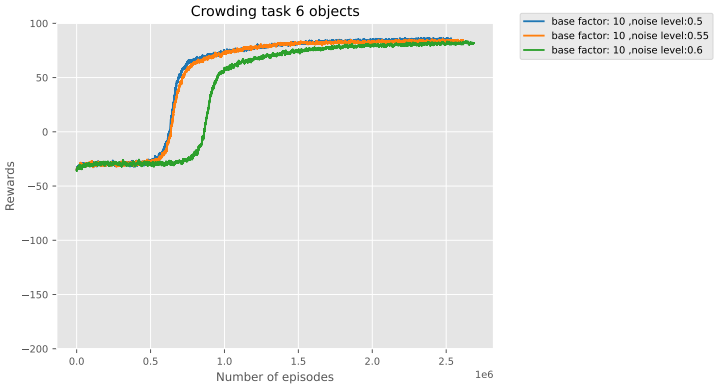

(a) Rewards during learning

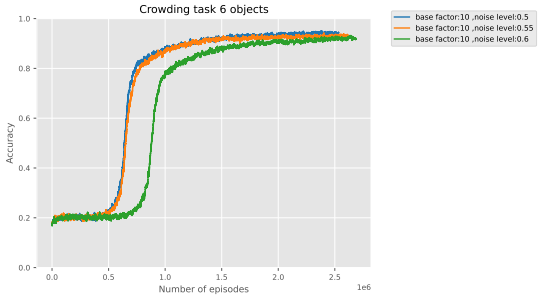

(b) Accuracy during learning

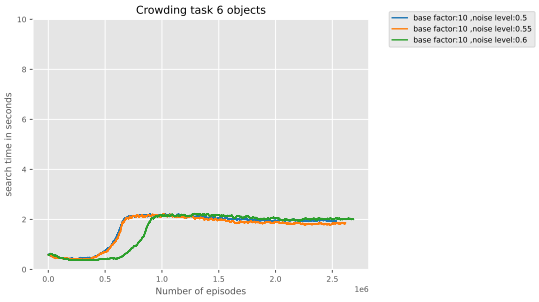

(c) Search time during learning

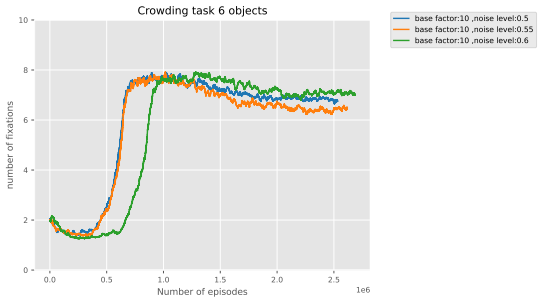

(d) Number of fixations during learning

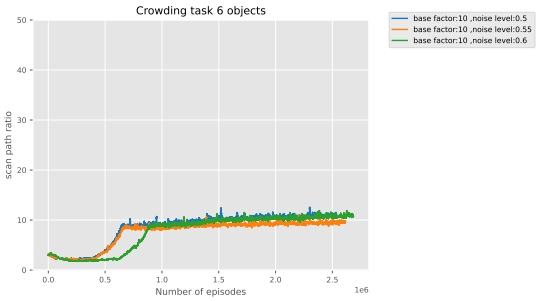

(e) Scan path ratio during learning

Figure A.6: Learning performance with Gaussian noise (six objects)

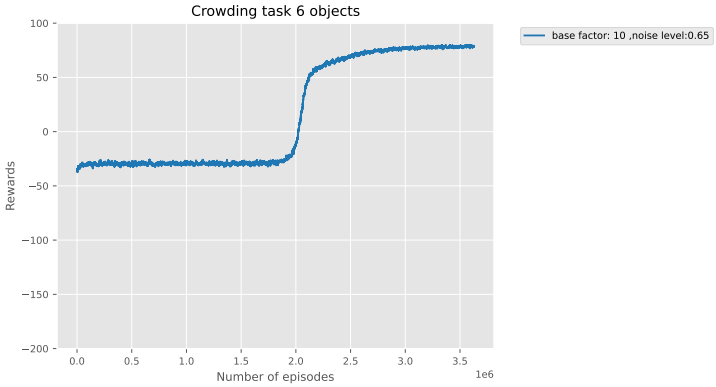

(a) Rewards during learning

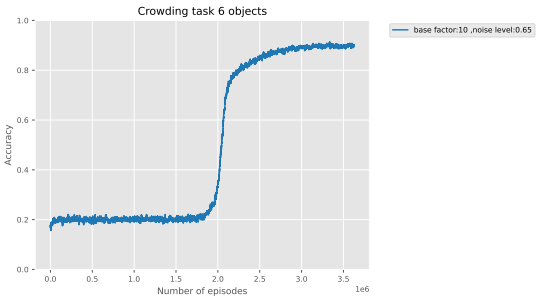

(b) Accuracy during learning

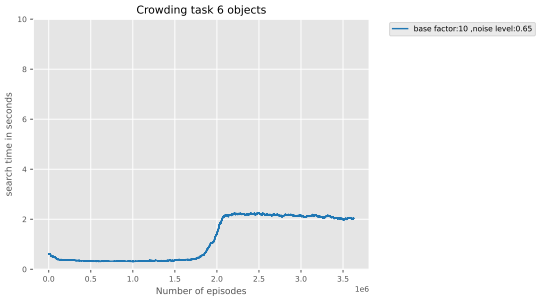

(c) Search time during learning

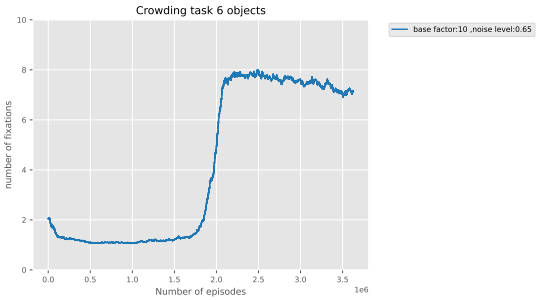

(d) Number of fixations during learning

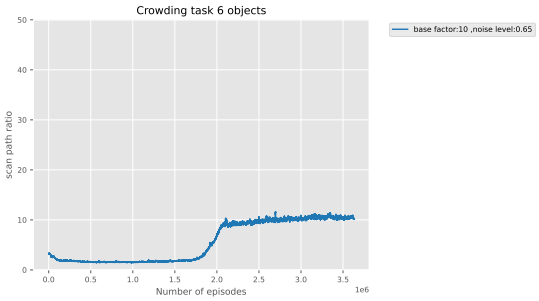

(e) Scan path ratio during learning

Figure A.7: Learning performance with Gaussian noise (six objects)

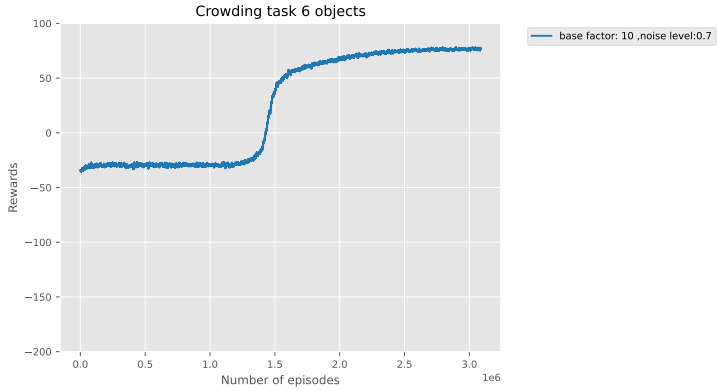

(a) Rewards during learning

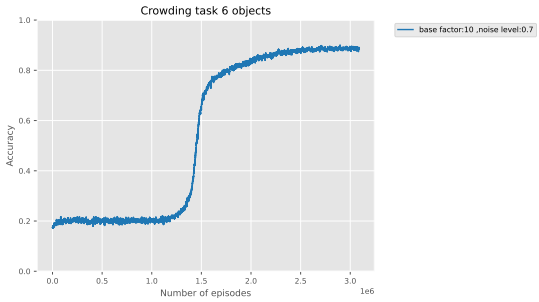

(b) Accuracy during learning

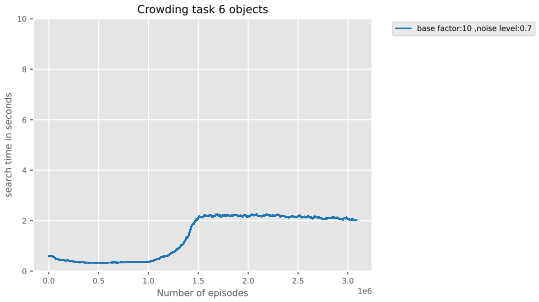

(c) Search time during learning

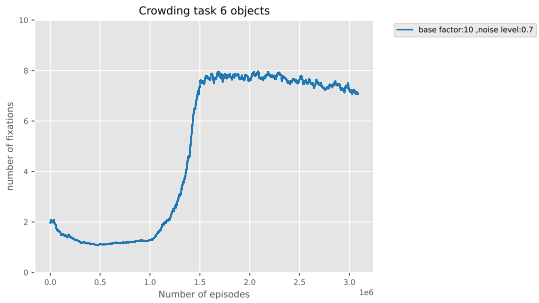

(d) Number of fixations during learning

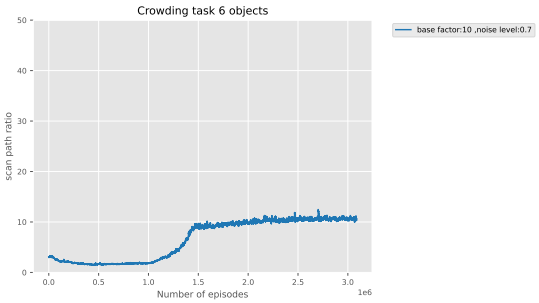

(e) Scan path ratio during learning

Figure A.8: Learning performance with Gaussian noise (six objects)

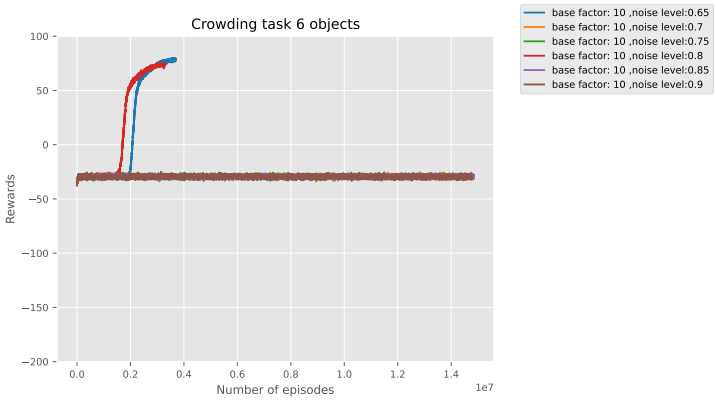

(a) Rewards during learning

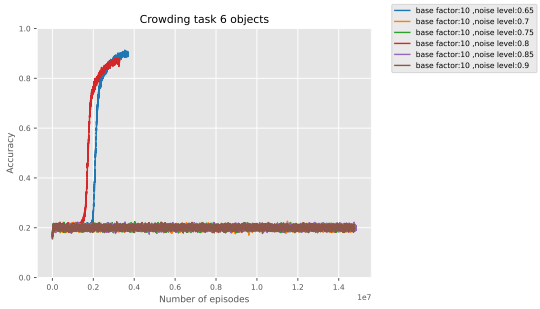

(b) Accuracy during learning

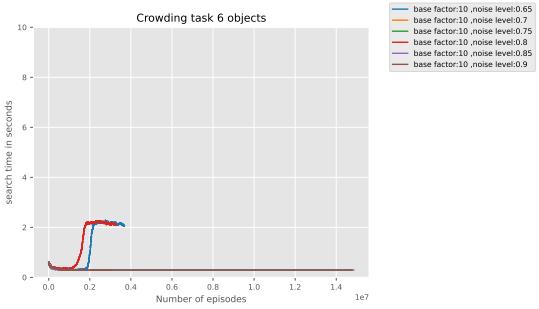

(c) Search time during learning

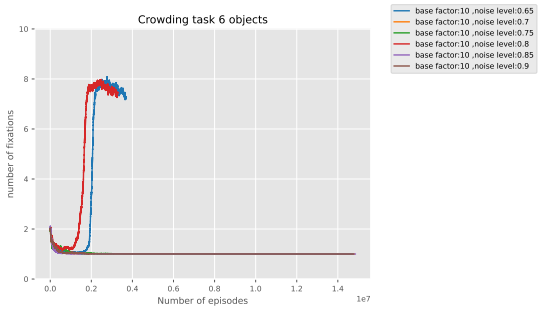

(d) Number of fixations during learning

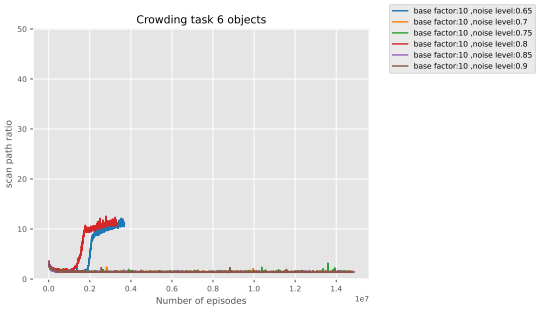

(e) Scan path ratio during learning

Figure A.9: Learning performance with Gaussian noise (six objects)

A.4 Six Objects: Model Performance (Blending)

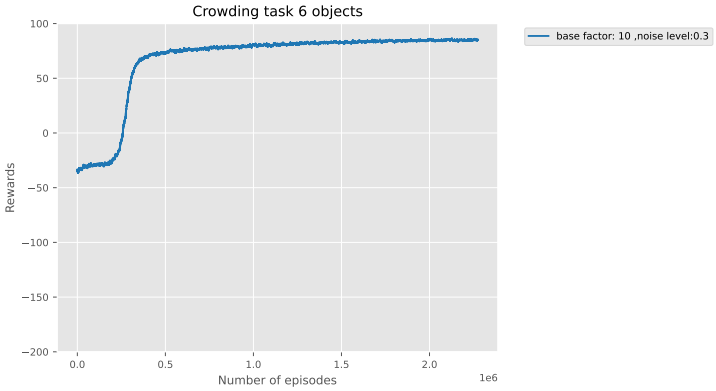

(a) Rewards during learning

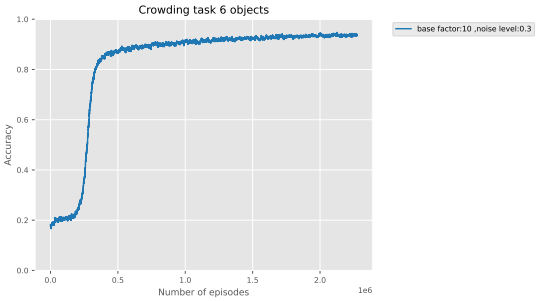

(b) Accuracy during learning

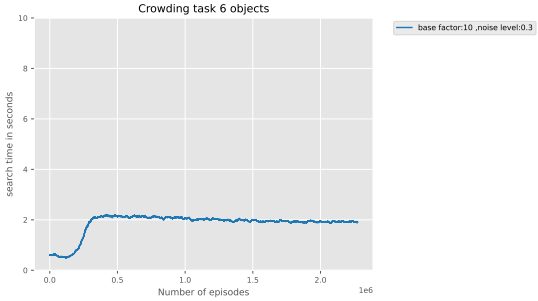

(c) Search time during learning

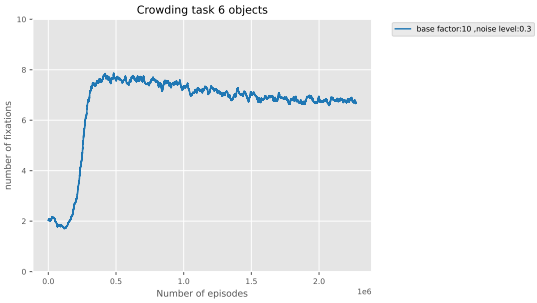

(d) Number of fixations during learning

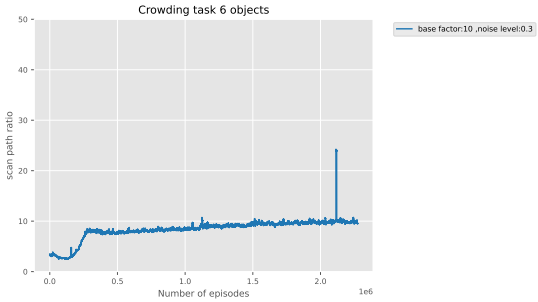

(e) Scan path ratio during learning

Figure A.10: Learning performance with Blending noise (six objects)

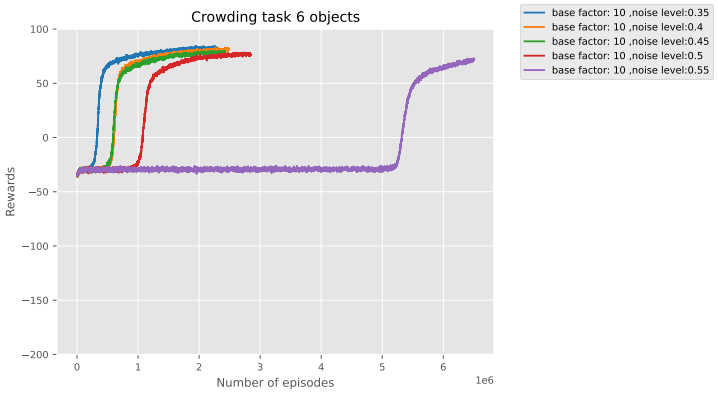

(a) Rewards during learning

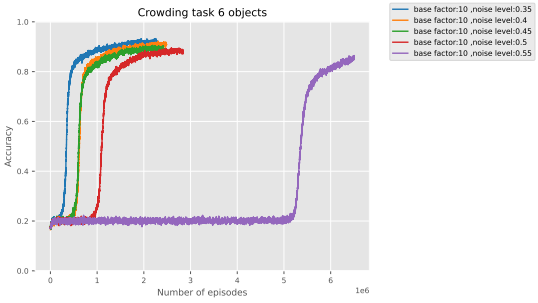

(b) Accuracy during learning

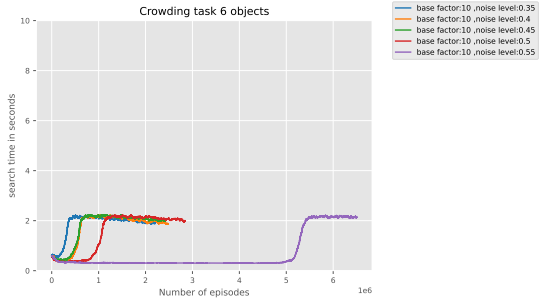

(c) Search time during learning

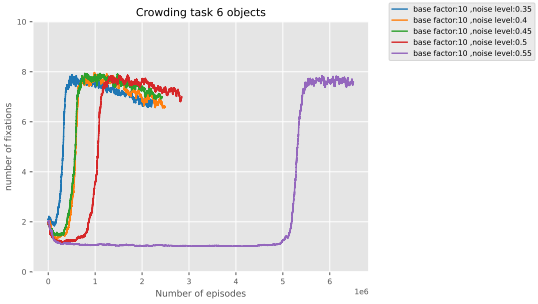

(d) Number of fixations during learning

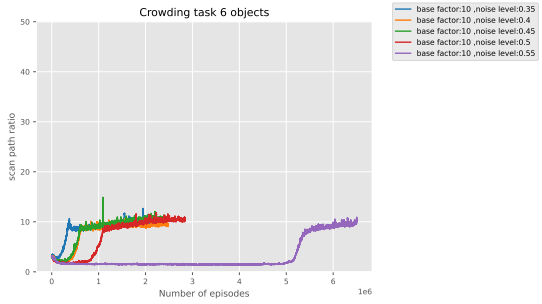

(e) Scan path ratio during learning

Figure A.11: Learning performance with Blending noise (six objects)

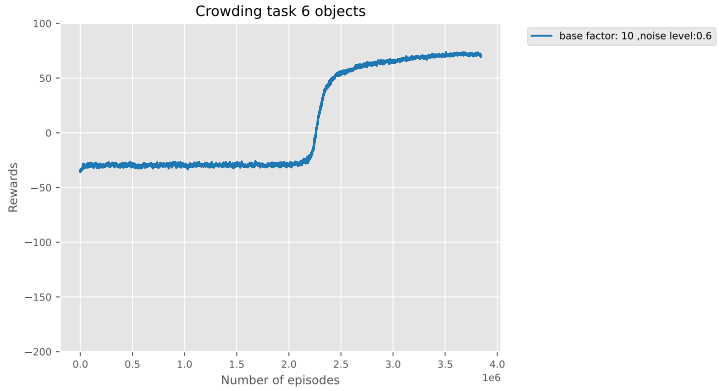

(a) Rewards during learning

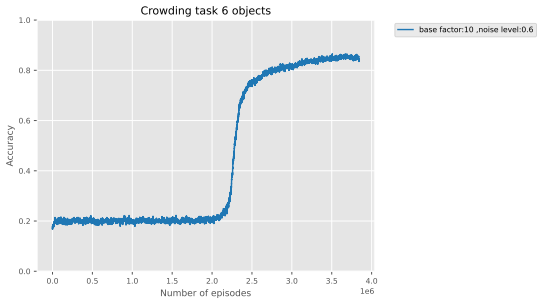

(b) Accuracy during learning

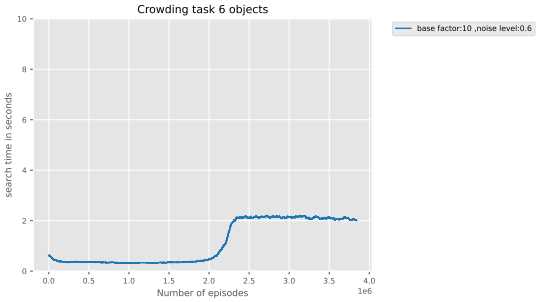

(c) Search time during learning

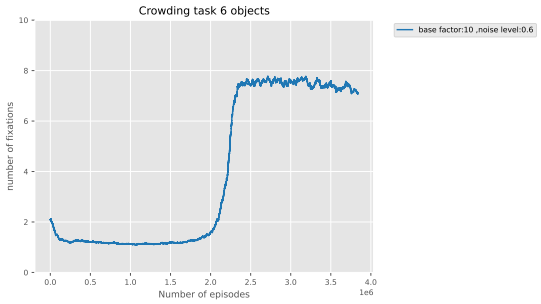

(d) Number of fixations during learning

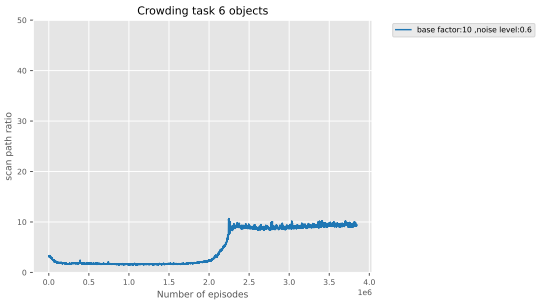

(e) Scan path ratio during learning

Figure A.12: Learning performance with Blending noise (six objects)

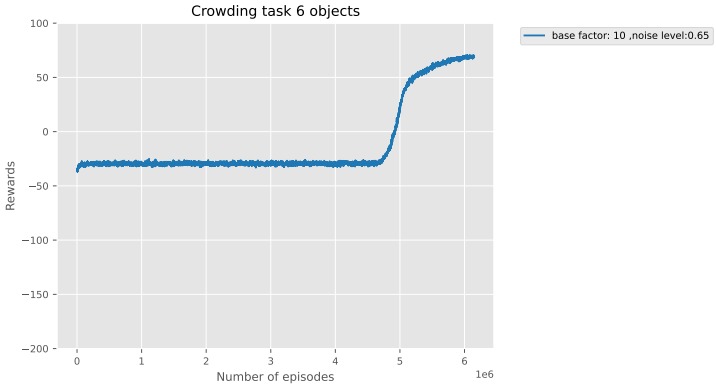

(a) Rewards during learning

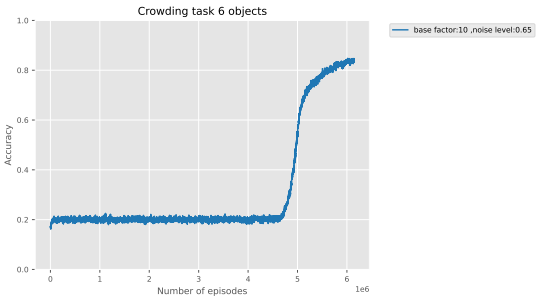

(b) Accuracy during learning

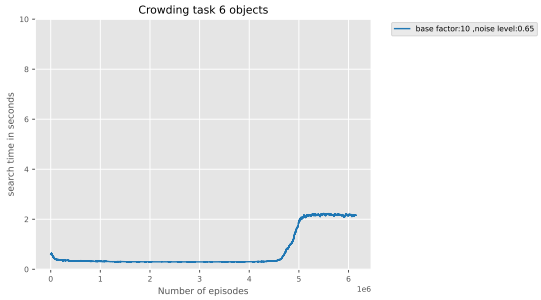

(c) Search time during learning

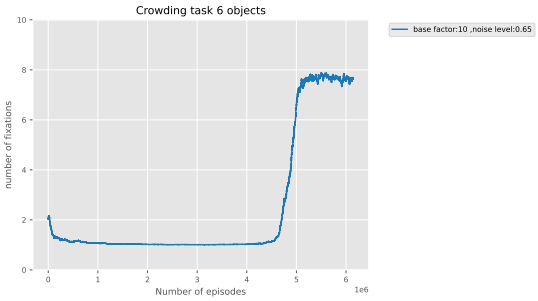

(d) Number of fixations during learning

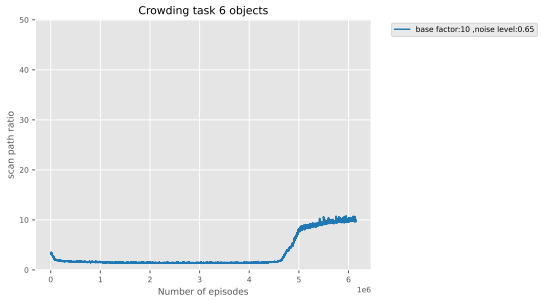

(e) Scan path ratio during learning

Figure A.13: Learning performance with Blending noise (six objects)

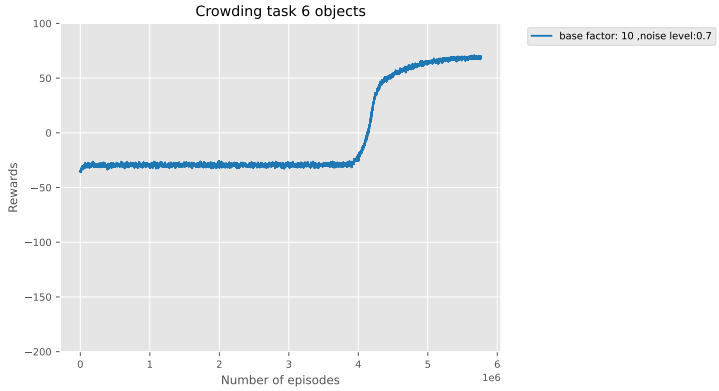

(a) Rewards during learning

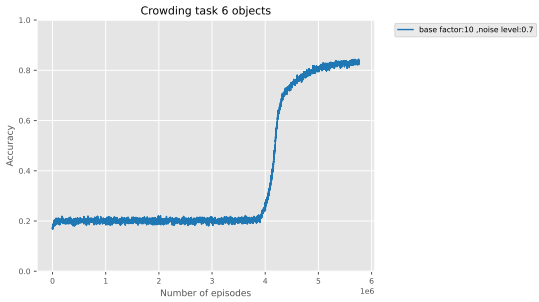

(b) Accuracy during learning

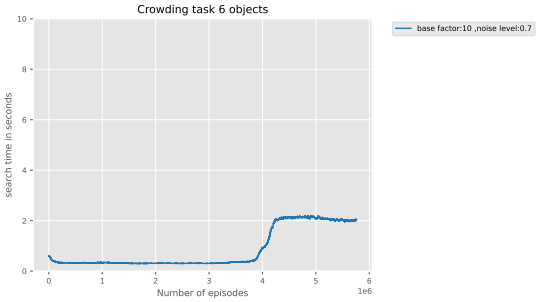

(c) Search time during learning

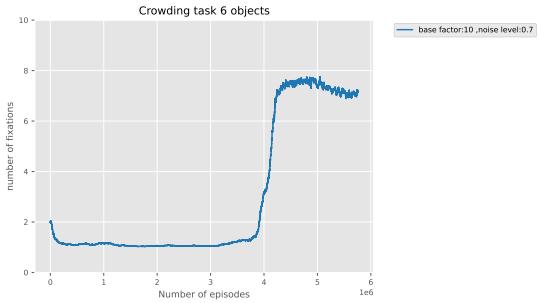

(d) Number of fixations during learning

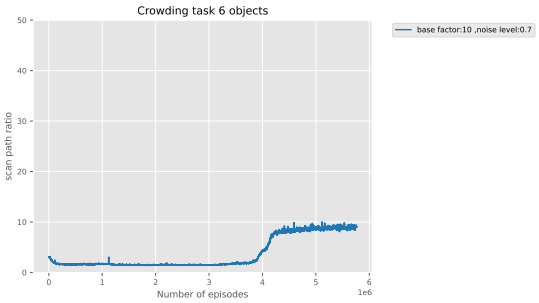

(e) Scan path ratio during learning

Figure A.14: Learning performance with Blending noise (six objects)

Appendix B

A Model of Visual Search with Cataract

B.1 Crowding Noise Histograms

B.1.1 Gaussian

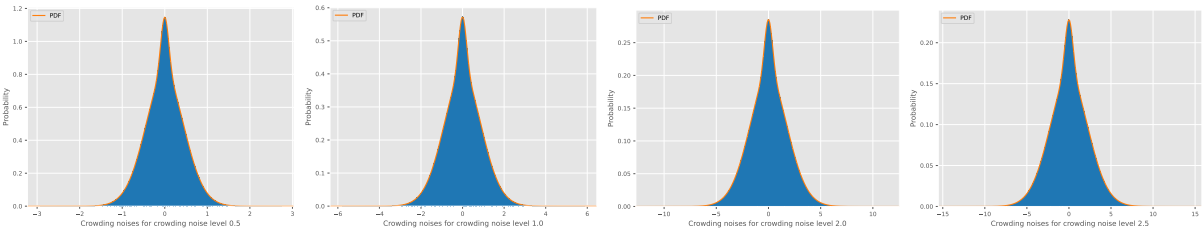

Crowding noise 1.0 Crowding noise 1.5 Crowding noise 2.0 Crowding noise 2.5

Crowding noise 3.0 Crowding noise 3.5

Figure B.1: Example of the crowding noise (Gaussian) in the cataract objects search task

B.1.2 Blending

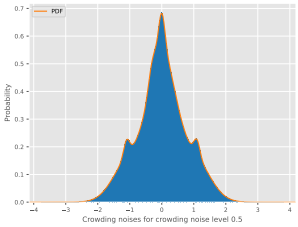

Crowding noise 1.0

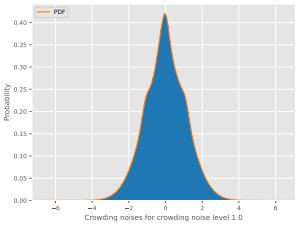

Crowding noise 1.5

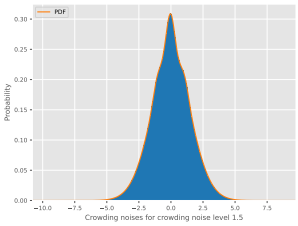

Crowding noise 2.0

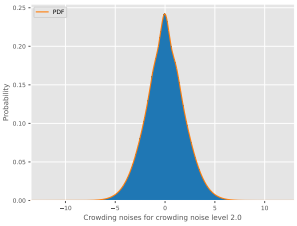

Crowding noise 2.5

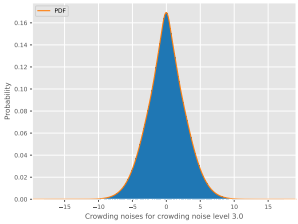

Crowding noise 3.0

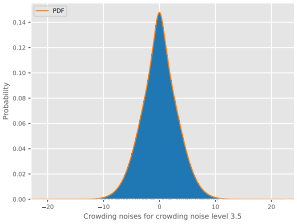

Crowding noise 3.5

Figure B.2: Example of the crowding noise (Blending) in the cataract objects search task

B.1.3 Task 1: Learning Performance (Blending)

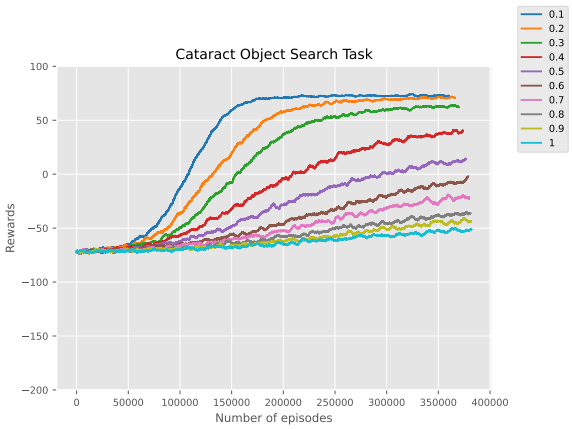

(a) Reward during learning

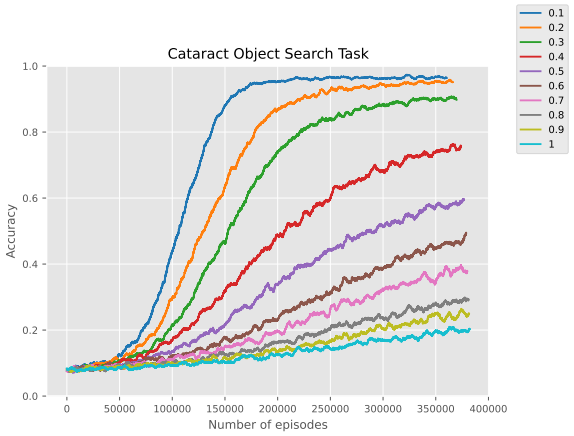

(b) Accuracy during learning

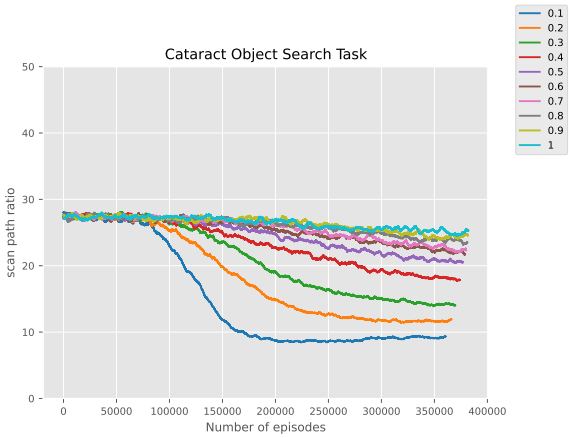

(c) Scan path ratio during learning

Figure B.3: The learning performance of the model with object search in cataracts - Blending noise.

B.2 Task 2: example visualisation

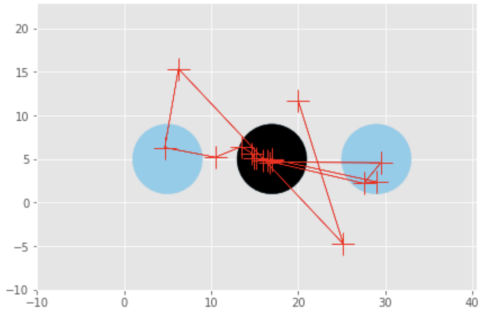

Figure B.4: An example of fixations outside the screen limits in the cataracts face recognition task (noise level 0.5)